

Image Processing Toolbox

For Use with MATLAB®

- Computation
- Visualization
- Programming

How to Contact The MathWorks:



www.mathworks.com Web
comp.soft-sys.matlab Newsgroup



support@mathworks.com Technical support
suggest@mathworks.com Product enhancement suggestions
bugs@mathworks.com Bug reports
doc@mathworks.com Documentation error reports
service@mathworks.com Order status, license renewals, passcodes
info@mathworks.com Sales, pricing, and general information



508-647-7000 Phone



508-647-7001 Fax



The MathWorks, Inc. Mail
3 Apple Hill Drive
Natick, MA 01760-2098

For contact information about worldwide offices, see the MathWorks Web site.

Image Processing Toolbox User's Guide

© COPYRIGHT 1993 - 2004 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

MATLAB, Simulink, Stateflow, Handle Graphics, and Real-Time Workshop are registered trademarks, and TargetBox is a trademark of The MathWorks, Inc.

Other product or brand names are trademarks or registered trademarks of their respective holders.

Printing History:	August 1993	First printing	Version 1
	May 1997	Second printing	Version 2
	April 2001	Third printing	Revised for Version 3.0
	June 2001	Online only	Revised for Version 3.1 (Release 12.1)
	July 2002	Online only	Revised for Version 3.2 (Release 13)
	May 2003	Fourth printing	Revised for Version 4.0 (Release 13.0.1)
	September 2003	Online only	Revised for Version 4.1 (Release 13SP1)
	June 2004	Online only	Revised for Version 4.2 (Release 14)
	August 2004	Online only	Revised for Version 5.0 (Release 14+)
	October 2004	Fifth printing	Revised for Version 5.0.1 (Release 14SP1)

Getting Started

1

What Is the Image Processing Toolbox?	1-2
Configuration Notes	1-2
Related Products	1-3
Example 1 — Some Basic Concepts	1-4
1. Read and Display an Image	1-4
2. Check How the Image Appears in the Workspace	1-5
3. Improving Image Contrast	1-6
4. Write the Image to a Disk File	1-8
5. Check the Contents of the Newly Written File	1-8
Example 2 — Advanced Topics	1-10
1. Read and Display an Image	1-10
2. Estimate the Value of Background Pixels	1-11
3. View the Background Approximation as a Surface	1-12
4. Create an Image with a Uniform Background	1-13
5. Adjust the Contrast in the Processed Image	1-14
6. Create a Binary Version of the Image	1-15
7. Determine the Number of Objects in the Image	1-17
8. Examine the Label Matrix	1-17
9. Display the Label Matrix as a Pseudocolor Indexed Image	1-18
10. Measure Object Properties in the Image	1-19
11. Compute Statistical Properties of Objects in the Image ..	1-20
Getting Help	1-22
Online Help	1-22
Image Processing Demos	1-23
MATLAB Newsgroup	1-23
Image Credits	1-24

Images in MATLAB and the Image Processing Toolbox . . .	2-2
Image Types in the Toolbox	2-3
Binary Images	2-4
Indexed Images	2-4
Intensity Images	2-5
Truecolor Images	2-6
Converting Between Image Types	2-11
Color Space Conversions	2-12
Converting Between Image Classes	2-13
Losing Information in Conversions	2-13
Converting Indexed Images	2-13
Multiframe Image Arrays	2-14
Reading and Writing Image Data	2-15
Reading a Graphics Image	2-15
Writing a Graphics Image	2-16
Querying a Graphics File	2-19
Converting Graphics File Formats	2-19
Working with DICOM Files	2-21
Reading Image Data from a DICOM File	2-21
Reading Metadata from a DICOM File	2-22
Writing Image Data or Metadata to a DICOM File	2-23
Example: Creating a New Series	2-25
Image Arithmetic	2-28
Image Arithmetic Saturation Rules	2-28
Nesting Calls to Image Arithmetic Functions	2-29
Coordinate Systems	2-30
Pixel Coordinates	2-30
Spatial Coordinates	2-31

Overview	3-2
Understanding Handle Graphics Object Property Settings ...	3-2
Using imshow to Display Images	3-4
Specifying the Initial Image Magnification	3-5
Controlling the Appearance of the Figure	3-6
Using the Image Tool to Explore Images	3-8
Opening the Image Tool	3-10
Specifying the Initial Image Magnification	3-11
Specifying the Colormap	3-11
Importing Image Data from the Workspace	3-13
Exporting the Image to the Workspace	3-14
Closing the Image Tool	3-14
Printing the Image in the Image Tool	3-15
Using Image Tool Navigation Aids	3-16
Overview Navigation	3-16
Panning the Image Displayed in the Image Tool	3-19
Zooming In and Out on an Image	3-19
Specifying the Magnification of the Image	3-20
Getting Information about the Pixels in an Image	3-22
Determining the Value of Individual Pixels	3-22
Getting the Display Range of an Image	3-24
Viewing Pixel Values with the Pixel Region Tool	3-25
Getting Information About an Image	3-29
Adjusting the Contrast and Brightness of an Image	3-31
Understanding Contrast Adjustment	3-32
Adjusting Contrast and Brightness	3-34
Using the Window/Level Tool	3-37
Autoscaling	3-39
Viewing Multiple Images	3-41
Displaying Each Image in a Separate Figure	3-41

Displaying Multiple Images in the Same Figure	3-42
Displaying Different Image Types	3-45
Displaying Indexed Images	3-45
Displaying Intensity Images	3-46
Displaying Binary Images	3-47
Displaying Truecolor Images	3-49
Special Display Techniques	3-51
Adding a Colorbar	3-51
Displaying All Frames of a Multiframe Image at Once	3-52
Converting a Multiframe Image to a Movie	3-54
Texture Mapping	3-55
Printing Images	3-56
Printing and Handle Graphics Object Properties	3-56
Setting Toolbox Display Preferences	3-58
Toolbox Preferences	3-58
Retrieving the Values of Toolbox Preferences	3-59
Setting the Values of Toolbox Preferences	3-59

Building GUIs with Modular Tools

4

Overview	4-2
Using Modular Tools	4-6
Displaying the Target Image	4-7
Specifying the Target Image	4-7
Specifying the Parent of a Modular Tool	4-11
Positioning the Modular Tools in a GUI	4-14
Example: Building a Pixel Information GUI	4-15
Adding Navigation Aids to a GUI	4-17
Making Connections for Interactivity	4-22
Creating Your Own Modular Tools	4-28

Spatial Transformations

5

Terminology	5-2
Interpolation	5-3
Interpolation Methods	5-3
Image Types	5-4
Image Resizing	5-5
Specifying the Size of the Output Image	5-5
Specifying the Interpolation Method	5-6
Using Filters to Prevent Aliasing	5-7
Image Rotation	5-8
Specifying the Interpolation Method	5-8
Specifying the Size of the Output Image	5-9
Image Cropping	5-10
Performing General Spatial Transformations	5-11
Specifying the Transformation Type	5-11
Performing the Transformation	5-13
Advanced Spatial Transformation Techniques	5-13

Image Registration

6

Terminology	6-2
Registering an Image	6-4
Point Mapping	6-4
Example: Registering to a Digital Orthophoto	6-5
Types of Supported Transformations	6-12
Selecting Control Points	6-14

Using the Control Point Selection Tool	6-14
Starting the Control Point Selection Tool	6-15
Viewing the Images	6-17
Specifying Matching Control Point Pairs	6-22
Saving Control Points	6-29
Using Correlation to Improve Control Points	6-31

Linear Filtering and Filter Design

7

Terminology	7-2
Linear Filtering	7-4
Convolution	7-4
Correlation	7-6
Filtering Using <code>imfilter</code>	7-7
Using Predefined Filter Types	7-15
Filter Design	7-17
FIR Filters	7-17
Frequency Transformation Method	7-18
Frequency Sampling Method	7-19
Windowing Method	7-20
Creating the Desired Frequency Response Matrix	7-22
Computing the Frequency Response of a Filter	7-23

Transforms

8

Terminology	8-2
Fourier Transform	8-3
Definition of Fourier Transform	8-3

Discrete Fourier Transform	8-8
Applications of the Fourier Transform	8-11
Discrete Cosine Transform	8-17
The DCT Transform Matrix	8-18
DCT and Image Compression	8-19
Radon Transform	8-21
Plotting the Radon Transform	8-23
Viewing the Radon Transform as an Image	8-25
Using the Radon Transform to Detect Lines	8-27
Inverse Radon Transform	8-29
Example: Reconstructing an Image from Parallel Projection Data	8-32
Fan-Beam Projection Data	8-36
Computing Fan-Beam Projection Data	8-37
Reconstructing an Image from Fan-Beam Projection Data . . .	8-39
Working with Fan-Beam Projection Data	8-40

Morphological Operations

9

Terminology	9-2
Dilation and Erosion	9-4
Understanding Dilation and Erosion	9-4
Structuring Elements	9-7
Dilating an Image	9-11
Eroding an Image	9-12
Combining Dilation and Erosion	9-14
Dilation- and Erosion-Based Functions	9-16
Morphological Reconstruction	9-19
Marker and Mask	9-19
Pixel Connectivity	9-23
Flood-Fill Operations	9-26
Finding Peaks and Valleys	9-29

Distance Transform	9-37
Objects, Regions, and Feature Measurement	9-40
Connected-Component Labeling	9-40
Selecting Objects in a Binary Image	9-42
Finding the Area of the Foreground of a Binary Image	9-42
Finding the Euler Number of a Binary Image	9-43
Lookup Table Operations	9-44

Analyzing and Enhancing Images

10

Terminology	10-2
Pixel Values and Statistics	10-3
Pixel Selection	10-3
Intensity Profile	10-5
Image Contours	10-9
Image Histogram	10-10
Summary Statistics	10-12
Region Property Measurement	10-12
Image Analysis	10-13
Edge Detection	10-13
Boundary Tracing	10-15
Line Detection Using the Hough Transform	10-20
Quadtree Decomposition	10-24
Texture Analysis	10-26
Texture Filter Functions	10-26
Gray-Level Co-Occurrence Matrix (GLCM)	10-29
Intensity Adjustment	10-36
Adjusting Intensity Values to a Specified Range	10-37
Histogram Equalization	10-41
Contrast-Limited Adaptive Histogram Equalization	10-43

Decorrelation Stretching	10-44
Noise Removal	10-49
Using Linear Filtering	10-49
Using Median Filtering	10-49
Using Adaptive Filtering	10-52

Region-Based Processing

11

Terminology	11-2
Specifying a Region of Interest	11-3
Selecting a Polygon	11-3
Other Selection Methods	11-4
Filtering a Region	11-6
Example: Filtering a Region in an Image	11-6
Specifying the Filtering Operation	11-7
Filling a Region	11-9

Image Deblurring

12

Terminology	12-2
Understanding Deblurring	12-3
Causes of Blurring	12-3
Deblurring Model	12-3
Using the Deblurring Functions	12-6
Deblurring with the Wiener Filter	12-7
Deblurring with a Regularized Filter	12-9

Deblurring with the Lucy-Richardson Algorithm	12-11
Deblurring with the Blind Deconvolution Algorithm	12-16
Creating Your Own Deblurring Functions	12-22

Avoiding Ringing in Deblurred Images	12-23
---	--------------

Color

13

Terminology	13-2
Working with Different Screen Bit Depths	13-3
Determining Screen Bit Depth	13-3
Choosing a Screen Bit Depth	13-4
Reducing the Number of Colors in an Image	13-6
Using rgb2ind	13-7
Reducing Colors in an Indexed Image	13-12
Dithering	13-13
Converting Color Data Between Color Spaces	13-15
Converting Between Device-Independent Color Spaces	13-15
Performing Profile-Based Conversions	13-19
Converting Between Device-Dependent Color Spaces	13-23

Neighborhood and Block Operations

14

Terminology	14-2
Block Processing Operations	14-3
Types of Block Processing Operations	14-3
Sliding Neighborhood Operations	14-4

Padding Borders	14-5
Linear and Nonlinear Filtering	14-5
Distinct Block Operations	14-8
Overlap	14-9
Column Processing	14-11
Sliding Neighborhoods	14-11
Distinct Blocks	14-12

Function Reference

15

Functions – Categorical List	15-2
Image Input, Output, and Display	15-3
Modular Interactive Tools	15-5
Spatial Transformation and Registration	15-7
Image Analysis and Statistics	15-8
Image Arithmetic	15-9
Image Enhancement and Restoration	15-10
Linear Filtering and Transforms	15-11
Morphological Operations	15-13
Region-Based, Neighborhood, and Block Processing	15-15
Colormap and Color Space Functions	15-16
Miscellaneous Functions	15-18
Functions – Alphabetical List	15-20

Index

Getting Started

This chapter contains two examples to get you started doing image processing using MATLAB® and the Image Processing Toolbox. The examples contain cross-references to other sections in the documentation manual that have in-depth discussions on the concepts presented in the examples.

What Is the Image Processing Toolbox? (p. 1-2)	Introduces the Image Processing Toolbox and its capabilities.
Example 1 — Some Basic Concepts (p. 1-4)	Guides you through an example of some of the basic image processing capabilities of the toolbox, including reading, writing, and displaying images
Example 2 — Advanced Topics (p. 1-10)	Guides you through some advanced image processing topics, including components labeling, object property measurement, image arithmetic, morphological image processing, and contrast enhancement
Getting Help (p. 1-22)	Provides pointers to additional sources of information
Image Credits (p. 1-24)	Provides information about the sources of the images used in the documentation.

What Is the Image Processing Toolbox?

The Image Processing Toolbox is a collection of functions that extend the capability of the MATLAB numeric computing environment. The toolbox supports a wide range of image processing operations, including

- Spatial image transformations
- Morphological operations
- Neighborhood and block operations
- Linear filtering and filter design
- Transforms
- Image analysis and enhancement
- Image registration
- Deblurring
- Region of interest operations

Many of the toolbox functions are MATLAB M-files, a series of MATLAB statements that implement specialized image processing algorithms. You can view the MATLAB code for these functions using the statement

```
type function_name
```

You can extend the capabilities of the Image Processing Toolbox by writing your own M-files, or by using the toolbox in combination with other toolboxes, such as the Signal Processing Toolbox and the Wavelet Toolbox.

For a list of the new features in this version, see the Release Notes documentation.

Configuration Notes

To determine if the Image Processing Toolbox is installed on your system, type this command at the MATLAB prompt.

```
ver
```

When you enter this command, MATLAB displays information about the version of MATLAB you are running, including a list of all toolboxes installed on your system and their version numbers.

For information about installing the toolbox, see the MATLAB Installation Guide for your platform.

For the most up-to-date information about system requirements, see the system requirements page, available in the products area at The MathWorks Web site (www.mathworks.com).

Related Products

The MathWorks provides several products that are relevant to the kinds of tasks you can perform with the Image Processing Toolbox and that extend the capabilities of MATLAB. For information about these related products, see www.mathworks.com/products/image/related.html.

Example 1 – Some Basic Concepts

This example introduces some basic image processing concepts, including reading and writing images, performing histogram equalization on an image, and getting information about an image. The example breaks this process into the following steps:

- Step 1: Read and display an image
- Step 2: Check how the image appears in the workspace
- Step 3: Improving image contrast
- Step 4: Write the image to a disk file
- Step 5: Get information about a graphics file

Before beginning with this example, you should already have installed the Image Processing Toolbox and have started MATLAB. If you are new to MATLAB, read the MATLAB Getting Started documentation to learn about basic MATLAB concepts.

1. Read and Display an Image

Clear the MATLAB workspace of any variables and close open figure windows.

```
clear, close all
```

To read an image, use the `imread` command. The example reads one of the sample images included with the Image Processing Toolbox, `pout.tif`, and stores it in an array named `I`.

```
I = imread('pout.tif');
```

`imread` infers from the file that the graphics file format is Tagged Image File Format (TIFF). For the list of supported graphics file formats, see the `imread` function reference documentation.

Now display the image. The toolbox includes two image display functions: `imshow` and `imshow`. `imshow` is the toolbox's fundamental image display function. `imshow` starts the Image Tool which presents an integrated environment for displaying images and performing some common image processing tasks. The Image Tool provides all the image display capabilities of

`imshow` but also provides access to several other tools for navigating and exploring images, such as scroll bars, the Pixel Region tool, Image Information tool, and the Contrast Adjustment tool. For more information, see Chapter 3, “Displaying and Exploring Images.”.

You can use either function to display an image. This example uses `imshow`.

```
imshow(I)
```



Intensity Image pout.tif

2. Check How the Image Appears in the Workspace

To see how the `imread` function stores the image data in the workspace, check the Workspace browser in the MATLAB desktop. The Workspace browser displays information about all the variables you create during a MATLAB session. The `imread` function returned the image data in the variable `I`, which is a 291-by-240 element array of `uint8` data. MATLAB can store images as `uint8`, `uint16`, or `double` arrays.

You can also get information about variables in the workspace by calling the `whos` command.

```
whos
Name      Size      Bytes  Class

      I      291x240    69840  uint8 array
```


```
Grand total is 69840 elements using 69840 bytes
```

For more information about image storage classes, see “Reading a Graphics Image” on page 2-15.

3. Improving Image Contrast

As you can see, `pout.tif` is a somewhat low contrast image. To see the distribution of intensities in `pout.tif`, you can create a histogram by calling the `imhist` function. (Precede the call to `imhist` with the `figure` command so that the histogram does not overwrite the display of the image `I` in the current figure window.)

```
figure, imhist(I)
```



Notice how the intensity range is rather narrow. It does not cover the potential range of `[0, 255]`, and is missing the high and low values that would result in good contrast.

The toolbox provides several ways to improve the contrast in an image. One way is to call the `histeq` function to spread the intensity values over the full range of the image, a process called *histogram equalization*.

```
I2 = histeq(I);
```

Display the new equalized image, `I2`, in a new figure window.


```
figure, imshow(I2)
```



Equalized Version of pout.tif

Call `imhist` again to create a histogram of the equalized image `I2`. If you compare the two histograms, the histogram of `I2` is more spread out than the histogram of `I1`.

```
figure, imhist(I2)
```



The toolbox includes several other functions that perform contrast adjustment, including the `imadjust` and `adapthisteq` functions. See “Intensity Adjustment” on page 10-36 for more information. In addition, the toolbox includes an interactive tool, called the Adjust Contrast tool, that you can use to adjust the contrast and brightness of an image interactively. To start this tool, call the `imcontrast` function. For more information, see “Adjusting the Contrast and Brightness of an Image” on page 3-31.

4. Write the Image to a Disk File

To write the newly adjusted image `I2` to a disk file, use the `imwrite` function. If you include the filename extension `'.png'`, the `imwrite` function writes the image to a file in Portable Network Graphics (PNG) format, but you can specify other formats.

```
imwrite (I2, 'pout2.png');
```

See the `imwrite` function reference page for a list of file formats it supports. See also “Writing a Graphics Image” on page 2-16 for a tutorial discussion on writing images using the Image Processing Toolbox.

5. Check the Contents of the Newly Written File

To see what `imwrite` wrote to the disk file, use the `imfinfo` function. This function returns information about the image in the file, such as its format, size, width, and height. See “Querying a Graphics File” on page 2-19 for more information about using `imfinfo`.

```
imfinfo('pout2.png')

ans =

    Filename: 'pout2.png'
  FileModDate: '29-Dec-2003 09:34:39'
    FileSize: 36938
      Format: 'png'
  FormatVersion: []
        Width: 240
        Height: 291
    BitDepth: 8
   ColorType: 'grayscale'
FormatSignature: [137 80 78 71 13 10 26 10]
```

```
        Colormap: []
        Histogram: []
        InterlaceType: 'none'
        Transparency: 'none'
SimpleTransparencyData: []
        BackgroundColor: []
        RenderingIntent: []
        Chromaticities: []
            Gamma: []
            XResolution: []
            YResolution: []
        ResolutionUnit: []
            XOffset: []
            YOffset: []
        OffsetUnit: []
        SignificantBits: []
        ImageModTime: '29 Dec 2003 14:34:39 +0000'
            Title: []
            Author: []
            Description: []
            Copyright: []
        CreationTime: []
            Software: []
        Disclaimer: []
            Warning: []
            Source: []
            Comment: []
        OtherText: []
```

Example 2 – Advanced Topics

This example introduces some advanced image processing concepts. The example calculates statistics about objects in the image but, before it performs these calculations, it preprocesses the image to achieve better results. The preprocessing involves creating a uniform background in the image and converting the image into a binary image. The example breaks this process into the following steps:

- Step 1: Read and display an image
- Step 2: Estimate the approximate value of background pixels
- Step 3: View the background approximation as a surface
- Step 4: Create an image with a uniform background
- Step 5: Adjust the contrast in the uniform image
- Step 6: Create a binary version of the image
- Step 7: Determine the number of objects in the image
- Step 8: Examine the label matrix
- Step 9: Display the label matrix as a pseudocolor indexed image
- Step 10: Measure properties of objects in the image
- Step 11: Compute statistics of objects in the image


1. Read and Display an Image

Clear the MATLAB workspace of any variables, close open figure windows, and close all open Image Tools.

```
clear, close all, imtool close all
```

Read and display the intensity image `rice.png`.

```
I = imread('rice.png');  
imshow(I)
```



Intensity Image rice.png

2. Estimate the Value of Background Pixels

In the sample image, the background illumination is brighter in the center of the image than at the bottom. In this step, the example uses a morphological opening operation to estimate the background illumination. An opening is an erosion followed by a dilation, using the same structuring element for both operations. The morphological opening has the effect of removing objects that cannot completely contain the structuring element. For more information about morphological image processing, see Chapter 9, “Morphological Operations.”

The example calls the `imopen` function to perform the morphological opening operation. Note the call to the `strel` function, which creates a disk-shaped structuring element with a radius of 15. To remove the rice grains from the image, the structuring element must be large enough so that it cannot fit entirely inside a single grain of rice.

```
background = imopen(I,strel('disk',15));
```

To see the estimated background image, type

```
figure, imshow(background)
```

3. View the Background Approximation as a Surface


Use the `surf` command to create a surface display of the background approximation background. The `surf` command creates colored parametric surfaces that enable you to view mathematical functions over a rectangular region. The `surf` function requires data of class `double`, however, so you first need to convert `background` using the `double` command.

```
figure, surf(double(background(1:8:end,1:8:end)),zlim([0 255]));  
set(gca,'ydir','reverse');
```

The example uses MATLAB indexing syntax to view only 1 out of 8 pixels in each direction; otherwise the surface plot would be too dense. The example also sets the scale of the plot to better match the range of the `uint8` data and reverses the y -axis of the display to provide a better view of the data (the pixels at the bottom of the image appear at the front of the surface plot).

In the surface display, `[0, 0]` represents the origin, or upper left corner of the image. The highest part of the curve indicates that the highest pixel values of `background` (and consequently `rice.png`) occur near the middle rows of the image. The lowest pixel values occur at the bottom of the image and are represented in the surface plot by the lowest part of the curve.

The surface plot is a Handle Graphics® object. You can use object properties to fine-tune its appearance. For information on working with MATLAB graphics, see the MATLAB graphics documentation.



4. Create an Image with a Uniform Background

To create a more uniform background, subtract the background image, `background`, from the original image, `I`.

```
I2 = imsubtract(I,background);
```

Because subtraction, like many MATLAB mathematical operations, is only supported for data of class `double`, you must use the Image Processing Toolbox image arithmetic `imsubtract` function.

Display the image with its more uniform background.

```
figure, imshow(I2)
```




Image with Uniform Background

5. Adjust the Contrast in the Processed Image

After subtraction, the image has a uniform background but is now a bit too dark. Use `imadjust` to adjust the contrast of the image.

```
I3 = imadjust(I2);
```

`imadjust` increases the contrast of the image by saturating 1% of the data at both low and high intensities of `I2` and by stretching the intensity values to fill the `uint8` dynamic range. See the reference page for `imadjust` for more information.

Display the adjusted image `I3`.

```
figure, imshow(I3);
```





Image After Intensity Adjustment

6. Create a Binary Version of the Image

Create a binary version of the image by using thresholding. The function `graythresh` automatically computes an appropriate threshold to use to convert the intensity image to binary. The `im2bw` function performs the conversion.

```
level = graythresh(I3);  
bw = im2bw(I3,level);  
figure, imshow(bw)
```



Binary Version of the Image

The binary image `bw` returned by `im2bw` is of class `logical`, as can be seen in this call to `whos`. The Image Processing Toolbox uses logical arrays to represent binary images. For more information, see “Binary Images” on page 2-4.

```
whos
```

MATLAB responds with

Name	Size	Bytes	Class
I	256x256	65536	uint8 array
I2	256x256	65536	uint8 array
I3	256x256	65536	uint8 array
background	256x256	65536	uint8 array
bw	256x256	65536	logical array
level	1x1	8	double array

```
Grand total is 327681 elements using 327688 bytes
```

7. Determine the Number of Objects in the Image

After converting the image to a binary image, you can use the `bwlabel` function to determine the number of grains of rice in the image. The `bwlabel` function labels all the components in the binary image `bw` and returns the number of components it finds in the image in the output value, `numObjects`.

```
[labeled,numObjects] = bwlabel(bw,4);
```

```
numObjects
```

```
ans =
```

```
101
```

The accuracy of the results depends on a number of factors, including

- The size of the objects
- Whether or not any objects are touching (in which case they might be labeled as one object)
- The accuracy of the approximated background
- The connectivity selected. The parameter 4, passed to the `bwlabel` function, means that pixels must touch along an edge to be considered connected. For more information about the connectivity of objects, see “Pixel Connectivity” on page 9-23.

8. Examine the Label Matrix

To better understand the label matrix returned by the `bwlabel` function, this step explores the pixel values in the image. There are several ways to get a closeup view of pixel values. For example, you can use `imcrop` to select a small portion of the image. Another way is to use toolbox Pixel Region tool to examine pixel values. Display the label matrix, using `imshow`,

```
figure, imshow(labeled);
```


Start the Pixel Region tool.

```
impixelregion
```

By default, it automatically associates itself with the image in the current figure. The Pixel Region tool draws a rectangle, called the *pixel region rectangle*, in the center of the visible part of the image. This rectangle defines

which pixels are displayed in the Pixel Region tool. As you move the rectangle, the Pixel Region tool updates the pixel values displayed in the window. For more information about using the toolbox modular interactive tools, see Chapter 4, “Building GUIs with Modular Tools.”

The following figure shows the Image Viewer with the Pixel Region rectangle positioned over the edges of two rice grains. Note how all the pixels in the rice grains have the values assigned by the `bwlabel` function and the background pixels have the value 0 (zero).



Examining the Label Matrix with the Pixel Region Tool

9. Display the Label Matrix as a Pseudocolor Indexed Image

A good way to view a label matrix is to display it as a pseudocolor indexed image. In the pseudocolor image, the number that identifies each object in the label matrix maps to a different color in the associated colormap matrix. The colors in the image make objects easier to distinguish.

To view a label matrix in this way, use the `label2rgb` function. Using this function, you can specify the colormap, the background color, and how objects in the label matrix map to colors in the colormap.

```
pseudo_color = label2rgb(labeled, @spring, 'c', 'shuffle');
imshow(pseudo_color);
```



Label Matrix Displayed as Pseudocolor Image

10. Measure Object Properties in the Image

The `regionprops` command measures object or region properties in an image and returns them in a structure array. When applied to an image with labeled components, it creates one structure element for each component.

This example uses `regionprops` to create a structure array containing some basic properties for `labeled`. When you set the `properties` parameter to `'basic'`, the `regionprops` function returns three commonly used measurements: area, centroid (or center of mass), and bounding box. The bounding box represents the smallest rectangle that can contain a region, or in this case, a grain of rice.

```
graindata = regionprops(labeled,'basic')
```

MATLAB responds with

```
graindata =

101x1 struct array with fields:
    Area
    Centroid
```

BoundingBox

To find the area of the 51st labeled component, access the Area field in the 51st element in the graindata structure array. Note that structure field names are case sensitive.

```
graindata(51).Area
```

returns the following results

```
ans =
```

```
140
```

To find the smallest possible bounding box and the centroid (center of mass) for the same component, use this code:

```
graindata(51).BoundingBox, graindata(51).Centroid
```

```
ans =
```

```
107.5000    4.5000    13.0000    20.0000
```

```
ans =
```

```
114.5000    15.4500
```

11. Compute Statistical Properties of Objects in the Image

Now use MATLAB functions to calculate some statistical properties of the thresholded objects. First use max to find the size of the largest grain. (In this example, the largest grain is actually two grains of rice that are touching.)

```
max([graindata.Area])
```

returns

```
ans =
```

```
404
```

Use the find command to return the component label of the grain of rice with this area.

```
biggrain = find([graindata.Area]==404)
```

returns

```
biggrain =
```

```
59
```

Find the mean of all the rice grain sizes.

```
mean([graindata.Area])
```


returns

```
ans =
```

```
175.0396
```

Make a histogram containing 20 bins that show the distribution of rice grain sizes. The histogram shows that the most common sizes for rice grains in this image are in the range of 150 to 250 pixels.

```
hist([graindata.Area],20)
```



Getting Help

For more information about the topics covered in these exercises, read the tutorial chapters that make up the remainder of this documentation. For reference information about any of the Image Processing Toolbox functions, see the online “Function Reference”, which complements the M-file help that is displayed in the MATLAB command window when you type

```
help functionname
```

For example,

```
help imtool
```

Online Help

The Image Processing Toolbox User’s Guide documentation is available online in both HTML and PDF formats. To access the HTML help, select **Help** from the menu bar of the MATLAB desktop. In the Help Navigator pane, click the **Contents** tab and expand the **Image Processing Toolbox** topic in the list.

To access the PDF help, click **Image Processing Toolbox** in the **Contents** tab of the Help browser and go to the link under “Printable Documentation (PDF).” (Note that to view the PDF help, you must have Adobe’s Acrobat Reader installed.)

Terminology

At the beginning of each chapter are glossaries of words you need to know to understand the information in the chapter. These tables clarify how we use terms that may be used in several different ways in image processing literature. For example:

- In the field of image processing, one word is sometimes used to describe more than one concept. For example the *resolution* of an image can describe the height and width of an image as a quantity of pixels in each direction, or it can describe the number of pixels per linear measure, such as 100 dots per inch.
- In the field of image processing, the same concepts are sometimes described by different terminology. For example, a *grayscale* image can also be called an *intensity* image.

Image Processing Demos

The Image Processing Toolbox is supported by a full complement of demo applications. These are very useful as templates for your own end-user applications, or for seeing how to use and combine your toolbox functions for powerful image analysis and enhancement.

To view all the Image Processing Toolbox demos, call the `iptdemos` function. This displays an HTML page in the MATLAB Help browser that lists all the Image Processing Toolbox demos.

You can also view this page by starting the MATLAB Help browser and clicking the **Demos** tab in the Help Navigator pane. From the list of products with demos, select the Image Processing Toolbox.

The toolbox demos are located under the subdirectory

```
matlabroot\toolbox\images\imdemos
```

where `matlabroot` represents your MATLAB installation directory.

MATLAB Newsgroup

If you read newsgroups on the Internet, you might be interested in the MATLAB newsgroup (`comp.soft-sys.matlab`). This newsgroup gives you access to an active MATLAB user community. It is an excellent way to seek advice and to share algorithms, sample code, and M-files with other MATLAB users.

Image Credits

This table lists the copyright owners of the images used in the Image Processing Toolbox documentation.

Image	Source
cameraman	Copyright Massachusetts Institute of Technology. Used with permission.
cell	Cancer cell from a rat's prostate, courtesy of Alan W. Partin, M.D., Ph.D., Johns Hopkins University School of Medicine.
circuit	Micrograph of 16-bit A/D converter circuit, courtesy of Steve Decker and Shujaat Nadeem, MIT, 1993.
concordaerial and westconcordaerial	Visible color aerial photographs courtesy of mPower3/Emerge.
concordorthophoto and westconcordorthophoto	Orthoregistered photographs courtesy of Massachusetts Executive Office of Environmental Affairs, MassGIS.
forest	Photograph of Carmanah Ancient Forest, British Columbia, Canada, courtesy of Susan Cohen.
LAN files	Permission to use Landsat™ data sets provided by Space Imaging, LLC, Denver, Colorado.
liftingbody	Picture of M2-F1 lifting body in tow, courtesy of NASA (Image number E-10962).
m83	M83 spiral galaxy astronomical image courtesy of Anglo-Australian Observatory, photography by David Malin.

Image	Source
moon	Copyright Michael Myers. Used with permission.
saturn	Voyager 2 image, 1981-08-24, NASA catalog #PIA01364.
solarspectra	Courtesy of Ann Walker. Used with permission.
tissue	Courtesy of Alan W. Partin, M.D., PhD., Johns Hopkins University School of Medicine.
trees	<i>Trees with a View</i> , watercolor and ink on paper, copyright Susan Cohen. Used with permission.

Introduction

This chapter introduces you to the fundamentals of image processing using MATLAB and the Image Processing Toolbox.

Images in MATLAB and the Image Processing Toolbox (p. 2-2)

How images are represented in MATLAB and the Image Processing Toolbox

Image Types in the Toolbox (p. 2-3)

Fundamental image types supported by the Image Processing Toolbox

Converting Between Image Types (p. 2-11)

Converting between the image types

Converting Between Image Classes (p. 2-13)

Converting image data from one class to another

Multiframe Image Arrays (p. 2-14)

Working with multiframe data

Reading and Writing Image Data (p. 2-15)

Reading data into the MATLAB workspace, writing data to graphics files, and getting information about image files

Working with DICOM Files (p. 2-21)

Reading and writing DICOM files

Image Arithmetic (p. 2-28)

Adding, subtracting, multiplying, and dividing images

Coordinate Systems (p. 2-30)

Pixel versus spatial coordinates

Images in MATLAB and the Image Processing Toolbox

The basic data structure in MATLAB is the *array*, an ordered set of real or complex elements. This object is naturally suited to the representation of *images*, real-valued ordered sets of color or intensity data.

MATLAB stores most images as two-dimensional arrays (i.e., matrices), in which each element of the matrix corresponds to a single *pixel* in the displayed image. (Pixel is derived from *picture element* and usually denotes a single dot on a computer display.) For example, an image composed of 200 rows and 300 columns of different colored dots would be stored in MATLAB as a 200-by-300 matrix. Some images, such as truecolor images, require a three-dimensional array, where the first plane in the third dimension represents the red pixel intensities, the second plane represents the green pixel intensities, and the third plane represents the blue pixel intensities.

This convention makes working with images in MATLAB similar to working with any other type of matrix data, and makes the full power of MATLAB available for image processing applications. For example, you can select a single pixel from an image matrix using normal matrix subscripting.

```
I(2,15)
```

This command returns the value of the pixel at row 2, column 15 of the image I.

Note MATLAB expresses pixel coordinates by row and column (r,c). This is the reverse of spatial coordinates that are expressed as width and height (x,y). In the reference pages for toolbox functions, when the syntax for a function uses r and c, it refers to the pixel coordinate system. When the syntax uses x and y, it refers to the spatial coordinate system.

Image Types in the Toolbox


The Image Processing Toolbox defines four basic types of images, summarized in the following table. These image types determine the way MATLAB interprets data matrix elements as pixel colors. The sections that follow provide more information about each image type. See also “Converting Between Image Types” on page 2-11.

Image Type	Interpretation
Binary	Logical array containing only 0's and 1's, interpreted as black and white, respectively. Also known as a <i>bilevel</i> image.
Indexed	<p>Array of class <code>logical</code>, <code>uint8</code>, <code>uint16</code>, <code>single</code>, or <code>double</code> whose pixel values are direct indices into a colormap. The colormap is an m-by-3 array of class <code>double</code>. Also known as a <i>pseudocolor</i> image.</p> <p>Note: For <code>single</code> or <code>double</code> arrays, integer values range from $[1, p]$. For <code>logical</code>, <code>uint8</code>, or <code>uint16</code> arrays, values range from $[0, p-1]$</p>
Intensity	<p>Array of class <code>uint8</code>, <code>uint16</code>, <code>int16</code>, <code>single</code>, or <code>double</code> whose pixel values specify intensity values. Also known as a <i>grayscale</i> image</p> <p>Note: For <code>single</code> or <code>double</code> arrays, values range from $[0, 1]$. For <code>uint8</code>, values range from $[0, 255]$. For <code>uint16</code>, values range from $[0, 65535]$. For <code>int16</code>, values range from $[-32768, 32767]$.</p>
Truecolor	<p>m-by-n-by-3 array of class <code>uint8</code>, <code>uint16</code>, <code>single</code>, or <code>double</code> whose pixel values specify intensity values. Also known as an <i>RGB</i> image.</p> <p>Note: For <code>single</code> or <code>double</code> arrays, values range from $[0, 1]$. For <code>uint8</code>, values range from $[0, 255]$. For <code>uint16</code>, values range from $[0, 65535]$.</p>

Binary Images

In a binary image, also known as a bilevel image, each pixel assumes one of only two discrete values: 1 or 0. A binary image is stored as a logical array.

The following figure shows a binary image with a close-up view of some of the pixel values.



Pixel Values in a Binary Image

Indexed Images

An indexed image consists of an array, called `X` in this documentation, and a colormap matrix, called `map`. The pixel values in the array are direct indices into a colormap.

The colormap matrix is an `m`-by-3 array of class `double` containing floating-point values in the range `[0,1]`. Each row of `map` specifies the red, green, and blue components of a single color. An indexed image uses direct mapping of pixel values to colormap values. The color of each image pixel is determined by using the corresponding value of `X` as an index into `map`.

A colormap is often stored with an indexed image and is automatically loaded with the image when you use the `imread` function. After you read the image and the colormap into the MATLAB workspace as separate variables, you must keep track of the association between the image and colormap. However, you are not limited to using the default colormap—you can use any colormap that you choose.

The relationship between the values in the image matrix and the colormap depends on the class of the image matrix. If the image matrix is of class `single` or `double`, it normally contains integer values 1 through p , where p is the length of the colormap. the value 1 points to the first row in the colormap, the value 2 points to the second row, and so on. If the image matrix is of class `logical`, `uint8` or `uint16`, the value 0 points to the first row in the colormap, the value 1 points to the second row, and so on.

The following figure illustrates the structure of an indexed image. In the figure, the image matrix is of class `double`, so the value 5 points to the fifth row of the colormap.




Image Courtesy of Susan Cohen

Pixel Values Index to Colormap Entries in Indexed Images


Intensity Images

An intensity image, also known as a grayscale image, is a data matrix, I , whose values represent intensities within some range. MATLAB stores an intensity image as a individual matrix, with each element of the matrix corresponding to one image pixel. The matrix can be of class `uint8`, `uint16`, `int16`, `single`, or

double. While intensity images are rarely saved with a colormap, MATLAB uses a colormap to display them.

For a matrix of class single or double, using the default grayscale colormap, the intensity 0 represents black and the intensity 1 represents white. For a matrix of type uint8, uint16, or int16, the intensity `intmin(class(I))` represents black and the intensity `intmax(class(I))` represents white.

The figure below depicts an intensity image of class double.



Pixel Values in an Intensity Image Define Gray Levels


Truecolor Images

A truecolor image, also known as an RGB image, is stored in MATLAB as an m -by- n -by-3 data array that defines red, green, and blue color components for each individual pixel. Truecolor images do not use a colormap. The color of each pixel is determined by the combination of the red, green, and blue intensities stored in each color plane at the pixel's location. Graphics file formats store truecolor images as 24-bit images, where the red, green, and blue components

are 8 bits each. This yields a potential of 16 million colors. The precision with which a real-life image can be replicated has led to the commonly used term truecolor image.

A truecolor array can be of class `uint8`, `uint16`, `single`, or `double`. In a truecolor array of class `single` or `double`, each color component is a value between 0 and 1. A pixel whose color components are (0,0,0) is displayed as black, and a pixel whose color components are (1,1,1) is displayed as white. The three color components for each pixel are stored along the third dimension of the data array. For example, the red, green, and blue color components of the pixel (10,5) are stored in `RGB(10,5,1)`, `RGB(10,5,2)`, and `RGB(10,5,3)`, respectively.

The following figure depicts a truecolor image of class `double`.




The Color Planes of an RGB Image

To determine the color of the pixel at (2,3), you would look at the RGB triplet stored in (2,3,1:3). Suppose (2,3,1) contains the value 0.5176, (2,3,2) contains 0.1608, and (2,3,3) contains 0.0627. The color for the pixel at (2,3) is

0.5176 0.1608 0.0627

To further illustrate the concept of the three separate color planes used in an RGB image, the code sample below creates a simple RGB image containing uninterrupted areas of red, green, and blue, and then creates one image for each of its separate color planes (red, green, and blue). It displays each color plane image separately, and also displays the original image.

```
RGB=reshape(ones(64,1)*reshape(jet(64),1,192),[64,64,3]);  
R=RGB(:,:,1);  
G=RGB(:,:,2);  
B=RGB(:,:,3);  
imshow(R)  
figure, imshow(G)  
figure, imshow(B)  
figure, imshow(RGB)
```



The Separated Color Planes of an RGB Image

Notice that each separated color plane in the figure contains an area of white. The white corresponds to the highest values (purest shades) of each separate color. For example, in the Red Plane image, the white represents the highest concentration of pure red values. As red becomes mixed with green or blue, gray pixels appear. The black region in the image shows pixel values that contain no red values, i.e., $R = 0$.

Converting Between Image Types

You might need to convert an image from one type to another. For example, if you want to filter a color image that is stored as an indexed image, you should first convert it to truecolor format. When you apply the filter to the truecolor image, MATLAB filters the intensity values in the image, as is appropriate. If you attempt to filter the indexed image, MATLAB simply applies the filter to the indices in the indexed image matrix, and the results might not be meaningful.

Note When you convert an image from one format to another, the resulting image might look different from the original. For example, if you convert a color indexed image to an intensity image, the resulting image is grayscale, not color.

The following table lists all the image type conversion functions in the Image Processing Toolbox.

Function	Description
dither	Create a binary image from a grayscale intensity image by dithering or create an indexed image from a truecolor image by dithering
gray2ind	Create an indexed image from a grayscale intensity image
grayslice	Create an indexed image from a grayscale intensity image by thresholding
im2bw	Create a binary image from a grayscale intensity image, indexed image, or truecolor image, based on a luminance threshold
ind2gray	Create a grayscale intensity image from an indexed image
ind2rgb	Create a truecolor image from an indexed image

Function	Description
mat2gray	Create a grayscale intensity image from data in a matrix, by scaling the data
rgb2gray	Create a grayscale intensity image from a truecolor image
rgb2ind	Create an indexed image from a truecolor image

You can also perform certain conversions just using MATLAB syntax. For example, you can convert an intensity image to truecolor format by concatenating three copies of the original matrix along the third dimension.

```
RGB = cat(3,I,I,I);
```

The resulting truecolor image has identical matrices for the red, green, and blue planes, so the image displays as shades of gray.

In addition to these standard conversion tools, there are some functions that return a different image type as part of the operation they perform. For example, the region-of-interest routines each return a binary image that you can use to mask an indexed or intensity image for filtering or for other operations.

Color Space Conversions

The Image Processing Toolbox represents colors as RGB values, either directly (in a truecolor image) or indirectly (in an indexed image). However, there are other methods for representing colors. For example, a color can be represented by its hue, saturation, and value components (HSV). Different methods for representing colors are called *color spaces*.

The toolbox provides a set of routines for converting between color spaces. The image processing functions themselves assume all color data is RGB, but you can process an image that uses a different color space by first converting it to RGB, and then converting the processed image back to the original color space. For more information about color space conversion routines, see Chapter 13, “Color.”

Converting Between Image Classes

You can convert `uint8` and `uint16` image data to `double` using the MATLAB `double` function. However, converting between classes changes the way MATLAB and the toolbox interpret the image data. If you want the resulting array to be interpreted properly as image data, you need to rescale or offset the data when you convert it.

For easier conversion of classes, use one of these toolbox functions: `im2uint8`, `im2uint16`, `im2int16`, `im2single`, or `im2double`. These functions automatically handle the rescaling and offsetting of the original data of any image class. For example, this command converts a double-precision RGB image with data in the range `[0,1]` to a `uint8` RGB image with data in the range `[0,255]`.

```
RGB2 = im2uint8(RGB1);
```

Losing Information in Conversions

When you convert to a class that uses fewer bits to represent numbers, you generally lose some of the information in your image. For example, a `uint16` intensity image is capable of storing up to 65,536 distinct shades of gray, but a `uint8` intensity image can store only 256 distinct shades of gray. When you convert a `uint16` intensity image to a `uint8` intensity image, `im2uint8` *quantizes* the gray shades in the original image. In other words, all values from 0 to 127 in the original image become 0 in the `uint8` image, values from 128 to 385 all become 1, and so on.

Converting Indexed Images

It is not always possible to convert an indexed image from one storage class to another. In an indexed image, the image matrix contains only indices into a colormap, rather than the color data itself, so no quantization of the color data is possible during the conversion.

For example, a `uint16` or `double` indexed image with 300 colors cannot be converted to `uint8`, because `uint8` arrays have only 256 distinct values. If you want to perform this conversion, you must first reduce the number of the colors in the image using the `imapprox` function. This function performs the quantization on the colors in the colormap, to reduce the number of distinct colors in the image. See “Reducing Colors in an Indexed Image” on page 13-12 for more information.

Multiframe Image Arrays

For some applications, you might need to work with collections of images related by time or view, such as magnetic resonance imaging (MRI) slices or movie frames.

The Image Processing Toolbox provides support for storing multiple images in the same array. Each separate image is called a *frame*. If an array holds multiple frames, they are concatenated along the fourth dimension. For example, an array with five 480-by-640 truecolor images would be 480-by-640-by-3-by-5. A similar multiframe intensity or indexed image would be 480-by-640-by-1-by-5.

Use the `cat` command to store separate images in one multiframe array. For example, if you have a group of images `A1`, `A2`, `A3`, `A4`, and `A5`, you can store them in a single array using

```
A = cat(4,A1,A2,A3,A4,A5)
```

You can also extract frames from a multiframe image. For example, if you have a multiframe image `MULTI`, this command extracts the third frame.

```
FRM3 = MULTI(:,:, :,3)
```

Note that, in a multiframe image array, each image must be the same size and have the same number of planes. In a multiframe indexed image, each image must also use the same colormap.

Multiframe Support Limitations

Many of the functions in the toolbox operate only on the first two or first three dimensions. You can still use four-dimensional arrays with these functions, but you must process each frame individually. For example, this call displays the seventh frame in the array `MULTI`.

```
imshow(MULTI(:,:, :,7))
```

If you pass an array to a function and the array has more dimensions than the function is designed to operate on, your results can be unpredictable. In some cases, the function simply processes the first frame of the array, but in other cases the operation does not produce meaningful results.

See the reference pages for information about how individual functions work with the dimensions of an image array.

Reading and Writing Image Data

This section describes how to read and write image data. Topics include

- Reading data stored in many standard graphics file formats
- Writing data to files in many standard graphics file formats
- Querying graphics image files for information stored in header fields
- Converting images between graphics file formats

For information about reading and writing data in Digital Imaging and Communications in Medicine (DICOM) file format, see “Working with DICOM Files” on page 2-21.

Reading a Graphics Image

The `imread` function reads an image from any supported graphics image file format, in any of the supported bit depths. Most image file formats use 8 bits to store pixel values. When these are read into memory, MATLAB stores them as class `uint8`. For file formats that support 16-bit data, PNG and TIFF, MATLAB stores the images as class `uint16`.

Note For indexed images, `imread` always reads the colormap into a matrix of class `double`, even though the image array itself may be of class `uint8` or `uint16`.

For example, this code reads a truecolor image into the MATLAB workspace as the variable `RGB`.

```
RGB = imread('football.jpg');
```

This code reads an indexed image with its associated colormap into the MATLAB workspace in two separate variables.

```
[X,map] = imread('trees.tif');
```

In these examples, `imread` infers the file format to use from the contents of the file. You can also specify the file format as an argument to `imread`. MATLAB supports many common graphics file formats, such as Microsoft Windows Bitmap (BMP), Graphics Interchange Format (GIF), Joint Photographic

Experts Group (JPEG), Portable Network Graphics (PNG), and Tagged Image File Format (TIFF) formats. For the latest information concerning the bit depths and/or image formats supported, see the reference page for the `imread` and `imformats` functions.

Reading Multiple Images from a Graphics File

MATLAB supports several graphics file formats, such as HDF and TIFF, that can contain multiple images. By default, `imread` imports only the first image from a file. To import additional images from the file, use the syntax supported by the file format.

For example, when used with TIFF files, you can use an index value with `imread` that identifies the image in the file you want to import. This example reads a series of 27 images from a TIFF file and stores the images in a four-dimensional array. You can use `imfinfo` to determine how many images are stored in the file.

```
mri = uint8(zeros(128,128,1,27)); % preallocate 4-D array

for frame=1:27
    [mri(:,:,,frame),map] = imread('mri.tif',frame);
end
```

When a file contains multiple images that are related in some way, such as a time sequence, you can store the images in MATLAB as a 4-D array. All the images must be the same size. For more information, see “Multiframe Image Arrays” on page 2-14.

Writing a Graphics Image

The function `imwrite` writes an image to a graphics file in one of the supported formats. The most basic syntax for `imwrite` takes the image variable name and a filename. If you include an extension in the filename, MATLAB infers the desired file format from it. (For more information, see the reference entry for the `imwrite` function.)

This example loads the indexed image `X` from a MAT-file, `clown.mat`, that contains the data matrix and the associated colormap and then writes the image to a BMP file.

```
load clown
whos
```

Name	Size	Bytes	Class
X	200x320	512000	double array
caption	2x1	4	char array
map	81x3	1944	double array

Grand total is 64245 elements using 513948 bytes

```
imwrite(X,map,'clown.bmp')
```

Specifying Additional Format-Specific Parameters

When using `imwrite` with some graphics formats, you can specify additional parameters. For example, with PNG files, you can specify the bit depth as an additional parameter. This example writes an intensity image `I` to a 4-bit PNG file.

```
imwrite(I,'clown.png','BitDepth',4);
```

This example writes an image `A` to a JPEG file, using an additional parameter to specify the compression quality parameter.

```
imwrite(A, 'myfile.jpg', 'Quality', 100);
```

For more information about the additional parameters associated with certain graphics formats, see the reference pages for `imwrite`.

Reading and Writing Binary Images in 1-Bit Format

In certain file formats, a binary image can be stored in a 1-bit format. If the file format supports it, MATLAB writes binary images as 1-bit images by default. When you read in a binary image in 1-bit format, MATLAB represents it in the workspace as a logical array.

This example reads in a binary image and writes it as a TIFF file. Because the TIFF format supports 1-bit images, the file is written to disk in 1-bit format.

```
BW = imread('text.png');
imwrite(BW,'test.tif');
```

To verify the bit depth of `test.tif`, call `imfinfo` and check the `BitDepth` field.

```
info = imfinfo('test.tif');
```

```
info.BitDepth
ans =

     1
```

Note When writing binary files, MATLAB sets the `ColorType` field to 'grayscale'.

Determining the Storage Class of the Output File

`imwrite` uses the following rules to determine the storage class used in the output image.

Storage Class of Image	Storage Class of Output Image File
logical	If the output image file format specified supports 1-bit images, <code>imwrite</code> creates a 1-bit image file. If the output image file format specified does not support 1-bit images, such as JPEG, <code>imwrite</code> converts the image to a class <code>uint8</code> intensity image.
uint8	If the output image file format specified supports unsigned 8-bit images, <code>imwrite</code> creates an unsigned 8-bit image file.
uint16	If the output image file format specified support unsigned 16-bit images (PNG or TIFF), <code>imwrite</code> creates an unsigned 16-bit image file. If the output image file format specified does not support 16-bit images, <code>imwrite</code> scales the image data to class <code>uint8</code> and creates an 8-bit image file.
int16	Partially supported; depends on file format.

Storage Class of Image	Storage Class of Output Image File
single	Partially supported; depends on file format.
double	MATLAB scales the image data to uint8 and creates an 8-bit image file, because most image file formats use 8 bits.

Querying a Graphics File

Note You can also get information interactively about an image displayed in the Image Tool — see “Getting Information About an Image” on page 3-29.

The `imfinfo` function enables you to obtain information about graphics files that are in any of the formats supported by the toolbox. The information you obtain depends on the type of file, but it always includes at least the following:

- Name of the file
- File format
- Version number of the file format
- File modification date
- File size in bytes
- Image width in pixels
- Image height in pixels
- Number of bits per pixel
- Image type: truecolor (RGB), intensity (grayscale), or indexed

See the reference entry for `imfinfo` for more information.

Converting Graphics File Formats

To change the graphics format of an image, use `imread` to read in the image and then save the image with `imwrite`, specifying the appropriate format.

To illustrate, this example uses the `imread` function to read an image in bitmap (BMP) format into the workspace. The example then writes the bitmap image to a file using Portable Network Graphics (PNG) format.

```
bitmap = imread('mybitmap.bmp','bmp');  
imwrite(bitmap,'mybitmap.png','png');
```

For the specifics of which bit depths are supported for the different graphics formats, and for how to specify the format type when writing an image to file, see the reference entries for `imread` and `imwrite`.

Working with DICOM Files

The Image Processing Toolbox includes support for working with image data in Digital Imaging and Communications in Medicine (DICOM) format. Topics covered include:

- “Reading Image Data from a DICOM File” on page 2-21
- “Reading Metadata from a DICOM File” on page 2-22
- “Writing Image Data or Metadata to a DICOM File” on page 2-23, including information about anonymizing a DICOM file

To see an example that reads both the image data and metadata from a DICOM file, modifies the image data, and writes the modified data to a new DICOM file, see “Example: Creating a New Series” on page 2-25. The example shows how to use the `dicomuid` function to generate a DICOM unique identifier, which you need to create a new series.

Reading Image Data from a DICOM File

To read image data from a DICOM file, use the `dicomread` function. The `dicomread` function reads files that comply with the DICOM specification but can also read certain common noncomplying files.

This example reads an image from a sample DICOM file included with the toolbox.

```
I = dicomread('CT-MON02-16-ankle.dcm');
```

To view the image data, use one of the toolbox image display functions, `imshow` or `imtool`. (Because the image data is signed 16-bit data, you must use the autoscaling syntax with either display function.)

```
imshow(I, 'DisplayRange', [])
```



Reading Metadata from a DICOM File

To read metadata from a DICOM file, use the `dicominfo` function.

The metadata in a DICOM files provides information, such as the size, dimensions, and bit depth of the image. In addition, the DICOM specification defines numerous other metadata fields that describe many other characteristics of the data, such as the modality used to create the data, the equipment settings used to capture the image, and information about the study.

In this example, `dicominfo` reads the metadata from a file, returning the information in a structure where every field in the structure is a specific piece of DICOM metadata.

```
info = dicominfo('CT-MON02-16-ankle.dcm');
```

```
info =
```

```
Filename: [1x47 char]  
FileModDate: '24-Dec-2000 19:54:47'  
FileSize: 525436  
Format: 'DICOM'  
FormatVersion: 3  
Width: 512
```

```

        Height: 512
        BitDepth: 16
        ColorType: 'grayscale'
        SelectedFrames: []
        FileStruct: [1x1 struct]
        StartOfPixelData: 1140
        MetaElementGroupLength: 192
        FileMetaInformationVersion: [2x1 double]
        MediaStorageSOPClassUID: '1.2.840.10008.5.1.4.1.1.7'
        MediaStorageSOPInstanceUID: [1x50 char]
        TransferSyntaxUID: '1.2.840.10008.1.2'
        ImplementationClassUID: '1.2.840.113619.6.5'
        .
        .
        .

```

You can use the metadata structure returned by `dicominfo` to specify the DICOM file you want to read using `dicomread`. For example, you can use this code to read metadata from the sample DICOM file and then pass the metadata to `dicomread` to read the image from the file.

```

info = dicominfo('CT-MON02-16-ankle.dcm');
I = dicomread(info);

```

Writing Image Data or Metadata to a DICOM File

To write image data or metadata to a file in DICOM format, use the `dicomwrite` function. This example writes the image `I` to the DICOM file `ankle.dcm`.

```

dicomwrite(I, 'h:\matlab\tmp\ankle.dcm')

```

When you write image data to a DICOM file, `dicomwrite` includes the minimum set of metadata fields required by the type of DICOM information object (IOD) you are creating. `dicomwrite` supports three types of DICOM IODs with full validation and supports many other types (e.g. X-ray, radiotherapy, nuclear medicine) without validation.

- Secondary capture (default)
- Magnetic resonance
- Computed tomography

You can also specify the metadata you want to write to the file by passing to `dicomwrite` an existing DICOM metadata structure that you retrieved using `dicominfo`.

```
info = dicominfo('CT-MON02-16-ankle.dcm');  
I = dicomread(info);  
dicomwrite(I, 'h:\matlab\tmp\ankle.dcm', info)
```

In this case, the `dicomwrite` function writes the relevant information in the metadata structure `info` to the new DICOM file.

Note that the metadata written to the file is not identical to the metadata in the `info` structure. When writing metadata to a file, there are certain fields that `dicomwrite` must update. For example, `dicomwrite` must update the file modification date in the new file. To illustrate, compare the instance ID in the original metadata with the ID in the new file.

```
info.SOPInstanceUID  
ans =  
  
1.2.840.113619.2.1.2411.1031152382.365.1.736169244
```

Using `dicominfo`, read the metadata from the newly written file and check the file modification date.

```
info2 = dicominfo('h:\matlab\tmp\ankle.dcm');  
  
info2.SOPInstanceUID  
  
ans =  
  
1.2.841.113411.2.1.2411.10311244477.365.1.63874544
```

Anonymizing a DICOM File

When using a DICOM file as part of a training set, blinded study, or a presentation, you might want to remove confidential patient information. To do this, use the `dicomanon` function.

The `dicomanon` function creates a new series and study values, changes some of the metadata, and then writes the file. For example, you could replace steps 4, 5, and 6 in the example in “Example: Creating a New Series” on page 2-25 with a call to the `dicomanon` function.

Example: Creating a New Series

When writing a modified image to a DICOM file, you might want to make the modified image the start of a new series. In the DICOM standard, images can be organized into series. When you write an image with metadata to a DICOM file, `dicomwrite` puts the image in the same series by default. To create a new series, you must assign a new DICOM unique identifier to the `SeriesInstanceUID` metadata field. This example illustrates this process:

- 1 Read an image from a DICOM file into the MATLAB workspace.

```
I = dicomread('CT-MON02-16-ankle.dcm');
```

To view the image, use either of the toolbox display functions, `imshow` or `imtool`. Because the DICOM image data is signed 16-bit data, you must use the autoscaling syntax.

```
imtool(I, 'DisplayRange', [])
```



- 2 Read the metadata from the same DICOM file.

```
info = dicominfo('CT-MON02-16-ankle.dcm');
```

To identify which series an image belongs to, view the value of the `SeriesInstanceUID` field.

```
info.SeriesInstanceUID  
  
ans =  
  
1.2.840.113619.2.1.2411.1031152382.365.736169244
```

- 3** You typically only start a new DICOM series when you modify the image in some way. This example removes all the text from the image.

The example finds the maximum and minimum values of all pixels in the image. The pixels that form the white text characters are set to the maximum pixel value.

```
max(I(:))  
ans =  
  
4080
```

```
min(I(:))  
  
ans =  
  
32
```

To remove these text characters, the example sets all pixels with the maximum value to the minimum value.

```
Imodified = I;  
Imodified(Imodified == 4080) = 32;
```

View the processed image.

```
imshow(Imodified)
```



- 4 Generate a new DICOM unique identifier (UID) using the `dicomuid` function. You need a new UID to write the modified image as a new series.

```
uid = dicomuid
```

```
uid =
```

```
1.3.6.1.4.1.9590.100.1.1.56461980611264497732341403390561061497
```

`dicomuid` is guaranteed to generate a unique UID.

- 5 Set the value of the `SeriesInstanceUID` field in the metadata associated with the original DICOM file to the generated value.

```
info.SeriesInstanceUID = uid;
```

- 6 Write the modified image to a new DICOM file, specifying the modified metadata structure, `info`, as an argument. Because you set the `SeriesInstanceUID` value, the image you write is part of a new series.

```
dicomwrite(Imodified, 'ankle_newseries.dcm', info);
```

To verify this operation, view the image and the `SeriesInstanceUID` metadata field in the new file.

Image Arithmetic

Image arithmetic is the implementation of standard arithmetic operations, such as addition, subtraction, multiplication, and division, on images. Image arithmetic has many uses in image processing both as a preliminary step in more complex operations and by itself. For example, image subtraction can be used to detect differences between two or more images of the same scene or object.

You can do image arithmetic using the MATLAB arithmetic operators. The Image Processing Toolbox also includes a set of functions that implement arithmetic operations for all numeric, nonsparse data types. The toolbox arithmetic functions accept any numeric data type, including `uint8`, `uint16`, and `double`, and return the result image in the same format. The functions perform the operations in double precision, on an element-by-element basis, but do not convert images to double-precision values in the MATLAB workspace. Overflow is handled automatically. The functions saturate return values to fit the data type. For details, see “Image Arithmetic Saturation Rules” on page 2-28.

Note On Intel architecture processors, the image arithmetic functions can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating their execution time. IPPL is only activated, however, when the data passed to these functions is of specific classes. See the reference pages for the individual arithmetic functions for more information.

Image Arithmetic Saturation Rules

The results of integer arithmetic can easily overflow the data type allotted for storage. For example, the maximum value you can store in `uint8` data is 255. Arithmetic operations can also result in fractional values, which cannot be represented using integer arrays.

MATLAB arithmetic operators and the Image Processing Toolbox arithmetic functions use these rules for integer arithmetic:

- Values that exceed the range of the integer type are saturated to that range.
- Fractional values are rounded.

For example, if the data type is `uint8`, results greater than 255 (including `Inf`) are set to 255. The following table lists some additional examples.

Result	Class	Truncated Value
300	<code>uint8</code>	255
-45	<code>uint8</code>	0
10.5	<code>uint8</code>	11

Nesting Calls to Image Arithmetic Functions

You can use the image arithmetic functions in combination to perform a series of operations. For example, to calculate the average of two images,

$$C = \frac{A + B}{2}$$

You could enter

```
I = imread('rice.png');
I2 = imread('cameraman.tif');
K = imdivide(imadd(I,I2), 2); % not recommended
```

When used with `uint8` or `uint16` data, each arithmetic function rounds and saturates its result before passing it on to the next operation. This can significantly reduce the precision of the calculation. A better way to perform this calculation is to use the `imlincomb` function. `imlincomb` performs all the arithmetic operations in the linear combination in double precision and only rounds and saturates the final result.

```
K = imlincomb(.5,I,.5,I2); % recommended
```


Coordinate Systems

Locations in an image can be expressed in various coordinate systems, depending on context. This section discusses the two main coordinate systems used in the Image Processing Toolbox and the relationship between them. These two coordinate systems are described in

- “Pixel Coordinates”
- “Spatial Coordinates” on page 2-31

Pixel Coordinates

Generally, the most convenient method for expressing locations in an image is to use pixel coordinates. In this coordinate system, the image is treated as a grid of discrete elements, ordered from top to bottom and left to right, as illustrated by the following figure.



The Pixel Coordinate System

For pixel coordinates, the first component r (the row) increases downward, while the second component c (the column) increases to the right. Pixel coordinates are integer values and range between 1 and the length of the row or column.

There is a one-to-one correspondence between pixel coordinates and the coordinates MATLAB uses for matrix subscripting. This correspondence makes the relationship between an image’s data matrix and the way the image


is displayed easy to understand. For example, the data for the pixel in the fifth row, second column is stored in the matrix element (5,2).

Spatial Coordinates

In the pixel coordinate system, a pixel is treated as a discrete unit, uniquely identified by a single coordinate pair, such as (5,2). From this perspective, a location such as (5.3,2.2) is not meaningful.

At times, however, it is useful to think of a pixel as a square patch. From this perspective, a location such as (5.3,2.2) *is* meaningful, and is distinct from (5,2). In this spatial coordinate system, locations in an image are positions on a plane, and they are described in terms of x and y (not r and c as in the pixel coordinate system).

The following figure illustrates the spatial coordinate system used for images. Notice that y increases downward.



The Spatial Coordinate System

This spatial coordinate system corresponds closely to the pixel coordinate system in many ways. For example, the spatial coordinates of the center point of any pixel are identical to the pixel coordinates for that pixel.

There are some important differences, however. In pixel coordinates, the upper left corner of an image is (1,1), while in spatial coordinates, this location by default is (0.5,0.5). This difference is due to the pixel coordinate system's being discrete, while the spatial coordinate system is continuous. Also, the upper left corner is always (1,1) in pixel coordinates, but you can specify a nondefault

origin for the spatial coordinate system. See “Using a Nondefault Spatial Coordinate System” on page 2-32 for more information.

Another potentially confusing difference is largely a matter of convention: the order of the horizontal and vertical components is reversed in the notation for these two systems. As mentioned earlier, pixel coordinates are expressed as (r,c) , while spatial coordinates are expressed as (x,y) . In the reference pages, when the syntax for a function uses r and c , it refers to the pixel coordinate system. When the syntax uses x and y , it refers to the spatial coordinate system.

Using a Nondefault Spatial Coordinate System

By default, the spatial coordinates of an image correspond with the pixel coordinates. For example, the center point of the pixel in row 5, column 3 has spatial coordinates $x=3$, $y=5$. (Remember, the order of the coordinates is reversed.) This correspondence simplifies many of the toolbox functions considerably. Several functions primarily work with spatial coordinates rather than pixel coordinates, but as long as you are using the default spatial coordinate system, you can specify locations in pixel coordinates.


In some situations, however, you might want to use a nondefault spatial coordinate system. For example, you could specify that the upper left corner of an image is the point $(19.0,7.5)$, rather than $(0.5,0.5)$. If you call a function that returns coordinates for this image, the coordinates returned will be values in this nondefault spatial coordinate system.

To establish a nondefault spatial coordinate system, you can specify the `XData` and `YData` image properties when you display the image. These properties are two-element vectors that control the range of coordinates spanned by the image. By default, for an image `A`, `XData` is `[1 size(A,2)]`, and `YData` is `[1 size(A,1)]`.

For example, if `A` is a 100 row by 200 column image, the default `XData` is `[1 200]`, and the default `YData` is `[1 100]`. The values in these vectors are actually the coordinates for the center points of the first and last pixels (not the pixel edges), so the actual coordinate range spanned is slightly larger; for instance, if `XData` is `[1 200]`, the x -axis range spanned by the image is `[0.5 200.5]`.

These commands display an image using nondefault XData and YData.

```
A = magic(5);  
x = [19.5 23.5];  
y = [8.0 12.0];  
image(A,'XData',x,'YData',y), axis image, colormap(jet(25))
```



For information about the syntax variations that specify nondefault spatial coordinates, see the reference page for `imshow`.

Displaying and Exploring Images

This chapter describes the image display and exploration tools provided by the Image Processing Toolbox.

Overview (p. 3-2)	Comparison of toolbox display functions
Using <code>imshow</code> to Display Images (p. 3-4)	How to use the <code>imshow</code> display function
Using the Image Tool to Explore Images (p. 3-8)	How to use the Image Tool integrated display and exploration environment
Using Image Tool Navigation Aids (p. 3-16)	Image Tool navigation aids including the Overview tool, panning, and zooming
Getting Information about the Pixels in an Image (p. 3-22)	Image Tool pixel information tools, including the Pixel Region tool and the Pixel Information tool
Getting Information About an Image (p. 3-29)	Image Tool's Image Information tool
Adjusting the Contrast and Brightness of an Image (p. 3-31)	Image Tool's Adjust Contrast tool
Viewing Multiple Images (p. 3-41)	Using <code>imshow</code> and <code>imshow</code> to view multiple images
Displaying Different Image Types (p. 3-45)	Using <code>imshow</code> and <code>imshow</code> with each image type
Special Display Techniques (p. 3-51)	Using the <code>colorbar</code> , <code>montage</code> , and <code>warp</code> functions
Printing Images (p. 3-56)	Print images from <code>imshow</code> and the Image Tool
Setting Toolbox Display Preferences (p. 3-58)	Setting toolbox preferences

Overview

The Image Processing Toolbox includes two display functions, `imshow` and `imshow1`. Both functions work within the Handle Graphics architecture: they create an image object and display it in an axes object contained in a figure object. The toolbox functions automatically set the values of certain figure, axes, and image object properties to control how the image data is displayed — see “Understanding Handle Graphics Object Property Settings” on page 3-2.

`imshow` is the toolbox's fundamental image display function. Use `imshow` when you want to display any of the different image types supported by the toolbox, such as intensity (grayscale), truecolor (RGB), binary, and indexed. For more information, see “Using `imshow` to Display Images” on page 3-4. The `imshow` function is also a key building block for image applications you might want to create using the toolbox modular tools. For more information, see Chapter 4, “Building GUIs with Modular Tools.”

The other toolbox display function, `imshow1`, launches the Image Tool, which presents an integrated environment for displaying images and performing some common image processing tasks. The Image Tool provides all the image display capabilities of `imshow` but also provides access to several other tools for navigating and exploring images, such as scroll bars, the Pixel Region tool, the Image Information tool, and the Adjust Contrast tool. For more information, see “Using the Image Tool to Explore Images” on page 3-8.

In general, using the toolbox functions to display images is preferable to using the MATLAB image display functions `image` and `imagesc`. The toolbox functions are easier to use and are optimized for displaying images.

Understanding Handle Graphics Object Property Settings

When you display an image, `imshow` and `imshow1` set the Handle Graphics properties that control how the image is displayed. The following table lists the relevant properties and their settings for each image type. The table uses standard toolbox terminology to refer to the various image types: `X` represents an indexed image, `I` represents an intensity image, `BW` represents a binary image, and `RGB` represents an RGB (or truecolor) image.

Note Both `imshow` and `imshow` can perform automatic scaling of image data. When called with the syntax `imshow(I, 'DisplayRange', [])`, and similarly for `imshow`, the functions set the axes `CLim` property to `[min(I(:)) max(I(:))]`. `CDataMapping` is always scaled for intensity images, so that the value `min(I(:))` is displayed using the first colormap color, and the value `max(I(:))` is displayed using the last colormap color.

Handle Graphics Property	Indexed Images	Intensity (Grayscale) Images	Binary Images	Truecolor (RGB) Images
<code>CData</code> (Image)	Set to the data in X	Set to the data in I	Set to data in BW	Set to data in RGB
<code>CDataMapping</code> (Image)	Set to 'direct'	Set to 'scaled'	Set to 'direct'	Ignored when <code>CData</code> is 3-D
<code>CLim</code> (Axes)	Does not apply	double: [0 1] uint8: [0 255] uint16: [0 65535]	Set to [0 1]	Ignored when <code>CData</code> is 3-D
<code>Colormap</code> (Figure)	Set to data in map	Set to grayscale colormap	Set to a grayscale colormap whose values range from black to white	Ignored when <code>CData</code> is 3-D

Using imshow to Display Images

You can use the `imshow` function to display an image that has already been imported into the MATLAB workspace or to display an image stored in a graphics file. For example, this code reads an image into the MATLAB workspace and then displays it in a MATLAB figure window.

```
moon = imread('moon.tif');  
imshow(moon);
```

The `imshow` function displays the image in a MATLAB figure window, as shown in the following figure.




Image Displayed in a Figure Window by imshow

The `imshow` filename syntax

```
imshow('moon.tif');
```

can be useful for scanning through images. Note, however, that when you use this syntax, the image data is not stored in the MATLAB workspace. If you want to bring the image into the workspace, you must use the `getimage` function, which retrieves the image data from the current Handle Graphics image object. For example,

```
moon = getimage;
```

assigns the image data from `moon.tif` to the variable `moon` if the figure window in which it is displayed is currently active.

For more information about using `imshow`, see these additional topics.

- “Specifying the Initial Image Magnification” on page 3-5
- “Controlling the Appearance of the Figure” on page 3-6

For more information about using `imshow` to display the various image types supported by the toolbox, see “Displaying Different Image Types” on page 3-45.

Specifying the Initial Image Magnification

By default, `imshow` attempts to display an image in its entirety at 100% magnification (one screen pixel for each image pixel). However, if an image is too large to fit in a figure window on the screen at 100% magnification, `imshow` scales the image to fit onto the screen and issues a warning message.

To override the default initial magnification behavior for a particular call to `imshow`, specify the `InitialMagnification` parameter. For example, to view an image at 150% magnification, use this code.

```
pout = imread('pout.tif');  
imshow(pout, 'InitialMagnification', 150)
```

`imshow` attempts to honor the magnification you specify. However, if the image does not fit on the screen at the specified magnification, `imshow` scales the image to fit and issues a warning message. You can also specify the text string `'fit'` as the initial magnification value. In this case, `imshow` scales the image to fit the current size of the figure window.

You can also change the default initial magnification behavior of `imshow` by setting the `ImshowInitialMagnification` toolbox preference. To make this preference persist between sessions, include the command to set the preference

in your `startup.m` file. To learn more about toolbox preferences, see “Setting the Values of Toolbox Preferences” on page 3-59.

When `imshow` scales an image, it uses interpolation to determine the values for screen pixels that do not directly correspond to elements in the image matrix. For more information, see “Interpolation” on page 5-3.

Controlling the Appearance of the Figure

By default, when `imshow` displays an image in a figure, it surrounds the image with a gray border and does not include a visible axes box. If you want to display an image without the gray border or include a visible axes box with tick labels, you must set toolbox preferences. (For more information about setting toolbox preferences, see “Setting Toolbox Display Preferences” on page 3-58.)

For example, to display an image without a border, set the `ImshowBorder` preference to `'tight'`. By default, this preference is set to `'loose'`, which causes the border to be included. This code sets the preference to suppress the border and then displays an image.

```
iptsetpref('ImshowBorder','tight')
imshow('moon.tif')
```

The following figure shows the same image displayed with and without the border. Note that the image is the same size but the figure window takes up less space on your screen.




Image Displayed With and Without a Border

Using the Image Tool to Explore Images

The Image Tool is an image display tool that also provides access to several other related tools, such as the Pixel Region tool, the Image Information tool, and the Adjust Contrast tool. The Image Tool also provides navigation aids that can help explore large images, such as scroll bars, the Overview tool, pan tool, and zoom buttons. The Image Tool presents an integrated environment for displaying images and performing common image processing tasks.

For example, this code reads an image into the MATLAB workspace and then displays it in the Image Tool.

```
moon = imread('moon.tif');  
imtool(moon);
```

The following figure shows the image displayed in the Image Tool, with all of the related tools active. For more information about using the Image Tool and related tools, see the topics in the following list.

- “Opening the Image Tool” on page 3-10
- “Specifying the Initial Image Magnification” on page 3-11
- “Closing the Image Tool” on page 3-14
- “Specifying the Colormap” on page 3-11
- “Importing Image Data from the Workspace” on page 3-13
- “Exporting the Image to the Workspace” on page 3-14
- “Closing the Image Tool” on page 3-14
- “Printing the Image in the Image Tool” on page 3-15
- “Using Image Tool Navigation Aids” on page 3-16
- “Getting Information about the Pixels in an Image” on page 3-22.
- “Getting Information About an Image” on page 3-29
- “Adjusting the Contrast and Brightness of an Image” on page 3-31
- “Displaying Different Image Types” on page 3-45




Image Tool and Related Tools

Opening the Image Tool

To start the Image Tool, use the `imtool` function. Once started, you can launch another Image Tool from within an Image Tool by using the **New** option from the **File** menu. The following shows some common `imtool` syntaxes. For complete information, see the `imtool` function reference page.

When called with no arguments,

```
imtool
```

opens an empty Image Tool window. You can bring an image into the Image Tool by using the **Import from Workspace** option from the **File** menu — see “Importing Image Data from the Workspace” on page 3-13.

You can also call `imtool` specifying the name of the MATLAB workspace variable that contains image data.

```
moon = imread('moon.tif');  
imtool(moon)
```

Alternatively, you can specify the name of the graphics file containing the image.

```
imtool('moon.tif');
```

This syntax can be useful for scanning through graphics files. Note, however, that when you use this syntax, the image data is not stored in a MATLAB workspace variable. To bring the image displayed in the Image Tool into the workspace, you must use the `getimage` function or the **Export from Workspace** option from the Image Tool **File** menu — see “Exporting the Image to the Workspace” on page 3-14.

Specifying the Initial Image Magnification

Like `imshow`, the `imshow` function attempts to display an image in its entirety at 100% magnification (one screen pixel for each image pixel). Unlike `imshow`, `imshow` always honors the specified numeric magnification, showing only a portion of the image if it is too big to fit in a figure on the screen and adding scroll bars to allow navigation to parts of the image that are not currently visible. If the specified magnification would make the image too large to fit on the screen, `imshow` scales the image to fit, without issuing a warning. This is the default behavior, specified by the `imshow` 'InitialMagnification' parameter value 'adaptive'.

To override this default initial magnification behavior for a particular call to `imshow`, specify the `InitialMagnification` parameter. For example, to view an image at 150% magnification, use this code.

```
pout = imread('pout.tif');  
imshow(pout, 'InitialMagnification', 150)
```

You can also specify the text string 'fit' as the initial magnification value. In this case, `imshow` scales the image to fit the default size of a figure window.

You can also change the default initial magnification behavior of `imshow` by setting the `ImshowInitialMagnification` toolbox preference. The magnification value you specify affects every call to `imshow` for the current MATLAB session. To make this preference persist between sessions, include the command to set the preference in your `startup.m` file. To learn more about toolbox preferences, see “Setting the Values of Toolbox Preferences” on page 3-59.

When `imshow` scales an image, it uses interpolation to determine the values for screen pixels that do not directly correspond to elements in the image matrix. For more information, see “Interpolation” on page 5-3.

Specifying the Colormap


A colormap is a matrix that can have any number of rows, but must have three columns. Each row in the colormap is interpreted as a color, with the first element specifying the intensity of red, the second green, and the third blue.

To specify the color map used to display an indexed image or an intensity image in the Image Tool, select the **Choose Colormap** option on the **Tools** menu. This activates the Choose Colormap tool, shown below. Using this tool you can select

one of the MATLAB colormaps or select a colormap variable from the MATLAB workspace.

When you select a colormap, the Image Tool executes the colormap function you specify and updates the image displayed. You can edit the colormap command in the **Evaluate Colormap** text box; for example, you can change the number of entries in the colormap (default is 256). You can enter your own colormap function in this field. Press **Enter** to execute the command.

When you choose a colormap, the image updates to use the new map. If you click **OK**, the Image Tool applies the colormap and closes the Choose Colormap tool. If you click **Cancel**, the image reverts to the previous colormap.




Choose Colormap Tool

Importing Image Data from the Workspace

To import image data from the MATLAB workspace into the Image Tool, use the **Import from Workspace** option on the Image Tool **File** menu. In the dialog box, shown below, you select the workspace variable that you want to import into the workspace.

The following figure shows the **Import from Workspace** dialog box. You can use the **Filter** menu to limit the images included in the list to certain image types, i.e., binary, intensity, RGB, or indexed.



Import from Workspace Dialog Box

Exporting the Image to the Workspace

To export the image displayed in the Image Tool to the MATLAB workspace, you can use the **Export to Workspace** option on the Image Tool **File** menu. In the dialog box, shown below, you specify the name you want to assign to the variable in the workspace.



Export Image to Workspace Dialog Box

You can also use the `getimage` function to bring image data from the Image Tool into the MATLAB workspace.

The `getimage` function retrieves the image data (CData) from the current Handle Graphics image object. Because, by default, the Image Tool does not make handles to objects visible, you must use the toolbox function `imgca` to get a handle to the image axes displayed in the Image Tool. For example,

```
moon = getimage(imgca);
```

assigns the image data from `moon.tif` to the variable `moon` if the figure window in which it is displayed is currently active.

Closing the Image Tool

To close the Image Tool window, use the **Close** button in the window title bar or select the **Close** option from the Image Tool **File** menu. You can also use the `imtool` function to return a handle to the Image Tool and use the handle to close the Image Tool. When you close the Image Tool, any related tools that are currently open also close.

Because the Image Tool does not make the handles to its figure objects visible, the Image Tool does not close when you call the `close all` command. If you want to close multiple Image Tools, use the `imtool close all` syntax or select **Close all** from the Image Tool **File** menu.

Printing the Image in the Image Tool

To print the image displayed in the Image Tool, select the **Print to Figure** option from the **File** menu. The Image Tool opens another figure window and displays the image. Use the **Print** option on the **File** menu of this figure window to print the image. See “Printing Images” on page 3-56 for more information.

Using Image Tool Navigation Aids

If an image is large or viewed at a large magnification, the Image Tool displays only a portion of the entire image. When this occurs, the Image Tool includes scroll bars to allow navigation around the image. In some cases, scroll bars might not be sufficient. To help navigate large images, the Image Tool includes the following navigation aids:

- Overview tool — Provides a view of the entire image to help you understand which portion of the image is currently displayed in the Image Tool. See “Overview Navigation” on page 3-16 for more information.
- Pan tool — Lets you click and grab the image displayed and move it in the Image Tool. See “Panning the Image Displayed in the Image Tool” on page 3-19 for more information.
- Zoom tools — Lets you zoom in or out on the image. See “Zooming In and Out on an Image” on page 3-19 for more information.
- Magnification Box — Lets you specify the exact magnification of the image. See “Specifying the Magnification of the Image” on page 3-20 for more information.

Overview Navigation

To get an overview of the image displayed in the Image Tool, use the Overview tool. The Overview tool displays a view of the entire image, scaled to fit, in a separate window. Superimposed over this view of the image is a rectangle, called the *detail rectangle*. The detail rectangle shows which part of the image is currently visible in the Image Tool window. You can change the portion of the image visible in the Image Tool by moving the detail rectangle over the image in the Overview tool.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

When the Image Tool starts, the Overview tool also starts, by default. You can also start the Overview tool by clicking the Overview button  in the Image Tool toolbar or by selecting the **Overview** option from the **Tools** menu in the Image Tool. The following figure shows the Image Tool with the Overview tool. For more information about using the Overview tool, see these topics:

- “Using the Overview Tool” on page 3-17

- “Specifying the Color of the Detail Rectangle” on page 3-18
- “Getting the Position and Size of the Detail Rectangle” on page 3-18
- “Printing the View of the Image in the Overview Tool” on page 3-18




Image Tool with Overview Tool

Using the Overview Tool

To use the Overview tool to explore an image displayed in the Image Tool, follow this procedure:

- 1 Start the Overview tool by clicking the Overview button  in the Image Tool toolbar or by selecting **Overview** from the **Tools** menu. The Overview

tool opens in a separate window containing a view of the entire image, scaled to fit.

The Image Tool opens the Overview tool, by default. If the Overview tool is already active, clicking the **Overview** button brings the tool to the front of the windows open on your screen.

- 2 Using the mouse, move the cursor into the detail rectangle. The cursor changes to the fleur shape, .
- 3 Press and hold the mouse button to drag the detail rectangle anywhere on the image. The Image Tool updates the view of the image to make the specified region visible.

Specifying the Color of the Detail Rectangle

By default, the color of the detail rectangle in the Overview tool is blue. You might want to change the color of the rectangle to achieve better contrast with the predominant color of the underlying image. To do this, right-click anywhere inside the boundary of the detail rectangle and select a color from the **Set Rectangle Color** option on the context menu.

Getting the Position and Size of the Detail Rectangle

To get the current position and size of the detail rectangle, right-click anywhere inside it and select **Copy Position** from the context menu. You can also access this option from the **Edit** menu of the Overview tool.

This option copies the position information to the clipboard. The position information is a vector of the form `[xmin ymin width height]`. To paste this position vector into the MATLAB workspace or another application, right-click and select **Paste** from the context menu.

Printing the View of the Image in the Overview Tool

You can print the view of the image displayed in the Overview tool. Select the **Print to Figure** option from the Overview tool **File** menu. See “Printing Images” on page 3-56 for more information.

Panning the Image Displayed in the Image Tool

To change the portion of the image displayed in the Image Tool, you can use the Drag image to pan button to move the image displayed in the window. This is called *panning* the image.

To pan an image displayed in the Image Tool,

- 1 Click the Drag image to pan button  in the toolbar or select **Pan** from the **Tools** menu.
- 2 Move the cursor over the image in the Image Tool, using the mouse. The cursor changes to an open-hand shape .
- 3 Press and hold the mouse button and drag the image in the Image Tool. When you drag the image, the cursor changes to the closed-hand shape .
- 4 To leave Pan mode, click the Drag image to pan button again to deselect it.

Note As you pan the image in the Image Tool, the Overview tool updates the position of the detail rectangle — see “Overview Navigation” on page 3-16.

Zooming In and Out on an Image

To enlarge an image to get a closer look or shrink an image to see the whole image in context, use the Zoom buttons on the toolbar. (You can also zoom in or out on an image by changing the magnification — see “Specifying the Magnification of the Image” on page 3-20.)

To zoom in or zoom out on an image,

- 1 Click the appropriate magnifying glass button in the Image Tool toolbar or select the **Zoom In** or **Zoom Out** option in the **Tools** menu.



- 2 Move the cursor over the image you want to zoom in or out on, using the mouse. The cursor changes to the appropriate magnifying glass icon. With each click, the Image Tool changes the magnification of the image, centering the new view of the image on the spot where you clicked.

When you zoom in or out on an image, the magnification value displayed in the magnification edit box changes and the **Overview** window updates the position of the detail rectangle.

- 3 To leave zoom mode, click the active zoom button again to deselect it.

Specifying the Magnification of the Image

To enlarge an image to get a closer look or to shrink an image to see the whole image in context, you can use the magnification edit box, shown in the following figure. (You can also use the Zoom buttons to enlarge or shrink an image. See “Zooming In and Out on an Image” on page 3-19 for more information.)




Image Tool Magnification Edit Box and Menu

To change the magnification of an image,

- 1** Move the cursor into the magnification edit box. The cursor changes to the text entry cursor.
- 2** Type a new value in the magnification edit box and press **Enter**. The Image Tool changes the magnification of the image and displays the new view in the window.

You can also specify a magnification by clicking the menu associated with the magnification edit box and selecting from a list of preset magnifications. If you choose the **Fit to Window** option, the Image Tool scales the image so that the entire image is visible.

Getting Information about the Pixels in an Image

Often, you need to get information about the pixels in an image, such as their location and value. The Image Tool provides several ways to get this information, including:

- Pixel Information tool — Displays the location and value of the pixel under the cursor in the lower left corner of the Image Tool window. See “Determining the Value of Individual Pixels” on page 3-22 for more information.
- Display Range tool — Displays the display range of the image in the lower right corner of the Image Tool window. See “Getting the Display Range of an Image” on page 3-24 for more information.
- Pixel Region tool — Displays an extreme close-up view of the pixels in a specific region of an image. See “Viewing Pixel Values with the Pixel Region Tool” on page 3-25 for more information.


Determining the Value of Individual Pixels

The Image Tool provides information about the location and value of individual pixels in an image. This information is displayed in the Pixel Information tool at the bottom left corner of the Image Tool window. The pixel value and location information represent the pixel under the current location of the cursor. The Image Tool updates this information as you move the cursor over the image.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

The following figure shows the Image Tool with pixel location and value displayed in the Pixel Information tool. For more information, see “Saving the Pixel Value and Location Information” on page 3-23.



Pixel Information in Image Tool

Saving the Pixel Value and Location Information

To save the pixel location and value information displayed, right-click a pixel in the image and choose the **Copy pixel info** option. The Image Tool copies the x - and y -coordinates and the pixel value to the clipboard.

To paste this position vector into the MATLAB workspace or another application, right-click and select **Paste** from the context menu.

Getting the Display Range of an Image


The Image Tool provides information about the display range of pixels in an intensity image. The display range is the value of the axes `CLim` property, which controls the mapping of image `CData` to the figure colormap. `CLim` is a two-element vector `[cmin cmax]` specifying the `CData` value to map to the first color in the colormap (`cmin`) and the `CData` value to map to the last color in the colormap (`cmax`). Data values in between are linearly scaled.

The Image Tool displays this information in the Display Range tool at the bottom right corner of the window. The Image Tool does not show the display range for indexed, RGB, or binary images.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

The following figure shows the Image Tool displaying the image with display range information.



Display Range Information in Image Tool

Viewing Pixel Values with the Pixel Region Tool

To view the values of pixels in a specific region of an image displayed in the Image Tool, use the Pixel Region tool. The Pixel Region tool superimposes a rectangle, called the *pixel region rectangle*, over the image displayed in the Image Tool. This rectangle defines the group of pixels that are displayed, in extreme close-up view, in the Pixel Region tool window.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

Start the Pixel Region tool by clicking the Pixel Region button  in the Image Tool toolbar or by selecting the **Pixel Region** option from the **Tools** menu in the Image Tool. The following figure shows the Image Tool with the Pixel Region tool. For more information about using the Pixel Region tool, see these additional topics:

- “Selecting a Region” on page 3-26
- “Customizing the View” on page 3-27
- “Determining the Location of the Pixel Region Rectangle” on page 3-28
- “Printing the View of the Image in the Pixel Region Tool” on page 3-28




Image Tool with Pixel Region Tool and Pixel Region Rectangle

Selecting a Region

To examine pixels in specific regions of an image, perform this procedure:

- 1 Start the Pixel Region tool by clicking the Pixel Region button  in the Image Tool toolbar or by selecting the **Pixel Region** option from the **Tools** menu. The Image Tool displays the pixel region rectangle  in the center of the target image and opens the Pixel Region tool.


Note Scrolling the image can move the pixel region rectangle off the part of the image that is currently displayed. To bring the pixel region rectangle back to the center of the part of the image that is currently visible, click the Pixel Region button again. For help finding the Pixel Region tool in large images, see “Determining the Location of the Pixel Region Rectangle” on page 3-28.

- 2 Using the mouse, position the pointer over the pixel region rectangle. The pointer changes to the fleur shape, .
- 3 Click the left mouse button and drag the pixel region rectangle to any part of the image. As you move the pixel region rectangle over the image, the Pixel Region tool updates the pixel values displayed. You can also move the pixel region rectangle by moving the scroll bars in the Pixel Region tool window.

Customizing the View

The pixel region rectangle defines the group of pixels that are displayed in the Pixel Region tool. To view a larger region, grab any side of the Pixel Region tool figure window and resize it, or use the zoom tools in the Pixel Region toolbar to zoom in or out on the image.

The Pixel Region tool displays the pixels at high magnification, overlaying each pixel with its numeric value. For RGB images, this information includes three numeric values, one for each band of the image. For indexed images, this information includes the index value and the associated RGB value. If you would rather not see the numeric values in the display, go to the Pixel Region tool **Edit** menu and clear the **Superimpose Pixel Values** option.



Pixel Region Tool Edit Menu


Determining the Location of the Pixel Region Rectangle

To determine the current location of the pixel region in the target image, you can use the pixel information given at the bottom of the tool. This information includes the x - and y -coordinates of pixels in the target image coordinate system.

You can also retrieve the current position of the pixel region rectangle by selecting the **Copy Position** option from the Pixel Region tool **Edit** menu. This option copies the position information to the clipboard. The position information is a vector of the form [xmin ymin width height].

To paste this position vector into the MATLAB workspace or another application, right-click and select **Paste** from the context menu.

The following figure shows these components of the Pixel Region tool.



Pixel Region Rectangle Location Information

Printing the View of the Image in the Pixel Region Tool

You can print the view of the image displayed in the Pixel Region tool. Select the **Print to Figure** option from the Pixel Region tool **File** menu. See “Printing Images” on page 3-56 for more information.

Getting Information About an Image

To get information about the image displayed in the Image Tool, use the Image Information tool. The Image Information tool can provide two types of information about an image:

- **Basic information** — Includes width, height, class, and image type. For intensity (grayscale) and indexed images, this information also includes the minimum and maximum intensity values.
- **Image metadata** — Displays all the metadata from the graphics file that contains the image. This is the same information returned by the `imfinfo` function or the `dicominfo` function.

Note The Image Information tool can display image metadata only when you specify the filename containing the image to Image Tool, e.g., `imtool('moon.tif')`.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

Start the Image Information tool by clicking the Image Information button  in the Image Tool toolbar or by selecting the **Image Information** option from the **Tools** menu in the Image Tool.

The following figure shows the Image Tool with the Image Information tool. In the figure, the Image Information tool displays both basic image information and image metadata.




Image Tool with Image Information Tool

Adjusting the Contrast and Brightness of an Image

To adjust the contrast or brightness of an image displayed in the Image Tool, use the Adjust Contrast tool. The Adjust Contrast tool provides an interactive way to adjust the contrast and brightness of an image by manipulating the display range of the image.

The Adjust Contrast tool displays a histogram of the image that is overlaid with a red rectangular box, called a *window*. The histogram represents the *dynamic range* of the image, i.e., the full grayscale resolution possible given the image storage class. The window represents the *display range* of the image, i.e., the actual range of pixel values in the image. An image where grayscale values in the display range do not take full advantage of the dynamic range has a dull, washed-out, low-contrast look. By changing the size of the window, you can make the display range take better advantage of the full dynamic range of the image. The Adjust Contrast tool also provides a way to adjust contrast and brightness using the mouse — see “Using the Window/Level Tool” on page 3-37.

For example, view an image in the Image Tool.

```
imtool('moon.tif')
```

Start the Adjust Contrast tool by clicking the Adjust Contrast button  in the Image Tool toolbar, or by selecting the **Adjust Contrast** option from the **Tools** menu in the Image Tool.

The following figure shows the Image Tool with the Adjust Contrast tool. For more information, see these additional topics:

- “Understanding Contrast Adjustment” on page 3-32
- “Adjusting Contrast and Brightness” on page 3-34
- “Autoscaling” on page 3-39




Image Tool with Adjust Contrast Tool


Understanding Contrast Adjustment

An image lacks contrast when there are no sharp differences between black and white. Brightness refers to the overall lightness or darkness of an image.

To change the contrast or brightness of an image, the Adjust Contrast tool performs *contrast stretching*. In this process, pixel values below a specified value are mapped to black and pixel values above a specified value are mapped to white. The result is a linear mapping of a subset of pixel values to the entire range of display intensities (dynamic range). This produces an image of higher

contrast by darkening pixels whose value is below a specified value and lightening pixels whose value is above a specified value. The Adjust Contrast tools adjust brightness by moving this window over the display range, without changing its size. In this way, the pixel values map to lighter or darker intensities.

The following figure shows this mapping. Note that the lower limit and upper limit mark the boundaries of the window, displayed graphically as the red box in the Adjust Contrast tool.



Display Range to Dynamic Range Mapping

The Adjust Contrast tool accomplishes this contrast stretching by modifying the `CLim` property of the axes object that contains the image. The `CLim` property controls the mapping of image pixel values to display intensities.

By default, the Image Tool sets the `CLim` property to the entire dynamic range available to the data type. For example, the dynamic range of an image of class `uint8` is from 0 to 255. When you use the Adjust Contrast tool, you change the contrast in the image by changing this mapping between image pixel values (display range) and the dynamic range. You create a window over the range that defines which pixels in the image map to the black in the dynamic range by shrinking the range from the bottom up.

Adjusting Contrast and Brightness

The Adjust Contrast tool provides several ways you can change the size of the window to change the mapping between pixel values and display intensities. This example illustrates these capabilities.

- 1 View an image in the Image Tool. This example opens the image `pout.tif`, which is a low-contrast image.

```
imtool('pout.tif')
```

- 2 Start the Adjust Contrast tool by clicking the Adjust Contrast button  in the Image Tool toolbar, or by selecting **Adjust Contrast** from the **Tools** menu in the Image Tool.

The following figure shows the image displayed in the Image Tool with the Adjust Contrast tool open in a separate window. In the figure, note how pixel values in the image histogram are clustered in the middle of the dynamic range. Note that the display range, shown in the lower right corner of the Image Tool, is the full dynamic range of the image.





Image with Default Pixel Value to Display Intensity Mapping

- 3 Adjust the contrast by changing the size of the window overlaid on the image histogram. The Adjust Contrast tool provides several ways to change this window:

- By grabbing one of the red handles on the right and left edges of the window and dragging it. If you shrink the window from the left, the image becomes darker. If you shrink the window from the right, the image becomes lighter.
- By specifying values in the **Minimum value** and **Maximum value** fields and pressing **Enter**.
- By clicking the dropper button associated with the minimum or maximum value fields. When you do this, the cursor becomes an eye dropper shape. Position this cursor over the pixel in the image you want to be the minimum (or maximum) value and click the mouse button.
- By specifying values in the **Window Width** and **Window Center** fields and pressing **Enter**. These values provide another way to specify the size of the window.

If, while making interactive adjustments to contrast and brightness, you want to return the image to its original state, click the **Reset Image** button. The following figure shows the Adjust Contrast tool with the window resized. For information about adjusting the contrast using the **Auto Scale** button, see “Autoscaling” on page 3-39.



Adjust Contrast tool with Window Resized

- 4 Adjust the brightness by moving the window to the left or right, without changing its size. Move the cursor anywhere in the window, and click and drag the window.

The following figure shows the `pout.tif` image after contrast adjustment. Note that the Image Tool updates the display range values displayed in the lower right corner of the Image Tool window.



Contrast Adjusted Image

Using the Window/Level Tool

The Adjust Contrast tool also lets you adjust contrast and brightness by clicking and dragging the mouse. In medical applications, this type of interactivity is often called adjusting the *window/level*.

When you start the Adjust Contrast tool and move the cursor over the image displayed in the Image Tool, the cursor changes shape to the Window/Level cursor . The following table summarizes the mouse motion.

Mouse Motion		Effect
Horizontally to the left		Shrinks the window from both sides
Horizontally to the right		Expands the window from both sides
Vertically up		Moves the window to the right over the histogram, increasing brightness.
Vertically down		Moves the window to the left over the image histogram, decreasing brightness.

The following example shows how to use the window/level capability to improve the contrast of an image.

- 1 Read an image from a sample DICOM file included with the toolbox.

```
I = dicomread('CT-MONO2-16-ankle.dcm');
```

- 2 View the image data using the Image Tool. Because the image data is signed 16-bit data, this example uses the autoscaling syntax.

```
imshow(I, 'DisplayRange', [])
```




- 3 Start the Adjust Contrast tool by clicking the Adjust Contrast button  in the Image Tool toolbar, or by selecting **Adjust Contrast** from the **Tools** menu in the Image Tool.
- 4 Move the cursor over the image. The cursor changes to the window/level shape .
- 5 Press and hold the left mouse button. As you move the window/level cursor horizontally to the left or right to adjust the contrast or move the cursor vertically up or down to change the brightness.

Autoscaling

You can perform automatic contrast adjustment by setting the axes CLim values to the minimum and maximum pixel values in the image. To do this, click the **Autoscaling** button in the Adjust Contrast tool. By default, the CLim minimum and maximum values are set to the full dynamic range of the image.

You can also trim outliers at the top and bottom of the image histogram. By default, the Adjust Contrast tool removes the top 1% and the bottom 1%, but you can specify other percentages. (You can perform this same operation from the command line using the `stretchlim` function.) The following figure shows the **Auto Scale Window**.



Adjust Contrast Tool Auto Scale Window

Viewing Multiple Images

If you specify a file that contains multiple images, `imshow` and `imtool` only display the first image in the file. To view all the images in the file, import the images into the MATLAB workspace by calling `imread`. See “Reading Multiple Images from a Graphics File” on page 2-16 for more information.

Some applications create collections of images related by time or view, such as magnetic resonance imaging (MRI) slices or frames of data acquired from a video stream. The Image Processing Toolbox supports these collections of images as four-dimensional arrays, where each separate image is called a frame and the frames are concatenated along the fourth dimension. All the frames in a multiframe image must be the same size.

Once the images are in the MATLAB workspace, there are two ways to display them using `imshow`:

- Displaying each image in a separate figure window
- Displaying multiple frames in a single figure window

To view all the frames in a multiframe image at once, you can also use the `montage` function. See “Displaying All Frames of a Multiframe Image at Once” on page 3-52 for more information.

Displaying Each Image in a Separate Figure

The simplest way to display multiple images is to display them in separate figure windows. MATLAB does not place any restrictions on the number of images you can display simultaneously.

The Image Tool can only display one image frame at a time. Each time you call `imtool`, it opens a new figure window. Use standard MATLAB indexing syntax to specify the frame to display.

```
imtool(multiframe_array(:,:, :,1));
```

In contrast, `imshow` always displays an image in the current figure. If you display two images in succession, the second image replaces the first image. To view multiple figures with `imshow`, use the `figure` command to explicitly create a new empty figure before calling `imshow` for the next image. For example, to view the first three frames in an array of intensity images `I`,

```
imshow(I(:,:, :,1))
```

```
figure, imshow(I(:,:,2))  
figure, imshow(I(:,:,3))
```

The Image Tool can only display one image frame at a time. Use standard MATLAB indexing syntax to specify the frame to display.

```
imtool(multiframe_array(:,:,1));
```

Displaying Multiple Images in the Same Figure

You can use the `imshow` function with the MATLAB `subplot` function or the MATLAB `subimage` function to display multiple images in a single figure window.

Note `imtool` does not support this capability.

Dividing a Figure Window into Multiple Display Regions. `subplot` divides a figure into multiple display regions. The syntax of `subplot` is

```
subplot(m,n,p)
```


This syntax divides the figure into an m -by- n matrix of display regions and makes the p th display region active.

Note When you use `subplot` to display multiple color images in one figure window, the images must share the colormap of the last image displayed. In some cases, as illustrated by the following example, the display results can be unacceptable. As an alternative, you can use the `subimage` function, described in “Using the `subimage` Function to Display Multiple Images” on page 3-43, or you can map all images to the same colormap as you load them.

For example, you can use this syntax to display two images side by side.

```
[X1,map1]=imread('forest.tif');  
[X2,map2]=imread('trees.tif');  
subplot(1,2,1), imshow(X1,map1)  
subplot(1,2,2), imshow(X2,map2)
```


In the figure, note how the first image displayed, X1, appears dark after the second image is displayed.



Two Images in Same Figure Using the Same Colormap

Using the `subimage` Function to Display Multiple Images. `subimage` converts images to RGB before displaying them and therefore circumvents the colormap sharing problem. This example uses `subimage` to display the forest and the trees images with better results.

```
[X1,map1]=imread('forest.tif');  
[X2,map2]=imread('trees.tif');  
subplot(1,2,1), subimage(X1,map1)  
subplot(1,2,2), subimage(X2,map2)
```



Two Images in Same Figure Using Separate Colormaps

Displaying Different Image Types

This section describes how to use `imshow` and `imtool` with the different types of images supported by the Image Processing Toolbox.

- Indexed images
- Intensity (grayscale) images
- Binary images
- RGB (truecolor) images

If you need help determining what type of image you are working with, see “Image Types in the Toolbox” on page 2-3.

Displaying Indexed Images

To display an indexed image, using either `imshow` or `imtool`, specify both the image matrix and the colormap. This documentation uses the variable name `X` to represent an indexed image in the workspace, and `map` to represent the colormap.

```
imshow(X,map)
```

or

```
imtool(X,map)
```

For each pixel in `X`, these functions display the color stored in the corresponding row of `map`. If the image matrix data is of class `double`, the value 1 points to the first row in the colormap, the value 2 points to the second row, and so on. However, if the image matrix data is of class `uint8` or `uint16`, the value 0 (zero) points to the first row in the colormap, the value 1 points to the second row, and so on. This offset is handled automatically by the `imtool` and `imshow` functions.

If the colormap contains a greater number of colors than the image, the functions ignore the extra colors in the colormap. If the colormap contains fewer colors than the image requires, the functions set all image pixels over the limits of the colormap’s capacity to the last color in the colormap. For example, if an image of class `uint8` contains 256 colors, and you display it with a colormap that contains only 16 colors, all pixels with a value of 15 or higher are displayed with the last color in the colormap.

Displaying Intensity Images

To display an intensity (grayscale) image, using either `imshow` or `imtool`, specify the image matrix as an argument. This documentation uses the variable name `I` to represent an intensity image in the workspace.

```
imshow(I)
```

or

```
imtool(I)
```

Both functions display the image by *scaling* the intensity values to serve as indices into a grayscale colormap.

If `I` is `double`, a pixel value of 0.0 is displayed as black, a pixel value of 1.0 is displayed as white, and pixel values in between are displayed as shades of gray. If `I` is `uint8`, then a pixel value of 255 is displayed as white. If `I` is `uint16`, then a pixel value of 65535 is displayed as white.

Intensity images are similar to indexed images in that each uses an `m-by-3` RGB colormap, but you normally do not specify a colormap for an intensity image. MATLAB displays intensity images by using a grayscale system colormap (where `R=G=B`). By default, the number of levels of gray in the colormap is 256 on systems with 24-bit color, and 64 or 32 on other systems. (See “Working with Different Screen Bit Depths” on page 13-3 for a detailed explanation.)

Displaying Intensity Images That Have Unconventional Ranges

In some cases, the image data you want to display as an intensity image might have a display range that is outside the conventional toolbox range (i.e., `[0,1]` for `single` or `double` arrays, `[0,255]` for `uint8` arrays, `[0,65535]` for `uint16` arrays, or `[-32767,32768]` for `int16` arrays). For example, if you filter an intensity image, some of the output data might fall outside the range of the original data.

To display unconventional range data as an image, you can specify the display range directly, using this syntax for both the `imshow` and `imtool` functions.

```
imshow(I, 'DisplayRange', [low high])
```


or

```
imtool(I, 'DisplayRange', [low high])
```

If you use an empty matrix (`[]`) for the display range, these functions scale the data automatically, setting low and high to the minimum and maximum values in the array.

The next example filters an intensity image, creating unconventional range data. The example calls `imshow` to display the image, using the automatic scaling option. If you execute this example, note the display range specified in the lower right corner of the Image Tool window.

```
I = imread('testpat1.png');  
J = filter2([1 2;-1 -2],I);  
imshow(J, 'DisplayRange', []);
```



Displaying Binary Images

In MATLAB, a binary image is of class `logical`. Binary images contain only 0's and 1's. Pixels with the value 0 are displayed as black; pixels with the value 1 are displayed as white.


Note For the toolbox to interpret the image as binary, it must be of class `logical`. Intensity images that happen to contain only 0's and 1's are not binary images.

To display a binary image, using either `imshow` or `imtool`, specify the image matrix as an argument. For example, this code reads a binary image into the MATLAB workspace and then displays the image. This documentation uses the variable name `BW` to represent a binary image in the workspace

```
BW = imread('circles.png');  
imshow(BW)
```

or

```
imtool(BW)
```




Changing the Display Colors of a Binary Image

You might prefer to invert binary images when you display them, so that 0 values are displayed as white and 1 values are displayed as black. To do this, use the NOT (`~`) operator in MATLAB. (In this figure, a box is drawn around the image to show the image boundary.) For example:

```
imshow(~BW)
```

or

```
imtool(~BW)
```




You can also display a binary image using the indexed image colormap syntax. For example, the following command specifies a two-row colormap that displays 0's as red and 1's as blue.

```
imshow(BW,[1 0 0; 0 0 1])
```

or

```
imtool(BW,[1 0 0; 0 0 1])
```



Displaying Truecolor Images

Truecolor images, also called RGB images, represent color values directly, rather than through a colormap. A truecolor image is an m -by- n -by-3 array. For each pixel (r, c) in the image, the color is represented by the triplet $(r, c, 1:3)$.

To display a truecolor image, using either `imshow` or `imtool`, specify the image matrix as an argument. For example, this code reads a truecolor image into the

MATLAB workspace and then displays the image. This documentation uses the variable name `RGB` to represent a truecolor image in the workspace

```
RGB = imread('peppers.png');  
imshow(RGB)
```

or

```
imtool(RGB)
```



Systems that use 24 bits per screen pixel can display truecolor images directly, because they allocate 8 bits (256 levels) each to the red, green, and blue color planes. On systems with fewer colors, `imshow` displays the image using a combination of color approximation and dithering. See “Working with Different Screen Bit Depths” on page 13-3 for more information.

Note If you display a color image and it appears in black and white, check if the image is an indexed image. With indexed images, you must specify the colormap associated with the image. For more information, see “Displaying Indexed Images” on page 3-45.

Special Display Techniques

In addition to `imshow` and `imtool`, the toolbox includes functions that perform specialized display operations, or exercise more direct control over the display format. These functions, together with the MATLAB graphics functions, provide a range of image display options.

This section includes the following topics:

- “Adding a Colorbar” on page 3-51
- “Displaying All Frames of a Multiframe Image at Once” on page 3-52
- “Converting a Multiframe Image to a Movie” on page 3-54
- “Texture Mapping” on page 3-55

Adding a Colorbar

To display an image with a colorbar that indicates the range of intensity values, first use the `imshow` function to display the image in a MATLAB figure window and then call the `colorbar` function to add the colorbar to the image.


When you add a colorbar to an axes object that contains an image object, the colorbar indicates the data values that the different colors in the image correspond to.

If you want to add a colorbar to an image displayed in the Image Tool, select the **Print to Figure** option from the Image Tool **File** menu. The Image Tool displays the image in a separate figure window to which you can add a colorbar.

Seeing the correspondence between data values and the colors displayed by using a colorbar is especially useful if you are displaying unconventional range data as an image, as described under “Displaying Intensity Images That Have Unconventional Ranges” on page 3-46.

In the example below, a grayscale image of class `uint8` is filtered, resulting in data that is no longer in the range `[0,255]`.

```
RGB = imread('saturn.png');
I = rgb2gray(RGB);
h = [1 2 1; 0 0 0; -1 -2 -1];
I2 = filter2(h,I);
imshow(I2,'DisplayRange',[0,255]), colorbar
```



Displaying All Frames of a Multiframe Image at Once

To view all the frames in a multiframe array at one time, use the `montage` function. `montage` divides a figure window into multiple display regions and displays each image in a separate region.

The syntax for `montage` is similar to the `imshow` syntax. To display a multiframe intensity image, the syntax is

```
montage(I)
```


To display a multiframe indexed image, the syntax is

```
montage(X, map)
```

Note All the frames in a multiframe indexed array must use the same colormap.

This example loads and displays all frames of a multiframe indexed image. The example initializes an array to hold the 27 frames in the multiframe image file and then loops, using `imread` to read a single frame from the image file at each iteration.

```
mri = uint8(zeros(128,128,1,27));  
  
for frame=1:27  
    [mri(:,:,frame),map] = imread('mri.tif',frame);  
end  
montage(mri,map);
```



All Frames of Multiframe Image Displayed in One Figure

`montage` displays the first frame in the first position of the first row, the next frame in the second position of the first row, and so on. `montage` arranges the frames so that they roughly form a square.

Converting a Multiframe Image to a Movie

To create a MATLAB movie from a multiframe image array, use the `immovie` function. This example creates a movie from a multiframe indexed image.

```
mov = immovie(X,map);
```

In the example, `X` is a four-dimensional array of images that you want to use for the movie.

You can play the movie in MATLAB using the `movie` function.

```
movie(mov);
```

This example loads the multiframe image `mri.tif` and makes a movie out of it. It won't do any good to show the results here, so try it out; it's fun to watch.

```
mri = uint8(zeros(128,128,1,27));  
for frame=1:27  
    [mri(:,:,,frame),map] = imread('mri.tif',frame);  
end  
  
mov = immovie(mri,map);  
movie(mov);
```

Note that `immovie` displays the movie as it is being created, so you actually see the movie twice. The movie runs much faster the second time (using `movie`).


Note To view a MATLAB movie, you must have MATLAB installed. To make a movie that can be run outside MATLAB, use the MATLAB `avifile` and `addframe` functions to create an AVI file. AVI files can be created using indexed and RGB images of classes `uint8` and `double`, and don't require a multiframe image.

Texture Mapping

When you use `imshow` or `imtool` to view an image, MATLAB displays the image in two dimensions. However, it is also possible to map an image onto a parametric surface, such as a sphere, or below a surface plot. The `warp` function creates these displays by *texture mapping* the image. Texture mapping is a process that maps an image onto a surface grid using interpolation.

This example texture-maps an image of a test pattern onto a cylinder.

```
[x,y,z] = cylinder;  
I = imread('testpat1.png');  
warp(x,y,z,I);
```



An Image Texture-Mapped onto a Cylinder

The image might not map onto the surface in the way that you expect. One way to modify the way the texture map appears is to change the settings of the `Xdir`, `Ydir`, and `Zdir` properties. For more information, see “Changing Axis Direction” in the MATLAB Graphics documentation.

For more information about texture mapping, see the reference entry for the `warp` function.

Printing Images

If you want to output a MATLAB image to use in another application (such as a word-processing program or graphics editor), use `imwrite` to create a file in the appropriate format. See “Writing a Graphics Image” on page 2-16 for details.

If you want to print an image, use `imshow` to display the image in a MATLAB figure window. If you are using the Image Tool, you must use the **Print to Figure** option on the Image Tool **File** menu. When you choose this option, the Image Tool opens a separate figure window and displays the image in it. You can access the standard MATLAB printing capabilities in this figure window. You can also use the **Print to Figure** option to print the image displayed in the Overview tool and the Pixel Region tool.

Once the image is displayed in a figure window, you can use either the MATLAB print command or the **Print** option from the **File** menu of the figure window to print the image. When you print from the figure window, the output includes nonimage elements such as labels, titles, and other annotations.

Printing and Handle Graphics Object Properties

The output reflects the settings of various properties of Handle Graphic objects. In some cases, you might need to change the settings of certain properties to get the results you want. Here are some tips that might be helpful when you print images:

- Image colors print as shown on the screen. This means that images are not affected by the figure object’s `InvertHardcopy` property.
- To ensure that printed images have the proper size and aspect ratio, set the figure object’s `PaperPositionMode` property to `auto`. When `PaperPositionMode` is set to `auto`, the width and height of the printed figure are determined by the figure’s dimensions on the screen. By default, the value of `PaperPositionMode` is `manual`. If you want the default value of `PaperPositionMode` to be `auto`, you can add this line to your `startup.m` file.

```
set(0, 'DefaultFigurePaperPositionMode', 'auto')
```

For detailed information about printing with **File/Print** or the `print` command (and for information about Handle Graphics), see “Printing and Exporting Figures with MATLAB” in the MATLAB Graphics documentation. For a

complete list of options for the print command, enter `help print` at the MATLAB command-line prompt or see the print command reference page in the MATLAB documentation.

Setting Toolbox Display Preferences

You can use Image Processing Toolbox preferences to control certain characteristics of how `imshow` and `imtool` display images on your screen. For example, using toolbox preferences, you can specify the initial magnification used by `imtool` and `imshow`.

This section

- Lists the preferences supported by the toolbox
- Describes how to get the current value of a preference using the `iptgetpref` function
- Describes how to set the value of a preference using the `iptsetpref` function

Toolbox Preferences

The Image Processing Toolbox supports several preferences that affect how `imshow` and `imtool` display images. The following table lists these preferences with brief descriptions. For detailed information about toolbox preferences and their values, see the `iptsetpref` reference page.

Toolbox Preference	Description
<code>ImshowBorder</code>	Controls whether <code>imshow</code> displays the figure window as larger than the image (leaving a border between the image axes and the edges of the figure), or the same size as the image (leaving no border).
<code>ImshowAxesVisible</code>	Controls whether <code>imshow</code> displays images with the axes box and tick labels.
<code>ImshowInitialMagnification</code>	Controls the magnification <code>imshow</code> uses when it initially displays an image. This preference can be overridden for a single call to <code>imshow</code> ; see “Specifying the Initial Image Magnification” on page 3-5 for more details.
<code>ImtoolInitialMagnification</code>	Controls the magnification the Image Tool uses when it initially displays an image. This preference can be overridden for a single call to <code>imtool</code> ; see “Specifying the Initial Image Magnification” on page 3-11 for more details.

Retrieving the Values of Toolbox Preferences

To determine the current value of a preference, use the `iptgetpref` function. This example uses `iptgetpref` to determine the value of the `imtoolInitialMagnification` preference.

```
iptgetpref('ImtoolInitialMagnification')
```

```
ans =
```

```
100
```

Preference names are case insensitive and can be abbreviated. For more information, see the `iptgetpref` reference page.

Setting the Values of Toolbox Preferences

To specify the value of a toolbox preference, use the `iptsetpref` function. This example calls `iptsetpref` to specify that `imshow` resize the figure window so that it fits tightly around displayed images.

```
iptsetpref('ImshowBorder', 'tight');
```

For detailed information about toolbox preferences and their values, see the `iptsetpref` reference page.

The value you specify lasts for the duration of the current MATLAB session. To preserve your preference settings from one session to the next, include the `iptsetpref` commands in your `startup.m` file.

Building GUIs with Modular Tools

This chapter describes how to use the toolbox modular tools to create custom image processing applications.

Overview (p. 4-2)

Lists the modular interactive tools

Using Modular Tools (p. 4-6)

Describes how to use the modular tools to create GUIs


Creating Your Own Modular Tools
(p. 4-28)



Describes the utility function the toolbox provides to help
you create your own modular tools


Overview


The toolbox includes several new modular interactive tools that you can activate from the command line and use with images displayed in a MATLAB figure window, called the *target image* in this documentation. The tools are modular because they can be used independently or in combination to create custom graphical user interfaces (GUIs) for image processing applications. The Image Tool uses these modular tools.

The following table lists the modular tools in alphabetical order. The table includes an illustration of each tool and the function you use to create it. For more information about how the tools operate, see “Using the Image Tool to Explore Images” on page 3-8. For more information about using tools to create GUIs, see “Using Modular Tools” on page 4-6.

Modular Tool	Example	Description
Adjust Contrast tool		<p>Displays a histogram of the target image and enables interactive adjustment of contrast and brightness by manipulation of the display range.</p> <p>Use the <code>imcontrast</code> function to create the tool in a separate figure window and associate it with an image.</p>
Display Range tool		<p>Displays a text string identifying the display range values of the associated image.</p> <p>Use the <code>imdisplayrange</code> function to create the tool, associate it with an image, and embed it in a figure or uipanel.</p>

Modular Tool	Example	Description																										
<p>Image Information tool</p>	 <p>The screenshot shows a window titled 'Image Information: image1.tif'. It contains a 'Properties' section with fields for 'Image Size', 'Image Type', 'Image Name', and 'Image Path'. Below this is a 'Metadata' section with a table of key-value pairs:</p> <table border="1"> <thead> <tr> <th>Key</th> <th>Value</th> </tr> </thead> <tbody> <tr><td>1. Title</td><td>1. Title</td></tr> <tr><td>2. Author</td><td>2. Author</td></tr> <tr><td>3. Subject</td><td>3. Subject</td></tr> <tr><td>4. Date</td><td>4. Date</td></tr> <tr><td>5. Description</td><td>5. Description</td></tr> <tr><td>6. Keywords</td><td>6. Keywords</td></tr> <tr><td>7. Comment</td><td>7. Comment</td></tr> <tr><td>8. Creation Date</td><td>8. Creation Date</td></tr> <tr><td>9. Modification Date</td><td>9. Modification Date</td></tr> <tr><td>10. Software</td><td>10. Software</td></tr> <tr><td>11. Original Name</td><td>11. Original Name</td></tr> <tr><td>12. Original Path</td><td>12. Original Path</td></tr> </tbody> </table>	Key	Value	1. Title	1. Title	2. Author	2. Author	3. Subject	3. Subject	4. Date	4. Date	5. Description	5. Description	6. Keywords	6. Keywords	7. Comment	7. Comment	8. Creation Date	8. Creation Date	9. Modification Date	9. Modification Date	10. Software	10. Software	11. Original Name	11. Original Name	12. Original Path	12. Original Path	<p>Displays basic attributes about the target image. If the name of the graphics file containing the image was specified, the tool includes any metadata the image file might contain.</p> <p>Use the <code>imageinfo</code> function to create the tool in a separate figure window and associate it with an image.</p>
Key	Value																											
1. Title	1. Title																											
2. Author	2. Author																											
3. Subject	3. Subject																											
4. Date	4. Date																											
5. Description	5. Description																											
6. Keywords	6. Keywords																											
7. Comment	7. Comment																											
8. Creation Date	8. Creation Date																											
9. Modification Date	9. Modification Date																											
10. Software	10. Software																											
11. Original Name	11. Original Name																											
12. Original Path	12. Original Path																											
<p>Magnification box</p>	 <p>The screenshot shows a simple rectangular box with a light gray background and a thin border. Inside the box, the text '100%' is centered.</p>	<p>Creates a text edit box containing the current magnification of the target image. Users can change the magnification of the image by entering a new magnification value.</p> <p>Use <code>immgbox</code> to create the tool, associate it with an image, and embed it in a figure or uipanel.</p> <p>Note: The target image must be contained in a scroll panel.</p>																										

Modular Tool	Example	Description
Overview tool		<p>Displays the target image in its entirety with the portion currently visible in the scroll panel outlined by a rectangle superimposed on the image. Moving the rectangle changes the portion of the target image that is currently visible in the scroll panel.</p> <p>Use <code>imoverview</code> to create the tool in a separate figure window and associate it with an image.</p> <p>Use <code>imoverviewpanel</code> to create the tool in a <code>uipanel</code> that can be embedded within another figure or <code>uipanel</code>.</p> <p>Note: The target image must be contained in a scroll panel.</p>
Pixel Information tool		<p>Displays information about the pixel the mouse is over in the target image.</p> <p>Use <code>impixelinfo</code> to create the tool, associate it with an image, and display it in a figure or <code>uipanel</code>.</p> <p>If you want to display only the pixel values, without the Pixel info label, use <code>impixelinfoval</code>.</p>

Modular Tool	Example	Description
Pixel Region tool		<p>Display pixel values for a specified region in the target image.</p> <p>Use <code>impixelregion</code> to create the tool in a separate figure window and associate it with an image.</p> <p>Use <code>impixelregionpanel</code> to create the tool as a uipanel that can be embedded within another figure or uipanel.</p>
Scroll Panel tool		<p>Display target image in a scrollable panel.</p> <p>Use <code>imscrollpanel</code> to add a scroll panel to an image displayed in a figure window.</p>

Using Modular Tools

To use the modular tools to create custom graphical user interfaces (GUIs) for image processing applications, follow this general procedure:

- 1** Display the target image in a figure window.

Image processing applications typically include the display of the target image, i.e., the image being processed. You can use the `imshow` function as the foundation for your GUI application. (You can also use the MATLAB `image` and `imagesc` functions.) See “Displaying the Target Image” on page 4-7 for more information.

- 2** Create the modular tool, specifying the target image.

The modular tools operate on an image. When you create a tool, you must associate it with a target image. Most of the tools associate themselves with the image in the current axes, by default. But you can specify the handle to a specific image object, or a handle to a figure, axes, or uipanel object that contains an image. See “Specifying the Target Image” on page 4-7 for more information.

Depending on how you designed your GUI, you might also want to specify the parent object of the modular tool. This is optional; by default, the tools either use the same parent as the target image or open in a separate figure window. If you want to change this default, you must specify the parent. See “Specifying the Parent of a Modular Tool” on page 4-11 for more information.

In addition, when you create custom GUIs, you might need to specify the position of the graphics objects in the GUI, including the modular tools. See “Positioning the Modular Tools in a GUI” on page 4-14 for more information.

- 3** Set up interactivity between the tool and the target image.

This is an optional step. The modular tools all set up their interactive connections to the target image automatically. However, your GUI might require some additional connectivity. See “Making Connections for Interactivity” on page 4-22.

The following sections provide more detail on these steps. For a complete illustration, see “Example: Building a Pixel Information GUI” on page 4-15.

Displaying the Target Image

As the foundation for any image processing GUI you create, use `imshow` to display the target image (or images) in a MATLAB figure window. (You can also use the MATLAB `image` or `imagesc` functions.) Once the image is displayed in the figure, you can associate any of the modular tools with the image displayed in the figure.

This example uses `imshow` to display an image in a figure window.

```
himage = imshow('pout.tif');
```

Because some of the modular tools add themselves to the figure window containing the image, make sure that the Image Processing Toolbox `ImshowBorder` preference is set to 'loose', if you are using the `imshow` function. By including a border, you ensure that the modular tools are not displayed over the image in the figure. This is the default setting.

Specifying the Target Image

To associate a modular tool with a target image displayed in a MATLAB figure window, create the tool using the appropriate tool creation function, specifying a handle to the target image as an argument. The function creates the tool and automatically sets up the interactivity connection between the tool and the target image that the tool requires.

This section covers the following topics:

- “Associating Modular Tools with the Default Target Image” on page 4-7
- “Associating Modular Tools with a Particular Image” on page 4-8
- “Getting the Handle of the Target Image” on page 4-10

Associating Modular Tools with the Default Target Image

By default, most of the modular tool creation functions support a no-argument syntax that uses the image in the current figure as the target image. If the current figure contains multiple images, the tools associate themselves with the first image in the figure object’s children (the last image created).

`impixelinfo`, `impixelinfoval` and `imshowarray` can work with multiple images in a figure.

For example, to use the Pixel Information tool with a target image, display the image in a figure window, using `imshow`, and then call the `impixelinfo` function to create the tool. In this example, the image in the current figure is the target image.

```
imshow('pout.tif');  
impixelinfo
```

The following figure shows the target image in a figure with the Pixel Information tool in the lower left corner of the window. The Pixel Information tool automatically sets up a connection to the target image: when you move the cursor over the image, the tool displays the x - and y -coordinates and value of the pixel under the cursor.




Figure Window with Pixel Information Tool

Associating Modular Tools with a Particular Image

You can specify the target image of the modular tool when you create it. Pass a handle to the target image as an argument to the modular tool creation

function. You can also specify a handle to a figure, axes, or uipanel object that contains the target image.

Continuing the example in the previous section, you might want to add the Display Range tool to the figure window that already contains the Pixel Information tool. To do this, call the `imshow` function, specifying the handle to the target image. You could also have specified the handle of the figure, axes, or uipanel object containing the target image.

```
himage = imshow('pout.tif');  
hpixelinfo = impixelinfo(himage)  
hdisplayrange = imshow(himage)
```

Note that the example retrieves handles to the uipanel objects created by the `impixelinfo` and `imshow` functions; both tools are uipanel objects. It can be helpful to get handles to the tools if you want to change their positioning. See “Positioning the Modular Tools in a GUI” on page 4-14 for more information.

The following figure shows the target image in a figure with the Pixel Information tool in the lower left corner and the Display Range tool in the lower right corner of the window. The Display Range tool automatically sets up a connection to the target image: when you move the cursor over the image (or images) in the figure, the Display Range tool shows the display range of the image.




Figure Window with Pixel Information and Display Range Tools

Getting the Handle of the Target Image

The examples in the previous section use the optional `imshow` syntax that returns a handle to the image displayed, `himage`. When creating GUIs with the modular tools, having a handle to the target image can be useful. You can get the handle when you first display the image, using this optional `imshow` syntax. You can also get a handle to the target image using the `imhandles` function. The `imhandles` function returns all the image objects that are children of a specified figure, axes, uipanel, or image object.

For example, `imshow` returns a handle to the image in this syntax.

```
hfig = figure;  
himage = imshow('moon.tif')  
himage =
```

```
152.0055
```

When you call the `imhandles` function, specifying a handle to the figure (or axes) containing the image, it returns a handle to the same image.

```
himage2 = imhandles(hfig)
himage2 =

    152.0055
```

Specifying the Parent of a Modular Tool

When you create a modular tool, in addition to specifying the target image, you can optionally specify the object that you want to be the parent of the tool. By specifying the parent, you determine where the tool appears on your screen. Using this syntax of the modular tool creation functions, you can add the tool to the figure window containing the target image, open the tool in a separate figure window, or create some other combination.

Specifying the parent is optional; the modular tools all have a default behavior. Some of the smaller tools, such as the Pixel Information tool, use the parent of the target image as their parent, inserting themselves in the same figure window as the target image. Other modular tools, such as the Pixel Region tool or the Overview tool, open in separate figures of their own.


Two of the tools, the Pixel Region tool and the Overview tool, have a separate creation function to provide this capability. Their primary creation functions, `imoverview` and `impixelregion`, open the tools in a separate figure window. To specify a different parent, you must use the `imoverviewpanel` and `impixelregionpanel` functions.

Note The Overview tool and the Pixel Region tool provide additional capabilities when created in their own figure windows. For example, both tools include zoom buttons that are not part of their `uipanel` versions.

This example shows the default behavior when you create the Pixel Region tool using the `impixelregion` function. The tool opens in a separate figure window, as shown in the following figure.

```
himage = imshow('pout.tif')
hpixelinfopanel = impixelinfo(himage)
hdisplayrangepanel = imdisplayrange(himage)
```

```
hpixreg = impixelregion(himage)
```



Target Image with Pixel Region Tool in Separate Window (Default)

To embed the Pixel Region tool in the same window as the target image, you must specify the handle of the target image's parent figure as the parent of the Pixel Region tool when you create it.


The following example creates a figure and an axes object, getting handles to both objects. The example needs these handles to perform some repositioning of the objects in the figure to ensure their visibility. See "Positioning the Modular Tools in a GUI" on page 4-14 for more information. The example then creates the modular tools, specifying the figure containing the target image as the parent of the Pixel Region tool. Note that the example uses the `impixelregionpanel` function to create the tool.

```

hfig = figure;
hax = axes('units','normalized','position',[0 .5 1 .5]);
himage = imshow('pout.tif')
hpixelinfopanel = impixelinfo(himage)
hdrangepanel = imdisplayrange(himage)
hpixreg = impixelregionpanel(hfig,himage)
set(hpixreg, 'Units','normalized','Position',[0 .08 1 .4])

```

The following figure shows the Pixel Region embedded in the same figure as the target image.



Target Image with Embedded Pixel Region Tool

Positioning the Modular Tools in a GUI

When you create the modular tools, they have default positioning behavior. For example, the `impixelinfo` function creates the tool as a uipanel object that is the full width of the figure window, positioned in the lower left corner of the target image figure window.

Because the modular tools are constructed from standard Handle Graphics objects, such as uipanel objects, you can use properties of the objects to change their default positioning or other characteristics.

For example, in “Specifying the Parent of a Modular Tool” on page 4-11, when the Pixel Region tool was embedded in the same figure window as the target image, the example had to reposition both the image object and the Pixel Region tool uipanel object to make them both visible in the figure window.

To reposition a modular tool or other graphics object, set the value of the `Position` property of the object. As the value of this property, you specify a four-element position vector `[left bottom width height]`, where `left` and `bottom` specify the distance from the lower left corner of the parent container object, such as a figure. The `width` and `height` specify the dimensions of the object. You must specify the units of these values. Many graphics object use normalized units that specify the relative position, not the exact location in pixels, to allow better resizing behavior. Get the value of the `Units` property of the object to determine how the object interprets the position vector.

For example, when you first create an embedded Pixel Region tool in a figure, it appears to take over the entire figure because, by default, the position vector is set to `[0 0 1 1]`, in normalized units. This position vector tells the tool to align itself with the bottom left corner of its parent and fill the entire object. To accommodate the image and the Pixel Information tool and Display Range tools, change the position of the Pixel Region so that it fills the lower half of the figure, leaving room at the bottom for the Pixel Information and Display Range tools. Here is the position vector for the Pixel Region tool.

```
set(hpixreg, 'Units','normalized','Position',[0 .08 1 .4])
```

To accommodate the Pixel Region tool, you must reposition the target image so that it takes the upper half of the window, using the following position vector. To reposition the image, you must specify the `Position` property of the axes object that contains it; image objects do not have a `Position` property.

```
set(hax,'Units','normalized','Position',[0 0.5 1 0.5])
```

Example: Building a Pixel Information GUI

This example shows how to use the tools to create a simple GUI. If your work typically requires information about the pixels in an image, you might want to create a custom image display function that provides this information. You can do this by using the three modular tools that provide pixel information:

- Pixel Information tool
- Display Range tool
- Pixel Region tool

This example creates a simple image display function that includes these tools. In this GUI, the Pixel Region tool is embedded in the same figure window as the target image. The example suppresses the figure window toolbar and menu bar, because the standard figure zoom tools are not compatible with the toolbox modular navigation tools — see “Adding Navigation Aids to a GUI” on page 4-17.

```
function my_pixinfotool(im)
% Create figure, setting up properties
hfig = figure('Toolbar','none',...
             'Menubar', 'none',...
             'Name','My Pixel Info Tool',...
             'NumberTitle','off',...
             'IntegerHandle','off');

% Create axes
% Reposition the image to accommodate the Pixel Region tool
hax = axes('Units','normalized',...
          'Position',[0 .5 1 .5]);

% Display image, getting handle to image
himage = imshow(im)

% Add Pixel Information tool to image figure window
hpixinfo = impixelinfo(himage)

% Add Display Range tool to image figure window
hdrange = imdisplayrange(himage)

% Add Pixel Region tool as panel in same figure
```

```
hpixreg = impixelregionpanel(hfig,himage)

% Reposition the Pixel Region tool to fit in the figure
% window, leaving room for the Pixel Information and
% Display Range tools.
set(hpixreg, 'units','normalized','position',[0 .08 1 .4])
```

To use the tool, pass it an image that is already in the MATLAB workspace.

```
pout = imread('pout.tif');
my_pixinfotool(pout)
```

The tool opens a figure window, displaying the image in the upper half and the three pixel information modular tools, the Pixel Information tool, Display Range tool, and the Pixel Region tool, in the lower half.



Custom Image Display Tool with Pixel Information

Adding Navigation Aids to a GUI

Note The toolbox modular navigation tools are incompatible with standard MATLAB figure window navigation tools. When using these tools in a GUI, suppress the toolbar and menu bar in the figure windows to avoid conflicts between the tools.

The toolbox includes several modular tools that you can use to add navigation aids to a GUI application:

- Scroll Panel
- Overview tool
- Magnification box

The Scroll Panel is the primary navigation tool; it is a prerequisite for the other navigation tools. When you display an image in a Scroll Panel, the tool displays only a portion of the image, if it is too big to fit into the figure window. When only a portion of the image is visible, the Scroll Panel adds horizontal and vertical scroll bars, to enable viewing of the parts of the image that are not currently visible.

Once you create a Scroll Panel, you can optionally add the other modular navigation tools: the Overview tool and the Magnification tool. The Overview tool displays a view of the entire image, scaled to fit, with a rectangle superimposed over it that indicates the part of the image that is currently visible in the scroll panel. The Magnification Box displays the current magnification of the image and can be used to change the magnification.

The following sections provide more details.


- “Understanding Scroll Panels” on page 4-18 — Adding a scroll panel to an image display changes the relationship of the graphics objects used in the display. This section provides some essential background.
- “Example: Building a Navigation GUI for Large Images” on page 4-20 — This section shows how to add a scroll panel to an image display.

Understanding Scroll Panels

When you display an image in a scroll panel, it changes the object hierarchy of your displayed image. This diagram illustrates the typical object hierarchy for an image displayed in an axes object in a figure object.

```
hfig = figure;  
himage = imshow('concordaerial.png');
```

The following figure shows this object hierarchy.




Object Hierarchy of Image Displayed in a Figure

When you call the `imscrollpanel` function to put the target image in a scrollable window, this object hierarchy changes. For example, this code adds a scroll panel to an image displayed in a figure window, specifying the parent of the scroll panel and the target image as arguments. The example suppresses the figure window toolbar and menu bar because they are not compatible with the scroll panel navigation tools.


```
hfig = figure('Toolbar','none',...  
             'Menubar','none');  
himage = imshow('concordaerial.png');  
hpanel = imscrollpanel(hfig,himage);
```

The following figure shows the object hierarchy after the call to `imscrollpanel`. Note how `imscrollpanel` inserts new objects (shaded in gray) into the hierarchy between the figure object and the axes object containing the image.



Object Hierarchy of Image Displayed in Scroll Panel

The following figure shows how these graphics objects appear in the scrollable image as it displayed on the screen.



Components of a Scroll Panel

Example: Building a Navigation GUI for Large Images

If your work typically requires that you view large images, you might want to create a custom image display function that includes the modular navigation tools.

This example creates a tool that accepts an image as an argument and displays the image in a scroll panel with an Overview tool and a Magnification box.

Note Because the toolbox scrollable navigation is incompatible with standard MATLAB figure window navigation tools, the example suppresses the toolbar and menu bar in the figure window.

```
function my_large_image_display(im)
```

```
% Create a figure without toolbar and menubar.
hfig = figure('Toolbar','none',...
             'Menubar', 'none',...
             'Name','My Large Image Display Tool',...
             'NumberTitle','off',...
             'IntegerHandle','off');

% Display the image in a figure with imshow.
himage = imshow(im)

% Add the scroll panel.
hpanel = imscrollpanel(hfig,himage)

% Position the scroll panel to accommodate the other tools.
set(hpanel,'Units','normalized','Position',[0 .1 1 .9])

% Add the Magnification box.
hMagBox = immagbox(hfig,himage);


% Position the Magnification box
pos = get(hMagBox,'Position');
set(hMagBox,'Position',[0 0 pos(3) pos(4)])

% Add the Overview tool.
hovervw = imoverview(himage)
```

To use the tool, pass it a large image that is already in the MATLAB workspace.

```
big_image = imread('peppers.png');
my_large_image_display(big_image)
```

The tool opens a figure window, displaying the image in a scroll panel with the other modular navigation tools: the Overview tool and the Magnification Box tool.



Custom Image Display Tool with Navigation Aids

Making Connections for Interactivity

When you create a modular tool and associate it with a target image, the tool automatically makes the necessary connections to the target image to do its job.

For example, the Pixel Information tool sets up a connection to the target image so that it can display the location and value of the pixel currently under the cursor. The Overview tool sets up a two-way connection to the target image:

- **Target image to the Overview tool** — If the visible portion of the image changes, by scrolling, panning, or by changing the magnification, the Overview tool changes the size and location of the detail rectangle to indicate the portion of the image that is now visible.

- **Overview tool to the target image** — If a user moves the detail rectangle in the Overview tool, the portion of the target image visible in the scroll panel is updated.

The modular tools accomplish this interactivity by using callback properties of the graphics objects. For example, the figure object supports a `WindowButtonMotionFcn` callback that executes whenever the mouse button is depressed.

In addition, some of the modular tools also support an application programmer interface (API). This API is a set of functions that let you get information about the tool as it operates and set up callbacks to get notification of events.

For example, the Magnification box supports a single API function: `setMagnification`. You can use this API function to set the magnification value displayed in the Magnification box. The Magnification box automatically notifies the scroll panel to change the magnification of the image based on the value. The scroll panel also supports an extensive set of API functions; see the `imscrollpanel` reference page for more information.

Example: Building an Image Comparison Tool

To illustrate how to use callbacks to make the connections required for interactions between tools, this example uses the Scroll Panel API to build a simple image comparison GUI. This custom tool displays two images side by side in scroll panels that are synchronized in location and magnification. The custom tool also includes an Overview tool and a Magnification box.

```
function my_image_compare_tool(left_image, right_image)

% Create the figure
hFig = figure('ToolBar','none',...
             'MenuBar','none',...
             'Name','My Image Compare Tool',...
             'NumberTitle','off',...
             'IntegerHandle','off');

% Display left image
subplot(121)
hImL = imshow(left_image);

% Display right image
```

```
subplot(122)
hImR = imshow(right_image);

% Create a scroll panel for left image
hSpL = imscrollpanel(hFig,hImL);
set(hSpL,'Units','normalized',...
    'Position',[0 0.1 .5 0.9])

% Create scroll panel for right image
hSpR = imscrollpanel(hFig,hImR);
set(hSpR,'Units','normalized',...
    'Position',[0.5 0.1 .5 0.9])

% Add a Magnification box
hMagBox = immagbox(hFig,hImL);
pos = get(hMagBox,'Position');
set(hMagBox,'Position',[0 0 pos(3) pos(4)])

%% Add an Overview tool
imoverview(hImL)

%% Get APIs from the scroll panels
apiL = iptgetapi(hSpL);
apiR = iptgetapi(hSpR)

%% Synchronize left and right scroll panels
apiL.setMagnification(apiR.getMagnification())
apiL.setVisibleLocation(apiR.setVisibleLocation())

% When magnification changes on left scroll panel,
% tell right scroll panel
apiL.addNewMagnificationCallback(apiR.setMagnification);


% When magnification changes on right scroll panel,
% tell left scroll panel
apiR.addNewMagnificationCallback(apiL.setMagnification);

% When location changes on left scroll panel,
% tell right scroll panel
apiL.addNewLocationCallback(apiR.setVisibleLocation);
```

```
% When location changes on right scroll panel,  
% tell left scroll panel  
apiR.addNewLocationCallback(apiL.setVisibleLocation);
```

The tool sets up a complex interaction between the scroll panels with just a few calls to Scroll Panel API functions. In the code, the tool specifies a callback function to execute every time the magnification changes. The function specified is the `setMagnification` API function of the other scroll panel. Thus, whenever the magnification changes in one of the scroll panels, the other scroll panel changes its magnification to match. The tool sets up a similar connection for position changes.

The following figure is a sequence diagram that shows the interaction between the two scroll panels set up by the comparison tool for both changes in magnification and location.




Scroll Panel Connections in Custom Image Comparison Tool

To use the image comparison tool, pass it two images as arguments.

```

left_image = imread('peppers.png');
right_image = edge(left_image(:,:,1),'canny');
my_image_compare_tool(left_image,right_image);
    
```

The tool opens a figure window, displaying the two images side by side, in separate scroll panels. The custom compare tool also includes an Overview tool and a Magnification box. When you move the detail rectangle in the Overview tool or change the magnification in one image, both images respond.



Custom Image Comparison Tool with Synchronized Scroll Panels

Creating Your Own Modular Tools

Because the toolbox uses an open architecture for the modular interactive tools, you can extend the toolbox by creating your own modular interactive tools, using standard Handle Graphics concepts and techniques. To help you create tools that integrate well with the existing modular interactive tools, the toolbox includes many utility functions that perform commonly needed tasks.

The utility functions can help check the input arguments to your tool, add callback functions to a callback list or remove them from a list, draw a rectangle over an image, and align figure windows in relation to a fixed window. The following table lists these utility functions in alphabetical order. See the function's reference page for more detailed information.

Utility Function	Description
<code>getimagemodel</code>	Retrieve image model objects from image handles
<code>getrangefromclass</code>	Get dynamic range of image based on its class
<code>imagemodel</code>	Access to properties of an image relevant to its display
<code>imattributes</code>	Return information about image attributes
<code>imgca</code>	Get handle to most recent current axis containing an image
<code>imgcf</code>	Get handle to most recent current figure containing an image
<code>imgetfile</code>	Displays Open Image dialog box
<code>imhandles</code>	Get all image handles within a handle
<code>impositionrect</code>	Create position rectangle that can be dragged
<code>iptaddcallback</code>	Add function handle to a callback list
<code>iptcheckconn</code>	Check validity of connectivity argument
<code>iptcheckhandle</code>	Check validity of image handle argument

Utility Function	Description
<code>iptcheckinput</code>	Check validity of input argument
<code>iptcheckmap</code>	Check validity of colormap argument
<code>iptchecknargin</code>	Check number of input arguments
<code>iptcheckstrs</code>	Check validity of string argument
<code>iptgetapi</code>	Get application programmer interface (API) for a handle
<code>ipticondir</code>	Return names of directories containing IPT and MATLAB icons
<code>iptnum2ordinal</code>	Convert positive integer to ordinal string
<code>iptremovecallback</code>	Delete function handle from callback list
<code>iptwindowalign</code>	Align figure windows

Spatial Transformations

This chapter describes the spatial transformation functions in the Image Processing Toolbox. Spatial transformations map pixel locations in an input image to new locations in an output image.

Terminology (p. 5-2)	Provides definitions of image processing terms used in this section
Interpolation (p. 5-3)	Defines interpolation, the process used to estimate the value of a pixel in an output image when that pixel does not appear in the input image
Image Resizing (p. 5-5)	Describes how to use the <code>imresize</code> function to change the size of an image
Image Rotation (p. 5-8)	Describes how to use the <code>imrotate</code> function to rotate an image
Image Cropping (p. 5-10)	Describes how to use the <code>imcrop</code> function to extract a rectangular portion of an image
Performing General Spatial Transformations (p. 5-11)	Describes the general spatial transformation capabilities of the toolbox

Terminology

An understanding of the following terms will help you to use this chapter.

Term	Definition
aliasing	Artifacts in an image that can appear as a result of reducing an image's size. When the size of an image is reduced, original pixels are downsampled to create fewer pixels. Aliasing that occurs as a result of size reduction normally appears as "stair-step" patterns (especially in high contrast images), or as "moire" (ripple-effect) patterns.
antialiasing	Any method for correcting aliasing (see above). The method discussed in this chapter is low-pass filtering (see below).
bicubic interpolation	Output pixel values are calculated from a weighted average of pixels in the nearest 4-by-4 neighborhood.
bilinear interpolation	Output pixel values are calculated from a weighted average of pixels in the nearest 2-by-2 neighborhood.
geometric operation	An operation that modifies the spatial relations between pixels in an image. Examples include resizing (growing or shrinking), rotating, and shearing.
interpolation	The process used to estimate an image value at a location in between image pixels.
nearest-neighbor interpolation	Output pixel values are assigned the value of the pixel that the point falls within. No other pixels are considered.

Interpolation

Interpolation is the process used to estimate an image value at a location in between image pixels. For example, if you resize an image so it contains more pixels than it did originally, the toolbox uses interpolation to determine the values for the additional pixels. The `imresize` and `imrotate` geometric functions use two-dimensional interpolation as part of the operations they perform. The `improfile` image analysis function also uses interpolation. See “Intensity Profile” on page 10-5 for information about this function.

Interpolation Methods

The Image Processing Toolbox provides three interpolation methods:

- Nearest-neighbor interpolation
- Bilinear interpolation
- Bicubic interpolation

The interpolation methods all work in a fundamentally similar way. In each case, to determine the value for an interpolated pixel, they find the point in the input image that the output pixel corresponds to. They then assign a value to the output pixel by computing a weighted average of some set of pixels in the vicinity of the point. The weightings are based on the distance each pixel is from the point.

The methods differ in the set of pixels that are considered:

- For nearest-neighbor interpolation, the output pixel is assigned the value of the pixel that the point falls within. No other pixels are considered.
- For bilinear interpolation, the output pixel value is a weighted average of pixels in the nearest 2-by-2 neighborhood.
- For bicubic interpolation, the output pixel value is a weighted average of pixels in the nearest 4-by-4 neighborhood.

The number of pixels considered affects the complexity of the computation. Therefore the bilinear method takes longer than nearest-neighbor interpolation, and the bicubic method takes longer than bilinear. However, the greater the number of pixels considered, the more accurate the effect is, so there is a tradeoff between processing time and quality.

Image Types

The functions that use interpolation take an argument that specifies the interpolation method. For most of these functions, the default method is nearest-neighbor interpolation. This method produces acceptable results for all image types, and is the only method that is appropriate for indexed images. For intensity and RGB images, however, you should generally specify bilinear or bicubic interpolation, because these methods produce better results than nearest-neighbor interpolation.

For RGB images, interpolation is performed on the red, green, and blue image planes individually.

For binary images, interpolation has effects that you should be aware of. If you use bilinear or bicubic interpolation, the computed values for the pixels in the output image will not all be 0 or 1. The effect on the resulting output image depends on the class of the input image:

- If the class of the input image is `double`, the output image is a grayscale image of class `double`. The output image is not binary, because it includes values other than 0 and 1.
- If the class of the input image is `uint8`, the output image is a binary image of class `uint8`. The interpolated pixel values are rounded off to 0 and 1 so the output image can be of class `uint8`.

Note For bicubic interpolation, you might need to clamp doubles to within the [0 1] range.

If you use nearest-neighbor interpolation, the result is always binary, because the values of the interpolated pixels are taken directly from pixels in the input image.

Image Resizing

To change the size of an image, use the `imresize` function. Using `imresize`, you can

- Specify the size of the output image
- Specify the interpolation method used
- Specify the filter to use to prevent aliasing

Specifying the Size of the Output Image

Using `imresize`, you can specify the size of the output image in two ways:

- By specifying the magnification factor to be used on the image
- By specifying the dimensions of the output image

Using the Magnification Factor

To enlarge an image, specify a magnification factor greater than 1. To reduce an image, specify a magnification factor between 0 and 1. For example, the command below increases the size of the image `I` by 1.25 times.

```
I = imread('circuit.tif');
J = imresize(I,1.25);
imshow(I)
figure, imshow(J)
```





Image Courtesy of Steve Decker and Shujaat Nadeem

Specifying the Size of the Output Image

You can specify the size of the output image by passing a vector that contains the number of rows and columns in the output image. The following command creates an output image, `Y`, with 100 rows and 150 columns.

```
Y = imresize(X,[100 150])
```

Note If the specified size does not produce the same aspect ratio as the input image, the output image is distorted.

Specifying the Interpolation Method

By default, `imresize` uses nearest-neighbor interpolation to determine the values of pixels in the output image, but you can specify other interpolation methods. This table lists the supported interpolation methods in order of complexity. See “Interpolation” on page 5-3 for more information about these methods.

Argument Value	Interpolation Method
'nearest'	Nearest-neighbor interpolation (the default)
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation

In this example, `imresize` uses the bilinear interpolation method.

```
Y = imresize(X,[100 150],'bilinear')
```

Using Filters to Prevent Aliasing

Reducing the size of an image can introduce artifacts, such as aliasing, in the output image because information is always lost when you reduce the size of an image. Aliasing appears as ripple patterns (called moiré patterns) in the output image.

When you reduce the size of the image using either bilinear or bicubic interpolation, `imresize` automatically applies a lowpass filter to the image before interpolation, to limit the impact of aliasing on the output image. You can specify the size of this filter or specify a different filter.

Note Even with lowpass filtering, resizing can introduce artifacts, because information is always lost when you reduce the size of an image.

The `imresize` function does not apply a lowpass filter if nearest-neighbor interpolation is used. Nearest-neighbor interpolation is primarily used for indexed images, and lowpass filtering is not appropriate for these images.

You can also specify a filter of your own creation. For more information about specifying a filter, see the reference page for `imresize`.

Image Rotation

To rotate an image, use the `imrotate` function. `imrotate` accepts two primary arguments:

- The image to be rotated
- The rotation angle

You specify the rotation angle in degrees. If you specify a positive value, `imrotate` rotates the image counterclockwise; if you specify a negative value, `imrotate` rotates the image clockwise. This example rotates the image `I` 35 degrees in the counterclockwise direction.

```
J = imrotate(I,35);
```

As optional arguments to `imrotate`, you can also specify

- The interpolation method
- The size of the output image



Specifying the Interpolation Method

By default, `imrotate` uses nearest-neighbor interpolation to determine the value of pixels in the output image, but you can specify other interpolation methods. This table lists the supported interpolation methods in order of complexity. See “Interpolation” on page 5-3 for more information about these methods.

Argument Value	Interpolation Method
'nearest'	Nearest-neighbor interpolation (the default)
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation

For example, these commands rotate an image 35° counterclockwise and use bilinear interpolation.

```
I = imread('circuit.tif');  
J = imrotate(I,35,'bilinear');  
imshow(I)  
figure, imshow(J)
```



Specifying the Size of the Output Image

By default, `imrotate` creates an output image large enough to include the entire original image. Pixels that fall outside the boundaries of the original image are set to 0 and appear as a black background in the output image. If you specify the text string 'crop' as an argument, `imrotate` crops the output image to be the same size as the input image. (See the reference page for `imrotate` for an example of cropping.)

Image Cropping


To extract a rectangular portion of an image, use the `imcrop` function. `imcrop` accepts two primary arguments:

- The image to be cropped
- The coordinates of a rectangle that defines the crop area

If you call `imcrop` without specifying the crop rectangle, you can specify the crop rectangle interactively. In this case, the cursor changes to crosshairs when it is over the image. Position the crosshairs over a corner of the crop region and press and hold the left mouse button. When you drag the crosshairs over the image you specify the rectangular crop region. `imcrop` draws a rectangle around the area you are selecting. When you release the mouse button, `imcrop` creates a new image from the selected region.

In this example, you display an image and call `imcrop`. The `imcrop` function displays the image in a figure window and waits for you to draw the cropping rectangle on the image. In the figure, the rectangle you select is shown in red. The example then calls `imshow` to view the cropped image.

```
imshow circuit.tif  
I = imcrop;  
imshow(I);
```



Performing General Spatial Transformations

To perform general two-dimensional (2-D) spatial transformations, use the `imtransform` function. (For information about performing advanced transformations, see “Advanced Spatial Transformation Techniques” on page 5-13.)

The `imtransform` function accepts two primary arguments:

- The image to be transformed
- A spatial transformation structure, called a `TFORM`, that specifies the type of transformation you want to perform

Specifying the Transformation Type

You specify the type of transformation you want to perform in a `TFORM` structure. There are two ways to create a `TFORM` structure:

- Using the `maketform` function
- Using the `cp2tform` function

Using `maketform`

When you use the `maketform` function, you can specify the type of transformation you want to perform. The following table lists the types of transformations `maketform` supports in alphabetical order.

Transformation	Description
'affine'	Transformation that can include translation, rotation, scaling, stretching, and shearing. Straight lines remain straight, and parallel lines remain parallel, but rectangles might become parallelograms.
'box'	Special case of an affine transformation where each dimension is shifted and scaled independently.
'composite'	Composition of two or more transformations.

Transformation	Description
'custom'	User-defined transformation, providing the forward and/or inverse functions that are called by <code>imtransform</code> .
'projective'	Transformation in which straight lines remain straight but parallel lines converge toward vanishing points. (The vanishing points can fall inside or outside the image — even at infinity.)

The 'custom' and 'composite' capabilities of `maketform` allow a virtually limitless variety of spatial transformations to be used with `imtransform` and/or `tformarray`.

Using `cp2tform`

You use `cp2tform` to create the `TFORM` when you want to perform a transformation that requires fitting of data points, such as a polynomial transformation. Chapter 6, “Image Registration,” explains how to use the `cp2tform` function to fit a 2-D transformation to a set of control points selected in a pair of images.


Note When used with `imtransform`, `TFORM` structures must define a 2-D spatial transformation. If an image contains more than two dimensions, such as an RGB image, the same 2-D transformation is automatically applied to all 2-D planes along the higher dimensions. To define an n-dimensional transformation, use the `tformarray` function.

Performing the Transformation

Once you define the transformation in a TFORM struct, you can perform the transformation by calling `imtransform`.

For example, this code uses `imtransform` to perform a projective transformation of a checkerboard image.

```
I = checkerboard(20,1,1);
figure; imshow(I)
T = maketform('projective',[1 1; 41 1; 41 41; 1 41],...
             [5 5; 40 5; 35 30; -10 30]);
R = makesampler('cubic','circular');
K = imtransform(I,T,R,'Size',[100 100],'XYScale',1);
figure, imshow(K)
```



The `imtransform` function options let you control many aspects of the transformation. For example, note how the transformed image appears to contain multiple copies of the original image. This is accomplished by using the 'Size' option, to make the output image larger than the input image, and then specifying a padding method that extends the input image by repeating the pixels in a circular pattern. The Image Processing Toolbox Image Transformation demos provide more examples of using the `imtransform` function, and related functions, to perform different types of spatial transformations.

Advanced Spatial Transformation Techniques

The following functions, when used in combination, provide a vast array of options for defining and working with 2-D, N-D, and mixed-D spatial transformations:

- `maketform`
- `fliptform`

- `tformfwd`
- `tforminv`
- `findbounds`
- `makesampler`
- `tformarray`
- `imtransform`

The `imtransform`, `findbounds`, and `tformarray` functions use the `tformfwd` and `tforminv` functions internally to encapsulate the forward transformations needed to determine the extent of an output image or array and/or to map the output pixels/array locations back to input locations. You can use `tformfwd` and `tforminv` to explore the geometric effects of a transformation by applying them to points and lines and plotting the results. They support a consistent handling of both image and pointwise data.

The previous example, “Performing the Transformation” on page 5-13, used the `makesampler` function with a standard interpolation method. You can also use it to obtain special effects or custom processing. For example, you could specify your own separable filtering/interpolation kernel, build a custom resampler around the MATLAB `interp2` or `interp3` functions, or even implement an advanced antialiasing technique.

And, as noted, you can use `tformarray` to work with arbitrary-dimensional array transformations. The arrays do not even need to have the same dimensions. The output can have either a lower or higher number of dimensions than the input.

For example, if you are sampling 3-D data on a 2-D slice or manifold, the input array might have a lower dimensionality. The output dimensionality might be higher, for example, if you combine multiple 2-D transformations into a single 2-D to 3-D operation.

Image Registration

This chapter describes the image registration capabilities of the Image Processing Toolbox. Image registration is the process of aligning two or more images of the same scene. Image registration is often used as a preliminary step in other image processing applications.

Terminology (p. 6-2)	Provides definitions of image processing terms used in this section
Registering an Image (p. 6-4)	Steps you through an example of the image registration process
Types of Supported Transformations (p. 6-12)	Lists the types of supported transformations
Selecting Control Points (p. 6-14)	Describes how to use the Control Point Selection Tool (<code>cpselect</code>) to select control points in pairs of images
Using Correlation to Improve Control Points (p. 6-31)	Describes how to use the <code>cpcorr</code> function to fine-tune your control point selections

Terminology

An understanding of the following terms will help you to use this chapter.

Term	Definition
aligned image	Output image after registration has been performed. The output image is derived by applying a transformation to the input image (see below) that brings it into alignment with the base image (see below).
base image	Image against which you compare the image to be registered. It is also often called the <i>reference</i> image.
control point pairs	Matching locations, also referred to as <i>landmarks</i> , in the input image and the base image.
distortion	Differences in one image as compared to another of the same subject. These differences might have occurred as a result of terrain relief and other changes in perspective when imaging the same scene from different viewpoints. Lens and other internal sensor distortions, or differences between sensors and sensor types, can also cause distortion.
global transformation	Transformation in which a single mathematical expression applies to an entire image.
input image	Image that you want to register. It is often called the <i>observed</i> image.

Term	Definition
local transformation	Transformation in which different mathematical expressions (usually differing in parameters rather than form) apply to different regions within an image.
spatial transformation	Mapping of locations of points in one image to new locations in another image.

Registering an Image

Image registration is the process of aligning two or more images of the same scene. Typically, one image, called the base image, is considered the reference to which the other images, called input images, are compared. The object of image registration is to bring the input image into alignment with the base image by applying a spatial transformation to the input image.

A spatial transformation maps locations in one image to new locations in another image. (For more details, see Chapter 5, “Spatial Transformations.”) Determining the parameters of the spatial transformation needed to bring the images into alignment is key to the image registration process.


Image registration is often used as a preliminary step in other image processing applications. For example, you can use image registration to align satellite images of the earth’s surface or images created by different medical diagnostic modalities (MRI and SPECT). After registration, you can compare features in the images to see how a river has migrated, how an area is flooded, or to see if a tumor is visible in an MRI or SPECT image.

Point Mapping

The Image Processing Toolbox provides tools to support point mapping to determine the parameters of the transformation required to bring an image into alignment with another image. In point mapping, you pick points in a pair of images that identify the same feature or landmark in the images. Then, a spatial mapping is inferred from the positions of these control points.

Note You might need to perform several iterations of this process, experimenting with different types of transformations, before you achieve a satisfactory result. In some cases, you might perform successive registrations, removing gross global distortions first, and then removing smaller local distortions in subsequent passes.

The following figure provides a graphic illustration of this process. This process is best understood by looking at an example. See “Example: Registering to a Digital Orthophoto” on page 6-5 for an extended example.



Overview of Image Registration Process

Example: Registering to a Digital Orthophoto

This example illustrates the steps involved in performing image registration using point mapping. These steps include:

- 1 Read the images into the MATLAB workspace.
- 2 Specify control point pairs in the images.
- 3 Save the control point pairs.

- 4 Fine-tune the control points using cross-correlation. (This is an optional step.)
- 5 Specify the type of transformation to be used and infer its parameters from the control point pairs.
- 6 Transform the unregistered image to bring it into alignment.

To illustrate this process, the example registers a digital aerial photograph to a digital orthophoto covering the same area. Both images are centered on the business district of West Concord, Massachusetts.

The aerial image is geometrically uncorrected: it includes camera perspective, terrain and building relief, and internal (lens) distortions, and it does not have any particular alignment or registration with respect to the earth.

The orthophoto, supplied by the Massachusetts Geographic Information System (MassGIS), has been orthorectified to remove camera, perspective, and relief distortions (via a specialized image transformation process). It is also georegistered (and geocoded)—the columns and rows of the digital orthophoto image are aligned to the axes of the Massachusetts State Plane coordinate system, each pixel center corresponds to a definite geographic location, and every pixel is 1 meter square in map units.

Step 1: Read the Images into MATLAB

In this example, the base image is `westconcordorthophoto.png`, the MassGIS georegistered orthophoto. It is a panchromatic (grayscale) image. The image to be registered is `westconcordaerial.png`, a digital aerial photograph supplied by mPower3/Emerge, and is a visible-color RGB image.

```
orthophoto = imread('westconcordorthophoto.png');  
figure, imshow(orthophoto)  
unregistered = imread('westconcordaerial.png');  
figure, imshow(unregistered)
```

You do not have to read the images into the MATLAB workspace. The `cpsselect` function accepts file specifications for grayscale images. However, if you want to use cross-correlation to tune your control point positioning, the images must be in the workspace.



Aerial Photo Image

Image Courtesy of mPower3/Emerge



Orthophoto Image

Image Courtesy of MassGIS

Step 2: Choose Control Points in the Images


The toolbox provides an interactive tool, called the Control Point Selection Tool, that you can use to pick pairs of corresponding control points in both images. Control points are landmarks that you can find in both images, like a road intersection, or a natural feature.

To start this tool, enter `cpselect` at the MATLAB prompt, specifying as arguments the input and base images.

Note The unregistered image is an RGB image. Because the Control Point Selection Tool only accepts grayscale images, the example passes only one plane of the color image to `cpselect`.

```
cpselect(unregistered(:,:,1),orthophoto)
```

The `cpselect` function displays two views of both the input image and the base image in which you can pick control points by pointing and clicking. For more information, see “Selecting Control Points” on page 6-14. This figure shows the Control Point Selection Tool with four pairs of control points selected. The number of control point pairs you pick is at least partially determined by the type of transformation you want to perform (specified in Step 5). See “Types of Supported Transformations” on page 6-12 for information about the minimum number of points required by each transformation.



Step 3: Save the Control Point Pairs to the MATLAB Workspace

In the Control Point Selection Tool, click the **File** menu and choose the **Save Points to Workspace** option. See “Saving Control Points” on page 6-29 for more information.

For example, the Control Point Selection Tool returns the following set of control points in the input image. These values represent spatial coordinates; the left column are x -coordinates, the right column are y -coordinates.

```
input_points =
    120.7086    93.9772
    319.2222    78.9202
    127.9838   291.6312
    352.0729   281.1445
```

Step 4: Fine-Tune the Control Point Pair Placement

This is an optional step that uses cross-correlation to adjust the position of the control points you selected with `cpselect`. See “Using Correlation to Improve Control Points” on page 6-31 for more information.

Note `cpcorr` can only adjust points for images that are the same scale and have the same orientation. Because the Concord image is rotated in relation to the base image, `cpcorr` cannot tune the control points. When it cannot tune the points, `cpcorr` returns the input points unmodified.

```
input_points_corr = cpcorr(input_points,base_points,...
                           unregistered(:,:,1),orthophoto)

input_points_corr =
    120.7086    93.9772
    319.2222    78.9202
    127.1046   289.8935
    352.0729   281.1445
```

Step 5: Specify the Type of Transformation and Infer Its Parameters

In this step, you pass the control points to the `cp2tform` function that determines the parameters of the transformation needed to bring the image into alignment. `cp2tform` is a data-fitting function that determines the transformation based on the geometric relationship of the control points. `cp2tform` returns the parameters in a geometric transformation structure, called a `TFORM` structure.

When you use `cp2tform`, you must specify the type of transformation you want to perform. The `cp2tform` function can infer the parameters for five types of

transformations. You must choose which transformation will correct the type of distortion present in the input image. See “Types of Supported Transformations” on page 6-12 for more information. Images can contain more than one type of distortion.

The predominant distortion in the aerial image of West Concord (the input image) results from the camera perspective. Ignoring terrain relief, which is minor in this area, image registration can correct for this using a projective transformation. The projective transformation also rotates the image into alignment with the map coordinate system underlying the base digital orthophoto image. (Given sufficient information about the terrain and camera, you could correct these other distortions at the same time by creating a composite transformation with `maketform`. See “Performing General Spatial Transformations” on page 5-11 for more information.)

```
mytform = cp2tform(input_points,base_points,'projective');
```

Step 6: Transform the Unregistered Image

As the final step in image registration, transform the input image to bring it into alignment with the base image. You use `imtransform` to perform the transformation, passing it the input image and the `TFORM` structure, which defines the transformation. `imtransform` returns the transformed image. For more information about using `imtransform`, see Chapter 5, “Spatial Transformations.”

```
registered = imtransform(unregistered,mytform)
```

Note `imtransform` applies the transformation defined in `mytform`, which is based on control points picked in only one plane of the RGB image, to all three planes of the input image.

Compare the transformed image to the base image to see how the registration came out.



Registered Image




Orthophoto Image



Types of Supported Transformations

The `cp2tform` function can infer the parameters for six types of transformations. This table lists the transformations in order of complexity, with examples of each type of distortion.

The first four transformations, 'linear conformal', 'affine', 'projective', and 'polynomial' are global transformations. In these transformations, a single mathematical expression applies to an entire image. The last two transformations, 'piecewise linear' and 'lwm' (local weighted mean), are local transformations. In these transformations, different mathematical expressions apply to different regions within an image.

When exploring how different transformations affect the images you are working with, try the global transformations first. If these transformations are not satisfactory, try the local transformations: the piecewise linear transformation first and then the local weighted mean transformation..

Transformation Type	Description	Minimum Control Points	Example
'linear conformal'	Use this transformation when shapes in the input image are unchanged, but the image is distorted by some combination of translation, rotation, and scaling. Straight lines remain straight, and parallel lines are still parallel.	2 pairs	
'affine'	Use this transformation when shapes in the input image exhibit shearing. Straight lines remain straight, and parallel lines remain parallel, but rectangles become parallelograms.	3 pairs	
'projective'	Use this transformation when the scene appears tilted. Straight lines remain straight, but parallel lines converge toward vanishing points (which might or might not fall within the image).	4 pairs	

Transformation Type	Description	Minimum Control Points	Example
'polynomial'	Use this transformation when objects in the image are curved. The higher the order of the polynomial, the better the fit, but the result can contain more curves than the base image.	6 pairs (order 2) 10 pairs (order 3) 16 pairs (order 4)	
'piecewise linear'	Use this transformation when parts of the image appear distorted differently.	4 pairs	
'lwm'	Use this transformation (local weighted mean), when the distortion varies locally and piecewise linear is not sufficient.	6 pairs (12 pairs recommended)	

Selecting Control Points

The toolbox includes an interactive tool that enables you to specify control points in the images you want to register. The tool displays the images side by side. When you are satisfied with the number and placement of the control points, you can save the control points.


Using the Control Point Selection Tool

To specify control points in a pair of images you want to register, use the Control Point Selection Tool, `cpselect`. The tool displays the image you want to register, called the *input* image, next to the image you want to compare it to, called the *base* image or *reference* image.

Specifying control points is a four-step process:

- 1** Start the tool, specifying the input image and the base image.
- 2** View the images, looking for visual elements that you can identify in both images. `cpselect` provides many ways to navigate around the image, panning and zooming to view areas of the image in more detail.
- 3** Specify matching control point pairs in the input image and the base image.
- 4** Save the control points in the MATLAB workspace.

The following figure shows the default appearance of the tool when you first start it.



Control Point Selection Tool

Starting the Control Point Selection Tool

To use the Control Point Selection Tool, enter the `cpselect` command at the MATLAB prompt. As arguments, specify the image you want to register (the input image), and the image you want to compare it to (the base image).

To illustrate, this code fragment reads an image into a variable, `moon_base`, in the MATLAB workspace. It then creates another version of the image with a

deliberate size distortion, called `moon_input`. This is the image that needs registration to remove the size distortion. The code then starts the `cpselect` tool, specifying the two images.


```
moon_base = imread('moon.tif');  
moon_input = imresize(moon_base, 1.2);  
cpselect(moon_input, moon_base);
```

The `cpselect` command has other optional arguments. For example, you can restart a control point selection session by including a `cpstruct` structure as the third argument. For more information about restarting sessions, see “Saving Control Points” on page 6-29. For complete details, see the `cpselect` reference page.

Default Views of the Images

When the Control Point Selection Tool starts, it contains four image display windows. The top two windows are called the **Detail** windows. These windows show a closeup view of a portion of the images you are working with. The input image is on the left and the base image is on the right. The two windows at the bottom of the interface are called the **Overview** windows. These windows show the images in their entirety, at the largest scale that fits the window. The input overview image is on the left and the base overview image is on the right.

Superimposed on the image in the **Overview** windows is a rectangle, called the detail rectangle. This rectangle defines the part of the image that is visible in the **Detail** window. By default, at startup, the detail rectangle covers one quarter of the entire image and is positioned over the center of the image.



Viewing the Images

By default, cpselect displays the entire base and input images in the **Overview** windows and displays a closeup view of a portion of these images in the **Detail** windows. However, to find visual elements that are common to both images, you might want to change the section of the image displayed in the detail view or zoom in on a part of the image to view it in more detail. The following sections describe the different ways to change your view of the images:

- “Using Scroll Bars to View Other Parts of an Image” on page 6-18
- “Using the Detail Rectangle to Change the View” on page 6-18
- “Panning the Image Displayed in the Detail Window” on page 6-18
- “Zooming In and Out on an Image” on page 6-19
- “Specifying the Magnification of the Images” on page 6-20

- “Locking the Relative Magnification of the Input and Base Images” on page 6-21

Using Scroll Bars to View Other Parts of an Image

To view parts of an image that are not visible in the **Detail** or **Overview** windows, use the scroll bars provided in each window.

As you scroll the image in the **Detail** window, note how the detail rectangle moves over the image in the **Overview** window. The position of the detail rectangle always shows the portion of the image in the **Detail** window.

Using the Detail Rectangle to Change the View

To get a closer view of any part of the image, move the detail rectangle in the **Overview** window over that section of the image. `cpselect` displays that section of the image in the **Detail** window at a higher magnification than the overview window.

To move the detail rectangle,

- 1 Click the Default Cursor button  in the toolbar.
- 2 Move the pointer into the detail rectangle. The cursor changes to the fleur shape, .
- 3 Press and hold the mouse button to drag the detail rectangle anywhere on the image.

Note As you move the detail rectangle over the image in the **Overview** window, the view of the image displayed in the **Detail** window changes.

Panning the Image Displayed in the Detail Window

To change the section of the image displayed in the **Detail** window, use the pan tool to move the image in the window.

To use the pan tool,

- 1 Click the Drag Images to Pan button  in the toolbar.

- 2 Move the pointer over the image in the **Detail** window. The cursor changes to the fleur shape,  .
- 3 Press and hold the mouse button and drag the image in the **Detail** window.


Note As you move the image in the **Detail** window, the detail rectangle in the **Overview** window moves.

Zooming In and Out on an Image

To enlarge an image to get a closer look or shrink an image to see the whole image in context, use the Zoom buttons on the button bar. (You can also zoom in or out on an image by changing the magnification. See “Specifying the Magnification of the Images” on page 6-20 for more information.)

To zoom in or zoom out on the base or input images,

- 1 Click the appropriate magnifying glass button.



- 2 Move the pointer over the image you want to zoom in or out on. The cursor changes to crosshairs,  .

You can zoom in or out on either the input or the base images, in either the **Detail** or **Overview** windows. To keep the relative magnifications of the base and input images synchronized, click the **Lock ratio** check box. See “Locking the Relative Magnification of the Input and Base Images” on page 6-21 for more information.

Note If you zoom in close on the image displayed in the **Overview** window, the detail rectangle might no longer be visible.

You can use the zoom tool in two ways:

- Position the cursor over a location in the image and click the mouse. With each click, `cpselect` changes the magnification of the image by a preset amount. (See “Specifying the Magnification of the Images” on page 6-20 for a list of some of these magnifications.) `cpselect` centers the new view of the image on the spot where you clicked.
- Alternatively, you can position the cursor over a location in the image and, while pressing and holding the mouse button, draw a rectangle defining the area you want to zoom in or out on. `cpselect` magnifies the image so that the chosen section fills the **Detail** window. `cpselect` resizes the detail rectangle in the **Overview** window as well.

Note When you zoom in or out on an image, notice how the magnification value changes.

Specifying the Magnification of the Images


To enlarge an image to get a closer look or to shrink an image to see the whole image in context, use the magnification edit box. (You can also use the **Zoom** buttons to enlarge or shrink an image. See “Zooming In and Out on an Image” on page 6-19 for more information.)

To change the magnification of an image,

- 1 Move the cursor into the magnification edit box of the window you want to change. The cursor changes to the text entry cursor.

Note Each **Detail** window and **Overview** window has its own magnification edit box.


- 2 Type a new value in the magnification edit box and press **Enter**, or click the menu associated with the edit box and choose from a list of preset magnifications. `cpselect` changes the magnification of the image and displays the new view in the appropriate window.



Locking the Relative Magnification of the Input and Base Images

To keep the relative magnification of the input and base images automatically synchronized in the **Detail** or **Overview** windows, click the **Lock Ratio** check box. The two **Detail** windows and the two **Overview** windows each have their own **Lock ratio** check boxes.

When the **Lock Ratio** check box is selected, `cpselect` changes the magnification of *both* the input and base images when you zoom in or out on either one of the images or specify a magnification value for either of the images.



Specifying Matching Control Point Pairs

The primary function of the Control Point Selection Tool is to enable you to pick control points in the image to be registered, the input image, and the image to which you are comparing it, the base image. When you start `cpselect`, the point selection tool is enabled, by default.

You specify control points by pointing and clicking in the input and base images, in either the **Detail** or the **Overview** windows. Each point you specify in the input image must have a match in the base image. The following sections describe the ways you can use the Control Point Selection Tool to choose control point pairs:

- “Picking Control Point Pairs Manually”
- “Using Control Point Prediction” on page 6-24

This section also describes how to move control points after you’ve created them and how to delete control points.

Picking Control Point Pairs Manually

To specify a pair of control points in your images,

- 1 Click the Control Point Selection button . Control point selection mode is active by default.
- 2 Position the cursor over a feature you have visually selected in any of the images displayed. The cursor changes to a pointing finger, .

You can pick control points in either of the **Detail** windows, input or base, or in either of the **Overview** windows, input or base. You also can work in either direction: input-to-base image, or base-to-input image.


- 3 Click the mouse button. `cpselect` places a control point symbol at the position you specified, in both the **Detail** window and the **Overview** window. (The appearance of the control point symbol indicates its current state. Initially, control points are in an active, unmatched state. See “Control Point States” on page 6-26 for more information.)

Note Depending on where in the image you pick control points, the symbol for the point might be visible in the **Overview** window, but not in the **Detail** window.

- 4** To create the match for this control point, move the cursor into the corresponding **Detail** or **Overview** window. For example, if you started in an input window, move the cursor to a base window.
- 5** Click the mouse button. `cpselect` places a control point symbol at the position you specified, in both the **Detail** and **Overview** windows. Because this control point completes a pair, the appearance of this symbol indicates an active, matched state. Note that the appearance of the first control point you selected (in step 3) also changes to an active, matched state.

You pick pairs of control points by moving from a view of the input image to a view of the base image, or vice versa. You can pick several control points in one view of the image, and then move to the corresponding window to locate their matches. To match an unmatched control point, select it to make it active, and then pick a point in the corresponding view window. When you select a match for a control point, the symbols for both points change to indicate their matched state. You can move or delete control points after you create them.

The following figure illustrates control points in several states.




Using Control Point Prediction

Instead of picking matching control points by moving the cursor between corresponding **Detail** or **Overview** windows, you can let the Control Point Selection Tool estimate the match for the control points you specify, automatically. The Control Point Selection Tool determines the position of the matching control point based on the geometric relationship of the previously selected control points.

Note By default, the Control Point Selection Tool does not include predicted points in the set of valid control points returned in `input_points` or `base_points`. To include predicted points, you must accept them by selecting the points and fine-tuning their position with the cursor. When you move a predicted point, the Control Point Selection Tool changes the symbol to indicate that it has changed to a standard control point. For more information, see “Moving Control Points” on page 6-27.

To illustrate point prediction, this figure shows four control points selected in the input image, where the points form the four corners of a square. (The control points selections in the figure do not attempt to identify any landmarks in the image.) The figure shows the picking of a fourth point, in the left window, and the corresponding predicted point in the right window. Note how the Control Point Selection Tool places the predicted point at the same location relative to the other control points, forming the bottom right corner of the square.



Note Because the Control Point Selection Tool predicts control point locations based on the locations of the previous control points, you cannot use point prediction until you have a minimum of two pairs of matched points. Until this minimum is met, the Control Point Prediction button is disabled.


To use control point prediction,

- 1 Click the Control Point Prediction button  .
- 2 Position the cursor anywhere in any of the images displayed. The cursor changes to a pointing finger,  .

You can pick control points in either of the **Detail** windows, input or base, or in either of the **Overview** windows, input or base. You also can work in either direction: input-to-base image or base-to-input image.

- 3 Click either mouse button. The Control Point Selection Tool places a control point symbol at the position you specified and places another control point symbol for a matching point in all the other windows. The symbol for the predicted point contains the letter “P,” indicating that it’s a predicted control point.

This figure illustrates predicted points in active unmatched, matched, and predicted states. For a complete description of all point states, see “Control Point States” on page 6-26.



Control Point States

The appearance of control point symbols indicates their current state. When you first pick a control point, its state is active and unmatched. When you pick

the match for a control point, the appearance of both symbols changes to indicate their matched status.

This table lists all the possible control point states with their symbols. cpselect displays this list in a separate window called a **Legend**. The Legend is visible by default, but you can control its visibility using the **Legend** option from the **View** menu.

Control Point States

Symbol	State	Description
	Active unmatched	The point is currently selected but does not have a matching point. This is the initial state of most points.
	Active matched	The point is currently selected and has a matching point.
	Active predicted	The point is a predicted point. If you move its position, the point changes to active matched state.
	Unmatched	The point is not selected and it is unmatched. You must select it before you can create its matching point.
	Matched	The point has a matching point.
	Predicted	This point was added by cpselect during point prediction.

Moving Control Points

To move a control point,


- 1 Click the Control Point Selection button  or the Default Cursor button .
- 2 Position the cursor over the control point you want to move.
- 3 Press and hold the mouse button and drag the control point. The state of the control point changes to active when you move it.

If you move a predicted control point, the state of the control point changes to a regular (nonpredicted) control point.

Deleting Control Points


To delete a control point, and optionally its matching point,

- 1 Click the Control Point Selection button  or the Default Cursor button .
- 2 Click the control point you want to delete. Its state changes to active. If the control point has a match, both points become active.
- 3 Delete the point (or points) using one of these methods:
 - Pressing the **Backspace** key
 - Pressing the **Delete** key
 - Choosing one of the delete options from the **Edit** menu
Using this menu you can delete individual points or pairs of matched points, in the input or base images.



Undoing and Redoing Control Point Selections

You can undo a deletion or series of deletions using the **Undo Delete** option on the cpselect **Edit** menu.




After undoing a deletion, you can delete the points again using the **Redo** option, also on the **Edit** menu.

Saving Control Points

After you specify control point pairs, you must save them in the MATLAB workspace to make them available for the next step in image registration, processing by `cp2tform`.

To save control points to the MATLAB workspace,

- 1 Select **File** on the Control Point Selection Tool menu bar.
- 2 Choose the **Save Points to Workspace** option. The Control Point Selection Tool displays this dialog box:



By default, the Control Point Selection Tool saves the x -coordinates and y -coordinates that specify the locations of the control points you selected in two arrays named `input_points` and `base_points`, although you can specify other names. These are n -by-2 arrays, where n is the number of valid control point pairs you selected. For example, this is an example of the `input_points` array if you picked four pairs of control points. The values in the left column represent the x -coordinates; the values in the right column represent the y -coordinates.


```
input_points =

    215.6667    262.3333
    225.7778    311.3333
    156.5556    340.1111
    270.8889    368.8889
```

Whenever you exit the Control Point Selection Tool, it asks if you want to save your control points.

Saving Your Control Point Selection Session

To save the current state of the Control Point Selection Tool, select the **Structure with all points** check box in the **Save Points to Workspace** dialog box.



This option saves the positions of all the control points you specified and their current states in a cpstruct structure.

```
cpstruct =

    inputPoints: [4x2 double]
    basePoints: [4x2 double]
    inputBasePairs: [4x2 double]
    ids: [4x1 double]
    inputIdPairs: [4x2 double]
    baseIdPairs: [4x2 double]
    isInputPredicted: [4x1 double]
    isBasePredicted: [4x1 double]
```

You can use the cpstruct to restart a control point selection session at the point where you left off.

This option is useful if you are picking many points over a long time and want to preserve unmatched and predicted points when you resume work. The Control Point Selection Tool does not include unmatched and predicted points in the input_points and base_points arrays.

To extract the arrays of valid control point coordinates from a cpstruct, use the cpstruct2pairs function.

Using Correlation to Improve Control Points

You might want to fine-tune the control points you selected using `cpselect`. Using cross-correlation, you can sometimes improve the points you selected by eye using the Control Point Selection Tool.

To use cross-correlation, pass sets of control points in the input and base images, along with the images themselves, to the `cpcorr` function.

```
input_pts_adj= cpcorr(input_points, base_points, input, base);
```

The `cpcorr` function defines 11-by-11 regions around each control point in the input image and around the matching control point in the base image, and then calculates the correlation between the values at each pixel in the region. Next, the `cpcorr` function looks for the position with the highest correlation value and uses this as the optimal position of the control point. The `cpcorr` function only moves control points up to 4 pixels based on the results of the cross-correlation.

Note Features in the two images must be at the same scale and have the same orientation. They cannot be rotated relative to each other.

If `cpcorr` cannot correlate some of the control points, it returns their values in `input_points` unmodified.

Linear Filtering and Filter Design

The Image Processing Toolbox provides a number of functions for designing and implementing two-dimensional linear filters for image data. This chapter describes these functions and how to use them effectively.

Terminology (p. 7-2)

Provides definitions of image processing terms used in this section

Linear Filtering (p. 7-4)

Provides an explanation of linear filtering and how it is implemented in the toolbox. This topic describes filtering in terms of the spatial domain, and is accessible to anyone doing image processing.

Filter Design (p. 7-17)

Discusses designing two-dimensional finite impulse response (FIR) filters. This section assumes you are familiar with working in the frequency domain.

Terminology

An understanding of the following terms will help you to use this chapter. Note that this table includes brief definitions of terms related to filter design; a detailed discussion of these terms and the theory behind filter design is outside the scope of this user's guide.

Term	Definition
convolution	Neighborhood operation in which each output pixel is a weighted sum of neighboring input pixels. The weights are defined by the convolution kernel. Image processing operations implemented with convolution include smoothing, sharpening, and edge enhancement.
convolution kernel	Matrix of weights used to perform convolution. A convolution kernel is a correlation kernel that has been rotated 180 degrees.
correlation	Neighborhood operation in which each output pixel is a weighted sum of neighboring input pixels. The weights are defined by the correlation kernel. Correlation is closely related mathematically to convolution.
correlation kernel	Matrix of weights used to perform correlation. The filter design functions in the Image Processing Toolbox return correlation kernels. A correlation kernel is a convolution kernel that has been rotated 180 degrees.
FIR filter	Filter whose response to a single point, or impulse, has finite extent. FIR stands for finite impulse response. An FIR filter can be implemented using convolution. All filter design functions in the Image Processing Toolbox return FIR filters.
frequency response	Mathematical function describing the gain of a filter in response to different input frequencies.
neighborhood operation	Operation in which each output pixel is computed from a set of neighboring input pixels. Convolution, dilation, and median filtering are examples of neighborhood operations.

Term	Definition
ripples	Oscillations around a constant value. The frequency response of a practical filter often has ripples where the frequency response of an ideal filter is flat.
window method	Filter design method that multiplies the ideal impulse response by a window function, which tapers the ideal impulse response. The resulting filter's frequency response approximates a desired frequency response.

Linear Filtering

Filtering is a technique for modifying or enhancing an image. For example, you can filter an image to emphasize certain features or remove other features.

Filtering is a *neighborhood operation*, in which the value of any given pixel in the output image is determined by applying some algorithm to the values of the pixels in the neighborhood of the corresponding input pixel. A pixel's neighborhood is some set of pixels, defined by their locations relative to that pixel. (See Chapter 14, "Neighborhood and Block Operations," for a general discussion of neighborhood operations.)

Linear filtering is filtering in which the value of an output pixel is a linear combination of the values of the pixels in the input pixel's neighborhood.

This section discusses linear filtering in MATLAB and the Image Processing Toolbox. It includes

- A description of filtering, using convolution and correlation
- A description of how to perform filtering using the `imfilter` function
- A discussion about using predefined filter types

See "Filter Design" on page 7-17 for information about how to design filters.

Convolution

Linear filtering of an image is accomplished through an operation called *convolution*. In convolution, the value of an output pixel is computed as a weighted sum of neighboring pixels. The matrix of weights is called the *convolution kernel*, also known as the *filter*.

For example, suppose the image is

```
A = [17  24  1  8  15
      23  5  7  14 16
       4  6 13 20 22
      10 12 19 21  3
      11 18 25  2  9]
```

and the convolution kernel is


$$h = \begin{bmatrix} 8 & 1 & 6 \\ 3 & 5 & 7 \\ 4 & 9 & 2 \end{bmatrix}$$

The following figure shows how to compute the (2,4) output pixel using these steps:

- 1 Rotate the convolution kernel 180 degrees about its center element.
- 2 Slide the center element of the convolution kernel so that it lies on top of the (2,4) element of A.
- 3 Multiply each weight in the rotated convolution kernel by the pixel of A underneath.
- 4 Sum the individual products from step 3.

Hence the (2,4) output pixel is

$$1 \cdot 2 + 8 \cdot 9 + 15 \cdot 4 + 7 \cdot 7 + 14 \cdot 5 + 16 \cdot 3 + 13 \cdot 6 + 20 \cdot 1 + 22 \cdot 8 = 575$$



Computing the (2,4) Output of Convolution


Correlation

The operation called *correlation* is closely related to convolution. In correlation, the value of an output pixel is also computed as a weighted sum of neighboring pixels. The difference is that the matrix of weights, in this case called the *correlation kernel*, is not rotated during the computation. The following figure shows how to compute the (2,4) output pixel of the correlation of A, assuming h is a correlation kernel instead of a convolution kernel, using these steps:

- 1 Slide the center element of the correlation kernel so that lies on top of the (2,4) element of A.
- 2 Multiply each weight in the correlation kernel by the pixel of A underneath.
- 3 Sum the individual products from step 3.

The (2,4) output pixel from the correlation is

$$1 \cdot 8 + 8 \cdot 1 + 15 \cdot 6 + 7 \cdot 3 + 14 \cdot 5 + 16 \cdot 7 + 13 \cdot 4 + 20 \cdot 9 + 22 \cdot 2 = 585$$




Computing the (2,4) Output of Correlation

Filtering Using `imfilter`

Filtering of images, either by correlation or convolution, can be performed using the toolbox function `imfilter`. This example filters an image with a 5-by-5 filter containing equal weights. Such a filter is often called an *averaging filter*.

```
I = imread('coins.png');  
h = ones(5,5) / 25;  
I2 = imfilter(I,h);  
imshow(I), title('Original Image');  
figure, imshow(I2), title('Filtered Image')
```



Original Image

Filtered Image

Data Types

The `imfilter` function handles data types similarly to the way the image arithmetic functions do, as described in “Image Arithmetic Saturation Rules” on page 2-28. The output image has the same data type, or numeric class, as the input image. The `imfilter` function computes the value of each output pixel using double-precision, floating-point arithmetic. If the result exceeds the range of the data type, the `imfilter` function truncates the result to that data type's allowed range. If it is an integer data type, `imfilter` rounds fractional values.

Because of the truncation behavior, you might sometimes want to consider converting your image to a different data type before calling `imfilter`. In this example, the output of `imfilter` has negative values when the input is of class `double`.

```
A = magic(5)

A =
    17    24     1     8    15
    23     5     7    14    16
     4     6    13    20    22
    10    12    19    21     3
    11    18    25     2     9

h = [-1 0 1]

h =
    -1     0     1

imfilter(A,h)

ans =
    24   -16   -16    14    -8
     5   -16     9     9   -14
     6     9    14     9  -20
    12     9     9   -16  -21
    18    14   -16   -16    -2
```

Notice that the result has negative values. Now suppose A is of class `uint8`, instead of `double`.

```
A = uint8(magic(5));
imfilter(A,h)

ans =
    24     0     0    14     0
     5     0     9     9     0
     6     9    14     9     0
    12     9     9     0     0
    18    14     0     0     0
```

Since the input to `imfilter` is of class `uint8`, the output also is of class `uint8`, and so the negative values are truncated to 0. In such cases, it might be appropriate to convert the image to another type, such as a signed integer type, `single`, or `double`, before calling `imfilter`.

Correlation and Convolution Options

The `imfilter` function can perform filtering using either correlation or convolution. It uses correlation by default, because the filter design functions, described in “Filter Design” on page 7-17, and the `fspecial` function, described in “Using Predefined Filter Types” on page 7-15, produce correlation kernels.

However, if you want to perform filtering using convolution instead, you can pass the string `'conv'` as an optional input argument to `imfilter`. For example:

```
A = magic(5);
h = [-1 0 1]
imfilter(A,h)    % filter using correlation


ans =
    24   -16   -16    14    -8
     5   -16     9     9   -14
     6     9    14     9   -20
    12     9     9   -16   -21
    18    14   -16   -16    -2

imfilter(A,h,'conv')    % filter using convolution

ans =
   -24    16    16   -14     8
    -5    16    -9    -9    14
    -6    -9   -14    -9    20
   -12    -9    -9    16    21
   -18   -14    16    16     2
```


Boundary Padding Options

When computing an output pixel at the boundary of an image, a portion of the convolution or correlation kernel is usually off the edge of the image, as illustrated in the following figure.



When the Values of the Kernel Fall Outside the Image


The `imfilter` function normally fills in these off-the-edge image pixels by assuming that they are 0. This is called zero padding and is illustrated in the following figure.




Zero Padding of Outside Pixels

When you filter an image, zero padding can result in a dark band around the edge of the image, as shown in this example.

```
I = imread('eight.tif');
h = ones(5,5) / 25;
I2 = imfilter(I,h);
imshow(I), title('Original Image');
figure, imshow(I2), title('Filtered Image with Black Border')
```



To eliminate the zero-padding artifacts around the edge of the image, `imfilter` offers an alternative boundary padding method called *border replication*. In border replication, the value of any pixel outside the image is determined by replicating the value from the nearest border pixel. This is illustrated in the following figure.



Replicated Boundary Pixels

To filter using border replication, pass the additional optional argument 'replicate' to `imfilter`.

```
I3 = imfilter(I,h,'replicate');  
figure, imshow(I3);  
title('Filtered Image with Border Replication')
```



Filtered Image with Border Replication

The `imfilter` function supports other boundary padding options, such as 'circular' and 'symmetric'. See the reference page for `imfilter` for details.

Multidimensional Filtering

The `imfilter` function can handle both multidimensional images and multidimensional filters. A convenient property of filtering is that filtering a three-dimensional image with a two-dimensional filter is equivalent to filtering each plane of the three-dimensional image individually with the same two-dimensional filter. This example shows how easy it is to filter each color plane of a truecolor image with the same filter:

1 Read in an RGB image and display it.

```
rgb = imread('peppers.png');  
imshow(rgb);
```



2 Filter the image and display it.

```
h = ones(5,5)/25;  
rgb2 = imfilter(rgb,h);  
figure, imshow(rgb2)
```



Relationship to Other Filtering Functions

MATLAB has several two-dimensional and multidimensional filtering functions. The function `filter2` performs two-dimensional correlation, `conv2` performs two-dimensional convolution, and `convn` performs multidimensional

convolution. Each of these filtering functions always converts the input to double, and the output is always double. These other filtering functions always assume the input is zero padded, and they do not support other padding options.

In contrast, the `imfilter` function does not convert input images to double. The `imfilter` function also offers a flexible set of boundary padding options, as described in “Boundary Padding Options” on page 7-10.

Using Predefined Filter Types

The `fspecial` function produces several kinds of predefined filters, in the form of correlation kernels. After creating a filter with `fspecial`, you can apply it directly to your image data using `imfilter`. This example illustrates applying an *unsharp masking* filter to an intensity image. The unsharp masking filter has the effect of making edges and fine detail in the image more crisp.

```
I = imread('moon.tif');
h = fspecial('unsharp');
I2 = imfilter(I,h);
imshow(I), title('Original Image')
figure, imshow(I2), title('Filtered Image')
```



Image Courtesy of Michael Myers

Original Image



Filtered Image

Filter Design

This section describes working in the frequency domain to design filters. Topics discussed include

- Finite impulse response (FIR) filters, the class of linear filter that the toolbox supports
- The frequency transformation method, which transforms a one-dimensional FIR filter into a two-dimensional FIR filter
- The frequency sampling method, which creates a filter based on a desired frequency response
- The windowing method, which multiplies the ideal impulse response with a window function to generate the filter
- Creating the desired frequency response matrix
- Computing the frequency response of a filter

This section assumes you are familiar with working in the frequency domain. This topic is discussed in many signal processing and image processing textbooks.

Note Most of the design methods described in this section work by creating a two-dimensional filter from a one-dimensional filter or window created using functions from the Signal Processing Toolbox. Although this toolbox is not required, you might find it difficult to design filters in the Image Processing Toolbox if you do not have the Signal Processing Toolbox as well.

FIR Filters

The Image Processing Toolbox supports one class of linear filter, the two-dimensional finite impulse response (FIR) filter. FIR filters have several characteristics that make them ideal for image processing in the MATLAB environment:

- FIR filters are easy to represent as matrices of coefficients.
- Two-dimensional FIR filters are natural extensions of one-dimensional FIR filters.

- There are several well-known, reliable methods for FIR filter design.
- FIR filters are easy to implement.
- FIR filters can be designed to have linear phase, which helps prevent distortion.

Another class of filter, the infinite impulse response (IIR) filter, is not as suitable for image processing applications. It lacks the inherent stability and ease of design and implementation of the FIR filter. Therefore, this toolbox does not provide IIR filter support.



Frequency Transformation Method

The frequency transformation method transforms a one-dimensional FIR filter into a two-dimensional FIR filter. The frequency transformation method preserves most of the characteristics of the one-dimensional filter, particularly the transition bandwidth and ripple characteristics. This method uses a *transformation matrix*, a set of elements that defines the frequency transformation.

The toolbox function `ftrans2` implements the frequency transformation method. This function's default transformation matrix produces filters with nearly circular symmetry. By defining your own transformation matrix, you can obtain different symmetries. (See Jae S. Lim, *Two-Dimensional Signal and Image Processing*, 1990, for details.)

The frequency transformation method generally produces very good results, as it is easier to design a one-dimensional filter with particular characteristics than a corresponding two-dimensional filter. For instance, the next example designs an optimal equiripple one-dimensional FIR filter and uses it to create a two-dimensional filter with similar characteristics. The shape of the one-dimensional frequency response is clearly evident in the two-dimensional response.

```
b = remez(10,[0 0.4 0.6 1],[1 1 0 0]);
h = ftrans2(b);
[H,w] = freqz(b,1,64,'whole');
colormap(jet(64))
plot(w/pi-1,fftshift(abs(H)))
figure, freqz2(h,[32 32])
```



One-Dimensional Frequency Response (left) and Corresponding Two-Dimensional Frequency Response (right)


Frequency Sampling Method

The frequency sampling method creates a filter based on a desired frequency response. Given a matrix of points that define the shape of the frequency response, this method creates a filter whose frequency response passes through those points. Frequency sampling places no constraints on the behavior of the frequency response between the given points; usually, the response ripples in these areas.

The toolbox function `fsamp2` implements frequency sampling design for two-dimensional FIR filters. `fsamp2` returns a filter `h` with a frequency response that passes through the points in the input matrix `Hd`. The example below creates an 11-by-11 filter using `fsamp2` and plots the frequency response of the resulting filter. (The `freqz2` function in this example calculates the two-dimensional frequency response of a filter. See “Computing the Frequency Response of a Filter” on page 7-23 for more information.)

```
Hd = zeros(11,11); Hd(4:8,4:8) = 1;
[f1,f2] = freqspace(11,'meshgrid');
mesh(f1,f2,Hd), axis([-1 1 -1 1 0 1.2]), colormap(jet(64))
h = fsamp2(Hd);
```

```
figure, freqz2(h,[32 32]), axis([-1 1 -1 1 0 1.2])
```



Desired Two-Dimensional Frequency Response (left) and Actual Two-Dimensional Frequency Response (right)

Notice the ripples in the actual frequency response, compared to the desired frequency response. These ripples are a fundamental problem with the frequency sampling design method. They occur wherever there are sharp transitions in the desired response.

You can reduce the spatial extent of the ripples by using a larger filter. However, a larger filter does not reduce the height of the ripples, and requires more computation time for filtering. To achieve a smoother approximation to the desired frequency response, consider using the frequency transformation method or the windowing method.

Windowing Method

The windowing method involves multiplying the ideal impulse response with a window function to generate a corresponding filter. Like the frequency sampling method, the windowing method produces a filter whose frequency response approximates a desired frequency response. The windowing method, however, tends to produce better results than the frequency sampling method.

The toolbox provides two functions for window-based filter design, `fwind1` and `fwind2`. `fwind1` designs a two-dimensional filter by using a two-dimensional


window that it creates from one or two one-dimensional windows that you specify. `fwind2` designs a two-dimensional filter by using a specified two-dimensional window directly.

`fwind1` supports two different methods for making the two-dimensional windows it uses:

- Transforming a single one-dimensional window to create a two-dimensional window that is nearly circularly symmetric, by using a process similar to rotation
- Creating a rectangular, separable window from two one-dimensional windows, by computing their outer product

The example below uses `fwind1` to create an 11-by-11 filter from the desired frequency response `Hd`. Here, the `hamming` function from the Signal Processing Toolbox is used to create a one-dimensional window, which `fwind1` then extends to a two-dimensional window.

```
Hd = zeros(11,11); Hd(4:8,4:8) = 1;
[f1,f2] = freqspace(11,'meshgrid');
mesh(f1,f2,Hd), axis([-1 1 -1 1 0 1.2]), colormap(jet(64))
h = fwind1(Hd,hamming(11));
figure, freqz2(h,[32 32]), axis([-1 1 -1 1 0 1.2])
```




Desired Two-Dimensional Frequency Response (left) and Actual Two-Dimensional Frequency Response (right)

Creating the Desired Frequency Response Matrix

The filter design functions `fsamp2`, `fwind2`, and `fwind2` all create filters based on a desired frequency response magnitude matrix. You can create an appropriate desired frequency response matrix using the `freqspace` function. `freqspace` returns correct, evenly spaced frequency values for any size response. If you create a desired frequency response matrix using frequency points other than those returned by `freqspace`, you might get unexpected results, such as nonlinear phase.

For example, to create a circular ideal lowpass frequency response with cutoff at 0.5, use

```
[f1,f2] = freqspace(25,'meshgrid');  
Hd = zeros(25,25); d = sqrt(f1.^2 + f2.^2) < 0.5;  
Hd(d) = 1;  
mesh(f1,f2,Hd)
```



Ideal Circular Lowpass Frequency Response

Note that for this frequency response, the filters produced by `fsamp2`, `fwind1`, and `fwind2` are real. This result is desirable for most image processing applications. To achieve this in general, the desired frequency response should be symmetric about the frequency origin ($f_1 = 0$, $f_2 = 0$).


Computing the Frequency Response of a Filter

The `freqz2` function computes the frequency response for a two-dimensional filter. With no output arguments, `freqz2` creates a mesh plot of the frequency response. For example, consider this FIR filter,

```
h =[0.1667    0.6667    0.1667
    0.6667   -3.3333    0.6667
    0.1667    0.6667    0.1667];
```

This command computes and displays the 64-by-64 point frequency response of `h`.

```
freqz2(h)
```



Frequency Response of a Two-Dimensional Filter

To obtain the frequency response matrix `H` and the frequency point vectors `f1` and `f2`, use output arguments

```
[H,f1,f2] = freqz2(h);
```

`freqz2` normalizes the frequencies `f1` and `f2` so that the value 1.0 corresponds to half the sampling frequency, or π radians.

For a simple `m`-by-`n` response, as shown above, `freqz2` uses the two-dimensional fast Fourier transform function `fft2`. You can also specify vectors of arbitrary frequency points, but in this case `freqz2` uses a slower algorithm.

See “Fourier Transform” on page 8-3 for more information about the fast Fourier transform and its application to linear filtering and filter design.

Transforms

The usual mathematical representation of an image is a function of two spatial variables: $f(x, y)$. The value of the function at a particular location (x, y) represents the intensity of the image at that point. The term *transform* refers to an alternative mathematical representation of an image.

This chapter defines several important transforms and shows examples of their application to image processing.

Terminology (p. 8-2)	Provides definitions of image processing terms used in this section
Fourier Transform (p. 8-3)	Defines the Fourier transform and some of its applications in image processing
Discrete Cosine Transform (p. 8-17)	Describes the discrete cosine transform (DCT) of an image and its application, particularly in image compression
Radon Transform (p. 8-21)	Describes how the Image Processing Toolbox <code>radon</code> function computes <i>projections</i> of an image matrix along specified directions
Fan-Beam Projection Data (p. 8-36)	Describes how the Image Processing Toolbox <code>radon</code> function computes <i>projections</i> of an image matrix along specified directions

Terminology

An understanding of the following terms will help you to use this chapter. Note that this table includes brief definitions of terms related to transforms; a detailed discussion of these terms and the theory behind transforms is outside the scope of this user's guide.

Term	Definition
discrete transform	Transform whose input and output values are discrete samples, making it convenient for computer manipulation. Discrete transforms implemented by MATLAB and the Image Processing Toolbox include the discrete Fourier transform (DFT) and the discrete cosine transform (DCT).
frequency domain	Domain in which an image is represented by a sum of periodic signals with varying frequency.
inverse transform	Operation that when performed on a transformed image produces the original image.
spatial domain	Domain in which an image is represented by intensities at given points in space. This is the most common representation for image data.
transform	Alternative mathematical representation of an image. For example, the Fourier transform is a representation of an image as a sum of complex exponentials of varying magnitudes, frequencies, and phases. Transforms are useful for a wide range of purposes, including convolution, enhancement, feature detection, and compression.

Fourier Transform

The Fourier transform is a representation of an image as a sum of complex exponentials of varying magnitudes, frequencies, and phases. The Fourier transform plays a critical role in a broad range of image processing applications, including enhancement, analysis, restoration, and compression.

This section includes the following subsections:

- “Definition of Fourier Transform”
- “Discrete Fourier Transform” on page 8-8, including a discussion of fast Fourier transform
- “Applications of the Fourier Transform” on page 8-11 (sample applications using Fourier transforms)

Definition of Fourier Transform

If $f(m, n)$ is a function of two discrete spatial variables m and n , then the *two-dimensional Fourier transform* of $f(m, n)$ is defined by the relationship

$$F(\omega_1, \omega_2) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} f(m, n) e^{-j\omega_1 m} e^{-j\omega_2 n}$$

The variables ω_1 and ω_2 are frequency variables; their units are radians per sample. $F(\omega_1, \omega_2)$ is often called the *frequency-domain* representation of $f(m, n)$. $F(\omega_1, \omega_2)$ is a complex-valued function that is periodic both in ω_1 and ω_2 , with period 2π . Because of the periodicity, usually only the range $-\pi \leq \omega_1, \omega_2 \leq \pi$ is displayed. Note that $F(0, 0)$ is the sum of all the values of $f(m, n)$. For this reason, $F(0, 0)$ is often called the *constant component* or *DC component* of the Fourier transform. (DC stands for direct current; it is an electrical engineering term that refers to a constant-voltage power source, as opposed to a power source whose voltage varies sinusoidally.)


The inverse two-dimensional Fourier transform is given by

$$f(m, n) = \frac{1}{4\pi^2} \int_{\omega_1=-\pi}^{\pi} \int_{\omega_2=-\pi}^{\pi} F(\omega_1, \omega_2) e^{j\omega_1 m} e^{j\omega_2 n} d\omega_1 d\omega_2$$

Roughly speaking, this equation means that $f(m, n)$ can be represented as a sum of an infinite number of complex exponentials (sinusoids) with different frequencies. The magnitude and phase of the contribution at the frequencies (ω_1, ω_2) are given by $F(\omega_1, \omega_2)$.


Visualizing the Fourier Transform

To illustrate, consider a function $f(m, n)$ that equals 1 within a rectangular region and 0 everywhere else. To simplify the diagram, $f(m, n)$ is shown as a continuous function, even though the variables m and n are discrete.



Rectangular Function


The following figure shows, as a mesh plot, the magnitude of the Fourier transform, $|F(\omega_1, \omega_2)|$, of the rectangular function shown in the preceding figure. The mesh plot of the magnitude is a common way to visualize the Fourier transform.



Magnitude Image of a Rectangular Function

The peak at the center of the plot is $F(0, 0)$, which is the sum of all the values in $f(m, n)$. The plot also shows that $F(\omega_1, \omega_2)$ has more energy at high horizontal frequencies than at high vertical frequencies. This reflects the fact that horizontal cross sections of $f(m, n)$ are narrow pulses, while vertical cross sections are broad pulses. Narrow pulses have more high-frequency content than broad pulses.


Another common way to visualize the Fourier transform is to display $\log|F(\omega_1, \omega_2)|$ as an image, as shown.



Log of the Fourier Transform of a Rectangular Function

Using the logarithm helps to bring out details of the Fourier transform in regions where $F(\omega_1, \omega_2)$ is very close to 0.

Examples of the Fourier transform for other simple shapes are shown below.



Fourier Transforms of Some Simple Shapes

Discrete Fourier Transform

Working with the Fourier transform on a computer usually involves a form of the transform known as the discrete Fourier transform (DFT). There are two principal reasons for using this form:

- The input and output of the DFT are both discrete, which makes it convenient for computer manipulations.
- There is a fast algorithm for computing the DFT known as the fast Fourier transform (FFT).

The DFT is usually defined for a discrete function $f(m, n)$ that is nonzero only over the finite region $0 \leq m \leq M - 1$ and $0 \leq n \leq N - 1$. The two-dimensional M -by- N DFT and inverse M -by- N DFT relationships are given by

$$F(p, q) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) e^{-j(2\pi/M)pm} e^{-j(2\pi/N)qn} \quad \begin{array}{l} p = 0, 1, \dots, M-1 \\ q = 0, 1, \dots, N-1 \end{array}$$

$$f(m, n) = \frac{1}{MN} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} F(p, q) e^{j(2\pi/M)pm} e^{j(2\pi/N)qn} \quad \begin{array}{l} m = 0, 1, \dots, M-1 \\ n = 0, 1, \dots, N-1 \end{array}$$

The values $F(p, q)$ are the DFT coefficients of $f(m, n)$. The zero-frequency coefficient, $F(0, 0)$, is often called the “DC component.” DC is an electrical engineering term that stands for direct current. (Note that matrix indices in MATLAB always start at 1 rather than 0; therefore, the matrix elements $f(1, 1)$ and $F(1, 1)$ correspond to the mathematical quantities $f(0, 0)$ and $F(0, 0)$, respectively.)

The MATLAB functions `fft`, `fft2`, and `fftn` implement the fast Fourier transform algorithm for computing the one-dimensional DFT, two-dimensional DFT, and N -dimensional DFT, respectively. The functions `ifft`, `ifft2`, and `ifftn` compute the inverse DFT.

Relationship to the Fourier Transform


The DFT coefficients $F(p, q)$ are samples of the Fourier transform $F(\omega_1, \omega_2)$.

$$F(p, q) = F(\omega_1, \omega_2) \Big|_{\substack{\omega_1 = 2\pi p/M \\ \omega_2 = 2\pi q/N}} \quad \begin{array}{l} p = 0, 1, \dots, M-1 \\ q = 0, 1, \dots, N-1 \end{array}$$

Example


- 1 Construct a matrix f that is similar to the function $f(m, n)$ in the example in “Definition of Fourier Transform” on page 8-3. Remember that $f(m, n)$ is equal to 1 within the rectangular region and 0 elsewhere. Use a binary image to represent $f(m, n)$.

```
f = zeros(30,30);
f(5:24,13:17) = 1;
imshow(f, 'notruesize')
```



- 2 Compute and visualize the 30-by-30 DFT of f with these commands.

```
F = fft2(f);
F2 = log(abs(F));
imshow(F2, [-1 5], 'notruesize'); colormap(jet); colorbar
```



Discrete Fourier Transform Computed Without Padding


This plot differs from the Fourier transform displayed in “Visualizing the Fourier Transform” on page 8-4. First, the sampling of the Fourier transform is much coarser. Second, the zero-frequency coefficient is displayed in the upper left corner instead of the traditional location in the center.

- 3 To obtain a finer sampling of the Fourier transform, add zero padding to f when computing its DFT. The zero padding and DFT computation can be performed in a single step with this command.

```
F = fft2(f,256,256);
```

This command zero-pads f to be 256-by-256 before computing the DFT.

```
imshow(log(abs(F)),[-1 5]); colormap(jet); colorbar
```



Discrete Fourier Transform Computed with Padding

- 4 The zero-frequency coefficient, however, is still displayed in the upper left corner rather than the center. You can fix this problem by using the function `fftshift`, which swaps the quadrants of `F` so that the zero-frequency coefficient is in the center.

```
F = fft2(f,256,256);
F2 = fftshift(F);
imshow(log(abs(F2)),[-1 5]); colormap(jet); colorbar
```

The resulting plot is identical to the one shown in “Visualizing the Fourier Transform” on page 8-4.

Applications of the Fourier Transform


This section presents a few of the many image processing-related applications of the Fourier transform.

Frequency Response of Linear Filters

The Fourier transform of the impulse response of a linear filter gives the frequency response of the filter. The function `freqz2` computes and displays a

filter's frequency response. The frequency response of the Gaussian convolution kernel shows that this filter passes low frequencies and attenuates high frequencies.

```
h = fspecial('gaussian');  
freqz2(h)
```



Frequency Response of a Gaussian Filter

See “Linear Filtering and Filter Design” on page 7-1 for more information about linear filtering, filter design, and frequency responses.

Fast Convolution

A key property of the Fourier transform is that the multiplication of two Fourier transforms corresponds to the convolution of the associated spatial functions. This property, together with the fast Fourier transform, forms the basis for a fast convolution algorithm.

Note The FFT-based convolution method is most often used for large inputs. For small inputs it is generally faster to use `imfilter`.

To illustrate, this example performs the convolution of A and B, where A is an M-by-N matrix and B is a P-by-Q matrix:

- 1 Create two matrices.

```
A = magic(3);
B = ones(3);
```

- 2 Zero-pad A and B so that they are at least (M+P-1)-by-(N+Q-1). (Often A and B are zero-padded to a size that is a power of 2 because `fft2` is fastest for these sizes.) The example pads the matrices to be 8-by-8.

```
A(8,8) = 0;
B(8,8) = 0;
```

- 3 Compute the two-dimensional DFT of A and B using `fft2`.

- 4 Multiply the two DFTs together.

- 5 Compute the inverse two-dimensional DFT of the result using `ifft2`.

The following code performs steps 3, 4, and 5 in the procedure.

```
C = ifft2(fft2(A).*fft2(B));
```

- 6 Extract the nonzero portion of the result and remove the imaginary part caused by roundoff error.

```
C = C(1:5,1:5);
C = real(C)
```

```
C =
```

```

 8.0000    9.0000   15.0000    7.0000    6.0000
11.0000   17.0000   30.0000   19.0000   13.0000
15.0000   30.0000   45.0000   30.0000   15.0000
 7.0000   21.0000   30.0000   23.0000    9.0000
 4.0000   13.0000   15.0000   11.0000    2.0000
```

Locating Image Features

The Fourier transform can also be used to perform correlation, which is closely related to convolution. Correlation can be used to locate features within an image; in this context correlation is often called *template matching*.

This example illustrates how to use correlation to locate occurrences of the letter “a” in an image containing text:

- 1 Read in the sample image.

```
bw = imread('text.png');
```

- 2 Create a template for matching by extracting the letter “a” from the image.

```
a = bw(32:45,88:98);
```

You can also create the template image by using the interactive version of `imcrop`, using the `pixval` function to determine the coordinates of features in an image.

The following figure shows both the original image and the template.

```
imshow(bw);  
figure, imshow(a);
```



Image (left) and the Template to Correlate (right)

- 3 Compute the correlation of the template image `a` with the original image `bw` by rotating the template image by 180° and then using the FFT-based convolution technique described in “Fast Convolution” on page 8-12. (Convolution is equivalent to correlation if you rotate the convolution kernel by 180° .) To match the template to the image, use the `fft2` and `ifft2` functions.

```
C = real(ifft2(fft2(bw) .* fft2(rot90(a,2),256,256)));
```

The following image shows the result of the correlation. Bright peaks in the image correspond to occurrences of the letter.

```
figure, imshow(C,[]) % Scale image to appropriate display range.
```



Correlated Image

- 4** To view the locations of the template in the image, find the maximum pixel value and then define a threshold value that is less than this maximum. The locations of these peaks are indicated by the white spots in the thresholded correlation image. (To make the locations easier to see in this figure, the thresholded image has been dilated to enlarge the size of the points.)


```
max(C(:))
```

```
ans =
```

```
68.0000
```

```
thresh = 60; % Use a threshold that's a little less than max.
```

```
figure, imshow(C > thresh)% Display showing pixels over  
threshold.
```



Correlated, Thresholded Image Showing Template Locations

Discrete Cosine Transform

The discrete cosine transform (DCT) represents an image as a sum of sinusoids of varying magnitudes and frequencies. The `dct2` function in the Image Processing Toolbox computes the two-dimensional discrete cosine transform (DCT) of an image. The DCT has the property that, for a typical image, most of the visually significant information about the image is concentrated in just a few coefficients of the DCT. For this reason, the DCT is often used in image compression applications. For example, the DCT is at the heart of the international standard lossy image compression algorithm known as JPEG. (The name comes from the working group that developed the standard: the Joint Photographic Experts Group.)

The two-dimensional DCT of an M -by- N matrix A is defined as follows.

$$B_{pq} = \alpha_p \alpha_q \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} A_{mn} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{array}{l} 0 \leq p \leq M-1 \\ 0 \leq q \leq N-1 \end{array}$$

$$x_p = \begin{cases} 1/\sqrt{M}, & p = 0 \\ \sqrt{2/M}, & 1 \leq p \leq M-1 \end{cases} \quad \alpha_q = \begin{cases} 1/\sqrt{N}, & q = 0 \\ \sqrt{2/N}, & 1 \leq q \leq N-1 \end{cases}$$

The values B_{pq} are called the *DCT coefficients* of A . (Note that matrix indices in MATLAB always start at 1 rather than 0; therefore, the MATLAB matrix elements $A(1,1)$ and $B(1,1)$ correspond to the mathematical quantities A_{00} and B_{00} , respectively.)

The DCT is an invertible transform, and its inverse is given by


$$A_{mn} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \alpha_p \alpha_q B_{pq} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{array}{l} 0 \leq m \leq M-1 \\ 0 \leq n \leq N-1 \end{array}$$

$$x_p = \begin{cases} 1/\sqrt{M}, & p = 0 \\ \sqrt{2/M}, & 1 \leq p \leq M-1 \end{cases} \quad \alpha_q = \begin{cases} 1/\sqrt{N}, & q = 0 \\ \sqrt{2/N}, & 1 \leq q \leq N-1 \end{cases}$$

The inverse DCT equation can be interpreted as meaning that any M-by-N matrix A can be written as a sum of MN functions of the form

$$\alpha_p \alpha_q \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{array}{l} 0 \leq p \leq M-1 \\ 0 \leq q \leq N-1 \end{array}$$

These functions are called the *basis functions* of the DCT. The DCT coefficients B_{pq} , then, can be regarded as the *weights* applied to each basis function. For 8-by-8 matrices, the 64 basis functions are illustrated by this image.



The 64 Basis Functions of an 8-by-8 Matrix

Horizontal frequencies increase from left to right, and vertical frequencies increase from top to bottom. The constant-valued basis function at the upper left is often called the *DC basis function*, and the corresponding DCT coefficient B_{00} is often called the *DC coefficient*.

The DCT Transform Matrix

The Image Processing Toolbox offers two different ways to compute the DCT. The first method is to use the function `dct2`. `dct2` uses an FFT-based algorithm for speedy computation with large inputs. The second method is to use the DCT *transform matrix*, which is returned by the function `dctmtx` and might be more

efficient for small square inputs, such as 8-by-8 or 16-by-16. The M-by-M transform matrix T is given by

$$T_{pq} = \begin{cases} \frac{1}{\sqrt{M}} & p = 0, \quad 0 \leq q \leq M - 1 \\ \sqrt{\frac{2}{M}} \cos \frac{\pi(2q+1)p}{2M} & 1 \leq p \leq M - 1, \quad 0 \leq q \leq M - 1 \end{cases}$$

For an M-by-M matrix A , T^*A is an M-by-M matrix whose columns contain the one-dimensional DCT of the columns of A . The two-dimensional DCT of A can be computed as $B=T^*A^*T'$. Since T is a real orthonormal matrix, its inverse is the same as its transpose. Therefore, the inverse two-dimensional DCT of B is given by $T' * B^* T$.

DCT and Image Compression

In the JPEG image compression algorithm, the input image is divided into 8-by-8 or 16-by-16 blocks, and the two-dimensional DCT is computed for each block. The DCT coefficients are then quantized, coded, and transmitted. The JPEG receiver (or JPEG file reader) decodes the quantized DCT coefficients, computes the inverse two-dimensional DCT of each block, and then puts the blocks back together into a single image. For typical images, many of the DCT coefficients have values close to zero; these coefficients can be discarded without seriously affecting the quality of the reconstructed image.

The example code below computes the two-dimensional DCT of 8-by-8 blocks in the input image, discards (sets to zero) all but 10 of the 64 DCT coefficients in each block, and then reconstructs the image using the two-dimensional inverse DCT of each block. The transform matrix computation method is used.

```
I = imread('cameraman.tif');
I = im2double(I);
T = dctmtx(8);
B = blkproc(I,[8 8], 'P1*x*P2',T,T');
mask = [1  1  1  1  0  0  0  0
        1  1  1  0  0  0  0  0
        1  1  0  0  0  0  0  0
        1  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0]
```

```
        0  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0];
B2 = blkproc(B,[8 8], 'P1.*x',mask);
I2 = blkproc(B2,[8 8], 'P1*x*P2',T',T);
imshow(I), figure, imshow(I2)
```




Image Courtesy of MIT



Although there is some loss of quality in the reconstructed image, it is clearly recognizable, even though almost 85% of the DCT coefficients were discarded. To experiment with discarding more or fewer coefficients, and to apply this technique to other images, try running the demo function `dctdemo`.

Radon Transform


The `radon` function in the Image Processing Toolbox computes *projections* of an image matrix along specified directions. A projection of a two-dimensional function $f(x,y)$ is a set of line integrals. The `radon` function computes the line integrals from multiple sources along parallel paths, or *beams*, in a certain direction. The beams are spaced 1 pixel unit apart. To represent an image, the `radon` function takes multiple, parallel-beam projections of the image from different angles by rotating the source around the center of the image. The following figure shows a single projection at a specified rotation angle.



Parallel-Beam Projection at Rotation Angle Theta

Note For information about creating projection data from line integrals along paths that radiate from a single source, called fan-beam projections, see “Fan-Beam Projection Data” on page 8-36. To convert parallel-beam projection data to fan-beam projection data, use the `para2fan` function.

For example, the line integral of $f(x,y)$ in the vertical direction is the projection of $f(x,y)$ onto the x -axis; the line integral in the horizontal direction is the projection of $f(x,y)$ onto the y -axis. The following figure shows horizontal and vertical projections for a simple two-dimensional function.



Horizontal and Vertical Projections of a Simple Function


Projections can be computed along any angle θ . In general, the Radon transform of $f(x,y)$ is the line integral of f parallel to the y' -axis

$$R_{\theta}(x') = \int_{-\infty}^{\infty} f(x' \cos \theta - y' \sin \theta, x' \sin \theta + y' \cos \theta) dy'$$

where

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

The following figure illustrates the geometry of the Radon transform.



Geometry of the Radon Transform

Plotting the Radon Transform


You can compute the Radon transform of an image I for the angles specified in the vector θ using the `radon` function with this syntax.

```
[R, xp] = radon(I, theta);
```


The columns of R contain the Radon transform for each angle in θ . The vector xp contains the corresponding coordinates along the x' -axis. The center pixel of I is defined to be `floor((size(I)+1)/2)`; this is the pixel on the x' -axis corresponding to $x' = 0$.

The commands below compute and plot the Radon transform at 0° and 45° of an image containing a single square object. xp is the same for all projection angles.

```
I = zeros(100,100);  
I(25:75, 25:75) = 1;  
imshow(I)
```




```
[R,xp] = radon(I,[0 45]);  
figure; plot(xp,R(:,1)); title('R_{0^o} (x\prime)')
```



Radon Transform of a Square Function at 0 Degrees

```
figure; plot(xp,R(:,2)); title('R_{45^o} (x\prime)')
```




Radon Transform of a Square Function at 45 Degrees

Viewing the Radon Transform as an Image

The Radon transform for a large number of angles is often displayed as an image. In this example, the Radon transform for the square image is computed at angles from 0° to 180° , in 1° increments.

```
theta = 0:180;
[R, xp] = radon(I, theta);
imagesc(theta, xp, R);
title('R_{\theta} (X\prime)');
xlabel('\theta (degrees)');
ylabel('X\prime');
set(gca, 'XTick', 0:20:180);
colormap(hot);
colorbar
```



Radon Transform Using 180 Projections

Using the Radon Transform to Detect Lines

The Radon transform is closely related to a common computer vision operation known as the Hough transform. You can use the `radon` function to implement a form of the Hough transform used to detect straight lines. The steps are

- 1 Compute a binary edge image using the `edge` function.

```
I = fitsread('solarspectra.fts');
I = mat2gray(I);
BW = edge(I);
imshow(I), figure, imshow(BW)
```





Image Courtesy of Ann Walker


Original Image



Edge Image

- 2 Compute the Radon transform of the edge image.


```
theta = 0:179;
[R, xp] = radon(BW, theta);
figure, imagesc(theta, xp, R); colormap(hot);
xlabel('\theta (degrees)'); ylabel('x\prime');
title('R_{\theta} (x\prime)');
colorbar
```



Radon Transform of an Edge Image

- 3** Find the locations of strong peaks in the Radon transform matrix. The locations of these peaks correspond to the locations of straight lines in the original image.

In the following figure, the strongest peaks in R correspond to $\theta = 1^\circ$ and $x' = -80$. The line perpendicular to that angle and located at $x' = -80$ is shown below, superimposed in red on the original image. The Radon transform geometry is shown in black. Notice that the other strong lines parallel to the red line also appear as peaks at $\theta = 1^\circ$ in the transform. Also, the lines perpendicular to this line appear as peaks at $\theta = 91^\circ$.



Radon Transform Geometry and the Strongest Peak (Red)

Inverse Radon Transform

The `iradon` function performs the inverse Radon transform, which is commonly used in tomography applications. This transform inverts the Radon transform (which was introduced in the previous section), and can therefore be used to reconstruct images from projection data.

As described in “Radon Transform” on page 8-21, given an image I and a set of angles θ , the radon function can be used to calculate the Radon transform.

```
R = radon(I,theta);
```

The function `iradon` can then be called to reconstruct the image I .


```
IR = iradon(R,theta);
```

In the example above, projections are calculated from the original image I . In most application areas, there is no *original image* from which projections are formed. For example, in X-ray absorption tomography, projections are formed by measuring the attenuation of radiation that passes through a physical specimen at different angles. The original image can be thought of as a cross section through the specimen, in which intensity values represent the density of the specimen. Projections are collected using special purpose hardware, and then an internal image of the specimen is reconstructed by `iradon`. This allows for noninvasive imaging of the inside of a living body or another opaque object.

The Radon transform reconstructs an image from parallel-beam projections. In *parallel-beam geometry*, each projection is formed by combining a set of line integrals through an image at a specific angle.

The following figure illustrates how parallel-beam geometry is applied in X-ray absorption tomography. Note that there is an equal number of n emitters and n sensors. Each sensor measures the radiation emitted from its corresponding emitter, and the attenuation in the radiation gives a measure of the integrated density, or mass, of the object. This corresponds to the line integral that is calculated in the Radon transform.

The parallel-beam geometry used in the figure is the same as the geometry that was described in “Radon Transform” on page 8-21. $f(x,y)$ denotes the brightness of the image and $R_{\theta}(x')$ is the projection at angle θ .



Parallel-Beam Projections Through an Object

Another geometry that is commonly used is *fan-beam* geometry, in which there is one source and n sensors. For more information, see “Fan-Beam Projection Data” on page 8-36. To convert parallel-beam projection data into fan-beam projection data, use the `para2fan` function.

Improving the Results

`iradon` uses the *filtered backprojection* algorithm to compute the inverse Radon transform. This algorithm forms an approximation of the image I based on the projections in the columns of R . A more accurate result can be obtained by using

more projections in the reconstruction. As the number of projections (the length of `theta`) increases, the reconstructed image `IR` more accurately approximates the original image `I`. The vector `theta` must contain monotonically increasing angular values with a constant incremental angle $\Delta\theta$. When the scalar $\Delta\theta$ is known, it can be passed to `iradon` instead of the array of `theta` values. Here is an example.

```
IR = iradon(R,Dtheta);
```

The filtered backprojection algorithm filters the projections in `R` and then reconstructs the image using the filtered projections. In some cases, noise can be present in the projections. To remove high frequency noise, apply a window to the filter to attenuate the noise. Many such windowed filters are available in `iradon`. The example call to `iradon` below applies a Hamming window to the filter. See the `iradon` reference page for more information.

```
IR = iradon(R,theta,'Hamming');
```

`iradon` also enables you to specify a normalized frequency, `D`, above which the filter has zero response. `D` must be a scalar in the range $[0,1]$. With this option, the frequency axis is rescaled so that the whole filter is compressed to fit into the frequency range $[0,D]$. This can be useful in cases where the projections contain little high-frequency information but there is high-frequency noise. In this case, the noise can be completely suppressed without compromising the reconstruction. The following call to `iradon` sets a normalized frequency value of 0.85.

```
IR = iradon(R,theta,0.85);
```

Example: Reconstructing an Image from Parallel Projection Data

The commands below illustrate how to reconstruct an image from parallel projection data. The test image is the Shepp-Logan head phantom, which can be generated by the Image Processing Toolbox function `phantom`. The phantom image illustrates many of the qualities that are found in real-world tomographic imaging of human heads. The bright elliptical shell along the exterior is analogous to a skull, and the many ellipses inside are analogous to brain features.

- 1 Create a Shepp-Logan head phantom image.

```
P = phantom(256);  
imshow(P)
```




- 2 Compute the Radon transform of the phantom brain for three different sets of theta values. R1 has 18 projections, R2 has 36 projections, and R3 has 90 projections.

```
theta1 = 0:10:170; [R1,xp] = radon(P,theta1);  
theta2 = 0:5:175; [R2,xp] = radon(P,theta2);  
theta3 = 0:2:178; [R3,xp] = radon(P,theta3);
```

- 3 Display a plot of one of the Radon transforms of the Shepp-Logan head phantom. The following figure shows R3, the transform with 90 projections.

```
figure, imagesc(theta3,xp,R3); colormap(hot); colorbar  
xlabel('\theta'); ylabel('x\prime');
```



Radon Transform of Head Phantom Using 90 Projections


Note how some of the features of the input image appear in this image of the transform. The first column in the Radon transform corresponds to a projection at 0° that is integrating in the vertical direction. The centermost column corresponds to a projection at 90° , which is integrating in the horizontal direction. The projection at 90° has a wider profile than the projection at 0° due to the larger vertical semiaxis of the outermost ellipse of the phantom.

- 4 Reconstruct the head phantom image from the projection data created in step 2 and display the results.

```
I1 = iradon(R1,10);
I2 = iradon(R2,5);
I3 = iradon(R3,2);
imshow(I1)
figure, imshow(I2)
```

```
figure, imshow(I3)
```

The following figure shows the results of all three reconstructions. Notice how image I1, which was reconstructed from only 18 projections, is the least accurate reconstruction. Image I2, which was reconstructed from 36 projections, is better, but it is still not clear enough to discern clearly the small ellipses in the lower portion of the image. I3, reconstructed using 90 projections, most closely resembles the original image. Notice that when the number of projections is relatively small (as in I1 and I2), the reconstruction can include some artifacts from the back projection.



I1



I2




I3

Inverse Radon Transforms of the Shepp-Logan Head Phantom

Fan-Beam Projection Data

The fanbeam function in the Image Processing Toolbox computes *projections* of an image matrix along specified directions. A projection of a two-dimensional function $f(x,y)$ is a set of line integrals. The fanbeam function computes the line integrals along paths that radiate from a single source, forming a fan shape. To represent an image, the fanbeam function takes multiple projections of the image from different angles by rotating the source around the center of the image. The following figure shows a single fan-beam projection at a specified rotation angle.



Fan-Beam Projection at Rotation Angle Theta

This section

- Describes how to use the fanbeam function to generate fan-beam projection data
- Describes how to reconstruct an image from fan-beam projection data
- Shows an example that creates a fan-beam projection of an image and then reconstructs the image from the fan-beam projection data

Note For information about creating projection data from line integrals along parallel paths, see “Radon Transform” on page 8-21. To convert fan-beam projection data to parallel-beam projection data, use the `fan2para` function.

Computing Fan-Beam Projection Data


To compute fan-beam projection data, use the `fanbeam` function. You specify as arguments an image and the distance between the vertex of the fan-beam projections and the center of rotation (the center pixel in the image). The `fanbeam` function determines the number of beams, based on the size of the image and the settings of `fanbeam` parameters.

By default, `fanbeam` positions the sensors along an arc at distance D from the center of rotation, and spaces the sensors at 1 degree intervals. Using the `FanSensorSpacing` parameter, you can specify a different angle between each beam. Using the `FanSensorGeometry` parameter, you can optionally specify that `fanbeam` position sensors along a straight line, rather than an arc. With this geometry, you specify the spacing between sensors in pixels. In this case, only the sensor aligned with the center pixel is distance D from the center of rotation.

`fanbeam` takes projections at different angles by rotating the source around the center pixel at 1 degree intervals. Using the `FanRotationIncrement` parameter you can specify a different rotation angle increment.


The following figures illustrate both these geometries.

The first figure illustrates geometry used by the `fanbeam` function when `FanSensorGeometry` is set to 'arc' (the default). Note how you specify the distance between sensors by specifying the angular spacing of the beams.



Fan-Beam Projection with Arc Geometry

The following figure illustrates geometry used by the `fanbeam` function when `FanSensorGeometry` is set to `'line'`. In this figure, note how you specify the position of the sensors by specifying the distance between them in pixels.



Fan-Beam Projection with Line Geometry

Reconstructing an Image from Fan-Beam Projection Data

To reconstruct an image from fan-beam projection data, use the `ifanbeam` function. With this function, you specify as arguments the projection data and the distance between the vertex of the fan-beam projections and the center of rotation when the projection data was created. For example, this code recreates the image I from the projection data P and distance D .

```
I = ifanbeam(P,D);
```

By default, the `ifanbeam` function assumes that the fan-beam projection data was created using the arc fan sensor geometry, with beams spaced at 1 degree angles and projections taken at 1 degree increments over a full 360 degree range. As with the `fanbeam` function, you can use `ifanbeam` parameters to specify other values for these characteristics of the projection data. Use the same values for these parameters that were used when the projection data was created. For more information about these parameters, see “Computing Fan-Beam Projection Data” on page 8-37.

The `ifanbeam` function converts the fan-beam projection data to parallel-beam projection data with the `fan2para` function, and then calls the `iradon` function to perform the image reconstruction. For this reason, the `ifanbeam` function supports certain `iradon` parameters, which it passes to the `iradon` function. See “Inverse Radon Transform” on page 8-29 for more information about the `iradon` function.

Working with Fan-Beam Projection Data

The commands below illustrate how to use `fanbeam` and `ifanbeam` to form projections from a sample image and then reconstruct the image from the projections. The test image is the Shepp-Logan head phantom, which can be generated by the Image Processing Toolbox function `phantom`. The phantom image illustrates many of the qualities that are found in real-world tomographic imaging of human heads.

- 1 Generate the test image and display it.

```
P = phantom(256);  
imshow(P)
```



- 2 Compute fan-beam projection data of the test image, using the `FanSensorSpacing` parameter to vary the sensor spacing. The example uses the fanbeam arc geometry, so you specify the spacing between sensors by specifying the angular spacing of the beams. The first call spaces the beams at 2 degrees; the second at 1 degree; and the third at 0.25 degrees. In each call, the distance between the center of rotation and vertex of the projections is constant at 250 pixels. In addition, `fanbeam` rotates the projection around the center pixel at 1 degree increments.

```
D = 250;
```


```
dsensor1 = 2;  
F1 = fanbeam(P,D, 'FanSensorSpacing', dsensor1);
```

```
dsensor2 = 1;  
F2 = fanbeam(P,D, 'FanSensorSpacing', dsensor2);
```

```
dsensor3 = 0.25  
[F3, sensor_pos3, fan_rot_angles3] = fanbeam(P,D,...  
'FanSensorSpacing', dsensor3);
```

- 3 Plot the projection data `F3`. Because `fanbeam` calculates projection data at rotation angles from 0 to 360 degrees, the same patterns occur at an offset of 180 degrees. The same features are being sampled from both sides. Compare this plot to the plot of the parallel-beam projection data of the head phantom on page 8-34.

```
figure, imagesc(fan_rot_angles3, sensor_pos3, F3)  
colormap(hot); colorbar  
xlabel('Fan Rotation Angle (degrees)')  
ylabel('Fan Sensor Position (degrees)')
```



- 4 Reconstruct the image from the fan-beam projection data using `ifanbeam`. In each reconstruction, match the fan sensor spacing with the spacing used when the projection data was created in step 2. The example uses the `OutputSize` parameter to constrain the output size of each reconstruction to be the same as the size of the original image `|P|`.

```
output_size = max(size(P));
```


```
Ifan1 = ifanbeam(F1,D,  
                'FanSensorSpacing',dsensor1,'OutputSize',output_size);  
figure, imshow>Ifan1)
```

```
Ifan2 = ifanbeam(F2,D,  
                'FanSensorSpacing',dsensor2,'OutputSize',output_size);  
figure, imshow>Ifan2)
```


```
Ifan3 = ifanbeam(F3,D,  
                'FanSensorSpacing',dsensor3,'OutputSize',output_size);
```

```
figure, imshow(Ifan3)
```

The following figure shows the result of each transform. Note how the quality of the reconstruction gets better as the number of beams in the projection increases. The first image, Ifan1, was created using 2 degree spacing of the beams; the second image, ifan2, was created using 1 degree spacing of the beams; the third image, ifan3, was created using 0.25 spacing of the beams.



Ifan1



Ifan2



Ifan3

Reconstruction of the Head Phantom Image from Fan-Beam Projections

Morphological Operations

Morphology is a technique of image processing based on shapes. The value of each pixel in the output image is based on a comparison of the corresponding pixel in the input image with its neighbors. By choosing the size and shape of the neighborhood, you can construct a morphological operation that is sensitive to specific shapes in the input image.

This chapter describes the Image Processing Toolbox morphological functions. You can use these functions to perform common image processing tasks, such as contrast enhancement, noise removal, thinning, skeletonization, filling, and segmentation.

Terminology (p. 9-2)	Provides definitions of image processing terms used in this section
Dilation and Erosion (p. 9-4)	Defines the two fundamental morphological operations, dilation and erosion, and some of the morphological image processing operations that are based on combinations of these operations
Morphological Reconstruction (p. 9-19)	Describes morphological reconstruction and the toolbox functions that use this type of processing
Distance Transform (p. 9-37)	Describes how to use the <code>bwdist</code> function to compute the distance transform of an image
Objects, Regions, and Feature Measurement (p. 9-40)	Describes functions that return information about a binary image
Lookup Table Operations (p. 9-44)	Describes functions that perform lookup table operations

Terminology

An understanding of the following terms will help you to use this chapter.

Term	Definition
background	In a binary image, pixels that are off, i.e., set to the value 0, are considered the background. When you view a binary image, the background pixels appear black.
connectivity	Criteria that describe how pixels in an image form a connected group. For example, a connected component is “8-connected” if diagonally adjacent pixels are considered to be touching; otherwise, it is “4-connected.” The toolbox supports 2-D as well as multidimensional connectivities. See “Pixel Connectivity” on page 9-23 for more information.
foreground	In a binary image, pixels that are on, i.e., set to the value 1, are considered the foreground. When you view a binary image, the foreground pixels appear white.
global maxima	Highest regional maxima in the image. See the entry for regional maxima in this table for more information.
global minima	Lowest regional minima in the image. See the entry for regional minima in this table for more information.
morphology	A broad set of image processing operations that process images based on shapes. Morphological operations apply a structuring element to an input image, creating an output image of the same size. The most basic morphological operations are dilation and erosion.
neighborhood	Set of pixels that are defined by their locations relative to the pixel of interest. A neighborhood can be defined by a structuring element or by specifying a connectivity.
object	Set of pixels in a binary image that form a connected group. In the context of this chapter, “object” and “connected component” are equivalent.

Term	Definition
packed binary image	Method of compressing binary images that can speed up the processing of the image.
regional maxima	Connected set of pixels of constant intensity from which it is impossible to reach a point with higher intensity without first descending; that is, a connected component of pixels with the same intensity value, t , surrounded by pixels that all have a value less than t .
regional minima	Connected set of pixels of constant intensity from which it is impossible to reach a point with lower intensity without first ascending; that is, a connected component of pixels with the same intensity value, t , surrounded by pixels that all have a value greater than t .
structuring element	Matrix used to define a neighborhood shape and size for morphological operations, including dilation and erosion. It consists of only 0's and 1's and can have an arbitrary shape and size. The pixels with values of 1 define the neighborhood.

Dilation and Erosion

Dilation and erosion are two fundamental morphological operations. Dilation adds pixels to the boundaries of objects in an image, while erosion removes pixels on object boundaries. The number of pixels added or removed from the objects in an image depends on the size and shape of the *structuring element* used to process the image.

The following sections

- Provide important background information about how the dilation and erosion functions operate
- Describe structuring elements and how to create them
- Describe how to perform a morphological dilation
- Describe how to perform a morphological erosion
- Describe some of the common operations that are based on dilation and erosion
- Describe toolbox functions that are based on dilation and erosion

To view an extended example that uses morphological processing to solve an image processing problem, see the Image Processing Toolbox watershed segmentation demo.


Understanding Dilation and Erosion

In the morphological dilation and erosion operations, the state of any given pixel in the output image is determined by applying a rule to the corresponding pixel and its neighbors in the input image. The rule used to process the pixels defines the operation as a dilation or an erosion. This table lists the rules for both dilation and erosion.

Rules for Dilation and Erosion


Operation	Rule
Dilation	The value of the output pixel is the <i>maximum</i> value of all the pixels in the input pixel's neighborhood. In a binary image, if any of the pixels is set to the value 1, the output pixel is set to 1.
Erosion	The value of the output pixel is the <i>minimum</i> value of all the pixels in the input pixel's neighborhood. In a binary image, if any of the pixels is set to 0, the output pixel is set to 0.

The following figure illustrates the dilation of a binary image. Note how the structuring element defines the neighborhood of the pixel of interest, which is circled. (See “Structuring Elements” on page 9-7 for more information.) The dilation function applies the appropriate rule to the pixels in the neighborhood and assigns a value to the corresponding pixel in the output image. In the figure, the morphological dilation function sets the value of the output pixel to 1 because one of the elements in the neighborhood defined by the structuring element is on.



Morphological Dilation of a Binary Image

The following figure illustrates this processing for a grayscale image. The figure shows the processing of a particular pixel in the input image. Note how the function applies the rule to the input pixel's neighborhood and uses the highest value of all the pixels in the neighborhood as the value of the corresponding pixel in the output image.



Morphological Dilation of a Grayscale Image

Processing Pixels at Image Borders (Padding Behavior)

Morphological functions position the origin of the structuring element, its center element, over the pixel of interest in the input image. For pixels at the edge of an image, parts of the neighborhood defined by the structuring element can extend past the border of the image.

To process border pixels, the morphological functions assign a value to these undefined pixels, as if the functions had padded the image with additional rows and columns. The value of these padding pixels varies for dilation and erosion operations. The following table describes the padding rules for dilation and erosion for both binary and grayscale images.

Rules for Padding Images

Operation	Rule
Dilation	<p>Pixels beyond the image border are assigned the minimum value afforded by the data type.</p> <p>For binary images, these pixels are assumed to be set to 0. For grayscale images, the minimum value for <code>uint8</code> images is 0.</p>
Erosion	<p>Pixels beyond the image border are assigned the <i>maximum</i> value afforded by the data type.</p> <p>For binary images, these pixels are assumed to be set to 1. For grayscale images, the maximum value for <code>uint8</code> images is 255.</p>

Note By using the minimum value for dilation operations and the maximum value for erosion operations, the toolbox avoids *border effects*, where regions near the borders of the output image do not appear to be homogeneous with the rest of the image. For example, if erosion padded with a minimum value, eroding an image would result in a black border around the edge of the output image.

Structuring Elements

An essential part of the dilation and erosion operations is the structuring element used to probe the input image. Two-dimensional, or *flat*, structuring elements consist of a matrix of 0's and 1's, typically much smaller than the image being processed. The center pixel of the structuring element, called the *origin*, identifies the pixel of interest — the pixel being processed. The pixels in the structuring element containing 1's define the *neighborhood* of the structuring element. These pixels are also considered in dilation or erosion processing. Three-dimensional, or *nonflat*, structuring elements use 0's and 1's to define the extent of the structuring element in the *x*- and *y*-planes and add height values to define the third dimension.


The Origin of a Structuring Element

The morphological functions use this code to get the coordinates of the origin of structuring elements of any size and dimension.

```
origin = floor((size(nhood)+1)/2)
```

(In this code `nhood` is the neighborhood defining the structuring element. Because structuring elements are MATLAB objects, you cannot use the size of the STREL object itself in this calculation. You must use the STREL `getnhood` method to retrieve the neighborhood of the structuring element from the STREL object. For information about other STREL object methods, see the `strel` function reference page.)

For example, the following illustrates a diamond-shaped structuring element.



Origin of a Diamond-Shaped Structuring Element

Creating a Structuring Element

The toolbox dilation and erosion functions accept structuring element objects, called STRELS. You use the `strel` function to create STRELS of any arbitrary size and shape. The `strel` function also includes built-in support for many common shapes, such as lines, diamonds, disks, periodic lines, and balls.

Note You typically choose a structuring element the same size and shape as the objects you want to process in the input image. For example, to find lines in an image, create a linear structuring element.

For example, this code creates a flat, diamond-shaped structuring element.

```
se = strel('diamond',3)
se =
```

```
Flat STREL object containing 25 neighbors.
Decomposition: 3 STREL objects containing a total of 13 neighbors
```

Neighborhood:

```
  0   0   0   1   0   0   0
  0   0   1   1   1   0   0
  0   1   1   1   1   1   0
  1   1   1   1   1   1   1
  0   1   1   1   1   1   0
  0   0   1   1   1   0   0
  0   0   0   1   0   0   0
```

Structuring Element Decomposition

To enhance performance, the `strel` function might break structuring elements into smaller pieces, a technique known as *structuring element decomposition*.

For example, dilation by an 11-by-11 square structuring element can be accomplished by dilating first with a 1-by-11 structuring element, and then with an 11-by-1 structuring element. This results in a theoretical speed improvement of a factor of 5.5, although in practice the actual speed improvement is somewhat less.

Structuring element decompositions used for the 'disk' and 'ball' shapes are approximations; all other decompositions are exact. Decomposition is not used with an arbitrary structuring element unless it is a flat structuring element whose neighborhood is all 1's.

To view the sequence of structuring elements used in a decomposition, use the STREL `getsequence` method. The `getsequence` function returns an array of the structuring elements that form the decomposition. For example, here are the structuring elements created in the decomposition of a diamond-shaped structuring element.

```
sel = strel('diamond',4)
sel =
```

```
Flat STREL object containing 41 neighbors.
Decomposition: 3 STREL objects containing a total of 13 neighbors
```

Neighborhood:

```

0 0 0 0 1 0 0 0 0
0 0 0 1 1 1 0 0 0
0 0 1 1 1 1 1 0 0
0 1 1 1 1 1 1 1 0
1 1 1 1 1 1 1 1 1
0 1 1 1 1 1 1 1 0
0 0 1 1 1 1 1 0 0
0 0 0 1 1 1 0 0 0
0 0 0 0 1 0 0 0 0

```

```

seq = getsequence(sel)
seq =
3x1 array of STREL objects

```

```

seq(1)
ans =
Flat STREL object containing 5 neighbors.

```

Neighborhood:

```

0 1 0
1 1 1
0 1 0

```

```

seq(2)
ans =
Flat STREL object containing 4 neighbors.

```

Neighborhood:

```

0 1 0
1 0 1
0 1 0

```

```

seq(3)
ans =
Flat STREL object containing 4 neighbors.

```

Neighborhood:

```

0 0 1 0 0

```

```

0   0   0   0   0
1   0   0   0   1
0   0   0   0   0
0   0   1   0   0

```

Dilating an Image

To dilate an image, use the `imdilate` function. The `imdilate` function accepts two primary arguments:

- The input image to be processed (grayscale, binary, or packed binary image)
- A structuring element object, returned by the `strel` function, or a binary matrix defining the neighborhood of a structuring element

`imdilate` also accepts two optional arguments: `PADOPT` and `PACKOPT`. The `PADOPT` argument affects the size of the output image. The `PACKOPT` argument identifies the input image as packed binary. (See the `bwpack` reference page for information about binary image packing.)

This example dilates a simple binary image containing one rectangular object.

```

BW = zeros(9,10);
BW(4:6,4:7) = 1
BW =
    0     0     0     0     0     0     0     0     0     0
    0     0     0     0     0     0     0     0     0     0
    0     0     0     0     0     0     0     0     0     0
    0     0     0     1     1     1     1     0     0     0
    0     0     0     1     1     1     1     0     0     0
    0     0     0     1     1     1     1     0     0     0
    0     0     0     0     0     0     0     0     0     0
    0     0     0     0     0     0     0     0     0     0
    0     0     0     0     0     0     0     0     0     0

```

To expand all sides of the foreground component, the example uses a 3-by-3 square structuring element object. (For more information about using the `strel` function, see “Structuring Elements” on page 9-7.)

```

SE = strel('square',3)
SE =

```

Flat STREL object containing 3 neighbors.

Neighborhood:

```

1   1   1
1   1   1
1   1   1

```

To dilate the image, pass the image BW and the structuring element SE to the `imdilate` function. Note how dilation adds a rank of 1's to all sides of the foreground object.

```
BW2 = imdilate(BW,SE)
```

```
BW2 =
```

```

0   0   0   0   0   0   0   0   0   0
0   0   0   0   0   0   0   0   0   0
0   0   1   1   1   1   1   1   0   0
0   0   1   1   1   1   1   1   0   0
0   0   1   1   1   1   1   1   0   0
0   0   1   1   1   1   1   1   0   0
0   0   1   1   1   1   1   1   0   0
0   0   0   0   0   0   0   0   0   0
0   0   0   0   0   0   0   0   0   0

```

Eroding an Image

To erode an image, use the `imerode` function. The `imerode` function accepts two primary arguments:

- The input image to be processed (grayscale, binary, or packed binary image)
- A structuring element object, returned by the `strel` function, or a binary matrix defining the neighborhood of a structuring element

`imerode` also accepts three optional arguments: `PADOPT`, `PACKOPT`, and `M`.

The `PADOPT` argument affects the size of the output image. The `PACKOPT` argument identifies the input image as packed binary. If the image is packed binary, `M` identifies the number of rows in the original image. (See the `bwpack` reference page for more information about binary image packing.)

The following example erodes the binary image `circbw.tif`:

- 1 Read the image into the MATLAB workspace.

```
BW1 = imread('circbw.tif');
```

- 2 Create a structuring element. The following code creates a diagonal structuring element object. (For more information about using the `strel` function, see “Structuring Elements” on page 9-7.)

```
SE = strel('arbitrary',eye(5));  
SE=
```

Flat STREL object containing 5 neighbors.

Neighborhood:


```
 1   0   0   0   0  
 0   1   0   0   0  
 0   0   1   0   0  
 0   0   0   1   0  
 0   0   0   0   1
```

- 3 Call the `imerode` function, passing the image `BW` and the structuring element `SE` as arguments.

```
BW2 = imerode(BW1,SE);
```

Notice the diagonal streaks on the right side of the output image. These are due to the shape of the structuring element.

```
imshow(BW1)  
figure, imshow(BW2)
```



Combining Dilation and Erosion

Dilation and erosion are often used in combination to implement image processing operations. For example, the definition of a morphological *opening* of an image is an erosion followed by a dilation, using the same structuring element for both operations. The related operation, morphological *closing* of an image, is the reverse: it consists of dilation followed by an erosion with the same structuring element.

The following section uses `imdilate` and `imerode` to illustrate how to implement a morphological opening. Note, however, that the toolbox already includes the `imopen` function, which performs this processing. The toolbox includes functions that perform many common morphological operations. See “Dilation- and Erosion-Based Functions” on page 9-16 for a complete list.

Morphological Opening

You can use morphological opening to remove small objects from an image while preserving the shape and size of larger objects in the image. For example, you can use the `imopen` function to remove all the circuit lines from the original circuit image, `circbw.tif`, creating an output image that contains only the rectangular shapes of the microchips.

To morphologically open the image, perform these steps:

- 1 Read the image into the MATLAB workspace.

```
BW1 = imread('circbw.tif');
```

- 2 Create a structuring element.

```
SE = strel('rectangle',[40 30]);
```

The structuring element should be large enough to remove the lines when you erode the image, but not large enough to remove the rectangles. It should consist of all 1's, so it removes everything but large contiguous patches of foreground pixels.

- 3 Erode the image with the structuring element.


```
BW2 = imerode(BW1,SE);  
imshow(BW2)
```

This removes all the lines, but also shrinks the rectangles.



- 4 To restore the rectangles to their original sizes, dilate the eroded image using the same structuring element, SE.

```
BW3 = imdilate(BW2,SE);  
imshow(BW3)
```



Dilation- and Erosion-Based Functions

This section describes two common image processing operations that are based on dilation and erosion:

- Skeletonization
- Perimeter determination

This table lists other functions in the toolbox that perform common morphological operations that are based on dilation and erosion. For more information about these functions, see their reference pages.

Dilation- and Erosion-Based Functions

Function	Morphological Definition
<code>bwhitmiss</code>	Logical AND of an image, eroded with one structuring element, and the image's complement, eroded with a second structuring element.
<code>imbothat</code>	Subtracts the original image from a morphologically closed version of the image. Can be used to find intensity troughs in an image.
<code>imclose</code>	Dilates an image and then erodes the dilated image using the same structuring element for both operations.


Dilation- and Erosion-Based Functions

Function	Morphological Definition (Continued)
<code>imopen</code>	Erodes an image and then dilates the eroded image using the same structuring element for both operations.
<code>imtophat</code>	Subtracts a morphologically opened image from the original image. Can be used to enhance contrast in an image.


Skeletonization

To reduce all objects in an image to lines, without changing the essential structure of the image, use the `bwmorph` function. This process is known as *skeletonization*.

```
BW1 = imread('circbw.tif');
BW2 = bwmorph(BW1, 'skel', Inf);
imshow(BW1)
figure, imshow(BW2)
```



Original Image



Skeletonization of Image

Perimeter Determination


The `bwperim` function determines the perimeter pixels of the objects in a binary image. A pixel is considered a perimeter pixel if it satisfies both of these criteria:

- The pixel is on.


- One (or more) of the pixels in its neighborhood is off.

For example, this code finds the perimeter pixels in a binary image of a circuit board.

```
BW1 = imread('circbw.tif');  
BW2 = bwperim(BW1);  
imshow(BW1)  
figure, imshow(BW2)
```



Original Image



Perimeters Determined

Morphological Reconstruction

Morphological reconstruction is another major part of morphological image processing. Based on dilation, morphological reconstruction has these unique properties:

- Processing is based on two images, a marker and a mask, rather than one image and a structuring element.
- Processing repeats until stability; i.e., the image no longer changes.
- Processing is based on the concept of connectivity, rather than a structuring element.

This section

- Provides background information about morphological reconstruction and describes how to use the `imreconstruct` function
- Describes how pixel connectivity affects morphological reconstruction
- Describes how to use the `imfill` function, which is based on morphological reconstruction
- Describes a group of functions, all based on morphological reconstruction, that process image extrema, i.e., the areas of high and low intensity in images

Marker and Mask

Morphological reconstruction processes one image, called the *marker*, based on the characteristics of another image, called the *mask*. The high points, or peaks, in the marker image specify where processing begins. The processing continues until the image values stop changing.

To illustrate morphological reconstruction, consider this simple image. It contains two primary regions, the blocks of pixels containing the values 14 and 18. The background is primarily all set to 10, with some pixels set to 11.

```
A = [10 10 10 10 10 10 10 10 10 10;
     10 14 14 14 10 10 11 10 11 10;
     10 14 14 14 10 10 10 11 10 10;
     10 14 14 14 10 10 11 10 11 10;
     10 10 10 10 10 10 10 10 10 10;
     10 11 10 10 10 18 18 18 10 10;
     10 10 10 11 10 18 18 18 10 10;
     10 10 11 10 10 18 18 18 10 10;
     10 11 10 11 10 10 10 10 10 10;
     10 10 10 10 10 10 11 10 10 10];
```

To morphologically reconstruct this image, perform these steps:

- 1 Create a marker image. As with the structuring element in dilation and erosion, the characteristics of the marker image determine the processing performed in morphological reconstruction. The peaks in the marker image should identify the location of objects in the mask image that you want to emphasize.

One way to create a marker image is to subtract a constant from the mask image, using `imsubtract`.

```
marker = imread('A.mat');
marker = imsubtract(marker, 2);
```

8	8	8	8	8	8	8	8	8	8
8	12	12	12	8	8	9	8	9	8
8	12	12	12	8	8	8	9	8	8
8	12	12	12	8	8	9	8	9	8
8	8	8	8	8	8	8	8	8	8
8	9	8	8	8	16	16	16	8	8
8	8	8	9	8	16	16	16	8	8
8	8	9	8	8	16	16	16	8	8
8	9	8	9	8	8	8	8	8	8
8	8	8	8	8	8	9	8	8	8

- 2 Call the `imreconstruct` function to morphologically reconstruct the image. In the output image, note how all the intensity fluctuations except the intensity peak have been removed.

```


recon = imreconstruct(marker, mask)
recon =
    10    10    10    10    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10

```

Understanding Morphological Reconstruction

Morphological reconstruction can be thought of conceptually as repeated dilations of the marker image until the contour of the marker image fits under the mask image. In this way, the peaks in the marker image “spread out”, or dilate.

This figure illustrates this processing in 1-D. Each successive dilation is constrained to lie underneath the mask. When further dilation ceases to change the image, processing stops. The final dilation is the reconstructed image. (Note: the actual implementation of this operation in the toolbox is done much more efficiently. See the `imreconstruct` reference page for more details.) The figure shows the successive dilations of the marker.



Repeated Dilations of Marker Image, Constrained by Mask

Pixel Connectivity

Morphological processing starts at the peaks in the marker image and spreads throughout the rest of the image based on the connectivity of the pixels. Connectivity defines which pixels are connected to other pixels.

For example, this binary image contains one foreground object—all the pixels that are set to 1. If the foreground is 4-connected, the image has one background object, and all the pixels are set to 0. However, if the foreground is 8-connected, the foreground makes a closed loop and the image has two separate background objects: the pixels in the loop and the pixels outside the loop.






0	0	0	0	0	0	0	0
0	1	1	1	1	1	0	0
0	1	0	0	0	1	0	0
0	1	0	0	0	1	0	0
0	1	0	0	0	1	0	0
0	1	1	1	1	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Defining Connectivity in an Image

The following table lists all the standard two- and three-dimensional connectivities supported by the toolbox. See these sections for more information:

- “Choosing a Connectivity”
- “Specifying Custom Connectivities”

Supported Connectivities

Two-Dimensional Connectivities		
4-connected	Pixels are connected if their edges touch. This means that a pair of adjoining pixels are part of the same object only if they are both on and are connected along the horizontal or vertical direction.	
8-connected	Pixels are connected if their edges or corners touch. This means that if two adjoining pixels are on, they are part of the same object, regardless of whether they are connected along the horizontal, vertical, or diagonal direction.	
Three-Dimensional Connectivities		
6-connected	Pixels are connected if their faces touch.	 6 faces
18-connected	Pixels are connected if their faces or edges touch.	 6 faces + 12 edges
26-connected	Pixels are connected if their faces, edges, or corners touch.	 6 faces + 12 edges + 8 corners

Choosing a Connectivity

The type of neighborhood you choose affects the number of objects found in an image and the boundaries of those objects. For this reason, the results of many morphology operations often differ depending upon the type of connectivity you specify.

For example, if you specify a 4-connected neighborhood, this binary image contains two objects; if you specify an 8-connected neighborhood, the image has one object.

```

0   0   0   0   0   0
0   1   1   0   0   0
0   1   1   0   0   0
0   0   0   1   1   0
0   0   0   1   1   0

```

Specifying Custom Connectivities

You can also define custom neighborhoods by specifying a 3-by-3-by-...-by-3 array of 0's and 1's. The 1-valued elements define the connectivity of the neighborhood relative to the center element.

For example, this array defines a “North/South” connectivity that has the effect of breaking up an image into independent columns.

```

CONN = [ 0 1 0; 0 1 0; 0 1 0 ]
CONN =
    0     1     0
    0     1     0
    0     1     0

```

Note Connectivity arrays must be symmetric about their center element. Also, you can use a 2-D connectivity array with a 3-D image; the connectivity affects each “page” in the 3-D image.

Flood-Fill Operations

The `imfill` function performs a *flood-fill* operation on binary and grayscale images. For binary images, `imfill` changes connected background pixels (0's) to foreground pixels (1's), stopping when it reaches object boundaries. For grayscale images, `imfill` brings the intensity values of dark areas that are surrounded by lighter areas up to the same intensity level as surrounding pixels. (In effect, `imfill` removes regional minima that are not connected to the image border. See “Finding Areas of High or Low Intensity” for more information.) This operation can be useful in removing irrelevant artifacts from images.

This section includes information about

- Specifying the connectivity in flood-fill operations
- Specifying the starting point for binary image fill operations
- Filling holes in binary or grayscale images

Specifying Connectivity

For both binary and grayscale images, the boundary of the fill operation is determined by the connectivity you specify.

Note `imfill` differs from the other object-based operations in that it operates on *background* pixels. When you specify connectivity with `imfill`, you are specifying the connectivity of the background, not the foreground.

The implications of connectivity can be illustrated with this matrix.

```
BW = [ 0  0  0  0  0  0  0  0;
       0  1  1  1  1  1  0  0;
       0  1  0  0  0  1  0  0;
       0  1  0  0  0  1  0  0;
       0  1  0  0  0  1  0  0;
       0  1  1  1  1  0  0  0;
       0  0  0  0  0  0  0  0;
       0  0  0  0  0  0  0  0];
```

If the background is 4-connected, this binary image contains two separate background elements (the part inside the loop and the part outside). If the

background is 8-connected, the pixels connect diagonally, and there is only one background element.

Specifying the Starting Point

For binary images, you can specify the starting point of the fill operation by passing in the location subscript or by using `imfill` in interactive mode, selecting starting pixels with a mouse. See the reference page for `imfill` for more information about using `imfill` interactively.

For example, if you call `imfill`, specifying the pixel `BW(4,3)` as the starting point, `imfill` only fills the inside of the loop because, by default, the background is 4-connected.

```
imfill(BW,[4 3])
```

```
ans =
    0     0     0     0     0     0     0     0
    0     1     1     1     1     1     0     0
    0     1     1     1     1     1     0     0
    0     1     1     1     1     1     0     0
    0     1     1     1     1     1     0     0
    0     1     1     1     1     0     0     0
    0     0     0     0     0     0     0     0
    0     0     0     0     0     0     0     0
```

If you specify the same starting point, but use an 8-connected background connectivity, `imfill` fills the entire image.

```
imfill(BW,[4 3],8)
```

```
ans =
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
    1     1     1     1     1     1     1     1
```

Filling Holes

A common use of the flood-fill operation is to fill holes in images. For example, suppose you have an image, binary or grayscale, in which the foreground objects represent spheres. In the image, these objects should appear as disks, but instead are donut shaped because of reflections in the original photograph. Before doing any further processing of the image, you might want to first fill in the “donut holes” using `imfill`.


Because the use of flood-fill to fill holes is so common, `imfill` includes special syntax to support it for both binary and grayscale images. In this syntax, you just specify the argument `'holes'`; you do not have to specify starting locations in each hole.

To illustrate, this example fills holes in a grayscale image of a spinal column.

```
[X,map] = imread('spine.tif');  
I = ind2gray(X,map);  
Ifill = imfill(I,'holes');  
imshow(I);figure, imshow(Ifill)
```



Original



After Filling Holes

Finding Peaks and Valleys

Grayscale images can be thought of in three dimensions: the x - and y -axes represent pixel positions and the z -axis represents the intensity of each pixel. In this interpretation, the intensity values represent elevations, as in a topographical map. The areas of high intensity and low intensity in an image, peaks and valleys in topographical terms, can be important morphological features because they often mark relevant image objects.

For example, in an image of several spherical objects, points of high intensity could represent the tops of the objects. Using morphological processing, these maxima can be used to identify objects in an image.


This section covers these topics:

- “Understanding the Maxima and Minima Functions”
- “Finding Areas of High or Low Intensity” on page 9-30
- “Suppressing Minima and Maxima” on page 9-31
- “Imposing a Minimum” on page 9-33

Understanding the Maxima and Minima Functions

An image can have multiple regional maxima or minima but only a single global maximum or minimum. Determining image peaks or valleys can be used to create marker images that are used in morphological reconstruction.

This figure illustrates the concept in 1-D.



Finding Areas of High or Low Intensity

The toolbox includes functions that you can use to find areas of high or low intensity in an image:

- The `imregionalmax` and `imregionalmin` functions identify *all* regional minima or maxima.
- The `imextendedmax` and `imextendedmin` functions identify all regional minima or maxima that are greater than or less than a specified threshold.

The functions accept a grayscale image as input and return a binary image as output. In the output binary image, the regional minima or maxima are set to 1; all other pixels are set to 0.

For example, this simple image contains two primary regional maxima, the blocks of pixels containing the value 13 and 18, and several smaller maxima, set to 11.

```
A = [10  10  10  10  10  10  10  10  10  10;
     10  13  13  13  10  10  11  10  11  10;
     10  13  13  13  10  10  10  11  10  10;
     10  13  13  13  10  10  11  10  11  10;
     10  10  10  10  10  10  10  10  10  10;
     10  11  10  10  10  18  18  18  10  10;
     10  10  10  11  10  18  18  18  10  10;
     10  10  11  10  10  18  18  18  10  10;
     10  11  10  11  10  10  10  10  10  10;
     10  10  10  10  10  10  11  10  10  10];
```

The binary image returned by `imregionalmax` pinpoints all these regional maxima.

```
B = imregionalmax(A)
```

```
B =
```

```

0  0  0  0  0  0  0  0  0  0
0  1  1  1  0  0  1  0  1  0
0  1  1  1  0  0  0  1  0  0
0  1  1  1  0  0  1  0  1  0
0  0  0  0  0  0  0  0  0  0
0  1  0  0  0  1  1  1  0  0
0  0  0  1  0  1  1  1  0  0
0  0  1  0  0  1  1  1  0  0
0  1  0  1  0  0  0  0  0  0
0  0  0  0  0  0  1  0  0  0

```

You might want only to identify areas of the image where the change in intensity is extreme; that is, the difference between the pixel and neighboring pixels is greater than (or less than) a certain threshold. For example, to find only those regional maxima in the sample image, A, that are at least two units higher than their neighbors, use `imextendedmax`.

```
B = imextendedmax(A,2)
```

```
B =
```

```

0  0  0  0  0  0  0  0  0  0
0  1  1  1  0  0  0  0  0  0
0  1  1  1  0  0  0  0  0  0
0  1  1  1  0  0  0  0  0  0
0  0  0  0  0  0  0  0  0  0
0  0  0  0  0  1  1  1  0  0
0  0  0  0  0  1  1  1  0  0
0  0  0  0  0  1  1  1  0  0
0  0  0  0  0  0  0  0  0  0
0  0  0  0  0  0  0  0  0  0

```

Suppressing Minima and Maxima

In an image, every small fluctuation in intensity represents a regional minimum or maximum. You might only be interested in significant minima or maxima and not in these smaller minima and maxima caused by background texture.

To remove the less significant minima and maxima but retain the significant minima and maxima, use the `imhmax` or `imhmin` function. With these functions,

you can specify a contrast criteria or threshold level, h , that suppresses all maxima whose height is less than h or whose minima are greater than h .

Note The `imregionalmin`, `imregionalmax`, `imextendedmin`, and `imextendedmax` functions return a binary image that marks the locations of the regional minima and maxima in an image. The `imhmax` and `imhmin` functions produce an altered image.

For example, this simple image contains two primary regional maxima, the blocks of pixels containing the value 14 and 18, and several smaller maxima, set to 11.

```
A = [10  10  10  10  10  10  10  10  10  10;
     10  14  14  14  10  10  11  10  11  10;
     10  14  14  14  10  10  10  11  10  10;
     10  14  14  14  10  10  11  10  11  10;
     10  10  10  10  10  10  10  10  10  10;
     10  11  10  10  10  18  18  18  10  10;
     10  10  10  11  10  18  18  18  10  10;
     10  10  11  10  10  18  18  18  10  10;
     10  11  10  11  10  10  10  10  10  10;
     10  10  10  10  10  10  11  10  10  10];
```


To eliminate all regional maxima except the two significant maxima, use `imhmax`, specifying a threshold value of 2. Note that `imhmax` only affects the maxima; none of the other pixel values are changed. The two significant maxima remain, although their heights are reduced.

```

B = imhmax(A,2)
B =
    10    10    10    10    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    12    12    12    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    16    16    16    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10

```

This figure takes the second row from the sample image to illustrate in 1-D how `imhmax` changes the profile of the image.



Imposing a Minimum

You can emphasize specific minima (dark objects) in an image using the `imimposemin` function. The `imimposemin` function uses morphological reconstruction to eliminate all minima from the image except the minima you specify.

To illustrate the process of imposing a minimum, this code creates a simple image containing two primary regional minima and several other regional minima.

```

mask = uint8(10*ones(10,10));
mask(6:8,6:8) = 2;
mask(2:4,2:4) = 7;

```

```
mask(3,3) = 5;
mask(2,9) = 9
mask(3,8) = 9
mask(9,2) = 9
mask(8,3) = 9
```

```
mask = 10  10  10  10  10  10  10  10  10  10
       10  7  7  7  10  10  10  10  9  10
       10  7  6  7  10  10  10  9  10  10
       10  7  7  7  10  10  10  10  10  10
       10  10  10  10  10  10  10  10  10  10
       10  10  10  10  10  2  2  2  10  10
       10  10  10  10  10  2  2  2  10  10
       10  10  9  10  10  2  2  2  10  10
       10  9  10  10  10  10  10  10  10  10
       10  10  10  10  10  10  10  10  10  10
```

Creating a Marker Image

To obtain an image that emphasizes the two deepest minima and removes all others, create a marker image that pinpoints the two minima of interest. You can create the marker image by explicitly setting certain pixels to specific values or by using other morphological functions to extract the features you want to emphasize in the mask image.

This example uses `imextendedmin` to get a binary image that shows the locations of the two deepest minima.


```
marker = imextendedmin(mask,1)
marker = 0  0  0  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0  0  0
        0  0  1  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0  0  0
        0  0  0  0  0  1  1  1  0  0
        0  0  0  0  0  1  1  1  0  0
        0  0  0  0  0  1  1  1  0  0
        0  0  0  0  0  0  0  0  0  0
        0  0  0  0  0  0  0  0  0  0
```

Applying the Marker Image to the Mask

Now use `imimposemin` to create new minima in the mask image at the points specified by the marker image. Note how `imimposemin` sets the values of pixels specified by the marker image to the lowest value supported by the datatype (0 for `uint8` values). `imimposemin` also changes the values of all the other pixels in the image to eliminate the other minima.

```
I = imimposemin(mask,marker)
I =
    11    11    11    11    11    11    11    11    11    11
    11     8     8     8    11    11    11    11    11    11
    11     8     0     8    11    11    11    11    11    11
    11     8     8     8    11    11    11    11    11    11
    11    11    11    11    11    11    11    11    11    11
    11    11    11    11    11     0     0     0    11    11
    11    11    11    11    11     0     0     0    11    11
    11    11    11    11    11     0     0     0    11    11
    11    11    11    11    11    11    11    11    11    11
    11    11    11    11    11    11    11    11    11    11
```

This figure illustrates in 1-D how `imimposemin` changes the profile of row 2 of the image.







Imposing a Minimum

Distance Transform





The distance transform provides a metric or measure of the separation of points in the image. The Image Processing Toolbox provides a function, `bwdist`, that calculates the distance between each pixel that is set to off (0) and the nearest nonzero pixel for binary images.

The `bwdist` function supports several distance metrics, listed in the following table.

Distance Metrics

Distance Metric	Description	Illustration
Euclidean	The Euclidean distance is the straight-line distance between two pixels.	 Image  Distance transform
City Block	The city block distance metric measures the path between the pixels based on a 4-connected neighborhood. Pixels whose edges touch are 1 unit apart; pixels diagonally touching are 2 units apart.	 Image  Distance transform

Distance Metrics (Continued)

Distance Metric	Description	Illustration
Chessboard	The chessboard distance metric measures the path between the pixels based on an 8-connected neighborhood. Pixels whose edges or corners touch are 1 unit apart.	  <p>Image Distance transform</p>
Quasi-Euclidean	The quasi-Euclidean metric measures the total Euclidean distance along a set of horizontal, vertical, and diagonal line segments.	  <p>Image Distance transform</p>

This example creates a binary image containing two intersecting circular objects.

```


center1 = -10;
center2 = -center1;
dist = sqrt(2*(2*center1)^2);
radius = dist/2 * 1.4;
lims = [floor(center1-1.2*radius) ceil(center2+1.2*radius)];
[x,y] = meshgrid(lims(1):lims(2));
bw1 = sqrt((x-center1).^2 + (y-center1).^2) <= radius;
bw2 = sqrt((x-center2).^2 + (y-center2).^2) <= radius;
bw = bw1 | bw2;
figure, imshow(bw), title('bw')

```



To compute the distance transform of the complement of the binary image, use the `bwdist` function. In the image of the distance transform, note how the centers of the two circular areas are white.

```
D = bwdist(~bw);  
figure, imshow(D,[]), title('Distance transform of ~bw')
```



Objects, Regions, and Feature Measurement

The toolbox includes several functions that return information about the features in a binary image, including

- Connected-component labeling, and using the label matrix to get statistics about an image
- Selecting objects in a binary image
- Finding the area of the foreground of a binary image
- Finding the Euler number of a binary image

Connected-Component Labeling

The `bwlabel` and the `bwlabeln` functions perform *connected-component labeling*, which is a method for identifying each object in a binary image. The `bwlabel` function supports 2-D inputs only; the `bwlabeln` function supports inputs of any dimension.

These functions return a matrix, called a *label matrix*. A label matrix is an image, the same size as the input image, in which the objects in the input image are distinguished by different integer values in the output matrix. For example, `bwlabel` can identify the objects in this binary image.

```
BW = [0  0  0  0  0  0  0  0;
      0  1  1  0  0  1  1  1;
      0  1  1  0  0  0  1  1;
      0  1  1  0  0  0  0  0;
      0  0  0  1  1  0  0  0;
      0  0  0  1  1  0  0  0;
      0  0  0  1  1  0  0  0;
      0  0  0  0  0  0  0  0];
```

```

X = bwlabel(BW,4)
X =
     0     0     0     0     0     0     0     0
     0     1     1     0     0     3     3     3
     0     1     1     0     0     0     3     3
     0     1     1     0     0     0     0     0
     0     0     0     2     2     0     0     0
     0     0     0     2     2     0     0     0
     0     0     0     2     2     0     0     0
     0     0     0     0     0     0     0     0

```

In the output matrix, the 1's represent one object, the 2's a second object, and the 3's a third. (If you had used 8-connected neighborhoods (the default), there would be only two objects, because the first and second objects would be a single object, connected along the diagonal.)

Viewing a Label Matrix


The label matrix returned by `bwlabel` or `bwlabeln` is of class `double`; it is not a binary image. One way to view it is to display it as a pseudocolor indexed image, using `label2rgb`. In the pseudocolor image, each number that identifies an object in the label matrix is used as an index value into the associated colormap matrix. When you view a label matrix as an RGB image, the objects in the image are easier to distinguish.

To illustrate this technique, this example uses `label2rgb` to view the label matrix `X`. The call to `label2rgb` specifies one of the standard MATLAB colormaps, `jet`. The third argument, `'k'`, specifies the background color (black).

```

X = bwlabel(BW1,4);
RGB = label2rgb(X, @jet, 'k');
imshow(RGB, 'notruesize')

```



Using Color to Distinguish Objects in a Binary Image

Selecting Objects in a Binary Image

You can use the `bwselect` function to select individual objects in a binary image. You specify pixels in the input image, and `bwselect` returns a binary image that includes only those objects from the input image that contain one of the specified pixels.

You can specify the pixels either noninteractively or with a mouse. For example, suppose you want to select objects in the image displayed in the current axes. You type

```
BW2 = bwselect;
```

The cursor changes to crosshairs when it is over the image. Click the objects you want to select; `bwselect` displays a small star over each pixel you click. When you are done, press **Return**. `bwselect` returns a binary image consisting of the objects you selected, and removes the stars.

See the reference page for `bwselect` for more information.

Finding the Area of the Foreground of a Binary Image

The `warea` function returns the area of a binary image. The area is a measure of the size of the foreground of the image. Roughly speaking, the area is the number of on pixels in the image.

`bwarea` does not simply count the number of pixels set to on, however. Rather, `bwarea` weights different pixel patterns unequally when computing the area. This weighting compensates for the distortion that is inherent in representing a continuous image with discrete pixels. For example, a diagonal line of 50 pixels is longer than a horizontal line of 50 pixels. As a result of the weighting `bwarea` uses, the horizontal line has area of 50, but the diagonal line has area of 62.5.

This example uses `bwarea` to determine the percentage area increase in `circbw.tif` that results from a dilation operation.

```
BW = imread('circbw.tif');
SE = ones(5);
BW2 = imdilate(BW,SE);
increase = (bwarea(BW2) - bwarea(BW)) / bwarea(BW);
increase =

    0.3456
```

See the reference page for `bwarea` for more information about the weighting pattern.

Finding the Euler Number of a Binary Image

The `bweuler` function returns the Euler number for a binary image. The Euler number is a measure of the topology of an image. It is defined as the total number of objects in the image minus the number of holes in those objects. You can use either 4- or 8-connected neighborhoods.

This example computes the Euler number for the circuit image, using 8-connected neighborhoods.

```
BW1 = imread('circbw.tif');
eul = bweuler(BW1,8)

eul =


    -85
```

In this example, the Euler number is negative, indicating that the number of holes is greater than the number of objects.

Lookup Table Operations

Certain binary image operations can be implemented most easily through lookup tables. A lookup table is a column vector in which each element represents the value to return for one possible combination of pixels in a neighborhood.

You can use the `makelut` function to create lookup tables for various operations. `makelut` creates lookup tables for 2-by-2 and 3-by-3 neighborhoods. This figure illustrates these types of neighborhoods. Each neighborhood pixel is indicated by an x, and the center pixel is the one with a circle.



For a 2-by-2 neighborhood, there are 16 possible permutations of the pixels in the neighborhood. Therefore, the lookup table for this operation is a 16-element vector. For a 3-by-3 neighborhood, there are 512 permutations, so the lookup table is a 512-element vector.

Once you create a lookup table, you can use it to perform the desired operation by using the `applylut` function.

The example below illustrates using lookup table operations to modify an image containing text. You begin by writing a function that returns 1 if three or more pixels in the 3-by-3 neighborhood are 1; otherwise, it returns 0. You then call `makelut`, passing in this function as the first argument, and using the second argument to specify a 3-by-3 lookup table.

```
f = inline('sum(x(:)) >= 3');
lut = makelut(f,3);
```

`lut` is returned as a 512-element vector of 1's and 0's. Each value is the output from the function for one of the 512 possible permutations.

You then perform the operation using `applylut`.

```
BW1 = imread('text.png');  
BW2 = applylut(BW1,lut);  
imshow(BW1)  
figure, imshow(BW2)
```





Image Before and After Applying Lookup Table Operation

For information about how `applylut` maps pixel combinations in the image to entries in the lookup table, see the reference page for `applylut`.

Note You cannot use `makelut` and `applylut` for neighborhoods of sizes other than 2-by-2 or 3-by-3. These functions support only 2-by-2 and 3-by-3 neighborhoods, because lookup tables are not practical for neighborhoods larger than 3-by-3. For example, a lookup table for a 4-by-4 neighborhood would have 65,536 entries.

Analyzing and Enhancing Images

This chapter describes the Image Processing Toolbox functions that support a range of standard image processing operations for analyzing and enhancing images.

Terminology (p. 10-2)	Provides definitions of image processing terms used in this section
Pixel Values and Statistics (p. 10-3)	Describes the toolbox functions that return information about the data values that make up an image
Image Analysis (p. 10-13)	Describes the toolbox functions that return information about the structure of an image
Texture Analysis (p. 10-26)	Describes the toolbox functions that return information about the texture of an image
Intensity Adjustment (p. 10-36)	Describes the toolbox functions used to improve an image by intensity adjustment
Noise Removal (p. 10-49)	Describes the toolbox functions used to improve an image by removing noise

Terminology

An understanding of the following terms will help you to use this chapter.

Term	Definition
adaptive filter	Filter whose properties vary across an image depending on the local characteristics of the image pixels.
contour	Path in an image along which the image intensity values are equal to a constant.
edge	Curve that follows a path of rapid change in image intensity. Edges are often associated with the boundaries of objects in a scene. Edge detection is used to identify the edges in an image.
property	Quantitative measurement of an image or image region. Examples of image region properties include centroid, bounding box, and area.
histogram	Graph used in image analysis that shows the distribution of intensities in an image. You can use the information in a histogram to choose an appropriate enhancement operation. For example, if an image histogram shows that the range of intensity values is small, you can use an intensity adjustment function to spread the values across a wider range.
noise	Errors in the image acquisition process that result in pixel values that do not reflect the true intensities of the real scene.
profile	Set of intensity values taken from regularly spaced points along a line segment or multiline path in an image. For points that do not fall on the center of a pixel, the intensity values are interpolated.
quadtree decomposition	Image analysis technique that partitions an image into homogeneous blocks.

Pixel Values and Statistics

The Image Processing Toolbox provides several functions that return information about the data values that make up an image. These functions return information about image data in various forms, including

- Data values for selected pixels (`pixval`, `impixel`)
- Data values along a path in an image (`improfile`)
- Contour plot of the image data (`imcontour`)
- Histogram of the image data (`imhist`)
- Summary statistics for the image data (`mean2`, `std2`, `corr2`)
- Feature measurements for image regions (`regionprops`)

Pixel Selection

The toolbox includes two functions that provide information about the values of pixels in an image.

- The `pixval` function interactively displays the data values for pixels as you move the cursor over the image. `pixval` can also display the Euclidean distance between two pixels.
- The `impixel` function returns the data values for a selected pixel or set of pixels. You can supply the coordinates of the pixels as input arguments, or you can select pixels using a mouse.

Note You can also get pixel value information interactively using the Image Tool — see “Getting Information about the Pixels in an Image” on page 3-22.

Getting Pixel Information with `pixval`

To use `pixval`, you first display an image and then enter the `pixval` command. `pixval` installs a black bar at the bottom of the figure, which displays the (x,y) coordinates for whatever pixel the cursor is currently over and the color data for that pixel.

If you click the image and hold down the mouse button while you move the cursor, `pixval` also displays the Euclidean distance between the point you

clicked and the current cursor location. `pixval` draws a line between these points to indicate the distance being measured. When you release the mouse button, the line and the distance display disappear.

Getting Pixel Information with `impixel`

`impixel` provides pixel information in a variable in the MATLAB workspace.

If you call `impixel` with no input arguments, the function associates itself with the image in the current axes. The cursor changes to crosshairs when it is over the image. You can then click the pixels of interest; `impixel` displays a small star over each pixel you select. When you are done selecting pixels, press **Return**. `impixel` returns the color values for the selected pixels, and the stars disappear.

This example illustrates how to use `impixel`:

- 1 Display an image.

```
imshow canoe.tif
```

- 2 Call `impixel` to select points.

```
vals = impixel
```

Click several points in the image to select pixels. When you are finished selecting points, press **Return**.



The `impixel` function returns the pixel values in `vals`.

```
vals =
```

```
0.1294    0.1294    0.1294
0.5176         0         0
0.7765    0.6118    0.4196
```

Intensity Profile

The `improfile` function calculates and plots the intensity values along a line segment or a multiline path in an image. You can supply the coordinates of the line segments as input arguments, or you can define the desired path using a mouse. In either case, `improfile` uses interpolation to determine the values of equally spaced points along the path. (By default, `improfile` uses nearest-neighbor interpolation, but you can specify a different method. For more information, see “Interpolation” on page 5-3.) `improfile` works best with intensity and RGB images.

For a single line segment, `improfile` plots the intensity values in a two-dimensional view. For a multiline path, `improfile` plots the intensity values in a three-dimensional view.

If you call `improfile` with no arguments, the cursor changes to crosshairs when it is over the image. You can then specify line segments by clicking the endpoints; `improfile` draws a line between each two consecutive points you select. When you finish specifying the path, press **Return**. `improfile` displays the plot in a new figure.

In this example, you call `improfile` and specify a single line with the mouse. In this figure, the line is shown in red, and is drawn from top to bottom.

```
I = fitsread('solarspectra.fts');  
imshow(I,[]);  
improfile
```





Image Courtesy of Ann Walker

`improfile` displays a plot of the data along the line. Notice the peaks and valleys and how they correspond to the light and dark bands in the image.



Plot Produced by `improfile`


The example below shows how `improfile` works with an RGB image. Use `imshow` to display the image in a figure window. Call `improfile` without any arguments and trace a line segment in the image interactively. In the figure, the black line indicates a line segment drawn from top to bottom. Double-click to end the line segment.

```
imshow peppers.png
improfile
```



RGB Image with Line Segment Drawn with improfile

The `improfile` function displays a plot of the intensity values along the line segment. The plot includes separate lines for the red, green, and blue intensities. In the plot, notice how low the blue values are at the beginning of the plot where the line traverses the orange pepper.



Plot of Intensity Values Along a Line Segment in an RGB Image


Image Contours

You can use the toolbox function `imcontour` to display a contour plot of the data in an intensity image. This function is similar to the `contour` function in MATLAB, but it automatically sets up the axes so their orientation and aspect ratio match the image.

This example displays an intensity image of grains of rice and a contour plot of the image data:


- 1 Read an intensity image and display it.

```
I = imread('rice.png');  
imshow(I)
```



2 Display a contour plot of the intensity image.

```
figure, imcontour(I,3)
```



You can use the `clabel` function to label the levels of the contours. See the description of `clabel` in the MATLAB Function Reference for details.


Image Histogram

An *image histogram* is a chart that shows the distribution of intensities in an indexed or intensity image. The image histogram function `imhist` creates this plot by making n equally spaced bins, each representing a range of data values. It then calculates the number of pixels within each range.

The following example displays an image of grains of rice and a histogram based on 64 bins. The histogram shows a peak at around 100, corresponding to the dark gray background in the image. For information about how to modify an image by changing the distribution of its histogram, see “Adjusting Intensity Values to a Specified Range” on page 10-37.


1 Read image and display it.

```
I = imread('rice.png');  
imshow(I)
```



2 Display histogram of image.

```
figure, imhist(I)
```



Summary Statistics

You can compute standard statistics of an image using the `mean2`, `std2`, and `corr2` functions. `mean2` and `std2` compute the mean and standard deviation of the elements of a matrix. `corr2` computes the correlation coefficient between two matrices of the same size.

These functions are two-dimensional versions of the `mean`, `std`, and `corrcoef` functions described in the MATLAB Function Reference.

Region Property Measurement

You can use the `regionprops` function to compute properties for image regions. For example, `regionprops` can measure such properties as the area, center of mass, and bounding box for a region you specify. See the reference page for `regionprops` for more information.

Image Analysis

Image analysis techniques return information about the structure of an image. This section describes toolbox functions that you can use for these image analysis techniques:

- “Edge Detection”
- “Boundary Tracing” on page 10-15
- “Line Detection Using the Hough Transform” on page 10-20
- “Quadtree Decomposition” on page 10-24

The toolbox also includes functions that return information about the texture of an image. See “Texture Analysis” on page 10-26 for more information.

Edge Detection

You can use the `edge` function to detect edges, which are those places in an image that correspond to object boundaries. To find edges, this function looks for places in the image where the intensity changes rapidly, using one of these two criteria:

- Places where the first derivative of the intensity is larger in magnitude than some threshold
- Places where the second derivative of the intensity has a zero crossing

`edge` provides a number of derivative estimators, each of which implements one of the definitions above. For some of these estimators, you can specify whether the operation should be sensitive to horizontal edges, vertical edges, or both. `edge` returns a binary image containing 1's where edges are found and 0's elsewhere.

The most powerful edge-detection method that `edge` provides is the Canny method. The Canny method differs from the other edge-detection methods in that it uses two different thresholds (to detect strong and weak edges), and includes the weak edges in the output only if they are connected to strong edges. This method is therefore less likely than the others to be fooled by noise, and more likely to detect true weak edges.

The following example illustrates the power of the Canny edge detector by showing the results of applying the Sobel and Canny edge detectors to the same image:


- 1 Read image and display it.

```
I = imread('coins.png');  
imshow(I)
```




- 2 Apply the Sobel and Canny edge detectors to the image and display them.

```
BW1 = edge(I, 'sobel');  
BW2 = edge(I, 'canny');  
imshow(BW1)  
figure, imshow(BW2)
```



Sobel Filter



Canny Filter

For an interactive demonstration of edge detection, try running `edgedemo`.

Boundary Tracing

The toolbox includes two functions you can use to find the boundaries of objects in a binary image:

- `bwtraceboundary`
- `bwboundaries`

The `bwtraceboundary` function returns the row and column coordinates of all the pixels on the border of an object in an image. You must specify the location of a border pixel on the object as the starting point for the trace.

The `bwboundaries` function returns the row and column coordinates of border pixels of all the objects in an image.

For both functions, the nonzero pixels in the binary image belong to an object and pixels with the value 0 (zero) constitute the background.

The following example uses `bwtraceboundary` to trace the border of an object in a binary image and then uses `bwboundaries` to trace the borders of all the objects in the image:


- 1 Read image and display it.

```
I = imread('coins.png');  
imshow(I)
```



- 2 Convert the image to a binary image. `bwtraceboundary` and `bwboundaries` only work with binary images.

```
BW = im2bw(I);  
imshow(BW)
```



- 3** Determine the row and column coordinates of a pixel on the border of the object you want to trace. `bwboundary` uses this point as the starting location for the boundary tracing.

```
dim = size(BW)
col = round(dim(2)/2) - 90;
row = min(find(BW(:,col)))
```

- 4** Call `bwtraceboundary` to trace the boundary from the specified point. As required arguments, you must specify a binary image, the row and column coordinates of the starting point, and the direction of the first step. The example specifies north ('N'). For information about this parameter, see “Choosing the First Step and Direction for Boundary Tracing” on page 10-18.

```
boundary = bwtraceboundary(BW,[row, col], 'N');
```

- 5** Display the original grayscale image and use the coordinates returned by `bwtraceboundary` to plot the border on the image.

```
imshow(I)
hold on;
plot(boundary(:,2),boundary(:,1), 'g', 'LineWidth',3);
```




- 6** To trace the boundaries of all the coins in the image, use the `bwboundaries` function. By default, `bwboundaries` finds the boundaries of all objects in an image, including objects inside other objects. In the binary image used in this example, some of the coins contain black areas that `bwboundaries` interprets as separate objects. To ensure that `bwboundaries` only traces the coins, use `imfill` to fill the area inside each coin.

```
BW_filled = imfill(BW, 'holes');
boundaries = bwboundaries(BW_filled);
```

`bwboundaries` returns a cell array, where each cell contains the row/column coordinates for an object in the image.

- 7** Plot the borders of all the coins on the original grayscale image using the coordinates returned by `bwboundaries`.

```
for k=1:10
    b = boundaries{k};
    plot(b(:,2),b(:,1),'g','LineWidth',3);
end
```



Choosing the First Step and Direction for Boundary Tracing


For certain objects, you must take care when selecting the border pixel you choose as the starting point and the direction you choose for the first step parameter (north, south, etc.).

For example, if an object contains a hole and you select a pixel on a thin part of the object as the starting pixel, you can trace the outside border of the object or the inside border of the hole, depending on the direction you choose for the first step. For filled objects, the direction you select for the first step parameter is not as important.

To illustrate, this figure shows the pixels traced when the starting pixel is on a thin part of the object and the first step is set to north and south. The connectivity is set to 8 (the default).


FirstStep = North

Direction = Clockwise

 = Boundary pixel = Starting point

FirstStep = South

Direction = Clockwise

**Impact of First Step and Direction Parameters on Boundary Tracing**

Line Detection Using the Hough Transform

The Image Processing Toolbox includes functions that support the Hough transform.

- `hough`
- `houghpeaks`
- `houghlines`

The `hough` function implements the Standard Hough Transform (SHT). The Hough transform is designed to detect lines, using the parametric representation of a line:

$$\rho = x \cdot \cos(\theta) + y \cdot \sin(\theta)$$

The variable `rho` is the distance from the origin to the line along a vector perpendicular to the line. `theta` is the angle between the x-axis and this vector. The `hough` function generates a parameter space matrix whose rows and columns correspond to these `rho` and `theta` values, respectively.

The `houghpeaks` function finds peak values in this space, which represent potential lines in the input image.

The `houghlines` function finds the endpoints of the line segments corresponding to peaks in the Hough transform and it automatically fills in small gaps.


The following example shows how to use these functions to detect lines in an image.

- 1 Read an image into the MATLAB workspace.

```
I = imread('circuit.tif');
```


- 2 For this example, rotate and crop the image.

```
rotI = imrotate(I,33,'crop');
```



- 3** Find the edges in the image.

```
BW = edge(rotI, 'canny');
```




- 4** Compute the Hough transform of the image using the hough function.

```
[H,theta,rho] = hough(BW);
```

- 5** Display the transform.

```
imshow(H,[],'XData',theta,'YData',rho,...  
       'InitialMagnification','fit');  
xlabel('\theta'), ylabel('\rho');  
axis on, axis normal, hold on;
```



- 6** Find the peaks in the Hough transform matrix, H , using the `houghpeaks` function.

```
P = houghpeaks(H,5,'threshold',ceil(0.3*max(H(:))));
```

- 7** Plot the peaks.

```
x = theta(P(:,2));  
y = rho(P(:,1));  
plot(x,y,'s','color','white');
```

- 8** Find lines in the image.


```
lines = houghlines(BW,theta,rho,P,'FillGap',5,'MinLength',7);
```

- 9** Create a plot that superimposes the lines on the original image.

```
figure, imshow(rotI), hold on  
max_len = 0;  
for k = 1:length(lines)  
    xy = [lines(k).point1; lines(k).point2];  
    plot(xy(:,1),xy(:,2),'LineWidth',2,'Color','green');  
  
    % Plot beginnings and ends of lines  
    plot(xy(1,1),xy(1,2),'x','LineWidth',2,'Color','yellow');  
    plot(xy(2,1),xy(2,2),'x','LineWidth',2,'Color','red');
```

```
% Determine the endpoints of the longest line segment
len = norm(lines(k).point1 - lines(k).point2);
if ( len > max_len)
    max_len = len;
    xy_long = xy;
end
end

% highlight the longest line segment
plot(xy_long(:,1),xy_long(:,2),'LineWidth',2,'Color','cyan');
```



Quadtree Decomposition

Quadtree decomposition is an analysis technique that involves subdividing an image into blocks that are more homogeneous than the image itself. This technique reveals information about the structure of the image. It is also useful as the first step in adaptive compression algorithms.

You can perform quadtree decomposition using the `qtdecomp` function. This function works by dividing a square image into four equal-sized square blocks, and then testing each block to see if it meets some criterion of homogeneity (e.g., if all the pixels in the block are within a specific dynamic range). If a block meets the criterion, it is not divided any further. If it does not meet the criterion, it is subdivided again into four blocks, and the test criterion is applied to those blocks. This process is repeated iteratively until each block meets the criterion. The result might have blocks of several different sizes.

Example: Performing Quadtree Decomposition

To illustrate, this example performs quadtree decomposition on a 512-by-512 intensity image. For an interactive demonstration of quadtree decomposition, run the demo `qtdemo`.

- 1 Read in the intensity image.

```
I = imread('liftingbody.png');
```

- 2 Specify the test criteria used to determine the homogeneity of each block in the decomposition. For example, the criterion might be this threshold calculation.

```
max(block(:)) - min(block(:)) <= 0.2
```

You can also supply `qtdecomp` with a function (rather than a threshold value) for deciding whether to split blocks; for example, you might base the decision on the variance of the block. See the reference page for `qtdecomp` for more information.

- 3 Perform this quadtree decomposition by calling the `qtdecomp` function, specifying the image and the threshold value as arguments.

```
S = qtdecomp(I,0.27)
```

You specify the threshold as a value between 0 and 1, regardless of the class of `I`. If `I` is `uint8`, `qtdecomp` multiplies the threshold value by 255 to

determine the actual threshold to use. If I is `uint16`, `qtdecomp` multiplies the threshold value by 65535.

`qtdecomp` first divides the image into four 256-by-256 blocks and applies the test criterion to each block. If a block does not meet the criterion, `qtdecomp` subdivides it and applies the test criterion to each block. `qtdecomp` continues to subdivide blocks until all blocks meet the criterion. Blocks can be as small as 1-by-1, unless you specify otherwise.

`qtdecomp` returns S as a sparse matrix, the same size as I . The nonzero elements of S represent the upper left corners of the blocks; the value of each nonzero element indicates the block size.

The following figure shows the original image and a representation of its quadtree decomposition. (To see how this representation was created, see the example on the `qtdecomp` reference page.) Each black square represents a homogeneous block, and the white lines represent the boundaries between blocks. Notice how the blocks are smaller in areas corresponding to large changes in intensity in the image.




Image Courtesy of NASA




Image and a Representation of Its Quadtree Decomposition

Texture Analysis

The toolbox supports a set of functions that you can use for texture analysis. Texture analysis refers to the characterization of regions in an image by their texture content. Texture analysis attempts to quantify intuitive qualities described by terms such as rough, smooth, silky, or bumpy as a function of the spatial variation in pixel intensities. In this sense, the roughness or bumpiness refers to variations in the intensity values, or gray levels.

Texture analysis is used in a variety of applications, including remote sensing, automated inspection, and medical image processing. Texture analysis can be used to find the texture boundaries, called texture segmentation. Texture analysis can be helpful when objects in an image are more characterized by their texture than by intensity, and traditional thresholding techniques cannot be used effectively.

The toolbox provides two types of texture functions:

- **Texture filter functions** — These functions use standard statistical measures to characterize the local texture of an image. See “Texture Filter Functions” on page 10-26 for more information.
- **Gray-level co-occurrence matrix** — These functions characterize the texture of an image by calculating how often pairs of pixel with specific values and in a specified spatial relationship occur in an image and then extracting statistical measures from this matrix. See “Gray-Level Co-Occurrence Matrix (GLCM)” on page 10-29 for more information

Texture Filter Functions

The toolbox includes three texture analysis functions that filter an image using standard statistical measures, such as range, standard deviation, and entropy. Entropy is a statistical measure of randomness. These statistics can characterize the texture of an image because they provide information about the local variability of the intensity values of pixels in an image.

For example, in areas with smooth texture, the range of values in the neighborhood around a pixel will be a small value; in areas of rough texture, the range will be larger. Similarly, calculating the standard deviation of pixels in a neighborhood can indicate the degree of variability of pixel values in that region.

The following sections provide additional information about the texture functions:

- “Understanding the Texture Filter Functions” on page 10-27
- “Example: Using the Texture Functions” on page 10-28

Understanding the Texture Filter Functions

The three statistical texture filtering functions are

`rangefilt` — Calculates the local range of an image

`stdfilt` — Calculates the local standard deviation of an image

`entropyfilt` — Calculates the local entropy of an intensity image

The functions all operate in a similar way: they define a neighborhood around the pixel of interest and calculate the statistic for that neighborhood.

This example shows how the `rangefilt` function operates on a simple array.

```
A = [ 1 2 3 4 5; 6 7 8 9 10; 11 12 13 14 15; 16 17 18 19 20 ]
```

```
A =
```

```

     1     2     3     4     5
     6     7     8     9    10
    11    12    13    14    15
    16    17    18    19    20
```


```
B = rangefilt(A)
```

```
B =
```

```

     6     7     7     7     6
    11    12    12    12    11
    11    12    12    12    11
     6     7     7     7     6
```

The following figure shows how the value of element $B(2,4)$ was calculated from $A(2,4)$. By default, the `rangefilt` function uses a 3-by-3 neighborhood but you can specify neighborhoods of different shapes and sizes.



Determining Pixel Values in Range Filtered Output Image

The `stdfilt` and `entropyfilt` functions operate similarly, defining a neighborhood around the pixel of interest and calculating the statistic for the neighborhood to determine the pixel value in the output image. The `stdfilt` function calculates the standard deviation of all the values in the neighborhood.

The `entropyfilt` function calculates the entropy of the neighborhood and assigns that value to the output pixel. Note that, by default, the `entropyfilt` function defines a 9-by-9 neighborhood around the pixel of interest. To calculate the entropy of an entire image, use the `entropy` function.

Example: Using the Texture Functions

The following example illustrates how the texture filter functions can detect regions of texture in an image. In the figure, the background is smooth; there is very little variation in the gray-level values. In the foreground, the surface contours of the coins exhibit more texture. In this image, foreground pixels have more variability and thus higher range values. Range filtering makes the edges and contours of the coins more visible.

To see an example of using filtering functions, view the [Texture Segmentation Using Texture Filters](#) demo.

1 Read in the image and display it.


```
I = imread('eight.tif');
```

```
imshow(I)
```



- 2 Filter the image with the `rangefilt` function and display the results. Note how range filtering highlights the edges and surface contours of the coins.

```
K = rangefilt(I);  
figure, imshow(K)
```



Gray-Level Co-Occurrence Matrix (GLCM)

The texture filter functions provide a statistical view of texture based on the image histogram. These functions can provide useful information about the texture of an image but cannot provide information about shape, i.e., the spatial relationships of pixels in an image.

Another statistical method that considers the spatial relationship of pixels is the gray-level co-occurrence matrix (GLCM), also known as the gray-level

spatial dependence matrix. The toolbox provides functions to create a GLCM and derive statistical measurements from it.

This section includes the following topics.

- “Creating a Gray-Level Co-Occurrence Matrix” on page 10-30
- “Specifying the Offsets” on page 10-31
- “Deriving Statistics from a GLCM” on page 10-32
- “Example: Plotting the Correlation” on page 10-33

Creating a Gray-Level Co-Occurrence Matrix


To create a GLCM, use the `graycomatrix` function. The `graycomatrix` function creates a gray-level co-occurrence matrix (GLCM) by calculating how often a pixel with the intensity (gray-level) value i occurs in a specific spatial relationship to a pixel with the value j . By default, the spatial relationship is defined as the pixel of interest and the pixel to its immediate right (horizontally adjacent), but you can specify other spatial relationships between the two pixels. Each element (i,j) in the resultant `glcm` is simply the sum of the number of times that the pixel with value i occurred in the specified spatial relationship to a pixel with value j in the input image.

Because the processing required to calculate a GLCM for the full dynamic range of an image is prohibitive, `graycomatrix` scales the input image. By default, `graycomatrix` uses scaling to reduce the number of intensity values in intensity image from 256 to eight. The number of gray levels determines the size of the GLCM. To control the number of gray levels in the GLCM and the scaling of intensity values, using the `NumLevels` and the `GrayLimits` parameters of the `graycomatrix` function. See the `graycomatrix` reference page for more information.

The gray-level co-occurrence matrix can reveal certain properties about the spatial distribution of the gray levels in the texture image. For example, if most of the entries in the GLCM are concentrated along the diagonal, the texture is coarse with respect to the specified offset. You can also derive several statistical measures from the GLCM. See “Deriving Statistics from a GLCM” on page 10-32 for more information.

To illustrate, the following figure shows how `graycomatrix` calculates the first three values in a GLCM. In the output GLCM, element (1,1) contains the value 1 because there is only one instance in the input image where two horizontally

adjacent pixels have the values 1 and 1, respectively. $glcm(1,2)$ contains the value 2 because there are two instances where two horizontally adjacent pixels have the values 1 and 2. Element (1,3) in the GLCM has the value 0 because there are no instances of two horizontally adjacent pixels with the values 1 and 3. `graycomatrix` continues processing the input image, scanning the image for other pixel pairs (i,j) and recording the sums in the corresponding elements of the GLCM.



Process Used to Create the GLCM

Specifying the Offsets


By default, the `graycomatrix` function creates a single GLCM, with the spatial relationship, or *offset*, defined as two horizontally adjacent pixels. However, a single GLCM might not be enough to describe the textural features of the input image. For example, a single horizontal offset might not be sensitive to texture with a vertical orientation. For this reason, `graycomatrix` can create multiple GLCMs for a single input image.

To create multiple GLCMs, specify an array of offsets to the `graycomatrix` function. These offsets define pixel relationships of varying direction and distance. For example, you can define an array of offsets that specify four directions (horizontal, vertical, and two diagonals) and four distances. In this case, the input image is represented by 16 GLCMs. When you calculate statistics from these GLCMs, you can take the average.

You specify these offsets as a p -by-2 array of integers. Each row in the array is a two-element vector, $[\text{row_offset}, \text{col_offset}]$, that specifies one offset. row_offset is the number of rows between the pixel of interest and its neighbor. col_offset is the number of columns between the pixel of interest and its neighbor. This example creates an offset that specifies four directions and 4 distances for each direction. For more information about specifying offsets, see the `graycomatrix` reference page.

```
offsets = [ 0 1; 0 2; 0 3; 0 4;...
           -1 1; -2 2; -3 3; -4 4;...
           -1 0; -2 0; -3 0; -4 0;...
           -1 -1; -2 -2; -3 -3; -4 -4];
```

The figure illustrates the spatial relationships of pixels that are defined by this array of offsets, where D represents the distance from the pixel of interest.



Deriving Statistics from a GLCM

After you create the GLCMs, you can derive several statistics from them using the `graycoprops` function. These statistics provide information about the texture of an image. The following table lists the statistics you can derive. You specify the statistics you want when you call the `graycoprops` function. For

detailed information about these statistics, see the graycoprops reference page.


Statistic	Description
Contrast	Measures the local variations in the gray-level co-occurrence matrix.
Correlation	Measures the joint probability occurrence of the specified pixel pairs.
Energy	Provides the sum of squared elements in the GLCM. Also known as uniformity or the angular second moment.
Homogeneity	Measures the closeness of the distribution of elements in the GLCM to the GLCM diagonal.

Example: Plotting the Correlation

This example shows how to create a set of GLCMs and derive statistics from them and illustrates how the statistics returned by graycoprops have a direct relationship to the original input image.

- 1 Read in an intensity image and display it. The example converts the truecolor (RGB) image to a grayscale intensity image and then rotates it 90° for this example.

```
circuitBoard = rot90(rgb2gray(imread('board.tif')));  
imshow(circuitBoard)
```



- 2 Define offsets of varying direction and distance. Because the image contains objects of a variety of shapes and sizes that are arranged in horizontal and vertical directions, the example specifies a set of horizontal offsets that only vary in distance.

```
offsets0 = [zeros(40,1) (1:40)'];
```

- 3 Create the GLCMs. Call the `graycomatrix` function specifying the offsets.


```
glcms = graycomatrix(circuitBoard,'Offset',offsets0)
```

- 4 Derive statistics from the GLCMs using the `graycoprops` function. The example calculates the contrast and correlation.


```
stats = graycoprops(glcms,'Contrast Correlation');
```

- 5 Plot correlation as a function of offset.

```
figure, plot([stats.Correlation]);
title('Texture Correlation as a function of offset');
xlabel('Horizontal Offset')
ylabel('Correlation')
```



The plot contains peaks at offsets 7, 15, 23, and 30. If you examine the input image closely, you can see that certain vertical elements in the image have a periodic pattern that repeats every seven pixels. The following figure shows the upper left corner of the image and points out where this pattern occurs.




Intensity Adjustment

Image enhancement techniques are used to improve an image, where “improve” is sometimes defined objectively (e.g., increase the signal-to-noise ratio), and sometimes subjectively (e.g., make certain features easier to see by modifying the colors or intensities).

Intensity adjustment is an image enhancement technique that maps an image’s intensity values to a new range. To illustrate, this figure shows a low-contrast image with its histogram. Notice in the histogram of the image how all the values gather in the center of the range.

```
I = imread('pout.tif');  
imshow(I)  
figure, imhist(I,64)
```



If you remap the data values to fill the entire intensity range [0, 255], you can increase the contrast of the image. The following sections describe several intensity adjustment techniques, including

- “Adjusting Intensity Values to a Specified Range” on page 10-37
- “Histogram Equalization” on page 10-41
- “Contrast-Limited Adaptive Histogram Equalization” on page 10-43
- “Decorrelation Stretching” on page 10-44

The functions described in this section apply primarily to intensity images. However, some of these functions can be applied to color images as well. For information about how these functions work with color images, see the reference pages for the individual functions.


Adjusting Intensity Values to a Specified Range

You can adjust the intensity values in an image using the `imadjust` function, where you specify the range of intensity values in the output image.

For example, this code increases the contrast in a low-contrast intensity image by remapping the data values to fill the entire intensity range [0, 255].

```
I = imread('pout.tif');
J = imadjust(I);
imshow(J)
figure, imhist(J,64)
```

This figure displays the adjusted image and its histogram. Notice the increased contrast in the image, and that the histogram now fills the entire range.



Adjusted Image and Its Histogram

Specifying the Adjustment Limits

You can optionally specify the range of the input values and the output values using `imadjust`. You specify these ranges in two vectors that you pass to `imadjust` as arguments. The first vector specifies the low- and high-intensity

values that you want to map. The second vector specifies the scale over which you want to map them.

Note Note that you must specify the intensities as values between 0 and 1 regardless of the class of I. If I is `uint8`, the values you supply are multiplied by 255 to determine the actual values to use; if I is `uint16`, the values are multiplied by 65535. To learn about an alternative way to set these limits automatically, see “Setting the Adjustment Limits Automatically” on page 10-39.

For example, you can decrease the contrast of an image by narrowing the range of the data. In the example below, the man’s coat is too dark to reveal any detail. `imadjust` maps the range [0,51] in the `uint8` input image to [128,255] in the output image. This brightens the image considerably, and also widens the dynamic range of the dark portions of the original image, making it much easier to see the details in the coat. Note, however, that because all values above 51 in the original image are mapped to 255 (white) in the adjusted image, the adjusted image appears washed out.

```
I = imread('cameraman.tif');  
J = imadjust(I,[0 0.2],[0.5 1]);  
imshow(I)  
figure, imshow(J)
```



Image Courtesy of MIT

Image After Remapping and Widening the Dynamic Range

Setting the Adjustment Limits Automatically

To use `imadjust`, you must typically perform two steps:

- 1 View the histogram of the image to determine the intensity value limits.
- 2 Specify these limits as a fraction between 0.0 and 1.0 so that you can pass them to `imadjust` in the `[low_in high_in]` vector.

For a more convenient way to specify these limits, use the `stretchlim` function. (The `imadjust` function uses `stretchlim` for its simplest syntax, `imadjust(I)`.)

This function calculates the histogram of the image and determines the adjustment limits automatically. The `stretchlim` function returns these values as fractions in a vector that you can pass as the `[low_in high_in]` argument to `imadjust`; for example:

```
I = imread('rice.png');  
J = imadjust(I,stretchlim(I),[0 1]);
```

By default, `stretchlim` uses the intensity values that represent the bottom 1% (0.01) and the top 1% (0.99) of the range as the adjustment limits. By trimming the extremes at both ends of the intensity range, `stretchlim` makes more room in the adjusted dynamic range for the remaining intensities. But you can specify other range limits as an argument to `stretchlim`. See the `stretchlim` reference page for more information.

Gamma Correction


`imadjust` maps low to bottom, and high to top. By default, the values between low and high are mapped linearly to values between bottom and top. For example, the value halfway between low and high corresponds to the value halfway between bottom and top.

`imadjust` can accept an additional argument that specifies the *gamma correction* factor. Depending on the value of gamma, the mapping between values in the input and output images might be nonlinear. For example, the value halfway between low and high might map to a value either greater than or less than the value halfway between bottom and top.

Gamma can be any value between 0 and infinity. If gamma is 1 (the default), the mapping is linear. If gamma is less than 1, the mapping is weighted toward

higher (brighter) output values. If gamma is greater than 1, the mapping is weighted toward lower (darker) output values.

The figure below illustrates this relationship. The three transformation curves show how values are mapped when gamma is less than, equal to, and greater than 1. (In each graph, the x -axis represents the intensity values in the input image, and the y -axis represents the intensity values in the output image.)



Plots Showing Three Different Gamma Correction Settings

The example below illustrates gamma correction. Notice that in the call to `imadjust`, the data ranges of the input and output images are specified as empty matrices. When you specify an empty matrix, `imadjust` uses the default range of [0,1]. In the example, both ranges are left empty; this means that gamma correction is applied without any other adjustment of the data.

```
[X,map] = imread('forest.tif')
I = ind2gray(X,map);
J = imadjust(I,[],[],0.5);
imshow(I)
figure, imshow(J)
```



Image Courtesy of Susan Cohen

Image Before and After Applying Gamma Correction

Histogram Equalization

The process of adjusting intensity values can be done automatically by the `histeq` function. `histeq` performs *histogram equalization*, which involves transforming the intensity values so that the histogram of the output image approximately matches a specified histogram. (By default, `histeq` tries to match a flat histogram with 64 bins, but you can specify a different histogram instead; see the reference page for `histeq`.)

This example illustrates using `histeq` to adjust an intensity image. The original image has low contrast, with most values in the middle of the intensity range. `histeq` produces an output image having values evenly distributed throughout the range.

```
I = imread('pout.tif');  
J = histeq(I);  
imshow(J)  
figure, imhist(J,64)
```





Image After Histogram Equalization with Its Histogram

`histeq` can return a 1-by-256 vector that shows, for each possible input value, the resulting output value. (The values in this vector are in the range [0,1], regardless of the class of the input image.) You can plot this data to get the transformation curve. For example:

```
I = imread('pout.tif');  
[J,T] = histeq(I);  
figure,plot((0:255)/255,T);
```



Notice how this curve reflects the histograms in the previous figure, with the input values mostly between 0.3 and 0.6, while the output values are distributed evenly between 0 and 1.

For an interactive demonstration of intensity adjustment, try running `imadjdemo`.

Contrast-Limited Adaptive Histogram Equalization

As an alternative to using `histeq`, you can perform contrast-limited adaptive histogram equalization (CLAHE) using the `adapthisteq` function. While `histeq` works on the entire image, `adapthisteq` operates on small regions in the image, called *tiles*. Each tile's contrast is enhanced, so that the histogram of the output region approximately matches a specified histogram. After performing the equalization, `adapthisteq` combines neighboring tiles using bilinear interpolation to eliminate artificially induced boundaries.

To avoid amplifying any noise that might be present in the image, you can use `adapthisteq` optional parameters to limit the contrast, especially in homogeneous areas.

To illustrate, this example uses `adapthisteq` to adjust the contrast in an intensity image. The original image has low contrast, with most values in the middle of the intensity range. `adapthisteq` produces an output image having values evenly distributed throughout the range.

```
I = imread('pout.tif');  
J = adapthisteq(I);  
imshow(I)  
figure, imshow(J)
```




Image After CLAHE Equalization with Its Histogram

Decorrelation Stretching

Decorrelation stretching enhances the color separation of an image with significant band-band correlation. The exaggerated colors improve visual interpretation and make feature discrimination easier. You apply decorrelation stretching with the `decorrstretch` function. See “Adding a Linear Contrast Stretch” on page 10-47 on how to add an optional linear contrast stretch to the decorrelation stretch.

The number of color bands, `NBANDS`, in the image is usually three. But you can apply decorrelation stretching regardless of the number of color bands.

The original color values of the image are mapped to a new set of color values with a wider range. The color intensities of each pixel are transformed into the

color eigenspace of the NBANDS-by-NBANDS covariance or correlation matrix, stretched to equalize the band variances, then transformed back to the original color bands.

To define the bandwise statistics, you can use the entire original image or, with the `subset` option, any selected subset of it. See the `decorrstretch` reference page.

Simple Decorrelation Stretching

You can apply decorrelation and stretching operations on the library of images available in the `imdemos` directory. The library includes a LANDSAT image of the Little Colorado River. In this example, you perform a simple decorrelation stretch on this image:

- 1 The image has seven bands, but just read in the three visible colors:

```
A = multibandread('littlecoriver.lan', [512, 512, 7], ...
    'uint8=>uint8', 128, 'bil', 'ieee-le', ...
    {'Band', 'Direct', [3 2 1]});
```


- 2 Then perform the decorrelation stretch:

```
B = decorrstretch(A);
```

- 3 Now view the results:

```
imshow(A); figure; imshow(B)
```


Compare the two images. The original has a strong violet (red-bluish) tint, while the transformed image has a somewhat expanded color range.



Little Colorado River Before (left) and After (right) Decorrelation Stretch

A color band scatterplot of the images shows how the bands are decorrelated and equalized:

```
rA = A(:,:,1); gA = A(:,:,2); bA = A(:,:,3);  
figure, plot3(rA(:),gA(:),bA(:),'.'); grid('on')  
xlabel('Red (Band 3)'); ylabel('Green (Band 2)'); ...  
    zlabel('Blue (Band 1)')  
rB = B(:,:,1); gB = B(:,:,2); bB = B(:,:,3);  
figure, plot3(rB(:),gB(:),bB(:),'.'); grid('on')  
xlabel('Red (Band 3)'); ylabel('Green (Band 2)'); ...  
    zlabel('Blue (Band 1)')
```




Color Scatterplot Before (left) and After (right) Decorrelation Stretch

Adding a Linear Contrast Stretch

Now try the same transformation, but with a linear contrast stretch applied after the decorrelation stretch:

```
imshow(A); C = decorrstretch(A,'Tol',0.01); figure; imshow(C)
```

Compare the transformed image to the original.



Little Colorado River After Decorrelation Stretch Followed by Linear Contrast Stretch

Adding the linear contrast stretch enhances the resulting image by further expanding the color range. In this case, the transformed color range is mapped within each band to a normalized interval between 0.01 and 0.99, saturating 2%.

See the `stretchlim` function reference page for more about `To1`. Without the `To1` option, `decorrstretch` applies no linear contrast stretch.

Note You can apply a linear contrast stretch as a separate operation after performing a decorrelation stretch, using `stretchlim` and `imadjust`. This alternative, however, often gives inferior results for `uint8` and `uint16` images, because the pixel values must be clamped to `[0 255]` (or `[0 65535]`). The `To1` option in `decorrstretch` circumvents this limitation.

Noise Removal

Digital images are prone to a variety of types of noise. There are several ways that noise can be introduced into an image, depending on how the image is created. For example:

- If the image is scanned from a photograph made on film, the film grain is a source of noise. Noise can also be the result of damage to the film, or be introduced by the scanner itself.
- If the image is acquired directly in a digital format, the mechanism for gathering the data (such as a CCD detector) can introduce noise.
- Electronic transmission of image data can introduce noise.

The toolbox provides a number of different ways to remove or reduce noise in an image. Different methods are better for different kinds of noise. The methods available include

- “Using Linear Filtering”
- “Using Median Filtering”
- “Using Adaptive Filtering” on page 10-52

To simulate the effects of some of the problems listed above, the toolbox provides the `imnoise` function, which you can use to *add* various types of noise to an image. The examples in this section use this function.

Using Linear Filtering

You can use linear filtering to remove certain types of noise. Certain filters, such as averaging or Gaussian filters, are appropriate for this purpose. For example, an averaging filter is useful for removing grain noise from a photograph. Because each pixel gets set to the average of the pixels in its neighborhood, local variations caused by grain are reduced.

See “Linear Filtering” on page 7-4 for more information.

Using Median Filtering

Median filtering is similar to using an averaging filter, in that each output pixel is set to an average of the pixel values in the neighborhood of the corresponding input pixel. However, with median filtering, the value of an output pixel is determined by the *median* of the neighborhood pixels, rather

than the mean. The median is much less sensitive than the mean to extreme values (called *outliers*). Median filtering is therefore better able to remove these outliers without reducing the sharpness of the image. The `medfilt2` function implements median filtering.

Note Median filtering is a specific case of *order-statistic filtering*, also known as *rank filtering*. For information about order-statistic filtering, see the reference page for the `ordfilt2` function.

The following example compares using an averaging filter and `medfilt2` to remove *salt and pepper* noise. This type of noise consists of random pixels' being set to black or white (the extremes of the data range). In both cases the size of the neighborhood used for filtering is 3-by-3.

1 Read in the image and display it.

```
I = imread('eight.tif');  
imshow(I)
```




2 Add noise to it.

```
J = imnoise(I,'salt & pepper',0.02);  
figure, imshow(J)
```



- 3** Filter the noisy image with an averaging filter and display the results.

```
K = filter2(fspecial('average',3),J)/255;  
figure, imshow(K)
```



- 4** Now use a median filter to filter the noisy image and display the results. Notice that `medfilt2` does a better job of removing noise, with less blurring of edges.

```
L = medfilt2(J,[3 3]);  
figure, imshow(K)  
figure, imshow(L)
```





Using Adaptive Filtering

The `wiener2` function applies a Wiener filter (a type of linear filter) to an image *adaptively*, tailoring itself to the local image variance. Where the variance is large, `wiener2` performs little smoothing. Where the variance is small, `wiener2` performs more smoothing.

This approach often produces better results than linear filtering. The adaptive filter is more selective than a comparable linear filter, preserving edges and other high-frequency parts of an image. In addition, there are no design tasks; the `wiener2` function handles all preliminary computations and implements the filter for an input image. `wiener2`, however, does require more computation time than linear filtering.

`wiener2` works best when the noise is constant-power (“white”) additive noise, such as Gaussian noise. The example below applies `wiener2` to an image of Saturn that has had Gaussian noise added. For an interactive demonstration of filtering to remove noise, try running `nrfilt demo`.

```
RGB = imread('saturn.png');  
I = rgb2gray(RGB);  
J = imnoise(I,'gaussian',0,0.005);  
K = wiener2(J,[5 5]);  
imshow(J)  
figure, imshow(K)
```



Original Image Courtesy of NASA

Noisy Version (left) and Filtered Version (right)

Region-Based Processing

This chapter describes operations that you can perform on a selected region of an image.

Terminology (p. 11-2)	Provides definitions of image processing terms used in this section
Specifying a Region of Interest (p. 11-3)	Describes how to specify a region of interest using the <code>roipoly</code> function
Filtering a Region (p. 11-6)	Describes how to apply a filter to a region using the <code>roifilt2</code> function
Filling a Region (p. 11-9)	Describes how to fill a region of interest using the <code>roifill</code> function

Terminology

An understanding of the following terms will help you to use this section.

Term	Definition
binary mask	Binary image with the same size as the image you want to process. The mask contains 1's for all pixels that are part of the region of interest, and 0's everywhere else.
filling a region	Process that fills a region of interest by interpolating the pixel values from the borders of the region. This process can be used to make objects in an image seem to disappear as they are replaced with values that blend in with the background area.
filtering a region	Process of applying a filter to a region of interest. For example, you can apply an intensity adjustment filter to certain regions of an image.
interpolation	Method used to estimate an image value at a location in between image pixels.
masked filtering	Operation that applies filtering only to the regions of interest in an image that are identified by a binary mask. Filtered values are returned for pixels where the binary mask contains 1's; unfiltered values are returned for pixels where the binary mask contains 0's.
region of interest	Portion of an image that you want to filter or perform some other operation on. You define a region of interest by creating a binary mask. There can be more than one region defined in an image. The regions can be geographic in nature, such as polygons that encompass contiguous pixels, or they can be defined by a range of intensities. In the latter case, the pixels are not necessarily contiguous.

Specifying a Region of Interest

A *region of interest* is a portion of an image that you want to filter or perform some other operation on. You define a region of interest by creating a *binary mask*, which is a binary image with the same size as the image you want to process. The mask contains 1's for all pixels that are part of the region of interest, and 0's everywhere else.

The following subsections discuss methods for creating binary masks:

- “Selecting a Polygon” on page 11-3
- “Other Selection Methods” on page 11-4 (using any binary mask or the `roicolor` function)

For an interactive demonstration of region-based processing, try running `roidemo`.

Selecting a Polygon

You can use the `roipoly` function to specify a polygonal region of interest. If you call `roipoly` with no input arguments, the cursor changes to crosshairs when it is over the image displayed in the current axes. You can then specify the vertices of the polygon by clicking points in the image with the mouse. When you are done selecting vertices, press **Return**; `roipoly` returns a binary image of the same size as the input image, containing 1's inside the specified polygon, and 0's everywhere else.


The example below illustrates using the interactive syntax of `roipoly` to create a binary mask. In the figure, the border of the selected region that was created using a mouse is shown in red.

```
I = imread('pout.tif');  
imshow(I)  
BW = roipoly;
```



Polygonal Region of Interest Selected Using `roipoly`

```
imshow(BW)
```



Binary Mask Created for the Region Shown in the Preceding Figure

You can also use `roipoly` noninteractively. See the reference page for `roipoly` for more information.

Other Selection Methods

`roipoly` provides an easy way to create a binary mask. However, you can use *any* binary image as a mask, provided that the binary image is the same size as the image being filtered.

For example, suppose you want to filter the intensity image `I`, filtering only those pixels whose values are greater than 0.5. You can create the appropriate mask with this command.

```
BW = (I > 0.5);
```

You can also use the `poly2mask` function to create a binary mask. Unlike the `roipoly` function, `poly2mask` does not require an input image. For more information, see the `poly2mask` reference page.

You can also use the `roicolor` function to define the region of interest based on a color or intensity range. For more information, see the reference page for `roicolor`.

Filtering a Region

You can use the `roifilt2` function to process a region of interest. When you call `roifilt2`, you specify an intensity image, a binary mask, and a filter. `roifilt2` filters the input image and returns an image that consists of filtered values for pixels where the binary mask contains 1's and unfiltered values for pixels where the binary mask contains 0's. This type of operation is called *masked filtering*.

Note `roifilt2` is best suited to operations that return data in the same range as in the original image, because the output image takes some of its data directly from the input image. Certain filtering operations can result in values outside the normal image data range (i.e., [0,1] for images of class `double`, [0,255] for images of class `uint8`, and [0,65535] for images of class `uint16`). For more information, see the reference page for `roifilt2`.

Example: Filtering a Region in an Image

This example uses masked filtering to increase the contrast of a specific region of an image:

- 1 Read in the image.

```
I = imread('pout.tif');
```

- 2 Create the mask.

This example uses the mask `BW` created in “Selecting a Polygon” on page 11-3. The region of interest specified by the mask is the logo on the girl's jacket.

- 3 Create the filter.

```
h = fspecial('unsharp');
```

- 4 Call `roifilt2`, specifying the image to be filtered, the mask, and the filter.

```
I2 = roifilt2(h,I,BW);  
imshow(I)  
figure, imshow(I2)
```



Image Before and After Using an Unsharp Filter on the Region of Interest

Specifying the Filtering Operation

`roifilt2` also enables you to specify your own function to operate on the region of interest. This example uses the `imadjust` function to lighten parts of an image:

- 1 Read in the image.

```
I = imread('cameraman.tif');
```

- 2 Create the mask. In this example, the mask is a binary image containing text. The mask image must be cropped to be the same size as the image to be filtered.

```
BW = imread('text.png');  
mask = BW(1:256,1:256);
```

- 3 Create the filter.

```
f = inline('imadjust(x,[],[],0.3)');
```

- 4 Call `roifilt2`, specifying the image to be filtered, the mask, and the filter. The resulting image, `I2`, has the text imprinted on it.

```
I2 = roifilt2(I,mask,f);  
imshow(I2)
```



Image Brightened Using a Binary Mask Containing Text

Filling a Region

You can use the `roifill` function to fill a region of interest, interpolating from the borders of the region. This function is useful for image editing, including removal of extraneous details or artifacts.

`roifill` performs the fill operation using an interpolation method based on Laplace's equation. This method results in the smoothest possible fill, given the values on the boundary of the region.

As with `roipoly`, you select the region of interest with the mouse. When you complete the selection, `roifill` returns an image with the selected region filled in.

This example shows how to use `roifill` to fill a region in an image.

- 1 Read an image into the MATLAB workspace and display it. Because the image is an indexed image, the example uses `ind2gray` to convert it to a grayscale image.

```
load trees
I = ind2gray(X,map);
imshow(I)
```

- 2 Call `roifill` and then use the mouse to select the region you want to fill (shown in red in the following figure). `roifill` returns the modified image in `I2`.

```
I2 = roifill;
```



- 3 Display the modified image. Note how the region of interest defined in the previous step has been filled.

```
imshow(I2)
```



Image Deblurring

This chapter describes how to deblur an image using the toolbox deblurring functions.

Terminology (p. 12-2)	Provides definitions of image processing terms used in this section
Understanding Deblurring (p. 12-3)	Defines deblurring and deconvolution
Using the Deblurring Functions (p. 12-6)	Provides step-by-step examples of using <code>deconvwnr</code> , <code>deconvreg</code> , <code>deconvlucy</code> , and <code>deconvblind</code> functions
Avoiding Ringing in Deblurred Images (p. 12-23)	Describes how to use the <code>edgetaper</code> function to avoid “ringing” in deblurred images

Terminology

An understanding of the following terms will help you to use this chapter.

Term	MATLAB Definition
deconvolution	Process of reversing the effect of convolution.
distortion operator	Operator that describes a process causing the acquired image to be different from the original scene. Distortion caused by a point spread function (see below) is just one type of distortion.
optical transfer function (OTF)	In the frequency domain, the OTF describes the response of a linear, position-invariant system to an impulse. The OTF is the Fourier transform of the point spread function (PSF).
point spread function (PSF)	In the spatial domain, the PSF describes the degree to which an optical system blurs (spreads) a point of light. The PSF is the inverse Fourier transform of the OTF.

Understanding Deblurring

This section provides some background on deblurring techniques. The section includes these topics:

- “Causes of Blurring”
- “Deblurring Model”

Causes of Blurring

The blurring, or degradation, of an image can be caused by many factors:

- Movement during the image capture process, by the camera or, when long exposure times are used, by the subject
- Out-of-focus optics, use of a wide-angle lens, atmospheric turbulence, or a short exposure time, which reduces the number of photons captured
- Scattered light distortion in confocal microscopy

Deblurring Model

A blurred or degraded image can be approximately described by this equation $\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}$, where

- g** The blurred image
- H** The distortion operator, also called the *point spread function* (PSF). This function, when convolved with the image, creates the distortion.
- f** The original true image
- n** Additive noise, introduced during image acquisition, that corrupts the image

Note The image **f** really doesn't exist. This image represents what you would have if you had perfect image acquisition conditions.

Importance of the PSF

Based on this model, the fundamental task of deblurring is to deconvolve the blurred image with the PSF that exactly describes the distortion.

Note The quality of the deblurred image is mainly determined by knowledge of the PSF.

To illustrate, this example takes a clear image and deliberately blurs it by convolving it with a PSF. The example uses the `fspecial` function to create a PSF that simulates a motion blur, specifying the length of the blur in pixels, (`LEN=31`), and the angle of the blur in degrees (`THETA=11`). Once the PSF is created, the example uses the `imfilter` function to convolve the PSF with the original image, `I`, to create the blurred image, `Blurred`. (To see how deblurring is the reverse of this process, using the same images, see “Deblurring with the Wiener Filter” on page 12-7.)

```
I = imread('peppers.png');  
I = I(60+[1:256],222+[1:256],:); % crop the image  
figure; imshow(I); title('Original Image');
```



```
LEN = 31;  
THETA = 11;  
PSF = fspecial('motion',LEN,THETA); % create PSF  
Blurred = imfilter(I,PSF,'circular','conv');  
figure; imshow(Blurred); title('Blurred Image');
```



Using the Deblurring Functions

The toolbox includes four deblurring functions, listed here in order of complexity:

<code>deconvwnr</code>	Implements deblurring using the Wiener filter
<code>deconvreg</code>	Implements deblurring using a regularized filter
<code>deconvlucy</code>	Implements deblurring using the Lucy-Richardson algorithm
<code>deconvblind</code>	Implements deblurring using the blind deconvolution algorithm

All the functions accept a PSF and the blurred image as their primary arguments. The `deconvwnr` function implements a least squares solution. The `deconvreg` function implements a constrained least squares solution, where you can place constraints on the output image (the smoothness requirement is the default). With either of these functions, you should provide some information about the noise to reduce possible noise amplification during deblurring.

The `deconvlucy` function implements an accelerated, damped Lucy-Richardson algorithm. This function performs multiple iterations, using optimization techniques and Poisson statistics. With this function, you do not need to provide information about the additive noise in the corrupted image.

The `deconvblind` function implements the blind deconvolution algorithm, which performs deblurring without knowledge of the PSF. When you call `deconvblind`, you pass as an argument your initial guess at the PSF. The `deconvblind` function returns a restored PSF in addition to the restored image. The implementation uses the same damping and iterative model as the `deconvlucy` function.

Note You might need to perform many iterations of the deblurring process, varying the parameters you specify to the deblurring functions with each iteration, until you achieve an image that, based on the limits of your information, is the best approximation of the original scene. Along the way, you must make numerous judgments about whether newly uncovered features in the image are features of the original scene or simply artifacts of the deblurring process.

For information about creating your own deblurring functions, see “Creating Your Own Deblurring Functions” on page 12-22. To avoid “ringing” in a deblurred image, you can use the `edgetaper` function to preprocess your image before passing it to the deblurring functions. See “Avoiding Ringing in Deblurred Images” on page 12-23 for more information.

Deblurring with the Wiener Filter

Use the `deconvwnr` function to deblur an image using the Wiener filter. Wiener deconvolution can be used effectively when the frequency characteristics of the image and additive noise are known, to at least some degree. In the absence of noise, the Wiener filter reduces to the ideal inverse filter.

This example deblurs the blurred image created in “Deblurring Model” on page 12-3, specifying the same PSF function that was used to create the blur. This example illustrates the importance of knowing the PSF, the function that caused the blur. When you know the exact PSF, the results of deblurring can be quite effective.

- 1 Read an image into the MATLAB workspace. (To speed the deblurring operation, the example also crops the image.)

```
I = imread('peppers.png');  
I = I(10+[1:256],222+[1:256],:);  
figure;imshow(I);title('Original Image');
```



2 Create a PSF.

```
LEN = 31;  
THETA = 11;  
PSF = fspecial('motion',LEN,THETA);
```

3 Create a simulated blur in the image.

```
Blurred = imfilter(I,PSF,'circular','conv');  
figure; imshow(Blurred);title('Blurred Image');
```



4 Deblur the image.

```
wnr1 = deconvwnr(Blurred,PSF);  
figure;imshow(wnr1);  
title('Restored, True PSF');
```



Refining the Result

You can affect the deconvolution results by providing values for the optional arguments supported by the `deconvwnr` function. Using these arguments you can specify the noise-to-signal power value and/or provide autocorrelation functions to help refine the result of deblurring. To see the impact of these optional arguments, view the Image Processing Toolbox deblurring demos.

Deblurring with a Regularized Filter

Use the `deconvreg` function to deblur an image using a regularized filter. A regularized filter can be used effectively when limited information is known about the additive noise.

To illustrate, this example simulates a blurred image by convolving a Gaussian filter PSF with an image (using `imfilter`). Additive noise in the image is simulated by adding Gaussian noise of variance V to the blurred image (using `imnoise`):

- 1 Read an image into the MATLAB workspace. The example uses cropping to reduce the size of the image to be deblurred. This is not a required step in deblurring operations.

```
I = imread('tissue.png');  
I = I(125+[1:256],1:256,:);  
figure; imshow(I); title('Original Image');
```




Image Courtesy Alan W. Partin

- 2 Create the PSF.

```
PSF = fspecial('gaussian',11,5);
```


- 3 Create a simulated blur in the image and add noise.

```
Blurred = imfilter(I,PSF,'conv');
```

```
V = .02;
```

```
BlurredNoisy = imnoise(Blurred,'gaussian',0,V);
```


```
figure;imshow(BlurredNoisy);title('Blurred and Noisy Image');
```



- 4 Use deconvreg to deblur the image, specifying the PSF used to create the blur and the noise power, NP.

```
NP = V*prod(size(I));
```

```
[reg1 LAGRA] = deconvreg(BlurredNoisy,PSF,NP);  
figure,imshow(reg1),title('Restored Image');
```



Refining the Result

You can affect the deconvolution results by providing values for the optional arguments supported by the `deconvreg` function. Using these arguments you can specify the noise power value, the range over which `deconvreg` should iterate as it converges on the optimal solution, and the regularization operator to constrain the deconvolution. To see the impact of these optional arguments, view the Image Processing Toolbox deblurring demos.

Deblurring with the Lucy-Richardson Algorithm

Use the `deconvlucy` function to deblur an image using the accelerated, damped, Lucy-Richardson algorithm. The algorithm maximizes the likelihood that the resulting image, when convolved with the PSF, is an instance of the blurred image, assuming Poisson noise statistics. This function can be effective when you know the PSF but know little about the additive noise in the image.

The `deconvlucy` function implements several adaptations to the original Lucy-Richardson maximum likelihood algorithm that address complex image restoration tasks. Using these adaptations, you can

- Reduce the effect of noise amplification on image restoration
- Account for nonuniform image quality (e.g., bad pixels, flat-field variation)
- Handle camera read-out and background noise
- Improve the restored image resolution by subsampling

The following sections provide more information about each of these adaptations.

Reducing the Effect of Noise Amplification

Noise amplification is a common problem of maximum likelihood methods that attempt to fit data as closely as possible. After many iterations, the restored image can have a speckled appearance, especially for a smooth object observed at low signal-to-noise ratios. These speckles do not represent any real structure in the image, but are artifacts of fitting the noise in the image too closely.

To control noise amplification, the `deconvlucy` function uses a damping parameter, `DAMPAR`. This parameter specifies the threshold level for the deviation of the resulting image from the original image, below which damping occurs. For pixels that deviate in the vicinity of their original values, iterations are suppressed.

Damping is also used to reduce *ringing*, the appearance of high-frequency structures in a restored image. Ringing is not necessarily the result of noise amplification. See “Avoiding Ringing in Deblurred Images” on page 12-23 for more information.

Accounting for Nonuniform Image Quality

Another complication of real-life image restoration is that the data might include bad pixels, or that the quality of the receiving pixels might vary with time and position. By specifying the `WEIGHT` array parameter with the `deconvlucy` function, you can specify that certain pixels in the image be ignored. To ignore a pixel, assign a weight of zero to the element in the `WEIGHT` array that corresponds to the pixel in the image.

The algorithm converges on predicted values for the bad pixels based on the information from neighborhood pixels. The variation in the detector response from pixel to pixel (the so-called flat-field correction) can also be accommodated by the `WEIGHT` array. Instead of assigning a weight of 1.0 to the good pixels, you can specify fractional values and weight the pixels according to the amount of the flat-field correction.

Handling Camera Read-Out Noise

Noise in charge coupled device (CCD) detectors has two primary components:

- Photon counting noise with a Poisson distribution

- Read-out noise with a Gaussian distribution

The Lucy-Richardson iterations intrinsically account for the first type of noise. You must account for the second type of noise; otherwise, it can cause pixels with low levels of incident photons to have negative values.

The `deconvlucy` function uses the `READOUT` input parameter to handle camera read-out noise. The value of this parameter is typically the sum of the read-out noise variance and the background noise (e.g., number of counts from the background radiation). The value of the `READOUT` parameter specifies an offset that ensures that all values are positive.

Handling Undersampled Images

The restoration of undersampled data can be improved significantly if it is done on a finer grid. The `deconvlucy` function uses the `SUBSMPL` parameter to specify the subsampling rate, if the PSF is known to have a higher resolution.


If the undersampled data is the result of camera pixel binning during image acquisition, the PSF observed at each pixel rate can serve as a finer grid PSF. Otherwise, the PSF can be obtained via observations taken at subpixel offsets or via optical modeling techniques. This method is especially effective for images of stars (high signal-to-noise ratio), because the stars are effectively forced to be in the center of a pixel. If a star is centered between pixels, it is restored as a combination of the neighboring pixels. A finer grid redirects the consequent spreading of the star flux back to the center of the star's image.

Example: Using the `deconvlucy` Function to Deblur an Image

To illustrate a simple use of `deconvlucy`, this example simulates a blurred, noisy image by convolving a Gaussian filter PSF with an image (using `imfilter`) and then adding Gaussian noise of variance V to the blurred image (using `imnoise`):

- 1 Read an image into the MATLAB workspace. (The example uses cropping to reduce the size of the image to be deblurred. This is not a required step in deblurring operations.)

```
I = imread('board.tif');  
I = I(50+[1:256],2+[1:256],:);  
figure;imshow(I);title('Original Image');
```



2 Create the PSF.

```
PSF = fspecial('gaussian',5,5);
```


3 Create a simulated blur in the image and add noise.

```
Blurred = imfilter(I,PSF,'symmetric','conv');
```

```
V = .002;
```

```
BlurredNoisy = imnoise(Blurred,'gaussian',0,V);
```


```
figure;imshow(BlurredNoisy);title('Blurred and Noisy Image');
```



- 4 Use `deconvlucy` to restore the blurred and noisy image, specifying the PSF used to create the blur, and limiting the number of iterations to 5 (the default is 10).

Note The `deconvlucy` function can return values in the output image that are beyond the range of the input image.

```
luc1 = deconvlucy(BlurredNoisy,PSF,5);  
figure; imshow(luc1);  
title('Restored Image');
```



Refining the Result

The `deconvlucy` function, by default, performs multiple iterations of the deblurring process. You can stop the processing after a certain number of iterations to check the result, and then restart the iterations from the point where processing stopped. To do this, pass in the input image as a cell array, for example, `{BlurredNoisy}`. The `deconvlucy` function returns the output image as a cell array that you can then pass as an input argument to `deconvlucy` to restart the deconvolution.

The output cell array contains these four elements:

Element	Description
output{1}	Original input image
output{2}	Image produced by the last iteration
output{3}	Image produced by the next to last iteration
output{4}	Internal information used by <code>deconvlucy</code> to know where to restart the process

The `deconvlucy` function supports several other optional arguments you can use to achieve the best possible result, such as specifying a damping parameter to handle additive noise in the blurred image. To see the impact of these optional arguments, view the Image Processing Toolbox deblurring demos.

Deblurring with the Blind Deconvolution Algorithm

Use the `deconvblind` function to deblur an image using the blind deconvolution algorithm. The algorithm maximizes the likelihood that the resulting image, when convolved with the resulting PSF, is an instance of the blurred image, assuming Poisson noise statistics. The blind deconvolution algorithm can be used effectively when no information about the distortion (blurring and noise) is known. The `deconvblind` function restores the image and the PSF simultaneously, using an iterative process similar to the accelerated, damped Lucy-Richardson algorithm.

The `deconvblind` function, just like the `deconvlucy` function, implements several adaptations to the original Lucy-Richardson maximum likelihood algorithm that address complex image restoration tasks. Using these adaptations, you can

- Reduce the effect of noise on the restoration
- Account for nonuniform image quality (e.g., bad pixels)
- Handle camera read-out noise

For more information about these adaptations, see “Deblurring with the Lucy-Richardson Algorithm” on page 12-11. In addition, the `deconvblind`

function supports PSF constraints that can be passed in through a user-specified function.

Example: Using the `deconvblind` Function to Deblur an Image

To illustrate blind deconvolution, this example creates a simulated blurred image and then uses `deconvblind` to deblur it. The example makes two passes at deblurring the image to show the effect of certain optional parameters on the deblurring operation:

- 1 Read an image into the MATLAB workspace.


```
I = imread('cameraman.tif');  
figure; imshow(I); title('Original Image');
```



Image Courtesy of MIT

- 2 Create the PSF.

```
PSF = fspecial('motion',13,45);  
figure; imshow(PSF,[],'notruesize'); title('Original PSF');
```



Original PSF

- 3 Create a simulated blur in the image.

```
Blurred = imfilter(I,PSF,'circ','conv');  
figure; imshow(Blurred); title('Blurred Image');
```



- 4 Deblur the image, making an initial guess at the size of the PSF.

To determine the size of the PSF, examine the blurred image and measure the width of a blur (in pixels) around an obviously sharp object. In the sample blurred image, you can measure the blur near the contour of the man's sleeve. Because the size of the PSF is more important than the values it contains, you can typically specify an array of 1's as the initial PSF.


The following figure shows a restoration where the initial guess at the PSF is the same size as the PSF that caused the blur. In a real application, you

might need to rerun `deconvblind`, experimenting with PSFs of different sizes, until you achieve a satisfactory result. The restored PSF returned by each deconvolution can also provide valuable hints at the optimal PSF size. See the Image Processing Toolbox deblurring demos for an example.

```
INITPSF = ones(size(PSF));
[J P]= deconvblind(Blurred,INITPSF,30);
figure; imshow(J); title('Restored Image');
figure; imshow(P,[],'notruesize');
title('Restored PSF');
```



Restored Image



Restored PSF

Although `deconvblind` was able to deblur the image to a great extent, the ringing around the sharp intensity contrast areas in the restored image is unsatisfactory. (The example eliminated edge-related ringing by using the 'circular' option with `imfilter` when creating the simulated blurred image in step 3.)

The next steps in the example repeat the deblurring process, attempting to achieve a better result by

- Eliminating high-contrast areas from the processing
- Specifying a better PSF

- 5 Create a WEIGHT array to exclude areas of high contrast from the deblurring operation. This can reduce contrast-related ringing in the result.

To exclude a pixel from processing, you create an array of the same size as the original image, and assign the value 0 to the pixels in the array that correspond to pixels in the original image that you want to exclude from processing. (See “Accounting for Nonuniform Image Quality” on page 12-12 for information about WEIGHT arrays.)

To create a WEIGHT array, the example uses a combination of edge detection and morphological processing to detect high-contrast areas in the image. Because the blur in the image is linear, the example dilates the image twice. (For more information about using these functions, see Chapter 9, “Morphological Operations.”) To exclude the image boundary pixels (a high-contrast area) from processing, the example uses padarray to assign the value 0 to all border pixels.

```
WEIGHT = edge(I,'sobel',.28);
se1 = strel('disk',1);
se2 = strel('line',13,45);
WEIGHT = ~imdilate(WEIGHT,[se1 se2]);
WEIGHT = padarray(WEIGHT(2:end-1,2:end-1),[2 2]);
figure; imshow(WEIGHT); title('Weight Array');
```



Weight Array

- 6 Refine the guess at the PSF. The reconstructed PSF P returned by the first pass at deconvolution shows a clear linearity, as shown in the figure in step

4. For the second pass, the example uses a new PSF, P1, which is the same as the restored PSF but with the small amplitude pixels set to 0.


```
P1 = P;  
P1(find(P1 < 0.01))=0;
```

- 7 Rerun the deconvolution, specifying the WEIGHT array and the modified PSF. Note how the restored image has much less ringing around the sharp intensity contrast areas than the result of the first pass (step 4).

```
[J2 P2] = deconvblind(Blurred,P1,50,[],WEIGHT);  
figure; imshow(J2);  
title('Newly Deblurred Image');  
figure; imshow(P2,[],'notruesize');  
title('Newly Reconstructed PSF');
```



Newly Deblurred Image



Newly Reconstructed PSF

Refining the Result

The `deconvblind` function, by default, performs multiple iterations of the deblurring process. You can stop the processing after a certain number of iterations to check the result, and then restart the iterations from the point where processing stopped. To use this feature, you must pass in both the blurred image and the PSF as cell arrays, for example, `{Blurred}` and `{INITPSF}`.

The `deconvblind` function returns the output image and the restored PSF as cell arrays. The output image cell array contains these four elements:

Element	Description
<code>output{1}</code>	Original input image
<code>output{2}</code>	Image produced by the last iteration
<code>output{3}</code>	Image produced by the next to last iteration
<code>output{4}</code>	Internal information used by <code>deconvlucy</code> to know where to restart the process

The PSF output cell array contains similar elements.


The `deconvblind` function supports several other optional arguments you can use to achieve the best possible result, such as specifying a damping parameter to handle additive noise in the blurred image. To see the impact of these optional arguments, as well as the functional option that allows you to place additional constraints on the PSF reconstruction, see the Image Processing Toolbox deblurring demos.

Creating Your Own Deblurring Functions


All the toolbox deblurring functions perform deconvolution in the frequency domain, where the process becomes a simple matrix multiplication. To work in the frequency domain, the deblurring functions must convert the PSF you provide into an optical transfer function (OTF), using the `psf2otf` function. The toolbox also provides a function to convert an OTF into a PSF, `otf2psf`. The toolbox makes these functions available in case you want to create your own deblurring functions. (In addition, to aid this conversion between PSFs and OTFs, the toolbox also makes the padding function `padarray` available.)

Avoiding Ringing in Deblurred Images

The discrete Fourier transform (DFT), used by the deblurring functions, assumes that the frequency pattern of an image is periodic. This assumption creates a high-frequency drop-off at the edges of images. In the figure, the shaded area represents the actual extent of the image; the unshaded area represents the assumed periodicity.



This high-frequency drop-off can create an effect called *boundary related ringing* in deblurred images. In this figure, note the horizontal and vertical patterns in the image.



To avoid ringing, use the `edgetaper` function to preprocess your images before passing them to the deblurring functions. The `edgetaper` function removes the high-frequency drop-off at the edge of an image by blurring the entire image and then replacing the center pixels of the blurred image with the original image. In this way, the edges of the image taper off to a lower frequency.

Color

This chapter describes the toolbox functions that help you work with color image data. Note that “color” includes shades of gray; therefore much of the discussion in this chapter applies to grayscale images as well as color images.

Terminology (p. 13-2)

Provides definitions of image processing terms used in this section

Working with Different Screen Bit Depths (p. 13-3)

Describes how to determine the screen bit depth of your system and provides recommendations if you can change the bit depth

Reducing the Number of Colors in an Image (p. 13-6)

Describes how to use `imapprox` and `rgb2ind` to reduce the number of colors in an image, including information about dithering

Converting Color Data Between Color Spaces (p. 13-15)

Defines the concept of image color space and describes how to convert images between color spaces

Terminology

An understanding of the following terms will help you to use this chapter.

Term	Definition
approximation	Method by which the software chooses replacement colors in the event that direct matches cannot be found. The methods of approximation discussed in this chapter are colormap mapping, uniform quantization, and minimum variance quantization.
indexed image	Image whose pixel values are direct indices into an RGB colormap. In MATLAB, an indexed image is represented by an array of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . The colormap is always an <code>m</code> -by-3 array of class <code>double</code> . This documentation often uses the variable name <code>X</code> to represent an indexed image in memory, and <code>map</code> to represent the colormap.
intensity image	Image consisting of intensity (grayscale) values. In MATLAB, intensity images are represented by an array of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . While intensity images are not stored with colormaps, MATLAB uses a system colormap to display them. This documentation often uses the variable name <code>I</code> to represent an intensity image in memory. This term is synonymous with the term <i>grayscale</i> .
RGB image	Image in which each pixel is specified by three values — one each for the red, blue, and green components of the pixel's color. In MATLAB, an RGB image is represented by an <code>m</code> -by- <code>n</code> -by-3 array of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . This documentation often uses the variable name <code>RGB</code> to represent an RGB image in memory.
screen bit depth	Number of bits per screen pixel.
screen color resolution	Number of distinct colors that can be produced by the screen.

Working with Different Screen Bit Depths

Most computer displays use 8, 16, or 24 bits per screen pixel. The number of bits per screen pixel determines the display's *screen bit depth*. The screen bit depth determines the *screen color resolution*, which is how many distinct colors the display can produce.

Regardless of the number of colors your system can display, MATLAB can store and process images with very high bit depths: 2^{24} colors for `uint8` RGB images, 2^{48} colors for `uint16` RGB images, and 2^{159} for `double` RGB images. These images are displayed best on systems with 24-bit color, but usually look fine on 16-bit systems as well. (For additional information about how MATLAB handles color, see the MATLAB graphics documentation.)

This section

- Describes how to determine your system's screen bit depth
- Provides guidelines for choosing a screen bit depth

Determining Screen Bit Depth

To determine the bit depth of your system's screen, enter this command at the MATLAB prompt.

```
get(0, 'ScreenDepth')  
ans =
```

32

The integer MATLAB returns represents the number of bits per screen pixel:

Value	Screen Bit Depth
8	8-bit displays support 256 colors. An 8-bit display can produce any of the colors available on a 24-bit display, but only 256 distinct colors can appear at one time. (There are 256 shades of gray available, but if all 256 shades of gray are used, they take up all the available color slots.)
16	16-bit displays usually use 5 bits for each color component, resulting in 32 (i.e., 2^5) levels each of red, green, and blue. This supports 32,768 (i.e., 2^{15}) distinct colors (of which 32 are shades of gray). Some systems use the extra bit to increase the number of levels of green that can be displayed. In this case, the number of different colors supported by a 16-bit display is actually 64,536 (i.e. 2^{16}).
24	24-bit displays use 8 bits for each of the three color components, resulting in 256 (i.e., 2^8) levels each of red, green, and blue. This supports 16,777,216 (i.e., 2^{24}) different colors. (Of these colors, 256 are shades of gray. Shades of gray occur where $R=G=B$.) The 16 million possible colors supported by 24-bit display can render a lifelike image.
32	32-bit displays use 24 bits to store color information and use the remaining 8 bits to store transparency data (alpha channel). For information about how MATLAB supports the alpha channel, see the section “Transparency” in the MATLAB graphics documentation.

Choosing a Screen Bit Depth

Depending on your system, you might be able to choose the screen bit depth you want to use. (There might be tradeoffs between screen bit depth and screen color resolution.) In general, 24-bit display mode produces the best results. If you need to use a lower screen bit depth, 16-bit is generally preferable to 8-bit. However, keep in mind that a 16-bit display has certain limitations, such as

- An image might have finer gradations of color than a 16-bit display can represent. If a color is unavailable, MATLAB uses the closest approximation.
- There are only 32 shades of gray available. If you are working primarily with grayscale images, you might get better display results using 8-bit display mode, which provides up to 256 shades of gray.

For information about reducing the number of colors used by an image, see “Reducing the Number of Colors in an Image” on page 13-6.

Reducing the Number of Colors in an Image

This section describes how to reduce the number of colors in an indexed or RGB image. A discussion is also included about dithering, which is used by the toolbox's color-reduction functions (see below). Dithering is used to increase the apparent number of colors in an image.

The table below summarizes the Image Processing Toolbox functions for color reduction.

Function	Purpose
<code>imapprox</code>	Reduces the number of colors used by an indexed image, enabling you to specify the number of colors in the new colormap.
<code>rgb2ind</code>	Converts an RGB image to an indexed image, enabling you to specify the number of colors to store in the new colormap.

On systems with 24-bit color displays, RGB (truecolor) images can display up to 16,777,216 (i.e., 2^{24}) colors. On systems with lower screen bit depths, RGB images are still displayed reasonably well, because MATLAB automatically uses color approximation and dithering if needed.

Indexed images, however, might cause problems if they have a large number of colors. In general, you should limit indexed images to 256 colors for the following reasons:

- On systems with 8-bit display, indexed images with more than 256 colors will need to be dithered or mapped and, therefore, might not display well.
- On some platforms, colormaps cannot exceed 256 entries.
- If an indexed image has more than 256 colors, MATLAB cannot store the image data in a `uint8` array, but generally uses an array of class `double` instead, making the storage size of the image much larger (each pixel uses 64 bits).
- Most image file formats limit indexed images to 256 colors. If you write an indexed image with more than 256 colors (using `imwrite`) to a format that does not support more than 256 colors, you will receive an error.

Using `rgb2ind`

`rgb2ind` converts an RGB image to an indexed image, reducing the number of colors in the process. This function provides the following methods for approximating the colors in the original image:

- Quantization
 - Uniform quantization
 - Minimum variance quantization
- Colormap mapping


The quality of the resulting image depends on the approximation method you use, the range of colors in the input image, and whether or not you use dithering. Note that different methods work better for different images. See “Dithering” on page 13-13 for a description of dithering and how to enable or disable it.

Quantization

Reducing the number of colors in an image involves *quantization*. The function `rgb2ind` uses quantization as part of its color reduction algorithm. `rgb2ind` supports two quantization methods: *uniform quantization* and *minimum variance quantization*.

An important term in discussions of image quantization is *RGB color cube*, which is used frequently throughout this section. The RGB color cube is a three-dimensional array of all of the colors that are defined for a particular data type. Since RGB images in MATLAB can be of type `uint8`, `uint16`, or `double`, three possible color cube definitions exist. For example, if an RGB image is of class `uint8`, 256 values are defined for each color plane (red, blue, and green), and, in total, there will be 2^{24} (or 16,777,216) colors defined by the color cube. This color cube is the same for all `uint8` RGB images, regardless of which colors they actually use.

The `uint8`, `uint16`, and `double` color cubes all have the same range of colors. In other words, the brightest red in a `uint8` RGB image appears the same as the brightest red in a `double` RGB image. The difference is that the `double` RGB color cube has many more shades of red (and many more shades of all colors). The following figure shows an RGB color cube for a `uint8` image.



RGB Color Cube for uint8 Images

Quantization involves dividing the RGB color cube into a number of smaller boxes, and then mapping all colors that fall within each box to the color value at the *center* of that box.

Uniform quantization and minimum variance quantization differ in the approach used to divide up the RGB color cube. With uniform quantization, the color cube is cut up into equal-sized boxes (smaller cubes). With minimum variance quantization, the color cube is cut up into boxes (not necessarily cubes) of different sizes; the sizes of the boxes depend on how the colors are distributed in the image.


Uniform Quantization. To perform uniform quantization, call `rgb2ind` and specify a *tolerance*. The tolerance determines the size of the cube-shaped boxes into which the RGB color cube is divided. The allowable range for a tolerance setting is `[0,1]`. For example, if you specify a tolerance of `0.1`, the edges of the boxes are one-tenth the length of the RGB color cube and the maximum total number of boxes is

$$n = (\text{floor}(1/\text{tol})+1)^3$$

The commands below perform uniform quantization with a tolerance of 0.1.

```
RGB = imread('peppers.png');
[x,map] = rgb2ind(RGB, 0.1);
```

The following figure illustrates uniform quantization of a uint8 image. For clarity, the figure shows a two-dimensional slice (or color plane) from the color cube where red=0 and green and blue range from 0 to 255. The actual pixel values are denoted by the centers of the x's.



Uniform Quantization on a Slice of the RGB Color Cube

After the color cube has been divided, all empty boxes are thrown out. Therefore, only one of the boxes is used to produce a color for the colormap. As shown earlier, the maximum length of a colormap created by uniform quantization can be predicted, but the colormap can be smaller than the prediction because `rgb2ind` removes any colors that do not appear in the input image.

Minimum Variance Quantization. To perform minimum variance quantization, call `rgb2ind` and specify the maximum number of colors in the output image's colormap. The number you specify determines the number of boxes into which

the RGB color cube is divided. These commands use minimum variance quantization to create an indexed image with 185 colors.

```
RGB = imread('peppers.png');  
[X,map] = rgb2ind(RGB,185);
```


Minimum variance quantization works by associating pixels into groups based on the variance between their pixel values. For example, a set of blue pixels might be grouped together because they have a small variance from the center pixel of the group.

In minimum variance quantization, the boxes that divide the color cube vary in size, and do not necessarily fill the color cube. If some areas of the color cube do not have pixels, there are no boxes in these areas.

While you set the number of boxes, n , to be used by `rgb2ind`, the placement is determined by the algorithm as it analyzes the color data in your image. Once the image is divided into n optimally located boxes, the pixels within each box are mapped to the pixel value at the center of the box, as in uniform quantization.

The resulting colormap usually has the number of entries you specify. This is because the color cube is divided so that each region contains at least one color that appears in the input image. If the input image uses fewer colors than the number you specify, the output colormap will have fewer than n colors, and the output image will contain all of the colors of the input image.

The following figure shows the same two-dimensional slice of the color cube as shown in the preceding figure (demonstrating uniform quantization). Eleven boxes have been created using minimum variance quantization.



Minimum Variance Quantization on a Slice of the RGB Color Cube

For a given number of colors, minimum variance quantization produces better results than uniform quantization, because it takes into account the actual data. Minimum variance quantization allocates more of the colormap entries to colors that appear frequently in the input image. It allocates fewer entries to colors that appear infrequently. As a result, the accuracy of the colors is higher than with uniform quantization. For example, if the input image has many shades of green and few shades of red, there will be more greens than reds in the output colormap. Note that the computation for minimum variance quantization takes longer than that for uniform quantization.

Colormap Mapping

If you specify an actual colormap to use, `rgb2ind` uses *colormap mapping* (instead of quantization) to find the colors in the specified colormap that best match the colors in the RGB image. This method is useful if you need to create images that use a fixed colormap. For example, if you want to display multiple indexed images on an 8-bit display, you can avoid color problems by mapping them all to the same colormap. Colormap mapping produces a good approximation if the specified colormap has similar colors to those in the RGB image. If the colormap does not have similar colors to those in the RGB image, this method produces poor results.

This example illustrates mapping two images to the same colormap. The colormap used for the two images is created on the fly using the MATLAB function `colorcube`, which creates an RGB colormap containing the number of colors that you specify. (`colorcube` always creates the same colormap for a given number of colors.) Because the colormap includes colors all throughout the RGB color cube, the output images can reasonably approximate the input images.

```
RGB1 = imread('autumn.tif');
RGB2 = imread('peppers.png');
X1 = rgb2ind(RGB1,colorcube(128));
X2 = rgb2ind(RGB2,colorcube(128));
```

Note The function `subimage` is also helpful for displaying multiple indexed images. For more information, see “Displaying Multiple Images in the Same Figure” on page 3-42 or the reference page for `subimage`.

Reducing Colors in an Indexed Image

Use `imapprox` when you need to reduce the number of colors in an indexed image. `imapprox` is based on `rgb2ind` and uses the same approximation methods. Essentially, `imapprox` first calls `ind2rgb` to convert the image to RGB format, and then calls `rgb2ind` to return a new indexed image with fewer colors.

For example, these commands create a version of the trees image with 64 colors, rather than the original 128.

```
load trees
[Y,newmap] = imapprox(X,map,64);
imshow(Y, newmap);
```

The quality of the resulting image depends on the approximation method you use, the range of colors in the input image, and whether or not you use dithering. Note that different methods work better for different images. See “Dithering” on page 13-13 for a description of dithering and how to enable or disable it.

Dithering

When you use `rgb2ind` or `imapprox` to reduce the number of colors in an image, the resulting image might look inferior to the original, because some of the colors are lost. `rgb2ind` and `imapprox` both perform *dithering* to increase the apparent number of colors in the output image. Dithering changes the colors of pixels in a neighborhood so that the average color in each neighborhood approximates the original RGB color.

For an example of how dithering works, consider an image that contains a number of dark orange pixels for which there is no exact match in the colormap. To create the appearance of this shade of orange, the Image Processing Toolbox selects a combination of colors from the colormap, that, taken together as a six-pixel group, approximate the desired shade of pink. From a distance, the pixels appear to be the correct shade, but if you look up close at the image, you can see a blend of other shades. This example loads a 24-bit image, and then use `rgb2ind` to create two indexed images with just eight colors each:

- 1 Read image and display it.

```
rgb=imread('onion.png');  
imshow(rgb);
```



- 2 Create an indexed image with eight colors and without dithering.

```
[X_no_dither,map]=rgb2ind(rgb,8,'nodither');  
figure, imshow(X_no_dither,map);
```



- 3** Create an indexed image using eight colors with dithering.

```
[X_dither,map]=rgb2ind(rgb,8,'dither');  
figure, imshow(X_dither,map);
```



Notice that the dithered image has a larger number of apparent colors but is somewhat fuzzy-looking. The image produced without dithering has fewer apparent colors, but an improved spatial resolution when compared to the dithered image. One risk in doing color reduction without dithering is that the new image can contain false contours.

Converting Color Data Between Color Spaces

The Image Processing Toolbox represents colors as RGB values, either directly (in an RGB image) or indirectly (in an indexed image, where the colormap is stored in RGB format). However, there are other models besides RGB for representing colors numerically. The various models are referred to as *color spaces* because most of them can be mapped into a 2-D, 3-D, or 4-D coordinate system; thus, a color specification is made up of coordinates in a 2-D, 3-D, or 4-D space.

The various color spaces exist because they present color information in ways that make certain calculations more convenient or because they provide a way to identify colors that is more intuitive. For example, the RGB color space defines a color as the percentages of red, green, and blue hues mixed together. Other color models describe colors by their hue (green), saturation (dark green), and luminance, or intensity.

The toolbox supports these color spaces by providing a means for converting color data from one color space to another through a mathematical transformation.

This section

- Describes how to convert color data between these color spaces
- Describes how to perform color space conversions using ICC profiles
- Describes some toolbox functions for converting between the RGB color space and three commonly used color spaces: YIQ, HSV, and YCbCr

Converting Between Device-Independent Color Spaces

The standard terms used to describe colors, such as hue, brightness, and intensity, are subjective and make comparisons difficult.

In 1931, the International Commission on Illumination, known by the acronym CIE, for *Commission Internationale de l'Éclairage*, studied human color perception and developed a standard, called the CIE XYZ. This standard defined a three-dimensional space where three values, called tristimulus values, define a color. This standard is still widely used today.

In the decades since that initial specification, the CIE has developed several additional color space specifications that attempt to provide alternative color

representations that are better suited to some purposes than XYZ. For example, in 1976, in an effort to get a perceptually uniform color space that could be correlated with the visual appearance of colors, the CIE created the $L^*a^*b^*$ color space.

The toolbox supports conversions between members of the CIE family of device-independent color spaces. In addition, the toolbox also supports conversions between these CIE color spaces and the sRGB color space. This color space was defined by an industry group to describe the characteristics of a typical PC monitor.

This section

- Lists the supported device-independent color spaces
- Provides an example of how to perform a conversion
- Provides guidelines about data type support of the various conversions

Supported Conversions

This table lists all the device-independent color spaces that the toolbox supports.

Color Space	Description	Supported Conversions
XYZ	The original, 1931 CIE color space specification.	xyY , uvl , $u'v'L$, and $L^*a^*b^*$
xyY	CIE specification that provides normalized chromaticity values. The capital Y value represents luminance and is the same as in XYZ.	XYZ
uvL	CIE specification that attempts to make the chromaticity plane more visually uniform. l is luminance and is the same as Y in XYZ.	XYZ
$u'v'L$	CIE specification in which u and v are rescaled to improve uniformity.	XYZ

Color Space	Description	Supported Conversions
$L^*a^*b^*$	CIE specification that attempts to make the luminance scale more perceptually uniform. L^* is a nonlinear scaling of L , normalized to a reference white point.	XYZ
L^*ch	CIE specification where c is chroma and h is hue. These values are a polar coordinate conversion of a^* and b^* in $L^*a^*b^*$.	$L^*a^*b^*$
$sRGB$	Standard adopted by major manufacturers that characterizes the average PC monitor.	XYZ and $L^*a^*b^*$

Example: Performing a Color Space Conversion

To illustrate a conversion between two device-independent color spaces, this example reads an RGB color image into the MATLAB workspace and converts the color data to the XYZ color space:

- 1 Import color space data. This example reads an RGB color image into the MATLAB workspace.


```
I_rgb = imread('peppers.png');
```
- 2 Create a color transformation structure. A color transformation structure defines the conversion between two color spaces. You use the `makecform` function to create the structure, specifying a transformation type string as an argument.

This example creates a color transformation structure that defines a conversion from RGB color data to XYZ color data.

```
C = makecform('srgb2xyz');
```
- 3 Perform the conversion. You use the `applycform` function to perform the conversion, specifying as arguments the color data you want to convert and

the color transformation structure that defines the conversion. The `applycform` function returns the converted data.

```
I_xyz = applycform(I_rgb,C);
```

```
whos
  Name          Size          Bytes  Class
  C              1x1            7744   struct array
  I_xyz          384x512x3     1179648 uint16 array
  I_rgb          384x512x3     589824  uint8 array
```

Color Space Data Encodings

When you convert between two device-independent color spaces, the data type used to encode the color data can sometimes change, depending on what encodings the color spaces support. In the preceding example, the original image is `uint8` data. The XYZ conversion is `uint16` data. The XYZ color space does not define a `uint8` encoding. The following table lists the data types that can be used to represent values in all the device-independent color spaces.

Color Space	Encodings
XYZ	uint16 or double
xyY	double
uvL	double
u'v'L	double
L*a*b*	uint8, uint16, or double
L*ch	double
sRGB	double

As the table indicates, certain color spaces have data type limitations. For example, the XYZ color space does not define a `uint8` encoding. If you convert 8-bit CIE LAB data into the XYZ color space, the data is returned in `uint16` format. If you want the returned XYZ data to be in the same format as the input LAB data, you can use one of the following toolbox color space format conversion functions.

- `lab2double`
- `lab2uint8`
- `lab2uint16`
- `xyz2double`
- `xyz2uint16`

Performing Profile-Based Conversions

The Image Processing Toolbox can perform color space conversions based on device profiles. This section includes the following topics:

- “Understanding Device Profiles” on page 13-19
- “Reading ICC Profiles” on page 13-20
- “Writing Profile Information to a File” on page 13-20
- “Example: Performing a Profile-Based Conversion” on page 13-21
- “Specifying the Rendering Intent” on page 13-22

Understanding Device Profiles

If two colors have the same CIE colorimetry, they will match *if viewed under the same conditions*. However, because color images are typically produced for a wide variety of viewing environments, it is necessary to go beyond simple application of the CIE system.

For this reason, the International Color Consortium (ICC) has defined a Color Management System (CMS) that provides a means for communicating color information among input, output, and display devices. The CMS uses device *profiles* that contain color information specific to a particular device. Vendors that support CMS provide profiles that characterize the color reproduction of their devices, and methods, called Color Management Modules (CMM), that interpret the contents of each profile and perform the necessary image processing.

Device profiles contain the information that color management systems need to translate color data between devices. Any conversion between color spaces is a mathematical transformation from some domain space to a range space. With profile-based conversions, the domain space is often called the *source space* and the range space is called the *destination space*. In the ICC color management model, profiles are used to represent the source and destination spaces.

For more information about color management systems, go to the International Color Consortium Web site, www.color.org.

Reading ICC Profiles

To read an ICC profile into the MATLAB workspace, use the `iccread` function. In this example, the function reads in the profile for the color space that describes color monitors.

```
P = iccread('sRGB.icm');
```

`iccread` returns the contents of the profile in the structure `P`. All profiles contain a header, a tag table, and a series of tagged elements. The header contains general information about the profile, such as the device class, the device color space, and the file size. The tagged elements, or tags, are the data constructs that contain the information used by the CMM. For more information about the contents of a profile, see the `iccread` function reference page.

Using `iccread`, you can read both Version 2 (ICC.1:2001-04) or Version 4 (ICC.1:2001-12) ICC profile formats. For detailed information about these specifications and their differences, visit the ICC web site, www.color.org.

Writing Profile Information to a File

To export ICC profile information from the MATLAB workspace to a file, use the `iccwrite` function. This example reads a profile into the MATLAB workspace and then writes the profile information out to a new file.

```
P = iccread('sRGB.icm');  
P_new = iccwrite(P, 'my_profile.icm');
```

`iccwrite` returns the profile it writes to the file in `P_new` because it can be different than the input profile `P`. For example, `iccwrite` updates the `Filename` field in `P` to match the name of the file specified as the second argument.

`iccwrite` checks the validity of the input profile structure. If any required fields are missing, `iccwrite` errors. To determine if a structure is a valid ICC profile, use the `isicc` function.

Using `iccwrite`, you can export profile information in both Version 2 (ICC.1:2001-04) or Version 4 (ICC.1:2001-12) ICC profile formats. The value of the `Version` field in the file profile header determines the format version. For

detailed information about these specifications and their differences, visit the ICC web site, www.color.org.

Example: Performing a Profile-Based Conversion

To illustrate a profile-based color space conversion, this section presents an example that converts color data from the RGB space of a monitor to the CMYK space of a printer. This conversion requires two profiles: a monitor profile and a printer profile. The source color space in this example is monitor RGB and the destination color space is printer CMYK:

- 1** Import RGB color space data. This example imports an RGB color image into the MATLAB workspace.

```
I_rgb = imread('peppers.png');
```

- 2** Read ICC profiles. Read the source and destination profiles into the MATLAB workspace. This example uses the sRGB profile as the source profile. The sRGB profile is an industry-standard color space that describes a color monitor.

```
inprof = iccread('sRGB.icm');
```

For the destination profile, the example uses a profile that describes a particular color printer. The printer vendor supplies this profile. (The following profile and several other useful profiles can be obtained as downloads from www.adobe.com.)

```
outprof = iccread('USSheetfedCoated.icc');
```

- 3** Create a color transformation structure. You must create a color transformation structure to define the conversion between the color spaces in the profiles. You use the `makecform` function to create the structure, specifying a transformation type string as an argument.

Note The color space conversion might involve an intermediate conversion into a device-independent color space, called the Profile Connection Space (PCS), but this is transparent to the user.

This example creates a color transformation structure that defines a conversion from RGB color data to CMYK color data.

```
C = makecform('icc',inprof,outprof);
```

- 4 Perform the conversion. You use the `applycform` function to perform the conversion, specifying as arguments the color data you want to convert and the color transformation structure that defines the conversion. The function returns the converted data.

```
I_cmyk = applycform(I_rgb,C);
```

- 5 Write the converted data to a file. To export the CMYK data, use the `imwrite` function, specifying the format as TIFF. If the format is TIFF and the data is an m-by-n-by-4 array, `imwrite` writes CMYK data to the file.

```
imwrite(I_cmyk,'pep_cmyk.tif','tif')
```

To verify that the CMYK data was written to the file, use `imfinfo` to get information about the file and look at the `PhotometricInterpretation` field.

```
info = imfinfo('pep_cmyk.tif');  
info.PhotometricInterpretation  
ans =  
    'CMYK'
```

Specifying the Rendering Intent

For most devices, the range of reproducible colors is much smaller than the range of colors represented by the PCS. It is for this reason that four rendering intents (or gamut mapping techniques) are defined in the profile format. Each one has distinct aesthetic and color-accuracy tradeoffs.

When you create a profile-based color transformation structure, you can specify the rendering intent for the source as well as the destination profiles. For more information, see the `makecform` reference information.

Converting Between Device-Dependent Color Spaces

The toolbox includes functions that you can use to convert RGB data to several common device-dependent color spaces, and vice versa:

- YIQ
- YCbCr
- Hue, saturation, value (HSV)

YIQ Color Space

The National Television Systems Committee (NTSC) defines a color space known as YIQ. This color space is used in televisions in the United States. One of the main advantages of this format is that grayscale information is separated from color data, so the same signal can be used for both color and black and white sets.

In the NTSC color space, image data consists of three components: luminance (Y), hue (I), and saturation (Q). The first component, luminance, represents grayscale information, while the last two components make up chrominance (color information).

The function `rgb2ntsc` converts `colormaps` or RGB images to the NTSC color space. `ntsc2rgb` performs the reverse operation.

For example, these commands convert an RGB image to NTSC format.

```
RGB = imread('peppers.png');  
YIQ = rgb2ntsc(RGB);
```

Because luminance is one of the components of the NTSC format, the RGB to NTSC conversion is also useful for isolating the gray level information in an image. In fact, the toolbox functions `rgb2gray` and `ind2gray` use the `rgb2ntsc` function to extract the grayscale information from a color image.

For example, these commands are equivalent to calling `rgb2gray`.

```
YIQ = rgb2ntsc(RGB);
```

```
I = YIQ(:, :, 1);
```

Note In the YIQ color space, I is one of the two color components, not the grayscale component.

YCbCr Color Space

The YCbCr color space is widely used for digital video. In this format, luminance information is stored as a single component (Y), and chrominance information is stored as two color-difference components (Cb and Cr). Cb represents the difference between the blue component and a reference value. Cr represents the difference between the red component and a reference value.

YCbCr data can be double precision, but the color space is particularly well suited to uint8 data. For uint8 images, the data range for Y is [16, 235], and the range for Cb and Cr is [16, 240]. YCbCr leaves room at the top and bottom of the full uint8 range so that additional (nonimage) information can be included in a video stream.

The function `rgb2ycbcr` converts colormaps or RGB images to the YCbCr color space. `ycbcr2rgb` performs the reverse operation.

For example, these commands convert an RGB image to YCbCr format.

```
RGB = imread('peppers.png');  
YCBCR = rgb2ycbcr(RGB);
```

HSV Color Space

The HSV color space (hue, saturation, value) is often used by people who are selecting colors (e.g., of paints or inks) from a color wheel or palette, because it corresponds better to how people experience color than the RGB color space does. The functions `rgb2hsv` and `hsv2rgb` convert images between the RGB and HSV color spaces.

As hue varies from 0 to 1.0, the corresponding colors vary from red through yellow, green, cyan, blue, magenta, and back to red, so that there are actually red values both at 0 and 1.0. As saturation varies from 0 to 1.0, the corresponding colors (hues) vary from unsaturated (shades of gray) to fully saturated (no white component). As value, or brightness, varies from 0 to 1.0, the corresponding colors become increasingly brighter.

The following figure illustrates the HSV color space.




Illustration of the HSV Color Space


The `rgb2hsv` function converts colormaps or RGB images to the HSV color space. `hsv2rgb` performs the reverse operation. These commands convert an RGB image to the HSV color space.

```
RGB = imread('peppers.png');
HSV = rgb2hsv(RGB);
```

For closer inspection of the HSV color space, the next block of code displays the separate color planes (hue, saturation, and value) of an HSV image.

```
RGB=reshape(ones(64,1)*reshape(jet(64),1,192),[64,64,3]);
HSV=rgb2hsv(RGB);
H=HSV(:,:,1);
S=HSV(:,:,2);
V=HSV(:,:,3);
imshow(H)
figure, imshow(S);
```

```
figure, imshow(V);  
figure, imshow(RGB);
```



The Separated Color Planes of an HSV Image

As the hue plane image in the preceding figure illustrates, hue values make a linear transition from high to low. If you compare the hue plane image against the original image, you can see that shades of deep blue have the highest values, and shades of deep red have the lowest values. (As stated previously, there are values of red on both ends of the hue scale. To avoid confusion, the sample image uses only the red values from the *beginning* of the hue range.)

Saturation can be thought of as the purity of a color. As the saturation plane image shows, the colors with the highest saturation have the highest values and are represented as white. In the center of the saturation image, notice the various shades of gray. These correspond to a mixture of colors; the cyans, greens, and yellow shades are mixtures of true colors. Value is roughly

equivalent to brightness, and you will notice that the brightest areas of the value plane correspond to the brightest colors in the original image.

Neighborhood and Block Operations

This chapter discusses these generic block processing functions. Topics covered include

Terminology (p. 14-2)	Provides definitions of image processing terms used in this section
Block Processing Operations (p. 14-3)	Provides an overview of the types of block processing operations supported by the toolbox
Sliding Neighborhood Operations (p. 14-4)	Defines sliding neighborhood operations and describes how you can use them to implement many types of filtering operations
Distinct Block Operations (p. 14-8)	Describes block operations
Column Processing (p. 14-11)	Describes how to process sliding neighborhoods or distinct blocks as columns

Terminology

An understanding of the following terms will help you to use this section.

Term	Definition
block operation	Operation in which an image is processed in blocks rather than all at once. The blocks have the same size across the image. Some operation is applied to one block at a time. The blocks are reassembled to form an output image.
border padding	Additional rows and columns temporarily added to the border(s) of an image when some of the blocks extend outside the image. The additional rows and columns normally contain zeros.
center pixel	Pixel at the center of a neighborhood.
column processing	Operation in which neighborhoods are reshaped into columns before processing in order to speed up computation time.
distinct block operation	Block operation in which the blocks do not overlap.
neighborhood operation	Operation in which each output pixel is computed from a set of neighboring input pixels. Convolution, dilation, and median filtering are examples of neighborhood operations. A neighborhood operation can also be called a sliding neighborhood operation.
overlap	Extra rows and columns of pixels outside a block whose values are taken into account when processing the block. These extra pixels cause distinct blocks to overlap one another. The <code>blkproc</code> function enables you to specify an overlap.

Block Processing Operations

Certain image processing operations involve processing an image in sections called *blocks*, rather than processing the entire image at once. The Image Processing Toolbox provides several functions for specific operations that work with blocks, for example, the `imdilate` function for image dilation. In addition, the toolbox provides more generic functions for processing an image in blocks. This section discusses these generic block processing functions.

To use one of the functions, you supply information about the size of the blocks, and specify a separate function to use to process the blocks. The block processing function does the work of breaking the input image into blocks, calling the specified function for each block, and reassembling the results into an output image.

Types of Block Processing Operations

Using these functions, you can perform various block processing operations, including *sliding neighborhood operations* and *distinct block operations*:

- In a sliding neighborhood operation, the input image is processed in a pixelwise fashion. That is, for each pixel in the input image, some operation is performed to determine the value of the corresponding pixel in the output image. The operation is based on the values of a block of neighboring pixels.
- In a distinct block operation, the input image is processed a block at a time. That is, the image is divided into rectangular blocks, and some operation is performed on each block individually to determine the values of the pixels in the corresponding block of the output image.


In addition, the toolbox provides functions for *column processing operations*. These operations are not actually distinct from block operations; instead, they are a way of speeding up block operations by rearranging blocks into matrix columns.

Note that even if you do not use these block processing functions, the information here might be useful to you, as it includes concepts fundamental to many areas of image processing. In particular, the discussion of sliding neighborhood operations is applicable to linear filtering and morphological operations. See Chapter 7, “Linear Filtering and Filter Design,” and Chapter 9, “Morphological Operations,” for information about these applications.

Sliding Neighborhood Operations

A sliding neighborhood operation is an operation that is performed a pixel at a time, with the value of any given pixel in the output image being determined by the application of an algorithm to the values of the corresponding input pixel's *neighborhood*. A pixel's neighborhood is some set of pixels, defined by their locations relative to that pixel, which is called the *center pixel*. The neighborhood is a rectangular block, and as you move from one element to the next in an image matrix, the neighborhood block slides in the same direction.

The following figure shows the neighborhood blocks for some of the elements in a 6-by-5 matrix with 2-by-3 sliding blocks. The center pixel for each neighborhood is marked with a dot.



Neighborhood Blocks in a 6-by-5 Matrix

The center pixel is the actual pixel in the input image being processed by the operation. If the neighborhood has an odd number of rows and columns, the center pixel is actually in the center of the neighborhood. If one of the dimensions has even length, the center pixel is just to the left of center or just above center. For example, in a 2-by-2 neighborhood, the center pixel is the upper left one.

For any m -by- n neighborhood, the center pixel is

$$\text{floor}((m + 1)/2)$$

In the 2-by-3 block shown in Figure , the center pixel is (1,2), or the pixel in the second column of the top row of the neighborhood.

To perform a sliding neighborhood operation,

- 1 Select a single pixel.
- 2 Determine the pixel's neighborhood.
- 3 Apply a function to the values of the pixels in the neighborhood. This function must return a scalar.
- 4 Find the pixel in the output image whose position corresponds to that of the center pixel in the input image. Set this output pixel to the value returned by the function.
- 5 Repeat steps 1 through 4 for each pixel in the input image.

For example, the function might be an averaging operation that sums the values of the neighborhood pixels and then divides the result by the number of pixels in the neighborhood. The result of this calculation is the value of the output pixel.

Padding Borders

As the neighborhood block slides over the image, some of the pixels in a neighborhood might be missing, especially if the center pixel is on the border of the image. For example, if the center pixel is the pixel in the upper left corner of the image, the neighborhoods include pixels that are not part of the image.

To process these neighborhoods, sliding neighborhood operations *pad* the borders of the image, usually with 0's. In other words, these functions process the border pixels by assuming that the image is surrounded by additional rows and columns of 0's. These rows and columns do not become part of the output image and are used only as parts of the neighborhoods of the actual pixels in the image.

Linear and Nonlinear Filtering

You can use sliding neighborhood operations to implement many kinds of filtering operations. One example of a sliding neighbor operation is convolution, which is used to implement linear filtering. MATLAB provides the `conv` and `filter2` functions for performing convolution, and the toolbox

provides the `imfilter` function. See Chapter 7, “Linear Filtering and Filter Design,” for more information about these functions.

In addition to convolution, there are many other filtering operations you can implement through sliding neighborhoods. Many of these operations are nonlinear in nature. For example, you can implement a sliding neighborhood operation where the value of an output pixel is equal to the standard deviation of the values of the pixels in the input pixel’s neighborhood.

You can use the `nlfilter` function to implement a variety of sliding neighborhood operations. `nlfilter` takes as input arguments an image, a neighborhood size, and a function that returns a scalar, and returns an image of the same size as the input image. The value of each pixel in the output image is computed by passing the corresponding input pixel’s neighborhood to the function. For example, this call computes each output pixel by taking the standard deviation of the values of the input pixel’s 3-by-3 neighborhood (that is, the pixel itself and its eight contiguous neighbors).

```
I2 = nlfilter(I,[3 3],'std2');
```

You can write an M-file to implement a specific function, and then use this function with `nlfilter`. For example, this command processes the matrix `I` in 2-by-3 neighborhoods with a function called `myfun.m`.


```
nlfilter(I,[2 3],@myfun);
```

`@myfun` is an example of a function handle. You can also use an inline function. For example:

```
f = inline('sqrt(min(x(:)))');  
I2 = nlfilter(I,[2 2],f);
```

The example below uses `nlfilter` to set each pixel to the maximum value in its 3-by-3 neighborhood.

```
I = imread('tire.tif');  
f = inline('max(x(:))');  
I2 = nlfilter(I,[3 3],f);  
imshow(I);  
figure, imshow(I2);
```



Each Output Pixel Set to Maximum Input Neighborhood Value

Many operations that `nlfilter` can implement run much faster if the computations are performed on matrix columns rather than rectangular neighborhoods. For information about this approach, see the reference page for `colfilt`.

Note `nlfilter` is an example of a “function function.” For more information on how to use this kind of function, see “Function Functions” in the MATLAB documentation. For more information on inline functions, see `inline` in the MATLAB Function Reference documentation. For more information on function handles, see `function_handle` in the MATLAB Function Reference documentation.

Distinct Block Operations

Distinct blocks are rectangular partitions that divide a matrix into m -by- n sections. Distinct blocks overlay the image matrix starting in the upper left corner, with no overlap. If the blocks don't fit exactly over the image, the toolbox adds zero padding so that they do. The following figure shows a 15-by-30 matrix divided into 4-by-8 blocks.




Image Divided into Distinct Blocks

The zero padding process adds 0's to the bottom and right of the image matrix, as needed. After zero padding, the matrix is size 16-by-32.

The function `blkproc` performs distinct block operations. `blkproc` extracts each distinct block from an image and passes it to a function you specify. `blkproc` assembles the returned blocks to create an output image.

For example, the command below processes the matrix `I` in 4-by-6 blocks with the function `myfun`.


```
I2 = blkproc(I,[4 6],@myfun);
```

You can specify the function as an inline function. For example:

```
f = inline('mean2(x)*ones(size(x))');  
I2 = blkproc(I,[4 6],f);
```

The example below uses `blkproc` to set every pixel in each 8-by-8 block of an image matrix to the average of the elements in that block.

```
I = imread('tire.tif');  
f = inline('uint8(round(mean2(x)*ones(size(x))))');  
I2 = blkproc(I,[8 8],f);  
imshow(I)  
figure, imshow(I2);
```



Notice that `inline` computes the mean of the block and then multiplies the result by a matrix of ones, so that the output block is the same size as the input block. As a result, the output image is the same size as the input image. `blkproc` does not require that the images be the same size; however, if this is the result you want, you must make sure that the function you specify returns blocks of the appropriate size.

Note `blkproc` is an example of a “function function.” For more information on how to use this kind of function, see the “Function Functions” section in the MATLAB documentation.

Overlap

When you call `blkproc` to define distinct blocks, you can specify that the blocks overlap each other, that is, you can specify extra rows and columns of pixels outside the block whose values are taken into account when processing the block. When there is an overlap, `blkproc` passes the expanded block (including the overlap) to the specified function.

The following figure shows the overlap areas for some of the blocks in a 15-by-30 matrix with 1-by-2 overlaps. Each 4-by-8 block has a one-row overlap above and below, and a two-column overlap on each side. In the figure, shading indicates the overlap. The 4-by-8 blocks overlay the image matrix starting in the upper left corner.




Image Divided into Distinct Blocks with Specified Overlaps

To specify the overlap, you provide an additional input argument to `blkproc`. To process the blocks in the figure above with the function `myfun`, the call is

```
B = blkproc(A,[4 8],[1 2],@myfun)
```

Overlap often increases the amount of zero padding needed. For example, in the figure, the original 15-by-30 matrix became a 16-by-32 matrix with zero padding. When the 15-by-30 matrix includes a 1-by-2 overlap, the padded matrix becomes an 18-by-36 matrix. The outermost rectangle in the figure delineates the new boundaries of the image after padding has been added to accommodate the overlap plus block processing. Notice that in the preceding figure, padding has been added to the left and top of the original image, not just to the right and bottom.

Column Processing

The toolbox provides functions that you can use to process sliding neighborhoods or distinct blocks as columns. This approach is useful for operations that MATLAB performs columnwise; in many cases, column processing can reduce the execution time required to process an image.

For example, suppose the operation you are performing involves computing the mean of each block. This computation is much faster if you first rearrange the blocks into columns, because you can compute the mean of every column with a single call to the mean function, rather than calling mean for each block individually.


You can use the `colfilt` function to implement column processing. This function

- 1 Reshapes each sliding or distinct block of an image matrix into a column in a temporary matrix
- 2 Passes the temporary matrix to a function you specify
- 3 Rearranges the resulting matrix back into the original shape

Sliding Neighborhoods

For a sliding neighborhood operation, `colfilt` creates a temporary matrix that has a separate column for each pixel in the original image. The column corresponding to a given pixel contains the values of that pixel's neighborhood from the original image.

The following figure illustrates this process. In this figure, a 6-by-5 image matrix is processed in 2-by-3 neighborhoods. `colfilt` creates one column for each pixel in the image, so there are a total of 30 columns in the temporary matrix. Each pixel's column contains the value of the pixels in its neighborhood, so there are six rows. `colfilt` zero-pads the input image as necessary. For example, the neighborhood of the upper left pixel in the figure has two zero-valued neighbors, due to zero padding.



colfilt Creates a Temporary Matrix for Sliding Neighborhood

The temporary matrix is passed to a function, which must return a single value for each column. (Many MATLAB functions work this way, for example, mean, median, std, sum, etc.) The resulting values are then assigned to the appropriate pixels in the output image.


colfilt can produce the same results as nlfilt with faster execution time; however, it might use more memory. The example below sets each output pixel to the maximum value in the input pixel's neighborhood, producing the same result as the nlfilt example shown in "Linear and Nonlinear Filtering" on page 14-5.

```
I2 = colfilt(I,[3 3],'sliding',@max);
```

Distinct Blocks

For a distinct block operation, colfilt creates a temporary matrix by rearranging each block in the image into a column. colfilt pads the original image with 0's, if necessary, before creating the temporary matrix.

The following figure illustrates this process. A 6-by-16 image matrix is processed in 4-by-6 blocks. colfilt first zero-pads the image to make the size 8-by-18 (six 4-by-6 blocks), and then rearranges the blocks into six columns of 24 elements each.



colfilt Creates a Temporary Matrix for Distinct Block Operation

After rearranging the image into a temporary matrix, `colfilt` passes this matrix to the function. The function must return a matrix of the same size as the temporary matrix. If the block size is m -by- n , and the image is mm -by- nn , the size of the temporary matrix is $(m*n)$ -by- $(\text{ceil}(mm/m)*\text{ceil}(nn/n))$. After the function processes the temporary matrix, the output is rearranged into the shape of the original image matrix.

This example sets all the pixels in each 8-by-8 block of an image to the mean pixel value for the block, producing the same result as the `blkproc` example in “Distinct Block Operations” on page 14-8.

```
I = im2double(imread('tire.tif'));  
f = inline('ones(64,1)*mean(x)');  
I2 = colfilt(I,[8 8],'distinct',f);
```

Notice that the inline function computes the mean of the block and then multiplies the result by a vector of ones, so that the output block is the same size as the input block. As a result, the output image is the same size as the input image.

Restrictions

You can use `colfilt` to implement many of the same distinct block operations that `blkproc` performs. However, `colfilt` has certain restrictions that `blkproc` does not:

- The output image must be the same size as the input image.
- The blocks cannot overlap.

For situations that do not satisfy these constraints, use `blkproc`.

Function Reference

This section describes the Image Processing Toolbox functions.

- Functions – Categorical List (p. 15-2) Contains a group of tables that organize the toolbox functions into category groups
- Functions – Alphabetical List (p. 15-20) Contains separate reference pages for each toolbox function

Functions – Categorical List

This section provides brief descriptions of all the functions in the Image Processing Toolbox. The functions are listed in tables in the following broad categories.

If you know the name of a function, use “Functions — Alphabetical List” to find the reference page.

Image Input, Output, and Display (p. 15-3)	Functions for importing, exporting, and displaying images and converting between image formats
Modular Interactive Tools (p. 15-5)	Functions for creating the modular interactive tools, with associated utility functions.
Spatial Transformation and Registration (p. 15-7)	Functions for performing spatial transformations and image registration
Image Analysis and Statistics (p. 15-8)	Functions for performing image analysis and getting pixel values and statistics
Image Arithmetic (p. 15-9)	Functions for performing arithmetic operations on images, such as adding, subtracting, multiplying, and dividing.
Image Enhancement and Restoration (p. 15-10)	Functions for image enhancement and restoration, such as deblurring
Linear Filtering and Transforms (p. 15-11)	Functions for creating and using linear filters and transforms
Morphological Operations (p. 15-13)	Functions for performing morphological image processing
Region-Based, Neighborhood, and Block Processing (p. 15-15)	Functions to define regions of interest and operate on these regions

Colormap and Color Space Functions (p. 15-16)	Functions for working with image color
Miscellaneous Functions (p. 15-18)	Functions that perform array operations, and set and get Image Processing Toolbox preferences

Image Input, Output, and Display

- Image Display (p. 15-3)
- Image File I/O (p. 15-3)
- Image Types and Type Conversions (p. 15-4)

Image Display

colorbar	Display color bar (MATLAB function)
image	Create and display image object (MATLAB function)
imagesc	Scale data and display as image (MATLAB function)
immovie	Make movie from multiframe indexed image
imshow	Display image in a MATLAB figure window
imtool	Display image in the Image Viewer
montage	Display multiple image frames as rectangular montage
subimage	Display multiple images in single figure
warp	Display image as texture-mapped surface

Image File I/O

dicomanon	Anonymize a DICOM file
dicomdict	Specify which DICOM data dictionary to use
dicominfo	Read metadata from a DICOM message
dicomread	Read a DICOM image
dicomuid	Generate DICOM unique identifier

<code>dicomwrite</code>	Write a DICOM image
<code>dicom-dict.txt</code>	Text file containing DICOM data dictionary
<code>imfinfo</code>	Return information about image file (MATLAB function)
<code>imread</code>	Read image file (MATLAB function)
<code>imwrite</code>	Write image file (MATLAB function)

Image Types and Type Conversions

<code>dither</code>	Convert image using dithering
<code>double</code>	Convert data to double precision (MATLAB function)
<code>gray2ind</code>	Convert intensity image to indexed image
<code>grayslice</code>	Create indexed image from intensity image by thresholding
<code>graythresh</code>	Compute global image threshold using Otsu's method
<code>im2bw</code>	Convert image to binary image by thresholding
<code>im2double</code>	Convert image array to double precision
<code>im2int16</code>	Convert image array to 16-bit signed integer
<code>im2java</code>	Convert image to instance of Java image object (MATLAB function)
<code>im2java2d</code>	Convert image to instance of Java buffered image object
<code>im2single</code>	Convert image array to single precision
<code>im2uint16</code>	Convert image array to 16-bit unsigned integers
<code>im2uint8</code>	Convert image array to 8-bit unsigned integers
<code>ind2gray</code>	Convert indexed image to intensity image
<code>ind2rgb</code>	Convert indexed image to RGB image
<code>int16</code>	Convert data to signed 16-bit integers (MATLAB function)
<code>label2rgb</code>	Convert a label matrix to an RGB image
<code>mat2gray</code>	Convert matrix to intensity image

<code>rgb2gray</code>	Convert RGB image or colormap to grayscale
<code>rgb2ind</code>	Convert RGB image to indexed image
<code>uint16</code>	Convert data to unsigned 16-bit integers (MATLAB function)
<code>uint8</code>	Convert data to unsigned 8-bit integers (MATLAB function)

Modular Interactive Tools

- Modular Tool Creation Functions (p. 15-5)
- Navigational tools (p. 15-6)
- Utility Functions (p. 15-6)

Modular Tool Creation Functions

<code>imageinfo</code>	Image Information tool
<code>imcontrast</code>	Adjust Contrast tool
<code>imshow_range</code>	Display range tool
<code>imshowinfo</code>	Pixel Information tool
<code>imshowinfoval</code>	Pixel Information tool, without text label
<code>imshowregion</code>	Pixel Region tool
<code>imshowregionpanel</code>	Pixel Region scroll panel

Navigational tools

<code>immagbox</code>	Image Information tool
<code>imoverview</code>	Adjust Contrast tool
<code>imoverviewpanel</code>	Display range tool
<code>imscrollpanel</code>	Pixel Information tool

Utility Functions

<code>axes2pix</code>	Convert axes coordinate to pixel coordinate
<code>getimage</code>	Get image data from axes
<code>getimagemodel</code>	Retrieve <code>imagemodel</code> objects from image handles
<code>imattributes</code>	Return information about image attributes
<code>imgca</code>	Get handle to current image axes
<code>imgcf</code>	Get handle to current image figure
<code>imggetfile</code>	Image Open File dialog box
<code>imhandles</code>	Get all image handles
<code>impositionrect</code>	Create position rectangle
<code>iptaddcallback</code>	Add function handle to callback list
<code>iptcheckhandle</code>	Check validity of handle
<code>iptgetapi</code>	Get Application Programmer Interface from a handle
<code>ipticondir</code>	Directories containing IPT and MATLAB icons
<code>iptremovecallback</code>	Delete function handle from callback list
<code>iptwindowalign</code>	Align figure windows
<code>trueysize</code>	Adjust display size of image

Spatial Transformation and Registration

- Spatial Transformations (p. 15-7)
- Image Registration (p. 15-7)

Spatial Transformations

checkerboard	Create checkerboard image
findbounds	Find output bounds for spatial transformation
fliptform	Flip the input and output roles of a TFORM structure
imcrop	Crop image
imresize	Resize image
imrotate	Rotate image
imtransform	Apply 2-D spatial transformation to image
makeresampler	Create resampling structure
maketform	Create geometric transformation structure
tformarray	Geometric transformation of a multidimensional array
tformfwd	Apply forward geometric transformation
tforminv	Apply inverse geometric transformation

Image Registration

cp2tform	Infer geometric transformation from control point pairs
cpcorr	Tune control point locations using cross-correlation
cpselect	Control point selection tool
cpstruct2pairs	Convert CPSTRUCT to valid pairs of control points
normxcorr2	Normalized two-dimensional cross-correlation

Image Analysis and Statistics

- Image Analysis (p. 15-8)
- Pixel Values and Statistics (p. 15-8)

Image Analysis

<code>bwboundaries</code>	Trace region boundaries in binary image
<code>bwtraceboundary</code>	Trace object in binary image
<code>edge</code>	Find edges in intensity image
<code>hough</code>	Hough transform
<code>houghlines</code>	Extract line segments based on the Hough transform
<code>houghpeaks</code>	Identify peaks in the Hough transform
<code>qtdecomp</code>	Perform quadtree decomposition
<code>qtgetblk</code>	Get block values in quadtree decomposition
<code>qtsetblk</code>	Set block values in quadtree decomposition

Texture Analysis

<code>entropy</code>	Entropy of an intensity image
<code>entropyfilt</code>	Local entropy of an intensity image
<code>graycomatrix</code>	Gray-level co-occurrence matrix
<code>graycoprops</code>	Properties of a gray-level co-occurrence matrix
<code>rangefilt</code>	Local range of an image
<code>stdfilt</code>	Local standard deviation of an image

Pixel Values and Statistics

<code>corr2</code>	Compute 2-D correlation coefficient
<code>imcontour</code>	Create contour plot of image data
<code>imhist</code>	Display histogram of image data
<code>impixel</code>	Determine pixel color values

<code>improfile</code>	Compute pixel-value cross-sections along line segments
<code>mean2</code>	Compute mean of matrix elements
<code>pixval</code>	Display information about image pixels
<code>regionprops</code>	Measure properties of image regions
<code>std2</code>	Compute standard deviation of matrix elements

Image Arithmetic

<code>imabsdiff</code>	Compute absolute difference of two images
<code>imadd</code>	Add two images, or add constant to image
<code>imcomplement</code>	Complement image
<code>imdivide</code>	Divide two images, or divide image by constant
<code>imlincomb</code>	Compute linear combination of images
<code>immultiply</code>	Multiply two images, or multiply image by constant
<code>imsubtract</code>	Subtract two images, or subtract constant from image

Image Enhancement and Restoration

- Image Enhancement (p. 15-10)
- Image Restoration (Deblurring) (p. 15-10)

Image Enhancement

<code>adapthisteq</code>	Perform adaptive histogram equalization using CLAHE
<code>decorrstretch</code>	Apply a decorrelation stretch to a multichannel image
<code>histeq</code>	Enhance contrast using histogram equalization
<code>imadjust</code>	Adjust image intensity values or colormap
<code>imnoise</code>	Add noise to an image
<code>intlut</code>	Compute new array values based on lookup table (LUT)
<code>medfilt2</code>	Perform 2-D median filtering
<code>ordfilt2</code>	Perform 2-D order-statistic filtering
<code>stretchlim</code>	Return a pair of intensities that can be used to increase the contrast of an image
<code>wiener2</code>	Perform 2-D adaptive noise-removal filtering

Image Restoration (Deblurring)

<code>deconvblind</code>	Restore image using blind deconvolution
<code>deconvlucy</code>	Restore image using accelerated Richardson-Lucy algorithm
<code>deconvreg</code>	Restore image using regularized filter
<code>deconvwnr</code>	Restore image using Wiener filter
<code>edgetaper</code>	Taper the discontinuities along the image edges
<code>otf2psf</code>	Convert optical transfer function to point spread function
<code>psf2otf</code>	Convert point spread function to optical transfer function

Linear Filtering and Transforms

- Linear Filtering (p. 15-11)
- Linear 2-D Filter Design (p. 15-11)
- Image Transforms (p. 15-11)

Linear Filtering

<code>conv2</code>	Perform 2-D convolution (MATLAB function)
<code>convmtx2</code>	Compute 2-D convolution matrix
<code>convn</code>	Perform N-D convolution (MATLAB function)
<code>filter2</code>	Perform 2-D filtering (MATLAB function)
<code>fspecial</code>	Create predefined filters
<code>imfilter</code>	Multidimensional image filtering

Linear 2-D Filter Design

<code>freqspace</code>	Determine 2-D frequency response spacing (MATLAB function.)
<code>freqz2</code>	Compute 2-D frequency response
<code>fsamp2</code>	Design 2-D FIR filter using frequency sampling
<code>ftrans2</code>	Design 2-D FIR filter using frequency transformation
<code>fwind1</code>	Design 2-D FIR filter using 1-D window method
<code>fwind2</code>	Design 2-D FIR filter using 2-D window method

Image Transforms

<code>dct2</code>	Compute 2-D discrete cosine transform
<code>dctmtx</code>	Compute discrete cosine transform matrix
<code>fan2para</code>	Convert fan-beam projection data to parallel-beam
<code>fanbeam</code>	Compute fan-beam transform
<code>fft2</code>	Compute 2-D fast Fourier transform (MATLAB function)

<code>fftn</code>	Compute N-D fast Fourier transform (MATLAB function)
<code>fftshift</code>	Reverse quadrants of output of FFT (MATLAB function)
<code>idct2</code>	Compute 2-D inverse discrete cosine transform
<code>ifft2</code>	Compute 2-D inverse fast Fourier transform (MATLAB function)
<code>ifftn</code>	Compute N-D inverse fast Fourier transform (MATLAB function)
<code>ifanbeam</code>	Compute inverse fan-beam transform
<code>iradon</code>	Compute inverse Radon transform
<code>para2fan</code>	Convert parallel-beam projections to fan-beam
<code>phantom</code>	Generate a head phantom image
<code>radon</code>	Compute Radon transform

Morphological Operations

- Intensity and Binary Images (p. 15-13)
- Binary Images (p. 15-14)
- Structuring Element (STREL) Creation and Manipulation (p. 15-15)

Intensity and Binary Images

<code>conndef</code>	Default connectivity array
<code>imbothat</code>	Perform bottom-hat filtering
<code>imclearborder</code>	Suppress light structures connected to image border
<code>imclose</code>	Close image
<code>imdilate</code>	Dilate image
<code>imerode</code>	Erode image
<code>imextendedmax</code>	Find extended-maxima transform
<code>imextendedmin</code>	Find extended-minima transform
<code>imfill</code>	Fill image regions
<code>imhmax</code>	Calculate H-maxima transform
<code>imhmin</code>	Calculate H-minima transform
<code>imimposemin</code>	Impose minima
<code>imopen</code>	Open image
<code>imreconstruct</code>	Perform morphological reconstruction
<code>imregionalmax</code>	Find regional maxima of image
<code>imregionalmin</code>	Find regional minima of image
<code>imtophat</code>	Perform tophat filtering
<code>watershed</code>	Find image watershed regions

Binary Images

<code>applylut</code>	Perform neighborhood operations using lookup tables
<code>bwarea</code>	Area of objects in binary image
<code>bwareaopen</code>	Binary area open; remove small objects
<code>bwdist</code>	Distance transform
<code>bweuler</code>	Euler number of binary image
<code>bwhitmiss</code>	Binary hit-and-miss operation
<code>bwlabel</code>	Label connected components in 2-D binary image
<code>bwlabeln</code>	Label connected components in N-D binary image
<code>bwmorph</code>	Perform morphological operations on binary image
<code>bwpack</code>	Pack binary image
<code>bwperim</code>	Find perimeter of objects in binary image
<code>bwselect</code>	Select objects in binary image
<code>bwulterode</code>	Ultimate erosion
<code>bwunpack</code>	Unpack a packed binary image
<code>imregionalmin</code>	Regional minima of image
<code>imtophat</code>	Perform tophat filtering
<code>makelut</code>	Construct lookup table for use with <code>applylut</code>

Structuring Element (STREL) Creation and Manipulation

getheight	Get the height of a structuring element
getneighbors	Get structuring element neighbor locations and heights
getnhood	Get structuring element neighborhood
getsequence	Extract sequence of decomposed structuring elements
isflat	Return true for flat structuring element
reflect	Reflect structuring element
strel	Create morphological structuring element
translate	Translate structuring element

Region-Based, Neighborhood, and Block Processing

- Region-Based Processing (p. 15-15)
- Neighborhood and Block Processing (p. 15-16)

Region-Based Processing

poly2mask	Convert region-of-interest polygon to mask
roicolor	Select region of interest, based on color
roifill	Smoothly interpolate within arbitrary region
roifilt2	Filter a region of interest
roipoly	Select polygonal region of interest

Neighborhood and Block Processing

<code>bestblk</code>	Choose block size for block processing
<code>blkproc</code>	Implement distinct block processing for image
<code>col2im</code>	Rearrange matrix columns into blocks
<code>colfilt</code>	Perform neighborhood operations using columnwise functions
<code>im2col</code>	Rearrange image blocks into columns
<code>nlfilter</code>	Perform general sliding-neighborhood operations

Colormap and Color Space Functions

- Colormap Manipulation (p. 15-16)
- Color Space Conversions (p. 15-16)

Colormap Manipulation

<code>brighten</code>	Brighten or darken colormap (MATLAB function)
<code>cmpermute</code>	Rearrange colors in colormap
<code>cmunique</code>	Find unique colormap colors and corresponding image
<code>colormap</code>	Set or get color lookup table (MATLAB function)
<code>imapprox</code>	Approximate indexed image by one with fewer colors
<code>rgbplot</code>	Plot RGB colormap components (MATLAB function)

Color Space Conversions

<code>applycform</code>	Apply device-independent color space transformation
<code>hsv2rgb</code>	Convert HSV values to RGB color space (MATLAB function)
<code>iccread</code>	Read ICC color profile
<code>iccwrite</code>	Write ICC color profile
<code>isicc</code>	Determine if ICC color profile is valid
<code>lab2double</code>	Convert $L^*a^*b^*$ color values to double

lab2uint16	Convert $L^*a^*b^*$ color values to uint16
lab2uint8	Convert $L^*a^*b^*$ color values to uint8
makecform	Create device-independent color space transform structure
ntsc2rgb	Convert NTSC values to RGB color space
rgb2hsv	Convert RGB values to HSV color space (MATLAB function)
rgb2ntsc	Convert RGB values to NTSC color space
rgb2ycbcr	Convert RGB values to YCbCr color space
whitepoint	Returns XYZ values of standard illuminants
xyz2double	Convert XYZ color values to double
xyz2uint16	Convert XYZ color values to uint16
ycbcr2rgb	Convert YCbCr values to RGB color space

Miscellaneous Functions

- Toolbox Preferences (p. 15-18)
- Interactive Mouse Utility Functions (p. 15-18)
- Array Operations (p. 15-19)
- Demos (p. 15-19)
- Performance (p. 15-19)

Toolbox Preferences

`iptgetpref` Get value of Image Processing Toolbox preference
`iptsetpref` Set value of Image Processing Toolbox preference

Toolbox Utility Functions

`getrangefromclass` Get dynamic range of image based on its class
`iptcheckconn` Check validity of connectivity argument
`iptcheckinput` Check validity of array
`iptcheckmap` Check validity of colormap
`iptchecknargin` Check number of input arguments
`iptcheckstrs` Check validity of strings
`iptnum2ordinal` Convert positive integer to ordinal string

Interactive Mouse Utility Functions

`getline` Select polyline with mouse
`getpts` Select points with mouse
`getrect` Select rectangle with mouse

Array Operations

`circshift` Shift array circularly (MATLAB function)
`padarray` Pad an array

Demos

`iptdemos` Display index of Image Processing Toolbox demos

Performance

`ipp1` Check for presence of Intel Performance Primitives Library (IPPL)

Functions – Alphabetical List

This section contains detailed descriptions of all toolbox functions. Each function reference page contains some or all of this information:

- The function name
- The purpose of the function
- The function syntax

All valid input argument and output argument combinations are shown. In some cases, an ellipsis (. . .) is used for the input arguments. This means that all preceding input argument combinations are valid for the specified output argument(s).

- A description of each argument
- A description of the function
- Additional remarks about usage
- An example of usage
- Related functions

Purpose Perform contrast-limited adaptive histogram equalization (CLAHE)

Syntax
`J = adapthisteq(I)`
`J = adapthisteq(I,param1,val1,param2,val2...)`

Description `J = adapthisteq(I)` enhances the contrast of the intensity image `I` by transforming the values using contrast-limited adaptive histogram equalization (CLAHE).

CLAHE operates on small regions in the image, called *tiles*, rather than the entire image. Each tile's contrast is enhanced, so that the histogram of the output region approximately matches the histogram specified by the 'Distribution' parameter. The neighboring tiles are then combined using bilinear interpolation to eliminate artificially induced boundaries. The contrast, especially in homogeneous areas, can be limited to avoid amplifying any noise that might be present in the image.

`J = adapthisteq(I,param1,val1,param2,val2...)` specifies any of the additional parameter/value pairs listed in the following table. Parameter names can be abbreviated, and case does not matter.

Parameter	Value
'NumTiles'	Two-element vector of positive integers specifying the number of tiles by row and column, [M N]. Both M and N must be at least 2. The total number of tiles is equal to M*N. Default: [8 8]
'ClipLimit'	Real scalar in the range [0 1] that specifies a contrast enhancement limit. Higher numbers result in more contrast. Default: 0.01

Parameter	Value
'NBins'	Positive integer scalar specifying the number of bins for the histogram used in building a contrast enhancing transformation. Higher values result in greater dynamic range at the cost of slower processing speed. Default: 256
'Range'	String specifying the range of the output image data. 'original' — Range is limited to the range of the original image, $[\min(I(:)) \max(I(:))]$. 'full' — Full range of the output image class is used. For example, for uint8 data, range is [0 255]. Default: 'full'
'Distribution'	String specifying the desired histogram shape for the image tiles. 'uniform' — Flat histogram 'rayleigh' — Bell-shaped histogram 'exponential' — Curved histogram Default: 'uniform'
'Alpha'	Nonnegative real scalar specifying a distribution parameter. Default: 0.4 Note: Only used when 'Distribution' is set to either 'rayleigh' or 'exponential'.

Class Support

Intensity image *I* can be of class uint8, uint16, int16, single, or double. The output image *J* has the same class as *I*.

Example

This example applies Contrast-Limited Adaptive Histogram Equalization (CLAHE) to an image and display the results.

```
I = imread('tire.tif');
A = adapthisteq(I,'clipLimit',0.02,'Distribution','rayleigh');
figure, imshow(I);
figure, imshow(A);
```

This example applies CLAHE to a color photograph.

```
[X MAP] = imread('shadow.tif');

% Convert indexed image to true-color (RGB) format
RGB = ind2rgb(X,MAP);

% Convert image to L*a*b* color space
cform2lab = makecform('srgb2lab');
LAB = applycform(RGB, cform2lab);

% Scale values to range from 0 to 1
L = LAB(:,:,1)/100;

% Perform CLAHE
LAB(:,:,1) = adapthisteq(L,'NumTiles',...
                        [8 8],'ClipLimit',0.005)*100;

% Convert back to RGB color space
cform2srgb = makecform('lab2srgb');
J = applycform(LAB, cform2srgb);

% Display the results
figure, imshow(RGB);
figure, imshow(J);
```

See Also

histeq

applycform

Purpose Apply color space transformation

Syntax `out = applycform(I,C)`

Description `out = applycform(I,C)` converts the color values in `I` to the color space specified in the color transformation structure `C`. The color transformation structure specifies various parameters of the transformation. See `makecform` for details.

If `I` is two-dimensional, each row is interpreted as a color. `I` typically has either three or four columns, depending on the input color space. `out` has the same number of rows and either three or four columns, depending on the output color space.

If `I` is three-dimensional, each row-column location is interpreted as a color, and `size(I,3)` is typically either three or four, depending on the input color space. `out` has the same number of rows and columns as `I`, and `size(out,3)` is either three or four, depending on the output color space.

Class Support `I` must be a real, nonsparse, finite array of class `uint8`, `uint16`, or `double`. The output array `out` has the same class and size as the input array, unless the output color space is `XYZ`. If the input is `XYZ` data of class `uint8`, the output is of class `uint16`, because there is no standard 8-bit encoding defined for `XYZ` color values.

Example Read in a color image that uses the RGB color space.

```
I = imread('peppers.png');
```

Create a color transformation structure that defines an RGB to `XYZ` conversion.

```
C = makecform('srgb2xyz');
```

Perform the transformation with `applycform`.

```
I_xyz = applycform(I,C);
```

See Also `lab2double`, `lab2uint8`, `lab2uint16`, `makecform`, `whitepoint`, `xyz2double`, `xyz2uint16`

Purpose Perform neighborhood operations on binary images using lookup tables

Syntax `A = applylut(BW,LUT)`

Description `A = applylut(BW,LUT)` performs a 2-by-2 or 3-by-3 neighborhood operation on binary image `BW` by using a lookup table (`LUT`). `LUT` is either a 16-element or 512-element vector returned by `makelut`. The vector consists of the output values for all possible 2-by-2 or 3-by-3 neighborhoods.

Class Support `BW` can be numeric or logical, and it must be real, two-dimensional, and nonsparse. `LUT` can be numeric or logical, and it must be a real vector with 16 or 512 elements. If all the elements of `LUT` are 0 or 1, then `A` is logical. If all the elements of `LUT` are integers between 0 and 255, then `A` is `uint8`. For all other cases, `A` is `double`.

Algorithm `applylut` performs a neighborhood operation on a binary image by producing a matrix of indices into `lut`, and then replacing the indices with the actual values in `lut`. The specific algorithm used depends on whether you use 2-by-2 or 3-by-3 neighborhoods.

2-by-2 Neighborhoods

For 2-by-2 neighborhoods, `length(lut)` is 16. There are four pixels in each neighborhood, and two possible states for each pixel, so the total number of permutations is $2^4 = 16$.

To produce the matrix of indices, `applylut` convolves the binary image `BW` with this matrix.

```

8     2
4     1

```

The resulting convolution contains integer values in the range [0,15]. `applylut` uses the central part of the convolution, of the same size as `BW`, and adds 1 to each value to shift the range to [1,16]. It then constructs `A` by replacing the values in the cells of the index matrix with the values in `lut` that the indices point to.

3-by-3 Neighborhoods

For 3-by-3 neighborhoods, `length(lut)` is 512. There are nine pixels in each neighborhood, and two possible states for each pixel, so the total number of permutations is $2^9 = 512$.

To produce the matrix of indices, `applylut` convolves the binary image `BW` with this matrix.

256	32	4
128	16	2
64	8	1

The resulting convolution contains integer values in the range [0,511]. `applylut` uses the central part of the convolution, of the same size as `BW`, and adds 1 to each value to shift the range to [1,512]. It then constructs `A` by replacing the values in the cells of the index matrix with the values in `lut` that the indices point to.

Example

In this example, you perform erosion using a 2-by-2 neighborhood. An output pixel is on only if all four of the input pixel's neighborhood pixels are on.

```
lut = makelut('sum(x(:)) == 4',2);  
BW = imread('text.png');  
BW2 = applylut(BW,lut);  
imshow(BW), figure, imshow(BW2)
```

The term watershed refers to a ridge that ...

... divides areas drained by different river systems.

The term watershed refers to a ridge that ...

... divides areas drained by different river systems.

See Also

`makelut`

Purpose	Interactive image contrast and brightness adjustment tool
Syntax	<code>pixelx = axes2pix(DIM, XDATA, AXESX)</code>
Description	<code>pixelx = axes2pix(DIM, XDATA, AXESX)</code> converts an axes coordinate into a pixel coordinate. For example if <code>pt = get(gca, 'CurrentPoint')</code> then <code>AXESX</code> could be <code>pt(1,1)</code> or <code>pt(1,2)</code> . <code>AXESX</code> must be in pixel coordinates. <code>XDATA</code> is a two-element vector returned by <code>get(image_handle, 'XData')</code> or <code>get(image_handle, 'YData')</code> . <code>DIM</code> is the number of image columns for the x coordinate, or the number of image rows for the y coordinate.
Class Support	<code>DIM</code> , <code>XDATA</code> , and <code>AXESX</code> can be double. The output is double.
Note	<code>axes2pix</code> performs minimal checking on the validity of <code>AXESX</code> , <code>DIM</code> , or <code>XDATA</code> . For example, <code>axes2pix</code> returns a negative coordinate if <code>AXESX</code> is less than <code>XDATA(1)</code> . The function calling <code>axes2pix</code> bears responsibility for error checking.
Examples	<p>Example with default <code>XData</code> and <code>YData</code>.</p> <pre>h = imshow('pout.tif'); [nrows,ncols] = size(get(h,'CData')); xdata = get(h,'XData') ydata = get(h,'YData') px = axes2pix(ncols,xdata,30) py = axes2pix(nrows,ydata,30)</pre> <p>Example with non-default <code>XData</code> and <code>YData</code>.</p> <pre>xdata = [10 100] ydata = [20 90] px = axes2pix(ncols,xdata,30) py = axes2pix(nrows,ydata,30)</pre>
See Also	<code>impxelinfo</code> , <code>bwselect</code> , <code>imfill</code> , <code>impxel</code> , <code>improfile</code> , <code>pixval</code> , <code>roipoly</code>

bestblk

Purpose Determine block size for block processing

Syntax `siz = bestblk([m n],k)`
`[mb,nb] = bestblk([m n],k)`

Description `siz = bestblk([m n],k)` returns, for an m-by-n image, the optimal block size for block processing. `k` is a scalar specifying the maximum row and column dimensions for the block; if the argument is omitted, it defaults to 100. The return value `siz` is a 1-by-2 vector containing the row and column dimensions for the block.

`[mb,nb] = bestblk([m n],k)` returns the row and column dimensions for the block in `mb` and `nb`, respectively.

Algorithm `bestblk` returns the optimal block size given `m`, `n`, and `k`. The algorithm for determining `siz` is

- If `m` is less than or equal to `k`, return `m`.
- If `m` is greater than `k`, consider all values between $\min(m/10, k/2)$ and `k`. Return the value that minimizes the padding required.

The same algorithm is then repeated for `n`.

Example

```
siz = bestblk([640 800],72)

siz =

    64    50
```

See Also `blkproc`

Purpose Implement distinct block processing for an image

Syntax

```
B = blkproc(A,[m n],fun)
B = blkproc(A,[m n],[mborder nborder],fun,...)
B = blkproc(A,'indexed',...)
```

Description `B = blkproc(A,[m n],fun)` processes the image `A` by applying the function `fun` to each distinct `m`-by-`n` block of `A`, padding `A` with 0's if necessary. `fun` is a function handle that accepts an `m`-by-`n` matrix, `x`, and returns a matrix, vector, or scalar `y`.

```
y = fun(x)
```

`blkproc` does not require that `y` be the same size as `x`. However, `B` is the same size as `A` only if `y` is the same size as `x`.

`B = blkproc(A,[m n],[mborder nborder],fun)` defines an overlapping border around the blocks. `blkproc` extends the original `m`-by-`n` blocks by `mborder` on the top and bottom, and `nborder` on the left and right, resulting in blocks of size $(m+2*mborder)$ -by- $(n+2*nborder)$. The `blkproc` function pads the border with 0's, if necessary, on the edges of `A`. The function `fun` should operate on the extended block.

The line below processes an image matrix as 4-by-6 blocks, each having a row border of 2 and a column border of 3. Because each 4-by-6 block has this 2-by-3 border, `fun` actually operates on blocks of size 8-by-12.

```
B = blkproc(A,[4 6],[2 3],fun)
```


`B = blkproc(A,'indexed',...)` processes `A` as an indexed image, padding with 0's if the class of `A` is `uint8` or `uint16`, or 1's if the class of `A` is `double`.

Class Support The input image `A` can be of any class supported by `fun`. The class of `B` depends on the class of the output from `fun`.

Example

This example uses `blkproc` to compute the 2-D DCT of each 8-by-8 block to the standard deviation of the elements in that block. In this example, `fun` is specified as a function handle created using `@`.

```
I = imread('cameraman.tif');
fun = @dct2;
J = blkproc(I,[8 8],fun);
imagesc(J), colormap(hot)
```



This example uses `blkproc` to set the pixels in each 16-by-16 block to the standard deviation of the elements in that block. In this example, `fun` is specified as an anonymous function.

```
I = imread('liftingbody.png');
fun = @(x) std2(x)*ones(size(x));
I2 = blkproc(I,[32 32],fun);
imshow(I), figure, imshow(I2,'DisplayRange',[])
```




Image Courtesy of NASA

See Also

`bestblk`, `colfilt`, `nlfilter`, `function_handle`


brighten

Purpose

Brighten or darken a colormap

brighten is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose	Compute the area of the objects in a binary image
Syntax	<code>total = bwarea(BW)</code>
Description	<code>total = bwarea(BW)</code> estimates the area of the objects in binary image <code>BW</code> . <code>total</code> is a scalar whose value corresponds roughly to the total number of on pixels in the image, but might not be exactly the same because different patterns of pixels are weighted differently.
Class Support	<code>BW</code> can be numeric or logical. For numeric input, any nonzero pixels are considered to be on. The return value <code>total</code> is of class <code>double</code> .
Algorithm	<p><code>bwarea</code> estimates the area of all of the on pixels in an image by summing the areas of each pixel in the image. The area of an individual pixel is determined by looking at its 2-by-2 neighborhood. There are six different patterns, each representing a different area:</p> <ul style="list-style-type: none">• Patterns with zero on pixels (area = 0)• Patterns with one on pixel (area = 1/4)• Patterns with two adjacent on pixels (area = 1/2)• Patterns with two diagonal on pixels (area = 3/4)• Patterns with three on pixels (area = 7/8)• Patterns with all four on pixels (area = 1) <p>Keep in mind that each pixel is part of four different 2-by-2 neighborhoods. This means, for example, that a single on pixel surrounded by off pixels has a total area of 1.</p>
Example	<p>This example computes the area in the objects of a 256-by-256 binary image.</p> <pre>BW = imread('circles.png'); imshow(BW);</pre>



```
bwarea(BW)
```

```
ans =
```

```
1.4187e+004
```

See Also

bweuler, bwperim

References

Pratt, William K., *Digital Image Processing*, New York, John Wiley & Sons, Inc., 1991, p. 634.

Purpose Binary area open; remove small objects

Syntax
 BW2 = bwareaopen(BW,P)
 BW2 = bwareaopen(BW,P,CONN)

Description BW2 = bwareaopen(BW,P) removes from a binary image all connected components (objects) that have fewer than P pixels, producing another binary image, BW2. The default connectivity is 8 for two dimensions, 26 for three dimensions, and conndef(ndims(BW),'maximal') for higher dimensions.

BW2 = bwareaopen(BW,P,CONN) specifies the desired connectivity. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Class Support BW can be a logical or numeric array of any dimension, and it must be nonsparse. The return value BW2 is of class logical.

Algorithm The basic steps are

- 1 Determine the connected components.
 L = bwlabeln(BW, CONN);

bwareaopen

- 2 Compute the area of each component.

```
S = regionprops(L, 'Area');
```


- 3 Remove small objects.

```
bw2 = ismember(L, find([S.Area] >= P));
```

Example

Read in the image and display it.

```
originalBW = imread('text.png');  
imshow(originalBW)
```



Remove all objects smaller than 50 pixels. Note the missing letters.

```
bwAreaOpenBW = bwareaopen(originalBW,50);  
figure, imshow(bwAreaOpenBW)
```



See Also


[bwlabel](#), [bwlabeln](#), [conndef](#), [regionprops](#)

Purpose Trace region boundaries in a binary image

Syntax

```
B = bwboundaries(BW)
B = bwboundaries(BW,CONN)
B = bwboundaries(BW,CONN,options)
[B L] = bwboundaries(...)
[B L N A] = bwboundaries()
```

Description `B = bwboundaries(BW)` traces the exterior boundaries of objects, as well as boundaries of holes inside these objects, in the binary image `BW`. `bwboundaries` also descends into the outermost objects (parents) and traces their children (objects completely enclosed by the parents). `BW` must be a binary image where nonzero pixels belong to an object and 0 pixels constitute the background. The following figure illustrates these components.



`bwboundaries` returns `B`, a `P`-by-1 cell array, where `P` is the number of objects and holes. Each cell in the cell array contains a `Q`-by-2 matrix. Each row in the matrix contains the row and column coordinates of a boundary pixel. `Q` is the number of boundary pixels for the corresponding region.

bwboundaries

`B = bwboundaries(BW,CONN)` specifies the connectivity to use when tracing parent and child boundaries. `CONN` can have either of the following scalar values.

Value	Meaning
4	4-connected neighborhood
8	8-connected neighborhood. This is the default.

`B = bwboundaries(BW,CONN,options)` specifies an optional argument, where `options` can have either of the following values:

Value	Meaning
'holes'	Search for both object and hole boundaries. This is the default.
'noholes'	Search only for object (parent and child) boundaries. This can provide better performance.

`[B,L] = bwboundaries(...)` returns the label matrix `L` as the second output argument. Objects and holes are labeled. `L` is a two-dimensional array of nonnegative integers that represent contiguous regions. The k th region includes all elements in `L` that have value k . The number of objects and holes represented by `L` is equal to $\max(L(:))$. The zero-valued elements of `L` make up the background.

`[B,L,N,A] = bwboundaries(...)` returns `N`, the number of objects found, and `A`, an adjacency matrix. The first `N` cells in `B` are object boundaries. `A` represents the parent-child-hole dependencies. `A` is a square, sparse, logical matrix with side of length $\max(L(:))$, whose rows and columns correspond to the positions of boundaries stored in `B`.

The boundaries enclosed by a `B{m}` as well as the boundary enclosing `B{m}` can both be found using `A` as follows:

```
enclosing_boundary = find(A(m,:));  
enclosed_boundaries = find(A(:,m));
```

Class Support BW can be logical or numeric and it must be real, 2-D, and nonsparse. L and N are double. A is sparse logical.

Example

Example 1

Read in and threshold an intensity image. Display the labeled objects using the jet colormap, on a gray background, with region boundaries outlined in white.

```
I = imread('rice.png');
BW = im2bw(I, graythresh(I));
[B,L] = bwboundaries(BW,'noholes');
imshow(label2rgb(L, @jet, [.5 .5 .5]))
hold on
for k = 1:length(B)
    boundary = B{k};
    plot(boundary(:,2), boundary(:,1), 'w', 'LineWidth', 2)
end
```

Example 2

Read in and display a binary image. Overlay the region boundaries on the image. Display text showing the region number (based on the label matrix) next to every boundary. Additionally, display the adjacency matrix using the MATLAB spy function.

After the image is displayed, use the zoom tool to read individual labels.

```
BW = imread('blobs.png');
[B,L,N,A] = bwboundaries(BW);
imshow(BW); hold on;
colors=['b' 'g' 'r' 'c' 'm' 'y'];
for k=1:length(B),
    boundary = B{k};
    cidx = mod(k,length(colors))+1;
    plot(boundary(:,2), boundary(:,1),
    colors(cidx),'LineWidth',2);
    %randomize text position for better visibility
    rndRow = ceil(length(boundary)/(mod(rand*k,7)+1));
    col = boundary(rndRow,2); row = boundary(rndRow,1);
    h = text(col+1, row-1, num2str(L(row,col)));
    set(h,'Color',colors(cidx),'FontSize',14,'FontWeight','bold');
end
```

bwboundaries

```
figure; spy(A);
```

Example 3

Display object boundaries in red and hole boundaries in green.

```
BW = imread('blobs.png');
[B,L,N] = bwboundaries(BW);
imshow(BW); hold on;
for k=1:length(B),
    boundary = B{k};
    if(k > N)
        plot(boundary(:,2), boundary(:,1), 'g','LineWidth',2);
    else
        plot(boundary(:,2), boundary(:,1), 'r','LineWidth',2);
    end
end
```

Example 4

Display parent boundaries in red (any empty row of the adjacency matrix belongs to a parent) and their holes in green.

```
BW = imread('blobs.png');
[B,L,N,A] = bwboundaries(BW);
imshow(BW); hold on;
for k=1:length(B),
    if(~sum(A(k,:)))
        boundary = B{k};
        plot(boundary(:,2), boundary(:,1), 'r','LineWidth',2);
        for l=find(A(:,k))'
            boundary = B{l};
            plot(boundary(:,2), boundary(:,1), 'g','LineWidth',2);
        end
    end
end
```

See Also

`bwtraceboundary`, `bwlabel`, `bwlabeln`

Purpose Distance transform

Syntax

```
D = bwdist(BW)
[D,L] = bwdist(BW)
[D,L] = bwdist(BW,METHOD)
```

Description `D = bwdist(BW)` computes the Euclidean distance transform of the binary image `BW`. For each pixel in `BW`, the distance transform assigns a number that is the distance between that pixel and the nearest nonzero pixel of `BW`. `bwdist` uses the Euclidean distance metric by default. `BW` can have any dimension. `D` is the same size as `BW`.

`[D,L] = bwdist(BW)` also computes the nearest-neighbor transform and returns it as label matrix `L`, which has the same size as `BW` and `D`. Each element of `L` contains the linear index of the nearest nonzero pixel of `BW`.

`[D,L] = bwdist(BW,METHOD)` computes the distance transform, where `METHOD` specifies an alternate distance metric. `METHOD` can take any of these values:

Method	Description
'chessboard'	In 2-D, the chessboard distance between (x_1, y_1) and (x_2, y_2) is $\max(x_1 - x_2 , y_1 - y_2)$
'cityblock'	In 2-D, the cityblock distance between (x_1, y_1) and (x_2, y_2) is $ x_1 - x_2 + y_1 - y_2 $

Method	Description
'euclidean'	In 2-D, the Euclidean distance between (x_1, y_1) and (x_2, y_2) is $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$ <p>This is the default method.</p>
'quasi-euclidean'	In 2-D, the quasi-Euclidean distance between (x_1, y_1) and (x_2, y_2) is $ x_1 - x_2 + (\sqrt{2} - 1) y_1 - y_2 , x_1 - x_2 > y_1 - y_2 $ $(\sqrt{2} - 1) x_1 - x_2 + y_1 - y_2 , \textit{otherwise}$

The METHOD string can be abbreviated.

Note `bwdist` uses fast algorithms to compute the true Euclidean distance transform, especially in the 2-D case. The other methods are provided primarily for pedagogical reasons. However, the alternative distance transforms are sometimes significantly faster for multidimensional input images, particularly those that have many nonzero elements.

Class Support

`BW` can be numeric or logical, and it must be nonsparse. `D` and `L` are double matrices with the same size as `BW`.

Example

Here is a simple example of the Euclidean distance transform.

```
bw = zeros(5,5); bw(2,2) = 1; bw(4,4) = 1
bw =
    0    0    0    0    0
    0    1    0    0    0
    0    0    0    0    0
    0    0    0    1    0
    0    0    0    0    0

[D,L] = bwdist(bw)
```


```
D =
    1.4142    1.0000    1.4142    2.2361    3.1623
    1.0000         0    1.0000    2.0000    2.2361
    1.4142    1.0000    1.4142    1.0000    1.4142
    2.2361    2.0000    1.0000         0    1.0000
    3.1623    2.2361    1.4142    1.0000    1.4142
```

```
L =
     7     7     7     7     7
     7     7     7     7    19
     7     7     7    19    19
     7     7    19    19    19
     7    19    19    19    19
```

In the nearest-neighbor matrix *L* the values 7 and 19 represent the position of the nonzero elements using linear matrix indexing. If a pixel contains a 7, its closest nonzero neighbor is at linear position 7.


This example compares the 2-D distance transforms for each of the supported distance methods. In the figure, note how the quasi-Euclidean distance transform best approximates the circular shape achieved by the Euclidean distance method.

```
bw = zeros(200,200); bw(50,50) = 1; bw(50,150) = 1;
bw(150,100) = 1;
D1 = bwdist(bw,'euclidean');
D2 = bwdist(bw,'cityblock');
D3 = bwdist(bw,'chessboard');
D4 = bwdist(bw,'quasi-euclidean');
figure
subplot(2,2,1), subimage(mat2gray(D1)), title('Euclidean')
hold on, imcontour(D1)
subplot(2,2,2), subimage(mat2gray(D2)), title('City block')
hold on, imcontour(D2)
subplot(2,2,3), subimage(mat2gray(D3)), title('Chessboard')
hold on, imcontour(D3)
subplot(2,2,4), subimage(mat2gray(D4)), title('Quasi-Euclidean')
hold on, imcontour(D4)
```



This example compares isosurface plots for the distance transforms of a 3-D image containing a single nonzero pixel in the center.

```
bw = zeros(50,50,50); bw(25,25,25) = 1;
D1 = bwdist(bw);
D2 = bwdist(bw,'cityblock');
D3 = bwdist(bw,'chessboard');
D4 = bwdist(bw,'quasi-euclidean');
figure
subplot(2,2,1), isosurface(D1,15), axis equal, view(3)
camlight, lighting gouraud, title('Euclidean')
subplot(2,2,2), isosurface(D2,15), axis equal, view(3)
camlight, lighting gouraud, title('City block')
subplot(2,2,3), isosurface(D3,15), axis equal, view(3)
camlight, lighting gouraud, title('Chessboard')
subplot(2,2,4), isosurface(D4,15), axis equal, view(3)
camlight, lighting gouraud, title('Quasi-Euclidean')
```



Algorithm

For two-dimensional Euclidean distance transforms, `bwdist` uses the second algorithm described in

Breu, Heinz, Joseph Gil, David Kirkpatrick, and Michael Werman, "Linear Time Euclidean Distance Transform Algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 17, No. 5, May 1995, pp. 529-533.

For higher dimensional Euclidean distance transforms, `bwdist` uses a nearest-neighbor search on an optimized kd-tree, as described in

Friedman, Jerome H., Jon Louis Bentley, and Raphael Ari Finkel, "An Algorithm for Finding Best Matches in Logarithmic Expected Time," *ACM Transactions on Mathematics Software*, Vol. 3, No. 3, September 1977, pp. 209-226.

For cityblock, chessboard, and quasi-Euclidean distance transforms, `bwdist` uses the two-pass, sequential scanning algorithm described in

Rosenfeld, A. and J. Pfaltz, "Sequential operations in digital picture processing," *Journal of the Association for Computing Machinery*, Vol. 13, No. 4, 1966, pp. 471-494.

The different distance measures are achieved by using different sets of weights in the scans, as described in

Paglieroni, David, "Distance Transforms: Properties and Machine Vision Applications," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, No. 1, January 1992, pp. 57-58.

See Also

watershed

Purpose Compute the Euler number of a binary image


Syntax `eul = bweuler(BW,n)`

Description `eul = bweuler(BW,n)` returns the Euler number for the binary image `BW`. The return value `eul` is a scalar whose value is the total number of objects in the image minus the total number of holes in those objects. The argument `n` can have a value of either 4 or 8, where 4 specifies 4-connected objects and 8 specifies 8-connected objects; if the argument is omitted, it defaults to 8.

Class Support `BW` can be numeric or logical and it must be real, nonsparse, and two-dimensional. The return value `eul` is of class `double`.

Example

```
BW = imread('circles.png');  
imshow(BW);
```



```
bweuler(BW)
```

```
ans =
```

```
-3
```

Algorithm `bweuler` computes the Euler number by considering patterns of convexity and concavity in local 2-by-2 neighborhoods. See [2] for a discussion of the algorithm used.

See Also `bwmorph`, `bwperim`

References

- [1] Horn, Berthold P. K., *Robot Vision*, New York, McGraw-Hill, 1986, pp. 73-77.
- [2] Pratt, William K., *Digital Image Processing*, New York, John Wiley & Sons, Inc., 1991, p. 633.

Purpose Binary hit-and-miss operation

Syntax
 BW2 = bwhitmiss(BW1,SE1,SE2)
 BW2 = bwhitmiss(BW1,INTERVAL)

Description BW2 = bwhitmiss(BW1,SE1,SE2) performs the hit-and-miss operation defined by the structuring elements SE1 and SE2. The hit-and-miss operation preserves pixels whose neighborhoods match the shape of SE1 and don't match the shape of SE2. SE1 and SE2 can be flat structuring element objects, created by `strel`, or neighborhood arrays. The neighborhoods of SE1 and SE2 should not have any overlapping elements. The syntax `bwhitmiss(BW1,SE1,SE2)` is equivalent to `imerode(BW1,SE1) & imerode(~BW1,SE2)`.

BW2 = bwhitmiss(BW1,INTERVAL) performs the hit-and-miss operation defined in terms of a single array, called an *interval*. An interval is an array whose elements can contain 1, 0, or -1. The 1-valued elements make up the domain of SE1, the -1-valued elements make up the domain of SE2, and the 0-valued elements are ignored. The syntax `bwhitmiss(INTERVAL)` is equivalent to `bwhitmiss(BW1,INTERVAL == 1, INTERVAL == -1)`.

Class Support BW1 can be a logical or numeric array of any dimension, and it must be nonsparse. BW2 is always a logical array with the same size as BW1. SE1 and SE2 must be flat STREL objects or they must be logical or numeric arrays containing 1's and 0's. INTERVAL must be an array containing 1's, 0's, and -1's.

Example This example performs the hit-and-miss operation on a binary image using an interval.

```
bw = [0 0 0 0 0 0
      0 0 1 1 0 0
      0 1 1 1 1 0
      0 1 1 1 1 0
      0 0 1 1 0 0
      0 0 1 0 0 0]

interval = [0 -1 -1
            1 1 -1
            0 1 0];
```

bwhitmiss

```
bw2 = bwhitmiss(bw,interval)
```

```
bw2 =
```

```
    0    0    0    0    0    0
    0    0    0    1    0    0
    0    0    0    0    1    0
    0    0    0    0    0    0
    0    0    0    0    0    0
    0    0    0    0    0    0
```

See Also

`imdilate`, `imerode`, `strel`

Purpose	Label connected components in a binary image
Syntax	<pre>L = bwlablel(BW,n) [L,num] = bwlablel(BW,n)</pre>
Description	<p><code>L = bwlablel(BW,n)</code> returns a matrix <code>L</code>, of the same size as <code>BW</code>, containing labels for the connected objects in <code>BW</code>. <code>n</code> can have a value of either 4 or 8, where 4 specifies 4-connected objects and 8 specifies 8-connected objects; if the argument is omitted, it defaults to 8.</p> <p>The elements of <code>L</code> are integer values greater than or equal to 0. The pixels labeled 0 are the background. The pixels labeled 1 make up one object, the pixels labeled 2 make up a second object, and so on.</p> <p><code>[L,num] = bwlablel(BW,n)</code> returns in <code>num</code> the number of connected objects found in <code>BW</code>.</p>
Remarks	<p><code>bwlablel</code> supports 2-D inputs only; <code>bwlableln</code> supports inputs of any dimension. In some cases, you might prefer to use <code>bwlableln</code> even for 2-D problems because it can be faster. If you have a 2-D input whose objects are relatively thick in the vertical direction, <code>bwlablel</code> is probably faster; otherwise <code>bwlableln</code> is probably faster.</p>
Class Support	<code>BW</code> can be logical or numeric, and it must be real, 2-D, and nonsparse. <code>L</code> is of class <code>double</code> .
Remarks	<p>You can use the MATLAB <code>find</code> function in conjunction with <code>bwlablel</code> to return vectors of indices for the pixels that make up a specific object. For example, to return the coordinates for the pixels in object 2,</p> <pre>[r,c] = find(bwlablel(BW)==2)</pre> <p>You can display the output matrix as a pseudocolor indexed image. Each object appears in a different color, so the objects are easier to distinguish than in the original image. See <code>label2rgb</code> for more information.</p>
Example	<p>This example illustrates using 4-connected objects. Notice objects 2 and 3; with 8-connected labeling, <code>bwlablel</code> would consider these a single object rather than two separate objects.</p>

bwlabel

```
BW = [1 1 1 0 0 0 0 0
      1 1 1 0 1 1 0 0
      1 1 1 0 1 1 0 0
      1 1 1 0 0 0 1 0
      1 1 1 0 0 0 1 0
      1 1 1 0 0 0 1 0
      1 1 1 0 0 1 1 0
      1 1 1 0 0 0 0 0];
```

```
L = bwlabel(BW,4)
```

```
L =
```

```
1 1 1 0 0 0 0 0
1 1 1 0 2 2 0 0
1 1 1 0 2 2 0 0
1 1 1 0 0 0 3 0
1 1 1 0 0 0 3 0
1 1 1 0 0 0 3 0
1 1 1 0 0 3 3 0
1 1 1 0 0 0 0 0
```

```
[r,c] = find(L==2);
rc = [r c]
```

```
rc =
```

```
2 5
3 5
2 6
3 6
```

Algorithm

`bwlabel` uses the general procedure outlined in reference [1], pp. 40-48:

- 1 Run-length encode the input image.
- 2 Scan the runs, assigning preliminary labels and recording label equivalences in a local equivalence table.
- 3 Resolve the equivalence classes.
- 4 Relabel the runs based on the resolved equivalence classes.

See Also

bweuler, bwlabeln, bwselect, label2rgb

Reference

[1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision, Volume I*, Addison-Wesley, 1992, pp. 28-48.

bwlabeln

Purpose Label connected components in N-D binary image

Syntax

```
L = bwlabeln(BW)
[L, NUM] = bwlabeln(BW)
[L, NUM] = bwlabeln(BW, CONN)
```

Description L = bwlabeln(BW) returns a label matrix L containing labels for the connected components in BW. BW can have any dimension; L is the same size as BW. The elements of L are integer values greater than or equal to 0. The pixels labeled 0 are the background. The pixels labeled 1 make up one object, the pixels labeled 2 make up a second object, and so on. The default connectivity is 8 for two dimensions, 26 for three dimensions, and conndef(ndims(BW), 'maximal') for higher dimensions.

[L, NUM] = bwlabeln(BW) returns in NUM the number of connected objects found in BW.

[L, NUM] = bwlabeln(BW, CONN) specifies the desired connectivity. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can also be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Remarks

bwlabe1 supports 2-D inputs only; bwlabe1n supports inputs of any dimension. In some cases, you might prefer to use bwlabe1n even for 2-D problems because it can be faster. If you have a 2-D input whose objects are relatively thick in the vertical direction, bwlabe1 is probably faster; otherwise bwlabe1n is probably faster.

Class Support

BW can be numeric or logical, and it must be real and nonsparse. L is of class double.

Example

```
BW = cat(3,[1 1 0; 0 0 0; 1 0 0],...
[0 1 0; 0 0 0; 0 1 0],...
[0 1 1; 0 0 0; 0 0 1])
```

```
bwlabe1n(BW)
```

```
ans(:,:,1) =
```

```
    1    1    0
    0    0    0
    2    0    0
```

```
ans(:,:,2) =
```

```
    0    1    0
    0    0    0
    0    2    0
```

```
ans(:,:,3) =
```

```
    0    1    1
    0    0    0
    0    0    2
```

Algorithm

bwlabe1n uses the following general procedure:

- 1 Scan all image pixels, assigning preliminary labels to nonzero pixels and recording label equivalences in a union-find table.

bwlabeln

- 2** Resolve the equivalence classes using the union-find algorithm [1].
- 3** Relabel the pixels based on the resolved equivalence classes.

See Also

`bwlabel`, `label2rgb`

Reference

[1] Sedgewick, Robert, *Algorithms in C*, 3rd Ed., Addison-Wesley, 1998, pp. 11-20.

Purpose Perform morphological operations on binary images

Syntax
 BW2 = bwmorph(BW,operation)
 BW2 = bwmorph(BW,operation,n)

Description BW2 = bwmorph(BW,operation) applies a specific morphological operation to the binary image BW.

BW2 = bwmorph(BW,operation,n) applies the operation n times. n can be Inf, in which case the operation is repeated until the image no longer changes.

operation is a string that can have one of the values listed below.

Operation	Description									
'bothat'	Performs the morphological “bottom hat” operation, which is closing (dilation followed by erosion) and subtracts the original image.									
'bridge'	Bridges unconnected pixels, that is, sets 0-valued pixels to 1 if they have two nonzero neighbors that are not connected. For example: <div style="text-align: center; margin-top: 10px;"> <table style="border: none;"> <tr> <td style="padding-right: 20px;">1 0 0</td> <td style="padding-right: 20px;"></td> <td style="padding-right: 20px;">1 1 0</td> </tr> <tr> <td>1 0 1</td> <td style="padding: 0 10px;">becomes</td> <td>1 1 1</td> </tr> <tr> <td>0 0 1</td> <td></td> <td>0 1 1</td> </tr> </table> </div>	1 0 0		1 1 0	1 0 1	becomes	1 1 1	0 0 1		0 1 1
1 0 0		1 1 0								
1 0 1	becomes	1 1 1								
0 0 1		0 1 1								
'clean'	Removes isolated pixels (individual 1's that are surrounded by 0's), such as the center pixel in this pattern. <div style="text-align: center; margin-top: 10px;"> <table style="border: none;"> <tr><td>0 0 0</td></tr> <tr><td>0 1 0</td></tr> <tr><td>0 0 0</td></tr> </table> </div>	0 0 0	0 1 0	0 0 0						
0 0 0										
0 1 0										
0 0 0										
'close'	Performs morphological closing (dilation followed by erosion).									

Operation	Description
'diag'	Uses diagonal fill to eliminate 8-connectivity of the background. For example: $\begin{array}{ccc} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{array} \quad \text{becomes} \quad \begin{array}{ccc} 0 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 0 \end{array}$
'dilate'	Performs dilation using the structuring element ones(3).
'erode'	Performs erosion using the structuring element ones(3)
'fill'	Fills isolated interior pixels (individual 0's that are surrounded by 1's), such as the center pixel in this pattern. $\begin{array}{ccc} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{array}$
'hbreak'	Removes H-connected pixels. For example: $\begin{array}{ccc} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{array} \quad \text{becomes} \quad \begin{array}{ccc} 1 & 1 & 1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{array}$
'majority'	Sets a pixel to 1 if five or more pixels in its 3-by-3 neighborhood are 1's; otherwise, it sets the pixel to 0.
'open'	Performs morphological opening (erosion followed by dilation).
'remove'	Removes interior pixels. This option sets a pixel to 0 if all its 4-connected neighbors are 1, thus leaving only the boundary pixels on.

Operation	Description															
'shrink'	With $n = \text{Inf}$, shrinks objects to points. It removes pixels so that objects without holes shrink to a point, and objects with holes shrink to a connected ring halfway between each hole and the outer boundary. This option preserves the Euler number															
'skel'	With $n = \text{Inf}$, removes pixels on the boundaries of objects but does not allow objects to break apart. The pixels remaining make up the image skeleton. This option preserves the Euler number.															
'spur'	Removes spur pixels. For example: <table style="margin-left: auto; margin-right: auto;"> <tr> <td style="padding-right: 20px;">0 0 0 0</td> <td></td> <td style="padding-right: 20px;">0 0 0 0</td> </tr> <tr> <td>0 0 0 0</td> <td></td> <td>0 0 0 0</td> </tr> <tr> <td>0 0 1 0</td> <td style="padding: 0 10px;">becomes</td> <td>0 0 0 0</td> </tr> <tr> <td>0 1 0 0</td> <td></td> <td>0 1 0 0</td> </tr> <tr> <td>1 1 0 0</td> <td></td> <td>1 1 0 0</td> </tr> </table>	0 0 0 0		0 0 0 0	0 0 0 0		0 0 0 0	0 0 1 0	becomes	0 0 0 0	0 1 0 0		0 1 0 0	1 1 0 0		1 1 0 0
0 0 0 0		0 0 0 0														
0 0 0 0		0 0 0 0														
0 0 1 0	becomes	0 0 0 0														
0 1 0 0		0 1 0 0														
1 1 0 0		1 1 0 0														
'thicken'	With $n = \text{Inf}$, thickens objects by adding pixels to the exterior of objects until doing so would result in previously unconnected objects being 8-connected. This option preserves the Euler number.															
'thin'	With $n = \text{Inf}$, thins objects to lines. It removes pixels so that an object without holes shrinks to a minimally connected stroke, and an object with holes shrinks to a connected ring halfway between each hole and the outer boundary. This option preserves the Euler number															
'tophat'	Performs morphological “top hat” operation, returning the image minus the morphological opening of the image.															


Class Support

The input image `BW` can be numeric or logical. It must be 2-D, real and nonspare. The output image `BW2` is of class `logical`.


Example

```
BW = imread('circles.png');
imshow(BW);
```


bwmorph



```
BW2 = bwmorph(BW, 'remove');  
figure, imshow(BW2)
```



```
BW3 = bwmorph(BW, 'skel', Inf);  
figure, imshow(BW3)
```

**See Also**

bweuler, bwperim, imdilate, imerode

References

- [1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision, Volume I*, Addison-Wesley, 1992.
- [2] Pratt, William K., *Digital Image Processing*, John Wiley & Sons, Inc., 1991.
- [3] Lam, L., Seong-Whan Lee, and Ching Y. Suen, "Thinning Methodologies-A Comprehensive Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol 14, No. 9, September 1992, page 879, bottom of first column through top of second column.


bwpack

Purpose Pack binary image

Syntax `BWP = bwpack(BW)`

Description `BWP = bwpack(BW)` packs the `uint8` binary image `BW` into the `uint32` array `BWP`, which is known as a *packed binary image*. Because each 8-bit pixel value in the binary image has only two possible values, 1 and 0, `bwpack` can map each pixel to a single bit in the packed output image.

`bwpack` processes the image pixels by column, mapping groups of 32 pixels into the bits of a `uint32` value. The first pixel in the first row corresponds to the least significant bit of the first `uint32` element of the output array. The first pixel in the 32nd input row corresponds to the most significant bit of this same element. The first pixel of the 33rd row corresponds to the least significant bit of the second output element, and so on. If `BW` is `M`-by-`N`, then `BWP` is `ceil(M/32)`-by-`N`. This figure illustrates how `bwpack` maps the pixels in a binary image to the bits in a packed binary image.



Binary image packing is used to accelerate some binary morphological operations, such as dilation and erosion. If the input to `imdilate` or `imerode` is a packed binary image, the functions use a specialized routine to perform the operation faster.

`bwunpack` is used to unpack packed binary images.

Class Support

BW can be logical or numeric, and it must be 2-D, real, and nonsparse. BWP is of class `uint32`.

Example

Pack, dilate, and unpack a binary image:

```
BW = imread('text.png');  
BWp = bwpack(BW);  
BWp_dilated = imdilate(BWp,ones(3,3),'ispacked');  
BW_dilated = bwunpack(BWp_dilated, size(BW,1));
```

See Also

`bwunpack`, `imdilate`, `imerode`

bwperim

Purpose Find perimeter pixels in binary image

Syntax
BW2 = bwperim(BW1)
BW2 = bwperim(BW1,CONN)

Description BW2 = bwperim(BW1) returns a binary image containing only the perimeter pixels of objects in the input image BW1. A pixel is part of the perimeter if it is nonzero and it is connected to at least one zero-valued pixel. The default connectivity is 4 for two dimensions, 6 for three dimensions, and conndef(ndims(BW), 'minimal') for higher dimensions.

BW2 = bwperim(BW1,CONN) specifies the desired connectivity. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood



Connectivity can also be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Class Support BW1 must be logical or numeric, and it must be nonsparse. BW2 is of class logical.

Example

```
BW1 = imread('circbw.tif');  
BW2 = bwperim(BW1,8);  
imshow(BW1)
```

```
figure, imshow(BW2)
```

**See Also**

`bwarea`, `bweuler`, `conndef`, `imfill`

bwselect

Purpose Select objects in a binary image

Syntax

```
BW2 = bwselect(BW,c,r,n)
BW2 = bwselect(BW,n)
[BW2,idx] = bwselect(...)

BW2 = bwselect(x,y,BW,xi,yi,n)
[x,y,BW2,idx,xi,yi] = bwselect(...)
```

Description `BW2 = bwselect(BW,c,r,n)` returns a binary image containing the objects that overlap the pixel (r,c) . r and c can be scalars or equal-length vectors. If r and c are vectors, `BW2` contains the sets of objects overlapping with any of the pixels $(r(k),c(k))$. n can have a value of either 4 or 8 (the default), where 4 specifies 4-connected objects and 8 specifies 8-connected objects. Objects are connected sets of on pixels (i.e., pixels having a value of 1).

`BW2 = bwselect(BW,n)` displays the image `BW` on the screen and lets you select the (r,c) coordinates using the mouse. If you omit `BW`, `bwselect` operates on the image in the current axes. Use normal button clicks to add points. Pressing **Backspace** or **Delete** removes the previously selected point. A shift-click, right-click, or double-click selects the final point; pressing **Return** finishes the selection without adding a point.

`[BW2,idx] = bwselect(...)` returns the linear indices of the pixels belonging to the selected objects.

`BW2 = bwselect(x,y,BW,xi,yi,n)` uses the vectors x and y to establish a nondefault spatial coordinate system for `BW1`. xi and yi are scalars or equal-length vectors that specify locations in this coordinate system.


`[x,y,BW2,idx,xi,yi] = bwselect(...)` returns the `XData` and `YData` in x and y , the output image in `BW2`, linear indices of all the pixels belonging to the selected objects in `idx`, and the specified spatial coordinates in xi and yi .

If `bwselect` is called with no output arguments, the resulting image is displayed in a new figure.


Class Support The input image `BW` can be logical or numeric and must be 2-D and nonsparse. The output image `BW2` is of class `logical`.

Example

```
BW1 = imread('text.png');  
c = [43 185 212];  
r = [38 68 181];  
BW2 = bwselect(BW1,c,r,4);  
imshow(BW1), figure, imshow(BW2)
```



The term watershed
refers to a ridge that ...
... divides areas
drained by different
river systems.

**See Also**

[bwlabel](#), [imfill](#), [impixel](#), [roipoly](#), [roifill](#)


bwtraceboundary

Purpose Trace object in a binary image

Syntax

```
B = bwtraceboundary(BW,P,fstep)
B = bwtraceboundary(BW,P,fstep,CONN)
B = bwtraceboundary(...,N,dir)
```

Description `B = bwtraceboundary(BW,P,fstep)` traces the outline of an object in binary image `bw`. Nonzero pixels belong to an object and 0 pixels constitute the background. `P` is a two-element vector specifying the row and column coordinates of the point on the object boundary where you want the tracing to begin. `fstep` is a string specifying the initial search direction for the next object pixel connected to `P`. You use strings such as 'N' for north, 'NE' for northeast, to specify the direction. The following figure illustrates all the possible values for `fstep`.



`bwtraceboundary` returns `B`, a `Q`-by-2 matrix, where `Q` is the number of boundary pixels for the region. `B` holds the row and column coordinates of the boundary pixels.

`B = bwtraceboundary(bw,P,fstep,CONN)` specifies the connectivity to use when tracing the boundary. `CONN` can have either of the following scalar values.

Value	Meaning
4	4-connected neighborhood Note: With this connectivity, <code>fstep</code> is limited to the following values: 'N', 'E', 'S', and 'W'.
8	8-connected neighborhood. This is the default.

`B = bwtraceboundary(...,N,dir)` specifies `n`, the maximum number of boundary pixels to extract, and `dir`, the direction in which to trace the boundary. When `N` is set to `Inf`, the default value, the algorithm identifies all the pixels on the boundary. `dir` can have either of the following values:

Value	Meaning
'clockwise'	Search in a clockwise direction. This is the default.
'counterclockwise'	Search in counterclockwise direction.

Class Support

`BW` can be logical or numeric and it must be real, 2-D, and nonsparse. `B`, `P`, `CONN`, and `N` are of class `double`. `dir` and `fstep` are strings.

Example

Read in and display a binary image. Starting from the top left, project a beam across the image searching for the first nonzero pixel. Use the location of that pixel as the starting point for the boundary tracing. Including the starting point, extract 50 pixels of the boundary and overlay them on the image. Mark the starting points with a green x. Mark beams that missed their targets with a red x.

```
BW = imread('blobs.png');  
imshow(BW,[]);  
s=size(BW);  
for row = 2:55:s(1)  
    for col=1:s(2)
```

bwtraceboundary

```
        if BW(row,col),
            break;
        end
    end

    contour = bwtraceboundary(BW, [row, col], 'W', 8, 50,...
        'counterclockwise');

    if(~isempty(contour))
        hold on;
        plot(contour(:,2),contour(:,1),'g','LineWidth',2);
        hold on;
        plot(col, row,'gx','LineWidth',2);
    else
        hold on; plot(col, row,'rx','LineWidth',2);
    end
end
```

See Also

[bwboundaries](#)

Purpose Ultimate erosion

Syntax
 BW2 = bwulterode(BW)
 BW2 = bwulterode(BW,METHOD,CONN)

Description BW2 = bwulterode(BW) computes the ultimate erosion of the binary image BW. The ultimate erosion of BW consists of the regional maxima of the Euclidean distance transform of the complement of BW. The default connectivity for computing the regional maxima is 8 for two dimensions, 26 for three dimensions, and conndef(ndims(BW), 'maximal') for higher dimensions.

BW2 = bwulterode(BW,METHOD,CONN) specifies the distance transform method and the regional maxima connectivity. METHOD can be one of the strings 'euclidean', 'cityblock', 'chessboard', and 'quasi-euclidean'.

CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for CONN a 3-by-3-by... - by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Class Support BW can be numeric or logical and it must be nonsparse. It can have any dimension. The return value BW2 is always a logical array.

bwulterode

Example

```
originalBW = imread('circles.png');  
imshow(originalBW)  
ultimateErosion = bwulterode(originalBW);  
figure, imshow(ultimateErosion)
```

See Also

`bwdist`, `conndef`, `imregionalmax`

Purpose Unpack binary image

Syntax `BW = bwunpack(BWP,M)`

Description `BW = bwunpack(BWP,M)` unpacks the packed binary image BWP. BWP is a `uint32` array. When it unpacks BWP, `bwunpack` maps the least significant bit of the first row of BWP to the first pixel in the first row of BW. The most significant bit of the first element of BWP maps to the first pixel in the 32nd row of BW, and so on. BW is M-by-N, where N is the number of columns of BWP. If M is omitted, its default value is `32*size(BWP,1)`.

Binary image packing is used to accelerate some binary morphological operations, such as dilation and erosion. If the input to `imdilate` or `imerode` is a packed binary image, the functions use a specialized routine to perform the operation faster.

`bwpack` is used to create packed binary images.

Class Support BWP is of class `uint32` and must be real, 2-D, and nonsparse. The return value BW is of class `uint8`.

Example Pack, dilate, and unpack a binary image.

```
bw = imread('text.png');  
bwp = bwpack(bw);  
bwp_dilated = imdilate(bwp,ones(3,3),'ispacked');  
bw_dilated = bwunpack(bwp_dilated, size(bw,1));
```

See Also `bwpack`, `imdilate`, `imerode`

checkerboard

Purpose Create checkerboard image

Syntax
I = checkerboard
I = checkerboard(N)
I = checkerboard(N,P,Q)

Description I = checkerboard creates an 8-by-8 square checkerboard image that has four identifiable corners. Each square has 10 pixels per side. The light squares on the left half of the checkerboard are white. The light squares on the right half of the checkerboard are gray.

I = checkerboard(N) creates a checkerboard image where each square has N pixels per side.

I = checkerboard(N,P,Q) creates a rectangular checkerboard where P specifies the number of rows and Q specifies the number of columns. If you omit Q, it defaults to P and the checkerboard is square.


Each row and column is made up of tiles. Each tile contains four squares, N pixels per side, defined as

TILE = [DARK LIGHT; LIGHT DARK] 

The light squares on the left half of the checkerboard are white. The light squares on the right half of the checkerboard are gray.

Example Create a checkerboard where the side of every square is 20 pixels in length.


```
I = checkerboard(20);  
imshow(I)
```



Create a rectangular checkerboard that is 2 tiles in height and 3 tiles wide.


```
J = checkerboard(10,2,3);
```

```
figure, imshow(J)
```



Create a black and white checkerboard.

```
K = (checkerboard > 0.5);  
figure, imshow(K)
```



See Also

[cp2tform](#), [imtransform](#), [maketform](#)

cmpermute

Purpose Rearrange the colors in a colormap

Syntax `[Y,newmap] = cmpermute(X,map)`
`[Y,newmap] = cmpermute(X,map,index)`

Description `[Y,newmap] = cmpermute(X,map)` randomly reorders the colors in `map` to produce a new colormap `newmap`. The `cmpermute` function also modifies the values in `X` to maintain correspondence between the indices and the colormap, and returns the result in `Y`. The image `Y` and associated colormap `newmap` produce the same image as `X` and `map`.

`[Y,newmap] = cmpermute(X,map,index)` uses an ordering matrix (such as the second output of `sort`) to define the order of colors in the new colormap.

Class Support The input image `X` can be of class `uint8` or `double`. `Y` is returned as an array of the same class as `X`.

Example To arrange a colormap in order by luminance, use

```
ntsc = rgb2ntsc(map);  
[dum,index] = sort(ntsc(:,1));  
[Y,newmap] = cmpermute(X,map,index);
```

This example orders a colormap by luminance.

```
load trees  
ntsc = rgb2ntsc(map);  
[dum,index] = sort(ntsc(:,1));  
[Y,newmap] = cmpermute(X,map,index);  
figure, imshow(X,map)  
figure, imshow(Y,newmap)
```

See Also `randperm`, `sort` in the MATLAB Function Reference

Purpose Find unique colormap colors and the corresponding image

Syntax

```
[Y,newmap] = cmunique(X,map)
[Y,newmap] = cmunique(RGB)
[Y,newmap] = cmunique(I)
```

Description [Y,newmap] = cmunique(X,map) returns the indexed image Y and associated colormap newmap that produce the same image as (X,map) but with the smallest possible colormap. The cmunique function removes duplicate rows from the colormap and adjusts the indices in the image matrix accordingly.

[Y,newmap] = cmunique(RGB) converts the true-color image RGB to the indexed image Y and its associated colormap newmap. The return value newmap is the smallest possible colormap for the image, containing one entry for each unique color in RGB. (Note that newmap might be very large, because the number of entries can be as many as the number of pixels in RGB.)

[Y,newmap] = cmunique(I) converts the intensity image I to an indexed image Y and its associated colormap newmap. The return value newmap is the smallest possible colormap for the image, containing one entry for each unique intensity level in I.

Class Support The input image can be of class uint8, uint16, or double. The class of the output image Y is uint8 if the length of newmap is less than or equal to 256. If the length of newmap is greater than 256, Y is of class double.

Example Use the magic function to create a sample 4-by-4 image that uses every value in the range between 1 and 16.

```
X = magic(4)

X =

    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1
```

cmunique

Concatenate two 8-entry grayscale colormaps created using the `gray` function. The resultant colormap, `map`, has 16 entries. Entries 9 through 16 are duplicates of entries 1 through 8.

```
map = [gray(8); gray(8)]
size(map)
```

```
ans =
```

```
16    3
```

Use `cmunique` to eliminate duplicate entries in the colormap.

```
[Y, newmap] = cmunique(X, map);
size(newmap)
```

```
ans =
```

```
8    3
```

`cmunique` adjusts the values in the original image `X` to index the new colormap.

```
Y =
```

```
7    1    2    4
4    2    1    7
0    6    5    3
3    5    6    0
```

View both images to verify that their appearance is the same.

```
imshow(X, map, 'InitialMagnification', 'fit')
figure, imshow(Y, newmap, 'InitialMagnification', 'fit')
```

See Also

`gray2ind`, `rgb2ind`

Purpose	Rearrange matrix columns into blocks
Syntax	<pre>A = col2im(B,[m n],[mm nn], block_type) A = col2im(B,[m n],[mm nn])</pre>
Description	<p>col2im rearranges matrix columns into blocks. block_type is a string with one of these values:</p> <ul style="list-style-type: none">• 'distinct' for m-by-n distinct blocks• 'sliding' for m-by-n sliding blocks (default) <p>A = col2im(B,[m n],[mm nn], 'distinct') rearranges each column of B into a distinct m-by-n block to create the matrix A of size mm-by-nn. If B = [A11(:) A12(:) A21(:) A22(:)], where each column has length m*n, then A = [A11 A12;A21 A22] where each Aij is m-by-n.</p> <p>A = col2im(B,[m n],[mm nn], 'sliding') rearranges the row vector B into a matrix of size (mm-m+1)-by-(nn-n+1). B must be a vector of size 1-by-(mm-m+1)*(nn-n+1). B is usually the result of processing the output of im2col(..., 'sliding') using a column compression function (such as sum).</p> <p>A = col2im(B,[m n],[mm nn]) uses the default block_type of 'sliding'.</p>
Class Support	B can be logical or numeric. The return value A is of the same class as B.
See Also	blkproc, colfilt, im2col, nlfilt

colfilt

Purpose

Perform neighborhood operations using columnwise functions

Syntax

```
B = colfilt(A,[m n],block_type,fun)
B = colfilt(A,[m n],[mblock nblock],block_type,fun)
B = colfilt(A,'indexed',...)
```

Description

`colfilt` processes distinct or sliding blocks as columns. `colfilt` can perform operations similar to `blkproc` and `nfilter`, but often executes much faster.

`B = colfilt(A,[m n],block_type,fun)` processes the image `A` by rearranging each `m`-by-`n` block of `A` into a column of a temporary matrix, and then applying the function `fun` to this matrix. `fun` must be a function handle. `colfilt` zero-pads `A`, if necessary.

Before calling `fun`, `colfilt` calls `im2col` to create the temporary matrix. After calling `fun`, `colfilt` rearranges the columns of the matrix back into `m`-by-`n` blocks using `col2im`.

`block_type` is a string that can have one of the values listed in this table.

Value	Description
'distinct'	Rearranges each <code>m</code> -by- <code>n</code> distinct block of <code>A</code> into a column in a temporary matrix, and then applies the function <code>fun</code> to this matrix. <code>fun</code> must return a matrix the same size as the temporary matrix. <code>colfilt</code> then rearranges the columns of the matrix returned by <code>fun</code> into <code>m</code> -by- <code>n</code> distinct blocks.
'sliding'	Rearranges each <code>m</code> -by- <code>n</code> sliding neighborhood of <code>A</code> into a column in a temporary matrix, and then applies the function <code>fun</code> to this matrix. <code>fun</code> must return a row vector containing a single value for each column in the temporary matrix. (Column compression functions such as <code>sum</code> return the appropriate type of output.) <code>colfilt</code> then rearranges the vector returned by <code>fun</code> into a matrix the same size as <code>A</code> .

`B = colfilt(A,[m n],[mblock nblock],block_type,fun)` processes the matrix `A` as above, but in blocks of size `mblock`-by-`nblock` to save memory. Note

that using the `[mblock nblock]` argument does not change the result of the operation.

`B = colfilt(A, 'indexed', ...)` processes `A` as an indexed image, padding with 0's if the class of `A` is `uint8` or `uint16`, or 1's if the class of `A` is `double` or `single`.

Note To save memory, the `colfilt` function might divide `A` into subimages and process one subimage at a time. This implies that `fun` may be called multiple times, and that the first argument to `fun` may have a different number of columns each time.

Class Support

The input image `A` can be of any class supported by `fun`. The class of `B` depends on the class of the output from `fun`.

Example

This example sets each output pixel to the mean value of the input pixel's 5-by-5 neighborhood.

```
I = imread('tire.tif');
imshow(I)
I2 = uint8(colfilt(I,[5 5],'sliding',@mean));
figure, imshow(I2)
```

See Also

`blkproc`, `col2im`, `function_handle`, `im2col`, `nlfilter`

colorbar

Purpose

Display a color bar

colorbar is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose Create connectivity array

Syntax `CONN = conndef(NUM_DIMS,TYPE)`

Description `CONN = conndef(NUM_DIMS,TYPE)` returns the connectivity array defined by `TYPE` for `NUM_DIMS` dimensions. `TYPE` can have either of the values listed in this table.

Value	Description
'minimal'	Defines a neighborhood whose neighbors are touching the central element on an (N-1)-dimensional surface, for the N-dimensional case.
'maximal'	Defines a neighborhood including neighbors that touch the central element in any way; it is <code>ones(repmat(3,1,NUM_DIMS))</code> .

Several Image Processing Toolbox functions use `conndef` to create the default connectivity input argument.

Examples

The minimal connectivity array for two dimensions includes the neighbors touching the central element along a line.

```
conn1 = conndef(2,'minimal')
```

```
conn1 =
     0     1     0
     1     1     1
     0     1     0
```

The minimal connectivity array for three dimensions includes all the neighbors touching the central element along a face.

```
conndef(3,'minimal')
```

```
ans(:,:,1) =
     0     0     0
     0     1     0
     0     0     0
```

conndef

```
ans(:,:,2) =  
    0    1    0  
    1    1    1  
    0    1    0
```

```
ans(:,:,3) =  
    0    0    0  
    0    1    0  
    0    0    0
```

The maximal connectivity array for two dimensions includes all the neighbors touching the central element in any way.

```
conn2 = conndef(2, 'maximal')
```

```
conn2 =  
    1    1    1  
    1    1    1  
    1    1    1
```

Purpose

Perform two-dimensional convolution

`conv2` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

convmtx2

Purpose Compute two-dimensional convolution matrix

Syntax `T = convmtx2(H,m,n)`
`T = convmtx2(H,[m n])`

Description `T = convmtx2(H,m,n)` or `T = convmtx2(H,[m n])` returns the convolution matrix `T` for the matrix `H`. If `X` is an `m`-by-`n` matrix, then `reshape(T*X(:),size(H)+[m n]-1)` is the same as `conv2(X,H)`.

Class Support The inputs are all of class `double`. The output matrix `T` is of class `sparse`. The number of nonzero elements in `T` is no larger than `prod(size(H))*m*n`.

See Also `conv2`
`convmtx` in the Signal Processing Toolbox User's Guide documentation

Purpose

Perform N-dimensional convolution

`convn` is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

corr2

Purpose Compute the two-dimensional correlation coefficient between two matrices

Syntax `r = corr2(A,B)`

Description `r = corr2(A,B)` computes the correlation coefficient between A and B, where A and B are matrices or vectors of the same size.

Class Support A and B can be numeric or logical. The return value r is a scalar double.

Algorithm `corr2` computes the correlation coefficient using

$$r = \frac{\sum_m \sum_n (A_{mn} - \bar{A})(B_{mn} - \bar{B})}{\sqrt{\left(\sum_m \sum_n (A_{mn} - \bar{A})^2\right)\left(\sum_m \sum_n (B_{mn} - \bar{B})^2\right)}}$$

where $\bar{A} = \text{mean2}(A)$, and $\bar{B} = \text{mean2}(B)$.

See Also `std2`

`corrcoef` in the MATLAB Function Reference

Purpose

Infer geometric transformation from control point pairs

Syntax

```
TFORM = cp2tform(input_points,base_points,transformtype)
TFORM = cp2tform(CPSTRUCT,transformtype)
TFORM = cp2tform(input_points,base_points,transformtype,parameter)
TFORM = cp2tform(CPSTRUCT,transformtype,parameter)
[TFORM,input_points,base_points] = cp2tform(CPSTRUCT,...)
[TFORM,input_points,base_points,input_points_bad,base_points_bad]
    = cp2tform(...,'piecewise linear')
```


Description



`TFORM = cp2tform(input_points,base_points,transformtype)` takes pairs of control points and uses them to infer a spatial transformation. `input_points` is an `m-by-2` double matrix containing the x - and y -coordinates of control points in the image you want to transform. `base_points` is an `m-by-2` double matrix containing the x - and y -coordinates of control points specified in the base image. The function returns a `TFORM` structure containing the spatial transformation.

`TFORM = cp2tform(CPSTRUCT,transformtype)` takes pairs of control points and uses them to infer a spatial transformation. `CPSTRUCT` is a structure that contains the control point matrices for the input and base images. You use the Control Point Selection Tool to create the `CPSTRUCT`.

`transformtype` specifies the type of spatial transformation to infer. This table lists all the transformation types supported by `cp2tform` in order of complexity.

The 'lwm' and 'polynomial' transform types can each take an optional, additional parameter. See the syntax descriptions that follow for details.

Transformation Type	Description	Minimum Control Points	Example
'linear conformal'	Use this transformation when shapes in the input image are unchanged, but the image is distorted by some combination of translation, rotation, and scaling. Straight lines remain straight, and parallel lines are still parallel.	2 pairs	
'affine'	Use this transformation when shapes in the input image exhibit shearing. Straight lines remain straight, and parallel lines remain parallel, but rectangles become parallelograms.	3 pairs	
'projective'	Use this transformation when the scene appears tilted. Straight lines remain straight, but parallel lines converge toward vanishing points that might or might not fall within the image.	4 pairs	
'polynomial'	Use this transformation when objects in the image are curved. The higher the order of the polynomial, the better the fit, but the result can contain more curves than the base image.	6 pairs (order 2) 10 pairs (order 3) 16 pairs (order 4)	

Transformation Type	Description	Minimum Control Points	Example
'piecewise linear'	Use this transformation when parts of the image appear distorted differently.	4 pairs	
'lwm'	Use this transformation (local weighted mean), when the distortion varies locally and piecewise linear is not sufficient.	6 pairs (12 pairs recommended)	

Note When `transformtype` is 'linear conformal', 'affine', 'projective', or 'polynomial', and `input_points` and `base_points` (or `CPSTRUCT`) have the minimum number of control points needed for a particular transformation, `cp2tform` finds the coefficients exactly. If `input_points` and `base_points` include more than the minimum number of points, `cp2tform` uses a least squares solution. For more information, see `mldivide`.

`TFORM = cp2tform(input_points,base_points,'polynomial',order)` returns a `TFORM` structure specifying a 'polynomial' transformation, where `order` specifies the order of the polynomial to use. `order` can be the scalar value 2, 3, or 4. If you omit `order`, it defaults to 3.

`TFORM = cp2tform(CPSTRUCT,'polynomial',order)` same as the previous syntax except that the control points are specified in a `CPSTRUCT`.

`TFORM = cp2tform(input_points,base_points,'lwm',N)` returns a `TFORM` structure specifying a 'lwm' transformation, where `N` specifies the number of points used to infer each polynomial. The radius of influence extends out to the furthest control point used to infer that polynomial. The `N` closest points are used to infer a polynomial of order 2 for each control point pair. If you omit `N`, it defaults to 12. `N` can be as small as 6, but making `N` small risks generating ill-conditioned polynomials.

`TFORM = cp2tform(CPSTRUCT, 'lwm', N)` same as the previous syntax except that the control points are specified in a `CPSTRUCT`.

`[TFORM, input_points, base_points] = cp2tform(CPSTRUCT, ...)` returns the control points that were actually used in the return values `input_points` and `base_points`. Unmatched and predicted points are not used. For more information, see `cpstruct2pairs`.

`[TFORM, input_points, base_points, input_points_bad, base_points_bad] = cp2tform(input_points, base_points, 'piecewise linear')` returns a `TFORM` structure specifying a 'piecewise linear' transformation. Returns the control points that were actually used in `input_points` and `base_points`, and returns the control points that were eliminated because they were middle vertices of degenerate fold-over triangles, in `input_points_bad` and `base_points_bad`.

`[TFORM, input_points, base_points, input_points_bad, base_points_bad] = cp2tform(CPSTRUCT, 'piecewise linear')` same as the previous syntax except that the control points are specified in a `CPSTRUCT`.

Algorithms

`cp2tform` uses the following general procedure:

- 1 Use valid pairs of control points to infer a spatial transformation or an inverse mapping from output space (x,y) to input space (u,v) according to `transformtype`.
- 2 Return `TFORM` structure containing spatial transformation.

The procedure varies depending on the `transformtype`.

Linear Conformal

Linear conformal transformations can include a rotation, a scaling, and a translation. Shapes and angles are preserved. Parallel lines remain parallel. Straight lines remain straight.

Let

```
sc = scale*cos(angle)
ss = scale*sin(angle)
```

```
[u v] = [x y 1] * [ sc -ss
```

```

    ss  sc
    tx  ty]

```

Solve for `sc`, `ss`, `tx`, `ty`.

```
t_lc = cp2tform(input_points,base_points,'linear conformal');
```

The coefficients of the inverse mapping are stored in `t_lc.tdata.Tinv`.

Since linear conformal transformations are a subset of affine transformations, `t_lc.forward_fcn` is `@affine_fwd` and `t_lc.inverse_fcn` is `@affine_inv`.

At least two control-point pairs are needed to solve for the four unknown coefficients.

Affine

In an affine transformation, the x and y dimensions can be scaled or sheared independently and there can be a translation. Parallel lines remain parallel. Straight lines remain straight. Linear conformal transformations are a subset of affine transformations.

For an affine transformation,

$$[u \ v] = [x \ y \ 1] * Tinv$$

`Tinv` is a 3-by-2 matrix. Solve for the six elements of `Tinv`.

```
t_affine = cp2tform(input_points,base_points,'affine');
```

The coefficients of the inverse mapping are stored in `t_affine.tdata.Tinv`.

At least three control-point pairs are needed to solve for the six unknown coefficients.

Projective

In a projective transformation, quadrilaterals map to quadrilaterals. Straight lines remain straight. Affine transformations are a subset of projective transformations.

For a projective transformation

$$[up \ vp \ wp] = [x \ y \ w] * Tinv$$

where

```
u = up/wp  
v = vp/wp
```

Tinv is a 3-by-3 matrix.

Assuming

```
Tinv = [ A D G;  
        B E H;  
        C F I ];  
  
u = (Ax + By + C)/(Gx + Hy + I)  
v = (Dx + Ey + F)/(Gx + Hy + I)
```

Solve for the nine elements of Tinv.

```
t_proj = cp2tform(input_points,base_points,'projective');
```

The coefficients of the inverse mapping are stored in `t_proj.tdata.Tinv`.

At least four control-point pairs are needed to solve for the nine unknown coefficients.

Polynomial

In a polynomial transformation, polynomial functions of x and y determine the mapping.

Second-Order Polynomials

For a second-order polynomial transformation,

```
[u v] = [1 x y x*y x^2 y^2] * Tinv
```

Both u and v are second-order polynomials of x and y . Each second-order polynomial has six terms. To specify all coefficients, Tinv has size 6-by-2.

```
t_poly_ord2 = cp2tform(input_points, base_points,'polynomial');
```

The coefficients of the inverse mapping are stored in `t_poly_ord2.tdata`.

At least six control-point pairs are needed to solve for the 12 unknown coefficients.

Third-Order Polynomials

For a third-order polynomial transformation:

$$[u \ v] = [1 \ x \ y \ x*y \ x^2 \ y^2 \ y*x^2 \ x*y^2 \ x^3 \ y^3] * T_{inv}$$

Both u and v are third-order polynomials of x and y . Each third-order polynomial has ten terms. To specify all coefficients, T_{inv} has size 10-by-2.

```
t_poly_ord3 = cp2tform(input_points, base_points, 'polynomial', 3);
```

The coefficients of the inverse mapping are stored in `t_poly_ord3.tdata`.

At least ten control-point pairs are needed to solve for the 20 unknown coefficients.

Fourth-Order Polynomials

For a fourth-order polynomial transformation:

$$[u \ v] = [1 \ x \ y \ x*y \ x^2 \ y^2 \ y*x^2 \ x*y^2 \ x^3 \ y^3 \ x^3*y \ x^2*y^2 \ x*y^3 \ x^4 \ y^4] * T_{inv}$$

Both u and v are fourth-order polynomials of x and y . Each fourth-order polynomial has 15 terms. To specify all coefficients, T_{inv} has size 15-by-2.

```
t_poly_ord4 = cp2tform(input_points, base_points, 'polynomial', 4);
```

The coefficients of the inverse mapping are stored in `t_poly_ord4.tdata`.

At least 15 control-point pairs are needed to solve for the 30 unknown coefficients.

Piecewise Linear

In a piecewise linear transformation, linear (affine) transformations are applied separately to each triangular region of the image [1].

- 1 Find a Delaunay triangulation of the base control points.
- 2 Using the three vertices of each triangle, infer an affine mapping from base to input coordinates.

Note At least four control-point pairs are needed. Four pairs result in two triangles with distinct mappings.

Local Weighted Mean

For each control point in `base_points`:

- 1 Find the N closest control points.
- 2 Use these N points and their corresponding points in `input_points` to infer a second-order polynomial.
- 3 Calculate the radius of influence of this polynomial as the distance from the center control point to the farthest point used to infer the polynomial (using `base_points`). [2]

Note At least six control-point pairs are needed to solve for the second-order polynomial. Ill-conditioned polynomials might result if too few pairs are used.

Example

```
I = checkerboard;
J = imrotate(I,30);
base_points = [11 11; 41 71];
input_points = [14 44; 70 81];
cpselect(J,I,input_points,base_points);

t = cp2tform(input_points,base_points,'linear conformal');
```

To recover angle and scale,

```
ss = t.tdata.Tinv(2,1); % ss = scale * sin(angle)
sc = t.tdata.Tinv(1,1); % sc = scale * cos(angle)
angle = atan2(ss,sc)*180/pi
scale = sqrt(ss*ss + sc*sc)
```

See Also

`cpcorr`, `cpselect`, `cpstruct2pairs`, `imtransform`

References

- [1] Goshtasby, Ardeshir, "Piecewise linear mapping functions for image registration," *Pattern Recognition*, Vol. 19, 1986, pp. 459-466.
- [2] Goshtasby, Ardeshir, "Image registration by local approximation methods," *Image and Vision Computing*, Vol. 6, 1988, pp. 255-261.

Purpose Tune control-point locations using cross correlation

Syntax `input_points = cpcorr(input_points_in,base_points_in,input,base)`

Description `input_points = cpcorr(input_points_in,base_points_in,input,base)` uses normalized cross-correlation to adjust each pair of control points specified in `input_points_in` and `base_points_in`.

`input_points_in` must be an M-by-2 double matrix containing the coordinates of control points in the input image. `base_points_in` is an M-by-2 double matrix containing the coordinates of control points in the base image.

`cpcorr` returns the adjusted control points in `input_points`, a double matrix the same size as `input_points_in`. If `cpcorr` cannot correlate a pair of control points, `input_points` contains the same coordinates as `input_points_in` for that pair.

`cpcorr` only moves the position of a control point by up to four pixels. Adjusted coordinates are accurate to one-tenth of a pixel. `cpcorr` is designed to get subpixel accuracy from the image content and coarse control-point selection.

Note `input` and `base` images must have the same scale for `cpcorr` to be effective.

`cpcorr` cannot adjust a point if any of the following occur:

- Points are too near the edge of either image.
- Regions of images around points contain `Inf` or `NaN`.
- Region around a point in input image has zero standard deviation.
- Regions of images around points are poorly correlated.

Class Support The images `input` and `base` can be numeric and must contain finite values. The control-point pairs are of class `double`.

Algorithm `cpcorr` uses the following general procedure.

For each control-point pair,

- 1 Extract an 11-by-11 template around the input control point and a 21-by-21 region around the base control point.
- 2 Calculate the normalized cross-correlation of the template with the region.
- 3 Find the absolute peak of the cross-correlation matrix.
- 4 Use the position of the peak to adjust the coordinates of the input control point.

Example

This example uses `cpcorr` to fine-tune control points selected in an image. Note the difference in the values of the `input_points` matrix and the `input_points_adj` matrix.

```
input = imread('onion.png');
base = imread('peppers.png');
input_points = [127 93; 74 59];
base_points = [323 195; 269 161];
input_points_adj = cpcorr(input_points,base_points,...
                          input(:,:,1),base(:,:,1))

input_points_adj =

    127.0000    93.0000
    71.0000    59.6000
```

See Also

`cp2tform`, `cpselect`, `imtransform`, `normxcorr2`

cpselect

Purpose Control Point Selection Tool

Syntax

```
cpselect(input,base)
cpselect(input,base,CPSTRUCT_IN )
cpselect(input,base,xyinput_in,xybase_in)
H = cpselect(input,base,...)
```

Description `cpselect(input,base)` starts the Control Point Selection Tool, a graphical user interface that enables you to select control points in two related images. `input` is the image that needs to be warped to bring it into the coordinate system of the base image. `input` and `base` can be either variables that contain images or strings that identify files containing grayscale images. The Control Point Selection Tool returns the control points in a `CPSTRUCT` structure. (For more information, see “Using the Control Point Selection Tool” in Chapter 6.)

`cpselect(input,base,CPSTRUCT_IN)` starts `cpselect` with an initial set of control points that are stored in `CPSTRUCT_IN`. This syntax allows you to restart `cpselect` with the state of control points previously saved in `CPSTRUCT_IN`.

`cpselect(input,base,xyinput_in,xybase_in)` starts `cpselect` with a set of initial pairs of control points. `xyinput_in` and `xybase_in` are m -by-2 matrices that store the input and base coordinates, respectively.

`H = cpselect(input,base,...)` returns a handle `H` to the tool. You can use the `close(H)` or `H.close` syntax to close the tool from the command line.

Class Support The input images can be of class `uint8`, `uint16`, `double`, or `logical`.

Algorithm `cpselect` uses the following general procedure for control-point prediction.

- 1 Find all valid pairs of control points.
- 2 Infer a spatial transformation between `input` and `base` control points using method that depends on the number of valid pairs, as follows:

2 pairs	Linear conformal
3 pairs	Affine
4 or more pairs	Projective

- 3 Apply spatial transformation to the new point to generate the predicted point.
- 4 Display predicted point.

Notes

Platform Support

cpselect requires Java and is not available on any platform that does not support Java. In addition, cpselect is not available on the Macintosh systems even with Java.

Memory Usage

To increase the amount of memory available to cpselect, you must put a file called 'java.opts' in your start-up directory.

Example

Start tool with saved images.

```
aerial = imread('westconcordaerial.png');  
cpselect(aerial(:, :, 1), 'westconcordorthophoto.png')
```

Start tool with workspace images and points.

```
I = checkerboard;  
J = imrotate(I,30);  
base_points = [11 11; 41 71];  
input_points = [14 44; 70 81];  
cpselect(J,I,input_points,base_points);
```

See Also

cpcorr, cp2tform, cpstruct2pairs, imtransform

cpstruct2pairs

Purpose Convert CPSTRUCT to valid pairs of control points

Syntax [input_points, base_points] = cpstruct2pairs(CPSTRUCT)

Description [input_points, base_points] = cpstruct2pairs(CPSTRUCT) takes a CPSTRUCT (produced by cpselect) and returns the arrays of coordinates of valid control point pairs in input_points and base_points. cpstruct2pairs eliminates unmatched points and predicted points.

Example Start the Control Point Selection Tool, cpselect, with saved images.

```
aerial = imread('westconcordaerial.png');  
cpselect(aerial(:,:,1), 'westconcordorthophoto.png')
```

Using cpselect, pick control points in the images. Select **Save to Workspace** from the **File** menu to save the points to the workspace. On the **Save** dialog box, check the **Structure with all points** check box and clear **Input points of valid pairs** and **Base points of valid pairs**. Click **OK**. Use cpstruct2pairs to extract the input and base points from the CPSTRUCT.


```
[input_points, base_points] = cpstruct2pairs(cpstruct);
```

See Also cp2tform, cpselect, imtransform

Purpose	Compute two-dimensional discrete cosine transform
Syntax	<pre>B = dct2(A) B = dct2(A,m,n) B = dct2(A,[m n])</pre>
Description	<p><code>B = dct2(A)</code> returns the two-dimensional discrete cosine transform of <code>A</code>. The matrix <code>B</code> is the same size as <code>A</code> and contains the discrete cosine transform coefficients $B(k_1, k_2)$.</p> <p><code>B = dct2(A,m,n)</code> or <code>B = dct2(A,[m n])</code> pads the matrix <code>A</code> with 0's to size <code>m</code>-by-<code>n</code> before transforming. If <code>m</code> or <code>n</code> is smaller than the corresponding dimension of <code>A</code>, <code>dct2</code> truncates <code>A</code>.</p>
Class Support	<code>A</code> can be numeric or logical. The returned matrix <code>B</code> is of class <code>double</code> .
Algorithm	<p>The discrete cosine transform (DCT) is closely related to the discrete Fourier transform. It is a separable linear transformation; that is, the two-dimensional transform is equivalent to a one-dimensional DCT performed along a single dimension followed by a one-dimensional DCT in the other dimension. The definition of the two-dimensional DCT for an input image <code>A</code> and output image <code>B</code> is</p> $B_{pq} = \alpha_p \alpha_q \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} A_{mn} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{array}{l} 0 \leq p \leq M-1 \\ 0 \leq q \leq N-1 \end{array}$ $\alpha_p = \begin{cases} 1/\sqrt{M}, & p = 0 \\ \sqrt{2/M}, & 1 \leq p \leq M-1 \end{cases} \quad \alpha_q = \begin{cases} 1/\sqrt{N}, & q = 0 \\ \sqrt{2/N}, & 1 \leq q \leq N-1 \end{cases}$ <p>where M and N are the row and column size of <code>A</code>, respectively. If you apply the DCT to real data, the result is also real. The DCT tends to concentrate information, making it useful for image compression applications.</p> <p>This transform can be inverted using <code>idct2</code>.</p>
Example	The commands below compute the discrete cosine transform for the autumn image. Notice that most of the energy is in the upper left corner.

dct2

```
RGB = imread('autumn.tif');  
I = rgb2gray(RGB);  
J = dct2(I);  
imshow(log(abs(J)),[]), colormap(jet(64)), colorbar
```



Now set values less than magnitude 10 in the DCT matrix to zero, and then reconstruct the image using the inverse DCT function `idct2`.

```
J(abs(J) < 10) = 0;  
K = idct2(J);  
imshow(I)  
figure, imshow(K,[0 255])
```

**See Also**

fft2, idct2, ifft2

References

- [1] Jain, Anil K., *Fundamentals of Digital Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1989, pp. 150-153.
- [2] Pennebaker, William B., and Joan L. Mitchell, *JPEG: Still Image Data Compression Standard*, Van Nostrand Reinhold, 1993.

dctmtx

Purpose	Compute discrete cosine transform matrix
Syntax	<code>D = dctmtx(n)</code>
Description	<code>D = dctmtx(n)</code> returns the n-by-n DCT (discrete cosine transform) matrix. $D*A$ is the DCT of the columns of A and $D'*A$ is the inverse DCT of the columns of A (when A is n-by-n).
Class Support	n is an integer scalar of class <code>double</code> . D is returned as a matrix of class <code>double</code> .
Remarks	<p>If A is square, the two-dimensional DCT of A can be computed as $D*A*D'$. This computation is sometimes faster than using <code>dct2</code>, especially if you are computing a large number of small DCTs, because D needs to be determined only once.</p> <p>For example, in JPEG compression, the DCT of each 8-by-8 block is computed. To perform this computation, use <code>dctmtx</code> to determine D, and then calculate each DCT using $D*A*D'$ (where A is each 8-by-8 block). This is faster than calling <code>dct2</code> for each individual block.</p>
Example	<pre>A = im2double(imread('rice.png')); D = dctmtx(size(A,1)); dct = D*A*D'; figure, imshow(dct)</pre>
See Also	<code>dct2</code>

Purpose Restore image using the blind deconvolution algorithm

Syntax

```
[J,PSF] = deconvblind(I,INITPSF)
[J,PSF] = deconvblind(I,INITPSF,NUMIT)
[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR)
[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR,WEIGHT)
[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR,WEIGHT,READOUT)
[J,PSF] = deconvblind(...,FUN)
```

Description [J,PSF] = deconvblind(I,INITPSF) deconvolves image I using the maximum likelihood algorithm, returning both the deblurred image J and a restored point-spread function PSF. The input array I and your initial guess at the PSF INITPSF can be numeric arrays or cell arrays. (Use cell arrays when you want to be able to perform additional deconvolutions that start where your initial deconvolution finished. The restored PSF is a positive array that is the same size as INITPSF, normalized so its sum adds up to 1.

Note The PSF restoration is affected strongly by the size of the initial guess INITPSF and less by the values it contains. For this reason, specify an array of 1's as your INITPSF.

To improve the restoration, deconvblind supports several optional parameters, described below. Use [] as a placeholder if you do not specify an intermediate parameter.

[J,PSF] = deconvblind(I,INITPSF,NUMIT) specifies the number of iterations (default is 10).

[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR) specifies the threshold deviation of the resulting image from the input image I (in terms of the standard deviation of Poisson noise) below which damping occurs. The iterations are suppressed for the pixels that deviate within the DAMPAR value from their original value. This suppresses the noise generation in such pixels, preserving necessary image details elsewhere. The default value is 0 (no damping).

deconvblind

[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR,WEIGHT) specifies which pixels in the input image I are considered in the restoration. By default, WEIGHT is a unit array, the same size as the input image. You can assign a value between 0.0 and 1.0 to elements in the WEIGHT array. The value of an element in the WEIGHT array determines how much the pixel at the corresponding position in the input image is considered. For example, to exclude a pixel from consideration, assign it a value of 0 in the WEIGHT array. You can adjust the weight value assigned to each pixel according to the amount of flat-field correction.

[J,PSF] = deconvblind(I,INITPSF,NUMIT,DAMPAR,WEIGHT,READOUT), where READOUT is an array (or a value) corresponding to the additive noise (e.g., background, foreground noise) and the variance of the read-out camera noise. READOUT has to be in the units of the image. The default value is 0.

[J,PSF] = deconvblind(...,FUN,P1,P2,...,PN), where FUN is a function describing additional constraints on the PSF. FUN must be a function handle.

FUN is called at the end of each iteration. FUN must accept the PSF as its first argument and can accept additional parameters P1, P2, ..., PN. The FUN function should return one argument, PSF, that is the same size as the original PSF and that satisfies the positivity and normalization constraints.

Note The output image J could exhibit ringing introduced by the discrete Fourier transform used in the algorithm. To reduce the ringing, use `I = edgetaper(I,PSF)` before calling `deconvblind`.

Resuming Deconvolution

You can use `deconvblind` to perform a deconvolution that starts where a previous deconvolution stopped. To use this feature, pass the input image I and the initial guess at the PSF, INITPSF, as cell arrays: {I} and {INITPSF}. When you do, the `deconvblind` function returns the output image J and the restored point-spread function, PSF, as cell arrays, which can then be passed as the

input arrays into the next `deconvblind` call. The output cell array `J` contains four elements:

`J{1}` contains `I`, the original image.

`J{2}` contains the result of the last iteration.

`J{3}` contains the result of the next-to-last iteration.

`J{4}` is an array generated by the iterative algorithm.

Class Support

`I` and `INITPSF` can be `uint8`, `uint16`, `int16`, `double`, or `single`. `DAMPAR` and `READOUT` must have the same class as the input image. Other inputs have to be `double`. The output image `J` (or the first array of the output cell) has the same class as the input image `I`. The output PSF is `double`.

Example

```
I = checkerboard(8);
PSF = fspecial('gaussian',7,10);
V = .0001;
BlurredNoisy = imnoise(imfilter(I,PSF),'gaussian',0,V);
WT = zeros(size(I));
WT(5:end-4,5:end-4) = 1;
INITPSF = ones(size(PSF));
FUN = @(PSF,P1) PSF + P1;
[J P]= deconvblind(BlurredNoisy,INITPSF,20,10*sqrt(V),WT,FUN,0);

subplot(221);imshow(BlurredNoisy);
title('A = Blurred and Noisy');
subplot(222);imshow(PSF,[]);
title('True PSF');
subplot(223);imshow(J);
title('Deblurred Image');
subplot(224);imshow(P,[]);
title('Recovered PSF');
```

See Also

`deconvlucy`, `deconvreg`, `deconvwnr`, `edgetaper`, `function_handle`, `imnoise`, `otf2psf`, `padarray`, `psf2otf`

deconvlucy

Purpose Restore image using the Lucy-Richardson algorithm

Syntax

```
J = deconvlucy(I,PSF)
J = deconvlucy(I,PSF,NUMIT)
J = deconvlucy(I,PSF,NUMIT,DAMPAR)
J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT)
J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT,READOUT)
J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT,READOUT,SUBSMPL)
```

Description `J = deconvlucy(I,PSF)` restores image `I` that was degraded by convolution with a point-spread function `PSF` and possibly by additive noise. The algorithm is based on maximizing the likelihood of the resulting image `J`'s being an instance of the original image `I` under Poisson statistics. The input array `I` can be a numeric array (the blurred image) or a cell array.

If `I` is a cell array, it can contain a single numerical array (the blurred image) or it can be the output from a previous run of `deconvlucy`.

When you pass a cell array to `deconvlucy` as input, it returns a 1-by-4 cell array `J`, where

- `J{1}` contains `I`, the original image.
- `J{2}` contains the result of the last iteration.
- `J{3}` contains the result of the next-to-last iteration.
- `J{4}` is an array generated by the iterative algorithm.

To improve the restoration, `deconvlucy` supports several optional parameters. Use `[]` as a placeholder if you do not specify an intermediate parameter.

`J = deconvlucy(I,PSF,NUMIT)` specifies the number of iterations the `deconvlucy` function performs. If this value is not specified, the default is 10.

`J = deconvlucy(I,PSF,NUMIT,DAMPAR)` specifies the threshold deviation of the resulting image from the image `I` (in terms of the standard deviation of Poisson noise) below which damping occurs. Iterations are suppressed for pixels that deviate beyond the `DAMPAR` value from their original value. This suppresses the noise generation in such pixels, preserving necessary image details elsewhere. The default value is 0 (no damping).

`J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT)` specifies the weight to be assigned to each pixel to reflect its recording quality in the camera. A bad pixel is excluded from the solution by assigning it zero weight value. Instead of giving a weight of unity for good pixels, you can adjust their weight according to the amount of flat-field correction. The default is a unit array of the same size as input image `I`.

`J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT,READOUT)` specifies a value corresponding to the additive noise (e.g., background, foreground noise) and the variance of the readout camera noise. `READOUT` has to be in the units of the image. The default value is 0.

`J = deconvlucy(I,PSF,NUMIT,DAMPAR,WEIGHT,READOUT,SUBSMPL)`, where `SUBSMPL` denotes subsampling and is used when the PSF is given on a grid that is `SUBSMPL` times finer than the image. The default value is 1.

Note The output image `J` could exhibit ringing introduced by the discrete Fourier transform used in the algorithm. To reduce the ringing, use `I = edgetaper(I,PSF)` before calling `deconvlucy`.

Class Support

`I` and `PSF` can be `uint8`, `uint16`, `int16`, `double`, or `single`. `DAMPAR` and `READOUT` must have the same class as the input image. Other inputs have to be `double`. The output image `J` (or the first array of the output cell) has the same class as the input image `I`.

Example

```
I = checkerboard(8);
PSF = fspecial('gaussian',7,10);
V = .0001;
BlurredNoisy = imnoise(imfilter(I,PSF),'gaussian',0,V);
WT = zeros(size(I));
WT(5:end-4,5:end-4) = 1;
J1 = deconvlucy(BlurredNoisy,PSF);
J2 = deconvlucy(BlurredNoisy,PSF,20,sqrt(V));
J3 = deconvlucy(BlurredNoisy,PSF,20,sqrt(V),WT);

subplot(221);imshow(BlurredNoisy);
title('A = Blurred and Noisy');
```

deconvlucy

```
subplot(222);imshow(J1);  
title('deconvlucy(A,PSF)');  
subplot(223);imshow(J2);  
title('deconvlucy(A,PSF,NI,DP)');  
subplot(224);imshow(J3);  
title('deconvlucy(A,PSF,NI,DP,WT)');
```

See Also

deconvblind, deconvreg, deconvwnr, otf2psf, padarray, psf2otf

Purpose	Restore image using a regularized filter
Syntax	<pre>J = deconvreg(I,PSF) J = deconvreg(I,PSF,NOISEPOWER) J = deconvreg(I,PSF,NOISEPOWER,LRANGE) J = deconvreg(I,PSF,NOISEPOWER,LRANGE,REGOP) [J, LAGRA] = deconvreg(I,PSF,...)</pre>
Description	<p>$J = \text{deconvreg}(I, \text{PSF})$ restores image I that was degraded by convolution with a point-spread function PSF and possibly by additive noise. The algorithm is a constrained optimum in a sense of least square error between the estimated and the true images under the requirement of preserving image smoothness.</p> <p>$J = \text{deconvreg}(I, \text{PSF}, \text{NOISEPOWER})$, where NOISEPOWER is the additive noise power. The default value is 0.</p> <p>$J = \text{deconvreg}(I, \text{PSF}, \text{NOISEPOWER}, \text{LRANGE})$, where LRANGE is a vector specifying range where the search for the optimal solution is performed. The algorithm finds an optimal Lagrange multiplier LAGRA within the LRANGE range. If LRANGE is a scalar, the algorithm assumes that LAGRA is given and equal to LRANGE; the NP value is then ignored. The default range is between [1e-9 and 1e9].</p> <p>$J = \text{deconvreg}(I, \text{PSF}, \text{NOISEPOWER}, \text{LRANGE}, \text{REGOP})$, where REGOP is the regularization operator to constrain the deconvolution. The default regularization operator is the Laplacian operator, to retain the image smoothness. The REGOP array dimensions must not exceed the image dimensions; any nonsingleton dimensions must correspond to the nonsingleton dimensions of PSF.</p> <p>$[J, \text{LAGRA}] = \text{deconvreg}(I, \text{PSF}, \dots)$ outputs the value of the Lagrange multiplier LAGRA in addition to the restored image J.</p>

Note The output image J could exhibit ringing introduced by the discrete Fourier transform used in the algorithm. To reduce the ringing, process the image with the `edgetaper` function prior to calling the `deconvreg` function. For example, $I = \text{edgetaper}(I, \text{PSF})$.

deconvreg

Class Support I can be of class uint8, uint16, int16, double, or single. Other inputs have to be of class double. J has the same class as I.

Example

```
I = checkerboard(8);
PSF = fspecial('gaussian',7,10);
V = .01;
BlurredNoisy = imnoise(imfilter(I,PSF),'gaussian',0,V);
NOISEPOWER = V*prod(size(I));
[J LAGRA] = deconvreg(BlurredNoisy,PSF,NOISEPOWER);

subplot(221); imshow(BlurredNoisy);
title('A = Blurred and Noisy');
subplot(222); imshow(J);
title('[J LAGRA] = deconvreg(A,PSF,NP)');
subplot(223); imshow(deconvreg(BlurredNoisy,PSF,[],LAGRA/10));
title('deconvreg(A,PSF,[],0.1*LAGRA)');
subplot(224); imshow(deconvreg(BlurredNoisy,PSF,[],LAGRA*10));
title('deconvreg(A,PSF,[],10*LAGRA)');
```

See Also deconvblind, deconvlucy, deconvwnr, otf2psf, padarray, psf2otf

Purpose Restore image using the Wiener filter

Syntax

```
J = deconvwnr(I,PSF)
J = deconvwnr(I,PSF,NSR)
J = deconvwnr(I,PSF,NCORR,ICORR)
```

Description `J = deconvwnr(I,PSF)` restores image `I` that was degraded by convolution with a point-spread function `PSF` and possibly by additive noise. The algorithm is optimal in a sense of least mean square error between the estimated and the true image, and uses the correlation matrixes of image and noise. In the absence of noise, the Wiener filter reduces to the ideal inverse filter.

`J = deconvwnr(I,PSF,NSR)`, where `NSR` is the noise-to-signal power ratio. `NSR` could be a scalar or an array of the same size as `I`. The default value is 0.

`J = deconvwnr(I,PSF,NCORR,ICORR)`, where `NCORR` and `ICORR` are the autocorrelation functions of the noise and the original image, respectively. `NCORR` and `ICORR` can be of any size or dimension not exceeding the original image. An `N`-dimensional `NCORR` or `ICORR` array corresponds to the autocorrelation within each dimension. A vector `NCORR` or `ICORR` represents an autocorrelation function in the first dimension if `PSF` is a vector. If `PSF` is an array, the 1-D autocorrelation function is extrapolated by symmetry to all nonsingleton dimensions of `PSF`. A scalar `NCORR` or `ICORR` represents the power of the noise or the image.

Note The output image `J` could exhibit ringing introduced by the discrete Fourier transform used in the algorithm. To reduce the ringing, process the image with the `edgetaper` function prior to calling the `deconvwnr` function. For example, `I = edgetaper(I,PSF)`

Class Support `I` can be of class `uint8`, `uint16`, `int16`, `double`, or `single`. Other inputs have to be of class `double`. `J` has the same class as `I`.

Example

```
I = checkerboard(8);
noise = 0.1*randn(size(I));
PSF = fspecial('motion',21,11);
Blurred = imfilter(I,PSF,'circular');
```

```
BlurredNoisy = im2uint8(Blurred + noise);

NSR = sum(noise(:).^2)/sum(I(:).^2);% noise-to-power ratio

NP = abs(fftn(noise)).^2;% noise power
NPOW = sum(NP(:))/prod(size(noise));
NCORR = fftshift(real(ifftn(NP)));% noise autocorrelation
function, centered

IP = abs(fftn(I)).^2;% original image power
IPOW = sum(IP(:))/prod(size(I));
ICORR = fftshift(real(ifftn(IP)));% image autocorrelation
function, centered
ICORR1 = ICORR(:,ceil(size(I,1)/2));

NSR = NPOW/IPOW;
subplot(221);imshow(BlurredNoisy,[]);
title('A = Blurred and Noisy');
subplot(222);imshow(deconvwnr(BlurredNoisy,PSF,NSR),[]);
title('deconvwnr(A,PSF,NSR)');
subplot(223);imshow(deconvwnr(BlurredNoisy,PSF,NCORR,ICORR),[])
;
title('deconvwnr(A,PSF,NCORR,ICORR)');
subplot(224);imshow(deconvwnr(BlurredNoisy,PSF,NPOW,ICORR1),[])
;
title('deconvwnr(A,PSF,NPOW,ICORR_1_D)');
```

See Also

deconvblind, deconvlucy, deconvreg, otf2psf, padarray, psf2otf

- Purpose** Apply a decorrelation stretch to a multichannel image
- Syntax**
S = decorrstretch(I)
S = decorrstretch(I,TOL)
- Description**
S = decorrstretch(I) applies a decorrelation stretch to a multichannel image I and returns the result in S. S has the same size and class as I. The mean and variance in each band are the same as in I.

S = decorrstretch(I,TOL) applies a contrast following the decorrelation stretch. The contrast stretch is controlled by TOL:
- TOL = [LOW_FRACT HIGH_FRACT] specifies the fraction of the image to saturate at low and high intensities.
 - If TOL is a scalar, LOW_FRACT = TOL, and HIGH_FRACT = 1 - TOL, which saturates equal fractions at low and high intensities.
- Notes**
The decorrelation stretch is normally applied to three band images (ordinary RGB images or RGB multispectral composite images), but decorrstretch works on an arbitrary number of bands.

The primary purpose of decorrelation stretch is visual enhancement. Small adjustments to TOL can strongly affect the visual appearance of the output.
- Class Support** The input image must be of class uint8, uint16, int16, single, or double.
- Example**

```
[X, map] = imread('forest.tif');  
S = decorrstretch(ind2rgb(X, map),'tol',0.01);  
figure, imshow(X,map)  
figure, imshow(S)
```
- See Also** imadjust, stretchlim

dicomanon

Purpose Anonymize a DICOM file

Syntax

```
dicomanon(file_in, file_out)
dicomanon(..., 'keep', FIELDS)
dicomanon(..., 'update', ATTRS)
```

Description `dicomanon(file_in, file_out)` removes confidential medical information from the DICOM file `file_in` and creates a new file `file_out` with the modified values. Image data and other attributes are unmodified.

`dicomanon(..., 'keep', FIELDS)` modifies all of the confidential data except for those listed in `FIELDS`, which is a cell array of field names. This syntax is useful for keeping metadata that does not uniquely identify the patient but is useful for diagnostic purposes (e.g., `PatientAge`, `PatientSex`, etc.).

Note Keeping certain fields might compromise patient confidentiality.

`dicomanon(..., 'update', ATTRS)` modifies the confidential data and updates particular confidential data. `ATTRS` is a structure. The field names of `ATTRS` are the attributes to preserve, and the structure values are the attribute values. Use this syntax to preserve the `Study/Series/Image` hierarchy or to replace a specific value with a more generic property (e.g., remove `PatientBirthDate` but keep a computed `PatientAge`).

For information about the fields that will be modified or removed, see DICOM Supplement 55 from <http://medical.nema.org/>.

Examples Remove all confidential metadata from a file.

```
dicomanon('patient.dcm', 'anonymized.dcm')
```

Create a training file.

```
dicomanon('tumor.dcm', 'tumor_anon.dcm', ...
          'keep', {'PatientAge', 'PatientSex', 'StudyDescription'})
```

Anonymize a series of images, keeping the hierarchy.

```
values.StudyInstanceUID = dicomuid;
```

```
values.SeriesInstanceseriesUID = dicomuid;

d = dir('*.dcm');
for p = 1:numel(d)
    dicomanon(d(p).name, sprintf('anon%d.dcm', p), ...
        'update', values)
end
```

See Also

dicominfo, dicomwrite

dicomdict

Purpose Get or set the active DICOM data dictionary

Syntax

```
dicomdict('set', dictionary)
dictionary = dicomdict('get')
dictionary = dicomdict('factory')
```

Description `dicomdict('set', dictionary)` sets the Digital Imaging and Communications in Medicine (DICOM) data dictionary to the value stored in `dictionary`, a string containing the filename of the dictionary. DICOM-related functions use this dictionary by default, unless a different dictionary is provided at the command line.

`dictionary = dicomdict('get')` returns a string containing the filename of the stored DICOM data dictionary.

`dictionary = dicomdict('factory')` resets the DICOM data dictionary to its default startup value.

Note The default data dictionary is a MAT-file, `dicom-dict.mat`. The toolbox also includes a text version of this default data dictionary, `dicom-dict.txt`. If you want to create your own DICOM data dictionary, open the `dicom-dict.txt` file in a text editor, modify it, and save it under another name.

Example

```
dictionary = dicomdict('get')

dictionary =

dicom-dict.mat
```

See Also `dicominfo`, `dicomread`, `dicomwrite`

Purpose	Read metadata from a DICOM message
Syntax	<pre>info = dicominfo(filename) info = dicominfo(filename, 'dictionary', D)</pre>
Description	<p><code>info = dicominfo(filename)</code> reads the metadata from the compliant Digital Imaging and Communications in Medicine (DICOM) file specified in the string <code>filename</code>.</p> <p><code>info = dicominfo(filename, 'dictionary', D)</code> uses the data dictionary file given in the string <code>D</code> to read the DICOM message. The file in <code>D</code> must be on the MATLAB search path. The default file is <code>dicom-dict.mat</code>.</p>
Examples	<pre>info = dicominfo('CT-MON02-16-ankle.dcm') info = Filename: [1x62 char] FileModDate: '18-Dec-2000 11:06:43' FileSize: 525436 Format: 'DICOM' FormatVersion: 3 Width: 512 Height: 512 BitDepth: 16 ColorType: 'grayscale' SelectedFrames: [] FileStruct: [1x1 struct] StartOfPixelData: 1140 FileMetaInformationGroupLength: 192 FileMetaInformationVersion: [2x1 uint8] MediaStorageSOPClassUID: '1.2.840.10008.5.1.4.1.1.7' . . .</pre>
See Also	<code>dicomdict</code> , <code>dicomread</code> , <code>dicomwrite</code> , <code>dicomuid</code>

dicomread

Purpose Read a DICOM image

Syntax

```
X = dicomread(filename)
X = dicomread(info)
[X,map] = dicomread(...)
[X,map,alpha] = dicomread(...)
[X,map,alpha,overlays] = dicomread(...)
[...] = dicomread(filename,'frames',v)
```

Description `X = dicomread(filename)` reads the image data from the compliant Digital Imaging and Communications in Medicine (DICOM) file `filename`. For single-frame grayscale images, `X` is an `M`-by-`N` array. For single-frame true-color images, `X` is an `M`-by-`N`-by-3 array. Multiframe images are always 4-D arrays.

`X = dicomread(info)` reads the image data from the message referenced in the DICOM metadata structure `info`. The `info` structure is produced by the `dicominfo` function.

`[X,map] = dicomread(...)` returns the image `X` and the colormap `map`. If `X` is a grayscale or true-color image, `map` is empty.

`[X,map,alpha] = dicomread(...)` returns the image `X`, the colormap `map`, and an alpha channel matrix for `X`. The values of `alpha` are 0 if the pixel is opaque; otherwise they are row indices into `map`. The RGB value in `map` should be substituted for the value in `X` to use alpha. `alpha` has the same height and width as `X` and is 4-D for a multiframe image.

`[X,map,alpha,overlays] = dicomread(...)` returns the image `X`, the colormap `map`, an alpha channel matrix for `X`, and any overlays from the DICOM file. Each overlay is a 1-bit black and white image with the same height and width as `X`. If multiple overlays are present in the file, `overlays` is a 4-D multiframe image. If no overlays are in the file, `overlays` is empty.

`[...] = dicomread(filename,'frames',v)` reads only the frames in the vector `v` from the image. `v` must be an integer scalar, a vector of integers, or the string 'all'. The default value is 'all'.

Class Support X can be uint8, int8, uint16, or int16. map must be double. alpha has the same size and type as X. overlays is a logical array.

Examples Retrieve the data matrix X and colormap matrix map and create a montage.

```
[X, map] = dicomread('US-PAL-8-10x-echo.dcm');  
montage(X, map);
```

Call dicomread with the information retrieved from the DICOM file using dicominfo. Because a DICOM image is a 16-bit image, the example uses the imshow autoscaling syntax to display the image.

```
info = dicominfo('CT-MON02-16-ankle.dcm');  
Y = dicomread(info);  
figure, imshow(Y, 'DisplayRange', []);
```

See Also dicomdict, dicominfo, dicomwrite

dicomuid

Purpose Generate a DICOM unique identifier

Syntax UID = dicomuid

Description UID = dicomuid creates a string UID containing a new DICOM unique identifier.

Multiple calls to dicomuid produce globally unique values. Two calls to dicomuid always return different values.

See Also dicominfo, dicomwrite

Purpose Write images as DICOM files

Syntax

```
dicomwrite(X, filename)
dicomwrite(X, map, filename)
dicomwrite(...,param1,value1,param2,value2,...)
dicomwrite(...,'ObjectType',IOD,...)
dicomwrite(...,'SOPClassUID',UID,...)
dicomwrite(...,meta_struct,...)
dicomwrite(...,info,...)
status = dicomwrite(...)
```

Description `dicomwrite(X, filename)` writes the binary, grayscale, or truecolor image `X` to the file `filename`, where `filename` is a string specifying the name of the Digital Imaging and Communications in Medicine (DICOM) file to create.

`dicomwrite(X,map,filename)` writes the indexed image `X` with colormap `map`.

`dicomwrite(...,param1,value1,param2,value2,...)` specifies additional metadata to write to the DICOM file or parameters that affect how the file is written. `param1` is a string containing the metadata attribute name or a `dicomwrite`-specific option. `value1` is the corresponding value for the attribute or option.

To find a list of the DICOM attributes that you can specify, see the data dictionary file, `dicom-dict.txt`, included with the Image Processing Toolbox.

The following table lists the options that you can specify, in alphabetical order. Default values are enclosed in braces ({}).

Option Name	Description
'CompressionMode'	String specifying the type of compression to use when storing the image. Possible values: { 'None' } 'JPEG lossless' 'JPEG lossy' 'RLE'
'CreateMode'	Specifies the method used for creating the data to put in the new file. Possible values: { 'Create' } — Verify input values and generate missing data values. 'Copy' — Copy all values from the input and do not generate missing values.
'Dictionary'	String specifying the name of a DICOM data dictionary.
'Endian'	String specifying the byte ordering of the file. 'Big' { 'Little' } Note: If VR is set to 'Explicit', 'Endian' must be 'Big'. dicomwrite ignores this value if CompressionMode or TransferSyntax is set.
'TransferSyntax'	A DICOM UID specifying the Endian and VR mode. Note: If If specified, dicomwrite ignores any values specified for the Endian, VR, and CompressionMode options. The TransferSyntax value encodes values for these options.

Option Name	Description
'VR'	String specifying whether the two-letter value representation (VR) code should be written to the file. 'explicit' — Write VR to file. { 'implicit' } — Infer from data dictionary. Note: If you specify the 'Endian' value 'Big', you must specify 'Explicit'.
'WritePrivate'	Logical value indicating whether private data should be written to the file. Possible values: true — Write private data to file. {false} — Do not write private data.

`dicomwrite(..., 'ObjectType', IOD, ...)` writes a file containing the necessary metadata for a particular type of DICOM Information Object (IOD). Supported IODs are

- 'Secondary Capture Image Storage' (default)
- 'CT Image Storage'
- 'MR Image Storage'

`dicomwrite(..., 'SOPClassUID', UID, ...)` provides an alternate method for specifying the IOD to create. UID is the DICOM unique identifier corresponding to one of the IODs listed above.

`dicomwrite(..., meta_struct, ...)` specifies optional metadata or file options in structure `meta_struct`. The names of fields in `meta_struct` must be the names of DICOM file attributes or options. The value of a field is the value you want to assign to the attribute or option.

`dicomwrite(..., info, ...)` specifies metadata in the metadata structure `info`, which is produced by the `dicominfo` function. For more information about this structure, see `dicominfo`.

`status = dicomwrite(...)` returns information about the metadata and the descriptions used to generate the DICOM file. This syntax can be useful when

dicomwrite

you specify an info structure that was created by `dicominfo` to the `dicomwrite` function. An info structure can contain many fields. If no metadata was specified, `dicomwrite` returns an empty matrix (`[]`).

The structure returned by `dicomwrite` contains these fields:

Field	Description
'BadAttribute'	The attribute's internal description is bad. It might be missing from the data dictionary or have incorrect data in its description.
'MissingCondition'	The attribute is conditional but no condition has been provided for when to use it.
'MissingData'	No data was provided for an attribute that must appear in the file.
'SuspectAttribute'	Data in the attribute does not match a list of enumerated values in the DICOM specification.

Example

This example reads a CT image from the sample DICOM file included with the toolbox. The example then writes the CT image to a file, creating a "secondary capture" image.

```
X = dicomread('CT-MONO2-16-ankle.dcm');  
dicomwrite(X, 'sc_file.dcm');
```

This example writes the CT image, `X`, to a DICOM file along with its metadata. You use the `dicominfo` function to retrieve metadata from a DICOM file.

```
metadata = dicominfo('CT-MONO2-16-ankle.dcm');  
dicomwrite(X, 'ct_file.dcm', metadata);
```

Copy all metadata from one file to another.

```
dicomwrite(X, 'ct_copy.dcm', metadata, 'CreateMode', 'copy');
```

See Also

`dicomdict`, `dicominfo`, `dicomread`, `dicomuid`

Purpose	Convert an image, increasing apparent color resolution by dithering
Syntax	<pre>X = dither(RGB, map) BW = dither(I)</pre>
Description	<p><code>X = dither(RGB, map)</code> creates an indexed image approximation of the RGB image in the array <code>RGB</code> by dithering the colors in the colormap <code>map</code>. The colormap cannot have more than 65,536 colors.</p> <p><code>X = dither(RGB, map, Qm, Qe)</code> creates an indexed image from <code>RGB</code>, specifying the parameters <code>Qm</code> and <code>Qe</code>. <code>Qm</code> specifies the number of quantization bits to use along each color axis for the inverse color map, and <code>Qe</code> specifies the number of quantization bits to use for the color space error calculations. If <code>Qe < Qm</code>, dithering cannot be performed, and an undithered indexed image is returned in <code>X</code>. If you omit these parameters, <code>dither</code> uses the default values <code>Qm = 5</code>, <code>Qe = 8</code>.</p> <p><code>BW = dither(I)</code> converts the intensity image in the matrix <code>I</code> to the binary (black and white) image <code>BW</code> by dithering.</p>
Class Support	<code>RGB</code> can be <code>uint8</code> , <code>uint16</code> , <code>single</code> , or <code>double</code> . <code>I</code> can be <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> . All other input arguments must be <code>double</code> . <code>BW</code> is logical. <code>X</code> is <code>uint8</code> , if it is an indexed image with 256 or fewer colors; otherwise, it is <code>uint16</code> .
Algorithm	<code>dither</code> increases the apparent color resolution of an image by applying Floyd-Steinberg's error diffusion dither algorithm.
Example	Convert intensity image to binary using dithering. <pre>I = imread('cameraman.tif'); BW = dither(I); imshow(I), figure, imshow(BW)</pre>
See Also	<code>rgb2ind</code>
References	[1] Floyd, R. W., and L. Steinberg, "An Adaptive Algorithm for Spatial Gray Scale," <i>International Symposium Digest of Technical Papers</i> , Society for Information Displays, 1975, p. 36.

[2] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 469-476.

Purpose

Convert data to double precision

`double` is a MATLAB built-in function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

edge

Purpose

Find edges in an intensity image

Syntax

```
BW = edge(I, 'sobel')
BW = edge(I, 'sobel', thresh)
BW = edge(I, 'sobel', thresh, direction)
[BW, thresh] = edge(I, 'sobel', ...)
```



```
BW = edge(I, 'prewitt')
BW = edge(I, 'prewitt', thresh)
BW = edge(I, 'prewitt', thresh, direction)
[BW, thresh] = edge(I, 'prewitt', ...)
```



```
BW = edge(I, 'roberts')
BW = edge(I, 'roberts', thresh)
[BW, thresh] = edge(I, 'roberts', ...)
```



```
BW = edge(I, 'log')
BW = edge(I, 'log', thresh)
BW = edge(I, 'log', thresh, sigma)
[BW, threshold] = edge(I, 'log', ...)
```



```
BW = edge(I, 'zerocross', thresh, h)
[BW, thresh] = edge(I, 'zerocross', ...)
```



```
BW = edge(I, 'canny')
BW = edge(I, 'canny', thresh)
BW = edge(I, 'canny', thresh, sigma)
[BW, threshold] = edge(I, 'canny', ...)
```

Description

edge takes an intensity image I as its input, and returns a binary image BW of the same size as I, with 1's where the function finds edges in I and 0's elsewhere.

edge supports six different edge-finding methods:

- The Sobel method finds edges using the Sobel approximation to the derivative. It returns edges at those points where the gradient of I is maximum.

- The Prewitt method finds edges using the Prewitt approximation to the derivative. It returns edges at those points where the gradient of I is maximum.
- The Roberts method finds edges using the Roberts approximation to the derivative. It returns edges at those points where the gradient of I is maximum.
- The Laplacian of Gaussian method finds edges by looking for zero crossings after filtering I with a Laplacian of Gaussian filter.
- The zero-cross method finds edges by looking for zero crossings after filtering I with a filter you specify.
- The Canny method finds edges by looking for local maxima of the gradient of I. The gradient is calculated using the derivative of a Gaussian filter. The method uses two thresholds, to detect strong and weak edges, and includes the weak edges in the output only if they are connected to strong edges. This method is therefore less likely than the others to be fooled by noise, and more likely to detect true weak edges.

The parameters you can supply differ depending on the method you specify. If you do not specify a method, edge uses the Sobel method.

Sobel Method

`BW = edge(I, 'sobel')` specifies the Sobel method.

`BW = edge(I, 'sobel', thresh)` specifies the sensitivity threshold for the Sobel method. edge ignores all edges that are not stronger than thresh. If you do not specify thresh, or if thresh is empty (`[]`), edge chooses the value automatically.

`BW = edge(I, 'sobel', thresh, direction)` specifies the direction of detection for the Sobel method. direction is a string specifying whether to look for 'horizontal' or 'vertical' edges or 'both' (the default).

`BW = edge(I, 'sobel', ..., options)` provides an optional string input. String 'nothinning' speeds up the operation of the algorithm by skipping the additional edge thinning stage. By default, or when 'thinning' string is specified, the algorithm applies edge thinning.

`[BW, thresh] = edge(I, 'sobel', ...)` returns the threshold value.

`[BW,thresh,gv,gh] = edge(I, 'sobel', ...)` returns vertical and horizontal edge responses to Sobel gradient operators. You can also use the following expressions to obtain gradient responses:

```
if ~(isa(I,'double') || isa(I,'single')); I = im2single(I); end
gh = imfilter(I,fspecial('sobel') /8,'replicate'); and
gv = imfilter(I,fspecial('sobel')'/8,'replicate');
```

Prewitt Method

`BW = edge(I, 'prewitt')` specifies the Prewitt method.

`BW = edge(I, 'prewitt', thresh)` specifies the sensitivity threshold for the Prewitt method. `edge` ignores all edges that are not stronger than `thresh`. If you do not specify `thresh`, or if `thresh` is empty (`[]`), `edge` chooses the value automatically.

`BW = edge(I, 'prewitt', thresh, direction)` specifies the direction of detection for the Prewitt method. `direction` is a string specifying whether to look for 'horizontal' or 'vertical' edges or 'both' (the default).

`[BW,thresh] = edge(I, 'prewitt', ...)` returns the threshold value.

Roberts Method

`BW = edge(I, 'roberts')` specifies the Roberts method.

`BW = edge(I, 'roberts', thresh)` specifies the sensitivity threshold for the Roberts method. `edge` ignores all edges that are not stronger than `thresh`. If you do not specify `thresh`, or if `thresh` is empty (`[]`), `edge` chooses the value automatically.

`[BW,thresh] = edge(I, 'roberts', ...)` returns the threshold value.

Laplacian of Gaussian Method

`BW = edge(I, 'log')` specifies the Laplacian of Gaussian method.

`BW = edge(I, 'log', thresh)` specifies the sensitivity threshold for the Laplacian of Gaussian method. `edge` ignores all edges that are not stronger than `thresh`. If you do not specify `thresh`, or if `thresh` is empty (`[]`), `edge` chooses the value automatically.

`BW = edge(I, 'log', thresh, sigma)` specifies the Laplacian of Gaussian method, using `sigma` as the standard deviation of the LoG filter. The default `sigma` is 2; the size of the filter is `n-by-n`, where `n = ceil(sigma*3)*2+1`.

`[BW, thresh] = edge(I, 'log', ...)` returns the threshold value.

Zero-Cross Method

`BW = edge(I, 'zerocross', thresh, h)` specifies the zero-cross method, using the filter `h`. `thresh` is the sensitivity threshold; if the argument is empty (`[]`), `edge` chooses the sensitivity threshold automatically.

`[BW, thresh] = edge(I, 'zerocross', ...)` returns the threshold value.

Canny Method

`BW = edge(I, 'canny')` specifies the Canny method.

`BW = edge(I, 'canny', thresh)` specifies sensitivity thresholds for the Canny method. `thresh` is a two-element vector in which the first element is the low threshold, and the second element is the high threshold. If you specify a scalar for `thresh`, this value is used for the high threshold and `0.4*thresh` is used for the low threshold. If you do not specify `thresh`, or if `thresh` is empty (`[]`), `edge` chooses low and high values automatically.

`BW = edge(I, 'canny', thresh, sigma)` specifies the Canny method, using `sigma` as the standard deviation of the Gaussian filter. The default `sigma` is 1; the size of the filter is chosen automatically, based on `sigma`.

`[BW, thresh] = edge(I, 'canny', ...)` returns the threshold values as a two-element vector.

Class Support

`I` is a nonsparse numeric array. `BW` is of class `logical`.

Remarks


For the 'log' and 'zerocross' methods, if you specify a threshold of 0, the output image has closed contours, because it includes all the zero crossings in the input image.

Example


Find the edges of an image using the Prewitt and Canny methods.

```
I = imread('circuit.tif');
BW1 = edge(I, 'prewitt');
BW2 = edge(I, 'canny');
```

```
imshow(BW1);  
figure, imshow(BW2)
```



Prewitt Filter



Canny Filter

See Also

`fspecial`

References

- [1] Canny, John, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-8, No. 6, 1986, pp. 679-698.
- [2] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 478-488.
- [3] Parker, James R., *Algorithms for Image Processing and Computer Vision*, New York, John Wiley & Sons, Inc., 1997, pp. 23-29.

Purpose Taper the discontinuities along the image edges

Syntax `J = edgetaper(I,PSF)`

Description `J = edgetaper(I,PSF)` blurs the edges of the input image `I` using the point spread function `PSF`. The output image `J` is the weighted sum of the original image `I` and its blurred version. The weighting array, determined by the autocorrelation function of `PSF`, makes `J` equal to `I` in its central region, and equal to the blurred version of `I` near the edges.

The `edgetaper` function reduces the ringing effect in image deblurring methods that use the discrete Fourier transform, such as `deconvwnr`, `deconvreg`, and `deconvlucy`.

Note The size of the `PSF` cannot exceed half of the image size in any dimension.

Class Support `I` and `PSF` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. `J` is of the same class as `I`.

Example

```
original = imread('cameraman.tif');
PSF = fspecial('gaussian',60,10);
edgesTapered = edgetaper(original,PSF);
figure, imshow(original,[]);
figure, imshow(edgesTapered,[]);
```

See Also `deconvlucy`, `deconvreg`, `deconvwnr`, `otf2psf`, `padarray`, `psf2otf`

entropy

Purpose Entropy of an intensity image

Syntax `E = entropy(I)`

Description `E = entropy(I)` returns `E`, a scalar value representing the entropy of intensity image `I`. Entropy is a statistical measure of randomness that can be used to characterize the texture of the input image. Entropy is defined as

$$-\sum(p.\log(p))$$

where `p` contains the histogram counts returned from `imhist`. By default, `entropy` uses two bins for logical arrays and 256 bins for `uint8`, `uint16`, or `double` arrays.

`I` can be a multidimensional image. If `I` has more than two dimensions, the `entropy` function treats it as a multidimensional intensity image and not as an RGB image.

Class Support `I` can be `logical`, `uint8`, `uint16`, or `double` and must be real, nonempty, and nonsparse. `E` is `double`.

Notes `entropy` converts any class other than `logical` to `uint8` for the histogram count calculation so that the pixel values are discrete and directly correspond to a bin value.

Example

```
I = imread('circuit.tif');  
J = entropy(I)
```

See Also `imhist`, `entropyfilt`, `blkproc`

References [1] Gonzalez, R.C., R.E. Woods, S.L. Eddins, *Digital Image Processing Using MATLAB*, New Jersey, Prentice Hall, 2003, Chapter 11.

Purpose Local entropy of an intensity image

Syntax `J = entropyfilt(I)`
`J = entropyfilt(I,NHOOD)`

Description `J = entropyfilt(I)` returns the array `J`, where each output pixel contains the entropy value of the 9-by-9 neighborhood around the corresponding pixel in the input image `I`. `I` can have any dimension. If `I` has more than two dimensions, the `entropyfilt` function treats it as a multidimensional intensity image and not as a truecolor (RGB) image. The output image `J` is the same size as the input image `I`.

For pixels on the borders of `I`, `entropyfilt` uses symmetric padding. In symmetric padding, the values of padding pixels are a mirror reflection of the border pixels in `I`.

`J = entropyfilt(I,NHOOD)` performs entropy filtering of the input image `I` where you specify the neighborhood in `NHOOD`. `NHOOD` is a multidimensional array of zeros and ones where the nonzero elements specify the neighbors. `NHOOD`'s size must be odd in each dimension.

By default, `entropyfilt` uses the neighborhood `true(9)`. `entropyfilt` determines the center element of the neighborhood by `floor((size(NHOOD) + 1)/2)`. To specify neighborhoods of various shapes, such as a disk, use the `strel` function to create a structuring element object and then use the `getnhood` function to extract the neighborhood from the structuring element object.

Class Support `I` can be `logical`, `uint8`, `uint16`, or `double`, and must be real and nonsparse. `NHOOD` can be `logical` or `numeric` and must contain zeros or ones. The output array `J` is of class `double`.

`entropyfilt` converts any class other than `logical` to `uint8` for the histogram count calculation so that the pixel values are discrete and directly correspond to a bin value.

Example

```
I = imread('circuit.tif');
J = entropyfilt(I);
imshow(I), figure, imshow(J,[]);
```

entropyfilt

See Also

entropy, imhist, rangefilt, stdfilt

References

[1] Gonzalez, R.C., R.E. Woods, S.L. Eddins, *Digital Image Processing Using MATLAB*, New Jersey, Prentice Hall, 2003, Chapter 11.

Purpose Compute parallel-beam projections from fan-beam tomography data

Syntax

```
P = fan2para(F,D)
P = fan2para(...,param1,val1,param2,val2,...)
[P,parallel_locations,parallel_rotation_angles] = fan2para(...)
```

Description `P = fan2para(F,D)` computes the parallel-beam data (sinogram) from the fan-beam data (sinogram) `F`. Each column of `F` contains the fan-beam spread angles at a single rotation angle. `D` is the distance from the fan-beam vertex to the center of rotation.

`fan2para` assumes the fan-beam spread angles are the same increments as the input rotation angles, split equally on either side of zero. The input rotation angles are assumed to be stepped in equal increments to cover $[0, 360)$ degrees. Output angles are calculated to cover $[0, 180)$ degrees in the same increments as the input.

`P = fan2para(...,param1,val1,param2,val2,...)` specifies parameters that control various aspects of the `fan2para` conversion, listed in the following table. Parameter names can be abbreviated, and case does not matter. Default values are in braces (`{}`).

Parameter	Description
'FanCoverage'	String specifying the range through which the beams are rotated. Possible values: {'cycle'} or 'minimal' See <code>ifanbeam</code> for details.
'FanRotationIncrement'	Positive real scalar specifying the increment of the rotation angle of the fan-beam projections, measured in degrees. Default value is 1.

fan2para

Parameter	Description
'FanSensorGeometry'	<p>String specifying how sensors are positioned.</p> <p>Possible values: { 'arc' } or 'line'</p> <p>See fanbeam for details.</p>
'FanSensorSpacing'	<p>Positive real scalar specifying the spacing of the fan beams. Interpretation of the value depends on the setting of 'FanSensorGeometry':</p> <p>If 'FanSensorGeometry' is 'arc', the value defines the angular spacing in degrees. Default value is 1.</p> <p>If 'FanSensorGeometry' is 'line', the value defines the linear spacing in pixels.</p>
'Interpolation'	<p>Text string specifying the type of interpolation used between the parallel-beam and fan-beam data.</p> <p>'nearest' — Nearest-neighbor</p> <p>{ 'linear' } — Linear</p> <p>'spline' — Piecewise cubic spline</p> <p>'pchip' — Piecewise cubic Hermite (PCHIP)</p> <p>'cubic' — Same as 'pchip'</p>
'ParallelCoverage'	<p>Text string specifying the range of rotation.</p> <p>Possible values: 'cycle' or { 'halfcycle' }.</p> <p>See para2fan for details.</p>

Parameter	Description
'ParallelRotationIncrement'	<p>Positive real scalar specifying the parallel-beam rotation angle increment, measured in degrees. Parallel beam angles are calculated to cover $[0,180)$ degrees with increment <code>PAR_ROT_INC</code>, where <code>PAR_ROT_INC</code> is the value of 'ParallelRotationIncrement'. $180/\text{PAR_ROT_INC}$ must be an integer.</p> <p>If 'ParallelRotationIncrement' is not specified, the increment is assumed to be the same as the increment of the fan-beam rotation angles.</p>
'ParallelSensorSpacing'	<p>Positive real scalar specifying the spacing of the parallel-beam sensors in pixels. The range of sensor locations is implied by the range of fan angles and is given by</p> $[D*\sin(\min(\text{FAN_ANGLES})), D*\sin(\max(\text{FAN_ANGLES}))]$ <p>If 'ParallelSensorSpacing' is not specified, the spacing is assumed to be uniform and is set to the minimum spacing implied by the fan angles and sampled over the range implied by the fan angles.</p>

```
[P,parallel_locations,parallel_rotation_angles] = fan2para(...)
returns the parallel-beam sensor locations in parallel_locations and
rotation angles in parallel_rotation_angles.
```

Class Support

The input arguments, `F` and `D`, can be double or single, and they must be nonparse. All other numeric inputs are double. The output `P` is double.

Example

Create synthetic parallel-beam data, derive fan-beam data, and then use the fan-beam data to recover the parallel-beam data.

```
ph = phantom(128);
theta = 0:179;
[Psynthetic, xp] = radon(ph, theta);
imshow(theta, xp, Psynthetic, [], 'n'), axis normal
title('Synthetic Parallel-Beam Data')
xlabel('\theta (degrees)')
ylabel('x''')
```

```
colormap(hot), colorbar
Fsynthetic = para2fan(Psynthetic,100,'FanSensorSpacing',1);
```

Recover original parallel-beam data.

```
[Precovered,Ploc,Pangles] = fan2para(Fsynthetic,100,...
                                     'FanSensorSpacing',1,...
                                     'ParallelSensorSpacing',1)
;
figure, imshow(Pangles,Ploc,Precovered,[],'n'), axis normal
title('Recovered Parallel-Beam Data')
xlabel('Rotation Angles (degrees)')
ylabel('Parallel Sensor Locations (pixels)')
colormap(hot), colorbar
```

See Also

fanbeam, ifanbeam, iradon, para2fan, phantom, radon

Purpose Compute fan-beam transform

Syntax

```
F = fanbeam(I,D)
F = fanbeam(...,param1,val1,param1,val2,...)
[F,sensor_positions,fan_rotation_angles] = fanbeam(...)
```

Description `F = fanbeam(I,D)` computes the fan-beam data (sinogram) `F` from the image `I`. `D` is the distance in pixels from the fan-beam vertex to the center of rotation. Each column of `F` contains the fan-beam sensor samples at one rotation angle. The sensors are assumed to have a one-degree angular spacing. The rotation angles are spaced equally to cover `[0:359]` degrees.

`F = fanbeam(...,param1,val1,param1,val2,...)` specifies parameters that control various aspects of the fan-beam projection, listed in the following table. Parameter names can be abbreviated, and case does not matter. Default values are enclosed in braces `{}`.

Parameter	Description
'FanRotationIncrement'	Positive real scalar specifying the increment of the rotation angle of the fan-beam projections, measured in degrees. Default value is 1.
'FanSensorGeometry'	Text string specifying how sensors are positioned. { 'arc' } — Sensors are spaced along a circular arc at distance <code>D</code> from the center of rotation. 'line' — Sensors are spaced equally along a line, the closest point of which is distance <code>D</code> from the center of rotation.
'FanSensorSpacing'	Positive real scalar specifying the spacing of the fan beams. Interpretation of the value depends on the setting of 'FanSensorGeometry': If 'FanSensorGeometry' is 'arc', the value defines the angular spacing in degrees. Default value is 1. If 'FanSensorGeometry' is 'line', the value defines the linear spacing in pixels.

fanbeam

[F,sensor_positions,fan_rotation_angles] = fanbeam(...) returns information about the position of sensors and rotation angles. If 'FanSensorGeometry' is 'arc', sensor_positions contains the fan-beam sensor measurement angles, measured in degrees. If 'FanSensorGeometry' is 'line', sensor_positions contains the fan-beam sensor positions along the line of sensors, measured in pixels. fan_rotation_angles contains rotation angles.

Class Support

I can be logical or numeric. All other numeric inputs and outputs can be double. None of the inputs can be sparse.

Example

```
ph = phantom(128);
imshow(ph)
[F,Floc,Fangles] = fanbeam(ph,250);
imshow(Fangles,Floc,F,[],'n'), axis normal
xlabel('Rotation Angles (degrees)')
ylabel('Sensor Positions (degrees)')
colormap(hot), colorbar
```

See Also

fan2para, ifanbeam, iradon, para2fan, phantom, radon

Reference

[1] Kak, A.C., & Slaney, M., *Principles of Computerized Tomographic Imaging*, IEEE Press, NY, 1988, pp. 92-93.

Purpose

Compute two-dimensional fast Fourier transform

fft2 is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose

Compute N-dimensional fast Fourier transform

fftn is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose

Shift zero-frequency component of fast Fourier transform to center of spectrum

`fftshift` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

filter2

Purpose

Perform two-dimensional linear filtering

`filter2` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose Find output bounds for spatial transformation

Syntax `outbounds = findbounds(TFORM,inbounds)`

Description `outbounds = findbounds(TFORM,inbounds)` estimates the output bounds corresponding to a given spatial transformation and a set of input bounds. TFORM is a spatial transformation structure as returned by `maketform`. `inbounds` is 2-by-`NUM_DIMS` matrix. The first row of `inbounds` specifies the lower bounds for each dimension, and the second row specifies the upper bounds. `NUM_DIMS` has to be consistent with the `ndims_in` field of TFORM.

`outbounds` has the same form as `inbounds`. It is an estimate of the smallest rectangular region completely containing the transformed rectangle represented by the input bounds. Since `outbounds` is only an estimate, it might not completely contain the transformed input rectangle.

Notes `imtransform` uses `findbounds` to compute the 'OutputBounds' parameter if the user does not provide it.

If TFORM contains a forward transformation (a nonempty `forward_fcn` field), then `findbounds` works by transforming the vertices of the input bounds rectangle and then taking minimum and maximum values of the result.

If TFORM does not contain a forward transformation, then `findbounds` estimates the output bounds using the Nelder-Mead optimization function `fminsearch`. If the optimization procedure fails, `findbounds` issues a warning and returns `outbounds = inbounds`.

Example

```
inbounds = [0 0; 1 1]
tform = maketform('affine',[2 0 0; .5 3 0; 0 0 1])
outbounds = findbounds(tform, inbounds)
```

See Also `cp2tform`, `imtransform`, `maketform`, `tformarray`, `tformfwd`, `tforminv`

fliptform

Purpose Flip the input and output roles of a TFORM structure

Syntax TFLIP = fliptform(T)

Description TFLIP = fliptform(T) creates a new spatial transformation structure, a TFORM struct, by flipping the roles of the inputs and outputs in an existing TFORM struct.

Example

```
T = maketform('affine', [.5 0 0; .5 2 0; 0 0 1]);
T2 = fliptform(T)
```

The following are equivalent:

```
x = tformfwd([-3 7],T)
x = tforminv([-3 7],T2)
```

See Also maketform, tformfwd, tforminv

Purpose

Determine frequency spacing for two-dimensional frequency response

freqspace is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

freqz2

Purpose Compute two-dimensional frequency response

Syntax

```
[H,f1,f2] = freqz2(h,n1,n2)
[H,f1,f2] = freqz2(h,[n2 n1])
[H,f1,f2] = freqz2(h)
[H,f1,f2] = freqz2(h,f1,f2)
[...] = freqz2(h,...,[dx dy])
[...] = freqz2(h,...,dx)
freqz2(...)
```

Description `[H,f1,f2] = freqz2(h,n1,n2)` returns `H`, the `n2`-by-`n1` frequency response of `h`, and the frequency vectors `f1` (of length `n1`) and `f2` (of length `n2`). `h` is a two-dimensional FIR filter, in the form of a computational molecule. `f1` and `f2` are returned as normalized frequencies in the range -1.0 to 1.0, where 1.0 corresponds to half the sampling frequency, or π radians.

`[H,f1,f2] = freqz2(h,[n2 n1])` returns the same result returned by `[H,f1,f2] = freqz2(h,n1,n2)`.

`[H,f1,f2] = freqz2(h)` uses `[n2 n1] = [64 64]`.

`[H,f1,f2] = freqz2(h,f1,f2)` returns the frequency response for the FIR filter `h` at frequency values in `f1` and `f2`. These frequency values must be in the range -1.0 to 1.0, where 1.0 corresponds to half the sampling frequency, or π radians.

`[...] = freqz2(h,...,[dx dy])` uses `[dx dy]` to override the intersample spacing in `h`. `dx` determines the spacing for the x dimension and `dy` determines the spacing for the y dimension. The default spacing is 0.5, which corresponds to a sampling frequency of 2.0.

`[...] = freqz2(h,...,dx)` uses `dx` to determine the intersample spacing in both dimensions.


With no output arguments, `freqz2(...)` produces a mesh plot of the two-dimensional magnitude frequency response.

Class Support The input matrix `h` can be of class `double` or of any integer class. All other inputs to `freqz2` must be of class `double`. All outputs are of class `double`.

Example

Use the window method to create a 16-by-16 filter, then view its frequency response using `freqz2`.

```
Hd = zeros(16,16);  
Hd(5:12,5:12) = 1;  
Hd(7:10,7:10) = 0;  
h = fwind1(Hd,bartlett(16));  
colormap(jet(64))  
freqz2(h,[32 32]); axis ([-1 1 -1 1 0 1])
```

**See Also**

`freqz` in the Signal Processing Toolbox User's Guide documentation

fsamp2

Purpose Design two-dimensional FIR filter using frequency sampling

Syntax
`h = fsamp2(Hd)`
`h = fsamp2(f1,f2,Hd,[m n])`

Description `fsamp2` designs two-dimensional FIR filters based on a desired two-dimensional frequency response sampled at points on the Cartesian plane.

`h = fsamp2(Hd)` designs a two-dimensional FIR filter with frequency response `Hd`, and returns the filter coefficients in matrix `h`. (`fsamp2` returns `h` as a computational molecule, which is the appropriate form to use with `filter2`.) The filter `h` has a frequency response that passes through points in `Hd`. If `Hd` is `m`-by-`n`, then `h` is also `m`-by-`n`.

`Hd` is a matrix containing the desired frequency response sampled at equally spaced points between -1.0 and 1.0 along the x and y frequency axes, where 1.0 corresponds to half the sampling frequency, or π radians.

$$\mathcal{I}_d(f_1, f_2) = H_d(\omega_1, \omega_2) \Big|_{\omega_1 = \pi f_1, \omega_2 = \pi f_2}$$

For accurate results, use frequency points returned by `freqspace` to create `Hd`. (See the entry for `freqspace` for more information.)


`h = fsamp2(f1,f2,Hd,[m n])` produces an `m`-by-`n` FIR filter by matching the filter response at the points in the vectors `f1` and `f2`. The frequency vectors `f1` and `f2` are in normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians. The resulting filter fits the desired response as closely as possible in the least squares sense. For best results, there must be at least `m*n` desired frequency points. `fsamp2` issues a warning if you specify fewer than `m*n` points.

Class Support The input matrix `Hd` can be of class `double` or of any integer class. All other inputs to `fsamp2` must be of class `double`. All outputs are of class `double`.

Example Use `fsamp2` to design an approximately symmetric two-dimensional bandpass filter with passband between 0.1 and 0.5 (normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians):


- 1 Create a matrix `Hd` that contains the desired bandpass response. Use `freqspace` to create the frequency range vectors `f1` and `f2`.

```
[f1,f2] = freqspace(21,'meshgrid');  
Hd = ones(21);  
r = sqrt(f1.^2 + f2.^2);  
Hd((r<0.1)|(r>0.5)) = 0;  
colormap(jet(64))  
mesh(f1,f2,Hd)
```



2 Design the filter that passes through this response.

```
h = fsamp2(Hd);  
freqz2(h)
```



Algorithm

`fsamp2` computes the filter h by taking the inverse discrete Fourier transform of the desired frequency response. If the desired frequency response is real and symmetric (zero phase), the resulting filter is also zero phase.

See Also

`conv2`, `filter2`, `freqspace`, `ftrans2`, `fwind1`, `fwind2`

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 213-217.

Purpose Create 2-D special filters

Syntax
`h = fspecial(type)`
`h = fspecial(type,parameters)`

Description `h = fspecial(type)` creates a two-dimensional filter `h` of the specified type. `fspecial` returns `h` as a correlation kernel, which is the appropriate form to use with `imfilter`. `type` is a string having one of these values.

Value	Description
'gaussian'	Gaussian lowpass filter
'sobel'	Sobel horizontal edge-emphasizing filter
'prewitt'	Prewitt horizontal edge-emphasizing filter
'laplacian'	Filter approximating the two-dimensional Laplacian operator
'log'	Laplacian of Gaussian filter
'average'	Averaging filter
'unsharp'	Unsharp contrast enhancement filter

`h = fspecial(type,parameters)` accepts a filter type plus additional modifying parameters particular to the type of filter chosen. If you omit these arguments, `fspecial` uses default values for the parameters.

The following list shows the syntax for each filter type. Where applicable, additional parameters are also shown.

- `h = fspecial('average',hsize)` returns an averaging filter `h` of size `hsize`. The argument `hsize` can be a vector specifying the number of rows and columns in `h`, or it can be a scalar, in which case `h` is a square matrix. The default value for `hsize` is `[3 3]`.
- `h = fspecial('disk',radius)` returns a circular averaging filter (pillbox) within the square matrix of side `2*radius+1`. The default radius is 5.

- `h = fspecial('gaussian', hsize, sigma)` returns a rotationally symmetric Gaussian lowpass filter of size `hsize` with standard deviation `sigma` (positive). `hsize` can be a vector specifying the number of rows and columns in `h`, or it can be a scalar, in which case `h` is a square matrix. The default value for `hsize` is `[3 3]`; the default value for `sigma` is 0.5.
- `h = fspecial('laplacian', alpha)` returns a 3-by-3 filter approximating the shape of the two-dimensional Laplacian operator. The parameter `alpha` controls the shape of the Laplacian and must be in the range 0.0 to 1.0. The default value for `alpha` is 0.2.
- `h = fspecial('log', hsize, sigma)` returns a rotationally symmetric Laplacian of Gaussian filter of size `hsize` with standard deviation `sigma` (positive). `hsize` can be a vector specifying the number of rows and columns in `h`, or it can be a scalar, in which case `h` is a square matrix. The default value for `hsize` is `[5 5]` and 0.5 for `sigma`.
- `h = fspecial('motion', len, theta)` returns a filter to approximate, once convolved with an image, the linear motion of a camera by `len` pixels, with an angle of `theta` degrees in a counterclockwise direction. The filter becomes a vector for horizontal and vertical motions. The default `len` is 9 and the default `theta` is 0, which corresponds to a horizontal motion of nine pixels.
- `h = fspecial('prewitt')` returns a 3-by-3 filter `h` (shown below) that emphasizes horizontal edges by approximating a vertical gradient. If you need to emphasize vertical edges, transpose the filter `h'`.

```
[ 1 1 1
  0 0 0
 -1 -1 -1]
```

To find vertical edges, or for x -derivatives, use `h'`.

- `h = fspecial('sobel')` returns a 3-by-3 filter `h` (shown below) that emphasizes horizontal edges using the smoothing effect by approximating a vertical gradient. If you need to emphasize vertical edges, transpose the filter `h'`.

```
[ 1 2 1
  0 0 0
 -1 -2 -1]
```

- `h = fspecial('unsharp', alpha)` returns a 3-by-3 unsharp contrast enhancement filter. The unsharp contrast enhancement filter enhances edges, and other high frequency components in an image, by subtracting a

smoothed (“unsharp”) version of an image from the original image. `fspecial` creates the unsharp filter from the negative of the Laplacian filter with parameter `alpha`. `alpha` controls the shape of the Laplacian and must be in the range 0.0 to 1.0. The default value for `alpha` is 0.2.

Class Support `h` is of class `double`.

Example

```
I = imread('cameraman.tif');
subplot(2,2,1);
imshow(I); title('Original Image');

H = fspecial('motion',20,45);
MotionBlur = imfilter(I,H,'replicate');
subplot(2,2,2);
imshow(MotionBlur);title('Motion Blurred Image');

H = fspecial('disk',10);
blurred = imfilter(I,H,'replicate');
subplot(2,2,3);
imshow(blurred); title('Blurred Image');

H = fspecial('unsharp');
sharpened = imfilter(I,H,'replicate');
subplot(2,2,4);
imshow(sharpened); title('Sharpened Image');
```



Image Courtesy of MIT

Original Image



Motion Blurred Image



Blurred Image



Sharpened Image

Algorithms

`fspecial` creates Gaussian filters using

$$h_g(n_1, n_2) = e^{-(n_1^2 + n_2^2)/(2\sigma^2)}$$

$$h(n_1, n_2) = \frac{h_g(n_1, n_2)}{\sum_{n_1} \sum_{n_2} h_g}$$

`fspecial` creates Laplacian filters using

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$$

$$\nabla^2 \approx \frac{4}{(\alpha + 1)} \begin{bmatrix} \frac{\alpha}{4} & \frac{1-\alpha}{4} & \frac{\alpha}{4} \\ \frac{1-\alpha}{4} & -1 & \frac{1-\alpha}{4} \\ \frac{\alpha}{4} & \frac{1-\alpha}{4} & \frac{\alpha}{4} \end{bmatrix}$$

fspecial creates Laplacian of Gaussian (LoG) filters using

$$h_g(n_1, n_2) = e^{-(n_1^2 + n_2^2)/(2\sigma^2)}$$

$$h(n_1, n_2) = \frac{(n_1^2 + n_2^2 - 2\sigma^2)h_g(n_1, n_2)}{2\pi\sigma^6 \sum_{n_1} \sum_{n_2} h_g}$$

fspecial creates averaging filters using

$$\text{ones}(n(1), n(2)) / (n(1) * n(2))$$

fspecial creates unsharp filters using

$$\frac{1}{(\alpha + 1)} \begin{bmatrix} -\alpha & \alpha - 1 & -\alpha \\ \alpha - 1 & \alpha + 5 & \alpha - 1 \\ -\alpha & \alpha - 1 & -\alpha \end{bmatrix}$$

See Also

conv2, edge, filter2, fsamp2, fwind1, fwind2, imfilter
del12 in the MATLAB Function Reference

ftrans2

Purpose Design two-dimensional FIR filter using frequency transformation

Syntax
`h = ftrans2(b,t)`
`h = ftrans2(b)`

Description `h = ftrans2(b,t)` produces the two-dimensional FIR filter `h` that corresponds to the one-dimensional FIR filter `b` using the transform `t`. (`ftrans2` returns `h` as a computational molecule, which is the appropriate form to use with `filter2`.) `b` must be a one-dimensional, odd-length (Type I) FIR filter such as can be returned by `fir1`, `fir2`, or `remez` in the Signal Processing Toolbox. The transform matrix `t` contains coefficients that define the frequency transformation to use. If `t` is `m`-by-`n` and `b` has length `Q`, then `h` is size $((m-1)*(Q-1)/2+1)$ -by- $((n-1)*(Q-1)/2+1)$.
`h = ftrans2(b)` uses the McClellan transform matrix `t`.

`t = [1 2 1; 2 -4 2; 1 2 1]/8;`

Remarks The transformation below defines the frequency response of the two-dimensional filter returned by `ftrans2`,

$$H(\omega_1, \omega_2) = B(\omega) \Big|_{\cos \omega = T(\omega_1, \omega_2)}$$

where $B(\omega)$ is the Fourier transform of the one-dimensional filter `b`,

$$B(\omega) = \sum_{n=-N}^N b(n)e^{-j\omega n}$$

and $T(\omega_1, \omega_2)$ is the Fourier transform of the transformation matrix `t`.

$$T(\omega_1, \omega_2) = \sum_{n_2} \sum_{n_1} t(n_1, n_2) e^{-j\omega_1 n_1} e^{-j\omega_2 n_2}$$

The returned filter `h` is the inverse Fourier transform of $H(\omega_1, \omega_2)$.


$$h(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2$$

Example

Use `ftrans2` to design an approximately circularly symmetric two-dimensional bandpass filter with passband between 0.1 and 0.6 (normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians):

- 1 Since `ftrans2` transforms a one-dimensional FIR filter to create a two-dimensional filter, first design a one-dimensional FIR bandpass filter using the Signal Processing Toolbox function `remez`.


```
colormap(jet(64))
b = remez(10,[0 0.05 0.15 0.55 0.65 1],[0 0 1 1 0 0]);
[H,w] = freqz(b,1,128,'whole');
plot(w/pi-1,fftshift(abs(H)))
```



- 2 Use `ftrans2` with the default McClellan transformation to create the desired approximately circularly symmetric filter.

```
h = ftrans2(b);
freqz2(h)
```

ftrans2



See Also

conv2, filter2, fsamp2, fwind1, fwind2

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 218-237.

Purpose	Design two-dimensional FIR filter using one-dimensional window method
Syntax	<pre>h = fwind1(Hd,win) h = fwind1(Hd,win1,win2) h = fwind1(f1,f2,Hd,...)</pre>
Description	<p>fwind1 designs two-dimensional FIR filters using the window method. fwind1 uses a one-dimensional window specification to design a two-dimensional FIR filter based on the desired frequency response Hd. fwind1 works with one-dimensional windows only; use fwind2 to work with two-dimensional windows.</p> <p><code>h = fwind1(Hd,win)</code> designs a two-dimensional FIR filter <code>h</code> with frequency response <code>Hd</code>. (fwind1 returns <code>h</code> as a computational molecule, which is the appropriate form to use with <code>filter2</code>.) fwind1 uses the one-dimensional window <code>win</code> to form an approximately circularly symmetric two-dimensional window using Huang's method. You can specify <code>win</code> using windows from the Signal Processing Toolbox, such as <code>boxcar</code>, <code>hamming</code>, <code>hanning</code>, <code>bartlett</code>, <code>blackman</code>, <code>kaiser</code>, or <code>chebwin</code>. If <code>length(win)</code> is <code>n</code>, then <code>h</code> is <code>n</code>-by-<code>n</code>.</p> <p><code>Hd</code> is a matrix containing the desired frequency response sampled at equally spaced points between -1.0 and 1.0 (in normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians) along the x and y frequency axes. For accurate results, use frequency points returned by <code>freqspace</code> to create <code>Hd</code>. (See the entry for <code>freqspace</code> for more information.)</p> <p><code>h = fwind1(Hd,win1,win2)</code> uses the two one-dimensional windows <code>win1</code> and <code>win2</code> to create a separable two-dimensional window. If <code>length(win1)</code> is <code>n</code> and <code>length(win2)</code> is <code>m</code>, then <code>h</code> is <code>m</code>-by-<code>n</code>.</p> <p><code>h = fwind1(f1,f2,Hd,...)</code> lets you specify the desired frequency response <code>Hd</code> at arbitrary frequencies (<code>f1</code> and <code>f2</code>) along the x- and y-axes. The frequency vectors <code>f1</code> and <code>f2</code> should be in the range -1.0 to 1.0, where 1.0 corresponds to half the sampling frequency, or π radians. The length of the windows controls the size of the resulting filter, as above.</p>
Class Support	The input matrix <code>Hd</code> can be of class <code>double</code> or of any integer class. All other inputs to <code>fwind1</code> must be of class <code>double</code> . All outputs are of class <code>double</code> .


fwind1

Example

Use `fwind1` to design an approximately circularly symmetric two-dimensional bandpass filter with passband between 0.1 and 0.5 (normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians):


- 1 Create a matrix `Hd` that contains the desired bandpass response. Use `freqspace` to create the frequency range vectors `f1` and `f2`.

```
[f1,f2] = freqspace(21,'meshgrid');  
Hd = ones(21);  
r = sqrt(f1.^2 + f2.^2);  
Hd((r<0.1)|(r>0.5)) = 0;  
colormap(jet(64))  
mesh(f1,f2,Hd)
```



- 2 Design the filter using a one-dimensional Hamming window.

```
h = fwind1(Hd,hamming(21));  
freqz2(h)
```



Algorithm

fwind1 takes a one-dimensional window specification and forms an approximately circularly symmetric two-dimensional window using Huang's method,

$$w(n_1, n_2) = w(t) \Big|_{t = \sqrt{n_1^2 + n_2^2}}$$

where $w(t)$ is the one-dimensional window and $w(n_1, n_2)$ is the resulting two-dimensional window.

Given two windows, fwind1 forms a separable two-dimensional window.

$$w(n_1, n_2) = w_1(n_1)w_2(n_2)$$

fwind1 calls fwind2 with H_d and the two-dimensional window. fwind2 computes h using an inverse Fourier transform and multiplication by the two-dimensional window.

$$h_d(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H_d(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2$$

$$h(n_1, n_2) = h_d(n_1, n_2)w(n_1, n_2)$$

fwind1

See Also


conv2, filter2, fsamp2, freqspace, ftrans2, fwind2

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990.


Purpose	Design two-dimensional FIR filter using two-dimensional window method
Syntax	<pre>h = fwind2(Hd,win) h = fwind2(f1,f2,Hd,win)</pre>
Description	<p>Use <code>fwind2</code> to design two-dimensional FIR filters using the window method. <code>fwind2</code> uses a two-dimensional window specification to design a two-dimensional FIR filter based on the desired frequency response <code>Hd</code>. <code>fwind2</code> works with two-dimensional windows; use <code>fwind1</code> to work with one-dimensional windows.</p> <p><code>h = fwind2(Hd,win)</code> produces the two-dimensional FIR filter <code>h</code> using an inverse Fourier transform of the desired frequency response <code>Hd</code> and multiplication by the window <code>win</code>. <code>Hd</code> is a matrix containing the desired frequency response at equally spaced points in the Cartesian plane. <code>fwind2</code> returns <code>h</code> as a computational molecule, which is the appropriate form to use with <code>filter2</code>. <code>h</code> is the same size as <code>win</code>.</p> <p>For accurate results, use frequency points returned by <code>freqspace</code> to create <code>Hd</code>. (See the entry for <code>freqspace</code> for more information.)</p> <p><code>h = fwind2(f1,f2,Hd,win)</code> lets you specify the desired frequency response <code>Hd</code> at arbitrary frequencies (<code>f1</code> and <code>f2</code>) along the x- and y-axes. The frequency vectors <code>f1</code> and <code>f2</code> should be in the range -1.0 to 1.0, where 1.0 corresponds to half the sampling frequency, or π radians. <code>h</code> is the same size as <code>win</code>.</p>
Class Support	The input matrix <code>Hd</code> can be of class <code>double</code> or of any integer class. All other inputs to <code>fwind2</code> must be of class <code>double</code> . All outputs are of class <code>double</code> .
Example	<p>Use <code>fwind2</code> to design an approximately circularly symmetric two-dimensional bandpass filter with passband between 0.1 and 0.5 (normalized frequency, where 1.0 corresponds to half the sampling frequency, or π radians):</p> <ol style="list-style-type: none"> 1 Create a matrix <code>Hd</code> that contains the desired bandpass response. Use <code>freqspace</code> to create the frequency range vectors <code>f1</code> and <code>f2</code>. <pre>[f1,f2] = freqspace(21,'meshgrid'); Hd = ones(21); r = sqrt(f1.^2 + f2.^2); Hd((r<0.1) (r>0.5)) = 0;</pre>

```
colormap(jet(64))  
mesh(f1,f2,Hd)
```




2 Create a two-dimensional Gaussian window using `fspecial`.

```
win = fspecial('gaussian',21,2);  
win = win ./ max(win(:)); % Make the maximum window value be 1.  
mesh(win)
```



3 Design the filter using the window from step 2.

```
h = fwind2(Hd,win);  
freqz2(h)
```



Algorithm

fwind2 computes h using an inverse Fourier transform and multiplication by the two-dimensional window w_{in} .

$$h_d(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H_d(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2$$

$$h(n_1, n_2) = h_d(n_1, n_2)w(n_1, n_2)$$

See Also

conv2, filter2, fsamp2, freqspace, ftrans2, fwind1

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 202-213.

Purpose	Get height of structuring element
Syntax	<code>H = getheight(SE)</code>
Description	<code>H = getheight(SE)</code> returns an array the same size as <code>getnhood(SE)</code> containing the height associated with each of the structuring element neighbors. <code>H</code> is all zeros for a flat structuring element.
Class Support	<code>SE</code> is a STREL object. <code>H</code> is of class double.
Example	<pre>se = strel(ones(3,3),magic(3)); getheight(se)</pre>
See Also	<code>strel</code> , <code>getnhood</code>

getimage

Purpose Get image data from axes

Syntax

```
A = getimage(h)
[x,y,A] = getimage(h)
[... ,A,flag] = getimage(h)
[...] = getimage
```

Description `A = getimage(h)` returns the first image data contained in the Handle Graphics object `h`. `h` can be a figure, axes, or image. `A` is identical to the image CData; it contains the same values and is of the same class (uint8, uint16, double, or logical) as the image CData. If `h` is not an image or does not contain an image, `A` is empty.

`[x,y,A] = getimage(h)` returns the image XData in `x` and the YData in `y`. XData and YData are two-element vectors that indicate the range of the *x*-axis and *y*-axis.

`[... ,A,flag] = getimage(h)` returns an integer flag that indicates the type of image `h` contains. This table summarizes the possible values for `flag`.

Flag	Type of Image
0	Not an image; A is returned as an empty matrix
1	Indexed image
2	Intensity image with values in standard range ([0,1] for single and double arrays, [0,255] for uint8 arrays, [0,65535] for uint16 arrays)
3	Intensity data, but not in standard range
4	RGB image
5	Binary image

`[...] = getimage` returns information for the current axes object. It is equivalent to `[...] = getimage(gca)`.

Class Support The output array `A` is of the same class as the image `CData`. All other inputs and outputs are of class `double`.

Note For `int16` and `single` images, the image data returned by `getimage` is of class `double`, not `int16` or `single`. This is because the `getimage` function gets the data from the image object's `CData` property and image objects store `int16` and `single` image data as class `double`.

For example, create an image object of class `int16`. If you retrieve the `CData` from the object and check its class, it returns `double`.

```
h = imshow(ones(10,'int16'));  
class(get(h,'CData'))
```

Therefore, if you get the image data using the `getimage` function, the data it returns is also of class `double`. The `flag` return value is set to 3.

```
[img,flag] = getimage(h);  
class(img)
```

The same is true for an image of class `single`. Getting the `CData` directly from the image object or by using `getimage`, the class of the returned data is `double`.

```
h = imshow(ones(10,'single'));  
class(get(h,'CData'))  
[img,flag] = getimage(h);  
class(img)
```

For images of class `single`, the `flag` return value is set to 2 because `single` and `double` share the same dynamic range.

Example After using `imshow` or `imtool` to display an image directly from a file, use `getimage` to get the image data into the workspace.

```
imshow rice.png  
I = getimage;  
  
imtool cameraman.tif  
I = getimage(imgca);
```

See Also `imshow`, `imtool`

getimagemodel

Purpose Retrieve image model object from image object

Syntax `imgmodel = getimagemodel(himage)`

Description `imgmodel = getimagemodel(himage)` returns the image model object associated with `himage`. `himage` must be a handle to an image object or an array of handles to image objects.

The return value `imgmodel` is an image model object. If `himage` is an array of handles to image objects, `imgmodel` is an array of image models.

If `himage` does not have an associated image model object, `getimagemodel` creates one.

Example

```
h = imshow('bag.png');  
imgmodel = getimagemodel(h);
```

See Also `imagemodel`

Purpose	Select polyline with mouse
Syntax	<pre>[x,y] = getline(fig) [x,y] = getline(ax) [x,y] = getline [x,y] = getline(...,'closed')</pre>
Description	<p>[x,y] = getline(fig) lets you select a polyline in the current axes of figure fig using the mouse. Coordinates of the polyline are returned in X and Y. Use normal button clicks to add points to the polyline. A shift-, right-, or double-click adds a final point and ends the polyline selection. Pressing Return or Enter ends the polyline selection without adding a final point. Pressing Backspace or Delete removes the previously selected point from the polyline.</p> <p>[x,y] = getline(ax) lets you select a polyline in the axes specified by the handle ax.</p> <p>[x,y] = getline is the same as [x,y] = getline(gcf).</p> <p>[x,y] = getline(...,'closed') animates and returns a closed polygon.</p>
See Also	getpts, getrect

getneighbors

Purpose Get structuring element neighbor locations and heights

Syntax `[offsets,heights] = getneighbors(SE)`

Description `[offsets,heights] = getneighbors(SE)` returns the relative locations and corresponding heights for each of the neighbors in the structuring element object SE.

`offsets` is a P-by-N array where P is the number of neighbors in the structuring element and N is the dimensionality of the structuring element. Each row of `offsets` contains the location of the corresponding neighbor, relative to the center of the structuring element.

`heights` is a P-element column vector containing the height of each structuring element neighbor.

Class Support SE is a STREL object. The return values `offsets` and `heights` are arrays of double-precision values.

Example

```
se = strel([1 0 1],[5 0 -5])
[offsets,heights] = getneighbors(se)
se =
Nonflat STREL object containing 2 neighbors.
```

```
Neighborhood:
    1    0    1
```

```
Height:
    5    0   -5
```

```
offsets =
    0   -1
    0    1
```

```
heights =
    5   -5
```

See Also `strel`, `getnhood`, `getheight`

Purpose	Get structuring element neighborhood
Syntax	<code>nhood = getnhood(SE)</code>
Description	<code>nhood = getnhood(SE)</code> returns the neighborhood associated with the structuring element SE.
Class Support	SE is a STREL object. nhood is a logical array.
Example	<pre>se = strel(eye(5)); nhood = getnhood(se)</pre>
See Also	<code>strel</code> , <code>getneighbors</code>

getpts

Purpose Select points with mouse

Syntax
`[x,y] = getpts(fig)`
`[x,y] = getpts(ax)`
`[x,y] = getpts`

Description `[x,y] = getpts(fig)` lets you choose a set of points in the current axes of figure `fig` using the mouse. Coordinates of the selected points are returned in `X` and `Y`.

Use normal button clicks to add points. A shift-, right-, or double-click adds a final point and ends the selection. Pressing **Return** or **Enter** ends the selection without adding a final point. Pressing **Backspace** or **Delete** removes the previously selected point.

`[x,y] = getpts(ax)` lets you choose points in the axes specified by the handle `ax`.

`[x,y] = getpts` is the same as `[x,y] = getpts(gcf)`.

See Also `getline`, `getrect`

Purpose	Get dynamic range of image based on its class
Syntax	<code>range = getrangefromclass(I)</code>
Description	<code>range = getrangefromclass(I)</code> returns the dynamic range of the image <code>I</code> , based on its class type. <code>range</code> is a two-element vector specifying the dynamic range.
Class Support	<code>I</code> can be <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>logical</code> , <code>single</code> , or <code>double</code> . <code>range</code> is of class <code>double</code> .
Note	For <code>single</code> and <code>double</code> data, <code>getrangefromclass</code> returns the range <code>[0 1]</code> to be consistent with the way <code>double</code> and <code>single</code> images are interpreted in MATLAB. For integer data, <code>getrangefromclass</code> returns the range of the class. For example, if the class is <code>uint8</code> , the dynamic range is <code>[0 255]</code> .
Example	This example reads in the 16-bit DICOM image and gets the range. <pre>CT = dicomread('CT-MONO2-16-ankle.dcm'); r = getrangefromclass(CT)</pre>
See Also	<code>intmin</code> , <code>intmax</code>

getrect

Purpose Select rectangle with mouse

Syntax

```
rect = getrect(fig)
rect = getrect(ax)
rect = getrect(fig)
```

Description `rect = getrect(fig)` lets you select a rectangle in the current axes of figure `fig` using the mouse. Coordinates of the rectangle are returned in `X` and `Y`.

Use the mouse to click and drag the desired rectangle. `rect` is a four-element vector with the form `[xmin ymin width height]`. To constrain the rectangle to be a square, use a shift- or right-click to begin the drag.

`rect = getrect(ax)` lets you select a rectangle in the axes specified by the handle `ax`.

See Also `getline`, `getpts`

Purpose Extract sequence of decomposed structuring elements

Syntax SEQ = getsequence(SE)

Description SEQ = getsequence(SE), where SE is a structuring element array, returns another structuring element array SEQ containing the individual structuring elements that form the decomposition of SE. SEQ is equivalent to SE, but the elements of SEQ have no decomposition.

Class Support SE and SEQ are arrays of STREL objects.

Example The `strel` function uses decomposition for square structuring elements larger than 3-by-3. Use `getsequence` to extract the decomposed structuring elements.

```
se = strel('square',5)
seq = getsequence(se)
se =
Flat STREL object containing 25 neighbors.
Decomposition: 2 STREL objects containing a total of 10 neighbors
```

```
Neighborhood:
  1   1   1   1   1
  1   1   1   1   1
  1   1   1   1   1
  1   1   1   1   1
  1   1   1   1   1
seq =
2x1 array of STREL objects
```

Use `imdilate` with the 'full' option to see that dilating sequentially with the decomposed structuring elements really does form a 5-by-5 square:

```
imdilate(1,seq,'full')
```

See Also `imdilate`, `imerode`, `strel`

gray2ind

Purpose Convert an intensity image to an indexed image

Syntax `[X,map] = gray2ind(I,n)`
`[X,map] = gray2ind(BW,n)`

Description `gray2ind` scales, then rounds, an intensity image to produce an equivalent indexed image.

`[X,map] = gray2ind(I,n)` converts the intensity image `I` to an indexed image `X` with colormap `gray(n)`. If `n` is omitted, it defaults to 64.

`[X,map] = gray2ind(BW,n)` converts the binary image `BW` to an indexed image `X` with colormap `gray(n)`. If `n` is omitted, it defaults to 2.

`n` must be an integer between 1 and 65536.

Class Support The input image `I` can be `logical`, `uint8`, `uint16`, `int16`, `single`, or `double` and must be a real and nonsparse. The image `I` can have any dimension. The class of the output image `X` is `uint8` if the colormap length is less than or equal to 256; otherwise it is `uint16`.

Example Convert an intensity image into an indexed image and then view the result.

```
I = imread('cameraman.tif');  
[X, map] = gray2ind(I, 16);  
imshow(X, map);
```

See Also `grayslice`, `ind2gray`, `mat2gray`

Purpose Gray-level co-occurrence matrix from an image

Syntax


```
glcm = graycomatrix(I)
glcms = graycomatrix(I,param1,val1,param2,val2,...)
[glcms,SI] = graycomatrix(...)
```

Description `glcms = graycomatrix(I)` creates a gray-level co-occurrence matrix (GLCM) from image `I`. `graycomatrix` creates the GLCM by calculating how often a pixel with gray-level (grayscale intensity) value i occurs horizontally adjacent to a pixel with the value j . (You can specify other pixel spatial relationships using the 'Offsets' parameter — see Parameters.) Each element (i,j) in `glcm` specifies the number of times that the pixel with value i occurred horizontally adjacent to a pixel with value j .

Because the processing required to calculate a GLCM for the full dynamic range of an image is prohibitive, `graycomatrix` scales the values in `I`. If `I` is a binary image, `graycomatrix` scales the image to two levels. If `I` is an intensity image, `graycomatrix` scales it to eight levels. The number of gray-levels determines the size of the GLCM. You can specify the number of gray-levels in the GLCM using the 'NumLevels' parameter, and the way that `graycomatrix` scales the values using the 'GrayLimits' parameter — see Parameters.

The following figure shows how `graycomatrix` calculates the value of `glcm(1,1)`. This element contains the value 1 because there is only one instance in the image where two, horizontally adjacent pixels have the values 1 and 1. As the figure shows, `glcm(1,2)` contains the value 2 because there are two instances where two, horizontally adjacent pixels have the values 1 and 2.

graycomatrix



`glcms = graycomatrix(I,param1,va11,param2,va12,...)` returns one or more gray-level co-occurrence matrices, depending on the values of the optional parameter/value pairs. Parameter names can be abbreviated, and case does not matter.

Parameters

The following table lists these parameters in alphabetical order.

Parameter	Description	Default
'GrayLimits'	A two-element vector, [low high], that specifies how the grayscale values in I are linearly scaled into gray levels. Grayscale values less than or equal to low are scaled to 1. Grayscale values greater than or equal to high are scaled to NumLevels. If graylimits is set to [], graycomatrix uses the minimum and maximum grayscale values in the image as limits, [min(I(:)) max(I(:))].	Minimum and maximum specified by class, e.g. double [0 1] int16 [-32768 32767]

Parameter	Description	Default										
'NumLevels'	An integer specifying the number of gray-levels to use when scaling the grayscale values in I. For example, if NumLevels is 8, graycomatrix scales the values in I so they are integers between 1 and 8. The number of gray-levels determines the size of the gray-level co-occurrence matrix (glcm).	8 (numeric) 2 (binary)										
'Offset'	<p>A p-by-2 array of integers specifying the distance between the pixel of interest and its neighbor. Each row in the array is a two-element vector, [row_offset, col_offset], that specifies the relationship, or <i>offset</i>, of a pair of pixels. row_offset is the number of rows between the pixel-of-interest and its neighbor. col_offset is the number of columns between the pixel-of-interest and its neighbor. Because the offset is often expressed as an angle, the following table lists the offset values that specify common angles, given the pixel distance D.</p> <table border="1"> <thead> <tr> <th>Angle</th> <th>Offset</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>[0 D]</td> </tr> <tr> <td>45</td> <td>[-D D]</td> </tr> <tr> <td>90</td> <td>[-D 0]</td> </tr> <tr> <td>135</td> <td>[-D -D]</td> </tr> </tbody> </table> <p>The figure illustrates the array: offset = [0 1; -1 1; -1 0; -1 -1]</p>	Angle	Offset	0	[0 D]	45	[-D D]	90	[-D 0]	135	[-D -D]	[0 1]
Angle	Offset											
0	[0 D]											
45	[-D D]											
90	[-D 0]											
135	[-D -D]											

graycomatrix

`[glcm, SI] = graycomatrix(...)` returns the scaled image, `SI`, used to calculate the gray-level co-occurrence matrix. The values in `SI` are between 1 and `NumLevels`.

Class Support `I` can be numeric or logical but must be two-dimensional, real, and nonsparse. `SI` is a double matrix having the same size as `I`. `glcms` is a 'NumLevels'-by-'NumLevels'-by-`P` double array where `P` is the number of offsets in 'Offset'.

Notes Another name for a gray-level co-occurrence matrix is a gray-level spatial dependence matrix.

`graycomatrix` ignores pixel pairs if either of the pixels contains a NaN.

`graycomatrix` replaces positive Infs with the value `NumLevels` and replaces negative Infs with the value 1.

`graycomatrix` ignores border pixels, if the corresponding neighbor pixel falls outside the image boundaries.

Example Calculate the gray-level co-occurrence matrix and return the scaled version of the image, `SI`, used by `graycomatrix` to generate the GLCM.

```
I = [ 1 1 5 6 8 8; 2 3 5 7 0 2; 0 2 3 5 6 7];  
[glcm,SI] = graycomatrix(I,'NumLevels',9,'G',[])
```

Calculate the gray-level co-occurrence matrix for the grayscale image, `circuit.tif`.

```
I = imread('circuit.tif');  
glcm = graycomatrix(I,'Offset',[2 0]);
```

See Also `graycoprops`

Purpose Properties of a gray-level co-occurrence matrix

Syntax stats = graycoprops(glcm, properties)

Description stats = graycoprops(glcm, properties) calculates the statistics specified in properties from the gray-level co-occurrence matrix glcm. glcm is an *m-by-n-by-p* array of valid gray-level co-occurrence matrices. If glcm is an array of GLCMs, stats is an array of statistics for each glcm.

graycoprops normalizes the gray-level co-occurrence matrix (GLCM) so that the sum of its elements is equal to 1. Each element (r,c) in the normalized GLCM is the joint probability occurrence of pixel pairs with a defined spatial relationship having gray level values r and c in the image. graycoprops uses the normalized GLCM to calculate properties.

properties can be a comma-separated list of strings, a cell array containing strings, the string 'all', or a space separated string. The property names can be abbreviated and are not case sensitive.

Property	Description	Formula
'Contrast'	Returns a measure of the intensity contrast between a pixel and its neighbor over the whole image. Range = [0 (size(GLCM,1) - 1)^2] Contrast is 0 for a constant image.	$\sum_{i,j} i - j ^2 p(i, j)$
'Correlation'	Returns a measure of how correlated a pixel is to its neighbor over the whole image. Range = [-1 1] Correlation is 1 or -1 for a perfectly positively or negatively correlated image. Correlation is NaN for a constant image.	$\sum_{i,j} \frac{(i - \mu_i)(j - \mu_j) p(i, j)}{\sigma_i \sigma_j}$

graycoprops

Property	Description	Formula
'Energy'	Returns the sum of squared elements in the GLCM. Range = [0 1] Energy is 1 for a constant image.	$\sum_{i,j} p(i,j)^2$
'Homogeneity'	Returns a value that measures the closeness of the distribution of elements in the GLCM to the GLCM diagonal. Range = [0 1] Homogeneity is 1 for a diagonal GLCM.	$\sum_{i,j} \frac{p(i,j)}{1 + i - j }$

stats is a structure with fields that are specified by properties. Each field contains a 1 x p array, where p is the number of gray-level co-occurrence matrices in GLCM. For example, if GLCM is an 8 x 8 x 3 array and properties is 'Energy', then stats is a structure containing the field Energy, which contains a 1 x 3 array.

Notes

Energy is also known as uniformity, uniformity of energy, and angular second moment.

Contrast is also known as variance and inertia.

Class Support

glcm can be logical or numeric, and it must contain real, non-negative, finite, integers. stats is a structure.

Example

```
GLCM = [0 1 2 3;1 1 2 3;1 0 2 0;0 0 0 3];  
stats = graycoprops(GLCM)  
  
I = imread('circuit.tif');  
GLCM2 = graycomatrix(I,'Offset',[2 0;0 2]);  
stats = graycoprops(GLCM2,{'contrast','homogeneity'})
```


See Also

graycomatrix

- Purpose** Create indexed image from intensity image using multilevel thresholding
- Syntax**
`X = grayscale(I,n)`
`X = grayscale(I,v)`
- Description**
`X = grayscale(I,n)` thresholds the intensity image `I` using threshold values $\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}$, returning an indexed image in `X`.
`X = grayscale(I,v)` thresholds the intensity image `I` using the values of `v`, where `v` is a vector of values between 0 and 1, returning an indexed image in `X`.
 You can view the thresholded image using `imshow(X,map)` with a colormap of appropriate length.

- Class Support** The input image `I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`, and must be nonsparse. Note that the threshold values are always between 0 and 1, even if `I` is of class `uint8` or `uint16`. In this case, each threshold value is multiplied by 255 or 65535 to determine the actual threshold to use.
 The class of the output image `X` depends on the number of threshold values, as specified by `n` or `length(v)`. If the number of threshold values is less than 256, then `X` is of class `uint8`, and the values in `X` range from 0 to `n` or `length(v)`. If the number of threshold values is 256 or greater, `X` is of class `double`, and the values in `X` range from 1 to `n+1` or `length(v)+1`.

- Example**
- ```
I = imread('snowflakes.png');
X = grayscale(I,16);
imshow(I)
figure, imshow(X,jet(16))
```



- See Also** `gray2ind`

# graythresh

---

**Purpose** Compute global image threshold using Otsu's method

**Syntax** `level = graythresh(I)`  
`[level EM] = graythresh(I)`

**Description** `level = graythresh(I)` computes a global threshold (`level`) that can be used to convert an intensity image to a binary image with `im2bw`. `level` is a normalized intensity value that lies in the range `[0, 1]`.

The `graythresh` function uses Otsu's method, which chooses the threshold to minimize the intraclass variance of the black and white pixels.

Multidimensional arrays are converted automatically to 2-D arrays using `reshape`. The `graythresh` function ignores any nonzero imaginary part of `I`.

`[level EM] = graythresh(I)` returns the effectiveness metric, `EM`, as the second output argument. The effectiveness metric is a value in the range `[0 1]` that indicates the effectiveness of the thresholding of the input image. The lower bound is attainable only by images having a single gray level, and the upper bound is attainable only by two-valued images.

**Class Support** The input image `I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double` and it must be nonsparse. The return value `level` is a double scalar. The effectiveness metric `EM` is a double scalar.

**Example**

```
I = imread('coins.png');
level = graythresh(I);
BW = im2bw(I,level);
imshow(BW)
```

**See Also** `im2bw`

**Reference** Otsu, N., "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 9, No. 1, 1979, pp. 62-66.

**Purpose** Enhance contrast using histogram equalization

**Syntax**

```
J = histeq(I,hgram)
J = histeq(I,n)
[J,T] = histeq(I,...)
```

```
newmap = histeq(X,map,hgram)
newmap = histeq(X,map)
[newmap,T] = histeq(X,...)
```

**Description** `histeq` enhances the contrast of images by transforming the values in an intensity image, or the values in the colormap of an indexed image, so that the histogram of the output image approximately matches a specified histogram.

`J = histeq(I,hgram)` transforms the intensity image `I` so that the histogram of the output intensity image `J` with `length(hgram)` bins approximately matches `hgram`. The vector `hgram` should contain integer counts for equally spaced bins with intensity values in the appropriate range: `[0, 1]` for images of class `double`, `[0, 255]` for images of class `uint8`, and `[0, 65535]` for images of class `uint16`. `histeq` automatically scales `hgram` so that `sum(hgram) = prod(size(I))`. The histogram of `J` will better match `hgram` when `length(hgram)` is much smaller than the number of discrete levels in `I`.

`J = histeq(I,n)` transforms the intensity image `I`, returning in `J` an intensity image with `n` discrete gray levels. A roughly equal number of pixels is mapped to each of the `n` levels in `J`, so that the histogram of `J` is approximately flat. (The histogram of `J` is flatter when `n` is much smaller than the number of discrete levels in `I`.) The default value for `n` is 64.

`[J,T] = histeq(I,...)` returns the grayscale transformation that maps gray levels in the intensity image `I` to gray levels in `J`.

`newmap = histeq(X,map,hgram)` transforms the colormap associated with the indexed image `X` so that the histogram of the gray component of the indexed image (`X,newmap`) approximately matches `hgram`. The `histeq` function returns the transformed colormap in `newmap`. `length(hgram)` must be the same as `size(map,1)`.

# histeq

---

`newmap = histeq(X,map)` transforms the values in the colormap so that the histogram of the gray component of the indexed image `X` is approximately flat. It returns the transformed colormap in `newmap`.

`[newmap,T] = histeq(X,...)` returns the grayscale transformation `T` that maps the gray component of `map` to the gray component of `newmap`.

## Class Support

For syntaxes that include an intensity image `I` as input, `I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image `J` has the same class as `I`.

For syntaxes that include an indexed image `X` as input, `X` can be of class `uint8`, `single`, or `double`; the output colormap is always of class `double`. The optional output `T` (the gray-level transform) is always of class `double`.

## Example


Enhance the contrast of an intensity image using histogram equalization.

```
I = imread('tire.tif');
J = histeq(I);
imshow(I)
figure, imshow(J)
```



Display the resulting histograms.

```
imhist(I,64)
figure; imhist(J,64)
```



## Algorithm

When you supply a desired histogram `hgram`, `histeq` chooses the grayscale transformation  $T$  to minimize

$$|c_1(T(k)) - c_0(k)|$$

where  $c_0$  is the cumulative histogram of  $A$ ,  $c_1$  is the cumulative sum of `hgram` for all intensities  $k$ . This minimization is subject to the constraints that  $T$  must be monotonic and  $c_1(T(a))$  cannot overshoot  $c_0(a)$  by more than half the distance between the histogram counts at  $a$ . `histeq` uses this transformation to map the gray levels in  $X$  (or the colormap) to their new values.

$$b = T(a)$$

If you do not specify `hgram`, `histeq` creates a flat `hgram`,

$$\text{hgram} = \text{ones}(1,n) * \text{prod}(\text{size}(A)) / n;$$

and then applies the previous algorithm.

## See Also

`brighten`, `imadjust`, `imhist`

# hough

**Purpose** Hough transform

**Syntax** [H, theta, rho] = hough(bw)

**Description** [H, theta, rho] = hough(BW) computes the Standard Hough Transform (SHT) of the binary image BW. You can use the hough function to detect lines in an image. The function returns H, the Hough transform matrix. theta (in degrees) and rho are the arrays of rho and theta values over which the Hough transform matrix was generated.

[H, theta, rho] = hough(BW,param1,val1,param2,val2) specifies parameter/value pairs, listed in the following table. Parameter names can be abbreviated, and case does not matter.

| Parameter         | Description                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'ThetaResolution' | Real scalar value between 0 and 90, exclusive, that specifies the spacing (in degrees) of the Hough transform bins along the theta axis. Default: 1. |
| 'RhoResolution'   | Real scalar value between 0 and norm(size(BW)), exclusive, that specifies the spacing of the Hough transform bins along the rho axis. Default: 1     |

**Notes** The hough function implements the Standard Hough Transform (SHT). The SHT uses the parametric representation of a line:

$$\rho = x \cdot \cos(\theta) + y \cdot \sin(\theta)$$

The variable rho is the distance from the origin to the line along a vector perpendicular to the line. theta is the angle between the x-axis and this vector. The hough function generates a parameter space matrix whose rows and columns correspond to rho and theta values respectively. Peak values in this space represent potential lines in the input image.

The Hough transform matrix, H, is NRHO-by-NTHETA where NRHO =  $2 \cdot \text{ceil}(\text{norm}(\text{size}(\text{BW})) / \text{RhoResolution}) - 1$ , and NTHETA =  $2 \cdot \text{ceil}(90 / \text{ThetaResolution})$ . Theta angle values are in the range [-90, 90) degrees and rho values range from -DIAGONAL to DIAGONAL where DIAGONAL =  $\text{RhoResolution} \cdot \text{ceil}(\text{norm}(\text{size}(\text{BW})) / \text{RhoResolution})$ . Note that

if  $90/D\theta$  is not an integer, the actual angle spacing will be  $90/\text{ceil}(90/D\theta)$ .

**Class Support** BW can be logical or numeric and it must be real, 2-D, and nonsparse.

**Example** Compute and display the Hough transform of an image

```
RGB = imread('gantrycrane.png');
I = rgb2gray(RGB); % convert to intensity
BW = edge(I,'canny'); % extract edges
[H,T,R] = hough(BW,'RhoResolution',0.5,'ThetaResolution',0.5);

% display the original image
subplot(2,1,1);
imshow(RGB);
title('gantrycrane.png');

% display the hough matrix
subplot(2,1,2);
imshow(imadjust(mat2gray(H)), 'XData',T, 'YData',R,...
 'InitialMagnification','fit');
title('Hough transform of gantrycrane.png');
xlabel('\theta'), ylabel('\rho');
axis on, axis normal, hold on;
colormap(hot);
```

**See also** houghpeaks, houghlines

# houghlines

---

**Purpose** Extract line segments based on the Hough transform

**Syntax** `lines = houghlines(BW, theta, rho, peaks)`

**Description** `lines = houghlines(BW, theta, rho, peaks)` extracts line segments in the image `BW` associated with particular bins in a Hough transform. `theta` and `rho` are vectors returned by function `hough`. `peaks` is a matrix returned by the `houghpeaks` function that contains the row and column coordinates of the Hough transform bins to use in searching for line segments.

The `houghlines` function returns `lines`, a structure array whose length equals the number of merged line segments found. Each element of the structure array has these fields:

| Field               | Description                                                                              |
|---------------------|------------------------------------------------------------------------------------------|
| <code>point1</code> | Two element vector [X Y] specifying the coordinates of the end-point of the line segment |
| <code>point2</code> | Two element vector [X Y] specifying the coordinates of the end-point of the line segment |
| <code>theta</code>  | Angle in degrees of the Hough transform bin                                              |
| <code>rho</code>    | rho axis position of the Hough transform bin                                             |

`lines = houghlines(...,param1,va11,param2,va12)` specifies parameter/value pairs, listed in the following table. Parameter names can be abbreviated, and case does not matter.

| Parameter   | Description                                                                                                                                                                                                                                                                                                   |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'FillGap'   | Positive real scalar value that specifies the distance between two line segments associated with the same Hough transform bin. When the distance between the line segments is less the value specified, the <code>houghlines</code> function merges the line segments into a single line segment. Default: 20 |
| 'MinLength' | Positive real scalar value that specifies whether merged lines should be kept or discarded. Lines shorter than the value specified are discarded. Default: 40                                                                                                                                                 |

## Class Support

`BW` can be logical or numeric and it must be real, 2-D, and nonsparse.

## Example

This example searches for line segments in an image and highlights the longest segment.

```
I = imread('circuit.tif');
rotI = imrotate(I,33,'crop');
BW = edge(rotI,'canny');
[H,T,R] = hough(BW);
imshow(H,[],'XData',T,'YData',R,'InitialMagnification','fit');
xlabel('\theta'), ylabel('\rho');
axis on, axis normal, hold on;
P = houghpeaks(H,5,'threshold',ceil(0.3*max(H(:)))));
x = T(P(:,2)); y = R(P(:,1));
plot(x,y,'s','color','white');
% Find lines and plot them
lines = houghlines(BW,T,R,P,'FillGap',5,'MinLength',7);
figure, imshow(rotI), hold on
max_len = 0;
for k = 1:length(lines)
 xy = [lines(k).point1; lines(k).point2];
```

# houghlines

---

```
plot(xy(:,1),xy(:,2),'LineWidth',2,'Color','green');

% Plot beginnings and ends of lines
plot(xy(1,1),xy(1,2),'x','LineWidth',2,'Color','yellow');
plot(xy(2,1),xy(2,2),'x','LineWidth',2,'Color','red');

% Determine the endpoints of the longest line segment
len = norm(lines(k).point1 - lines(k).point2);
if (len > max_len)
 max_len = len;
 xy_long = xy;
end
end

% highlight the longest line segment
plot(xy_long(:,1),xy_long(:,2),'LineWidth',2,'Color','cyan');
```

## See also

[hough](#), [houghpeaks](#)

**Purpose** Identify peaks in the Hough transform

**Syntax** `peaks = houghpeaks(H,numpeaks)`

**Description** `peaks = houghpeaks(H,numpeaks)` locates peaks in the Hough transform matrix, `H`, generated by the `hough` function. `numpeaks` specifies the maximum number of peaks to identify.

The function returns `peaks`, a `Q`-by-2 matrix, where `Q` can range from 0 to `numpeaks`. `Q` holds the row and column coordinates of the peaks. If `numpeaks` is omitted, it defaults to 1.

`peaks = houghpeaks(...,param1,va11,param2,va12)` specifies parameter/value pairs, listed in the following table. Parameter names can be abbreviated, and case does not matter.

| Parameter   | Description                                                                                                                                                                                                                                                                                                   |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'Threshold' | Nonnegative scalar value that specifies the threshold at which values of <code>H</code> are considered to be peaks. Threshold can vary from 0 to Inf. Default is <code>0.5*max(H(:))</code> .                                                                                                                 |
| 'NHoodSize' | Two-element vector of positive odd integers: <code>[M N]</code> . 'NHoodSize' specifies the size of the suppression neighborhood. This is the neighborhood around each peak that is set to zero after the peak is identified. Default: smallest odd values greater than or equal to <code>size(H)/50</code> . |

**Class Support** `H` is the output of the `hough` function. `numpeaks` is a positive integer scalar.

**Example** Locate and display two peaks in the Hough transform of the rotated circuit.tif image.

```
I = imread('circuit.tif');
BW = edge(imrotate(I,50,'crop'),'canny');
[H,T,R] = hough(BW);
P = houghpeaks(H,2);
```

# houghpeaks

---

```
imshow(H,[],'XData',T,'YData',R,'InitialMagnification','fit');
xlabel('\theta'), ylabel('\rho');
axis on, axis normal, hold on;
plot(T(P(:,2)),R(P(:,1)),'s','color','white');
```

## See also

hough, houghlines

**Purpose**

Convert hue-saturation-value (HSV) values to RGB color space

hsv2rgb is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

# iccread

---

**Purpose** Read ICC profile

**Syntax** `P = iccread(filename)`

**Description** `P = iccread(filename)` reads the International Color Consortium (ICC) color profile information from the file specified by `filename`. The file can be either an ICC profile file or a TIFF file containing an embedded ICC profile. To determine if a TIFF file contains an embedded ICC profile, use the `imfinfo` function to get information about the file and look for the `ICCProfileOffset` field.

`iccread` returns the profile information in the structure `P`, a 1-by-1 structure array whose fields contain the data structures (called tags) defined in the ICC specification. `iccread` can read profiles that conform with either Version 2 (ICC.1:2001-04) or Version 4 (ICC.1:2001-12) of the ICC specification. For more information about ICC profiles, visit the ICC web site, [www.color.org](http://www.color.org).

ICC profiles provide color management systems with the information necessary to convert color data between native device color spaces and device independent color spaces, called the Profile Connection Space (PCS). You can use the profile as the source or destination profile with the `makecform` function to compute color space transformations.

The number of fields in `P` depends on the profile class (input, output, or display) and the choices made by the profile creator. `iccread` returns all the tags for a given profile, both public and private. Private tags and certain public tags are left as encoded `uint8` data. The following table lists fields that are found in any profile structure generated by `iccread`, in the order they appear in the structure.

| Field     | Data Type           | Description              |
|-----------|---------------------|--------------------------|
| Header    | 1-by-1 struct array | Profile header fields    |
| TagTable  | n-by-3 cell array   | Profile tag table        |
| Copyright | Text string         | Profile copyright notice |

| Field           | Data Type           | Description                                                                                                                                                                                                                                                       |
|-----------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description     | 1-by-1 struct array | The Plain field in this structure contains a text string describing the profile.                                                                                                                                                                                  |
| MediaWhitepoint | double array        | XYZ tristimulus values of the device's media white point                                                                                                                                                                                                          |
| PrivateTags     | m-by-2 cell array   | Contents of all the private tags or tags not defined in the ICC specifications. The tag signatures are in the first column, and the contents of the tags are in the second column. Note that iccread leaves the contents of these tags in unsigned 8-bit encoding |
| Filename        | Text string         | Name of the file containing the profile                                                                                                                                                                                                                           |

Additionally, P might contain one or more of the following transforms:

- Three-component, matrix-based transform: A simple transform that is often used to transform between the RGB and XYZ color spaces. If this transform is present, P contains a field called MatTRC.
- N-component LUT-based transform: A transform that is used for transforming between color spaces that have a more complex relationship. This type of transform is found in any of the following fields in P:

|       |       |          |       |
|-------|-------|----------|-------|
| AToB0 | BToA0 | Preview0 | Gamut |
| AToB1 | BToA1 | Preview1 |       |
| AToB2 | BToA2 | Preview2 |       |

## Example

The example reads the ICC profile that describes a typical PC computer monitor.

```
P = iccread('sRGB.icm')
```

```
P =
```

# iccread

---

```
Header: [1x1 struct]
TagTable: {17x3 cell}
Copyright: 'Copyright (c) 1999 Hewlett-Packard Company'
Description: [1x1 struct]
MediaWhitePoint: [0.9505 1 1.0891]
MediaBlackPoint: [0 0 0]
DeviceMfgDesc: [1x1 struct]
DeviceModelDesc: [1x1 struct]
ViewingCondDesc: [1x1 struct]
ViewingConditions: [1x1 struct]
Luminance: [76.0365 80 87.1246]
Measurement: [1x36 uint8]
Technology: [115 105 103 32 0 0 0 0 67 82 84 32]
MatTRC: [1x1 struct]
PrivateTags: {}
Filename: 'sRGB.icm'
```

The profile header provides general information about the profile such as its class, color space, and PCS. For example, to determine the source color space, view the `ColorSpace` field in the Header structure.

```
P.Header.ColorSpace
```

```
ans =
```

```
RGB
```

## See Also

`applycform`, `iccwrite`, `isicc`, `makecform`

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>     | Write ICC color profile to disk file                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax</b>      | <code>P_new = iccwrite(P, filename)</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p><code>P_new = iccwrite(P, filename)</code> writes the International Color Consortium (ICC) color profile data in structure <code>P</code> to the file specified by <code>filename</code>.</p> <p><code>P</code> is a structure representing an ICC profile in the data format returned by <code>iccread</code> and used by <code>makecform</code> and <code>applycform</code> to compute color-space transformations. <code>P</code> must contain all the tags and fields required by the ICC profile specification. Some fields may be inconsistent, however, because of interactive changes to the structure. For instance, the tag table may not be correct because tags may have been added, deleted, or modified since the tag table was constructed. <code>iccwrite</code> makes any necessary corrections to the profile structure before writing it to the file and returns this corrected structure in <code>P_new</code>.</p> <p><code>iccwrite</code> can write the color profile data using either Version 2 (ICC.1:2001-04) or Version 4 (ICC.1:2001-12) of the ICC specification, depending on the value of the <code>Version</code> field in the file profile header. If any required fields are missing, <code>iccwrite</code> errors. For more information about ICC profiles, visit the ICC web site, <a href="http://www.color.org">www.color.org</a>.</p> |
|                    | <hr/> <p><b>Note</b> <code>iccwrite</code> does not perform automatic conversions from one version of the ICC specification to another. Such conversions have to be done manually, by adding fields or modifying fields. Use <code>isicc</code> to validate a profile.</p> <hr/>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Note</b>        | Portions of the implementation of <code>iccwrite</code> are derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Example</b>     | <p>This example reads a profile into the MATLAB workspace and then exports it in a new file.</p> <pre>P = iccread('monitor.icm'); pmon = iccwrite(P, 'monitor2.icm');</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>See Also</b>    | <code>applycform</code> , <code>iccread</code> , <code>isicc</code> , <code>makecform</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

# idct2

---

**Purpose** Compute two-dimensional inverse discrete cosine transform

**Syntax**  
B = idct2(A)  
B = idct2(A,m,n)  
B = idct2(A,[m n])

**Description** B = idct2(A) returns the two-dimensional inverse discrete cosine transform (DCT) of A.

B = idct2(A,m,n) or B = idct2(A,[m n]) pads A with 0's to size m-by-n before transforming. If [m n] < size(A), idct2 crops A before transforming.

For any A, idct2(dct2(A)) equals A to within roundoff error.

**Class Support** The input matrix A can be of class double or of any numeric class. The output matrix B is of class double.

**Algorithm** idct2 computes the two-dimensional inverse DCT using

$$A_{mn} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \alpha_p \alpha_q B_{pq} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{array}{l} 0 \leq m \leq M-1 \\ 0 \leq n \leq N-1 \end{array}$$

$$\alpha_p = \begin{cases} 1/\sqrt{M}, & p = 0 \\ \sqrt{2/M}, & 1 \leq p \leq M-1 \end{cases} \quad \alpha_q = \begin{cases} 1/\sqrt{N}, & q = 0 \\ \sqrt{2/N}, & 1 \leq q \leq N-1 \end{cases}$$

**Example** Create a DCT matrix.

```
RGB = imread('autumn.tif');
I = rgb2gray(RGB);
J = dct2(I);
imshow(log(abs(J)),[]), colormap(jet), colorbar
```

Set values less than magnitude 10 in the DCT matrix to zero, then reconstruct the image using the inverse DCT function idct2.

```
J(abs(J)<10) = 0;
K = idct2(J);
figure, imshow(I)
```

---

```
figure, imshow(K,[0 255])
```

**See Also**

dct2, dctmtx, fft2, ifft2

**References**

- [1] Jain, A. K., *Fundamentals of Digital Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1989, pp. 150-153.
- [2] Pennebaker, W. B., and J. L. Mitchell, *JPEG: Still Image Data Compression Standard*, New York, Van Nostrand Reinhold, 1993.

# ifanbeam

---

**Purpose** Compute inverse fan-beam transform

**Syntax**

```
I = ifanbeam(F,D)
I = ifanbeam(...,param1,val1,param2,val2,...)
[I,H] = ifanbeam(...)
```

**Description** `I = ifanbeam(F,D)` reconstructs the image `I` from projection data in the two-dimensional array `F`. Each column of `F` contains fan-beam projection data at one rotation angle. `ifanbeam` assumes that the center of rotation is the center point of the projections, which is defined as `ceil(size(F,1)/2)`.

The fan-beam spread angles are assumed to be the same increments as the input rotation angles split equally on either side of zero. The input rotation angles are assumed to be stepped in equal increments to cover `[0:359]` degrees.

`D` is the distance from the fan-beam vertex to the center of rotation.

`I = ifanbeam(...,param1,val1,param2,val2,...)` specifies parameters that control various aspects of the `ifanbeam` reconstruction, described in the following table. Parameter names can be abbreviated, and case does not matter. Default values are in braces (`{}`).

| Parameter              | Description                                                                                                                                                                                                                               |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'FanCoverage'          | String specifying the range through which the beams are rotated.<br><code>{'cycle'}</code> — Rotate through the full range <code>[0,360)</code> .<br><code>'minimal'</code> — Rotate the minimum range necessary to represent the object. |
| 'FanRotationIncrement' | Positive real scalar specifying the increment of the rotation angle of the fan-beam projections, measured in degrees. See <code>fanbeam</code> for details.                                                                               |
| 'FanSensorGeometry'    | String specifying how sensors are positioned. See <code>fanbeam</code> for details.                                                                                                                                                       |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'FanSensorSpacing' | Positive real scalar specifying the spacing of the fan-beam sensors. Interpretation of the value depends on the setting of 'FanSensorGeometry'. See fanbeam for details.                                                                                                                                                                                                                                                                                                                       |
| 'Filter'           | String specifying the name of a filter. See iradon for details.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 'FrequencyScaling' | Scalar in the range (0,1] that modifies the filter by rescaling its frequency axis. See iradon for details.                                                                                                                                                                                                                                                                                                                                                                                    |
| 'Interpolation'    | String specifying an interpolation method. See iradon for details.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 'OutputSize'       | <p>Positive scalar specifying the number of rows and columns in the reconstructed image.</p> <p>If 'OutputSize' is not specified, ifanbeam determines the size automatically.</p> <p>If you specify 'OutputSize', ifanbeam reconstructs a smaller or larger portion of the image, but does not change the scaling of the data.</p> <p><b>Note:</b> If the projections were calculated with the fanbeam function, the reconstructed image might not be the same size as the original image.</p> |

[I,H] = ifanbeam(...) returns the frequency response of the filter in the vector H.

**Notes**

ifanbeam converts the fan-beam data to parallel beam projections and then uses the filtered back projection algorithm to perform the inverse Radon transform. The filter is designed directly in the frequency domain and then multiplied by the FFT of the projections. The projections are zero-padded to a power of 2 before filtering to prevent spatial domain aliasing and to speed up the FFT.

**Class Support**

The input arguments, F and D, can be double or single. All other numeric input arguments must be double. The output arguments are double.

**Example**

```
ph = phantom(128);
```

# ifanbeam

---

```
d = 100;
F = fanbeam(ph,d);
I = ifanbeam(F,d,'FanSensorSpacing',0.5);
imshow(ph), figure, imshow(I);
```

## See Also

fan2para, fanbeam, iradon, para2fan, phantom, radon

## References

[1] Kak, A. C., and M. Slaney, *Principles of Computerized Tomographic Imaging*, New York, NY, IEEE Press, 1988.

**Purpose**

Compute two-dimensional inverse fast Fourier transform

`ifft2` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference pages.

# ifftn

---

## **Purpose**

Compute N-dimensional inverse fast Fourier transform

`ifftn` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference pages.

**Purpose** Convert an image to a binary image, based on threshold

**Syntax**

```
BW = im2bw(I,level)
BW = im2bw(X,map,level)
BW = im2bw(RGB,level)
```

**Description** `im2bw` produces binary images from indexed, intensity, or RGB images. To do this, it converts the input image to grayscale format (if it is not already an intensity image), and then uses thresholding to convert this grayscale image to binary. The output binary image `BW` has values of 1 (white) for all pixels in the input image with luminance greater than `level` and 0 (black) for all other pixels. (Note that you specify `level` in the range [0,1], regardless of the class of the input image.)

`BW = im2bw(I,level)` converts the intensity image `I` to black and white.

`BW = im2bw(X,map,level)` converts the indexed image `X` with colormap `map` to black and white.

`BW = im2bw(RGB,level)` converts the RGB image `RGB` to black and white.

---

**Note** The function `graythresh` can be used to compute the level argument automatically.

---

**Class Support** The input image can be of class `uint8`, `uint16`, `single`, `int16`, or `double`, and must be nonsparse. The output image `BW` is of class `logical`.

**Example**

```
load trees
BW = im2bw(X,map,0.4);
imshow(X,map), figure, imshow(BW)
```



Image Courtesy of Susan Cohen



**See Also**

graythresh, ind2gray, rgb2gray

**Purpose** Rearrange image blocks into columns

**Syntax**

```
B = im2col(A,[m n],block_type)
B = im2col(A,[m n])
B = im2col(A,'indexed',...)
```

**Description** `B = im2col(A,[m n],block_type)` rearranges image blocks into columns. `block_type` is a string that can have one of these values. The default value is enclosed in braces (`{}`).

| Value       | Description                                                                                                                                                                                                                                                                                                                                     |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 'distinct'  | Rearranges each distinct $m$ -by- $n$ block in the image $A$ into a column of $B$ . <code>im2col</code> pads $A$ with 0's, if necessary, so its size is an integer multiple of $m$ -by- $n$ . If $A = [A_{11} \ A_{12}; A_{21} \ A_{22}]$ , where each $A_{i,j}$ is $m$ -by- $n$ , then $B = [A_{11}(:) \ A_{12}(:) \ A_{21}(:) \ A_{22}(:)]$ . |
| {'sliding'} | Converts each sliding $m$ -by- $n$ block of $A$ into a column of $B$ , with no zero padding. $B$ has $m*n$ rows and contains as many columns as there are $m$ -by- $n$ neighborhoods of $A$ . If the size of $A$ is $[mm \ nn]$ , then the size of $B$ is $(m*n)$ -by- $((mm-m+1)*(nn-n+1))$ .                                                  |

For the sliding block case, each column of  $B$  contains the neighborhoods of  $A$  reshaped as `nhood(:)` where `nhood` is a matrix containing an  $m$ -by- $n$  neighborhood of  $A$ . `im2col` orders the columns of  $B$  so that they can be reshaped to form a matrix in the normal way. For example, suppose you use a function, such as `sum(B)`, that returns a scalar for each column of  $B$ . You can directly store the result in a matrix of size  $(mm-m+1)$ -by- $(nn-n+1)$ , using these calls.

```
B = im2col(A,[m n],'sliding');
C = reshape(sum(B),mm-m+1,nn-n+1);
```

`B = im2col(A,'indexed',...)` processes  $A$  as an indexed image, padding with 0's if the class of  $A$  is `uint8`, or 1's if the class of  $A$  is `double`.

**Class Support** The input image  $A$  can be numeric or logical. The output matrix  $B$  is of the same class as the input image.

# im2col

---

## See Also

`blkproc`, `col2im`, `colfilt`, `nlfilter`

**Purpose** Convert image to double precision

**Syntax**

```
I2 = im2double(I)
RGB2 = im2double(RGB)
I = im2double(BW)
X2 = im2double(X, 'indexed')
```

**Description** `im2double` takes an image as input, and returns an image of class `double`. If the input image is of class `double`, the output image is identical to it. If the input image is not `double`, `im2double` returns the equivalent image of class `double`, rescaling or offsetting the data as necessary.

`I2 = im2double(I)` converts the intensity image `I` to double precision, rescaling the data if necessary.

`RGB2 = im2double(RGB)` converts the truecolor image `RGB` to double precision, rescaling the data if necessary.

`I = im2double(BW)` converts the binary image `BW` to a double-precision intensity image.

`X2 = im2double(X, 'indexed')` converts the indexed image `X` to double precision, offsetting the data if necessary.

**Class Support** Intensity and truecolor images can be `uint8`, `uint16`, `double`, `logical`, `single`, or `int16`. Indexed images can be `uint8`, `uint16`, `double` or `logical`. Binary input images must be `logical`. The output image is `double`.

**Example**

```
I1 = reshape(uint8(linspace(1,255,25)),[5 5])
I2 = im2double(I1)
```

**See Also** `double`, `im2single`, `im2int16`, `im2uint8`, `im2uint16`.

# im2int16

---

**Purpose** Convert image to 16-bit signed integers

**Syntax**

```
I2 = im2int16(I)
RGB2 = im2int16(RGB)
I = im2int16(BW)
```

**Description** `im2int16` takes an image as input and returns an image of class `int16`. If the input image is of class `int16`, the output image is identical to it. If the input image is not of class `int16`, `im2int16` returns the equivalent image of class `int16`, rescaling or offsetting the data as necessary.

`I2 = im2int16(I)` converts the intensity image `I` to `int16`, rescaling the data if necessary.

`RGB2 = im2int16(RGB)` converts the truecolor image `RGB` to `int16`, rescaling the data if necessary.

`I = im2int16(BW)` converts the binary image `BW` to an `int16` intensity image, changing false-valued elements to -32768 and true-valued elements to 32767.

**Class Support** Intensity and truecolor images can be `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. Binary input images must be `logical`. The output image is `int16`.

**Example**

```
I = reshape(linspace(0,1,20),[5 4])
I2 = im2int16(I)
```

**See Also** `im2double`, `im2single`, `im2uint8`, `im2uint16`, `int16`

**Purpose**

Convert image to Java image

`im2java` is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference pages.

# im2java2d

---

**Purpose** Convert image to Java buffered image

**Syntax**  
`jimage = im2java2d(I)`  
`jimage = im2java2d(X,MAP)`

**Description** `jimage = im2java2d(I)` converts the image `I` to an instance of the Java image class `java.awt.image.BufferedImage`. The image `I` can be an intensity (grayscale), RGB, or binary image.

`jimage = im2java2d(X,MAP)` converts the indexed image `X` with colormap `MAP` to an instance of the Java class `java.awt.image.BufferedImage`.

---

**Note** The `im2java2d` function works with the Java 2D API. The `im2java` function works with the Java Abstract Windowing Toolkit (AWT).

---

**Class Support** Intensity, indexed, and RGB input images can be of class `uint8`, `uint16`, or `double`. Binary input images must be of class `logical`.

**Example** This example reads an image into the MATLAB workspace and then uses `im2java2d` to convert it into an instance of the Java class `java.awt.image.BufferedImage`.

```
I = imread('moon.tif');
javaImage = im2java2d(I);
frame = javax.swing.JFrame;
icon = javax.swing.ImageIcon(javaImage);
label = javax.swing.JLabel(icon);
frame.getContentPane.add(label);
frame.pack
frame.show
```

**Purpose** Convert image to single precision

**Syntax**

```
I2 = im2single(I)
RGB2 = im2single(RGB)
I = im2single(BW)
X2 = im2single(X, 'indexed')
```

**Description** `im2single` takes an image as input and returns an image of class `single`. If the input image is of class `single`, the output image is identical to it. If the input image is not of class `single`, `im2single` returns the equivalent image of class `single`, rescaling or offsetting the data as necessary.

`I2 = im2single(I)` converts the intensity image `I` to `single`, rescaling the data if necessary.

`RGB2 = im2single(RGB)` converts the truecolor image `RGB` to `single`, rescaling the data if necessary.

`I = im2single(BW)` converts the binary image `BW` to a single-precision intensity image.

`X2 = im2single(X, 'indexed')` converts the indexed image `X` to single precision, offsetting the data if necessary.

**Class Support** Intensity and truecolor images can be `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. Indexed images can be `uint8`, `uint16`, `double` or `logical`. Binary input images must be `logical`. The output image is `single`.

**Example**

```
I = reshape(uint8(linspace(1,255,25)),[5 5])
I2 = im2single(I)
```

**See Also** `im2double`, `im2int16`, `im2uint8`, `im2uint16`, `single`

# im2uint16

---

**Purpose** Convert image to 16-bit unsigned integers

**Syntax**

```
I2 = im2uint16(I)
RGB2 = im2uint16(RGB)
I = im2uint16(BW)
X2 = im2uint16(X, 'indexed')
```

**Description** `im2uint16` takes an image as input and returns an image of class `uint16`. If the input image is of class `uint16`, the output image is identical to it. If the input image is not of class `uint16`, `im2uint16` returns the equivalent image of class `uint16`, rescaling or offsetting the data as necessary.

`I2 = im2uint16(I)` converts the intensity image `I` to `uint16`, rescaling the data if necessary.

`RGB2 = im2uint16(RGB)` converts the truecolor image `RGB` to `uint16`, rescaling the data if necessary.

`I = im2uint16(BW)` converts the binary image `BW` to a `uint16` intensity image, changing 1-valued elements to 65535.

`X2 = im2uint16(X, 'indexed')` converts the indexed image `X` to `uint16`, offsetting the data if necessary. If `X` is of class `double`, `max(X(:))` must be 65536 or less.

**Class Support** Intensity and truecolor images can be `uint8`, `uint16`, `double`, `logical`, `single`, or `int16`. Indexed images can be `uint8`, `uint16`, `double`, or `logical`. Binary input images must be `logical`. The output image is `uint16`.

**Example**

```
I = reshape(linspace(0,1,20),[5 4])
I2 = im2uint16(I)
```

**See Also** `im2uint8`, `double`, `im2double`, `uint8`, `uint16`, `imapprox`

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | Convert image to 8-bit unsigned integers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax</b>        | <pre>I2 = im2uint8(I) RGB2 = im2uint8(RGB) I = im2uint8(BW) X2 = im2uint8(X, 'indexed')</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>   | <p>im2uint8 takes an image as input and returns an image of class uint8. If the input image is of class uint8, the output image is identical to it. If the input image is not of class uint8, im2uint8 returns the equivalent image of class uint8, rescaling or offsetting the data as necessary.</p> <p>I2 = im2uint8(I) converts the intensity image I to uint8, rescaling the data if necessary.</p> <p>RGB2 = im2uint8(RGB) converts the truecolor image RGB to uint8, rescaling the data if necessary.</p> <p>I = im2uint8(BW) converts the binary image BW to a uint8 intensity image, changing 1-valued elements to 255</p> <p>X2 = im2uint8(X, 'indexed') converts the indexed image X to uint8, offsetting the data if necessary. Note that it is not always possible to convert an indexed image to uint8. If X is of class double, the maximum value of X must be 256 or less; if X is of class uint16, the maximum value of X must be 255 or less.</p> |
| <b>Class Support</b> | Intensity and truecolor images can be uint8, uint16, double, logical, single, or int16. Indexed images can be uint8, uint16, double, or logical. Binary input images must be logical. The output image is uint8.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Example</b>       | <pre>I = reshape(uint8(linspace(0,255,255)), [5 5]) I2 = im2uint8(I)</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>See Also</b>      | im2double, im2int16, im2single, im2uint16, uint8                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

# imabsdiff

---

**Purpose** Compute absolute difference of two images

**Syntax** `Z = imabsdiff(X,Y)`

**Description** `Z = imabsdiff(X,Y)` subtracts each element in array `Y` from the corresponding element in array `X` and returns the absolute difference in the corresponding element of the output array `Z`. `X` and `Y` are real, nonsparse numeric arrays with the same class and size. `Z` has the same class and size as `X` and `Y`. If `X` and `Y` are integer arrays, elements in the output that exceed the range of the integer type are truncated.

If `X` and `Y` are double arrays, you can use the expression `abs(X-Y)` instead of this function.

---

**Note** On Intel architecture processors, `imabsdiff` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only if arrays `X`, `Y`, and `Z` are of class `logical`, `uint8`, or `single`, and are of the same class.

---

## Examples

This example calculates the absolute difference between two `uint8` arrays. Note that the absolute value prevents negative values from being rounded to zero in the result, as they are with `imsubtract`.

```
X = uint8([255 10 75; 44 225 100]);
Y = uint8([50 50 50; 50 50 50]);
Z = imabsdiff(X,Y)
```

```
Z =
 205 40 25
 6 175 50
```

Display the absolute difference between a filtered image and the original.

```
I = imread('cameraman.tif');
J = uint8(filter2(fspecial('gaussian'), I));
K = imabsdiff(I,J);
imshow(K,[]) % [] = scale data automatically
```

**See Also**

`imadd`, `imcomplement`, `imdivide`, `imlincomb`, `immultiply`, `imsubtract`, `ippl`

# imadd

---

**Purpose** Add two images, or add a constant to an image

**Syntax** `Z = imadd(X,Y)`

**Description** `Z = imadd(X,Y)` adds each element in array `X` with the corresponding element in array `Y` and returns the sum in the corresponding element of the output array `Z`. `X` and `Y` are real, nonsparse numeric arrays with the same size and class, or `Y` is a scalar double. `Z` has the same size and class as `X`, unless `X` is logical, in which case `Z` is double.

If `X` and `Y` are integer arrays, elements in the output that exceed the range of the integer type are truncated, and fractional values are rounded.

If `X` and `Y` are double arrays, you can use the expression `X+Y` instead of this function.

---

**Note** On Intel architecture processors, `imadd` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated if arrays `X`, `Y`, and `Z` are of class `logical`, `uint8`, or `single` and are of the same class. IPPL is also activated if `Y` is a double scalar and arrays `X` and `Z` are `uint8`, `int16`, or `single` and are of the same class.

---

**Examples** Add two `uint8` arrays. Note the truncation that occurs when the values exceed 255.

```
X = uint8([255 0 75; 44 225 100]);
Y = uint8([50 50 50; 50 50 50]);
Z = imadd(X,Y)
Z =
```

```
 255 50 125
 94 255 150
```

Add two images together and specify an output class.

```
I = imread('rice.png');
J = imread('cameraman.tif');
K = imadd(I,J,'uint16');
imshow(K,[])
```

Add a constant to an image.

```
I = imread('rice.png');
J = imadd(I,50);
subplot(1,2,1), imshow(I)
subplot(1,2,2), imshow(J)
```

**See Also**

`imabsdiff`, `imcomplement`, `imdivide`, `imlincomb`, `immultiply`, `imsubtract`,  
`ippl`

# imadjust

---

**Purpose** Adjust image intensity values or colormap

**Syntax**

```
J = imadjust(I)
J = imadjust(I,[low_in; high_in],[low_out; high_out])
J = imadjust(...,gamma)
newmap = imadjust(map,[low_in high_in],[low_out high_out],gamma)
RGB2 = imadjust(RGB1,...)
```

**Description** `J = imadjust(I)` maps the values in intensity image `I` to new values in `J` such that 1% of data is saturated at low and high intensities of `I`. This increases the contrast of the output image `J`. This syntax is equivalent to `imadjust(I,stretchlim(I))`.

`J = imadjust(I,[low_in; high_in],[low_out; high_out])` maps the values in intensity image `I` to new values in `J` such that values between `low_in` and `high_in` map to values between `low_out` and `high_out`. Values below `low_in` and above `high_in` are clipped; that is, values below `low_in` map to `low_out`, and those above `high_in` map to `high_out`. You can use an empty matrix (`[]`) for `[low_in high_in]` or for `[low_out high_out]` to specify the default of `[0 1]`.

`J = imadjust(I,[low_in; high_in],[low_out; high_out],gamma)` maps the values in intensity image `I` to new values in `J`, where `gamma` specifies the shape of the curve describing the relationship between the values in `I` and `J`. If `gamma` is less than 1, the mapping is weighted toward higher (brighter) output values. If `gamma` is greater than 1, the mapping is weighted toward lower (darker) output values. If you omit the argument, `gamma` defaults to 1 (linear mapping).

`newmap = imadjust(map,[low_in; high_in],[low_out; high_out],gamma)` transforms the colormap associated with an indexed image. If `low_in`, `high_in`, `low_out`, `high_out`, and `gamma` are scalars, then the same mapping applies to red, green, and blue components. Unique mappings for each color component are possible when

- `low_in` and `high_in` are both 1-by-3 vectors.

- `low_out` and `high_out` are both 1-by-3 vectors, or `gamma` is a 1-by-3 vector.

The rescaled colormap `newmap` is the same size as `map`.

`RGB2 = imadjust(RGB1, ...)` performs the adjustment on each image plane (red, green, and blue) of the RGB image `RGB1`. As with the colormap adjustment, you can apply unique mappings to each plane.

---

**Note** If `high_out < low_out`, the output image is reversed, as in a photographic negative.

---

## Class Support

For syntax variations that include an input image (rather than a colormap), the input image can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image has the same class as the input image. For syntax variations that include a colormap, the input and output colormaps are of class `double`.

## Example

Adjust a low-contrast grayscale image.

```
I = imread('pout.tif');
J = imadjust(I);
imshow(I), figure, imshow(J)
```



Adjust the grayscale image, specifying the contrast limits.

```
K = imadjust(I,[0.3 0.7],[]);
figure, imshow(K)
```


Adjust an RGB image.

```
RGB1 = imread('football.jpg');
```

# imadjust

---

```
RGB2 = imadjust(RGB1,[.2 .3 0; .6 .7 1],[,]);
imshow(RGB1), figure, imshow(RGB2)
```



## See Also

brighten, histeq, stretchlim

**Purpose** Create Image Information tool

**Syntax**

```
imageinfo
imageinfo(h)
imageinfo(filename)
imageinfo(info)
imageinfo(himage,filename)
imageinfo(himage,info)
hfig = imageinfo(...)
```

**Description** `imageinfo` creates an Image Information tool associated with the image in the current figure. The tool displays in a separate figure information about the basic attributes of the target image. `imageinfo` gets the image attributes by querying the image object's CData.

The following table lists the basic image information included in the Image Information tool display. Note that the tool contains either four or six fields, depending on the type of image.

| Attribute Name  | Value                                                                                                                                                                                                                                                                                   |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Width (columns) | Number of columns in the image                                                                                                                                                                                                                                                          |
| Height (rows)   | Number of rows in the image                                                                                                                                                                                                                                                             |
| Class           | Data type used by the image, such as <code>uint8</code> .<br><br><b>Note:</b> For <code>single</code> or <code>int16</code> images, <code>imageinfo</code> returns a class value of <code>double</code> , because image objects convert CData of these classes to <code>double</code> . |
| Image type      | One of the image types identified by the Image Processing Toolbox: <code>'intensity'</code> , <code>'truecolor'</code> , <code>'binary'</code> , or <code>'indexed'</code> .                                                                                                            |

# imageinfo

---

| Attribute Name    | Value                                                                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minimum intensity | For intensity images, this value represents the lowest intensity value of any pixel.<br>For indexed images, this value represents the lowest index value into a color map.<br>Not included for 'binary' or 'truecolor' images.   |
| Maximum intensity | For intensity images, this value represents the highest intensity value of any pixel.<br>For indexed images, this value represents the highest index value into a color map.<br>Not included for 'binary' or 'truecolor' images. |

`imageinfo(h)` creates an Image Information tool associated with `h`, where `h` is a handle to a figure, axes, or image object.

`imageinfo(filename)` creates an Image Information tool containing image metadata from the graphics file `filename`. The image does not have to be displayed in a figure window. `filename` can be any file type that has been registered with an information function in the file formats registry, `imformats`, so its information can be read by `imfinfo`. `filename` can also be a DICOM file with information readable by `dicominfo`.

`imageinfo(info)` creates an Image Information tool containing the image metadata in the structure `info`. `info` is a structure returned by the functions `imfinfo` or `dicominfo`, or `info` can be a user-created structure.

`imageinfo(himage, filename)` creates an Image Information tool containing information about the basic attributes of the image specified by the handle `himage` and the image metadata from the graphics file `filename`.

`imageinfo(himage, info)` creates an Image Information tool containing information about the basic attributes of the image specified by the handle `himage` and the image metadata in the structure `info`.

`hfig=imageinfo(...)` returns a handle to the Image Information tool figure.

**Example**

```
imageinfo('peppers.png')

h = imshow('bag.png');
info = imfinfo('bag.png');
imageinfo(h,info);

imshow('trees.tif');
imageinfo;
```

**See Also**

`dicominfo`, `imattributes`, `imfinfo`, `imformats`, `imtool`

# imagemodel

---

**Purpose** Access to properties of an image relevant to its display

**Syntax** `imgmodel = imagemodel(himage)`

**Description** `imgmodel = imagemodel(himage)` create an image model object associated with the target image `himage`. `himage` is a handle to an image object or an array of handles to image objects.

`imagemodel` returns an image model object or, if `himage` is an array of image objects, an array of image model objects.

An image model object stores information about an image such as class, type, display range, width, height, minimum intensity value and maximum intensity value.

The image model object supports methods that you can use to access this information, get information about the pixels in an image, and perform special text formatting. The following lists these methods with a brief description. Use `methods(imgmodel)` to get a list of image model methods.

| Method                       | Description                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>getClassType</code>    | Returns a string indicating the class of the image.<br><br><code>str = getClassType(imgmodel)</code><br><br>where <code>imgmodel</code> is a valid image model and <code>str</code> is a text string, such as 'uint8'.                                                                                                                                                              |
| <code>getDisplayRange</code> | Returns a double array containing the minimum and maximum values of the display range for an intensity image. For image types other than intensity, the value returned is an empty array.<br><br><code>disp_range = getDisplayRange(imgmodel)</code><br><br>where <code>imgmodel</code> is a valid image model and <code>disp_range</code> is an array of doubles, such as [0 255]. |

| Method          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| getImageHeight  | <p>Returns a double scalar containing the number of rows.</p> <pre data-bbox="630 373 1105 399">height = getImageHeight(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>height</code> is a double scalar.</p>                                                                                                                                                                                                                                                                                                             |
| getImageType    | <p>Returns a text string indicating the image type.</p> <pre data-bbox="630 590 1031 616">str = getImageType(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>str</code> is one of the text strings 'intensity', 'truecolor', 'binary', or 'indexed'.</p>                                                                                                                                                                                                                                                                  |
| getImageWidth   | <p>Returns a double scalar containing the number of columns.</p> <pre data-bbox="615 841 1061 868">width = getImageWidth(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>width</code> is a double scalar.</p>                                                                                                                                                                                                                                                                                                             |
| getMinIntensity | <p>Returns the minimum value in the image calculated as <code>min(Image(:))</code>. For an intensity image, the value returned is the minimum intensity. For an indexed image, the value returned is the minimum index. For any other image type, the value returned is an empty array.</p> <pre data-bbox="615 1197 1105 1223">minval = getMinIntensity(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>minval</code> is a numeric value. The class of <code>minval</code> depends on the class of the target image.</p> |

# imagemodel

| Method                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>getMaxIntensity</code>    | <p>Returns the maximum value in the image calculated as <code>max(Image(:))</code>. For an intensity image, the value returned is the maximum intensity. For an indexed image, the value returned is the maximum index. For any other image type, the value returned is an empty array.</p> <pre>maxval = getMaxIntensity(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>maxval</code> is a numeric value. The class of <code>maxval</code> depends on the class of the target image.</p> |
| <code>getNumberFormatFcn</code> | <p>Returns the handle to a function that converts a numeric value into a string.</p> <pre>fun = getNumberFormatFcn(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model. <code>fun</code> is a handle to a function that accepts a numeric value and returns the value as a text string. For example, you can use this function to convert the numeric return value of the <code>getPixelValue</code> method into a text string.</p> <pre>str = fun(getPixelValue(imgmodel,100,100))</pre>                |
| <code>getPixelInfoString</code> | <p>Returns a text string containing value of the pixel at the location specified by row and column.</p> <pre>str = getPixelInfoString(imgmodel,row,column)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>row</code> and <code>column</code> are numeric scalar values. <code>str</code> is a character array. For example, for an RGB image, the method returns a text string such as <code>'[66 35 60]'</code>.</p>                                                                               |

| Method                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><code>getPixelRegionFormatFcn</code></p>   | <p>Returns a handle to a function that formats the value of a pixel into a text string.</p> <pre data-bbox="609 404 1169 430">fun = getPixelRegionFormatFcn(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model. <code>fun</code> is a handle to a function that accepts the location (<code>row</code>, <code>column</code>) of a pixel in the target image and returns the value of the pixel as a specially formatted text string. For example, when used with an RGB image, this function returns a text string of the form 'R:000 G:000 B:000' where 000 is the actual pixel value.</p> <pre data-bbox="602 704 877 730">str = fun(100,100)</pre> |
| <p><code>getPixelValue</code></p>             | <p>Returns the value of the pixel at the location specified by <code>row</code> and <code>column</code> as a numeric array.</p> <pre data-bbox="609 855 1199 881">val = getPixelValue(imgmodel,row, column)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>row</code> and <code>column</code> are numeric scalar values. The class of <code>val</code> depends on the class of the target image.</p>                                                                                                                                                                                                                                              |
| <p><code>getDefaultPixelInfoString</code></p> | <p>Returns a text string indicating the type of information returned in a pixel information string. This string can be used in place of actual pixel information values.</p> <pre data-bbox="609 1177 1199 1203">str = getDefaultPixelInfoString(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model. Depending on the image type, <code>str</code> can be the text string 'Intensity', '[R G B]', 'BW', or '&lt;Index&gt; [R G B]'.</p>                                                                                                                                                                                                               |

# imagemodel

| Method                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>getDefaultPixelRegionString</code> | <p>Returns a text string indicating the type of information displayed in the Pixel Region tool for each image type. This string can be used in place of actual pixel values.</p> <pre>str = getDefaultPixelRegionString(imgmodel)</pre> <p>where <code>imgmodel</code> is a valid image model. Depending on the image type, <code>str</code> can be the text string <code>'000'</code>, <code>'R:000 G:000 B:000'</code>, <code>'0'</code>, or <code>'&lt;000&gt; R:0.00 G:0.00 B:0.00'</code>.</p> |
| <code>getScreenPixelRGBValue</code>      | <p>Returns the screen display value of the pixel at the location specified by <code>ROW</code> and <code>COLUMN</code> as a double array.</p> <pre>val = getScreenPixelRGBValue(imgmodel,row, col)</pre> <p>where <code>imgmodel</code> is a valid image model and <code>row</code> and <code>col</code> are numeric scalar values. <code>val</code> is an array of doubles, such as <code>[0.2 0.5 0.3]</code>.</p>                                                                                |

## Note

`imagemodel` works by querying the image object's `CData`. For a single or `int16` image, the image object converts its `CData` to `double`. For example, in the case of `h = imshow(int16(ones(10)))`, `class(get(h, 'CData'))` returns `'double'`. Therefore, `getClassType(imgmodel)` returns `'double'`.

## Examples

This example creates an image model.

```
h = imshow('peppers.png');
im = imagemodel(h);

figure,subplot(1,2,1)
h1 = imshow('hestain.png');
subplot(1,2,2)
h2 = imshow('coins.png');
im = imagemodel([h1 h2]);
```

## See Also

`getimagemodel`

**Purpose** Approximate indexed image by one with fewer colors

**Syntax**

```
[Y,newmap] = imapprox(X,map,n)
[Y,newmap] = imapprox(X,map,tol)
Y = imapprox(X,map,newmap)
[...] = imapprox(...,dither_option)
```

**Description** `[Y,newmap] = imapprox(X,map,n)` approximates the colors in the indexed image `X` and associated colormap `map` by using minimum variance quantization. `imapprox` returns indexed image `Y` with colormap `newmap`, which has at most `n` colors.

`[Y,newmap] = imapprox(X,map,tol)` approximates the colors in `X` and `map` through uniform quantization. `newmap` contains at most  $(\text{floor}(1/\text{tol})+1)^3$  colors. `tol` must be between 0 and 1.0.

`Y = imapprox(X,map,newmap)` approximates the colors in `map` by using colormap mapping to find the colors in `newmap` that best match the colors in `map`.

`Y = imapprox(...,dither_option)` enables or disables dithering. `dither_option` is a string that can have one of these values. The default value is enclosed in braces `{}`.

| Value      | Description                                                                                           |
|------------|-------------------------------------------------------------------------------------------------------|
| {'dither'} | Dithers, if necessary, to achieve better color resolution at the expense of spatial resolution.       |
| 'nodither' | Maps each color in the original image to the closest color in the new map. No dithering is performed. |

**Class Support** The input image `X` can be of class `uint8`, `uint16`, or `double`. The output image `Y` is of class `uint8` if the length of `newmap` is less than or equal to 256. If the length of `newmap` is greater than 256, `Y` is of class `double`.

**Algorithm** `imapprox` uses `rgb2ind` to create a new colormap that uses fewer colors.

**Example** Approximate the indexed image `trees.tif` by another indexed image containing only 16 colors.

# imapprox

---

```
[X, map] = imread('trees.tif');
[Y, newmap] = imapprox(X, map, 16);
imshow(Y, newmap)
```

## See Also

[cmunique](#), [dither](#), [rgb2ind](#)

**Purpose** Return information about image attributes

**Syntax**

```
attrs = imattributes
attrs = imattributes(himage)
attrs = imattributes(imgmodel)
```

**Description** `attrs = imattributes` returns information about an image in the current figure. If the current figure does not contain an image, `imattributes` returns an empty array.

`attrs = imattributes(himage)` returns information about the image specified by `himage`, a handle to an image object. `imattributes` gets the image attributes by querying the image object's `CData`.

`imattributes` returns image attribute information in `attrs`, a 4-by-2 or 6-by-2 cell array, depending on the image type. The first column of the cell array contains the name of the attribute as a text string. The second column contains the value of the attribute, also represented as a text string. The following table lists these attributes in the order they appear in the cell array.

| Attribute Name  | Value                                                                                                                                                                                                                                                                                   |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Width (columns) | Number of columns in the image                                                                                                                                                                                                                                                          |
| Height (rows)   | Number of rows in the image                                                                                                                                                                                                                                                             |
| Class           | Data type used by the image, such as <code>uint8</code> .<br><br><b>Note:</b> For single or <code>int16</code> images, <code>imageinfo</code> returns a class value of <code>double</code> , because image objects convert <code>CData</code> of these classes to <code>double</code> . |
| Image type      | One of the image types identified by the Image Processing Toolbox: <code>'intensity'</code> , <code>'truecolor'</code> , <code>'binary'</code> , or <code>'indexed'</code> .                                                                                                            |

# imattributes

| Attribute Name    | Value                                                                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Minimum intensity | For intensity images, this value represents the lowest intensity value of any pixel.<br>For indexed images, this value represents the lowest index value into a color map.<br>Not included for 'binary' or 'truecolor' images.   |
| Maximum intensity | For intensity images, this value represents the highest intensity value of any pixel.<br>For indexed images, this value represents the highest index value into a color map.<br>Not included for 'binary' or 'truecolor' images. |

`attrs = imattributes(imgmodel)` returns information about the image represented by the image model object, `imgmodel`.

## Example

Retrieve the attributes of an intensity image.

```
h = imshow('liftingbody.png');
attrs = imattributes(h)
attrs =

 'Width (columns)' '512'
 'Height (rows)' '512'
 'Class' 'uint8'
 'Image type' 'intensity'
 'Minimum intensity' '0'
 'Maximum intensity' '255'
```

Creates an image model object for a truecolor image and then get the attributes from the image model object.

```
h = imshow('gantrycrane.png');
im = imagemodel(h);
attrs = imattributes(im)
attrs =
```

|                   |             |
|-------------------|-------------|
| 'Width (columns)' | '400'       |
| 'Height (rows)'   | '264'       |
| 'Class'           | 'uint8'     |
| 'Image type'      | 'truecolor' |

## See Also

`imagemodel`

# imbothat

---

**Purpose** Perform bottom-hat filtering

**Syntax**  
`IM2 = imbothat(IM,SE)`  
`IM2 = imbothat(IM,NHOOD)`

**Description** `IM2 = imbothat(IM,SE)` performs morphological bottom-hat filtering on the grayscale or binary input image, `IM`, returning the filtered image, `IM2`. The argument `SE` is a structuring element returned by the `strel` function. `SE` must be a single structuring element object, not an array containing multiple structuring element objects.

`IM2 = imbothat(IM,NHOOD)` performs morphological bottom-hat filtering where `NHOOD` is an array of 0's and 1's that specifies the size and shape of the structuring element. This is equivalent to `imbothat(IM,strel(NHOOD))`.

**Class Support** `IM` can be numeric or logical and must be nonsparse. The output image has the same class as the input image. If the input is binary (logical), then the structuring element must be flat.

**Example** Top-hat filtering and bottom-hat filtering can be used together to enhance contrast in an image.

- 1 Read the image into the MATLAB workspace.

```
I = imread('pout.tif');
imshow(I)
```



- 2 Create disk-shaped structuring element, needed for morphological processing.

```
se = strel('disk',3);
```

- 3 Add the original image I to the top-hat filtered image, and then subtract the bottom-hat filtered image.

```
J = imsubtract(imadd(I,imtophat(I,se)), imbothat(I,se));
figure, imshow(J)
```



## See Also

`imtophat`, `strel`

# imclearborder

---

**Purpose** Suppress light structures connected to image border

**Syntax**  
IM2 = imclearborder(IM)  
IM2 = imclearborder(IM,CONN)

**Description** IM2 = imclearborder(IM) suppresses structures that are lighter than their surroundings and that are connected to the image border. IM can be an intensity or binary image. The output image, IM2, is intensity or binary, respectively. The default connectivity is 8 for two dimensions, 26 for three dimensions, and `conndef(ndims(BW), 'maximal')` for higher dimensions.

---

**Note** For intensity images, imclearborder tends to reduce the overall intensity level in addition to suppressing border structures.

---

IM2 = imclearborder(IM,CONN) specifies the desired connectivity. CONN can have any of the following scalar values.

---

| Value                                   | Meaning                   |
|-----------------------------------------|---------------------------|
| <b>Two-dimensional connectivities</b>   |                           |
| 4                                       | 4-connected neighborhood  |
| 8                                       | 8-connected neighborhood  |
| <b>Three-dimensional connectivities</b> |                           |
| 6                                       | 6-connected neighborhood  |
| 18                                      | 18-connected neighborhood |
| 26                                      | 26-connected neighborhood |

---

Connectivity can also be defined in a more general way for any dimension by using for CONN a 3-by-3-by- ... -by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

**Note** A pixel on the edge of the input image might not be considered to be a border pixel if a nondefault connectivity is specified. For example, if `conn = [0 0 0; 1 1 1; 0 0 0]`, elements on the first and last row are not considered to be border pixels because, according to that connectivity definition, they are not connected to the region outside the image.

## Class Support

IM can be a numeric or logical array of any dimension, and it must be nonsparse and real. IM2 has the same class as IM.

## Example

The following examples use this simple binary image to illustrate the effect of `imclearborder` when you specify different connectivities.

```
BW =
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 1 0 0 1 1 1 0 0 0
 0 1 0 1 1 1 0 0 0
 0 0 0 1 1 1 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
```

Using a 4-connected neighborhood, the pixel at (5,2) is not considered connected to the border pixel (4,1), so it is not cleared.

```
BWc1 = imclearborder(BW,4)
BWc1 =
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 1 1 1 0 0 0
 0 1 0 1 1 1 0 0 0
 0 0 0 1 1 1 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0
```

# imclearborder

---

Using an 8-connected neighborhood, pixel (5, 2) is considered connected to pixel (4, 1) so both are cleared.

```
BWc2 = imclearborder(BW,8)
```

```
BWc2 =
```

```
0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0
0 0 0 1 1 1 0 0 0
0 0 0 1 1 1 0 0 0
0 0 0 1 1 1 0 0 0
0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0
```

## Algorithm

`imclearborder` uses morphological reconstruction where

- Mask image is the input image.
- Marker image is zero everywhere except along the border, where it equals the mask image.

## See Also

`conndef`

## Reference


[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer, 1999, pp. 164-165.

- Purpose** Close an image
- Syntax**
- ```
IM2 = imclose(IM,SE)
IM2 = imclose(IM,NHOOD)
```
- Description**
- `IM2 = imclose(IM,SE)` performs morphological closing on the grayscale or binary image `IM`, returning the closed image, `IM2`. The structuring element, `SE`, must be a single structuring element object, as opposed to an array of objects.
- `IM2 = imclose(IM,NHOOD)` performs closing with the structuring element `strel(NHOOD)`, where `NHOOD` is an array of 0's and 1's that specifies the structuring element neighborhood.
- Class Support** `IM` can be any numeric or logical class and any dimension, and must be nonsparse. If `IM` is logical, then `SE` must be flat. `IM2` has the same class as `IM`.

Example This example uses `imclose` to join the circles in the image together by filling in the gaps between them and by smoothening their outer edges.

- 1 Read the image into the MATLAB workspace and view it.

```
originalBW = imread('circles.png');
imshow(originalBW);
```



- 2 Create a disk-shaped structuring element. Use a disk structuring element to preserve the circular nature of the object. Specify a radius of 10 pixels so that the largest gap gets filled.

```
se = strel('disk',10);
```

imclose

- 3 Perform a morphological close operation on the image.

```
closeBW = imclose(originalBW,se);  
figure, imshow(closeBW)
```



See Also

`imdilate`, `imerode`, `imopen`, `strel`

Purpose Complement image

Syntax `IM2 = imcomplement(IM)`

Description `IM2 = imcomplement(IM)` computes the complement of the image `IM`. `IM` can be a binary, intensity, or RGB image. `IM2` has the same class and size as `IM`.

In the complement of a binary image, zeros become ones and ones become zeros; black and white are reversed. In the complement of an intensity or RGB image, each pixel value is subtracted from the maximum pixel value supported by the class (or 1.0 for double-precision images) and the difference is used as the pixel value in the output image. In the output image, dark areas become lighter and light areas become darker.

If `IM` is an intensity or RGB image of class `double`, you can use the expression `1-IM` instead of this function. If `IM` is a binary image, you can use the expression `~IM` instead of this function.

Examples Create the complement of a `uint8` array.

```
X = uint8([ 255 10 75; 44 225 100]);
X2 = imcomplement(X)
X2 =
     0    245    180
    211     30    155
```


Reverse black and white in a binary image.

```
bw = imread('text.png');
bw2 = imcomplement(bw);
subplot(1,2,1),imshow(bw)
subplot(1,2,2),imshow(bw2)
```


Create the complement of an intensity image.

```
I = imread('glass.png');
J = imcomplement(I);
imshow(I), figure, imshow(J)
```

imcomplement



Original Image



Complement Image

See Also

`imabsdiff`, `imadd`, `imdivide`, `imlincomb`, `immultiply`, `imsubtract`

Purpose Create a contour plot of image data

Syntax

```
imcontour(I)
imcontour(I,n)
imcontour(I,v)
imcontour(x,y,...)
imcontour(...,LineStyle)
[C,h] = imcontour(...)
```

Description

`imcontour(I)` draws a contour plot of the intensity image `I`, automatically setting up the axes so their orientation and aspect ratio match the image.

`imcontour(I,n)` draws a contour plot of the intensity image `I`, automatically setting up the axes so their orientation and aspect ratio match the image. `n` is the number of equally spaced contour levels in the plot; if you omit the argument, the number of levels and the values of the levels are chosen automatically.

`imcontour(I,v)` draws a contour plot of `I` with contour lines at the data values specified in vector `v`. The number of contour levels is equal to `length(v)`.

`imcontour(x,y,...)` uses the vectors `x` and `y` to specify the x - and y -axis limits.

`imcontour(...,LineStyle)` draws the contours using the line type and color specified by `LineStyle`. Marker symbols are ignored.


`[C,h] = imcontour(...)` returns the contour matrix `C` and a vector of handles to the objects in the plot. (The objects are actually patches, and the lines are the edges of the patches.) You can use the `clabel` function with the contour matrix `C` to add contour labels to the plot.

Class Support The input image can be of class `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`.

Example

```
I = imread('circuit.tif');
imcontour(I,3)
```

imcontour



See Also

`clabel`, `contour`, `LineStyleSpec` in the MATLAB Function Reference

Purpose Adjust contrast tool

Syntax
`imcontrast`
`imcontrast(h)`
`h = imcontrast(h)`

Description `imcontrast` creates an Adjust Contrast tool associated with the intensity image in the current figure, called the target image. If the figure contains multiple images, the target image is the first intensity image in the figure (last created). `imcontrast` creates the tool in a separate figure window.

The Adjust Contrast tool is an interactive contrast and brightness adjustment tool. When an image contains more gray values than can be displayed at one time, you can use this tool to choose which values to display. For images with limited contrast, you can use this tool to stretch the image histogram to fill the dynamic range.

Note The Adjust Contrast tool only works with intensity (grayscale) images. `imcontrast` adjusts the `CLim` property of the axes containing the image and does not work with RGB images. Changing the `CLim` of indexed images may produce unexpected results.

`imcontrast(h)` creates the Adjust Contrast tool associated with the image specified by the handle `h`. `h` can be a handle to a figure, axes, uipanel, or image object. If `h` is a handle to a figure or axes, `imcontrast` associates the tool with the first image found in the first figure or axes.

`h=imcontrast(...)` returns a handle to the Adjust Contrast tool figure.

When there are more possible gray values in the image than can be displayed at one time, this tool helps you interactively pick which values to display. For images with limited contrast this tool also helps you stretch the image histogram to fill the dynamic range.

Remarks The Adjust Contrast tool presents a scaled histogram of pixel values (overly represented pixel values are truncated for clarity). Dragging on the left red bar in the histogram display changes the minimum value. The minimum value

imcontrast

(and any value less than the minimum) displays as black. Dragging on the right red bar in the histogram changes the maximum value. The maximum value (and any value greater than the maximum) displays as white. Values in between the red bars display as intermediate shades of gray.

Together the minimum and maximum values create a "window". Stretching the window reduces contrast. Shrinking the window increases contrast. Changing the center of the window changes the brightness of the image. It is possible to manually enter the minimum, maximum, width, and center values for the window. Changing one value automatically updates the other values and the image.

Window/Level Interactivity

Clicking and dragging the mouse within the target image interactively changes the image's window values. Dragging the mouse horizontally from left to right changes the window width (i.e., contrast). Dragging the mouse vertically up and down changes the window center (i.e., brightness). Holding down the CTRL key when clicking accelerates changes. Holding down the SHIFT key slows the rate of change. Keys must be pressed before clicking and dragging.

Example

```
imshow('pout.tif')
imcontrast(gca)
```

See Also

`imadjust`, `imtool`, `stretchlim`

Purpose

Crop an image

Syntax

```
I2 = imcrop(I)
X2 = imcrop(X,map)
RGB2 = imcrop(RGB)

I2 = imcrop(I,rect)
X2 = imcrop(X,map,rect)
RGB2 = imcrop(RGB,rect)

[...] = imcrop(x,y,...)
[A,rect] = imcrop(...)
[x,y,A,rect] = imcrop(...)
```

Description

`imcrop` crops an image to a specified rectangle. In the syntaxes below, `imcrop` displays the input image and waits for you to specify the crop rectangle with the mouse.

```
I2 = imcrop(I)
X2 = imcrop(X,map)
RGB2 = imcrop(RGB)
```

If you omit the input arguments, `imcrop` operates on the image in the current axes.

To specify the rectangle,

- For a single-button mouse, press the mouse button and drag to define the crop rectangle. Finish by releasing the mouse button.
- For a two- or three-button mouse, press the left mouse button and drag to define the crop rectangle. Finish by releasing the mouse button.

If you hold down the **Shift** key while dragging, or if you press the right mouse button on a two- or three-button mouse, `imcrop` constrains the bounding rectangle to be a square.

When you release the mouse button, `imcrop` returns the cropped image in the supplied output argument. If you do not supply an output argument, `imcrop` displays the output image in a new figure.

You can also specify the cropping rectangle noninteractively, using these syntaxes

```
I2 = imcrop(I,rect)
X2 = imcrop(X,map,rect)
RGB2 = imcrop(RGB,rect)
```

`rect` is a four-element vector with the form `[xmin ymin width height]`; these values are specified in spatial coordinates.

To specify a nondefault spatial coordinate system for the input image, precede the other input arguments with two, two-element vectors specifying the `XData` and `YData`. For example:

```
[...] = imcrop(x,y,...)
```

If you supply additional output arguments, `imcrop` returns information about the selected rectangle and the coordinate system of the input image. For example:

```
[A,rect] = imcrop(...)
[x,y,A,rect] = imcrop(...)
```

`A` is the output image. `x` and `y` are the `XData` and `YData` of the input image.

Class Support

If you specify `RECT` as an input argument, the input image can be logical or numeric, and must be real and nonsparse. `RECT` is of class `double`.

If you do not specify `RECT` as an input argument, `imcrop` calls `imshow`. `imshow` expects `I` to be `logical`, `uint8`, `uint16`, `int16`, `single`, or `double`. `RGB` can be `uint8`, `int16`, `uint16`, `single`, or `double`. `X` can be `logical`, `uint8`, `uint16`, `single`, or `double`. The input image must be real and nonsparse.

If you specify an image as an input argument, the output image has the same class as the input image.

If you don't specify an image as an input argument, i.e., you call `imcrop` with no input arguments, the output image has the same class as the input image except for the `int16` or `single`. The output image is `double` if the input image is `int16` or `single`.


Remarks

Because `rect` is specified in terms of spatial coordinates, the width and height elements of `rect` do not always correspond exactly with the size of the output

image. For example, suppose `rect` is `[20 20 40 30]`, using the default spatial coordinate system. The upper-left corner of the specified rectangle is the center of the pixel (20,20) and the lower-right corner is the center of the pixel (50,60). The resulting output image is 31-by-41, not 30-by-40, because the output image includes all pixels in the input image that are completely *or partially* enclosed by the rectangle.

Example

```
I = imread('circuit.tif');  
I2 = imcrop(I,[75 68 130 112]);  
imshow(I), figure, imshow(I2)
```



See Also

`zoom`

imdilate

Purpose Dilate image

Syntax

```
IM2 = imdilate(IM,SE)
IM2 = imdilate(IM,NHOOD)
IM2 = imdilate(IM,SE,PACKOPT)
IM2 = imdilate(...,PADOPT)
```

Description `IM2 = imdilate(IM,SE)` dilates the grayscale, binary, or packed binary image `IM`, returning the dilated image, `IM2`. The argument `SE` is a structuring element object, or array of structuring element objects, returned by the `strel` function.

If `IM` is logical and the structuring element is flat, `imdilate` performs binary dilation; otherwise, it performs grayscale dilation. If `SE` is an array of structuring element objects, `imdilate` performs multiple dilations of the input image, using each structuring element in `SE` in succession.

`IM2 = imdilate(IM,NHOOD)` dilates the image `IM`, where `NHOOD` is a matrix of 0's and 1's that specifies the structuring element neighborhood. This is equivalent to the syntax `imdilate(IM,strel(NHOOD))`. The `imdilate` function determines the center element of the neighborhood by `floor((size(NHOOD)+1)/2)`.

`IM2 = imdilate(IM,SE,PACKOPT)` or `imdilate(IM,NHOOD,PACKOPT)` specifies whether `IM` is a packed binary image. `PACKOPT` can have either of the following values. Default value is enclosed in braces `{}`.

Value	Description
'ispacked'	IM is treated as a packed binary image as produced by <code>bwpack</code> . IM must be a 2-D <code>uint32</code> array and <code>SE</code> must be a flat 2-D structuring element. If the value of <code>PACKOPT</code> is 'ispacked', <code>PADOPT</code> must be 'same'.
{ 'notpacked' }	IM is treated as a normal array.

`IM2 = imdilate(..., PADOPT)` specifies the size of the output image. `PADOPT` can have either of the following values. Default value is enclosed in braces (`{}`).

Value	Description
<code>{ 'same' }</code>	Make the output image the same size as the input image. If the value of <code>PACKOPT</code> is <code>'ispacked'</code> , <code>PADOPT</code> must be <code>'same'</code> .
<code>'full'</code>	Compute the full dilation.

`PADOPT` is analogous to the optional `SHAPE` argument to the `conv2` and `filter2` functions.

Class Support

`IM` can be logical or numeric and must be real and nonsparse. It can have any dimension. If `IM` is logical, `SE` must be flat. The output has the same class as the input. If the input is packed binary, then the output is also packed binary.

Examples

This example dilates a binary image with a vertical line structuring element.

```
bw = imread('text.png');
se = strel('line',11,90);
bw2 = imdilate(bw,se);
imshow(bw), title('Original')
figure, imshow(bw2), title('Dilated')
```



This example dilates a grayscale image with a rolling ball structuring element.

```
I = imread('cameraman.tif');
```

imdilate

```
se = strel('ball',5,5);  
I2 = imdilate(I,se);  
imshow(I), title('Original')  
figure, imshow(I2), title('Dilated')
```



To determine the domain of the composition of two flat structuring elements, dilate the scalar value 1 with both structuring elements in sequence, using the 'full' option.

```
se1 = strel('line',3,0)  
se1 =  
  
Flat STREL object containing 3 neighbors.  
Neighborhood:  
    1    1    1  
  
se2 = strel('line',3,90)  
se2 =  
  
Flat STREL object containing 3 neighbors.  
Neighborhood:  
    1  
    1  
    1  
  
composition = imdilate(1,[se1 se2],'full')
```

```
composition =  
    1    1    1  
    1    1    1  
    1    1    1
```

Algorithm

imdilate automatically takes advantage of the decomposition of a structuring element object (if it exists). Also, when performing binary dilation with a structuring element object that has a decomposition, imdilate automatically uses binary image packing to speed up the dilation.

Dilation using bit packing is described in [2].

See Also

bwpack, bwunpack, conv2, filter2, imclose, imerode, imopen, strel

References

[1] Haralick, R.M., and L. G. Shapiro, *Computer and Robot Vision*, Vol. I, Addison-Wesley, 1992, pp. 158-205.

[2] van den Boomgaard and van Balen, "Image Transforms Using Bitmapped Binary Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, No. 3, May, 1992, pp. 254-258.

imshow


Purpose Display range tool

Syntax

```
imshow  
imshow(h)  
imshow(hparent,himage)  
hpanel=imshow(...)
```

Description `imshow` creates a display range tool in the current figure. The display range tool shows the display range of the intensity image or images in the figure.

The tool is a `uipanel` object, positioned in the lower-right corner of the figure. It contains the text string “Display range:” followed by the display range values for the image, as shown in the following figure.



For an indexed, truecolor, or binary image, the display range is not applicable and is set to empty (`[]`).

`imshow(h)` Creates a display range tool in the figure specified by the handle `h`, where `h` is a handle to an image, axes, `uipanel`, or figure object. Axes, `uipanel`, or figure objects must contain at least one image object.

`imshow(hparent,himage)` creates a display range tool in `hparent` that shows the display range of `himage`. `himage` is a handle to an image or an array of image handles. `hparent` is a handle to the figure or `uipanel` object that contains the display range tool.

`hpanel=imshow(...)` returns a handle to the display range tool `uipanel`.

Note The display range tool can work with multiple images in a figure. When the cursor is not in an image in a figure, the display range tool displays the text string, `[black white]`.

Examples Display an image and include the Display Range tool.

```
imshow('bag.png');  
imshow;
```

Import a 16-bit DICOM image and display it with its default range and scaled range in the same figure.

```
dcm = dicomread('CT-MON02-16-ankle.dcm');  
subplot(1,2,1), imshow(dcm);  
subplot(1,2,2), imshow(dcm, []);  
imshow;
```

See also

`imtool`

imdivide

Purpose Divide one image into another, or divide an image by a constant

Syntax `Z = imdivide(X,Y)`

Description `Z = imdivide(X,Y)` divides each element in the array `X` by the corresponding element in array `Y` and returns the result in the corresponding element of the output array `Z`. `X` and `Y` are real, nonsparse numeric arrays with the same size and class, or `Y` can be a scalar double. `Z` has the same size and class as `X` and `Y`.

If `X` is an integer array, elements in the output that exceed the range of integer type are truncated, and fractional values are rounded.

If `X` and `Y` are double arrays, you can use the expression `X ./ Y` instead of this function.

Note On Intel architecture processors, `imdivide` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only if arrays `X` and `Y` are of class `uint8`, `int16`, or `single` and are of the same size and class.

Example Divide two `uint8` arrays. Note that fractional values greater than or equal to 0.5 are rounded up to the nearest integer.

```
X = uint8([ 255 10 75; 44 225 100]);
Y = uint8([ 50 20 50; 50 50 50 ]);
Z = imdivide(X,Y)
Z =
     5     1     2
     1     5     2
```

Estimate and divide out the background of the rice image.

```
I = imread('rice.png');
background = imopen(I,strel('disk',15));
Ip = imdivide(I,background);
imshow(Ip,[])
```

Divide an image by a constant factor.

```
I = imread('rice.png');  
J = imdivide(I,2);  
subplot(1,2,1), imshow(I)  
subplot(1,2,2), imshow(J)
```

See Also

`imabsdiff`, `imadd`, `imcomplement`, `imlincomb`, `immultiply`, `imsubtract`, `ippl`

imerode

Purpose Erode image

Syntax

```
IM2 = imerode(IM,SE)
IM2 = imerode(IM,NHOOD)
IM2 = imerode(IM,SE,PACKOPT,M)
IM2 = imerode(...,PADOPT)
```

Description `IM2 = imerode(IM,SE)` erodes the grayscale, binary, or packed binary image `IM`, returning the eroded image `IM2`. The argument `SE` is a structuring element object or array of structuring element objects returned by the `strel` function.

If `IM` is logical and the structuring element is flat, `imerode` performs binary dilation; otherwise it performs grayscale erosion. If `SE` is an array of structuring element objects, `imerode` performs multiple erosions of the input image, using each structuring element in `SE` in succession.

`IM2 = imerode(IM,NHOOD)` erodes the image `IM`, where `NHOOD` is an array of 0's and 1's that specifies the structuring element neighborhood. This is equivalent to the syntax `imerode(IM,strel(NHOOD))`. The `imerode` function determines the center element of the neighborhood by `floor((size(NHOOD)+1)/2)`

`IM2 = imerode(IM,SE,PACKOPT,M)` or `imerode(IM,NHOOD,PACKOPT,M)` specifies whether `IM` is a packed binary image and, if it is, provides the row dimension `M` of the original unpacked image. `PACKOPT` can have either of the following values. Default value is enclosed in braces (`{}`).

Value	Description
'ispacked'	<code>IM</code> is treated as a packed binary image as produced by <code>bwpack</code> . <code>IM</code> must be a 2-D <code>uint32</code> array and <code>SE</code> must be a flat 2-D structuring element.
{ 'notpacked' }	<code>IM</code> is treated as a normal array.

If `PACKOPT` is 'ispacked', you must specify a value for `M`.

`IM2 = imerode(...,PADOPT)` specifies the size of the output image. `PADOPT` can have either of the following values. Default value is enclosed in braces (`{}`).

Value	Description
<code>{ 'same' }</code>	Make the output image the same size as the input image. If the value of <code>PACKOPT</code> is <code>'ispacked'</code> , <code>PADOPT</code> must be <code>'same'</code> .
<code>'full'</code>	Compute the full erosion.

`PADOPT` is analogous to the `SHAPE` input to the `CONV2` and `FILTER2` functions.



Class Support

`IM` can be numeric or logical and it can be of any dimension. If `IM` is logical and the structuring element is flat, the output image is logical; otherwise the output image has the same class as the input. If the input is packed binary, then the output is also packed binary.

Examples

This example erodes a binary image with a disk structuring element.

```
originalBW = imread('circles.png');
se = strel('disk',11);
erodedBW = imerode(originalBW,se);
imshow(originalBW), figure, imshow(erodedBW)
```



This example erodes a grayscale image with a rolling ball.

```
I = imread('cameraman.tif');
se = strel('ball',5,5);
```

```
I2 = imerode(I,se);  
imshow(I), title('Original')  
figure, imshow(I2), title('Eroded')
```



Algorithm Notes

`imerode` automatically takes advantage of the decomposition of a structuring element object (if a decomposition exists). Also, when performing binary dilation with a structuring element object that has a decomposition, `imerode` automatically uses binary image packing to speed up the dilation.

Erosion using bit packing is described in [2].

See Also

`bwpack`, `bwunpack`, `conv2`, `filter2`, `imclose`, `imdilate`, `imopen`, `strel`

References

[1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Vol. I, Addison-Wesley, 1992, pp. 158-205.

[2] van den Boomgaard and van Balen, "Image Transforms Using Bitmapped Binary Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, No. 3, May, 1992, pp. 254-258.

Purpose Extended-maxima transform

Syntax
`BW = imextendedmax(I,H)`
`BW = imextendedmax(I,H,CONN)`

Description `BW = imextendedmax(I,H)` computes the extended-maxima transform, which is the regional maxima of the H-maxima transform. H is a nonnegative scalar. Regional maxima are connected components of pixels with the same intensity value, t , whose external boundary pixels all have a value less than t .

By default, `imextendedmax` uses 8-connected neighborhoods for 2-D images and 26-connected neighborhoods for 3-D images. For higher dimensions, `imextendedmax` uses `conndef(ndims(I), 'maximal')`.

`BW = imextendedmax(I,H,CONN)` computes the extended-maxima transform, where `CONN` specifies the connectivity. `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood


Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support I can be of any nonsparse numeric class and any dimension. BW has the same size as I and is always logical.

imextendedmax

Example

```
I = imread('glass.png');  
BW = imextendedmax(I,80);  
imshow(I), figure, imshow(BW)
```



Original Image



Extended Maxima Image

See Also

conndef, imextendedmin, imreconstruct

Reference

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

Purpose Extended-minima transform

Syntax
`BW = imextendedmin(I,h)`
`BW = imextendedmin(I,h,CONN)`

Description `BW = imextendedmin(I,h)` computes the extended-minima transform, which is the regional minima of the H-minima transform. `h` is a nonnegative scalar. Regional minima are connected components of pixels with the same intensity value, t , whose external boundary pixels all have a value greater than t .

By default, `imextendedmin` uses 8-connected neighborhoods for 2-D images, and 26-connected neighborhoods for 3-D images. For higher dimensions, `imextendedmin` uses `conndef(ndims(I), 'maximal')`.

`BW = imextendedmin(I,h,CONN)` computes the extended-minima transform, where `CONN` specifies the connectivity. `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood


Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support `I` can be of any nonsparse numeric class and any dimension. `BW` has the same size as `I` and is always logical.

imextendedmin

Example

```
I = imread('glass.png');  
BW = imextendedmin(I,50);  
imshow(I), figure, imshow(BW)
```



Original Image



Extended Minima Image

See Also

`conndef`, `imextendedmax`, `imreconstruct`

Reference

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

Purpose

Fill image regions

Syntax

```
BW2 = imfill(BW,locations)
BW2 = imfill(BW,'holes')
I2 = imfill(I)

BW2 = imfill(BW)
[BW2 locations] = imfill(BW)

BW2 = imfill(BW,locations,CONN)
BW2 = imfill(BW,CONN,'holes')
I2 = imfill(I,CONN)
```

Description

`BW2 = imfill(BW,locations)` performs a flood-fill operation on background pixels of the binary image `BW`, starting from the points specified in `locations`. If `locations` is a `P`-by-1 vector, it contains the linear indices of the starting locations. If `locations` is a `P`-by-`ndims`(`BW`) matrix, each row contains the array indices of one of the starting locations.

`BW2 = imfill(BW, 'holes')` fills holes in the binary image `BW`. A hole is a set of background pixels that cannot be reached by filling in the background from the edge of the image.

`I2 = imfill(I)` fills holes in the intensity image `I`. In this case, a hole is an area of dark pixels surrounded by lighter pixels.

Interactive Use

`BW2 = imfill(BW)` displays the binary image `BW` on the screen and lets you select the starting locations using the mouse. Click the mouse button to add points. Press **Backspace** or **Delete** to remove the previously selected point. A shift-click, right-click, or double-click selects a final point and then starts the fill operation; pressing **Return** finishes the selection without adding a point.

Note `imfill` supports interactive use only for 2-D images.

`[BW2,locations] = imfill(BW)` lets you select the starting points selected using the mouse, returning the locations of points in `locations`. `locations` is a vector of linear indices into the input image.

Specifying Connectivity

By default, `imfill` uses 4-connected background neighbors for 2-D inputs and 6-connected background neighbors for 3-D inputs. For higher dimensions the default background connectivity is determined by using `conndef(NUM_DIMS, 'minimal')`. You can override the default connectivity with these syntaxes:

```
BW2 = imfill(BW,LOCATIONS,CONN)
BW2 = imfill(BW,CONN,'holes')
I2 = imfill(I,CONN)
```

To override the default connectivity and interactively specify the starting locations, use this syntax:

```
BW2 = imfill(BW,0,CONN)
```

`CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support

The input image can be numeric or logical, and it must be real and nonsparse. It can have any dimension. The output image has the same class as the input image.

Examples


Fill in the background of a binary image from a specified starting location.

```
BW1 = logical([1 0 0 0 0 0 0 0
               1 1 1 1 1 0 0 0
               1 0 0 0 1 0 1 0
               1 0 0 0 1 1 1 0
               1 1 1 1 0 1 1 1
               1 0 0 1 1 0 1 0
               1 0 0 0 1 0 1 0
               1 0 0 0 1 1 1 0]);
```

```
BW2 = imfill(BW1,[3 3],8)
```

Fill in the holes of a binary image.

```
BW4 = im2bw(imread('coins.png'));
BW5 = imfill(BW4,'holes');
imshow(BW4), figure, imshow(BW5)
```



Original Image


Filled Image

Fill in the holes of an intensity image.

```
I = imread('tire.tif');
I2 = imfill(I,'holes');
figure, imshow(I), figure, imshow(I2)
```



Original Image



Filled Image

Algorithm

`imfill` uses an algorithm based on morphological reconstruction [1].

See Also

`bwselect`, `imreconstruct`, `roifill`

Reference

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 173-174.

Purpose Multidimensional image filtering

Syntax
`B = imfilter(A,H)`
`B = imfilter(A,H,option1,option2,...)`

Description `B = imfilter(A,H)` filters the multidimensional array `A` with the multidimensional filter `H`. The array `A` can be a nonsparse numeric array of any class and dimension. The result `B` has the same size and class as `A`.

Each element of the output `B` is computed using double-precision floating point. If `A` is an integer array, then output elements that exceed the range of the integer type are truncated, and fractional values are rounded.

`B = imfilter(A,H,option1,option2,...)` performs multidimensional filtering according to the specified options. Option arguments can have the following values.

Boundary Options

Option	Description
<code>X</code>	Input array values outside the bounds of the array are implicitly assumed to have the value <code>X</code> . When no boundary option is specified, <code>imfilter</code> uses <code>X = 0</code> .
<code>'symmetric'</code>	Input array values outside the bounds of the array are computed by mirror-reflecting the array across the array border.
<code>'replicate'</code>	Input array values outside the bounds of the array are assumed to equal the nearest array border value.
<code>'circular'</code>	Input array values outside the bounds of the array are computed by implicitly assuming the input array is periodic.

Output Size Options

Option	Description
'same'	The output array is the same size as the input array. This is the default behavior when no output size options are specified.
'full'	The output array is the full filtered result, and so is larger than the input array.

Correlation and Convolution Options

Option	Description
'corr'	<code>imfilter</code> performs multidimensional filtering using correlation, which is the same way that <code>filter2</code> performs filtering. When no correlation or convolution option is specified, <code>imfilter</code> uses correlation.
'conv'	<code>imfilter</code> performs multidimensional filtering using convolution.

N-D convolution is related to N-D correlation by a reflection of the filter matrix.

Note On Intel architecture processors, `imfilter` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only if `A` and `H` are both two-dimensional and `A` is of class `uint8`, `int16`, or `single`.

Examples

Read a color image into the workspace and view it.

```
originalRGB = imread('peppers.png');  
imshow(originalRGB)
```

Create a filter, `h`, that can be used to approximate linear camera motion.

```
h = fspecial('motion', 50, 45);
```

Apply the filter, using `imfilter`, to the image `rgb` to create a new image, `rgb2`.

```
filteredRGB = imfilter(originalRGB, h);  
figure, imshow(filteredRGB)
```

Note that `imfilter` is more memory efficient than some other filtering operations in that it outputs an array of the same data type as the input image array. In this example, the output is an array of `uint8`.

```
whos rgb2  
Name          Size          Bytes  Class  
  
h             37x37         10952  double array  
rgb          384x512x3    589824  uint8 array  
rgb2         384x512x3    589824  uint8 array
```

This example specifies the `replicate` boundary option.

```
boundaryReplicateRGB = imfilter(originalRGB, h, 'replicate');  
figure, imshow(boundaryReplicateRGB)
```

See Also

`conv2`, `convn`, `filter2`, `fspecial`, `ippl`

imfinfo

Purpose

Information about graphics file

`imfinfo` is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference pages.

Purpose	Get handle to most recent current axis containing an image
Syntax	<pre>hax = imgca hax = imgca(hfig)</pre>
Description	<p><code>hax = imgca</code> returns the handle of the of the most recent current axis that contains an image. The current axis may be in a regular figure window or in an Image Tool window.</p> <p>If no figure contains an axis that contains an image, <code>imgca</code> creates a new axis.</p> <p><code>hax = imgca(hfig)</code> returns the handle to the current axis that contains an image in the specified figure (it need not be the current figure).</p>
Note	<code>imgca</code> can be useful in getting the handle to the Image Tool axis. Because the Image Tool turns graphics object handle visibility off, you cannot retrieve a handle to the tool figure using <code>gca</code> .
Example	<p>Label the coins in the image, compute their centroids, and superimpose the centroid locations on the image. View the results using <code>imtool</code> and <code>imgca</code>.</p> <pre>I = imread('coins.png'); figure, imshow(I) bw = im2bw(I, graythresh(getimage)); figure, imshow(bw) bw2 = imfill(bw, 'holes'); L = bwlabel(bw2); s = regionprops(L, 'centroid'); centroids = cat(1, s.Centroid);</pre> <p>Display original image <code>I</code> and superimpose centroids.</p> <pre>imtool(I) hold(imgca, 'on') plot(imgca, centroids(:,1), centroids(:,2), 'r*') hold(imgca, 'off')</pre>
See also	<code>gca</code> , <code>gcf</code> , <code>imgcf</code>

imgcf

Purpose Get handle to most recent current figure containing an image

Syntax `hfig = imgcf`

Description `hfig = imgcf` returns the handle of the most recent current figure that contains an image. The figure may be a regular figure window that contains at least one image or an Image Tool window.

If none of the figures currently open contains an image, `imgcf` creates a new figure.

Note `imgcf` can be useful in getting the handle to the Image Tool figure window. Because the Image Tool turns graphics object handle visibility off, you cannot retrieve a handle to the tool figure using `gcf`.

Example

```
imtool rice.png
cmap = copper(256);
set(imgcf, 'Colormap', cmap)
```

See also `gca`, `gcf`, `imgca`

Purpose Displays Open Image dialog box

Syntax [filename, user_canceled] = imgetfile

Description [filename, user_canceled] = imgetfile displays the Open Image dialog box for the user to fill in, and returns the full path to the file selected in filename. If the user presses the **Cancel** button, user_canceled will be TRUE. Otherwise, user_canceled will be FALSE.

The Open Image dialog box is modal; it blocks the MATLAB command line until the user responds. The file types listed in the dialog are all formats listed in imformats plus DICOM.

See Also imformats, imtool, uigetfile

imhandles

Purpose Get all image handles within a handle

Syntax `himage = imhandles(H)`

Description `himage = imhandles(H)` takes a graphics handle `h` as an input and returns all of the image handles whose ancestor is `h`. `h` can be an array of valid figure, axes, uipanel, or image handles.

`himage` is an array of image handles.

Note `imhandles` errors if the image objects in `himage` do not have the same figure as their parent.

Examples This example returns the handle to the image object in the current figure.

```
figure, imshow('moon.tif');  
himage = imhandles(gcf)
```

This example displays two images in a figure and uses `imhandles` to get handles to both of the image objects in the figure.

```
subplot(1,2,1), imshow('autumn.tif');  
subplot(1,2,2), imshow('glass.png');  
himages = imhandles(gca)
```

See Also `imgca`, `imgcf`

Purpose Display histogram of image data

Syntax

```
imhist(I,n)
imhist(X,map)
[counts,x] = imhist(...)
```

Description `imhist(I)` displays a histogram for the intensity image `I` above a grayscale colorbar. The number of bins in the histogram is specified by the image type. If `I` is a grayscale image, `imhist` uses a default value of 256 bins. If `I` is a binary image, `imhist` uses 2 bins.

`imhist(I,n)` displays a histogram where `n` specifies the number of bins used in the histogram. `n` also specifies the length of the colorbar. If `I` is a binary image, `n` can only have the value 2.

`imhist(X,map)` displays a histogram for the indexed image `X`. This histogram shows the distribution of pixel values above a colorbar of the colormap `map`. The colormap must be at least as long as the largest index in `X`. The histogram has one bin for each entry in the colormap.

`[counts,x] = imhist(...)` returns the histogram counts in `counts` and the bin locations in `x` so that `stem(x,counts)` shows the histogram. For indexed images, it returns the histogram counts for each colormap entry; the length of `counts` is the same as the length of the colormap.

Note For intensity images, the `n` bins of the histogram are each half-open intervals of width $A/(n-1)$. In particular, for intensity images that are not `int16`, the p th bin is the half-open interval $A(p-1.5)/(n-1) \leq x < A(p-0.5)/(n-1)$, where x is the intensity value. For `int16` intensity images, the p th bin is the half-open interval $A(p-1.5)/(n-1) - 32768 \leq x < A(p-0.5)/(n-1) - 32768$, where x is the intensity value. The scale factor A depends on the image class. A is 1 if the intensity image is `double` or `single`, A is 255 if the intensity image is `uint8`, and A is 65535 if the intensity image is `uint16` or `int16`.


imhist

Class Support

An input intensity image can be of class `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. An input indexed image can be of class `uint8`, `uint16`, `single`, `double`, or `logical`.

Example

```
I = imread('pout.tif');  
imhist(I)
```



See Also

`histeq`

`hist` in the MATLAB Function Reference

Purpose H-maxima transform

Syntax
 $I2 = \text{imhmax}(I, h)$
 $I2 = \text{imhmax}(I, h, \text{CONN})$

Description $I2 = \text{imhmax}(I, h)$ suppresses all maxima in the intensity image I whose height is less than h , where h is a scalar.

Regional maxima are connected components of pixels with the same intensity value, t , whose external boundary pixels all have a value less than t .

By default, `imhmax` uses 8-connected neighborhoods for 2-D images, and 26-connected neighborhoods for 3-D images. For higher dimensions, `imhmax` uses `conndef(ndims(I), 'maximal')`.

$I2 = \text{imhmax}(I, h, \text{CONN})$ computes the H-maxima transform, where `CONN` specifies the connectivity. `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support I can be of any nonsparse numeric class and any dimension. $I2$ has the same size and class as I .

imhmax

Example

```
a = zeros(10,10);  
a(2:4,2:4) = 3; % maxima 3 higher than surround  
a(6:8,6:8) = 8; % maxima 8 higher than surround  
b = imhmax(a,4); % only the maxima higher than 4 survive.
```

See Also

conndef, imhmin, imreconstruct

Reference

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

Purpose H-minima transform

Syntax
`I2 = imhmin(I,h)`
`I2 = imhmin(I,h,CONN)`

Description `I2 = imhmin(I,h)` suppresses all minima in the intensity image `I` whose depth is less than `h`, where `h` is a scalar.

Regional minima are connected components of pixels with a constant intensity value, and whose external boundary pixels all have a higher value.

By default, `imhmin` uses 8-connected neighborhoods for 2-D images, and 26-connected neighborhoods for 3-D images. For higher dimensions, `imhmin` uses `conndef(ndims(I), 'maximal')`.

`I2 = imhmin(I,h,CONN)` computes the H-minima transform, where `CONN` specifies the connectivity. `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support `I` can be of any nonsparse numeric class and any dimension. `I2` has the same size and class as `I`.

Example

Create a sample image with two regional minima.

```
a = 10*ones(10,10);  
a(2:4,2:4) = 7;  
a(6:8,6:8) = 2
```

a =

```
10  10  10  10  10  10  10  10  10  10  
10   7   7   7  10  10  10  10  10  10  
10   7   7   7  10  10  10  10  10  10  
10   7   7   7  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10   2   2   2  10  10  
10  10  10  10  10   2   2   2  10  10  
10  10  10  10  10   2   2   2  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10
```

Suppress all minima below a specified value. Note how the region with pixel valued 7 disappears in the transformed image.

```
b = imhmin(a,4)
```

b =

```
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10   6   6   6  10  10  
10  10  10  10  10   6   6   6  10  10  
10  10  10  10  10   6   6   6  10  10  
10  10  10  10  10  10  10  10  10  10  
10  10  10  10  10  10  10  10  10  10
```

See Also

conndef, imhmax, imreconstruct

Reference

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

Purpose Impose minima

Syntax `I2 = imimposemin(I,BW)`
 `I2 = imimposemin(I,BW,CONN)`

Description `I2 = imimposemin(I,BW)` modifies the intensity image `I` using morphological reconstruction so it only has regional minima wherever `BW` is nonzero. `BW` is a binary image the same size as `I`.

By default, `imimposemin` uses 8-connected neighborhoods for 2-D images and 26-connected neighborhoods for 3-D images. For higher dimensions, `imimposemin` uses `conndef(ndims(I), 'minimum')`.

`I2 = imimposemin(I,BW,CONN)` specifies the connectivity, where `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can also be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support `I` can be of any nonsparse numeric class and any dimension. `BW` must be a nonsparse numeric array with the same size as `I`. `I2` has the same size and class as `I`.


imimposemin

Example

Modify an image so that it only has regional minima at one location.

- 1 Read an image and display it. This image is called the *mask* image.

```
mask = imread('glass.png');  
imshow(mask)
```




- 2 Create the marker image that will be used to process the mask image.

The example creates a binary image that is the same size as the mask image and sets a small area of the binary image to 1. These pixels define the location in the mask image where a regional minimum will be imposed.

```
marker = false(size(mask));  
marker(65:70,65:70) = true;
```

To show where these pixels of interest fall on the original image, this code superimposes the marker over the mask. The small white square marks the spot. This code is not essential to the impose minima operation.


```
J = mask;  
J(marker) = 255;  
figure, imshow(J); title('Marker Image Superimposed on Mask');
```



- 3 Impose the regional minimum on the input image using the `imimposemin` function.

The `imimposemin` function uses morphological reconstruction of the mask image with the marker image to impose the minima at the specified location. Note how all the dark areas of the original image, except the marked area, are lighter.


```
K = imimposemin(mask,marker);
figure, imshow(K);
```




- 4 To illustrate how this operation removes all minima in the original image except the imposed minimum, compare the regional minima in the original image with the regional minimum in the processed image. These calls to `imregionalmin` return binary images that specify the locations of all the regional minima in both images.

```
BW = imregionalmin(mask);
figure, imshow(BW); title('Regional Minima in Original Image')
BW2 = imregionalmin(K);
figure, imshow(BW2), title('Regional Minima After Processing');
```

imimposemin



Regional Minima in Original Image



Regional Minima After Processing

Algorithm

`imimposemin` uses a technique based on morphological reconstruction.

See Also

`conndef`, `imreconstruct`, `imregionalmin`

Purpose

Compute linear combination of images

Syntax

```
Z = imlincomb(K1,A1,K2,A2,...,Kn,An)
Z = imlincomb(K1,A1,K2,A2,...,Kn,An,K)
Z = imlincomb(..., output_class)
```

Description

`Z = imlincomb(K1,A1,K2,A2,...,Kn,An)` computes

$$K1*A1 + K2*A2 + \dots + Kn*An$$

where `K1`, `K2`, through `Kn` are real, double scalars and `A1`, `A2`, through `An` are real, nonsparse, numeric arrays with the same class and size. `Z` has the same class and size as `A1`.

`Z = imlincomb(K1,A1,K2,A2,...,Kn,An,K)` computes

$$K1*A1 + K2*A2 + \dots + Kn*An + K$$

where `imlincomb` adds `K`, a real, double scalar, to the sum of the products of `K1` through `Kn` and `A1` through `An`.

`Z = imlincomb(...,output_class)` lets you specify the class of `Z`. `output_class` is a string containing the name of a numeric class.

When performing a series of arithmetic operations on a pair of images, you can achieve more accurate results if you use `imlincomb` to combine the operations, rather than nesting calls to the individual arithmetic functions, such as `imadd`. When you nest calls to the arithmetic functions, and the input arrays are of an integer class, each function truncates and rounds the result before passing it to the next function, thus losing accuracy in the final result. `imlincomb` computes each element of the output `Z` individually, in double-precision floating point. If `Z` is an integer array, `imlincomb` truncates elements of `Z` that exceed the range of the integer type and rounds off fractional values.

On Intel architecture processors, `imlincomb` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only in the following cases:

$$Z = \text{imlincomb}(1.0, A1, 1.0, A2)$$

$$Z = \text{imlincomb}(1.0, A1, -1.0, A2)$$

$$Z = \text{imlincomb}(-1.0, A1, 1.0, A2)$$

imlincomb

```
Z = imlincomb( 1.0 , A1, K)
```

where A1, A2, and Z are of class uint8, int16, or single and are of the same class.

Examples

Example 1

Scale an image by a factor of 2.

```
I = imread('cameraman.tif');
J = imlincomb(2,I);
imshow(J)
```

Example 2

Form a difference image with the zero value shifted to 128.

```
I = imread('cameraman.tif');
J = uint8(filter2(fspecial('gaussian'), I));
K = imlincomb(1,I,-1,J,128); % K(r,c) = I(r,c) - J(r,c) + 128
figure, imshow(K)
```

Example 3

Add two images with a specified output class.

```
I = imread('rice.png');
J = imread('cameraman.tif');
K = imlincomb(1,I,1,J,'uint16');
figure, imshow(K,[])
```

Example 4

To illustrate how `imlincomb` performs all the arithmetic operations before truncating the result, compare the results of calculating the average of two arrays, X and Y, using nested arithmetic functions and then using `imlincomb`.

In the version that uses nested arithmetic functions, `imadd` adds 255 and 50 and truncates the result to 255 before passing it to `imdivide`. The average returned in `Z(1,1)` is 128.

```
X = uint8([ 255 10 75; 44 225 100]);
Y = uint8([ 50 20 50; 50 50 50 ]);
Z = imdivide(imadd(X,Y),2)
Z =
```

```
128    15    63
 47   128    75
```

`imlincomb` performs the addition and division in double precision and only truncates the final result. The average returned in `Z2(1,1)` is 153.

```
Z2 = imlincomb(.5,X,.5,Y)
Z2 =
 153    15    63
  47   138    75
```

See Also

`imadd`, `imcomplement`, `imdivide`, `immultiply`, `imsubtract`

immagbox

Purpose Magnification box for scroll panel

Syntax `hbox = immagbox(hparent,himage)`

Description `hbox = immagbox(hparent,himage)` creates a Magnification box for the image displayed in a scroll panel created by `imscrollpanel`. `hparent` is a handle to the figure or `uipanel` object that will contain the Magnification box. `himage` is a handle to the target image (the image in the scroll panel). `immagbox` returns `hbox`, which is a handle to the Magnification box `uicontrol` object

A Magnification box is an editable text box `uicontrol` that contains the current magnification of the target image. When you enter a new value in the magnification box, the magnification of the target image changes. When the magnification of the target image changes for any reason, the magnification box updates the magnification value.

API Functions A magnification box contains a structure of function handles, called an API. You can use the functions in this API to manipulate magnification box. To retrieve this structure, use the `iptgetapi` function.

```
api = iptgetapi(hbox)
```

The API for the Magnification box includes the following function.

Field	Description
<code>setMagnification</code>	Sets the magnification in units of screen pixels per image pixel. <code>api.setMagnification(new_mag)</code> where <code>new_mag</code> is a scalar magnification factor. Multiply <code>new_mag</code> by 100 to get percent magnification. For example if you call <code>api.setMagnification(2)</code> , the magnification box will show the string '200%'.

Example This example adds a magnification box to a scrollable image. Because the toolbox scrollable navigation is incompatible with standard MATLAB figure window navigation tools, the example suppresses the toolbar and menu bar in

the figure window. The example positions the scroll panel in the figure window to allow room for the magnification box.

```

hFig = figure('Toolbar','none',...
             'Menubar','none');
hIm = imshow('pears.png');
hSP = imscrollpanel(hFig,hIm);
set(hSP,'Units','normalized',...
     'Position',[0 .1 1 .9])

hMagBox = immagbox(hFig,hIm);
pos = get(hMagBox,'Position');
set(hMagBox,'Position',[0 0 pos(3) pos(4)])

```

Change the magnification of the image in the scroll panel, using the scroll panel API function `setMagnification`. Notice how the magnification box updates.

```

apiSP = iptgetapi(hSP);
apiSP.setMagnification(2)

```

See also

`imscrollpanel`, `iptgetapi`

immovie

Purpose Make movie from multiframe indexed image

Syntax

```
mov = immovie(X,map)
mov = immovie(RGB)
```

Description `mov = immovie(X,map)` returns the movie structure array `mov` from the images in the multiframe indexed image `X` with the colormap `map`. As it creates the movie array, `immovie` displays the movie frames on the screen. You can play the movie using the MATLAB `movie` function. For details about the movie structure array, see the reference page for `getframe`.

`X` comprises multiple indexed images, all having the same size and all using the colormap `map`. `X` is an `m-by-n-by-1-by-k` array, where `k` is the number of images.

`mov = immovie(RGB)` returns the movie structure array `mov` from the images in the multiframe, truecolor image `RGB`.

`RGB` comprises multiple truecolor images, all having the same size. `RGB` is an `m-by-n-by-3-by-k` array, where `k` is the number of images.

Remarks You can also use the MATLAB function `avifile` to make movies from images. The `avifile` function creates AVI files. To convert an existing MATLAB movie into an AVI file, use the `movie2avi` function.

Class Support An indexed image can be `uint8`, `uint16`, `single`, `double`, or `logical`. A truecolor image can be `uint8`, `uint16`, `single`, or `double`. `mov` is a MATLAB movie structure.

Example

```
load mri
mov = immovie(D,map);
movie(mov,3)
```

See Also `avifile`, `getframe`, `montage`, `movie`, `movie2avi`

Purpose Multiply two images, or multiply an image by a constant

Syntax `Z = immultiply(X,Y)`

Description `Z = immultiply(X,Y)` multiplies each element in array `X` by the corresponding element in array `Y` and returns the product in the corresponding element of the output array `Z`.

If `X` and `Y` are real numeric arrays with the same size and class, then `Z` has the same size and class as `X`. If `X` is a numeric array and `Y` is a scalar double, then `Z` has the same size and class as `X`.

If `X` is logical and `Y` is numeric, then `Z` has the same size and class as `Y`. If `X` is numeric and `Y` is logical, then `Z` has the same size and class as `X`.

`immultiply` computes each element of `Z` individually in double-precision floating point. If `X` is an integer array, then elements of `Z` exceeding the range of the integer type are truncated, and fractional values are rounded.

If `X` and `Y` are double arrays, you can use the expression `X.*Y` instead of this function.

Note On Intel architecture processors, `immultiply` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only if arrays `X`, `Y`, and `Z` are of class `logical`, `uint8`, or `single`, and are of the same class.

Example Multiply an image by itself. Note how the example converts the class of the image from `uint8` to `uint16` before performing the multiplication to avoid truncating the results.

```
I = imread('moon.tif');
I16 = uint16(I);
J = immultiply(I16,I16);
imshow(I), figure, imshow(J)
```

Scale an image by a constant factor:

```
I = imread('moon.tif');
```

immultiply

```
J = immultiply(I,0.5);  
subplot(1,2,1), imshow(I)  
subplot(1,2,2), imshow(J)
```

See also

[imabsdiff](#), [imadd](#), [imcomplement](#), [imdivide](#), [imlincomb](#), [imsubtract](#), [ippl](#)

Purpose Add noise to an image

Syntax `J = imnoise(I,type)`
`J = imnoise(I,type,parameters)`

Description `J = imnoise(I,type)` adds noise of a given type to the intensity image `I`. `type` is a string that can have one of these values.

Value	Description
'gaussian'	Gaussian white noise with constant mean and variance
'localvar'	Zero-mean Gaussian white noise with an intensity-dependent variance
'poisson'	Poisson noise
'salt & pepper'	On and off pixels
'speckle'	Multiplicative noise

Depending on which type of noise you specify, `imnoise` accepts additional parameters. If you omit these arguments, `imnoise` uses default values for the parameters. All numerical parameters are normalized; they correspond to operations with images with intensities ranging from 0 to 1.

`J = imnoise(I,'gaussian',m,v)` adds Gaussian white noise of mean `m` and variance `v` to the image `I`. The default is zero mean noise with 0.01 variance.

`J = imnoise(I,'localvar',V)` adds zero-mean, Gaussian white noise of local variance `V` to the image `I`. `V` is an array of the same size as `I`.

`J = imnoise(I,'localvar',image_intensity,var)` adds zero-mean, Gaussian noise to an image `I`, where the local variance of the noise, `var`, is a function of the image intensity values in `I`. The `image_intensity` and `var` arguments are vectors of the same size, and `plot(image_intensity,var)` plots the functional relationship between noise variance and image intensity. The `image_intensity` vector must contain normalized intensity values ranging from 0 to 1.

`J = imnoise(I, 'poisson')` generates Poisson noise from the data instead of adding artificial noise to the data. In order to respect Poisson statistics, the intensities of `uint8` and `uint16` images must correspond to the number of photons (or any other quanta of information). Double-precision images are used when the number of photons per pixel can be much larger than 65535 (but less than 10^{12}); the intensity values vary between 0 and 1 and correspond to the number of photons divided by 10^{12} .

`J = imnoise(I, 'salt & pepper', d)` adds salt and pepper noise to the image `I`, where `d` is the noise density. This affects approximately `d*prod(size(I))` pixels. The default is 0.05 noise density.

`J = imnoise(I, 'speckle', v)` adds multiplicative noise to the image `I`, using the equation $J = I + n * I$, where `n` is uniformly distributed random noise with mean 0 and variance `v`. The default for `v` is 0.04.

Note The mean and variance parameters for 'gaussian', 'localvar', and 'speckle' noise types are always specified as if the image were of class `double` in the range [0, 1]. If the input image is of class `uint8` or `uint16`, the `imnoise` function converts the image to `double`, adds noise according to the specified type and parameters, and then converts the noisy image back to the same class as the input.

Class Support

`I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image `J` is of the same class as `I`. If `I` has more than two dimensions it is treated as a multidimensional intensity image and not as an RGB image.

Example

```
I = imread('eight.tif');
J = imnoise(I, 'salt & pepper', 0.02);
imshow(I)
figure, imshow(J)
```



See Also

rand, randn in the MATLAB Function Reference

imopen

Purpose Open an image

Syntax
IM2 = imopen(IM,SE)
IM2 = imopen(IM,NHOOD)

Description IM2 = imopen(IM,SE) performs morphological opening on the grayscale or binary image IM with the structuring element SE. The argument SE must be a single structuring element object, as opposed to an array of objects.


IM2 = imopen(IM,NHOOD) performs opening with the structuring element strel(NHOOD), where NHOOD is an array of 0's and 1's that specifies the structuring element neighborhood.

Class Support IM can be any numeric or logical class and any dimension, and must be nonsparse. If IM is logical, then SE must be flat. IM2 has the same class as IM.

Example This example uses imopen to filter out the smaller objects in an image.

1 Read the image into the MATLAB workspace and display it.

```
I = imread('snowflakes.png');  
imshow(I)
```



2 Create a disk-shaped structuring element with a radius of 5 pixels.

```
se = strel('disk',5);
```

3 Remove snowflakes having a radius less than 5 pixels by opening it with the disk-shaped structuring element created in step 2.

```
I_opened = imopen(I,se);  
figure, imshow(I_opened,[])
```



See Also

`imclose`, `imdilate`, `imerode`, `strel`

imoverview

Purpose Overview navigation tool for image displayed in scroll panel

Syntax
`imoverview(himage)`
`hfig = imoverview(himage)`

Description `imoverview(himage)` creates an Overview tool associated with the image specified by the handle `himage`, called the target image. The target image must be contained in a scroll panel created by `imscrollpanel`.

The Overview tool is a navigation aid for images displayed in a scroll panel. `imoverview` creates the tool in a separate figure window that displays the target image in its entirety, scaled to fit. Over this scaled version of the image, the tool draws a rectangle, called the detail rectangle, that shows the portion of the target image that is currently visible in the scroll panel. To view portions of the image that are not currently visible in the scroll panel, move the detail rectangle in the Overview tool.

`hfig = imoverview(...)` returns a handle to the Overview tool figure.

Note To create an Overview tool that can be embedded in an existing figure or `uipanel` object, use `imoverviewpanel`.

Example This example creates a figure, disabling the toolbar and menubar, because the toolbox navigation tools are not compatible with the standard MATLAB zoom and pan tools. The example then creates a scroll panel in the figure and uses scroll panel API functions to set the magnification.

```
hFig = figure('Toolbar','none',...  
             'Menubar','none');  
hIm = imshow('tape.png');  
hSP = imscrollpanel(hFig,hIm);  
api = iptgetapi(hSP);  
api.setMagnification(2) % 2X = 200%  
imoverview(hIm)
```

See Also `imoverviewpanel`, `imscrollpanel`

Purpose	Overview navigation tool uipanel for image displayed in scroll panel
Syntax	<code>hpanel = imoverviewpanel(hparent,himage)</code>
Description	<p><code>hpanel=imoverviewpanel(hparent,himage)</code> creates creates an Overview tool panel associated with the image specified by the handle <code>himage</code>, called the target image. <code>himage</code> must be contained in a scroll panel created by <code>imscrollpanel</code>. <code>hparent</code> is a handle to the figure or uipanel object that will contain the Overview tool panel. <code>imoverviewpanel</code> returns <code>hpanel</code>, a handle to the Overview tool uipanel object</p> <p>The Overview tool is a navigation aid for images displayed in a scroll panel. <code>imoverviewpanel</code> creates the tool in a uipanel object that can be embedded in a figure or uipanel object. The tool displays the target image in its entirety, scaled to fit. Over this scaled version of image, the tool draws a rectangle, called the detail rectangle, that shows the portion of the target image that is currently visible in the scroll panel. To view portions of the image that are not currently visible in the scroll panel, move the detail rectangle in the Overview tool.</p>
Note	To create an Overview tool in a separate figure, use <code>imoverview</code> . When created using <code>imoverview</code> , the Overview tool includes zoom-in and zoom-out buttons.
Example	<p>This example creates an Overview navigation tool that is embedded in the same figure that contains the target image.</p> <pre>hFig = figure('Toolbar','none',... 'Menubar','none'); hIm = imshow('tissue.png'); hSP = imscrollpanel(hFig,hIm); set(hSP,'Units','normalized',... 'Position',[0 .5 1 .5]) hOvPanel = imoverviewpanel(hFig,hIm); set(hOvPanel,'Units','Normalized',... 'Position',[0 0 1 .5])</pre>
See also	<code>imoverview</code> , <code>imscrollpanel</code>

impixel

Purpose Determine pixel color values

Syntax

```
P = impixel(I)
P = impixel(X,map)
P = impixel(RGB)
```

```
P = impixel(I,c,r)
P = impixel(X,map,c,r)
P = impixel(RGB,c,r)
[c,r,P] = impixel(...)
```

```
P = impixel(x,y,I,xi,yi)
P = impixel(x,y,X,map,xi,yi)
P = impixel(x,y,RGB,xi,yi)
[xi,yi,P] = impixel(x,y,...)
```

Description

`impixel` returns the red, green, and blue color values of specified image pixels. In the syntaxes below, `impixel` displays the input image and waits for you to specify the pixels with the mouse.

```
P = impixel(I)
P = impixel(X,map)
P = impixel(RGB)
```

If you omit the input arguments, `impixel` operates on the image in the current axes.

Use normal button clicks to select pixels. Press **Backspace** or **Delete** to remove the previously selected pixel. A shift-click, right-click, or double-click adds a final pixel and ends the selection; pressing **Return** finishes the selection without adding a pixel.

When you finish selecting pixels, `impixel` returns an m -by-3 matrix of RGB values in the supplied output argument. If you do not supply an output argument, `impixel` returns the matrix in `ans`.

You can also specify the pixels noninteractively, using these syntaxes.

```
P = impixel(I,c,r)
P = impixel(X,map,c,r)
P = impixel(RGB,c,r)
```

`r` and `c` are equal-length vectors specifying the coordinates of the pixels whose RGB values are returned in `P`. The k th row of `P` contains the RGB values for the pixel $(r(k), c(k))$.

If you supply three output arguments, `impixel` returns the coordinates of the selected pixels. For example,

```
[c,r,P] = impixel(...)
```

To specify a nondefault spatial coordinate system for the input image, use these syntaxes.

```
P = impixel(x,y,I,xi,yi)
P = impixel(x,y,X,map,xi,yi)
P = impixel(x,y,RGB,xi,yi)
```

`x` and `y` are two-element vectors specifying the image `XData` and `YData`. `xi` and `yi` are equal-length vectors specifying the spatial coordinates of the pixels whose RGB values are returned in `P`. If you supply three output arguments, `impixel` returns the coordinates of the selected pixels.

```
[xi,yi,P] = impixel(x,y,...)
```

Class Support

The input image can be of class `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. All other inputs are of class `double`.

If the input is `double`, the output `P` is `double`. For all other input classes the output is `single`. The rest of the outputs are `double`.

Remarks

`impixel` works with indexed, intensity, and RGB images. `impixel` always returns pixel values as RGB triplets, regardless of the image type:

- For an RGB image, `impixel` returns the actual data for the pixel. The values are either `uint8` integers or `double` floating-point numbers, depending on the class of the image array.
- For an indexed image, `impixel` returns the RGB triplet stored in the row of the colormap that the pixel value points to. The values are `double` floating-point numbers.
- For an intensity image, `impixel` returns the intensity value as an RGB triplet, where `R=G=B`. The values are either `uint8` integers or `double` floating-point numbers, depending on the class of the image array.

impixel

Example

```
RGB = imread('peppers.png');  
c = [12 146 410];  
r = [104 156 129];  
pixels = impixel(RGB,c,r)
```

```
pixels =
```

```
    62    34    63  
   166    54    60  
    59    28    47
```

See Also

`improfile`, `pixval`


Purpose Pixel information tool

Syntax

```
impixelinfo
impixelinfo(h)
impixelinfo(hparent,himage)
hpanel = impixelinfo(...)
```

Description `impixelinfo` creates a pixel information tool in the current figure. The pixel information tool displays information about the pixel in an image that the cursor is positioned over. The tool displays pixel information for all the images in a figure..

The tool is a uipanel object, positioned in the lower-left corner of the figure. It contains the text string “Pixel info:” followed by the x and y coordinates of the pixel and its value, as shown in the following figure.



The information displayed depends on the image type. The following table shows the information displayed for each image type. If the cursor is outside of the image area in the figure, the pixel information tool displays the default pixel information string.

Image Type	Default Pixel Information String	Example
Intensity	(X,Y) Intensity	(13,30) 82
Indexed	(X,Y) <index> [R G B]	(2,6) <4> [0.29 0.05 0.32]
Binary	(X,Y) BW	(12,1) 0
Truecolor	(X,Y) [R G B]	(19,10) [15 255 10]

If you want to display the pixel information without the “Pixel Info” label, use the `impixelinfoval` function.

impixelinfo

`impixelinfo(h)` creates a pixel information tool in the figure specified by `h`, where `h` is a handle to an image, axes, uipanel, or figure object. Axes, uipanel, or figure objects must contain at least one image object.

`impixelinfo(hparent,himage)` creates a pixel information tool in `hparent` that provides information about the pixels in `himage`. `himage` is a handle to an image or an array of image handles. `hparent` is a handle to the figure or uipanel object that contains the pixel information tool.

`hpanel=impixelinfo(...)` returns a handle to the pixel information tool uipanel.

Note

To copy the pixel information string to the clipboard, right-click while the cursor is positioned over a pixel. In the context menu displayed, choose **Copy pixel info**.

Examples

This example displays an image and adds a pixel information tool to the figure. The example shows how you can change the position of the tool in the figure using properties of the tool uipanel object.

```
h = imshow('hestain.png');
hp = impixelinfo;
set(hp, 'Position', [150 290 300 20]);
```

This example shows how the pixel information tool works in a figure containing multiple images of different types.

```
figure
subplot(1,2,1), imshow('liftingbody.png');
subplot(1,2,2), imshow('autumn.tif');
impixelinfo;
```

See Also

`impixelinfoval`, `impixelregion`, `imshow`, `imtool`

Purpose Pixel information tool without the text label

Syntax `hcontrol = impixelinfoval(hparent,himage)`

Description `hcontrol = impixelinfoval(hparent,himage)` creates a pixel information tool in `hparent` that provides information about the pixels in the image specified by `himage`. `hparent` is a handle to a figure or uipanel object. `himage` can be a handle to an image or an array of image handles.

The pixel information tool displays information about the pixel in an image that the cursor is positioned over. The tool displays pixel information for all the images in a figure.

When created with `impixelinfo`, the tool is a uipanel object, positioned in the lower-left corner of the figure, that contains the text label "Pixel Info:" followed by the the x - and y -coordinates of the pixel and its value. When created with `impixelinfoval`, the tool is a uicontrol object positioned in the lower-left corner of the figure, that displays the pixel information without the text label, as shown in the following figure.

X and Y coordinates	Pixel Value
(167, 251)	114

The information displayed depends on the image type. See `impixelinfo` for details.

To copy the pixel value string to the clipboard, right-click while the cursor is positioned over a pixel. In the context menu displayed, choose **Copy pixel info**.

Example This example shows how to add a pixel value tool to a figure. Note how you can change the style and size of the font used to display the value in the tool using standard Handle Graphics commands.

```
ankle = dicomread('CT-MON02-16-ankle.dcm');
h = imshow(ankle,[]);
hText = impixelinfoval(gcf,h);
set(hText,'FontWeight','bold')
set(hText,'FontSize',10)
```

impixelinfoval

See also


`impixelinfo`

Purpose Pixel Region tool

Syntax
`impixelregion`
`impixelregion(h)`
`hfig = impixelregion(...)`

Description `impixelregion` creates a Pixel Region display tool associated with the image displayed in the current figure, called the target image.

The Pixel Region tool opens a separate figure window containing an extreme close-up view of a small region of pixels in the target image, shown in the following figure. The tool superimposes the numeric value of the pixel over each pixel. To define the region being examined, the tool overlays a rectangle on the target image, called the pixel region rectangle. To view pixels in a different region, click and drag the rectangle over the target image.



impixelregion

`impixelregion(h)` creates a Pixel Region tool associated with the object specified by the handle `h`. `h` can be a handle to a figure, axes, uipanel, or image object. If `h` is a handle to an axes or figure, `impixelregion` associates the tool with the first image found in the axes or figure.

`hfig=impixelregion(...)` returns `hfig`, a handle of the Pixel Region tool figure.

Note

To create a Pixel Region tool that can be embedded in an existing figure window or uipanel, use `impixelregionpanel`.

Example

This example displays an image and then creates a Pixel Region tool associated with the image.

```
imshow peppers.png
impixelregion
```

See Also

`impixelinfo`, `impixelregionpanel`, `imtool`

Purpose Pixel Region tool panel

Syntax `hpanel = impixelregionpanel(hparent,himage)`

Description `hpanel = impixelregionpanel(hparent,himage)` creates a Pixel Region tool panel associated with the image specified by the handle `himage`, called the target image. This is the image whose pixels are to be displayed. `hparent` is the handle to the figure or `uipanel` object that will contain the Pixel Region tool panel. `hpanel` is the handle to the Pixel Region tool scroll panel.

The Pixel Region tool is a `uipanel` object that contains an extreme close-up view of a small region of pixels in the target image. The tool superimposes the numeric value of the pixel over each pixel. To define the region being examined, the tool overlays a rectangle on the target image, called the pixel region rectangle. To view pixels in a different region, click and drag the rectangle over the target image.

Note To create a Pixel Region tool in a separate figure window, use `impixelregion`.

Example

```
himage = imshow('peppers.png');  
hfigure = figure;  
hpanel = impixelregionpanel(hfigure, himage);
```

Set the panel's position to the lower-left quadrant of the figure.

```
set(hpanel, 'Position', [0 0 .5 .5])
```

See Also `impixelregion`, `impositionrect`, `imtool`

impositionrect

Purpose Create draggable position rectangle

Syntax `H = impositionrect(hparent,position)`

Description `H = impositionrect(hparent,position)` creates a position rectangle on the object specified by `hparent`. The function returns `H`, a handle to the position rectangle, which is an `hggroup` object. `hparent` specifies the `hggroup`'s parent, which is typically an axes object, but can also be any other object that can be the parent of an `hggroup`. `position` is a four-element position vector that specifies the initial location of the rectangle. `position` has the form `[XMIN YMIN WIDTH HEIGHT]`

All measurements are in units specified by the `Units` property axes object. When you do not specify the `position` argument, `impositionrect` uses `[0 0 1 1]` as the default value.

Remarks A position rectangle can be dragged interactively using the mouse. When the position rectangle occupies a small number of screen pixels, its appearance changes to aid visibility.

The position rectangle has a context menu associated with it that you can use to copy the current position to the clipboard and change the color used to display the rectangle.

API Function Syntaxes A position rectangle contains a structure of function handles, called an API, that can be used to manipulate it. To retrieve this structure from the position rectangle, use the `iptgetapi` function.

`API = iptgetapi(H)`

The following lists the functions in the position rectangle API in the order

they appear in the API structure.

Function	Description
setPosition	<p>Sets the position rectangle to a new position.</p> <pre>api.setPosition(new_position)</pre> <p>where <code>new_position</code> is a four-element position vector.</p>
getPosition	<p>Returns the current position of the position rectangle.</p> <pre>position = api.getPosition()</pre> <p><code>position</code> is a four-element position vector.</p>
delete	<p>Deletes the position rectangle associated with the API.</p> <pre>api.delete()</pre>
setColor	<p>Sets the color used to draw the position rectangle.</p> <pre>api.setColor(new_color)</pre> <p>where <code>new_color</code> is a three-element vector specifying an RGB value.</p>
addNewPositionCallback	<p>Adds the function handle <code>fun</code> to the list of new-position callback functions.</p> <pre>id = api.addNewPositionCallback(fun)</pre> <p>Whenever the position rectangle changes its position, each function in the list is called with the syntax:</p> <pre>fun(position)</pre> <p>The return value, <code>id</code>, is used only with <code>removeNewPositionCallback</code>.</p>

impositionrect

Function	Description
<code>removeNewPositionCallback</code>	<p>Removes the corresponding function from the new-position callback list.</p> <pre>api.removeNewPositionCallback(id)</pre> <p>where <code>id</code> is the identifier returned by <code>api.addNewPositionCallback</code></p>
<code>setDragConstraintCallback</code>	<p>Sets the drag constraint function to be the specified function handle, <code>fcn</code>.</p> <pre>api.setDragConstraintCallback(fcn)</pre> <p>Whenever the position rectangle is moved because of a mouse drag, the constraint function is called using the syntax:</p> <pre>constrained_position = fcn(new_position)</pre> <p>where <code>new_position</code> is a four-element position vector. This allows a client, for example, to control where the position rectangle may be dragged.</p>

Examples

Display in the command window the updated position of the position rectangle as it moves in the axes.

```
close all, plot(1:10)
h = impositionrect(gca, [4 4 2 2]);
api = iptgetapi(h);
api.addNewPositionCallback(@(p) disp(p));
```

Constrain the position rectangle to move only up and down.

```
close all, plot(1:10)
h = impositionrect(gca, [4 4 2 2]);
api = getappdata(h, 'API');
api.setDragConstraintCallback(@(p) [4 p(2:4)]);
```

Specify the color of the position rectangle.

```
close all, plot(1:10)
h = impositionrect(gca, [4 4 2 2]);
api = iptgetapi(h, 'API');
api.setColor([1 0 0]);
```

When the position rectangle occupies only a few pixels on the screen, the rectangle is drawn in a different style to increase its visibility.

```
close all, imshow cameraman.tif
h = impositionrect(gca, [100 100 10 10]);
```

See Also

[iptgetapi](#)

improfile

Purpose Compute pixel-value cross-sections along line segments

Syntax

```
c = improfile
c = improfile(n)

c = improfile(I,xi,yi)
c = improfile(I,xi,yi,n)

[cx,cy,c] = improfile(...)
[cx,cy,c,xi,yi] = improfile(...)

[...] = improfile(x,y,I,xi,yi)
[...] = improfile(x,y,I,xi,yi,n)

[...] = improfile(...,method)
```

Description

`improfile` computes the intensity values along a line or a multiline path in an image. `improfile` selects equally spaced points along the path you specify, and then uses interpolation to find the intensity value for each point. `improfile` works with grayscale intensity images and RGB images.

If you call `improfile` with one of these syntaxes, it operates interactively on the image in the current axes.

```
c = improfile
c = improfile(n)
```

`n` specifies the number of points to compute the intensity value for. If you do not provide this argument, `improfile` chooses a value for `n`, roughly equal to the number of pixels the path traverses.

You specify the line or path using the mouse, by clicking points in the image. Press **Backspace** or **Delete** to remove the previously selected point. A shift-click, right-click, or double-click adds a final point and ends the selection; pressing **Return** finishes the selection without adding a point. When you finish selecting points, `improfile` returns the interpolated data values in `c`. `c` is an `n`-by-1 vector if the input is a grayscale intensity image, or an `n`-by-1-by-3 array if the input is an RGB image.

If you omit the output argument, `improfile` displays a plot of the computed intensity values. If the specified path consists of a single line segment, `improfile` creates a two-dimensional plot of intensity values versus the distance along the line segment; if the path consists of two or more line segments, `improfile` creates a three-dimensional plot of the intensity values versus their x - and y -coordinates.

You can also specify the path noninteractively, using these syntaxes.

```
c = improfile(I,xi,yi)
c = improfile(I,xi,yi,n)
```

`xi` and `yi` are equal-length vectors specifying the spatial coordinates of the endpoints of the line segments.

You can use these syntaxes to return additional information.

```
[cx,cy,c] = improfile(...)
[cx,cy,c,xi,yi] = improfile(...)
```

`cx` and `cy` are vectors of length `n`, containing the spatial coordinates of the points at which the intensity values are computed.

To specify a nondefault spatial coordinate system for the input image, use these syntaxes.

```
[...] = improfile(x,y,I,xi,yi)
[...] = improfile(x,y,I,xi,yi,n)
```

`x` and `y` are two-element vectors specifying the image `XData` and `YData`.

`[...] = improfile(...,method)` uses the specified interpolation method. `method` is a string that can have one of these values. The default value is enclosed in braces (`{}`).

Value	Description
{'nearest'}	Nearest-neighbor interpolation
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation


improfile

Class Support

The input image can be `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. All other inputs and outputs must be `double`.

Example

```
I = imread('liftingbody.png');  
x = [19 427 416 77];  
y = [96 462 37 33];  
improfile(I,x,y),grid on;
```



See Also

`impixel`

`interp2` in the MATLAB Function Reference

Purpose

Read image from graphics file

`imread` is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

imreconstruct

Purpose Morphological reconstruction

Syntax
IM = imreconstruct(MARKER, MASK)
IM = imreconstruct(MARKER, MASK, CONN)

Description IM = imreconstruct(MARKER, MASK) performs morphological reconstruction of the image MARKER under the image MASK. MARKER and MASK can be two intensity images or two binary images with the same size. The returned image IM is an intensity or binary image, respectively. MARKER must be the same size as MASK, and its elements must be less than or equal to the corresponding elements of MASK.

By default, imreconstruct uses 8-connected neighborhoods for 2-D images and 26-connected neighborhoods for 3-D images. For higher dimensions, imreconstruct uses conndef(ndims(I), 'maximal').

IM = imreconstruct(MARKER, MASK, CONN) performs morphological reconstruction with the specified connectivity. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Morphological reconstruction is the algorithmic basis for several other Image Processing Toolbox functions, including `imclearborder`, `imextendedmax`, `imextendedmin`, `imfill`, `imhmax`, `imhmin`, and `imimposemin`.

Class Support `MARKER` and `MASK` must be nonsparse numeric or logical arrays with the same class and any dimension. `IM` is of the same class as `MARKER` and `MASK`.

Algorithm `imreconstruct` uses the fast hybrid grayscale reconstruction algorithm described in [1].

See Also `imclearborder`, `imextendedmax`, `imextendedmin`, `imfill`, `imhmax`, `imhmin`, `imimposemin`

Reference [1] Vincent, L., "Morphological Grayscale Reconstruction in Image Analysis: Applications and Efficient Algorithms," *IEEE Transactions on Image Processing*, Vol. 2, No. 2, April, 1993, pp. 176-201.

imregionalmax

Purpose Find regional maxima

Syntax `BW = imregionalmax(I)`
`BW = imregionalmax(I,CONN)`

Description `BW = imregionalmax(I)` finds the regional maxima of `I`. `imregionalmax` returns the binary image `BW` that identifies the locations of the regional maxima in `I`. `BW` is the same size as `I`. In `BW`, pixels that are set to 1 identify regional maxima; all other pixels are set to 0.

Regional maxima are connected components of pixels with the same intensity value, t , whose external boundary pixels all have a value less than t .

By default, `imregionalmax` uses 8-connected neighborhoods for 2-D images and 26-connected neighborhoods for 3-D images. For higher dimensions, `imregionalmax` uses `conndef(ndims(I), 'maximal')`.

`BW = imregionalmax(I,CONN)` computes the regional maxima of `I` using the specified connectivity. `CONN` can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for `CONN` a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of `CONN`. Note that `CONN` must be symmetric about its center element.

Class Support I can be any nonsparse, numeric class and any dimension. BW is logical.

Example Create a sample image with several regional maxima.

```
A = 10*ones(10,10);
A(2:4,2:4) = 22;
A(6:8,6:8) = 33;
A(2,7) = 44;
A(3,8) = 45;
A(4,9) = 44;
A =
    10    10    10    10    10    10    10    10    10    10
    10    22    22    22    10    10    44    10    10    10
    10    22    22    22    10    10    10    45    10    10
    10    22    22    22    10    10    10    10    44    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    33    33    33    10    10
    10    10    10    10    10    33    33    33    10    10
    10    10    10    10    10    33    33    33    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10
```

Find the regional maxima.

```
regmax = imregionalmax(A)
regmax =
     0     0     0     0     0     0     0     0     0     0
     0     1     1     1     0     0     0     0     0     0
     0     1     1     1     0     0     0     1     0     0
     0     1     1     1     0     0     0     0     0     0
     0     0     0     0     0     0     0     0     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     0     0     0     0     0
     0     0     0     0     0     0     0     0     0     0
```

See Also `conndef`, `imreconstruct`, `imregionalmin`

imregionalmin

Purpose Find regional minima

Syntax
BW = imregionalmin(I)
BW = imregionalmin(I,CONN)

Description BW = imregionalmin(I) computes the regional minima of I. The output binary image BW has value 1 corresponding to the pixels of I that belong to regional minima and 0 otherwise. BW is the same size as I.

Regional minima are connected components of pixels with the same intensity value, t , whose external boundary pixels all have a value greater than t .

By default, imregionalmin uses 8-connected neighborhoods for 2-D images and 26-connected neighborhoods for 3-D images. For higher dimensions, imregionalmin uses conndef(ndims(I), 'maximal').

BW = imregionalmin(I,CONN) specifies the desired connectivity. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Class Support I can be any nonsparse, numeric class and any dimension. BW is logical.

Example

Create a 10-by-10 pixel sample image that contains two regional minima.

```
A = 10*ones(10,10);
A(2:4,2:4) = 2;
A(6:8,6:8) = 7;
A =
    10    10    10    10    10    10    10    10    10    10
    10     2     2     2    10    10    10    10    10    10
    10     2     2     2    10    10    10    10    10    10
    10     2     2     2    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10     7     7     7    10    10
    10    10    10    10    10     7     7     7    10    10
    10    10    10    10    10     7     7     7    10    10
    10    10    10    10    10    10    10    10    10    10
    10    10    10    10    10    10    10    10    10    10
```

Pass the sample image A to `imregionalmin`. The function returns a binary image, the same size as A, in which pixels with the value 1 represent the regional minima in A. `imregionalmin` sets all other pixels in to zero (0).

```
B = imregionalmin(A)
B =
     0     0     0     0     0     0     0     0     0     0
     0     1     1     1     0     0     0     0     0     0
     0     1     1     1     0     0     0     0     0     0
     0     1     1     1     0     0     0     0     0     0
     0     0     0     0     0     0     0     0     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     1     1     1     0     0
     0     0     0     0     0     0     0     0     0     0
     0     0     0     0     0     0     0     0     0     0
```

See Also

`conndef`, `imreconstruct`, `imregionalmax`

imresize

Purpose Resize an image

Syntax

```
B = imresize(A,m)
B = imresize(A,m,method)
B = imresize(A,[mrows ncols],method)

B = imresize(...,method,n)
B = imresize(...,method,h)
```

Description

`B = imresize(A,m)` returns an image `B` that is `m` times the size of `A`, using nearest-neighbor interpolation. `A` can be an indexed image, grayscale image, RGB, or binary image. If `m` is between 0 and 1.0, `B` is smaller than `A`. If `m` is greater than 1.0, `B` is larger than `A`.

`B = imresize(A,m,method)` returns an image that is `m` times the size of `A` using the interpolation method specified by `method`. `method` is a string that can have one of these values. The default value is enclosed in braces (`{}`).

Value	Description
{ 'nearest' }	Nearest-neighbor interpolation
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation

`B = imresize(A,[mrows ncols],method)` returns an image of the size specified by `[mrows ncols]`. If the specified size does not produce the same aspect ratio as the input image has, the output image is distorted.

When the specified output size is smaller than the size of the input image, and `method` is `'bilinear'` or `'bicubic'`, `imresize` applies a lowpass filter before interpolation to reduce aliasing. The default filter size is 11-by-11.

You can specify a different order for the default filter using

```
B = imresize(...,method,n)
```

`n` is an integer scalar specifying the size of the filter, which is `n`-by-`n`. If `n` is 0 (zero), `imresize` omits the filtering step.

You can also specify your own filter using this syntax.

```
B = imresize(...,method,h)
```

`h` is any two-dimensional FIR filter (such as those returned by `ftrans2`, `fwind1`, `fwind2`, or `fsamp2`).

Class Support

The input image `A` can be numeric or logical and it must be nonsparse. The output image `B` is of the same class as the input image.

See Also

`imrotate`, `imtransform`, `tformarray`

`interp2` in the MATLAB Function Reference

imrotate

Purpose Rotate an image

Syntax

```
B = imrotate(A,angle)
B = imrotate(A,angle,method)
B = imrotate(A,angle,method,bbox)
```

Description B = imrotate(A,angle) rotates the image A by angle degrees in a counterclockwise direction, using the nearest-neighbor interpolation. To rotate the image clockwise, specify a negative angle.

B = imrotate(A,angle,method) rotates the image A by angle degrees in a counterclockwise direction, using the interpolation method specified by method. method is a string that can have one of these values. The default value is enclosed in braces ({}).

Value	Description
{ 'nearest' }	Nearest-neighbor interpolation
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation Note: Bicubic interpolation can produce pixel values outside the original range.

B = imrotate(A,angle,method,bbox) rotates the image A through angle degrees. The bbox argument specifies the bounding box of the returned image. bbox is a string that can have one of these values. The default value is enclosed in braces ({}).

Value	Description
'crop'	Output image B includes only the central portion of the rotated image and is the same size as A.
{ 'loose' }	Output image B includes the whole rotated image and is generally larger than the input image A. imrotate sets pixels in areas outside the original image to zero.


Class Support

The input image A can be numeric or logical and it must be nonsparse. The output image B is of the same class as the input image.


Example

This example reads solar spectra image, stored in FITS format, and rotates the image to bring it into horizontal alignment. A rotation of -1 degree is all that is required.

```
I = fitsread('solarspectra.fts');  
I = mat2gray(I);  
J = imrotate(I,-1,'bilinear','crop');  
imshow(I)  
figure, imshow(J)
```



Original Image
Image Courtesy Ann Walker



Rotated Image

See Also

[imcrop](#), [imresize](#), [imtransform](#), [tformarray](#)

imscrollpanel

Purpose Scroll panel for interactive image navigation

Syntax `hpanel = imscrollpanel(hparent, himage)`

Description `hpanel = imscrollpanel(hparent, himage)` creates a scroll panel containing the target image (the image to be navigated).. `himage` is a handle to the target image. `hparent` is a handle to the figure or uipanel that will contain the new scroll panel. `hpanel` is the handle to the scroll panel, which is a uipanel object.

A scroll panel makes an image scrollable. If the size or magnification makes an image too large to display in a figure on the screen, the scroll panel displays a portion of the image at 100% magnification (one screen pixel represents one image pixel). The scroll panel adds horizontal and vertical scroll bars to enable navigation around the image.

`imscrollpanel` changes the object hierarchy of the target image. Instead of the familiar figure->axes->image object hierarchy, `imscrollpanel` inserts several uipanel and uicontrol objects between the figure and the axes object.

API functions A scroll panel contains a structure of function handles, called an API. You can use the functions in this API to manipulate the scroll panel. To retrieve this structure, use the `iptgetapi` function, as in the following example.

```
api = iptgetapi(hpanel)
```

This table lists the scroll panel API functions, in the order they appear in the structure.

Scroll Panel API Function	Description
setMagnification	<p>Sets the magnification in units of screen pixels per image pixel.</p> <pre>mag = api.getMagnification(new_mag)</pre> <p>where new_mag is a scalar magnification factor.</p>
getMagnification	<p>Returns the current magnification factor in units of screen pixels per image pixel.</p> <pre>mag = api.getMagnification()</pre> <p>Multiply mag by 100 to convert to percentage. For example if mag=2, that's equivalent to 200% magnification.</p>
setMagnificationAndCenter	<p>Changes the magnification and makes the point cx,cy appear in the center of the scroll panel. This operation is equivalent to a simultaneous zoom and recenter.</p> <pre>api.setMagnificationAndCenter(mag,cx,cy)</pre>
findFitMag	<p>Returns the magnification factor that would make the image just fit in the scroll panel.</p> <pre>mag = api.findFitMag()</pre>
setVisibleLocation	<p>Moves the target image so that the specified location is visible. Scrollbars update.</p> <pre>api.setVisibleLocation(xmin,ymin)</pre> <pre>api.setVisibleLocation([xmin ymin])</pre>

imscrollpanel

Scroll Panel API Function	Description
getVisibleLocation	<p>Returns the location of the currently visible portion of the target image.</p> <pre>loc = api.getVisibleLocation()</pre> <p>where loc is a vector [xmin ymin].</p>
getVisibleImageRect	<p>Returns the current visible portion of the image.</p> <pre>r = api.getVisibleImageRect()</pre> <p>where r is a rectangle [xmin ymin width height].</p>
addNewMagnificationCallback	<p>Adds the function handle fcn to the list of new-magnification callback functions.</p> <pre>id = api.addNewMagnificationCallback(fcn)</pre> <p>Whenever the scroll panel magnification changes, each function in the list is called with the syntax:</p> <pre>fcn(mag)</pre> <p>where mag is a scalar magnification factor.</p> <p>The return value, id, is used only with removeNewMagnificationCallback.</p>
removeNewMagnificationCallback	<p>Removes the corresponding function from the new-magnification callback list.</p> <pre>api.removeNewMagnificationCallback(id)</pre> <p>where id is the identifier returned by addNewMagnificationCallback.</p>

Scroll Panel API Function	Description
addNewLocationCallback	<p>Adds the function handle <code>fcn</code> to the list of new-location callback functions.</p> <pre>id = api.addNewLocationCallback(fcn)</pre> <p>Whenever the scroll panel location changes, each function in the list is called with the syntax:</p> <pre>fcn(loc)</pre> <p>where <code>loc</code> is <code>[xmin ymin]</code>.</p> <p>The return value, <code>id</code>, is used only with <code>removeNewLocationCallback</code>.</p>
removeNewLocationCallback	<p>Removes the corresponding function from the new-location callback list.</p> <pre>api.removeNewLocationCallback(id)</pre> <p>where <code>id</code> is the identifier returned by <code>addNewLocationCallback</code>.</p>

Note

Scrollbar navigation as provided by `imscrollpanel` is incompatible with the default MATLAB figure navigation buttons (pan, zoom in, zoom out). The corresponding menu items and toolbar buttons should be removed in a custom GUI that includes a scrollable `uipanel` created by `imscrollpanel`.

When you run `imscrollpanel`, it appears to take over the entire figure because, by default, an `hpanel` object has `'Units'` set to `'normalized'` and `'Position'` set to `[0 0 1 1]`. If you want to see other children of `hparent` while using your new scroll panel, you must manually set the `'Position'` property of `hpanel`.

Example

This example creates a scroll panel with a magnification box and an overview tool.

1 Create a scroll panel

```
hFig = figure('Toolbar','none',...
             'Menubar','none');
```

imscrollpanel

```
hIm = imshow('saturn.png');
hSP = imscrollpanel(hFig,hIm);
set(hSP,'Units','normalized',...
     'Position',[0 .1 1 .9])
```

- 2 Add a magnification box and an overview tool.

```
hMagBox = immagbox(hFig,hIm);
pos = get(hMagBox,'Position');
set(hMagBox,'Position',[0 0 pos(3) pos(4)])
imoverview(hIm)
```

- 3 Get the scroll panel API to programmatically control the view.

```
api = iptgetapi(hSP);
```

- 4 Get the current magnification and position.

```
mag = api.getMagnification()
r = api.getVisibleImageRect()
```

- 5 View the top left corner of the image.

```
api.setVisibleLocation(0.5,0.5)
```

- 6 Change the magnification to the value that just fits.

```
api.setMagnification(api.findFitMag())
```

- 7 Zoom in to 1600% on the dark spot.

```
api.setMagnificationAndCenter(16,306,800)
```

See also

`immagbox`, `imoverview`, `imoverviewpanel`, `imtool`, `iptgetapi`

Purpose

Display image

Syntax

```
imshow  
imshow(I)  
imshow(I,[low high])  
imshow(RGB)  
imshow(BW)  
imshow(X,map)  
imshow(filename)  
himage = imshow(...)  
imshow(...,param1,val1,param2,val2)
```

Description

`imshow` displays the file chooser dialog box so you can select an image file interactively.

`imshow(I)` displays the intensity image `I`.

`imshow(I,[low high])` displays `I` as a grayscale intensity image, specifying the display range for `I`. The value `low` (and any value less than `low`) displays as black; the value `high` (and any value greater than `high`) displays as white. Values in between are displayed as intermediate shades of gray, using the default number of gray levels. If you use an empty matrix (`[]`) for `[low high]`, `imshow` uses `[min(I(:)) max(I(:))]`; that is, the minimum value in `I` is displayed as black, and the maximum value is displayed as white.

`imshow(RGB)` displays the truecolor image `RGB`.

`imshow(BW)` displays the binary image `BW`. `imshow` displays pixels with the value 0 (zero) as black and pixels with the value 1 as white.

`imshow(X,map)` displays the indexed image `X` with the colormap `map`. A color map matrix may have any number of rows, but it must have exactly 3 columns. Each row is interpreted as a color, with the first element specifying the intensity of red light, the second green, and the third blue. Color intensity can be specified on the interval 0.0 to 1.0.

`imshow(filename)` displays the image stored in the graphics file `filename`. The file must contain an image that can be read by `imread` or `dicomread`. `imshow` calls `imread` or `dicomread` to read the image from the file, but does not store the image data in the MATLAB workspace. If the file contains multiple images,

imshow

the first one will be displayed. The file must be in the current directory or on the MATLAB path.

`himage = imshow(...)` returns the handle to the image object created by `imshow`.

`imshow(...,param1,val1,param2,val2,...)` displays the image, specifying parameters and corresponding values that control various aspects of the image display. The following table lists all `imshow` parameters. Parameter names can be abbreviated, and case does not matter..

Parameter	Value
'DisplayRange'	Two-element vector [LOW HIGH] that controls the display range of an intensity image. See above for more details about how to set [LOW HIGH]. Note: Including the parameter name is optional, except when the image is specified by a filename. The syntax <code>imshow(I,[LOW HIGH])</code> is equivalent to <code>imshow(I,'DisplayRange',[LOW HIGH])</code> . However, the 'DisplayRange' parameter must be specified when calling <code>imshow</code> with a filename, for example <code>imshow(filename,'DisplayRange'[LOW HIGH])</code> .
'InitialMagnification'	Any scalar value, or the text string 'fit' that specifies the initial magnification used to display the image. When set to 100, <code>imshow</code> displays the image at 100% magnification (one screen pixel for each image pixel). When set to 'fit', <code>imshow</code> scales the entire image to fit in the window. Initially, <code>imshow</code> always displays the entire image. If the magnification value is large enough that the image would be too big to display on the screen, <code>imshow</code> warns and displays the image at the largest magnification that fits on the screen. By default, the initial magnification parameter is set to the value returned by <code>iptgetpref('ImshowInitialMagnification')</code> .

Parameter	Value
'XData'	Two-element vector that establishes a nondefault spatial coordinate system by specifying the image XData. The value can have more than two elements, but only the first and last elements are actually used..
'YData'	Two-element vector that establishes a nondefault spatial coordinate system by specifying the image YData. The value can have more than two elements, but only the first and last elements are actually used..

Class Support

A truecolor image can be `uint8`, `uint16`, `double`, or `single`. An indexed image can be `logical`, `uint8`, `double`, or `single`. An intensity image can be `logical`, `uint8`, `double`, `single`, `int16`, or `uint16`. The input image must be nonsparse.

If your image is `int16` or `single`, the `CData` in the resulting image object will be `double`. For all other classes, the `CData` matches the input image class..

Related Toolbox Preferences

You can use the `iptsetpref` function to set several toolbox preferences that modify the behavior of `imshow`.

- `'ImshowBorder'` controls whether `imshow` displays the image with a border around it.
- `'ImshowAxesVisible'` controls whether `imshow` displays the image with the axes box and tick labels.
- `'ImshowInitialMagnification'` controls the initial magnification for image display, unless you override it in a particular call by specifying `imshow(..., 'InitialMagnification', initial_mag)`.

For more information about these preferences, see `iptsetpref`.

Remarks

`imshow` is the toolbox's fundamental image display function, optimizing figure, axes, and image object property settings for image display. `imtool` provides all the image display capabilities of `imshow` but also provides access to several other tools for navigating and exploring images, such as the Pixel Region tool, Image Information tool, and the Adjust Contrast tool. `imtool` presents an integrated environment for displaying images and performing some common image processing tasks.

imshow

Examples

Display an image from a file

```
imshow('board.tif')
```

Display an indexed image

```
[X,map] = imread('trees.tif');  
imshow(X,map)
```

Display an intensity image

```
I = imread('cameraman.tif');  
imshow(I)
```

Display an intensity image, adjust the display range

```
h = imshow(I,[0 80]);
```

See Also

`imread`, `imtool`, `iptgetpref`, `iptsetpref`, `subimage`, `truesize`, `warp`
`image`, `imagesc` in the MATLAB Function Reference

Purpose Subtract one image from another, or subtract a constant from an image

Syntax `Z = imsubtract(X,Y)`

Description `Z = imsubtract(X,Y)` subtracts each element in array `Y` from the corresponding element in array `X` and returns the difference in the corresponding element of the output array `Z`. `X` and `Y` are real, nonsparse numeric arrays of the same size and class, or `Y` is a double scalar. The array returned, `Z`, has the same size and class as `X` unless `X` is logical, in which case `Z` is double.

If `X` is an integer array, then elements of the output that exceed the range of the integer type are truncated, and fractional values are rounded.

If `X` and `Y` are double arrays, then you can use the expression `X - Y` instead of this function.

Note On Intel architecture processors, `imsubtract` can take advantage of the Intel Performance Primitives Library (IPPL), thus accelerating its execution time. IPPL is activated only if array `X` is of class `uint8`, `int16`, or `single`.

Examples Subtract two `uint8` arrays. Note that negative results are rounded to 0.

```
X = uint8([ 255 10 75; 44 225 100]);
Y = uint8([ 50 50 50; 50 50 50 ]);
Z = imadd(X,Y)
Z =
```

```
205    0   25
   0  175   50
```

Estimate and subtract the background of an image:

```
I = imread('rice.png');
background = imopen(I,stre1('disk',15));
Ip = imsubtract(I,background);
imshow(Ip,[])
```

Subtract a constant value from an image:

imsubtract

```
I = imread('rice.png');  
Iq = imsubtract(I,50);  
figure, imshow(I), figure, imshow(Iq)
```

See Also

[imabsdiff](#), [imadd](#), [imcomplement](#), [imdivide](#), [imlincomb](#), [immultiply](#), [ippl](#)

Purpose Display image in the Image Tool

Syntax

```
imtool
imtool(I)
imtool(I,[LOW HIGH])
imtool(RGB)
imtool(BW)
imtool(X,map)
imtool(filename)
hfig = imtool(...)
imtool close all
imtool(...,param1,val1,param2,val2,...)
```

Description `imtool` opens a new Image Tool in an empty state. Use the **File** menu options **Open** or **Import from Workspace** to choose an image for display.

`imtool(I)` displays the intensity image `I`.

`imtool(I,[LOW HIGH])` displays `I` as a grayscale intensity image, specifying the display range for `I`. The value `LOW` (and any value less than `LOW`) is displayed as black, the value `HIGH` (and any value greater than `HIGH`) is displayed as white. Values in between are displayed as intermediate shades of gray. `imtool` uses the default number of gray levels. If you use an empty matrix (`[]`) for `[LOW HIGH]`, `imtool` uses `[min(I(:)) max(I(:))]`; the minimum value in `I` is displayed as black, and the maximum value is displayed as white.

`imtool(RGB)` displays the truecolor image `RGB`.

`imtool(BW)` displays the binary image `BW`. Values of 0 is displayed as black, and values of 1 is displayed as white.

`imtool(X,map)` displays the indexed image `X` with colormap `map`.

`imtool(filename)` displays the image contained in the graphics file `filename`. The file must contain an image that can be read by `imread` or `dicomread`. `imtool` calls `imread` or `dicomread` to read the image from the file, but the image data is not stored in the MATLAB workspace. If the file contains multiple images, the first one is displayed. The file must be in the current directory or on the MATLAB path.

imtool

`hfig = imtool(...)` returns a handle `Hfig` to the figure created by `imtool`.
`close(Hfig)` closes the image tool.

`imtool close all` closes all image tools.

`imtool(...,param1,val1,param2,val2,...)` displays the image, specifying parameters and corresponding values that control various aspects of the image display. The following table lists all `imshow` parameters. Parameter names can be abbreviated, and case does not matter.

Parameter	Value
'DisplayRange'	Two-element vector [LOW HIGH] that controls the display range of an intensity image. See above for more details about how to set [LOW HIGH]. Note: Including the parameter name is optional, except when the image is specified by a filename. The syntax <code>imtool(I,[LOW HIGH])</code> is equivalent to <code>imtool(I,'DisplayRange',[LOW HIGH])</code> . However, the 'DisplayRange' parameter must be specified when calling <code>imtool</code> with a filename, as in the syntax <code>imtool(filename,'DisplayRange'[LOW HIGH])</code> .
'InitialMagnification'	One of two text strings: 'adaptive' or 'fit' or a numeric scalar value that specifies the initial magnification used to display the image. When set to 'adaptive', the entire image is visible on initial display. If the image is too large to display on the screen, <code>imtool</code> displays the image at the largest magnification that fits on the screen. When set to 'fit', <code>imtool</code> scales the entire image to fit in the window. When set to a numeric value, the value specifies the magnification as a percentage. For example, if you specify 100, the image is displayed at 100% magnification (one screen pixel for each image pixel). By default, the initial magnification parameter is set to the value returned by <code>iptgetpref('ImtoolInitialMagnification')</code> .

Class Support

A truecolor image can be `uint8`, `uint16`, `double`, or `single`. An indexed image can be `logical`, `uint8`, `double`, or `single`. An intensity image can be `logical`, `uint8`, `double`, `single`, `int16`, or `uint16`. The input image must be nonsparse.

**Related
Toolbox
Preferences**

You can use the `iptsetpref` function to set the following toolbox preference that modifies the behavior of `imtool`.

- `'ImtoolInitialMagnification'` controls the initial magnification for image display, unless you override it in a particular call by specifying `imtool(..., 'InitialMagnification', initial_mag)`.

For more information about these preferences, see the reference entry for `iptsetpref`.

Remarks

`imshow` is the toolbox's fundamental image display function, optimizing figure, axes, and image object property settings for image display. `imtool` provides all the image display capabilities of `imshow` but also provides access to several other tools for navigating and exploring images, such as the Pixel Region tool, Image Information tool, and the Adjust Contrast tool. `imtool` presents an integrated environment for displaying images and performing some common image processing tasks.

Examples

Display an image from a file.

```
imtool('board.tif')
```

Display an indexed image.

```
[X,map] = imread('trees.tif');  
imtool(X,map)
```

Display an intensity image.

```
I = imread('cameraman.tif');  
imtool(I)
```

Display an intensity image, adjusting the display range.

```
h = imtool(I,[0 80]);  
close(h)
```

imtool

See Also

`getimage`, `imageinfo`, `imcontrast`, `imshow`, `imdisplayrange`, `imgetfile`,
`imoverview`, `impixelinfo`, `impixelregion`, `imread`, `imshow`, `iptgetpref`,
`ipticondir`, `iptsetpref`, `iptwindowalign`

Purpose Perform top-hat filtering

Syntax
`IM2 = imtophat(IM,SE)`
`IM2 = imtophat(IM,NHOOD)`

Description `IM2 = imtophat(IM,SE)` performs morphological top-hat filtering on the grayscale or binary input image `IM` using the structuring element `SE`, where `SE` is returned by `strel`. `SE` must be a single structuring element object, not an array containing multiple structuring element objects.


`IM2 = imtophat(IM,NHOOD)`, where `NHOOD` is an array of 0's and 1's that specifies the size and shape of the structuring element, is the same as `imtophat(IM,strel(NHOOD))`.

Class Support `IM` can be numeric or logical and must be nonsparse. The output image `IM2` has the same class as the input image. If the input is binary (`logical`), the structuring element must be flat.

Example You can use top-hat filtering to correct uneven illumination when the background is dark. This example uses top-hat filtering with a disk-shaped structuring element to remove the uneven background illumination from an image.

1 Read an image into the MATLAB workspace.


```
I = imread('rice.png');
imshow(I)
```



imtophat


- 2 Create the structuring element and perform top-hat filtering of the image.

```
se = strel('disk',12);  
J = imtophat(I,se);  
figure, imshow(J)
```



- 3 Use `imadjust` to improve the visibility of the result.

```
K = imadjust(J);  
figure, imshow(K)
```



See Also

`imbothat`, `strel`

Purpose Apply 2-D spatial transformation to image

Syntax

```
B = imtransform(A, TFORM)
B = imtransform(A, TFORM, INTERP)
[B, XDATA, YDATA] = imtransform(...)
[B, XDATA, YDATA] = imtransform(..., param1, val1, param2, val2, ...)
```

Description `B = imtransform(A, TFORM)` transforms the image `A` according to the 2-D spatial transformation defined by `TFORM`, which is a spatial transformation structure (`TFORM`) as returned by `maketform` or `cp2tform`. If `ndims(A) > 2`, such as for an RGB image, then the same 2-D transformation is automatically applied to all 2-D planes along the higher dimensions.

When you use this syntax, `imtransform` automatically shifts the origin of your output image to make as much of the transformed image visible as possible. If you are using `imtransform` to do image registration, this syntax is not likely to give you the results you expect; you might want to set `'XData'` and `'YData'` explicitly.

`B = imtransform(A, TFORM, INTERP)` specifies the form of interpolation to use. `INTERP` can have one of these values. The default value is enclosed in braces (`{}`).

Value	Description
<code>'bicubic'</code>	Bicubic interpolation
<code>{'bilinear'}</code>	Bilinear interpolation
<code>'nearest'</code>	Nearest-neighbor interpolation

Alternatively, `INTERP` can be a `RESAMPLER` structure returned by `makeresampler`. This option allows more control over how resampling is performed.

`[B, XDATA, YDATA] = imtransform(...)` returns the location of the output image `B` in the output `X-Y` space. `XDATA` and `YDATA` are two-element vectors. The elements of `XDATA` specify the x -coordinates of the first and last columns of `B`. The elements of `YDATA` specify the y -coordinates of the first and last rows of `B`. Normally, `imtransform` computes `XDATA` and `YDATA` automatically so that `B`

imtransform

contains the entire transformed image A. However, you can override this automatic computation; see below.

`[B,XDATA,YDATA] = imtransform(...,param1,val1,param2,val2,...)` specifies parameters that control various aspects of the spatial transformation. This table lists all the parameters you can specify. Note that parameter names can be abbreviated and are not case sensitive.

Parameter	Description
'UData' 'VData'	Both of these parameters are two-element real vectors. 'UData' and 'VData' specify the spatial location of the image A in the 2-D input space U-V. The two elements of 'UData' give the <i>u</i> -coordinates (horizontal) of the first and last columns of A, respectively. The two elements of 'VData' give the <i>v</i> -coordinates (vertical) of the first and last rows of A, respectively. The default values for 'UData' and 'VData' are <code>[1 size(A,2)]</code> and <code>[1 size(A,1)]</code> , respectively.
'XData' 'YData'	Both of these parameters are two-element real vectors. 'XData' and 'YData' specify the spatial location of the output image B in the 2-D output space X-Y. The two elements of 'XData' give the <i>x</i> -coordinates (horizontal) of the first and last columns of B, respectively. The two elements of 'YData' give the <i>y</i> -coordinates (vertical) of the first and last rows of B, respectively.
	If 'XData' and 'YData' are not specified, <code>imtransform</code> estimates values for them that will completely contain the entire transformed output image.

Parameter	Description
'XYScale'	A one- or two-element real vector. The first element of 'XYScale' specifies the width of each output pixel in X-Y space. The second element (if present) specifies the height of each output pixel. If 'XYScale' has only one element, then the same value is used for both width and height. If 'XYScale' is not specified but 'Size' is, then 'XYScale' is computed from 'Size', 'XData', and 'YData'. If neither 'XYScale' nor 'Size' is provided, then the scale of the input pixels is used for 'XYScale'.
'Size'	A two-element vector of nonnegative integers. 'Size' specifies the number of rows and columns of the output image B. For higher dimensions, the size of B is taken directly from the size of A. In other words, $\text{size}(B, k) = \text{size}(A, k)$ for $k > 2$. If 'Size' is not specified, then it is computed from 'XData', 'YData', and 'XYScale'.

imtransform

Parameter	Description												
'FillValues'	<p>An array containing one or several fill values. Fill values are used for output pixels when the corresponding transformed location in the input image is completely outside the input image boundaries. If A is 2-D, 'FillValues' must be a scalar. However, if A's dimension is greater than two, then 'FillValues' can be an array whose size satisfies the following constraint: <code>size(fill_values,k)</code> must equal either <code>size(A,k+2)</code> or 1.</p> <p>For example, if A is a uint8 RGB image that is 200-by-200-by-3, then possibilities for 'FillValues' include</p> <table><tbody><tr><td>0</td><td>Fill with black</td></tr><tr><td>[0;0;0]</td><td>Fill with black</td></tr><tr><td>255</td><td>Fill with white</td></tr><tr><td>[255;255;255]</td><td>Fill with white</td></tr><tr><td>[0;0;255]</td><td>Fill with blue</td></tr><tr><td>[255;255;0]</td><td>Fill with yellow</td></tr></tbody></table>	0	Fill with black	[0;0;0]	Fill with black	255	Fill with white	[255;255;255]	Fill with white	[0;0;255]	Fill with blue	[255;255;0]	Fill with yellow
0	Fill with black												
[0;0;0]	Fill with black												
255	Fill with white												
[255;255;255]	Fill with white												
[0;0;255]	Fill with blue												
[255;255;0]	Fill with yellow												
	<p>If A is 4-D with size 200-by-200-by-3-by-10, then 'FillValues' can be a scalar, 1-by-10, 3-by-1, or 3-by-10.</p>												

Notes

- When you do not specify the output-space location for B using 'XData' and 'YData', `imtransform` estimates them automatically using the function `findbounds`. For some commonly used transformations, such as affine or projective, for which a forward mapping is easily computable, `findbounds` is fast. For transformations that do not have a forward mapping, such as the polynomial ones computed by `cp2tform`, `findbounds` can take significantly longer. If you can specify 'XData' and 'YData' directly for such transformations, `imtransform` might run noticeably faster.
- The automatic estimate of 'XData' and 'YData' using `findbounds` is not guaranteed in all cases to completely contain all the pixels of the transformed input image.

- The output values XDATA and YDATA might not exactly equal the input 'XData' and 'YData' parameters. This can happen either because of the need for an integer number of rows and columns, or if you specify values for 'XData', 'YData', 'XYScale', and 'Size' that are not entirely consistent. In either case, the first element of XDATA and YDATA always equals the first element of 'XData' and 'YData', respectively. Only the second elements of XDATA and YDATA might be different.
- imtransform assumes spatial-coordinate conventions for the transformation TFORM. Specifically, the first dimension of the transformation is the horizontal or *x*-coordinate, and the second dimension is the vertical or *y*-coordinate. Note that this is the reverse of the array subscripting convention in MATLAB.
- TFORM must be a 2-D transformation to be used with imtransform. For arbitrary-dimensional array transformations, see tformarray.

Class Support

The input image A can be of any nonsparse numeric class, real or complex, or it can be of class logical. The class of B is the same as the class of A.

Example

Example 1

Apply a horizontal shear to an intensity image.

```
I = imread('cameraman.tif');
tform = maketform('affine',[1 0 0; .5 1 0; 0 0 1]);
J = imtransform(I,tform);
imshow(I), figure, imshow(J)
```

Example 2

A projective transformation can map a square to a quadrilateral. In this example, set up an input coordinate system so that the input image fills the unit square and then transform the image into the quadrilateral with vertices (0 0), (1 0), (1 1), (0 1) to the quadrilateral with vertices (-4 2), (-8 3), (-3 -5), (6 3). Fill with gray and use bicubic interpolation. Make the output size the same as the input size.

```
I = imread('cameraman.tif');
udata = [0 1]; vdata = [0 1]; % input coordinate system
tform = maketform('projective',[ 0 0; 1 0; 1 1; 0 1],...
                    [-4 2; -8 -3; -3 -5; 6 3]);
[B,xdata,ydata] = imtransform(I, tform, 'bicubic', ...
```

imtransform

```
        'udata', udata,...
        'vdata', vdata,...
        'size', size(I),...
        'fill', 128);
subplot(1,2,1), imshow(udata,vdata,I), axis on
subplot(1,2,2), imshow(xdata,ydata,B), axis on
```

Example 3

Register an aerial photo to an orthophoto.

Read in the aerial photo.

```
unregistered = imread('westconcordaerial.png');
figure, imshow(unregistered)
```

Read in the orthophoto.

```
figure, imshow('westconcordorthophoto.png')
```

Load control points that were previously picked.

```
load westconcordpoints
```

Create a transformation structure for a projective transformation.

```
t_concord = cp2tform(input_points,base_points,'projective');
```

Get the width and height of the orthophoto and perform the transformation.

```
info = imfinfo('westconcordorthophoto.png');

registered = imtransform(unregistered,t_concord,...
    'XData',[1 info.Width], 'YData',[1 info.Height]);
figure, imshow(registered)
```

See Also

checkerboard, cp2tform, imresize, imrotate, maketform, makesampler, tformarray

Purpose

Display image in the image tool

Note This function is obsolete and may be removed in future versions. Use `imtool` instead.

imwrite

Purpose

Write image to graphics file

`imwrite` is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose Convert an indexed image to an intensity image

Syntax `I = ind2gray(X,map)`

Description `I = ind2gray(X,map)` converts the image `X` with colormap `map` to an intensity image `I`. `ind2gray` removes the hue and saturation information from the input image while retaining the luminance.

Class Support `X` can be of class `uint8`, `uint16`, `single`, or `double`. `map` is `double`. `I` is of the same class as `X`.

Example

```
load trees
I = ind2gray(X,map);
imshow(X,map)
figure, imshow(I)
```



Image Courtesy of Susan Cohen

Algorithm `ind2gray` converts the colormap to NTSC coordinates using `rgb2ntsc`, and sets the hue and saturation components (I and Q) to zero, creating a gray colormap. `ind2gray` then replaces the indices in the image `X` with the corresponding grayscale intensity values in the gray colormap.

See Also `gray2ind`, `imshow`, `imtool`, `rgb2ntsc`

ind2rgb

Purpose	Convert an indexed image to an RGB image
Syntax	<code>RGB = ind2rgb(X,map)</code>
Description	<code>RGB = ind2rgb(X,map)</code> converts the matrix <code>X</code> and corresponding colormap <code>map</code> to RGB (truecolor) format.
Class Support	<code>X</code> can be of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . <code>RGB</code> is an <code>m-by-n-by-3</code> array of class <code>double</code> .
See Also	<code>ind2gray</code> , <code>rgb2ind</code>

Purpose	Convert integer values using lookup table
Syntax	<code>B = intlut(A, LUT)</code>
Description	<p><code>B = intlut(A, LUT)</code> converts values in array <code>A</code> based on lookup table <code>LUT</code> and returns these new values in array <code>B</code>.</p> <p>For example, if <code>A</code> is a vector whose kth element is equal to <code>alpha</code>, then <code>B(k)</code> is equal to the <code>LUT</code> value corresponding to <code>alpha</code>, i.e., <code>LUT(alpha+1)</code>.</p>
Class Support	<code>A</code> can be <code>uint8</code> , <code>uint16</code> , or <code>int16</code> . If <code>A</code> is <code>uint8</code> , <code>LUT</code> must be a <code>uint8</code> vector with 256 elements. If <code>A</code> is <code>uint16</code> or <code>int16</code> , <code>LUT</code> must be a vector with 65536 elements that has the same class as <code>A</code> . <code>B</code> has the same size and class as <code>A</code> .
Example	<pre>A = uint8([1 2 3 4; 5 6 7 8; 9 10 11 12]) LUT = repmat(uint8([0 150 200 255]),1,64); B = intlut(A, LUT)</pre>
See Also	<code>ind2gray</code> , <code>rgb2ind</code>

ippl

Purpose	Check for presence of Intel Performance Primitives Library (IPPL)
Syntax	<pre>TF = ippl [TF B] = ippl</pre>
Description	<p>The Intel Performance Primitives Library (IPPL) provides a collection of basic functions used in signal and image processing. The IPPL takes advantage of the parallelism of the Single-Instruction, Multiple-Data (SIMD) instructions that make up the core of the MMX technology and Streaming SIMD Extensions. These instructions are available only on the Intel architecture processors. IPPL is used by some of the Image Processing Toolbox functions to accelerate their execution time.</p> <p><code>TF = ippl</code> returns true (1) if IPPL is available and false (0) otherwise.</p> <p><code>[TF B] = ippl</code> returns an additional column cell array <code>B</code>. Each row of <code>B</code> contains a string describing a specific IPPL module.</p> <p>When IPPL is available, the Image Processing Toolbox image arithmetic functions (<code>imabsdiff</code>, <code>imadd</code>, <code>imsubtract</code>, <code>imdivide</code>, <code>immultiply</code>, and <code>imlincomb</code>) and the <code>imfilter</code> function take advantage of it. Toolbox functions that use these functions also benefit.</p>
Notes	<p>IPPL is utilized only for some data types and only under specific conditions. See the help sections of the functions listed above for detailed information on when IPPL is activated.</p> <p>The IPPL function is likely to change.</p>
See Also	<code>imabsdiff</code> , <code>imadd</code> , <code>imdivide</code> , <code>imfilter</code> , <code>imlincomb</code> , <code>immultiply</code> , <code>imsubtract</code>

Purpose Add function handle to callback list

Syntax ID = iptaddcallback(h,callback,func_handle)

Description ID = iptaddcallback(h,callback,func_handle) adds the function handle func_handle to the list of functions to be called when the callback specified by callback executes. callback is a string specifying the name of a callback property of the Handle Graphics object specified by the handle h.

iptaddcallback returns a unique callback identifier, ID, that can be used with iptremovecallback to remove the function from the callback list.

iptaddcallback can be useful when you need to notify more than one tool about the same callback event for a single object.

Note Callback functions that have already been added to an object using the set command continue to work after you call iptaddcallback. The first time you call iptaddcallback for a given object and callback, the function checks to see if a different callback function is already installed. If a callback is already installed, iptaddcallback replaces that callback function with the iptaddcallback callback processor, and then adds the preexisting callback function to the iptaddcallback list.

Example This example creates a figure and registers two callback functions. Whenever MATLAB detects mouse motion over the figure, function handles f1 and f2 are called in the order in which they were added to the list.

```
h = figure;  
f1 = @(varargin) disp('Callback 1');  
f2 = @(varargin) disp('Callback 2');  
iptaddcallback(h, 'WindowButtonMotionFcn', f1);  
iptaddcallback(h, 'WindowButtonMotionFcn', f2);
```

See Also iptremovecallback

iptcheckconn

Purpose Check validity of connectivity argument

Syntax `iptcheckconn(CONN, func_name, var_name, arg_pos)`

Description `iptcheckconn(CONN, func_name, var_name, arg_pos)` checks whether CONN is a valid connectivity argument. If it is invalid, the function issues a formatted error message.

A connectivity argument can be one of the following scalar values: 1, 4, 6, 8, 18, or 26. A connectivity argument can also be a 3-by-3-by- ... -by-3 array of 0's and 1s. The central element of a connectivity array must be nonzero and the array must be symmetric about its center.

`func_name` is a string that specifies the name used in the formatted error message to identify the function checking the connectivity argument.

`var_name` is a string that specifies the name used in the formatted error message to identify the argument being checked.

`arg_pos` is a positive integer that indicates the position of the argument being checked in the function argument list. `iptcheckconn` converts this value to an ordinal number and includes this information in the formatted error message.

Class Support CONN must be of class `double` or `logical` and must be real and nonsparse.

Example To trigger this error message, this example creates a 4-by-4 array and passes it as the connectivity argument.

```
iptcheckconn(eye(4), 'func_name', 'var_name', 2)
```

See Also `iptnum2ordinal`

Purpose Check validity of handle

Syntax `iptcheckhandle(H,valid_types,func_name,var_name,arg_pos)`

Description `iptcheckhandle(H,valid_types,func_name,var_name,arg_pos)` checks the validity of the handle `H` and issues a formatted error message if the handle is invalid. `H` must be a handle to a single figure, `uipanel`, `hggroup`, `axes`, or image object.

`valid_types` is a cell array of strings specifying the set of Handle Graphics object types to which `H` is expected to belong. For example, if you specify `valid_types` as `{'uipanel','figure'}`, `H` can be a handle to either a `uipanel` object or a figure object.

`func_name` is a string that specifies the name used in the formatted error message to identify the function checking the handle.

`var_name` is a string that specifies the name used in the formatted error message to identify the argument being checked.

`arg_pos` is a positive integer that indicates the position of the argument being checked in the function argument list. `iptcheckhandle` converts this value to an ordinal number and includes this information in the formatted error message.

Example To trigger the error message, create a figure that does not contain an axes object and then check for a valid axes handle.

```
fig = figure; % create figure without an axes
iptcheckhandle(fig,{'axes'},'my_function','my_variable',2)
```

The following shows the format of the error message and indicates which parts you can customize using `iptcheckhandle` arguments.

iptcheckhandle

func_name arg_pos var_name
 ↓ ↓ ↓
Function MY_FUNCTION expected its second input argument, my_variable,
to be a handle of one of these types:

axes ←————— valid_types

Instead, its type was: figure.

See Also

iptcheckinput, iptcheckmap, iptchecknargin, iptcheckstrs,
iptnum2ordinal

- Purpose** Check validity of array and issue error message if invalid
- Syntax** `iptcheckinput(A,classes,attributes,func_name,var_name,arg_pos)`
- Description** `iptcheckinput(A,classes,attributes,func_name,var_name,arg_pos)` checks the validity of the array `A` and issues a formatted error message if it is invalid.
- `classes` is a cell array of strings specifying the set of classes to which `A` is expected to belong. For example, if you specify `classes` as `{'logical' 'cell'}`, `A` is required to be either a logical array or a cell array. The string `'numeric'` is interpreted as an abbreviation for the classes `uint8`, `uint16`, `uint32`, `int8`, `int16`, `int32`, `single`, and `double`.
- `attributes` is a cell array of strings specifying the set of attributes that `A` must satisfy. For example, if `attributes` is `{'real' 'nonempty' 'finite'}`, `A` must be real and nonempty, and it must contain only finite values. The following table lists the supported attributes in alphabetical order.
- | | | | |
|----------------------|-----------------------------|-----------------------|---------------------|
| <code>2d</code> | <code>nonemptyvector</code> | <code>odd</code> | <code>twod</code> |
| <code>column</code> | <code>nonnan</code> | <code>positive</code> | <code>vector</code> |
| <code>even</code> | <code>nonnegative</code> | <code>real</code> | |
| <code>finite</code> | <code>nonsparse</code> | <code>row</code> | |
| <code>integer</code> | <code>nonzero</code> | <code>scalar</code> | |
- `func_name` is a string that specifies the name used in the formatted error message to identify the function checking the input.
- `var_name` is a string that specifies the name used in the formatted error message to identify the argument being checked.
- `arg_pos` is a positive integer that indicates the position of the argument being checked in the function argument list. `iptcheckinput` converts this value to an ordinal number and includes this information in the formatted error message.

Example Create a three-dimensional array.

```
A = [ 1 2 3; 4 5 6 ];
B = [ 7 8 9; 10 11 12];
```

iptcheckinput

```
C = cat(3,A,B);  
iptcheckinput(F,{'numeric'},{'2d'},'func_name','var_name',2)
```

The following shows the format of the error message and indicates which parts you can customize using `iptcheckinput` arguments.

func_name arg_pos var_name

↓ ↓ ↓

Function FUNC_NAME expected its second input, var_name, to be two-dimensional.

↑

attributes

See Also

`iptcheckhandle`, `iptcheckmap`, `iptchecknargin`, `iptcheckstrs`, `iptnum2ordinal`

Purpose Check validity of colormap

Syntax `iptcheckmap(map, func_name, var_name, arg_pos)`

Description `iptcheckmap(map, func_name, var_name, arg_pos)` checks the validity of the MATLAB colormap and issues an error message if it is invalid.

`func_name` is a string that specifies the name used in the formatted error message to identify the function checking the colormap.

`var_name` is a string that specifies the name used in the formatted error message to identify the argument being checked.

`arg_pos` is a positive integer that indicates the position of the argument being checked in the function argument list. `iptcheckmap` converts this value to an ordinal number and includes this information in the formatted error message.

Example

```
bad_map = ones(10);  
iptcheckmap(bad_map, 'func_name', 'var_name', 2)
```

The following shows the format of the error message and indicates which parts you can customize using `iptcheckmap` arguments.

func_name	arg_pos	var_name
↓	↓	↓
Function FUNC_NAME expected its second input argument, var_name, to be a valid colormap. Valid colormaps must be nonempty, double, 2-D matrices with 3 columns.		

See Also

`iptcheckhandle`, `iptcheckinput`, `iptchecknargin`, `iptcheckstrs`, `iptnum2ordinal`

iptchecknargin

Purpose Check number of input arguments

Syntax `iptchecknargin(low,high,num_inputs,func_name)`

Description `iptchecknargin(low,high,num_inputs,func_name)` checks whether `num_inputs` is in the range indicated by `low` and `high`. If not, `iptchecknargin` issues a formatted error message.

`low` should be a scalar nonnegative integer.

`high` should be a scalar nonnegative integer or `Inf`.

`func_name` is a string that specifies the name used in the formatted error message to identify the function checking the handle.

Example Create a function and use `iptchecknargin` to check that the number of arguments passed to the function is within the expected range.

```
function test_function(varargin)
    iptchecknargin(1,3,nargin,mfilename);
```

Trigger the error message by executing the function at the MATLAB command line, specifying more than the expected number of arguments.

```
test_function(1,2,3,4)
```

See Also `iptcheckhandle`, `iptcheckinput`, `iptcheckmap`, `iptcheckstrs`, `iptnum2ordinal`

Purpose Check validity of option string

Syntax `out=iptcheckstrs(in,valid_strs,func_name,var_name,arg_pos)`

Description `out=iptcheckstrs(in,valid_strs,func_name,var_name,arg_pos)` checks the validity of the option string `in`. It returns the matching string in `valid_strs` in `out`. `iptcheckstrs` looks for a case-insensitive, nonambiguous match between `in` and the strings in `valid_strs`.

`valid_strs` is a cell array containing strings.

`func_name` is a string that specifies the name used in the formatted error message to identify the function checking the strings.

`var_name` is a string that specifies the name used in the formatted error message to identify the argument being checked.

`arg_pos` is a positive integer that indicates the position of the argument being checked in the function argument list. `iptcheckstrs` converts this value to an ordinal number and includes this information in the formatted error message.

Example To trigger this error message, define a cell array of some text strings and pass in another string that isn't in the cell array.

```
iptcheckstrs('option3',{'option1','option2'},...
            'func_name','var_name',2)
```

The following shows the format of the error message and indicates which parts you can customize using `iptcheckhandle` arguments.

```

      func_name           arg_pos           var_name
      ↓                 ↓                 ↓
Function FUNC_NAME expected its second input argument, var_name,
to match one of these strings:

option1, option2   ← valid_strs
```

The input, 'option3', did not match any of the valid strings.

iptcheckstrs

See Also

`iptcheckhandle`, `iptcheckinput`, `iptcheckmap`, `iptchecknargin`,
`iptnum2ordinal`

Purpose	Display index of Image Processing Toolbox demos
Syntax	<code>iptdemos</code>
Description	<code>iptdemos</code> displays the HTML page that lists all the Image Processing demos. <code>iptdemos</code> displays the page in the MATLAB Help browser.

iptgetapi

Purpose Get application programmer interface (API) for handle

Syntax API = iptgetapi(h)

Description API = iptgetapi(h) returns the API structure associated with handle h if there is one. Otherwise, iptgetapi returns an empty array.

For more information about handle APIs, see the help for `immagbox`, `impositionrect`, or `imscrollpanel`.

Example

```
hFig = figure('Toolbar','none',...
             'Menubar','none');
hIm = imshow('tape.png');
hSP = imscrollpanel(hFig,hIm);
api = iptgetapi(hSP);
api.setMagnification(2) % 2X = 200%
```

See Also `immagbox`, `impositionrect`, `imscrollpanel`

Purpose Return Image Processing Toolbox preference

Syntax `prefs = iptgetpref`
`value = iptgetpref(prefname)`

Description `prefs = iptgetpref` returns a structure containing all the Image Processing Toolbox preferences with their current values. Each field in the structure has the name of an Image Processing Toolbox preference. See `iptsetpref` for a list.

`value = iptgetpref(prefname)` returns the value of the Image Processing Toolbox preference specified by the string `prefname`. See `iptsetpref` for a complete list of valid preference names. Preference names are not case sensitive and can be abbreviated.

Example

```
value = iptgetpref('ImshowAxesVisible')  
  
value =  
  
off
```

See Also `imshow`, `iptsetpref`

ipticondir

Purpose Return names of directories containing IPT and MATLAB icons

Syntax [D1 D2] = ipticondir

Description [D1, D2] = imicondir returns the names of the directories containing the Image Processing Toolbox icons (D1) and the MATLAB icons (D2).

Example

```
[iptdir, MATLABdir] = ipticondir  
dir(iptdir)
```

See Also imtool

Purpose Convert positive integer to ordinal string

Syntax `string = iptnum2ordinal(number)`

Description `string = iptnum2ordinal(number)` converts the positive integer number to the ordinal text string `string`.

Example The following example returns the string 'fourth'.

```
str = iptnum2ordinal(4)
```

The following example returns the string '23rd'.

```
str = iptnum2ordinal(23)
```

iptremovecallback

Purpose Delete function handle from callback list

Syntax `iptremovecallback(h,callback,ID)`

Description `iptremovecallback(h,callback,ID)` deletes a callback from the list of callbacks created by `imaddcallback` for the object with handle `h` and the associated callback string `callback`. `ID` is the identifier of the callback to be deleted. This `ID` is returned by `iptaddcallback` when you add the function handle to the callback list.

Example This example registers three callbacks for a particular figure and then removes one of them.

Register the callbacks and try them interactively.

```
h = figure;
f1 = @(varargin) disp('Callback 1');
f2 = @(varargin) disp('Callback 2');
f3 = @(varargin) disp('Callback 3');
id1 = iptaddcallback(h, 'WindowButtonMotionFcn', f1);
id2 = iptaddcallback(h, 'WindowButtonMotionFcn', f2);
id3 = iptaddcallback(h, 'WindowButtonMotionFcn', f3);
```

Remove one of the callbacks and then move the mouse over the figure again. Whenever MATLAB detects mouse motion over the figure, function handles `f1` and `f3` are called in that order.

```
iptremovecallback(h, 'WindowButtonMotionFcn', id2);
```

See Also `iptaddcallback`

Purpose Set Image Processing Toolbox preferences or display valid values

Syntax `iptsetpref(prefname)`
`iptsetpref(prefname,value)`

Description `iptsetpref(prefname)` displays the valid values for the Image Processing Toolbox preference specified by `prefname`.

`iptsetpref(prefname,value)` sets the Image Processing Toolbox preference specified by the string `prefname` to the value specified by `value`. The setting persists until the end of the current MATLAB session, or until you change the setting. (To make the value persist between sessions, put the command in your `startup.m` file.)

iptsetpref

This table describes the available preferences. Note that the preference names are case insensitive and can be abbreviated. The default value is enclosed in braces ({}).

Preference Name	Description
'ImshowBorder'	<p>Controls whether <code>imshow</code> includes a border around the image in the figure window. Possible values:</p> <p>{ 'loose' } — Include a border between the image and the edges of the figure window, thus leaving room for axes labels, titles, etc.</p> <p>'tight' — Adjust the figure size so that the image entirely fills the figure.</p> <p>Note: There can still be a border if the image is very small, or if there are other objects besides the image and its axes in the figure.</p>
'ImshowAxesVisible'	<p>Controls whether <code>imshow</code> displays images with the axes box and tick labels. Possible values:</p> <p>'on' — Include axes box and tick labels.</p> <p>{ 'off' } — Do not include axes box and tick labels.</p>

Preference Name	Description
'ImshowInitialMagnification'	<p>Controls the initial magnification of the image displayed by <code>imshow</code>. Possible values:</p> <p>Any numeric value — <code>imshow</code> interprets numeric values as a percentage. The default value is 100. One hundred percent magnification means that there should be one screen pixel for every image pixel.</p> <p>'fit' — Scale the image so that it fits into the window in its entirety.</p> <p>You can override this preference by specifying the 'InitialMagnification' parameter when you call <code>imshow</code>, or by calling the <code>truesize</code> function manually after displaying the image.</p>
'ImtoolInitialMagnification'	<p>Controls the initial magnification of the image displayed by <code>imtool</code>. Possible values:</p> <p>{ 'adaptive' } — Display the entire image. If the image is too large to display on the screen at 100% magnification, display the image at the largest magnification that fits on the screen. This is the default.</p> <p>Any numeric value — Specify the magnification as a percentage. One hundred percent magnification means that there should be one screen pixel for every image pixel.</p> <p>'fit' — Scale the image so that it fits into the window in its entirety.</p> <p>You can override this this preference by specifying the 'InitialMagnification' parameter when you call <code>imtool</code>.</p>

Example

```
iptsetpref('ImshowBorder','tight')
```

iptsetpref

See Also

`imshow`, `imshow`, `iptgetpref`, `truesize`


`axis` in the MATLAB Function Reference

Purpose Align figure windows

Syntax `iptwindowalign(fixed_fig, fixed_fig_edge, ...
moving_fig, moving_fig_edge)`

Description `iptwindowalign(fixed_fig, fixed_fig_edge, moving_fig, ...
moving_fig_edge)` moves the figure `moving_fig` to align it with the figure `fixed_fig`. `moving_fig` and `fixed_fig` are handles to figure objects.

`fixed_fig_edge` and `moving_fig_edge` describe the alignment of the figures in relation to their edges and can take any of the following values: 'left', 'right', 'hcenter', 'top', 'bottom', or 'vcenter'. 'hcenter' means center horizontally and 'vcenter' means center vertically. The following figure shows these alignments.



Notes The two specified locations must be consistent in terms of their direction. For example, you cannot specify 'left' for `fixed_fig_edge` and 'bottom' for `moving_fig_edge`.

iptwindowalign

`iptwindowalign` constrains the position adjustment of `moving_fig` to keep it entirely visible on the screen.

Examples

Create two figures: `fig1` and `fig2`.

```
fig1 = figure;  
fig2 = figure;
```

Move `fig2` so its left edge is aligned with the right edge of `fig1`.

```
iptwindowalign(fig1, 'right', fig2, 'left');
```

Move `fig2` so its top edge is aligned with `fig1`'s bottom edge, and then move it so the two figures are vertically centered.

```
iptwindowalign(fig1, 'bottom', fig2, 'top');  
iptwindowalign(fig1, 'vcenter', fig2, 'vcenter')
```

See Also

`imtool`

Purpose Compute inverse Radon transform

Syntax
`I = iradon(R,theta)`
`I = iradon(R,theta,interp,filter,frequency_scaling,output_size)`
`[I,H] = iradon(...)`

Description `I = iradon(R,theta)` reconstructs the image `I` from projection data in the two-dimensional array `R`. The columns of `R` are parallel beam projection data. `iradon` assumes that the center of rotation is the center point of the projections, which is defined as `ceil(size(R,1)/2)`.

`theta` describes the angles (in degrees) at which the projections were taken. It can be either a vector containing the angles or a scalar specifying `D_theta`, the incremental angle between projections. If `theta` is a vector, it must contain angles with equal spacing between them. If `theta` is a scalar specifying `D_theta`, the projections were taken at angles `theta = m*D_theta`, where `m = 0,1,2,...,size(R,2)-1`. If the input is the empty matrix (`[]`), `D_theta` defaults to `180/size(R,2)`.

`iradon` uses the filtered back-projection algorithm to perform the inverse Radon transform. The filter is designed directly in the frequency domain and then multiplied by the FFT of the projections. The projections are zero-padded to a power of 2 before filtering to prevent spatial domain aliasing and to speed up the FFT.

`I = iradon(P,theta,interp,filter,frequency_scaling,output_size)` specifies parameters to use in the inverse Radon transform. You can specify any combination of the last four arguments. `iradon` uses default values for any of these arguments that you omit.

`interp` specifies the type of interpolation to use in the back projection. The available options are listed in order of increasing accuracy and computational complexity. The default value is enclosed in braces (`{}`).

Value	Description
'nearest'	Nearest-neighbor interpolation

Value	Description
{ 'linear' }	Linear interpolation
'spline'	Spline interpolation

`filter` specifies the filter to use for frequency domain filtering. `filter` can be any of the strings that specify standard filters. The default value is enclosed in braces (`{}`).

Value	Description
{ 'Ram-Lak' }	Cropped Ram-Lak or ramp filter. The frequency response of this filter is $ f $. Because this filter is sensitive to noise in the projections, one of the filters listed below might be preferable. These filters multiply the Ram-Lak filter by a window that deemphasizes high frequencies.
'Shepp-Logan'	Multiplies the Ram-Lak filter by a sinc function
'Cosine'	Multiplies the Ram-Lak filter by a cosine function
'Hamming'	Multiplies the Ram-Lak filter by a Hamming window
'Hann'	Multiplies the Ram-Lak filter by a Hann window

`frequency_scaling` is a scalar in the range (0,1] that modifies the filter by rescaling its frequency axis. The default is 1. If `frequency_scaling` is less than 1, the filter is compressed to fit into the frequency range $[0, \text{frequency_scaling}]$, in normalized frequencies; all frequencies above `frequency_scaling` are set to 0.

`output_size` is a scalar that specifies the number of rows and columns in the reconstructed image. If `output_size` is not specified, the size is determined from the length of the projections.

$$n = 2 * \text{floor}(\text{size}(R,1) / (2 * \text{sqrt}(2)))$$

If you specify `output_size`, `iradon` reconstructs a smaller or larger portion of the image but does not change the scaling of the data. If the projections were

calculated with the `radon` function, the reconstructed image might not be the same size as the original image.

`[I,H] = iradon(...)` returns the frequency response of the filter in the vector `H`.

Class Support

`R` can be double or single. All other numeric input arguments must be of class double. `I` has the same class as `R`. `H` is double.

Example

```
P = phantom(128);  
R = radon(P,0:179);  
I = iradon(R,0:179,'nearest','Hann');  
imshow(P), figure, imshow(I)
```



Algorithm

`iradon` uses the filtered back projection algorithm to perform the inverse Radon transform. The filter is designed directly in the frequency domain and then multiplied by the FFT of the projections. The projections are zero-padded to a power of 2 before filtering to prevent spatial domain aliasing and to speed up the FFT.

See Also

`fan2para`, `fanbeam`, `ifanbeam`, `para2fan`, `phantom`, `radon`

References

[1] Kak, A. C., and M. Slaney, *Principles of Computerized Tomographic Imaging*, New York, NY, IEEE Press, 1988.

isbw

Purpose Return true for a binary image

Note This function is obsolete and may be removed in future versions. Use `islogical` instead.

Syntax `flag = isbw(A)`

Description `flag = isbw(A)` returns 1 if *A* is a binary image and 0 otherwise.
The input image *A* is considered to be a binary image if it is a nonsparse logical array.

Class Support The input image *A* can be any MATLAB array.

See Also `isind`, `isgray`, `isrgb`

Purpose	Return true for flat structuring element
Syntax	TF = isflat(SE)
Description	TF = isflat(SE) returns true (1) if the structuring element SE is flat; otherwise it returns false (0). If SE is an array of STREL objects, then TF is the same size as SE.
Class Support	SE is a STREL object. TF is a double-precision value.
See Also	strel

isgray

Purpose Return true for intensity image

Note This function is obsolete and may be removed in future versions.

Syntax `flag = isgray(A)`

Description `flag = isgray(A)` returns 1 if *A* is a grayscale intensity image and 0 otherwise.

`isgray` uses these criteria to decide whether *A* is an intensity image:

- If *A* is of class `double`, all values must be in the range [0,1], and the number of dimensions of *A* must be 2.
- If *A* is of class `uint16` or `uint8`, the number of dimensions of *A* must be 2.

Note A four-dimensional array that contains multiple intensity images returns 0, not 1.

Class Support The input image *A* can be of class `logical`, `uint8`, `uint16`, or `double`.

See Also `isbw`, `isind`, `isrgb`

Purpose Return true for valid ICC color profile

Syntax TF = isicc(P)

Description TF = isicc(P) returns True if structure P is a valid ICC color profile; otherwise False.

isicc checks if P has a complete set of the tags required for an ICC profile. P must contain a Header field, which in turn must contain a Version field and a DeviceClass field. These fields, and others, are used to determine the set of required tags according to the ICC Profile Specification, either Version 2 (ICC.1:2001-04) or Version 4 (ICC.1:2001-12), which are available at www.color.org. The set of required tags is given in Section 6.3 in either version.

Example Read in a profile and test its validity.

```
P = iccread('sRGB.icm');
```

```
TF = isicc(P)
```

```
TF =
```

```
1
```

Create a MATLAB structure and test its validity.

```
S.name = 'Any Student';
```

```
S.score = 83;
```

```
S.grade = 'B+'
```

```
TF = isicc(S)
```

```
TF =
```

```
0
```

See Also applycform, iccread, iccwrite, makecform

isind

Purpose Return true for an indexed image

Note This function is obsolete and may be removed in future versions.

Syntax `flag = isind(A)`

Description `flag = isind(A)` returns 1 if A is an indexed image and 0 otherwise.

`isind` uses these criteria to determine if A is an indexed image:

- If A is of class `double`, all values in A must be integers greater than or equal to 1, and the number of dimensions of A must be 2.
- If A is of class `uint8` or `uint16`, the number of dimensions of A must be 2.

Note A four-dimensional array that contains multiple indexed images returns 0, not 1.

Class Support A can be of class `logical`, `uint8`, `uint16`, or `double`.

See Also `isbw`, `isgray`, `isrgb`

Purpose Return true for an RGB image

Note This function is obsolete and may be removed in future versions.

Syntax `flag = isrgb(A)`

Description `flag = isrgb(A)` returns 1 if A is an RGB truecolor image and 0 otherwise. `isrgb` uses these criteria to determine whether A is an RGB image:

- If A is of class `double`, all values must be in the range [0,1], and A must be m-by-n-by-3.
- If A is of class `uint16` or `uint8`, A must be m-by-n-by-3.

Note A four-dimensional array that contains multiple RGB images returns 0, not 1.

Class Support A can be of class `logical`, `uint8`, `uint16`, or `double`.

See Also `isbw`, `isgray`, `isind`

lab2double

Purpose Convert $L^*a^*b^*$ data to double

Syntax `labd = lab2double(lab)`

Description `labd = lab2double(lab)` converts an M-by-3 or M-by-N-by-3 array of $L^*a^*b^*$ color values to class double. The output array `labd` has the same size as `lab`.

The Image Processing Toolbox follows the convention that double-precision $L^*a^*b^*$ arrays contain 1976 CIE $L^*a^*b^*$ values. $L^*a^*b^*$ arrays that are `uint8` or `uint16` follow the convention in the ICC profile specification (ICC.1:2001-4, www.color.org) for representing $L^*a^*b^*$ values as unsigned 8-bit or 16-bit integers. The ICC encoding convention is illustrated by these tables.

Value (L^*)	uint8 Value	uint16 Value
0.0	0	0
100.0	255	65280
$100.0 + (25500/65280)$	None	65535

Value (a^* or b^*)	uint8 Value	uint16 Value
-128.0	0	0
0.0	128	32768
127.0	255	65280
$127.0 + (255/256)$	None	65535

Class Support `lab` is a `uint8`, `uint16`, or `double` array that must be real and nonsparse. `labd` is `double`.

Example Convert full intensity neutral color (white) from `uint8` to double.

```
lab2double(uint8([255 128 128]))  
ans =
```

100 0 0

See Also

`applycform`, `lab2uint8`, `lab2uint16`, `makecform`, `whitepoint`, `xyz2double`,
`xyz2uint16`

lab2uint16

Purpose Convert $L^*a^*b^*$ data to uint16

Syntax lab16 = lab2uint16(lab)

Description lab16 = lab2uint16(lab) converts an M-by-3 or M-by-N-by-3 array of $L^*a^*b^*$ color values to uint16. lab16 has the same size as lab.

The Image Processing Toolbox follows the convention that double-precision $L^*a^*b^*$ arrays contain 1976 CIE $L^*a^*b^*$ values. $L^*a^*b^*$ arrays that are uint8 or uint16 follow the convention in the ICC profile specification (ICC.1:2001-4, www.color.org) for representing $L^*a^*b^*$ values as unsigned 8-bit or 16-bit integers. The ICC encoding convention is illustrated by these tables.

Value (L^*)	uint8 Value	uint16 Value
0.0	0	0
100.0	255	65280
100.0 + (25500/65280)	None	65535

Value (a^* or b^*)	uint8 Value	uint16 Value
-128.0	0	0
0.0	128	32768
127.0	255	65280
127.0 + (255/256)	None	65535

Class Support lab can be a uint8, uint16, or double array that must be real and nonsparse. lab16 is of class uint16.

Example Convert full intensity neutral color (white) from double to uint16.

```
lab2uint16(100 0 0)
ans =

    65280 32768 32768
```

See Also

`applycform`, `lab2double`, `lab2uint8`, `makecform`, `whitepoint`, `xyz2double`,
`xyz2uint16`

lab2uint8

Purpose Convert $L^*a^*b^*$ data to uint8

Syntax `lab8 = lab2uint8(lab)`

Description `lab8 = lab2uint8(lab)` converts an M-by-3 or M-by-N-by-3 array of $L^*a^*b^*$ color values to uint8. `lab8` has the same size as `lab`.

The Image Processing Toolbox follows the convention that double-precision $L^*a^*b^*$ arrays contain 1976 CIE $L^*a^*b^*$ values. $L^*a^*b^*$ arrays that are uint8 or uint16 follow the convention in the ICC profile specification (ICC.1:2001-4, www.color.org) for representing $L^*a^*b^*$ values as unsigned 8-bit or 16-bit integers. The ICC encoding convention is illustrated by these tables.

Value (L^*)	uint8 Value	uint16 Value
0.0	0	0
100.0	255	65280
$100.0 + (25500/65280)$	None	65535

Value (a^* or b^*)	uint8 Value	uint16 Value
-128.0	0	0
0.0	128	32768
127.0	255	65280
$127.0 + (255/256)$	None	65535

Class Support `lab` is a uint8, uint16, or double array that must be real and nonsparse. `lab8` is uint8.

Example Convert full intensity neutral color (white) from double to uint8.

```
lab2uint8(100 0 0)
ans =

    255    128    128
```

See Also

`applycform`, `lab2double`, `lab2uint16`, `makecform`, `whitepoint`, `xyz2double`,
`xyz2uint16`

label2rgb

Purpose Convert a label matrix into an RGB image

Syntax

```
RGB = label2rgb(L)
RGB = label2rgb(L,map)
RGB = label2rgb(L,map,zerocolor)
RGB = label2rgb(L,map,zerocolor,order)
```

Description `RGB = label2rgb(L)` converts a label matrix `L`, such as those returned by `bwlabel` or `watershed`, into an RGB color image for the purpose of visualizing the labeled regions. The `label2rgb` function determines the color to assign to each object based on the number of objects in the label matrix and range of colors in the colormap. The `label2rgb` function picks colors from the entire range.

`RGB = label2rgb(L,map)` defines the colormap `map` to be used in the RGB image. `map` can have any of the following values:

- *n*-by-3 colormap matrix
- String containing the name of a MATLAB colormap function, such as 'jet' or 'gray' (See `colormap` for a list of supported colormaps.)
- Function handle of a colormap function, such as `@jet` or `@gray`

If you do not specify `map`, the default value is 'jet'.

`RGB = label2rgb(L,map,zerocolor)` defines the RGB color of the elements labeled 0 (zero) in the input label matrix `L`. As the value of `zerocolor`, specify an RGB triple or one of the strings listed in this table.

Value	Color
'b'	Blue
'c'	Cyan
'g'	Green
'k'	Black
'm'	Magenta
'r'	Red

Value	Color
'w'	White
'y'	Yellow

If you do not specify `zerocolor`, the default value for zero-labeled elements is `[1 1 1]` (white).


`RGB = label2rgb(L, map, zerocolor, order)` controls how `label2rgb` assigns colormap colors to regions in the label matrix. If `order` is `'noshuffle'` (the default), `label2rgb` assigns colormap colors to label matrix regions in numerical order. If `order` is `'shuffle'`, `label2rgb` assigns colormap colors pseudorandomly.

Class Support

The input label matrix `L` can have any nonsparse, numeric class. It must contain finite, nonnegative integers. The output of `label2rgb` is of class `uint8`.

Example

```
I = imread('rice.png');
figure, imshow(I), title('original image')
BW = im2bw(I, graythresh(I));
L = bwlabel(BW);
RGB = label2rgb(L);
RGB2 = label2rgb(L, 'spring', 'c', 'shuffle');
imshow(RGB), figure, imshow(RGB2)
```



See Also

`bwlabel`, `colormap`, `ismember`, `watershed`

makecform

Purpose Create a color transformation structure

Syntax

```
C = makecform(type)
C = makecform(type, 'whitepoint', WP)
C = makecform('icc', src_profile, dest_profile)
C = makecform('icc', src_profile, dest_profile,
    'SourceRenderingIntent', src_intent, 'DestRenderingIntent',
    dest_intent)
C = makecform('clut', profile, LUTtype)
C = makecform('mattrc', MatTrc, 'Direction', direction)
```

Description

`C = makecform(type)` creates the color transformation structure `C` that defines the color space conversion specified by `type`. To perform the transformation, pass the color transformation structure as an argument to the `applycform` function.

The `type` argument specifies one of the conversions listed in the following table. `makecform` supports conversions between members of the family of device-independent color spaces defined by the CIE, *Commission Internationale de l'Éclairage* (International Commission on Illumination). In addition, `makecform` supports conversions to and from the *sRGB* standard. For a list of the abbreviations used by the Image Processing Toolbox for each color space, see the Remarks section of this reference page.

Type	Description
'lab2lch'	Convert from $L^*a^*b^*$ to the L^*ch color space.
'lab2srgb' ¹	Convert from $L^*a^*b^*$ to the <i>srgb</i> color space.
'lab2xyz' ¹	Convert from $L^*a^*b^*$ to the <i>XYZ</i> color space.
'lch2lab'	Convert from L^*ch to the $L^*a^*b^*$ color space.
'srgb2lab' ¹	Convert from <i>srgb</i> to the $L^*a^*b^*$ color space.
'srgb2xyz'	Convert from <i>srgb</i> to the <i>XYZ</i> color space.
'upvp12xyz'	Convert from $u'v'L$ to the <i>XYZ</i> color space.
'uv12xyz'	Convert from uvL to the <i>XYZ</i> color space.

Type	Description
'xyz2xyy'	Convert from xyY to the XYZ color space.
'xyz2lab' ¹	Convert from XYZ to the $L^*a^*b^*$ color space.
'xyz2srgb'	Convert from XYZ to the $srgb$ color space.
'xyz2upvpl'	Convert from XYZ to the u^*v^*L color space.
'xyz2uvl'	Convert from XYZ to the uvL color space.
'xyz2xy1'	Convert from XYZ to the xyY color space.

¹For the 'xyz2lab', 'lab2xyz', 'srgb2lab', and 'lab2srgb' transforms, you can optionally specify the value of the reference illuminant, known as the white point. Use the syntax

```
C = makecform(type, 'WhitePoint', WP)
```

where WP is a 1-by-3 vector of XYZ values scaled so that Y = 1. The default is the CIE illuminant D50 as specified in the International Color Consortium specification ICC.1:2001-04. You can use the `whitepoint` function to create the WP vector.

`C = makecform('icc', src_profile, dest_profile)` creates a color transform based on two ICC profiles. `src_profile` and `dest_profile` are ICC profile structures returned by `iccread`.

`C = makecform('icc', src_profile, dest_profile, 'SourceRenderingIntent', src_intent, 'DestRenderingIntent', DEST_INTENT)` creates a color transform based on two ICC color profiles, `src_profile` and `dest_profile`, specifying rendering intent arguments for the source, `src_intent`, and the destination, `dest_intent`, profiles.

Rendering intents specify the style of reproduction that should be used when these profiles are combined. For most devices, the range of reproducible colors is much smaller than the range of colors represented by the PCS. Rendering intents define gamut mapping techniques. Possible values for these rendering

makecform

intents are listed below. Each rendering intent has distinct aesthetic and color-accuracy tradeoffs.

Value	Description
'AbsoluteColorimetric'	Maps all out-of-gamut colors to the nearest gamut surface while maintaining the relationship of all in-gamut colors. This absolute rendering contains color data that is relative to a perfectly reflecting diffuser.
'Perceptual' (default)	Employs vendor-specific gamut mapping techniques for optimizing the range of producible colors of a given device. The objective is to provide the most aesthetically pleasing result even though the relationship of the in-gamut colors might not be maintained. This media-relative rendering contains color data that is relative to the device's white point.
'RelativeColorimetric'	Maps all out-of-gamut colors to the nearest gamut surface while maintaining the relationship of all in-gamut colors. This media-relative rendering contains color data that is relative to the device's white point.
'Saturation'	Employs vendor-specific gamut mapping techniques for maximizing the saturation of device colors. This rendering is generally used for simple business graphics such as bar graphs and pie charts. This media-relative rendering contains color data that is relative to the device's white point.

`C = makecform('clut', profile, LUTtype)` creates the color transformation structure `C` based on a color lookup table (CLUT) contained in an ICC color profile. `profile` is an ICC profile structure returned by `iccread`. `LUTtype`

specifies which clut in the profile structure is to be used. It can be one of these strings.

LUT Type	Description
'AToB0'	Contains the components of a 16- or 8-bit LUTtag that transforms device colors to PCS colors using the perceptual rendering.
'AToB1'	Contains the components of a 16- or 8-bit LUTtag that transforms device colors to PCS colors using the relative rendering.
'AToB2'	Contains the components of a 16- or 8-bit LUTtag that transforms device colors to PCS colors using the saturation rendering.
'BToA0'	Contains the components of a 16- or 8-bit LUTtag that transforms PCS colors to device colors using the perceptual rendering.
'BToA1'	Contains the components of a 16- or 8-bit LUTtag that transforms PCS colors to device colors using the colorimetric rendering.
'BToA2'	Contains the components of a 16- or 8-bit LUTtag that transforms PCS colors to device colors using the saturation rendering.
'Gamut'	Contains the components of a 16- or 8-bit LUTtag that determines which PCS colors are out of gamut for a given device.
'Preview0'	Contains the components of a 16- or 8-bit Preview LUTtag that transforms PCS colors to the PCS colors available for soft proofing using the perceptual rendering.

makecform

LUT Type	Description
'Preview1'	Contains the components of a 16- or 8-bit Preview LUTtag that transforms PCS colors to the PCS colors available for soft proofing using the relative colorimetric rendering.
'Preview2'	Contains the components of a 16- or 8-bit Preview LUTtag that transforms PCS colors to the PCS colors available for soft proofing using the saturation rendering.

`C = makecform('mattrc', MatTrc, 'Direction', direction)` creates the color transformation structure `C` based on a Matrix/Tone Reproduction Curve (MatTRC) model, contained in an ICC color profile. `direction` can be either 'forward' or 'inverse' and specifies whether the MatTRC is to be applied in the forward or inverse direction. For more information, see section 6.3.1.2 of the International Color Consortium specification ICC.1:2001-04 (www.color.org).

Remarks

The Image Processing Toolbox uses the following abbreviations to represent color spaces.

Abbreviation	Description
xyz	1931 CIE XYZ tristimulus values (2° observer)
xy1	1931 CIE xyY chromaticity values (2° observer)
uv1	1960 CIE uvL values
upvp1	1976 CIE the $u'v'L$ values
lab	1976 CIE $L^*a^*b^*$ values
lch	Polar transformation of CIE $L^*a^*b^*$ values, where c = chroma and h = hue
srgb	Standard computer monitor RGB values, (IEC 61966-2-1)

Example

Convert RGB image to $L^*a^*b^*$, assuming input image is uint8.

```
rgb = imread('peppers.png');  
cform = makecform('srgb2lab');  
lab = applycform(rgb,cform);
```

See Also

[applycform](#), [iccread](#), [iccwrite](#), [isicc](#), [lab2double](#), [lab2uint16](#), [lab2uint8](#), [whitepoint](#), [xyz2double](#), [xyz2uint16](#)

makelut

Purpose Construct a lookup table for use with `applylut`

Syntax `lut = makelut(fun,n)`

Description `lut = makelut(fun,n)` returns a lookup table for use with `applylut`. `fun` is a function that accepts an n -by- n matrix of 1's and 0's as input and return a scalar. `n` can be either 2 or 3. `makelut` creates `lut` by passing all possible 2-by-2 or 3-by-3 neighborhoods to `fun`, one at a time, and constructing either a 16-element vector (for 2-by-2 neighborhoods) or a 512-element vector (for 3-by-3 neighborhoods). The vector consists of the output from `fun` for each possible neighborhood.

`fun` must be a function handle.

Class Support `lut` is returned as a vector of class `double`.

Example In this example, the function returns 1 (true) if the number of 1's in the neighborhood is 2 or greater, and returns 0 (false) otherwise. `makelut` then uses the function to construct a lookup table for 2-by-2 neighborhoods.

```
f = @(x) (sum(x(:)) >= 2);  
lut = makelut(f,2)
```

```
lut =  
  
    0  
    0  
    0  
    1  
    0  
    1  
    1  
    1  
    1  
    0  
    1  
    1  
    1  
    1  
    1  
    1  
    1
```

1

See Also

applylut

makeresampler

Purpose Create resampling structure

Syntax `R = makeresampler(interpolant, padmethod)`

Description `R = makeresampler(interpolant, padmethod)` creates a separable resampler structure for use with `tformarray` and `imtransform`.

The `interpolant` argument specifies the interpolating kernel that the separable resampler uses. In its simplest form, `interpolant` can have any of the following strings as a value.

Interpolant	Description
'cubic'	Cubic interpolation
'linear'	Linear interpolation
'nearest'	Nearest-neighbor interpolation

If you are using a custom interpolating kernel, you can specify `interpolant` as a cell array in either of these forms:

`{half_width, positive_half}` `half_width` is a positive scalar designating the half width of a symmetric interpolating kernel. `positive_half` is a vector of values regularly sampling the kernel on the closed interval `[0 positive_half]`.

`{half_width, interp_fcn}` `interp_fcn` is a function handle that returns interpolating kernel values, given an array of input values in the interval `[0 positive_half]`.

To specify the interpolation method independently along each dimension, you can combine both types of `interpolant` specifications. The number of elements in the cell array must equal the number of transform dimensions. For example, if you specify this value for `interpolant`

```
{'nearest', 'linear', {2 KERNEL_TABLE}}
```

the resampler uses nearest-neighbor interpolation along the first transform dimension, linear interpolation along the second dimension, and a custom table-based interpolation along the third.

The `padmethod` argument controls how the resampler interpolates or assigns values to output elements that map close to or outside the edge of the input array. The following table lists all the possible values of `padmethod`.

Pad Method	Description
'bound'	Assigns values from the fill value array to points that map outside the array and repeats border elements of the array for points that map inside the array (same as 'replicate'). When <code>interpolant</code> is 'nearest', this pad method produces the same results as 'fill'. 'bound' is like 'fill', but avoids mixing fill values and input image values.
'circular'	Pads array with circular repetition of elements within the dimension. Same as <code>padarray</code> .
'fill'	Generates an output array with smooth-looking edges (except when using nearest-neighbor interpolation). For output points that map near the edge of the input array (either inside or outside), it combines input image and fill values. When <code>interpolant</code> is 'nearest', this pad method produces the same results as 'bound'.
'replicate'	Pads array by repeating border elements of array. Same as <code>padarray</code> .
'symmetric'	Pads array with mirror reflections of itself. Same as <code>padarray</code> .

In the case of 'fill', 'replicate', 'circular', or 'symmetric', the resampling performed by `tformarray` or `imtransform` occurs in two logical steps:

- 1 Pad the array `A` infinitely to fill the entire input transform space.

makeresampler

- 2 Evaluate the convolution of the padded A with the resampling kernel at the output points specified by the geometric map.

Each nontransform dimension is handled separately. The padding is virtual, (accomplished by remapping array subscripts) for performance and memory efficiency. If you implement a custom resampler, you can implement these behaviors.

Custom Resamplers

The syntaxes described above construct a resampler structure that uses the separable resampler function that ships with the Image Processing Toolbox. It is also possible to create a resampler structure that uses a user-written resampler by using this syntax:

```
R = makeresampler(PropertyName,PropertyValue,...)
```

The makeresampler function supports the following properties.

Property	Description
'Type'	Can have the value 'separable' or 'custom' and must always be supplied. If 'Type' is 'separable', the only other properties that can be specified are 'Interpolant' and 'PadMethod', and the result is equivalent to using the makeresampler(interpolant,padmethod) syntax. If 'Type' is 'custom', you must specify the 'NDims' and 'ResampleFcn' properties and, optionally, the 'CustomData' property.
'PadMethod'	See the padmethod argument for more information.
'Interpolant'	See the interpolant argument for more information.
'NDims'	Positive integer indicating the dimensionality the custom resampler can handle. Use a value of Inf to indicate that the custom resampler can handle any dimension. If 'Type' is 'custom', NDims is required.

Property	Description
'ResampleFcn'	<p>Handle to a function that performs the resampling. The function is called with the following interface.</p> <pre>B = resample_fcn(A,M,TDIMS_A,TDIMS_B,FSIZE_A,FSIZE_B,F,R)</pre> <p>See the help for <code>tformarray</code> for information about the inputs <code>A</code>, <code>TDIMS_A</code>, <code>TDIMS_B</code>, and <code>F</code>. The argument <code>M</code> is an array that maps the transform subscript space of <code>B</code> to the transform subscript space of <code>A</code>. If <code>A</code> has <code>N</code> transform dimensions ($N = \text{length}(\text{TDIMS_A})$) and <code>B</code> has <code>P</code> transform dimensions ($P = \text{length}(\text{TDIMS_B})$), then $\text{ndims}(M) = P + 1$, if $N > 1$ and P if $N == 1$, and $\text{size}(M, P + 1) = N$.</p> <p>The first <code>P</code> dimensions of <code>M</code> correspond to the output transform space, permuted according to the order in which the output transform dimensions are listed in <code>TDIMS_B</code>. (In general <code>TDIMS_A</code> and <code>TDIMS_B</code> need not be sorted in ascending order, although such a limitation might be imposed by specific resamplers.) Thus, the first <code>P</code> elements of <code>size(M)</code> determine the sizes of the transform dimensions of <code>B</code>. The input transform coordinates to which each point is mapped are arrayed across the final dimension of <code>M</code>, following the order given in <code>TDIMS_A</code>. <code>M</code> must be double. <code>FSIZE_A</code> and <code>FSIZE_B</code> are the full sizes of <code>A</code> and <code>B</code>, padded with 1's as necessary to be consistent with <code>TDIMS_A</code>, <code>TDIMS_B</code>, and <code>size(A)</code>.</p>
'CustomData'	User-defined.

Example

Stretch an image in the *y*-direction using a separable resampler that applies cubic interpolation in the *y*-direction and nearest-neighbor interpolation in the *x*-direction. (This is equivalent to, but faster than, applying bicubic interpolation.)

```
A = imread('moon.tif');
resamp = makeresampler({'nearest','cubic'},'fill');
stretch = maketform('affine',[1 0; 0 1.3; 0 0]);
B = imtransform(A,stretch,resamp);
```

See Also

`imtransform`, `tformarray`

maketform

Purpose Create geometric transformation structure

Syntax `T = maketform(transformtype,...)`

Description `T = maketform(transformtype,...)` creates a multidimensional spatial transformation structure (called a TFORM struct) that can be used with the `tformfwd`, `tforminv`, `fliptform`, `imtransform`, or `tformarray` functions. `transformtype` can be any of the following spatial transformation types. `maketform` supports a special syntax for each transformation type. See the following sections for information about these syntaxes.

Transform Type	Description
'affine'	Affine transformation in 2-D or N-D
'projective'	Projective transformation in 2-D or N-D
'custom'	User-defined transformation that can be N-D to M-D
'box'	Independent affine transformation (scale and shift) in each dimension
'composite'	Composition of an arbitrary number of more basic transformations

Transform Types

Affine

`T = maketform('affine',A)` builds a TFORM struct `T` for an `N`-dimensional affine transformation. `A` is a nonsingular real $(N+1)$ -by- $(N+1)$ or $(N+1)$ -by- N matrix. If `A` is $(N+1)$ -by- $(N+1)$, the last column of `A` must be `[zeros(N,1);1]`. Otherwise, `A` is augmented automatically, such that its last column is `[zeros(N,1);1]`. The matrix `A` defines a forward transformation such that `tformfwd(U,T)`, where `U` is a 1-by-`N` vector, returns a 1-by-`N` vector `X`, such that $X = U * A(1:N,1:N) + A(N+1,1:N)$. `T` has both forward and inverse transformations.

`T = maketform('affine',U,X)` builds a TFORM struct `T` for a two-dimensional affine transformation that maps each row of `U` to the corresponding row of `X`.

The `U` and `X` arguments are each 3-by-2 and define the corners of input and output triangles. The corners cannot be collinear.

Projective

`T = maketform('projective',A)` builds a TFORM struct for an `N`-dimensional projective transformation. `A` is a nonsingular real $(N+1)$ -by- $(N+1)$ matrix. $A(N+1,N+1)$ cannot be 0. The matrix `A` defines a forward transformation such that `tformfwd(U,T)`, where `U` is a 1-by-`N` vector, returns a 1-by-`N` vector `X`, such that $X = W(1:N)/W(N+1)$, where $W = [U \ 1] * A$. The transformation structure `T` has both forward and inverse transformations.

`T = maketform('projective',U,X)` builds a TFORM struct `T` for a two-dimensional projective transformation that maps each row of `U` to the corresponding row of `X`. The `U` and `X` arguments are each 4-by-2 and define the corners of input and output quadrilaterals. No three corners can be collinear.

Custom

`T = maketform('custom',NDIMS_IN,NDIMS_OUT,...
FORWARD_FCN, INVERSE_FCN, TDATA)` builds a custom TFORM struct `T` based on user-provided function handles and parameters. `NDIMS_IN` and `NDIMS_OUT` are the numbers of input and output dimensions. `FORWARD_FCN` and `INVERSE_FCN` are function handles to forward and inverse functions. Those functions must support the following syntaxes:

Forward function: `X = FORWARD_FCN(U,T)`

Inverse function: `U = INVERSE_FCN(X,T)`

where `U` is a `P`-by-`NDIMS_IN` matrix whose rows are points in the transformation's input space, and `X` is a `P`-by-`NDIMS_OUT` matrix whose rows are points in the transformation's output space. The `TDATA` argument can be any MATLAB array and is typically used to store parameters of the custom transformation. It is accessible to `FORWARD_FCN` and `INVERSE_FCN` via the `tdata` field of `T`. Either `FORWARD_FCN` or `INVERSE_FCN` can be empty, although at least `INVERSE_FCN` must be defined to use `T` with `tformarray` or `imtransform`.

Box

`T = maketform('box',tsize,LOW,HIGH)` or

`T = maketform('box',INBOUNDS, OUTBOUNDS)` builds an `N`-dimensional affine

TFORM struct **T**. The **tsize** argument is an **N**-element vector of positive integers. **LOW** and **HIGH** are also **N**-element vectors. The transformation maps an input box defined by the opposite corners **ones(1,N)** and **tsize** or, alternatively, by corners **INBOUNDS(1,:)** and **INBOUND(2,:)** to an output box defined by the opposite corners **LOW** and **HIGH** or **OUTBOUNDS(1,:)** and **OUTBOUNDS(2,:)**. **LOW(K)** and **HIGH(K)** must be different unless **tsize(K)** is 1, in which case the affine scale factor along the **K**th dimension is assumed to be 1.0. Similarly, **INBOUNDS(1,K)** and **INBOUNDS(2,K)** must be different unless **OUTBOUNDS(1,K)** and **OUTBOUNDS(2,K)** are the same, and vice versa. The 'box' TFORM is typically used to register the row and column subscripts of an image or array to some world coordinate system.

Composite

T = maketform('composite',**T1**,**T2**,...,**TL**) or
T = maketform('composite', [**T1 T2 ... TL**]) builds a TFORM struct **T** whose forward and inverse functions are the functional compositions of the forward and inverse functions of **T1**, **T2**, ..., **TL**.

For example, if **L** = 3, then **tformfwd(U,T)** is the same as **tformfwd(tformfwd(tformfwd(U,T3),T2),T1)**. The components **T1** through **TL** must be compatible in terms of the numbers of input and output dimensions. **T** has a defined forward transform function only if all the component transforms have defined forward transform functions. **T** has a defined inverse transform function only if all the component functions have defined inverse transform functions.

Example

Make and apply an affine transformation.

```
T = maketform('affine',[.5 0 0; .5 2 0; 0 0 1]);  
tformfwd([10 20],T)  
I = imread('cameraman.tif');  
I2 = imtransform(I,T);  
imshow(I), figure, imshow(I2)
```

See Also

tformfwd, tforminv, fliptform, imtransform, tformarray

Purpose Convert a matrix to a grayscale intensity image

Syntax
`I = mat2gray(A,[amin amax])`
`I = mat2gray(A)`


Description `I = mat2gray(A,[amin amax])` converts the matrix `A` to the intensity image `I`. The returned matrix `I` contains values in the range 0.0 (black) to 1.0 (full intensity or white). `amin` and `amax` are the values in `A` that correspond to 0.0 and 1.0 in `I`.

`I = mat2gray(A)` sets the values of `amin` and `amax` to the minimum and maximum values in `A`.

Class Support The input array `A` can be logical or numeric. The output image `I` is double.

Example

```
I = imread('rice.png');  
J = filter2(fspecial('sobel'),I);  
K = mat2gray(J);  
imshow(I), figure, imshow(K)
```



See Also `gray2ind`

mean2

Purpose Compute the mean of the elements of a matrix

Syntax `B = mean2(A)`

Description `B = mean2(A)` computes the mean of the values in A.

Class Support The input image A can be numeric or logical. The output image B is a scalar of class double.

Algorithm `mean2` computes the mean of an array A using `mean(A(:))`.

See Also `std2`

`mean`, `std` in the MATLAB Function Reference

Purpose	Perform two-dimensional median filtering
Syntax	<pre>B = medfilt2(A, [m n]) B = medfilt2(A) B = medfilt2(A, 'indexed', ...)</pre>
Description	<p>Median filtering is a nonlinear operation often used in image processing to reduce “salt and pepper” noise. Median filtering is more effective than convolution when the goal is to simultaneously reduce noise and preserve edges.</p> <p><code>B = medfilt2(A, [m n])</code> performs median filtering of the matrix <code>A</code> in two dimensions. Each output pixel contains the median value in the <code>m</code>-by-<code>n</code> neighborhood around the corresponding pixel in the input image. <code>medfilt2</code> pads the image with 0's on the edges, so the median values for the points within $[m\ n]/2$ of the edges might appear distorted.</p> <p><code>B = medfilt2(A)</code> performs median filtering of the matrix <code>A</code> using the default 3-by-3 neighborhood.</p> <p><code>B = medfilt2(A, 'indexed', ...)</code> processes <code>A</code> as an indexed image, padding with 0's if the class of <code>A</code> is <code>uint8</code>, or 1's if the class of <code>A</code> is <code>double</code>.</p>
Class Support	The input image <code>A</code> can be of class <code>logical</code> , <code>uint8</code> , <code>uint16</code> , or <code>double</code> (unless the 'indexed' syntax is used, in which case <code>A</code> cannot be of class <code>uint16</code>). The output image <code>B</code> is of the same class as <code>A</code> .

Note For information about performance considerations, see `ordfilt2`.

Remarks	<p>If the input image <code>A</code> is of an integer class, all the output values are returned as integers. If the number of pixels in the neighborhood (i.e., $m*n$) is even, some of the median values might not be integers. In these cases, the fractional parts are discarded. Logical input is treated similarly.</p> <p>For example, suppose you call <code>medfilt2</code> using 2-by-2 neighborhoods, and the input image is a <code>uint8</code> array that includes this neighborhood.</p>
----------------	---

medfilt2

4 8

`medfilt2` returns an output value of 4 for this neighborhood, although the true median is 4.5.

Example

This example adds salt and pepper noise to an image, then restores the image using `medfilt2`.

```
I = imread('eight.tif');  
J = imnoise(I, 'salt & pepper', 0.02);  
K = medfilt2(J);  
imshow(J), figure, imshow(K)
```



Algorithm

`medfilt2` uses `ordfilt2` to perform the filtering.

See Also


`filter2`, `ordfilt2`, `wiener2`

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 469-476.

Purpose	Display multiple image frames as a rectangular montage
Syntax	<pre>montage(I) montage(BW) montage(X,map) montage(RGB) h = montage(...)</pre>
Description	<p>montage displays all the frames of a multiframe image array in a single image object, arranging the frames so that they roughly form a square.</p> <p>montage(I) displays the k frames of the intensity image array I. I is m-by-n-by-1-by-k.</p> <p>montage(BW) displays the k frames of the binary image array BW. BW is m-by-n-by-1-by-k.</p> <p>montage(X,map) displays the k frames of the indexed image array X, using the colormap map for all frames. X is m-by-n-by-1-by-k.</p> <p>montage(RGB) displays the k frames of the truecolor image array RGB. RGB is m-by-n-by-3-by-k.</p> <p>h = montage(...) returns the handle to the image object.</p>
Class Support	<p>An intensity image can be logical, uint8, uint16, int16, single, or double. An indexed image can be logical, uint8, uint16, single, or double. The map must be double. A truecolor image can be uint8, uint16, single, or double. The output is a handle to the graphics objects produced by this function.</p>
Example	<pre>load mri montage(D,map)</pre>

montage



See Also

[imovie](#)

Purpose	Perform general sliding-neighborhood operations
Syntax	<pre>B = nlfilt(A,[m n],fun) B = nlfilt(A,'indexed',...)</pre>
Description	<p><code>B = nlfilt(A,[m n],fun)</code> applies the function <code>fun</code> to each <code>m</code>-by-<code>n</code> sliding block of <code>A</code>. <code>fun</code> is a function that accepts an <code>m</code>-by-<code>n</code> matrix as input and returns a scalar result.</p> <pre>c = fun(x)</pre> <p><code>fun</code> must be a function handle.</p> <p><code>c</code> is the output value for the center pixel in the <code>m</code>-by-<code>n</code> block <code>x</code>. <code>nlfilt</code> calls <code>fun</code> for each pixel in <code>A</code>. <code>nlfilt</code> zero-pads the <code>m</code>-by-<code>n</code> block at the edges, if necessary.</p> <p><code>B = nlfilt(A,'indexed',...)</code> processes <code>A</code> as an indexed image, padding with 1's if <code>A</code> is of class <code>double</code> and 0's if <code>A</code> is of class <code>uint8</code>.</p>
Class Support	The input image <code>A</code> can be of any class supported by <code>fun</code> . The class of <code>B</code> depends on the class of the output from <code>fun</code> .
Remarks	<code>nlfilt</code> can take a long time to process large images. In some cases, the <code>colfilt</code> function can perform the same operation much faster.
Example	<p>This example produces the same result as calling <code>medfilt2</code> with a 3-by-3 neighborhood.</p> <pre>A = imread('cameraman.tif'); fun = @(x) median(x(:)); B = nlfilt(A,[3 3],fun); imshow(A), figure, imshow(B)</pre>
See Also	<code>blkproc</code> , <code>colfilt</code> , <code>function_handle</code>

normxcorr2

Purpose Normalized two-dimensional cross-correlation

Syntax `C = normxcorr2(TEMPLATE,A)`

Description `C = normxcorr2(TEMPLATE,A)` computes the normalized cross-correlation of the matrices `TEMPLATE` and `A`. The matrix `A` must be larger than the matrix `TEMPLATE` for the normalization to be meaningful. The values of `TEMPLATE` cannot all be the same. The resulting matrix `C` contains the correlation coefficients, which can range in value from -1.0 to 1.0.

Class Support The input matrices can be of class `uint8`, `uint16`, or `double`.

Algorithm `normxcorr2` uses the following general procedure:

- 1 Calculate cross-correlation in the spatial or the frequency domain, depending on size of images.
- 2 Calculate local sums by precomputing running sums. [1]
- 3 Use local sums to normalize the cross-correlation to get correlation coefficients. [2]

Example

```
T = .2*ones(11); % make light gray plus on dark gray background
T(6,3:9) = .6;
T(3:9,6) = .6;
BW = T>0.5; % make white plus on black background
imshow(BW), title('Binary')
figure, imshow(T), title('Template')

% make new image that offsets template T
T_offset = .2*ones(21);
offset = [3 5]; % shift by 3 rows, 5 columns
T_offset( (1:size(T,1))+offset(1), (1:size(T,2))+offset(2) ) = T;
imshow(T_offset), title('Offset Template')

% cross-correlate BW and T_offset to recover offset
cc = normxcorr2(BW,T_offset);
[max_cc, imax] = max(abs(cc(:)));
[ypeak, xpeak] = ind2sub(size(cc),imax(1));
corr_offset = [ (ypeak-size(T,1)) (xpeak-size(T,2)) ];
isequal(corr_offset,offset) % 1 means offset was recovered
```

See Also

corrcoef

References

[1] Lewis, J. P., "Fast Normalized Cross-Correlation," *Industrial Light & Magic*,
<<http://www.idiom.com/~zilla/Papers/nvisionInterface/nip.html>>

[2] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Volume II, Addison-Wesley, 1992, pp. 316-317.

ntsc2rgb

Purpose Convert NTSC values to RGB color space

Syntax
`rgbmap = ntsc2rgb(yiqmap)`
`RGB = ntsc2rgb(YIQ)`

Description `rgbmap = ntsc2rgb(yiqmap)` converts the *m*-by-3 NTSC (television) color values in `yiqmap` to RGB color space. If `yiqmap` is *m*-by-3 and contains the NTSC luminance (*Y*) and chrominance (*I* and *Q*) color components as columns, then `rgbmap` is an *m*-by-3 matrix that contains the red, green, and blue values equivalent to those colors. Both `rgbmap` and `yiqmap` contain intensities in the range 0 to 1.0. The intensity 0 corresponds to the absence of the component, while the intensity 1.0 corresponds to full saturation of the component.

`RGB = ntsc2rgb(YIQ)` converts the NTSC image `YIQ` to the equivalent truecolor image `RGB`.

`ntsc2rgb` computes the RGB values from the NTSC components using

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1.000 & 0.956 & 0.621 \\ 1.000 & -0.272 & -0.647 \\ 1.000 & -1.106 & 1.703 \end{bmatrix} \begin{bmatrix} Y \\ I \\ Q \end{bmatrix}$$

Class Support The input image or colormap must be of class `double`. The output is of class `double`.

Example Convert RGB image to NTSC and back.

```
RGB = imread('board.tif');  
NTSC = rgb2ntsc(RGB);  
RGB2 = ntsc2rgb(NTSC);
```

See Also `rgb2ntsc`, `rgb2ind`, `ind2rgb`, `ind2gray`

Purpose	Perform two-dimensional order-statistic filtering
Syntax	<pre>B = ordfilt2(A,order,domain) B = ordfilt2(A,order,domain,S) B = ordfilt2(...,padopt)</pre>
Description	<p><code>B = ordfilt2(A,order,domain)</code> replaces each element in <code>A</code> by the orderth element in the sorted set of neighbors specified by the nonzero elements in <code>domain</code>.</p> <p><code>B = ordfilt2(A,order,domain,S)</code>, where <code>S</code> is the same size as <code>domain</code>, uses the values of <code>S</code> corresponding to the nonzero values of <code>domain</code> as additive offsets.</p> <p><code>B = ordfilt2(...,padopt)</code> controls how the matrix boundaries are padded. Set <code>padopt</code> to 'zeros' (the default) or 'symmetric'. If <code>padopt</code> is 'zeros', <code>A</code> is padded with 0's at the boundaries. If <code>padopt</code> is 'symmetric', <code>A</code> is symmetrically extended at the boundaries.</p>
Class Support	The class of <code>A</code> can be logical, uint8, uint16, or double. The class of <code>B</code> is the same as the class of <code>A</code> , unless the additive offset form of <code>ordfilt2</code> is used, in which case the class of <code>B</code> is double.
Remarks	<p><code>domain</code> is equivalent to the structuring element used for binary image operations. It is a matrix containing only 1's and 0's; the 1's define the neighborhood for the filtering operation.</p> <p>For example, <code>B = ordfilt2(A,5,ones(3,3))</code> implements a 3-by-3 median filter; <code>B = ordfilt2(A,1,ones(3,3))</code> implements a 3-by-3 minimum filter; and <code>B = ordfilt2(A,9,ones(3,3))</code> implements a 3-by-3 maximum filter. <code>B = ordfilt2(A,1,[0 1 0; 1 0 1; 0 1 0])</code> replaces each element in <code>A</code> by the minimum of its north, east, south, and west neighbors.</p> <p>The syntax that includes <code>S</code> (the matrix of additive offsets) can be used to implement grayscale morphological operations, including grayscale dilation and erosion.</p> <p>Performance Considerations</p> <p>When working with large domain matrices that do not contain any zero-valued elements, <code>ordfilt2</code> can achieve higher performance if <code>A</code> is in an integer data</p>

ordfilt2

format (uint8, int8, uint16, int16). The gain in speed is larger for uint8 and int8 than for the 16-bit data types. For 8-bit data formats, the domain matrix must contain seven or more rows. For 16-bit data formats, the domain matrix must contain three or more rows and 520 or more elements.

Example

This examples uses a maximum filter with a [5 5] neighborhood. This is equivalent to `imdilate(image, strel('square', 5))`.

```
A = imread('snowflakes.png');
B = ordfilt2(A, 25, true(5));
figure, imshow(A), figure, imshow(B)
```

See Also

`medfilt2`

Reference

- [1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Volume I, Addison-Wesley, 1992.
- [2] Huang, T.S., G.J.Yang, and G.Y.Tang. "A fast two-dimensional median filtering algorithm.", *IEEE transactions on Acoustics, Speech and Signal Processing*, Vol ASSP 27, No. 1, February 1979.

- Purpose** Convert optical transfer function to point-spread function
- Syntax**
PSF = otf2psf(OTF)
PSF = otf2psf(OTF,OUTSIZE)
- Description**
PSF = otf2psf(OTF) computes the inverse Fast Fourier Transform (IFFT) of the optical transfer function (OTF) array and creates a point-spread function (PSF), centered at the origin. By default, the PSF is the same size as the OTF.
PSF = otf2psf(OTF,OUTSIZE) converts the OTF array into a PSF array, where OUTSIZE specifies the size of the output point-spread function. The size of the output array must not exceed the size of the OTF array in any dimension.
To center the PSF at the origin, otf2psf circularly shifts the values of the output array down (or to the right) until the (1,1) element reaches the central position, then it crops the result to match dimensions specified by OUTSIZE.
Note that this function is used in image convolution/deconvolution when the operations involve the FFT.
- Class Support** OTF can be any nonsparse, numeric array. PSF is of class double.
- Example**

```
PSF = fspecial('gaussian',13,1);  
OTF = psf2otf(PSF,[31 31]); % PSF --> OTF  
PSF2 = otf2psf(OTF,size(PSF)); % OTF --> PSF2  
subplot(1,2,1); surf(abs(OTF)); title('|OTF|');  
axis square; axis tight  
subplot(1,2,2); surf(PSF2); title('Corresponding PSF');  
axis square; axis tight
```
- See Also** psf2otf, circshift, padarray

padarray

Purpose Pad an array

Syntax

```
B = padarray(A,padsizel)
B = padarray(A,padsizel,padval)
B = padarray(A,padsizel,padval,direction)
```

Description B = padarray(A,padsizel) pads array A with 0's (zeros). padsizel is a vector of positive integers that specifies both the amount of padding to add and the dimension along which to add it. The value of an element in the vector specifies the amount of padding to add. The order of the element in the vector specifies the dimension along which to add the padding.

For example, a padsizel value of [2 3] means add 2 elements of padding along the first dimension and 3 elements of padding along the second dimension. By default, padarray adds padding before the first element and after the last element along the specified dimension.

B = padarray(A,padsizel,padval) pads array A where padval specifies the value to use as the pad value. padarray uses the value 0 (zero) as the default. padval can be a scalar that specifies the pad value directly or one of the following text strings that specifies the method padarray uses to determine the values of the elements added as padding.

Value	Meaning
'circular'	Pad with circular repetition of elements within the dimension.
'replicate'	Pad by repeating border elements of array.
'symmetric'	Pad array with mirror reflections of itself.

`B = padarray(A,padsize,padval,direction)` pads `A` in the direction specified by the string `direction`. `direction` can be one of the following strings. The default value is enclosed in braces (`{}`).

Value	Meaning
<code>{ 'both' }</code>	Pads before the first element and after the last array element along each dimension. This is the default.
<code>'post'</code>	Pad after the last array element along each dimension.
<code>'pre'</code>	Pad before the first array element along each dimension.

Class Support

When padding with a constant value, `A` can be numeric or logical. When padding using the `'circular'`, `'replicate'`, or `'symmetric'` methods, `A` can be of any class. `B` is of the same class as `A`.

Example

Example 1

Add three elements of padding to the beginning of a vector. The padding elements, indicated by the gray shading, contain mirror copies of the array elements.

```
a = [ 1 2 3 4 ];
b = padarray(a,[0 3],'symmetric','pre')
b ==
    3    2    1    1    2    3    4
```

Example 2

Add three elements of padding to the end of the first dimension of the array and two elements of padding to the end of the second dimension. The example uses the value of the last array element as the padding value.

```
A = [1 2; 3 4];
B = padarray(A,[3 2],'replicate','post')
```

B =

1	2	2	2
3	4	4	4
3	4	4	4
3	4	4	4
3	4	4	4

Example 3

Add three elements of padding to the vertical and horizontal dimensions of a three-dimensional array. Use default values for the pad value and direction.

A = [1 2; 3 4];

B = [5 6; 7 8];

C = cat(3,A,B)

C(:,:,1) =

1	2
3	4

C(:,:,2) =

5	6
7	8

D = padarray(C,[3 3])

D(:,:,1) ==

0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	1	2	0	0	0
0	0	0	3	4	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

```
D(:,:,2) ==
```

0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	5	6	0	0	0
0	0	0	7	8	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

See Also`circshift, imfilter`

para2fan

Purpose

Compute fan-beam projections from parallel-beam tomography data

Syntax

```
F = para2fan(P,D)
I = para2fan(...,param1,val1,param2,val2,...)
[F,fan_positions,fan_rotation_angles] = fan2para(...)
```

Description

`F = para2fan(P,D)` computes the fan-beam data (sinogram) `F` from the parallel-beam data (sinogram) `P`. Each column of `P` contains the parallel-beam sensor samples at one rotation angle. `D` is the distance in pixels from the center of rotation to the center of the sensors.

The sensors are assumed to have a one-pixel spacing. The parallel-beam rotation angles are assumed to be spaced equally to cover $[0,180]$ degrees. The calculated fan-beam rotation angles cover $[0,360)$ with the same spacing as the parallel-beam rotation angles. The calculated fan-beam angles are equally spaced with the spacing set to the smallest angle implied by the sensor spacing.

`I = para2fan(...,param1,val1,param2,val2,...)` specifies parameters that control various aspects of the `para2fan` conversion. Parameter names can be abbreviated, and case does not matter. Default values are enclosed in braces like this: `{default}`. Parameters include

Parameter	Description
'FanCoverage'	String specifying the range through which the beams are rotated. Possible values: {'cycle'} or 'minimal' See <code>ifanbeam</code> for details.
'FanRotationIncrement'	Positive real scalar specifying the rotation angle increment of the fan-beam projections in degrees. If 'FanCoverage' is 'cycle', 'FanRotationIncrement' must be a factor of 360. If 'FanRotationIncrement' is not specified, then it is set to the same spacing as the parallel-beam rotation angles.

Parameter	Description
'FanSensorGeometry'	Text string specifying how sensors are positioned. Possible values: { 'arc' } or 'line' See fanbeam for details.
'FanSensorSpacing'	Positive real scalar specifying the spacing of the fan beams. Interpretation of the value depends on the setting of 'FanSensorGeometry': If 'FanSensorGeometry' is 'arc', the value defines the angular spacing in degrees. Default value is 1. If 'FanSensorGeometry' is 'line', the value defines the linear spacing in pixels. If 'FanSensorSpacing' is not specified, the default is the smallest value implied by 'ParallelSensorSpacing' such that If 'FanSensorGeometry' is 'arc', 'FanSensorSpacing' is $180/\text{PI} \cdot \text{ASIN}(\text{PSPACE}/D)$ where PSPACE is the value of 'ParallelSensorSpacing'. If 'FanSensorGeometry' is 'line', 'FanSensorSpacing' is $D \cdot \text{ASIN}(\text{PSPACE}/D)$
'Interpolation'	Text string specifying the type of interpolation used between the parallel-beam and fan-beam data. 'nearest' — Nearest-neighbor { 'linear' } — Linear 'spline' — Piecewise cubic spline 'pchip' — Piecewise cubic Hermite (PCHIP) 'cubic' — Same as 'pchip'

para2fan

Parameter	Description
'ParallelCoverage'	Text string specifying the range of rotation. 'cycle' — Parallel data covers 360 degrees { 'halfcycle' } — Parallel data covers 180 degrees
'ParallelRotationIncrement'	Positive real scalar specifying the parallel-beam rotation angle increment, measured in degrees. Parallel beam angles are calculated to cover [0,180) degrees with increment PAR_ROT_INC, where PAR_ROT_INC is the value of 'ParallelRotationIncrement'. 180/PAR_ROT_INC must be an integer. If 'ParallelRotationIncrement' is not specified, the increment is assumed to be the same as the increment of the fan-beam rotation angles.
'ParallelSensorSpacing'	Positive real scalar specifying the spacing of the parallel-beam sensors in pixels. The range of sensor locations is implied by the range of fan angles and is given by $[D*\sin(\min(\text{FAN_ANGLES})), D*\sin(\max(\text{FAN_ANGLES}))]$ If 'ParallelSensorSpacing' is not specified, the spacing is assumed to be uniform and is set to the minimum spacing implied by the fan angles and sampled over the range implied by the fan angles.

[F,fan_positions,fan_rotation_angles] = fan2para(...) If 'FanSensorGeometry' is 'arc', fan_positions contains the fan-beam sensor measurement angles. If 'FanSensorGeometry' is 'line', fan_positions contains the fan-beam sensor positions along the line of sensors. fan_rotation_angles contains rotation angles.

Class Support

P and D can be double or single, and must be nonsparse. The other numeric input arguments must be double. The output arguments are double.

Example

Generate parallel-beam projections

```
ph = phantom(128);  
theta = 0:180;
```

```
[P,xp] = radon(ph,theta);  
imshow(theta,xp,P,[],'n'), axis normal  
title('Parallel-Beam Projections')  
xlabel('\theta (degrees)')  
ylabel('x''')  
colormap(hot), colorbar
```

Convert to fan-beam projections

```
[F,Fpos,Fangles] = para2fan(P,100);  
figure, imshow(Fangles,Fpos,F,[],'n'), axis normal  
title('Fan-Beam Projections')  
xlabel('\theta (degrees)')  
ylabel('Sensor Locations (degrees)')  
colormap(hot), colorbar
```

See Also

fan2para, fanbeam, iradon, ifanbeam, phantom, radon

phantom

Purpose Generate a head phantom image

Syntax

```
P = phantom(def, n)
P = phantom(E, n)
[P, E] = phantom(...)
```

Description `P = phantom(def, n)` generates an image of a head phantom that can be used to test the numerical accuracy of `radon` and `iradon` or other two-dimensional reconstruction algorithms. `P` is a grayscale intensity image that consists of one large ellipse (representing the brain) containing several smaller ellipses (representing features in the brain).

`def` is a string that specifies the type of head phantom to generate. Valid values are

- 'Shepp-Logan' — Test image used widely by researchers in tomography
- 'Modified Shepp-Logan' (default) — Variant of the Shepp-Logan phantom in which the contrast is improved for better visual perception

`n` is a scalar that specifies the number of rows and columns in `P`. If you omit the argument, `n` defaults to 256.

`P = phantom(E, n)` generates a user-defined phantom, where each row of the matrix `E` specifies an ellipse in the image. `E` has six columns, with each column containing a different parameter for the ellipses. This table describes the columns of the matrix.

Column	Parameter	Meaning
Column 1	A	Additive intensity value of the ellipse
Column 2	a	Length of the horizontal semiaxis of the ellipse
Column 3	b	Length of the vertical semiaxis of the ellipse
Column 4	x0	x-coordinate of the center of the ellipse

Column	Parameter	Meaning
Column 5	y0	y-coordinate of the center of the ellipse
Column 6	phi	Angle (in degrees) between the horizontal semiaxis of the ellipse and the x-axis of the image

For purposes of generating the phantom, the domains for the x - and y -axes span $[-1,1]$. Columns 2 through 5 must be specified in terms of this range.

$[P, E] = \text{phantom}(\dots)$ returns the matrix E used to generate the phantom.

Class Support

All inputs and all outputs must be of class `double`.

Remarks

For any given pixel in the output image, the pixel's value is equal to the sum of the additive intensity values of all ellipses that the pixel is a part of. If a pixel is not part of any ellipse, its value is 0.

The additive intensity value A for an ellipse can be positive or negative; if it is negative, the ellipse will be darker than the surrounding pixels. Note that, depending on the values of A , some pixels can have values outside the range $[0,1]$.

Example

```
P = phantom('Modified Shepp-Logan',200);
imshow(P)
```

phantom



Reference

[1] Jain, Anil K., *Fundamentals of Digital Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1989, p. 439.

See Also

radon, iradon

Purpose Display information about image pixels

Syntax

```
pixval on
pixval off
pixval
pixval(fig,option)
pixval(ax,option)
pixval(H,option)
```

Description `pixval on` turns on interactive display of information about image pixels in the current figure. `pixval` installs a black bar at the bottom of the figure, which displays the (x,y) coordinates for whatever pixel the cursor is currently over and the color information for that pixel. If the image is binary or intensity, the color information is a single intensity value. If the image is indexed or RGB, the color information is an RGB triplet. The values displayed are the actual data values, regardless of the class of the image array, or whether the data is in normal image range.

If you click the image and hold down the mouse button while you move the cursor, `pixval` also displays the Euclidean distance between the point you clicked and the current cursor location. `pixval` draws a line between these points to indicate the distance being measured. When you release the mouse button, the line and the distance display disappear.

You can move the display bar by clicking it and dragging it to another place in the figure.

`pixval off` turns interactive display off in the current figure. You can also turn off the display by clicking the button on the right side of the display bar.

`pixval` toggles interactive display on or off in the current figure.

`pixval(fig,option)` applies the `pixval` command to the figure specified by `fig`. `option` is a string containing 'on' or 'off'.

`pixval(ax,option)` applies the `pixval` command to the figure that contains the axes `ax`. `option` is a string containing 'on' or 'off'.

`pixval(H,option)` applies the `pixval` command to the figure that contains the image object `H`. `option` is a string containing 'on' or 'off'.

pixval

Example `figure, imshow peppers.png`
 `pixval`

See Also `impixel, improfile`

Purpose Convert region polygon to region mask

Syntax `BW = poly2mask(x,y,m,n)`

Purpose `BW = poly2mask(x,y,m,n)` computes a binary region-of-interest mask `BW` from a region-of-interest polygon represented by the vectors `x` and `y`. The size of `BW` is `m`-by-`n`. Pixels in `BW` that are inside the polygon `(x,y)` are set to 1; pixels outside the polygon are set to 0 (zero). The class of `BW` is `logical`.

`poly2mask` closes the polygon automatically if it isn't already closed.

Example

```
x = [63 186 54 190 63];
y = [60 60 209 204 60];
bw = poly2mask(x,y,256,256);
imshow(bw)
hold on
plot(x,y,'b','LineWidth',2)
hold off
```

Create a mask using random points.


```
x = 256*rand(1,4);
y = 256*rand(1,4);
x(end+1) = x(1);
y(end+1) = y(1);
bw = poly2mask(x,y,256,256);
imshow(bw)
hold on
plot(x,y,'b','LineWidth',2)
hold off
```

See Also `roipoly`

Algorithm When creating a region-of-interest mask, `poly2mask` must determine which pixels are included in the region. This determination can be difficult when pixels on the edge of a region are only partially covered by the border line. The following figure illustrates a triangular region of interest, examining in close-up one of the vertices of the ROI. The figure shows how pixels can be partially covered by the border of a region-of-interest.


Dividing Pixels into a 5-by-5 Subpixel Grid

The following figure shows the pixel that contains the vertex of the ROI shown previously with this 5-by-5 subpixel grid.




Adjusting the Vertices to the Subpixel Grid

poly2mask adjusts each vertex of the polygon so that the vertex lies on the subpixel grid. Note how poly2mask rounds up x and y coordinates to find the nearest grid corner. This creates a second, modified polygon, slightly smaller, in this case, than the original ROI. A portion is shown in the following figure.



Drawing a Path Between the Adjusted Vertices


poly2mask forms a path from each adjusted vertex to the next, following the edges of the subpixel grid. In the following figure, a portion of this modified polygon is shown by the thick dark lines.



Determining Which Pixels Are in the Region

poly2mask uses the following rule to determine which border pixels are inside the polygon: if the pixel's central subpixel is inside the modified polygon, the pixel is inside the region.

In the following figure, the central subpixels of pixels on the ROI border are shaded a dark gray color. Pixels inside the polygon are shaded a lighter gray. Note that the pixel containing the vertex is not part of the ROI because its center pixel is not inside the modified polygon.



psf2otf

Purpose Convert point-spread function to optical transfer function

Syntax OTF = psf2otf(PSF)
OTF = psf2otf(PSF,OUTSIZE)

Description OTF = psf2otf(PSF) computes the fast Fourier transform (FFT) of the point-spread function (PSF) array and creates the optical transfer function array, OTF, that is not influenced by the PSF off-centering. By default, the OTF array is the same size as the PSF array.

OTF = psf2otf(PSF,OUTSIZE) converts the PSF array into an OTF array, where OUTSIZE specifies the size of the OTF array. OUTSIZE cannot be smaller than the PSF array size in any dimension.

To ensure that the OTF is not altered because of PSF off-centering, psf2otf postpads the PSF array (down or to the right) with 0's to match dimensions specified in OUTSIZE, then circularly shifts the values of the PSF array up (or to the left) until the central pixel reaches (1,1) position.

Note that this function is used in image convolution/deconvolution when the operations involve the FFT.

Class Support PSF can be any nonsparse, numeric array. OTF is of class double.

Example

```
PSF = fspecial('gaussian',13,1);
OTF = psf2otf(PSF,[31 31]); % PSF --> OTF
subplot(1,2,1); surf(PSF); title('PSF');
axis square; axis tight
subplot(1,2,2); surf(abs(OTF)); title('Corresponding |OTF|');
axis square; axis tight
```

See Also otf2psf, circshift, padarray

Purpose Perform quadtree decomposition

Syntax

```
S = qtdecomp(I)
S = qtdecomp(I, threshold)
S = qtdecomp(I, threshold, mindim)
S = qtdecomp(I, threshold, [mindim maxdim])
S = qtdecomp(I, fun)
```

Description qtdecomp divides a square image into four equal-sized square blocks, and then tests each block to see if it meets some criterion of homogeneity. If a block meets the criterion, it is not divided any further. If it does not meet the criterion, it is subdivided again into four blocks, and the test criterion is applied to those blocks. This process is repeated iteratively until each block meets the criterion. The result can have blocks of several different sizes.

`S = qtdecomp(I)` performs a quadtree decomposition on the intensity image `I` and returns the quadtree structure in the sparse matrix `S`. If `S(k, m)` is nonzero, then `(k, m)` is the upper left corner of a block in the decomposition, and the size of the block is given by `S(k, m)`. By default, qtdecomp splits a block unless all elements in the block are equal.

`S = qtdecomp(I, threshold)` splits a block if the maximum value of the block elements minus the minimum value of the block elements is greater than `threshold`. `threshold` is specified as a value between 0 and 1, even if `I` is of class `uint8` or `uint16`. If `I` is `uint8`, the `threshold` value you supply is multiplied by 255 to determine the actual threshold to use; if `I` is `uint16`, the `threshold` value you supply is multiplied by 65535.

`S = qtdecomp(I, threshold, mindim)` will not produce blocks smaller than `mindim`, even if the resulting blocks do not meet the threshold condition.

`S = qtdecomp(I, threshold, [mindim maxdim])` will not produce blocks smaller than `mindim` or larger than `maxdim`. Blocks larger than `maxdim` are split even if they meet the threshold condition. `maxdim/mindim` must be a power of 2.

`S = qtdecomp(I, fun)` uses the function `fun` to determine whether to split a block. qtdecomp calls `fun` with all the current blocks of size `m`-by-`m` stacked into an `m`-by-`m`-by-`k` array, where `k` is the number of `m`-by-`m` blocks. `fun` returns a logical `k`-element vector, whose values are 1 if the corresponding block should

be split, and 0 otherwise. (For example, if $k(3)$ is 0, the third m -by- m block should not be split.) `fun` must be a `function_handle`.

Class Support

For the syntaxes that do not include a function, the input image can be of class `logical`, `uint8`, `uint16`, `int16`, `single`, or `double`. For the syntaxes that include a function, the input image can be of any class supported by the function. The output matrix is always of class `sparse`.

Remarks

`qtdecomp` is appropriate primarily for square images whose dimensions are a power of 2, such as 128-by-128 or 512-by-512. These images can be divided until the blocks are as small as 1-by-1. If you use `qtdecomp` with an image whose dimensions are not a power of 2, at some point the blocks cannot be divided further. For example, if an image is 96-by-96, it can be divided into blocks of size 48-by-48, then 24-by-24, 12-by-12, 6-by-6, and finally 3-by-3. No further division beyond 3-by-3 is possible. To process this image, you must set `mindim` to 3 (or to 3 times a power of 2); if you are using the syntax that includes a function, the function must return 0 at the point when the block cannot be divided further.

Example

```
I = uint8([1 1 1 2 3 6 6;...
           1 1 2 1 4 5 6 8;...
           1 1 1 1 7 7 7 7;...
           1 1 1 1 6 6 5 5;...
           20 22 20 22 1 2 3 4;...
           20 22 22 20 5 4 7 8;...
           20 22 20 20 9 12 40 12;...
           20 22 20 20 13 14 15 16]);
```

```
S = qtdecomp(I,.05);
disp(full(S));
```

View the block representation of quadtree decomposition.

```
I = imread('liftingbody.png');
S = qtdecomp(I,.27);
blocks = repmat(uint8(0),size(S));

for dim = [512 256 128 64 32 16 8 4 2 1];
    numblocks = length(find(S==dim));
    if (numblocks > 0)
```

```
values = repmat(uint8(1),[dim dim numblocks]);
values(2:dim,2:dim,:) = 0;
blocks = qtsetblk(blocks,S,dim,values);
end
end

blocks(end,1:end) = 1;
blocks(1:end,end) = 1;

imshow(I), figure, imshow(blocks,[])
```

The following figure shows the original image and a representation of the quadtree decomposition of the image.




Image Courtesy of NASA

See Also

`function_handle`, `qtgetblk`, `qtsetblk`

qtgetblk

Purpose Get block values in quadtree decomposition

Syntax `[vals,r,c] = qtgetblk(I,S,dim)`
`[vals,idx] = qtgetblk(I,S,dim)`

Description `[vals,r,c] = qtgetblk(I,S,dim)` returns in `vals` an array containing the `dim`-by-`dim` blocks in the quadtree decomposition of `I`. `S` is the sparse matrix returned by `qtdecomp`; it contains the quadtree structure. `vals` is a `dim`-by-`dim`-by-`k` array, where `k` is the number of `dim`-by-`dim` blocks in the quadtree decomposition; if there are no blocks of the specified size, all outputs are returned as empty matrices. `r` and `c` are vectors containing the row and column coordinates of the upper left corners of the blocks.

`[vals,idx] = qtgetblk(I,S,dim)` returns in `idx` a vector containing the linear indices of the upper left corners of the blocks.

Class Support `I` can be of class `logical`, `uint8`, `uint16`, `int16`, `single`, or `double`. `S` is of class `sparse`.

Remarks The ordering of the blocks in `vals` matches the columnwise order of the blocks in `I`. For example, if `vals` is 4-by-4-by-2, `vals(:, :, 1)` contains the values from the first 4-by-4 block in `I`, and `vals(:, :, 2)` contains the values from the second 4-by-4 block.

Example

```
I = [1  1  1  1  2  3  6  6
      1  1  2  1  4  5  6  8
      1  1  1  1 10 15  7  7
      1  1  1  1 20 25  7  7
     20 22 20 22  1  2  3  4
     20 22 22 20  5  6  7  8
     20 22 20 20  9 10 11 12
     22 22 20 20 13 14 15 16];
```

```
S = qtdecomp(I,5);
```

```
[vals,r,c] = qtgetblk(I,S,4)
```

See Also `qtdecomp`, `qtsetblk`

Purpose	Set block values in quadtree decomposition
Syntax	<code>J = qtsetblk(I,S,dim,vals)</code>
Description	<code>J = qtsetblk(I,S,dim,vals)</code> replaces each <code>dim</code> -by- <code>dim</code> block in the quadtree decomposition of <code>I</code> with the corresponding <code>dim</code> -by- <code>dim</code> block in <code>vals</code> . <code>S</code> is the sparse matrix returned by <code>qtdecomp</code> ; it contains the quadtree structure. <code>vals</code> is a <code>dim</code> -by- <code>dim</code> -by- <code>k</code> array, where <code>k</code> is the number of <code>dim</code> -by- <code>dim</code> blocks in the quadtree decomposition.
Class Support	<code>I</code> can be of class <code>logical</code> , <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> . <code>S</code> is of class <code>sparse</code> .
Remarks	The ordering of the blocks in <code>vals</code> must match the columnwise order of the blocks in <code>I</code> . For example, if <code>vals</code> is 4-by-4-by-2, <code>vals(:, :, 1)</code> contains the values used to replace the first 4-by-4 block in <code>I</code> , and <code>vals(:, :, 2)</code> contains the values for the second 4-by-4 block.
Example	<pre> I = [1 1 1 1 2 3 6 6 1 1 2 1 4 5 6 8 1 1 1 1 10 15 7 7 1 1 1 1 20 25 7 7 20 22 20 22 1 2 3 4 20 22 22 20 5 6 7 8 20 22 20 20 9 10 11 12 22 22 20 20 13 14 15 16]; S = qtdecomp(I,5); newvals = cat(3,zeros(4),ones(4)); J = qtsetblk(I,S,4,newvals) </pre>
See Also	<code>qtdecomp</code> , <code>qtgetblk</code>

radon

Purpose Radon transform

Syntax `R = radon(I,theta)`
`[R,xp] = radon(...)`

Description `R = radon(I,theta)` returns the Radon transform `R` of the intensity image `I` for the angle `theta` degrees.

The Radon transform is the projection of the image intensity along a radial line oriented at a specific angle. If `theta` is a scalar, `R` is a column vector containing the Radon transform for `theta` degrees. If `theta` is a vector, `R` is a matrix in which each column is the Radon transform for one of the angles in `theta`. If you omit `theta`, it defaults to `0:179`.

`[R,xp] = radon(...)` returns a vector `xp` containing the radial coordinates corresponding to each row of `R`.

The radial coordinates returned in `xp` are the values along the x' -axis, which is oriented at `theta` degrees counterclockwise from the x -axis. The origin of both axes is the center pixel of the image, which is defined as


$$\text{floor}((\text{size}(I)+1)/2)$$

For example, in a 20-by-30 image, the center pixel is (10,15).

Class Support `I` can be of class `double`, `logical`, or any integer class. All other inputs and outputs are of class `double`.

Example

```
iptsetpref('ImshowAxesVisible','on')
I = zeros(100,100);
I(25:75, 25:75) = 1;
theta = 0:180;
[R,xp] = radon(I,theta);
imshow(R,[],'Xdata',theta,'Ydata',xp,...
        'InitialMagnification','fit')
xlabel('\theta (degrees)')
ylabel('x')
colormap(hot), colorbar
```

**See Also**

fan2para, fanbeam, ifanbeam, iradon, para2fan, phantom


References

Bracewell, Ronald N., *Two-Dimensional Imaging*, Englewood Cliffs, NJ, Prentice Hall, 1995, pp. 505-537.


Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 42-45.

Algorithm

The Radon transform of an image is the sum of the Radon transforms of each individual pixel, the superposition principle.




The algorithm first divides pixels in the image into four parts and projects each subdivision separately.



Each pixel's contribution is proportionally split into the two nearest bins, according to the distance between the projected location and the bin centers. If the projection hits the center point, the bin on the axes gets the full .25 % weight of the pixel. If the projection hits the border between two bins, the bins share half the .25 %. If the projection falls between a midpoint and a bin border, the pixel

radon



Each pixel's contribution is proportionally split into the two nearest bins, according to the distance between the projected location and the bin centers.

Purpose	Local range of an image
Syntax	<pre>J = rangefilt(I) J = rangefilt(I,NHOOD)</pre>
Description	<p><code>J = rangefilt(I)</code> returns the array <code>J</code>, where each output pixel contains the range value (maximum value - minimum value) of the 3-by-3 neighborhood around the corresponding pixel in the input image <code>I</code>. <code>I</code> can have any dimension. The output image <code>J</code> is the same size as the input image <code>I</code>.</p> <p><code>J = rangefilt(I,NHOOD)</code> performs range filtering of the input image <code>I</code> where you specify the neighborhood in <code>NHOOD</code>. <code>NHOOD</code> is a multidimensional array of zeros and ones where the nonzero elements specify the neighborhood for the range filtering operation. <code>NHOOD</code>'s size must be odd in each dimension.</p> <p>By default, <code>rangefilt</code> uses the neighborhood <code>true(3)</code>. <code>rangefilt</code> determines the center element of the neighborhood by <code>floor((size(NHOOD) + 1)/2)</code>. For information about specifying neighborhoods, see Notes.</p>
Class Support	<p><code>I</code> can be logical or numeric and must be real and nonsparse. <code>NHOOD</code> can be logical or numeric and must contain zeros or ones.</p> <p>The output image <code>J</code> is the same class as <code>I</code>, except for signed integer data types. The output class for signed data types is the corresponding unsigned integer data type. For example, if the class of <code>I</code> is <code>int8</code>, then the class of <code>J</code> is <code>uint8</code>.</p>
Notes	<p><code>rangefilt</code> uses the morphological functions <code>imdilate</code> and <code>imerode</code> to determine the maximum and minimum values in the specified neighborhood. Consequently, <code>rangefilt</code> uses the padding behavior of these morphological functions.</p> <p>In addition, to specify neighborhoods of various shapes, such as a disk, use the <code>strel</code> function to create a structuring element object and then use the <code>getnhood</code> function to extract the neighborhood from the structuring element object.</p>
Examples	<p>(2-D) Identify the two flying objects by measuring the local range.</p> <pre>I = imread('liftingbody.png'); J = rangefilt(I);</pre>

rangefilt

```
imshow(I), figure, imshow(J);
```

(3-D) Quantify land cover changes in an RGB image. The example first converts the image to $L^*a^*b^*$ colorspace to separate intensity information into a single plane of the image, and then calculates the local range in each layer.

```
I = imread('autumn.tif');  
cform = makecform('srgb2lab');  
LAB = applycform(I, cform);  
rLAB = rangefilt(LAB);  
  
imshow(I);  
figure, imshow(rLAB(:,:,1),[]);  
figure, imshow(rLAB(:,:,2),[]);  
figure, imshow(rLAB(:,:,3),[]);
```

See Also

entropyfilt, getnhood, imdilate, imerode, stdfilt, strel

Purpose	Reflect structuring element
Syntax	<code>SE2 = reflect(SE)</code>
Description	<code>SE2 = reflect(SE)</code> reflects a structuring element through its center. The effect is the same as if you rotated the structuring element's domain 180 degrees around its center (for a 2-D structuring element). If <code>SE</code> is an array of structuring element objects, then <code>reflect(SE)</code> reflects each element of <code>SE</code> , and <code>SE2</code> has the same size as <code>SE</code> .
Class Support	<code>SE</code> and <code>SE2</code> are STREL objects.
Example	<pre>se = strel([0 0 1; 0 0 0; 0 0 0]) se2 = reflect(se) se = Flat STREL object containing 1 neighbor. Neighborhood: 0 0 1 0 0 0 0 0 0 se2 = Flat STREL object containing 1 neighbor. Neighborhood: 0 0 0 0 0 0 1 0 0</pre>
See Also	<code>strel</code>

regionprops

Purpose Measure properties of image regions

Syntax STATS = regionprops(L,properties)

Description STATS = regionprops(L,properties) measures a set of properties for each labeled region in the label matrix L. Positive integer elements of L correspond to different regions. For example, the set of elements of L equal to 1 corresponds to region 1; the set of elements of L equal to 2 corresponds to region 2; and so on. The return value STATS is a structure array of length $\max(L(:))$. The fields of the structure array denote different measurements for each region, as specified by properties.

properties can be a comma-separated list of strings, a cell array containing strings, the single string 'all', or the string 'basic'. This table lists the set of valid property strings. Property strings are case insensitive and can be abbreviated.

'Area'	'EquivDiameter'	'MajorAxisLength'
'BoundingBox'	'EulerNumber'	'MinorAxisLength'
'Centroid'	'Extent'	'Orientation'
'ConvexArea'	'Extrema'	'PixelIdxList'
'ConvexHull'	'FilledArea'	'PixelList'
'ConvexImage'	'FilledImage'	'Solidity'
'Eccentricity'	'Image'	

If properties is the string 'all', then all the preceding measurements are computed. If properties is not specified or if it is the string 'basic', then these measurements are computed: 'Area', 'Centroid', and 'BoundingBox'.

Definitions 'Area' — Scalar; the actual number of pixels in the region. (This value might differ slightly from the value returned by bwarea, which weights different patterns of pixels differently.)


'BoundingBox' — 1-by-ndims(L)*2 vector; the smallest rectangle containing the region. BoundingBox is [ul_corner width], where

ul_corner is in the form [x y z ...] and specifies the upper left corner of the bounding box

width is in the form [x_width y_width ...] and specifies the width of the bounding box along each dimension

'Centroid' — 1-by-ndims(L) vector; the center of mass of the region. Note that the first element of Centroid is the horizontal coordinate (or x -coordinate) of the center of mass, and the second element is the vertical coordinate (or y -coordinate). All other elements of Centroid are in order of dimension.

This figure illustrates the centroid and bounding box. The region consists of the white pixels; the green box is the bounding box, and the red dot is the centroid.



'MajorAxisLength' — Scalar; the length (in pixels) of the major axis of the ellipse that has the same normalized second central moments as the region. This property is supported only for 2-D input label matrices.

'MinorAxisLength' — Scalar; the length (in pixels) of the minor axis of the ellipse that has the same normalized second central moments as the region. This property is supported only for 2-D input label matrices.

'Eccentricity' — Scalar; the eccentricity of the ellipse that has the same second-moments as the region. The eccentricity is the ratio of the distance between the foci of the ellipse and its major axis length. The value is between 0 and 1. (0 and 1 are degenerate cases; an ellipse whose eccentricity is 0 is actually a circle, while an ellipse whose eccentricity is 1 is a line segment.) This property is supported only for 2-D input label matrices.


'Orientation' — Scalar; the angle (in degrees) between the x -axis and the major axis of the ellipse that has the same second-moments as the region. This property is supported only for 2-D input label matrices.

regionprops

This figure illustrates the axes and orientation of the ellipse. The left side of the figure shows an image region and its corresponding ellipse. The right side shows the same ellipse, with features indicated graphically; the solid blue lines are the axes, the red dots are the foci, and the orientation is the angle between the horizontal dotted line and the major axis.



'Image' — Binary image (logical) of the same size as the bounding box of the region; the on pixels correspond to the region, and all other pixels are off.



Original Image, Containing a Single Region





Image Returned

'FilledImage' — Binary image (logical) of the same size as the bounding box of the region. The on pixels correspond to the region, with all holes filled in.

'FilledArea' — Scalar; the number of on pixels in FilledImage.



Original Image, Containing a Single Region




Image Returned

'ConvexHull' — p-by-2 matrix; the smallest convex polygon that can contain the region. Each row of the matrix contains the x - and y -coordinates of one vertex of the polygon. This property is supported only for 2-D input label matrices.

'ConvexImage' — Binary image (logical); the convex hull, with all pixels within the hull filled in (i.e., set to on). (For pixels that the boundary of the hull passes through, regionprops uses the same logic as `roipoly` to determine whether the pixel is inside or outside the hull.) The image is the size of the bounding box of the region. This property is supported only for 2-D input label matrices.


'ConvexArea' — Scalar; the number of pixels in 'ConvexImage'. This property is supported only for 2-D input label matrices.

'EulerNumber' — Scalar; equal to the number of objects in the region minus the number of holes in those objects. This property is supported only for 2-D input label matrices.

'Extrema' — 8-by-2 matrix; the extrema points in the region. Each row of the matrix contains the x - and y -coordinates of one of the points. The format of the vector is [top-left top-right right-top right-bottom bottom-right bottom-left left-bottom left-top]. This property is supported only for 2-D input label matrices.

regionprops

This figure illustrates the extrema of two different regions. In the region on the left, each extrema point is distinct; in the region on the right, certain extrema points (e.g., top-left and left-top) are identical.



'EquivDiameter' — Scalar; the diameter of a circle with the same area as the region. Computed as $\sqrt{4 * \text{Area} / \pi}$. This property is supported only for 2-D input label matrices.

'Solidity' — Scalar; the proportion of the pixels in the convex hull that are also in the region. Computed as $\text{Area} / \text{ConvexArea}$. This property is supported only for 2-D input label matrices.

'Extent' — Scalar; the proportion of the pixels in the bounding box that are also in the region. Computed as the Area divided by the area of the bounding box. This property is supported only for 2-D input label matrices.

'PixelIdxList' — p-element vector containing the linear indices of the pixels in the region.

'PixelList' — p-by-ndims(L) matrix; the actual pixels in the region. Each row of the matrix has the form $[x \ y \ z \ \dots]$ and specifies the coordinates of one pixel in the region.

Class Support

The input label matrix L can have any numeric class.

Remarks

Using the Comma-Separated List Syntax

The comma-separated list syntax for structure arrays is very useful when you work with the output of regionprops. For example, for a field that contains a scalar, you can use this syntax to create a vector containing the value of this field for each region in the image.

For instance, if `stats` is a structure array with field `Area`, then the following two expressions are equivalent:

```
stats(1).Area, stats(2).Area, ..., stats(end).Area
```

and

```
stats.Area
```

Therefore, you can use these calls to create a vector containing the area of each region in the image.

```
stats = regionprops(L, 'Area');  
allArea = [stats.Area];
```

`allArea` is a vector of the same length as the structure array `stats`.

Selecting Regions Based on Certain Criteria

The function `ismember` is useful in conjunction with `regionprops` for selecting regions based on certain criteria. For example, these commands create a binary image containing only the regions whose area is greater than 80.

```
idx = find([stats.Area] > 80);  
BW2 = ismember(L,idx);
```

Performance Considerations

Most of the measurements take very little time to compute. The exceptions are these, which can take significantly longer, depending on the number of regions in `L`:

- 'ConvexHull'
- 'ConvexImage'
- 'ConvexArea'
- 'FilledImage'

Note that computing certain groups of measurements takes about the same amount of time as computing just one of them because `regionprops` takes advantage of intermediate computations used in both computations. Therefore, it is fastest to compute all the desired measurements in a single call to `regionprops`.

Working with Binary Images

You must convert a binary image to a label matrix before calling `regionprops`.

Two common ways to convert a binary image to a label matrix are by using the `bwlabel` function,

```
L = bwlabel(BW);
```

or using the `double` function,

```
L = double(BW);
```

Note, however, that these functions produce different but equally valid label matrices from the same binary image.

For example, given the following logical matrix, `BW`,

```
1 1 0 0 0 0
1 1 0 0 0 0
0 0 0 0 0 0
0 0 0 0 1 1
0 0 0 0 1 1
```

`bwlabel` creates a label matrix containing two contiguous regions labeled by the integer values 1 and 2.

```
mylabel = bwlabel(BW)
```

```
mylabel =
```

```
1 1 0 0 0 0
1 1 0 0 0 0
0 0 0 0 0 0
0 0 0 0 2 2
0 0 0 0 2 2
```

The `double` function creates a label matrix containing one discontinuous region labeled by the integer value 1.

```
mylabel2 = double(BW)
```

```
mylabel2 =
```

```
1 1 0 0 0 0
```

```

1     1     0     0     0     0
0     0     0     0     0     0
0     0     0     0     1     1
0     0     0     0     1     1

```

Because each result is legitimately desirable in certain situations, `regionprops` does not attempt to perform either type of conversion on binary images and instead requires that you convert them using either method.

Example

```

BW = imread('text.png');
L = bwlabel(BW);
stats = regionprops(L, 'all');
stats(23)
ans =

    Area: 90
    Centroid: [69.7556 89.8667]
    BoundingBox: [64.5000 83.5000 11 13]
    SubarrayIdx: {1x2 cell}
    MajorAxisLength: 14.5814
    MinorAxisLength: 11.8963
    Eccentricity: 0.5783
    Orientation: -89.2740
    ConvexHull: [19x2 double]
    ConvexImage: [13x11 logical]
    ConvexArea: 121
    Image: [13x11 logical]

    FilledImage: [13x11 logical]
    FilledArea: 98
    EulerNumber: 0
    Extrema: [8x2 double]
    EquivDiameter: 10.7047
    Solidity: 0.7438
    Extent: 0.6294
    PixelIdxList: [90x1 double]
    PixelList: [90x2 double]

```

See Also

`bwlabel`, `bwlabeln`, `ismember`, `watershed`
`ismember` (MATLAB function)

rgb2gray

Purpose Convert an RGB image or colormap to grayscale

Syntax `I = rgb2gray(RGB)`
`newmap = rgb2gray(map)`

Description `rgb2gray` converts RGB images to grayscale by eliminating the hue and saturation information while retaining the luminance.

`I = rgb2gray(RGB)` converts the truecolor image `RGB` to the grayscale intensity image `I`.

`newmap = rgb2gray(map)` returns a grayscale colormap equivalent to `map`.

Class Support If the input is an RGB image, it can be of class `uint8`, `uint16`, `single`, or `double`. The output image `I` is of the same class as the input image. If the input is a colormap, the input and output colormaps are both of class `double`.

Example Convert an RGB image to a grayscale image.

```
I = imread('board.tif');  
J = rgb2gray(I);  
figure, imshow(I), figure, imshow(J);
```

Convert the colormap to a grayscale colormap.

```
[X,map] = imread('trees.tif');  
gmap = rgb2gray(map);  
figure, imshow(X,map), figure, imshow(X,gmap);
```

Algorithm `rgb2gray` converts the RGB values to NTSC coordinates, sets the hue and saturation components to zero, and then converts back to RGB color space.

See Also `ind2gray`, `ntsc2rgb`, `rgb2ind`, `rgb2ntsc`

Purpose

Convert RGB values to hue-saturation-value (HSV) color space

rgb2hsv is a function in MATLAB. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference pages.

rgb2ind

Purpose

Convert an RGB image to an indexed image

Syntax

```
[X,map] = rgb2ind(RGB,n)
X = rgb2ind(RGB,map)
[X,map] = rgb2ind(RGB,tol)
[...] = rgb2ind(...,dither_option)
```

Description

rgb2ind converts RGB images to indexed images using one of three different methods:

- Uniform quantization
- Minimum variance quantization
- Colormap approximation

For all these methods, rgb2ind also dithers the image unless you specify 'nodither' for dither_option.

[X,map] = rgb2ind(RGB,n) converts the RGB image to an indexed image X using minimum variance quantization. map contains at most n colors. n must be less than or equal to 65536.

X = rgb2ind(RGB,map) converts the RGB image to an indexed image X with colormap map by matching colors in RGB with the nearest color in the colormap map. size(map,1) must be less than or equal to 65536.

[X,map] = rgb2ind(RGB,tol) converts the RGB image to an indexed image X using uniform quantization. map contains at most $(\text{floor}(1/\text{tol})+1)^3$ colors. tol must be between 0.0 and 1.0.

[...] = rgb2ind(...,dither_option) enables or disables dithering. dither_option is a string that can have one of these values:

- 'dither' (default) dithers, if necessary, to achieve better color resolution at the expense of spatial resolution.
- 'nodither' maps each color in the original image to the closest color in the new map. No dithering is performed.

Note The values in the resultant image `X` are indexes into the colormap `map` and cannot be used in mathematical processing, such as filtering operations.

Class Support

The input image can be of class `uint8`, `uint16`, `single`, or `double`. If the length of `map` is less than or equal to 256, the output image is of class `uint8`. Otherwise, the output image is of class `uint16`.

Remarks

If you specify `tol`, `rgb2ind` uses uniform quantization to convert the image. This method involves cutting the RGB color cube into smaller cubes of length `tol`. For example, if you specify a `tol` of 0.1, the edges of the cubes are one-tenth the length of the RGB cube. The total number of small cubes is

$$n = (\text{floor}(1/\text{tol})+1)^3$$

Each cube represents a single color in the output image. Therefore, the maximum length of the colormap is `n`. `rgb2ind` removes any colors that don't appear in the input image, so the actual colormap can be much smaller than `n`.

If you specify `n`, `rgb2ind` uses minimum variance quantization. This method involves cutting the RGB color cube into smaller boxes (not necessarily cubes) of different sizes, depending on how the colors are distributed in the image. If the input image actually uses fewer colors than the number you specify, the output colormap is also smaller.

If you specify `map`, `rgb2ind` uses colormap mapping, which involves finding the colors in `map` that best match the colors in the RGB image.

Example

```
RGB = imread('peppers.png');  
[X,map] = rgb2ind(RGB,128);  
imshow(X,map)
```

rgb2ind



See Also

`cmunique`, `dither`, `imapprox`, `ind2rgb`, `rgb2gray`

Purpose	Convert RGB values to NTSC color space
Syntax	<pre>YIQmap = rgb2ntsc(rgbmap) YIQ = rgb2ntsc(RGB)</pre>
Description	<p><code>YIQmap = rgb2ntsc(rgbmap)</code> converts the <i>m</i>-by-3 RGB values in <code>rgbmap</code> to NTSC color space. <code>YIQmap</code> is an <i>m</i>-by-3 matrix that contains the NTSC luminance (<i>Y</i>) and chrominance (<i>I</i> and <i>Q</i>) color components as columns that are equivalent to the colors in the RGB colormap.</p> <p><code>YIQ = rgb2ntsc(RGB)</code> converts the truecolor image <code>RGB</code> to the equivalent NTSC image <code>YIQ</code>.</p>
Remarks	<p>In the NTSC color space, the luminance is the grayscale signal used to display pictures on monochrome (black and white) televisions. The other components carry the hue and saturation information.</p> <p><code>rgb2ntsc</code> defines the NTSC components using</p> $\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$
Class Support	RGB can be of class <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> . <code>RGBMAP</code> can be <code>double</code> . The output is <code>double</code> .
See Also	<code>ntsc2rgb</code> , <code>rgb2ind</code> , <code>ind2rgb</code> , <code>ind2gray</code>

rgb2ycbcr

Purpose Convert RGB values to YCbCr color space

Syntax `ycbcrmap = rgb2ycbcr(rgbmap)`
`YCBCR = rgb2ycbcr(RGB)`

Description `ycbcrmap = rgb2ycbcr(rgbmap)` converts the RGB values in `rgbmap` to the YCbCr color space. `ycbcrmap` is an m -by-3 matrix that contains the YCbCr luminance (Y) and chrominance (Cb and Cr) color components as columns. Each row represents the equivalent color to the corresponding row in the RGB colormap.

`YCBCR = rgb2ycbcr(RGB)` converts the truecolor image RGB to the equivalent image in the YCbCr color space.

Class Support If the input is an RGB image, it can be of class `uint8`, `uint16`, or `double`; the output image is of the same class as the input image. If the input is a colormap, the input and output colormaps are both of class `double`.

See Also `ntsc2rgb`, `rgb2ntsc`, `ycbcr2rgb`

Purpose

Plot colormap

rgbplot is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

roicolor

Purpose Select region of interest, based on color

Syntax
`BW = roicolor(A,low,high)`
`BW = roicolor(A,v)`

Description `roicolor` selects a region of interest within an indexed or intensity image and returns a binary image. (You can use the returned image as a mask for masked filtering using `roifilt2`.)

`BW = roicolor(A,low,high)` returns a region of interest selected as those pixels that lie within the colormap range [low high].

$$BW = (A \geq low) \ \& \ (A \leq high)$$


`BW` is a binary image with 0's outside the region of interest and 1's inside.

`BW = roicolor(A,v)` returns a region of interest selected as those pixels in `A` that match the values in vector `v`. `BW` is a binary image with 1's where the values of `A` match the values of `v`.

Class Support The input image `A` must be numeric. The output image `BW` is of class `logical`.

Example

```
I = imread('rice.png');  
BW = roicolor(I,128,255);  
imshow(I);  
figure, imshow(BW)
```



See Also `roifilt2`, `roipoly`

Purpose Smoothly interpolate within an arbitrary image region

Syntax `J = roifill(I,c,r)`

`J = roifill(I)`

`J = roifill(I,BW)`

`[J,BW] = roifill(...)`

`J = roifill(x,y,I,xi,yi)`

`[x,y,J,BW,xi,yi] = roifill(...)`

Description

`roifill` fills in a specified polygon in an intensity image. It smoothly interpolates inward from the pixel values on the boundary of the polygon by solving Laplace's equation. `roifill` can be used, for example, to erase small objects in an image.

`J = roifill(I,c,r)` fills in the polygon specified by `c` and `r`, which are equal-length vectors containing the row-column coordinates of the pixels on vertices of the polygon. The k th vertex is the pixel $(r(k), c(k))$.

`J = roifill(I)` displays the image `I` on the screen and lets you specify the polygon using the mouse. If you omit `I`, `roifill` operates on the image in the current axes. Use normal button clicks to add vertices to the polygon. Pressing **Backspace** or **Delete** removes the previously selected vertex. A shift-click, right-click, or double-click adds a final vertex to the selection and then starts the fill; pressing **Return** finishes the selection without adding a vertex.

`J = roifill(I,BW)` uses `BW` (a binary image the same size as `I`) as a mask. `roifill` fills in the regions in `I` corresponding to the nonzero pixels in `BW`. If there are multiple regions, `roifill` performs the interpolation on each region independently.

`[J,BW] = roifill(...)` returns the binary mask used to determine which pixels in `I` get filled. `BW` is a binary image the same size as `I` with 1's for pixels corresponding to the interpolated region of `I` and 0's elsewhere.

`J = roifill(x,y,I,xi,yi)` uses the vectors `x` and `y` to establish a nondefault spatial coordinate system. `xi` and `yi` are equal-length vectors that specify polygon vertices as locations in this coordinate system.

roifill

`[x,y,J,BW,xi,yi] = roifill(...)` returns the XData and YData in `x` and `y`, the output image in `J`, the mask image in `BW`, and the polygon coordinates in `xi` and `yi`. `xi` and `yi` are empty if the `roifill(I,BW)` form is used.

If `roifill` is called with no output arguments, the resulting image is displayed in a new figure.

Class Support

The input image `I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The input binary mask `BW` can be any numeric class or `logical`. The output binary mask `BW` is always `logical`. The output image `J` is of the same class as `I`. All other inputs and outputs are of class `double`.

Example

```
I = imread('eight.tif');  
c = [222 272 300 270 221 194];  
r = [21 21 75 121 121 75];  
J = roifill(I,c,r);  
imshow(I)  
figure, imshow(J)
```



See Also

`roifilt2`, `roipoly`

Purpose	Filter a region of interest in an image
Syntax	<pre>J = roifilt2(h,I,BW) J = roifilt2(I,BW,fun)</pre>
Description	<p><code>J = roifilt2(h,I,BW)</code> filters the data in <code>I</code> with the two-dimensional linear filter <code>h</code>. <code>BW</code> is a binary image the same size as <code>I</code> that is used as a mask for filtering. <code>roifilt2</code> returns an image that consists of filtered values for pixels in locations where <code>BW</code> contains 1's, and unfiltered values for pixels in locations where <code>BW</code> contains 0's. For this syntax, <code>roifilt2</code> calls <code>filter2</code> to implement the filter.</p> <p><code>J = roifilt2(I,BW,fun)</code> processes the data in <code>I</code> using the function <code>fun</code>. The result <code>J</code> contains computed values for pixels in locations where <code>BW</code> contains 1's, and the actual values in <code>I</code> for pixels in locations where <code>BW</code> contains 0's.</p> <p><code>fun</code> must be a function handle.</p>
Class Support	For the syntax that includes a filter <code>h</code> , the input image can be logical or numeric, and the output array <code>J</code> has the same class as the input image. For the syntax that includes a function, <code>I</code> can be of any class supported by <code>fun</code> , and the class of <code>J</code> depends on the class of the output from <code>fun</code> .
Example	<p>This example continues the <code>roipoly</code> example.</p> <pre>I = imread('eight.tif'); c = [222 272 300 270 221 194]; r = [21 21 75 121 121 75]; BW = roipoly(I,c,r); h = fspecial('unsharp'); J = roifilt2(h,I,BW); imshow(J), figure, imshow(J)</pre>



See Also

`imfilter`, `filter2`, `function_handle`, `roipoly`

Purpose Select a polygonal region of interest

Syntax `BW = roipoly(I,c,r)`
`BW = roipoly(I)`

`BW = roipoly(x,y,I,xi,yi)`
`[BW,xi,yi] = roipoly(...)`
`[x,y,BW,xi,yi] = roipoly(...)`

Description Use `roipoly` to select a polygonal region of interest within an image. `roipoly` returns a binary image that you can use as a mask for masked filtering.

`BW = roipoly(I,c,r)` returns the region of interest selected by the polygon described by vectors `c` and `r`. `BW` is a binary image the same size as `I` with 0's outside the region of interest and 1's inside.

`BW = roipoly(I)` displays the image `I` on the screen and lets you specify the polygon using the mouse. If you omit `I`, `roipoly` operates on the image in the current axes. Use normal button clicks to add vertices to the polygon. Pressing **Backspace** or **Delete** removes the previously selected vertex. A shift-click, right-click, or double-click adds a final vertex to the selection and then starts the fill; pressing **Return** finishes the selection without adding a vertex.

`BW = roipoly(x,y,I,xi,yi)` uses the vectors `x` and `y` to establish a nondefault spatial coordinate system. `xi` and `yi` are equal-length vectors that specify polygon vertices as locations in this coordinate system.

`[BW,xi,yi] = roipoly(...)` returns the polygon coordinates in `xi` and `yi`. Note that `roipoly` always produces a closed polygon. If the points specified describe a closed polygon (i.e., if the last pair of coordinates is identical to the first pair), the length of `xi` and `yi` is equal to the number of points specified. If the points specified do not describe a closed polygon, `roipoly` adds a final point having the same coordinates as the first point. (In this case the length of `xi` and `yi` is one greater than the number of points specified.)

`[x,y,BW,xi,yi] = roipoly(...)` returns the `XData` and `YData` in `x` and `y`, the mask image in `BW`, and the polygon coordinates in `xi` and `yi`.

If `roipoly` is called with no output arguments, the resulting image is displayed in a new figure.

roipoly

Class Support

The input image *I* can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image *BW* is of class `logical1`. All other inputs and outputs are of class `double`.

Remarks

For any of the `roipoly` syntaxes, you can replace the input image *I* with two arguments, *m* and *n*, that specify the row and column dimensions of an arbitrary image. For example, these commands create a 100-by-200 binary mask.

```
c = [112 112 79 79];  
r = [37 66 66 37];  
BW = roipoly(100,200,c,r);
```

If you specify *m* and *n* with an interactive form of `roipoly`, an *m*-by-*n* black image is displayed, and you use the mouse to specify a polygon within this image.

Example

```
I = imread('eight.tif');  
c = [222 272 300 270 221 194];  
r = [21 21 75 121 121 75];  
BW = roipoly(I,c,r);  
imshow(I)  
figure, imshow(BW)
```



See Also

`roifilt2`, `roicolor`, `roifill`, `poly2mask`

Purpose	Compute the standard deviation of the elements of a matrix
Syntax	<code>b = std2(A)</code>
Description	<code>b = std2(A)</code> computes the standard deviation of the values in A.
Class Support	A can be numeric or logical. B is a scalar of class double.
Algorithm	<code>std2</code> computes the standard deviation of the array A using <code>std(A(:))</code> .
See Also	<code>corr2</code> , <code>mean2</code> <code>std</code> , <code>mean</code> in the MATLAB Function Reference

stdfilt

Purpose Local standard deviation of an image

Syntax

```
J = stdfilt(I)
J = stdfilt(I,NHOOD)
```

Description `J = stdfilt(I)` returns the array `J`, where each output pixel contains the standard deviation of the 3-by-3 neighborhood around the corresponding pixel in the input image `I`. `I` can have any dimension. The output image `J` is the same size as the input image `I`.

For pixels on the borders of `I`, `stdfilt` uses symmetric padding. In symmetric padding, the values of padding pixels are a mirror reflection of the border pixels in `I`.

`J = stdfilt(I,NHOOD)` calculates the local standard deviation of the input image `I`, where you specify the neighborhood in `NHOOD`. `NHOOD` is a multidimensional array of zeros and ones where the nonzero elements specify the neighbors. `NHOOD`'s size must be odd in each dimension.

By default, `stdfilt` uses the neighborhood ones(3). `stdfilt` determines the center element of the neighborhood by `floor((size(NHOOD) + 1)/2)`.

Class Support `I` can be logical or numeric and must be real and nonsparse. `NHOOD` can be logical or numeric and must contain zeros and/or ones. `J` is of class double.

Notes To specify neighborhoods of various shapes, such as a disk, use the `strel` function to create a structuring element object and then use the `getnhood` function to extract the neighborhood from the structuring element object.

Example

```
I = imread('circuit.tif');
J = stdfilt(I);
imshow(I);
figure, imshow(J,[]);
```

See also `entropyfilt`, `getnhood`, `rangefilt`, `std2`, `strel`

Purpose Create morphological structuring element


Syntax SE = strel(shape,parameters)

Description SE = strel(shape,parameters) creates a structuring element, SE, of the type specified by shape. This table lists all the supported shapes. Depending on shape, strel can take additional parameters. See the syntax descriptions that follow for details about creating each type of structuring element.

Flat Structuring Elements	
'arbitrary'	'pair'
'diamond'	'periodicline'
'disk'	'rectangle'
'line'	'square'
'octagon'	

Nonflat Structuring Elements	
'arbitrary'	'ball'

SE = strel('arbitrary',NHOOD) creates a flat structuring element where NHOOD specifies the neighborhood. NHOOD is a matrix containing 1's and 0's; the location of the 1's defines the neighborhood for the morphological operation. The center (or *origin*) of NHOOD is its center element, given by floor((size(NHOOD)+1)/2). You can omit the 'arbitrary' string and just use strel(NHOOD).




SE = strel('arbitrary',NHOOD,HEIGHT) creates a nonflat structuring element, where NHOOD specifies the neighborhood. HEIGHT is a matrix the same size as NHOOD containing the height values associated with each nonzero element of NHOOD. The HEIGHT matrix must be real and finite valued. You can omit the 'arbitrary' string and just use strel(NHOOD,HEIGHT).

SE = strel('ball',R,H,N) creates a nonflat, ball-shaped structuring element (actually an ellipsoid) whose radius in the X-Y plane is R and whose height is H. Note that R must be a nonnegative integer, H must be a real scalar, and N must be an even nonnegative integer. When N is greater than 0, the ball-shaped structuring element is approximated by a sequence of N nonflat, line-shaped structuring elements. When N equals 0, no approximation is used, and the structuring element members consist of all pixels whose centers are no greater than R away from the origin. The corresponding height values are determined from the formula of the ellipsoid specified by R and H. If N is not specified, the default value is 8.

Note Morphological operations run much faster when the structuring element uses approximations ($N > 0$) than when it does not ($N = 0$).


SE = strel('diamond',R) creates a flat, diamond-shaped structuring element, where R specifies the distance from the structuring element origin to the points of the diamond. R must be a nonnegative integer scalar.



SE = strel('disk',R,N) creates a flat, disk-shaped structuring element, where R specifies the radius. R must be a nonnegative integer. N must be 0, 4, 6, or 8. When N is greater than 0, the disk-shaped structuring element is approximated by a sequence of N periodic-line structuring elements. When N


equals 0, no approximation is used, and the structuring element members consist of all pixels whose centers are no greater than R away from the origin. If N is not specified, the default value is 4.

Note Morphological operations run much faster when the structuring element uses approximations ($N > 0$) than when it does not ($N = 0$). However, structuring elements that do not use approximations ($N = 0$) are not suitable for computing granulometries. Sometimes it is necessary for `strel` to use two extra line structuring elements in the approximation, in which case the number of decomposed structuring elements used is $N + 2$.




`SE = strel('line', LEN, DEG)` creates a flat, linear structuring element, where `LEN` specifies the length, and `DEG` specifies the angle (in degrees) of the line, as measured in a counterclockwise direction from the horizontal axis. `LEN` is


approximately the distance between the centers of the structuring element members at opposite ends of the line.




`SE = strel('octagon',R)` creates a flat, octagonal structuring element, where `R` specifies the distance from the structuring element origin to the sides of the octagon, as measured along the horizontal and vertical axes. `R` must be a nonnegative multiple of 3.




`SE = strel('pair',OFFSET)` creates a flat structuring element containing two members. One member is located at the origin. The second member's location is specified by the vector `OFFSET`. `OFFSET` must be a two-element vector of integers.




`SE = strel('periodicline',P,V)` creates a flat structuring element containing $2*P+1$ members. V is a two-element vector containing integer-valued row and column offsets. One structuring element member is located at the origin. The other members are located at $1*V, -1*V, 2*V, -2*V, \dots, P*V, -P*V$.



`SE = strel('rectangle',MN)` creates a flat, rectangle-shaped structuring element, where MN specifies the size. MN must be a two-element vector of nonnegative integers. The first element of MN is the number of rows in the structuring element neighborhood; the second element is the number of columns.



`SE = strel('square',W)` creates a square structuring element whose width is W pixels. W must be a nonnegative integer scalar.



Notes

For all shapes except 'arbitrary', structuring elements are constructed using a family of techniques known collectively as *structuring element decomposition*. The principle is that dilation by some large structuring elements can be computed faster by dilation with a sequence of smaller structuring elements. For example, dilation by an 11-by-11 square structuring element can be accomplished by dilating first with a 1-by-11 structuring element and then

with an 11-by-1 structuring element. This results in a theoretical performance improvement of a factor of 5.5, although in practice the actual performance improvement is somewhat less. Structuring element decompositions used for the 'disk' and 'ball' shapes are approximations; all other decompositions are exact.

Methods

This table lists the methods supported by the STREL object.

Method	Description
getheight	Get height of structuring element
getneighbors	Get structuring element neighbor locations and heights
getnhood	Get structuring element neighborhood
getsequence	Extract sequence of decomposed structuring elements
isflat	Return true for flat structuring element
reflect	Reflect structuring element
translate	Translate structuring element

Example

```
se1 = strel('square',11)      % 11-by-11 square
se2 = strel('line',10,45)    % line, length 10, angle 45 degrees
se3 = strel('disk',15)       % disk, radius 15
se4 = strel('ball',15,5)     % ball, radius 15, height 5
```

Algorithm

The method used to decompose diamond-shaped structuring elements is known as “logarithmic decomposition” [1].

The method used to decompose disk structuring elements is based on the technique called “radial decomposition using periodic lines” [2], [3]. For details, see the `MakeDiskStrel` subfunction in `toolbox/images/images/@strel/strel.m`.

The method used to decompose ball structuring elements is the technique called “radial decomposition of sphere” [2].

See Also

imdilate, imerode

References

- [1] van den Boomgard, Rein, and Richard van Balen, "Methods for Fast Morphological Image Transforms Using Bitmapped Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, No. 3, May 1992, pp. 252-254.
- [2] Adams, Rolf, "Radial Decomposition of Discs and Spheres," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 55, No. 5, September 1993, pp. 325-332.
- [3] Jones, Ronald, and Pierre Soille, "Periodic lines: Definition, cascades, and application to granulometrie," *Pattern Recognition Letters*, Vol. 17, 1996, pp. 1057-1063.

stretchlim

Purpose Find limits to contrast stretch an image

Syntax `LOW_HIGH = stretchlim(I,TOL)`
`LOW_HIGH = stretchlim(RGB,TOL)`

Description `LOW_HIGH = stretchlim(I,TOL)` returns a pair of intensities that can be used by `imadjust` to increase the contrast of an image.

`TOL = [LOW_FRACT HIGH_FRACT]` specifies the fraction of the image to saturate at low and high intensities.

If `TOL` is a scalar, `LOW_FRACT = TOL`, and `HIGH_FRACT = 1 - LOW_FRACT`, which saturates equal fractions at low and high intensities.

If you omit the argument, `TOL` defaults to `[0.01 0.99]`, saturating 2%.

If `TOL = 0`, `LOW_HIGH = [min(I(:)) max(I(:))]`.

`LOW_HIGH = stretchlim(RGB,TOL)` returns a 2-by-3 matrix of intensity pairs to saturate each plane of the RGB image. `TOL` specifies the same fractions of saturation for each plane.

Class Support The input image can be of class `uint8`, `uint16`, `int16`, `double`, or `single`. The output intensities returned, `LOW_HIGH`, are of class `double` and have values between 0 and 1.

Example

```
I = imread('pout.tif');
J = imadjust(I,stretchlim(I),[]);
imshow(I), figure, imshow(J)
```



See Also

brighten, histeq, imadjust

subimage

Purpose Display multiple images in the single figure

Syntax

```
subimage(X,map)
subimage(I)
subimage(BW)
subimage(RGB)
subimage(x,y,...)
h = subimage(...)
```

Description You can use `subimage` in conjunction with `subplot` to create figures with multiple images, even if the images have different colormaps. `subimage` works by converting images to true color for display purposes, thus avoiding colormap conflicts.

`subimage(X,map)` displays the indexed image `X` with colormap `map` in the current axes.

`subimage(I)` displays the intensity image `I` in the current axes.

`subimage(BW)` displays the binary image `BW` in the current axes.

`subimage(RGB)` displays the truecolor image `RGB` in the current axes.

`subimage(x,y,...)` displays an image using a nondefault spatial coordinate system.

`h = subimage(...)` returns a handle to an image object.

Class Support The input image can be of class `logical`, `uint8`, `uint16`, or `double`.

Example

```
load trees
[X2,map2] = imread('forest.tif');
subplot(1,2,1), subimage(X,map)
subplot(1,2,2), subimage(X2,map2)
```

See Also

`imshow`

`subplot` in the MATLAB Function Reference

Purpose Spatial transformation of a multidimensional array

Syntax `B = tformarray(A,T,R,TDIMS_A,TDIMS_B,TSIZE_B,TMAP_B,F)`

Description `B = tformarray(A,T,R,TDIMS_A,TDIMS_B,TSIZE_B,TMAP_B,F)` applies a spatial transformation to array `A` to produce array `B`. The `tformarray` function is like `imtransform`, but is intended for problems involving higher-dimensioned arrays or mixed input/output dimensionality, or requiring greater user control or customization. (Anything that can be accomplished with `imtransform` can be accomplished with a combination of `maketform`, `makeresampler`, `findbounds`, and `tformarray`; but for many tasks involving 2-D images, `imtransform` is simpler.)

This table provides a brief description of all the input arguments. See the following section for more detail about each argument. (Click an argument in the table to move to the appropriate section.)

Argument	Description
A	Input array or image
T	Spatial transformation structure, called a TFORM, typically created with <code>maketform</code>
R	Resampler structure, typically created with <code>makeresampler</code>
TDIMS_A	Row vector listing the input transform dimensions
TDIMS_B	Row vector listing the output transform dimensions
TSIZE_B	Output array size in the transform dimensions
TMAP_B	Array of point locations in output space; can be used as an alternative way to specify a spatial transformation
F	Array of fill values

`A` can be any nonsparse numeric array, and can be real or complex.

T is a TFORM structure that defines a particular spatial transformation. For each location in the output transform subscript space (as defined by TDIMS_B and TSIZE_B), tformarray uses T and the function tforminv to compute the corresponding location in the input transform subscript space (as defined by TDIMS_A and size(A)).

If T is empty, tformarray operates as a direct resampling function, applying the resampler defined in R to compute values at each transform space location defined in TMAP_B (if TMAP_B is nonempty), or at each location in the output transform subscript grid.

R is a structure that defines how to interpolate values of the input array at specified locations. R is usually created with makeresampler, which allows fine control over how to interpolate along each dimension, as well as what input array values to use when interpolating close to the edge of the array.

TDIMS_A and TDIMS_B indicate which dimensions of the input and output arrays are involved in the spatial transformation. Each element must be unique, and must be a positive integer. The entries need not be listed in increasing order, but the order matters. It specifies the precise correspondence between dimensions of arrays A and B and the input and output spaces of the transformer T. length(TDIMS_A) must equal T.ndims_in, and LENGTH(TDIMS_B) must equal T.ndims_out.

For example, if T is a 2-D transformation, TDIMS_A = [2 1], and TDIMS_B = [1 2], then the column dimension and row dimension of A correspond to the first and second transformation input-space dimensions, respectively. The row and column dimensions of B correspond to the first and second output-space dimensions, respectively.

TSIZE_B specifies the size of the array B along the output-space transform dimensions. Note that the size of B along nontransform dimensions is taken directly from the size of A along those dimensions. If, for example, T is a 2-D transformation, size(A) = [480 640 3 10], TDIMS_B is [2 1], and TSIZE_B is [300 200], then size(B) is [200 300 3].

TMAP_B is an optional array that provides an alternative way of specifying the correspondence between the position of elements of B and the location in output transform space. TMAP_B can be used, for example, to compute the result of an image warp at a set of arbitrary locations in output space. If TMAP_B is not empty, then the size of TMAP_B takes the form

[D1 D2 D3 ... DN L]

where N equals `length(TDIMS_B)`. The vector [D1 D2 ... DN] is used in place of `TSIZE_B`. If `TMAP_B` is not empty, then `TSIZE_B` should be [].

The value of L depends on whether or not `T` is empty. If `T` is not empty, then L is `T.ndims_out`, and each L -dimension point in `TMAP_B` is transformed to an input-space location using `T`. If `T` is empty, then L is `length(TDIMS_A)`, and each L -dimensional point in `TMAP_B` is used directly as a location in input space.

`F` is a double-precision array containing fill values. The fill values in `F` can be used in three situations:

- When a separable resampler is created with `makeresampler` and its `padmethod` is set to either 'fill' or 'bound'.
- When a custom resampler is used that supports the 'fill' or 'bound' pad methods (with behavior that is specific to the customization).
- When the map from the transform dimensions of `B` to the transform dimensions of `A` is deliberately undefined for some points. Such points are encoded in the input transform space by NaNs in either `TMAP_B` or in the output of `TFORMINV`.

In the first two cases, fill values are used to compute values for output locations that map outside or near the edges of the input array. Fill values are copied into `B` when output locations map well outside the input array. See `makeresampler` for more information about 'fill' and 'bound'.

`F` can be a scalar (including NaN), in which case its value is replicated across all the nontransform dimensions. `F` can also be a nonscalar, whose size depends on `size(A)` in the nontransform dimensions. Specifically, if K is the J th nontransform dimension of `A`, then `size(F,J)` must be either `size(A,K)` or 1. As a convenience to the user, `tformarray` replicates `F` across any dimensions with unit size such that after the replication `size(F,J)` equals `size(A,K)`.

For example, suppose `A` represents 10 RGB images and has size 200-by-200-by-3-by-10, `T` is a 2-D transformation, and `TDIMS_A` and `TDIMS_B` are both [1 2]. In other words, `tformarray` will apply the same 2-D transform to each color plane of each of the 10 RGB images. In this situation you have several options for `F`:

tformarray

- F can be a scalar, in which case the same fill value is used for each color plane of all 10 images.
- F can be a 3-by-1 vector, [R G B]'. Then R, G, and B are used as the fill values for the corresponding color planes of each of the 10 images. This can be interpreted as specifying an RGB fill color, with the same color used for all 10 images.
- F can be a 1-by-10 vector. This can be interpreted as specifying a different fill value for each of 10 images, with that fill value being used for all three color planes.
- F can be a 3-by-10 matrix, which can be interpreted as supplying a different RGB fill color for each of the 10 images.

Class Support

A can be any nonsparse numeric array, and can be real or complex. It can also be of class `logical`.

Example

Create a 2-by-2 checkerboard image where each square is 20 pixels wide, then transform it with a projective transformation. Use a `pad` method of 'circular' when creating a resampler, so that the output appears to be a perspective view of an infinite checkerboard. Swap the output dimensions. Specify a 100-by-100 output image. Leave `TMAP_B` empty, since `TSIZE_B` is specified. Leave the fill value empty, since it won't be needed.

```
I = checkerboard(20,1,1);
figure; imshow(I)
T = maketform('projective',[1 1; 41 1; 41 41; 1 41],...
             [5 5; 40 5; 35 30; -10 30]);
R = makesampler('cubic','circular');
J = tformarray(I,T,R,[1 2],[2 1],[100 100],[[],[]]);
figure; imshow(J)
```

See Also

`findbounds`, `imtransform`, `makesampler`, `maketform`

Purpose Apply forward spatial transformation

Syntax

```
[X,Y] = tformfwd(T,U,V)
[X1,X2,X3,...] = tformfwd(T,U1,U2,U3,...)
X = tformfwd(T,U)
[X1,X2,X3,...] = tformfwd(T,U)
X = tformfwd(T,U1,U2,U3,...)
```

Description `[X,Y] = tformfwd(T,U,V)` applies the 2D-to-2D spatial transformation defined in `T` to coordinate arrays `U` and `V`, mapping the point `[U(k) V(k)]` to the point `[X(k) Y(k)]`.

`T` is a `TFORM` struct created with `maketform`, `fliptform`, or `cp2tform`. Both `T.ndims_in` and `T.ndims_out` must equal 2. `U` and `V` are typically column vectors matching in length. In general, `U` and `V` can have any dimensionality, but must have the same size. In any case, `X` and `Y` will have the same size as `U` and `V`.

`[X1,X2,X3,...] = tformfwd(T,U1,U2,U3,...)` applies the `ndims_in`-to-`ndims_out` spatial transformation defined in `TFORM` structure `T` to the coordinate arrays `U1,U2,...,UNDIMS_IN` (where `NDIMS_IN = T.ndims_in` and `NDIMS_OUT = T.ndims_out`). The number of output arguments must equal `NDIMS_OUT`. The transformation maps the point

$$[U1(k) \ U2(k) \ \dots \ UNDIMS_IN(k)]$$

to the point

$$[X1(k) \ X2(k) \ \dots \ XNDIMS_OUT(k)].$$

`U1,U2,U3,...` can have any dimensionality, but must be the same size.

`X1,X2,X3,...` must have this size also.

`X = tformfwd(T,U)` applies the `ndims_in`-to-`ndims_out` spatial transformation defined in `TFORM` structure `T` to each row of `U`, where `U` is an `M`-by-`NDIMS_IN` matrix. It maps the point `U(k,:)` to the point `X(k,:)`. `X` is an `M`-by-`NDIMS_OUT` matrix.

`X = tformfwd(T,U)`, where `U` is an `(N+1)`-dimensional array, maps the point `U(k1,k2,...,kN,:)` to the point `X(k1,k2,...,kN,:)`. `size(U,N+1)` must equal

tformfwd

NDIMS_IN. X is an $(N+1)$ -dimensional array, with $\text{size}(X, I)$ equal to $\text{size}(U, I)$ for $I = 1, \dots, N$ and $\text{size}(X, N+1)$ equal to NDIMS_OUT.

$[X_1, X_2, X_3, \dots] = \text{tformfwd}(T, U)$ maps an $(N+1)$ -dimensional array to NDIMS_OUT equally sized N -dimensional arrays.

$X = \text{tformfwd}(T, U_1, U_2, U_3, \dots)$ maps NDIMS_IN N -dimensional arrays to one $(N+1)$ -dimensional array.

Note

$X = \text{tformfwd}(U, T)$ is an older form of the two-argument syntax that remains supported for backward compatibility.

Example

Create an affine transformation that maps the triangle with vertices $(0,0)$, $(6,3)$, $(-2,5)$ to the triangle with vertices $(-1,-1)$, $(0,-10)$, $(4,4)$.

```
u = [ 0  6  -2]';  
v = [ 0  3  5]';  
x = [-1  0  4]';  
y = [-1 -10  4]';  
tform = maketform('affine',[u v],[x y]);
```

Validate the mapping by applying tformfwd. Results should equal $[x, y]$

```
[xm, ym] = tformfwd(tform, u, v)
```

See Also

cp2tform, fliptform, maketform, tforminv

Purpose Apply inverse spatial transformation

Syntax $U = \text{tforminv}(X, T)$

Description $[U, V] = \text{tforminv}(T, X, Y)$ applies the 2D-to-2D inverse transformation defined in TFORM structure T to coordinate arrays X and Y , mapping the point $[X(k) \ Y(k)]$ to the point $[U(k) \ V(k)]$. Both $T.\text{ndims_in}$ and $T.\text{ndims_out}$ must equal 2. X and Y are typically column vectors matching in length. In general, X and Y can have any dimensionality, but must have the same size. In any case, U and V will have the same size as X and Y .

$[U1, U2, U3, \dots] = \text{tforminv}(T, X1, X2, X3, \dots)$ applies the NDIMS_OUT-to-NDIMS_IN inverse transformation defined in TFORM structure T to the coordinate arrays $X1, X2, \dots, X\text{NDIMS_OUT}$ (where $\text{NDIMS_IN} = T.\text{ndims_in}$ and $\text{NDIMS_OUT} = T.\text{ndims_out}$). The number of output arguments must equal NDIMS_IN . The transformation maps the point

$$[X1(k) \ X2(k) \ \dots \ X\text{NDIMS_OUT}(k)]$$

to the point

$$[U1(k) \ U2(k) \ \dots \ \text{NDIMS_IN}(k)].$$

$X1, X2, X3, \dots$ can have any dimensionality, but must be the same size.

$U1, U2, U3, \dots$ have this size also.

$U = \text{tforminv}(T, X)$ applies the NDIMS_OUT-to-NDIMS_IN inverse transformation defined in TFORM structure T to each row of X , where X is an M -by-NDIMS_OUT matrix. It maps the point $X(k, :)$ to the point $U(k, :)$. U is an M -by-NDIMS_IN matrix.

$U = \text{tforminv}(T, X)$, where X is an $(N+1)$ -dimensional array, maps the point $X(k1, k2, \dots, kN, :)$ to the point $U(k1, k2, \dots, kN, :)$. $\text{size}(X, N+1)$ must equal NDIMS_OUT . U is an $(N+1)$ -dimensional array, with $\text{size}(U, I)$ equal to $\text{size}(X, I)$ for $I = 1, \dots, N$ and $\text{size}(U, N+1)$ equal to NDIMS_IN .

$[U1, U2, U3, \dots] = \text{tforminv}(T, X)$ maps an $(N+1)$ -dimensional array to NDIMS_IN equally-sized N -dimensional arrays.

tforminv

`U = tforminv(T,X1,X2,X3,...)` maps `NDIMS_OUT` N-dimensional arrays to one (N+1)-dimensional array.

Note

`U = tforminv(X,T)` is an older form of the two-argument syntax that remains supported for backward compatibility.

Example

Create an affine transformation that maps the triangle with vertices (0,0), (6,3), (-2,5) to the triangle with vertices (-1,-1), (0,-10), (4,4).

```
u = [ 0  6 -2]';  
v = [ 0  3  5]';  
x = [-1  0  4]';  
y = [-1 -10  4]';  
tform = maketform('affine',[u v],[x y]);
```

Validate the mapping by applying `tforminv`. Results should equal `[u, v]`.

```
[um, vm] = tforminv(tform, x, y)
```

See Also

`cp2tform`, `tforminv`, `maketform`, `fliptform`

Purpose Translate structuring element

Syntax SE2 = translate(SE,V)

Description SE2 = reflect(SE,V) translates a structuring element SE in N-D space. V is an N-element vector containing the offsets of the desired translation in each dimension.

Class Support SE and SE2 are STREL objects; V is a vector of double-precision values.

Example Dilating with a translated version of `strel(1)` is a way to translate the input image in space. This example translates the `cameraman.tif` image down and to the right by 25 pixels.

```
I = imread('cameraman.tif');  
se = translate(strel(1), [25 25]);  
J = imdilate(I,se);  
imshow(I), title('Original')  
figure, imshow(J), title('Translated');
```



See Also `strel`, `reflect`

truesize

Purpose Adjust display size of an image

Syntax `truesize(fig,[mrows mcols])`
`truesize(fig)`

Description `truesize(fig,[mrows ncols])` adjusts the display size of an image. `fig` is a figure containing a single image or a single image with a colorbar. `[mrows ncols]` is a 1-by-2 vector that specifies the requested screen area in pixels that the image should occupy.

`truesize(fig)` uses the image height and width for `[mrows ncols]`. This results in the display's having one screen pixel for each image pixel.

If you omit the figure argument, `truesize` works on the current figure.

Remarks If the 'TruesizeWarning' toolbox preference is 'on', `truesize` displays a warning if the image is too large to fit on the screen. (The entire image is still displayed, but at less than true size.) If 'TruesizeWarning' is 'off', `truesize` does not display the warning. Note that this preference applies even when you call `truesize` indirectly, such as through `imshow`.

See Also `imshow`, `iptsetpref`, `iptgetpref`

Purpose

Convert data to unsigned 16-bit integers

`uint16` is a MATLAB built-in function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

uint8

Purpose

Convert data to unsigned 8-bit integers

`uint8` is a MATLAB built-in function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

Purpose Compute new values of A based on lookup table LUT

Note uintlut is an obsolete version of intlut and may be removed in a future version of the toolbox.

Syntax B = uintlut(A,LUT)

Purpose B = uintlut(A,LUT) creates an array B containing new values of A based on the lookup table LUT. For example, if A is a vector whose *k*th element is equal to alpha, then B(*k*) is equal to the LUT value corresponding to alpha, that is, LUT(alpha+1).

Class Support A must be uint8 or uint16. If A is uint8, then LUT must be a uint8 vector with 256 elements. If A is uint16, then LUT must be a uint16 vector with 65536 elements. B has the same size and class as A.

Example

```
A = uint8([1 2 3 4; 5 6 7 8;9 10 11 12]);
LUT = repmat(uint8([0 150 200 255]),1,64);
B = uintlut(A,LUT);
imshow(A), figure, imshow(B);
```

See Also `impixel`, `improfile`

warp

Purpose Display an image as a texture-mapped surface

Syntax

```
warp(X,map)
warp(I,n)
warp(BW)
warp(RGB)
warp(z,...)
warp(x,y,z,...)
h = warp(...)
```

Description `warp(X,map)` displays the indexed image `X` with colormap `map` as a texture map on a simple rectangular surface.

`warp(I,n)` displays the intensity image `I` with grayscale colormap of length `n` as a texture map on a simple rectangular surface.

`warp(BW)` displays the binary image `BW` as a texture map on a simple rectangular surface.

`warp(RGB)` displays the RGB image in the array `RGB` as a texture map on a simple rectangular surface.

`warp(z,...)` displays the image on the surface `z`.

`warp(x,y,z,...)` displays the image on the surface `(x,y,z)`.


`h = warp(...)` returns a handle to a texture-mapped surface.

Class Support The input image can be of class `logical`, `uint8`, `uint16`, or `double`.

Remarks Texture-mapped surfaces are generally rendered more slowly than images.

Example This example texture maps an image of a test pattern onto a cylinder.

```
[x,y,z] = cylinder;
I = imread('testpat1.png');
warp(x,y,z,I);
```



See Also

`imshow`

`image`, `imagesc`, `surf` in the MATLAB Function Reference

watershed

Purpose Find image watershed regions

Syntax L = watershed(A)
L = watershed(A,CONN)

Description L = watershed(A) computes a label matrix identifying the watershed regions of the input matrix A, which can have any dimension. The elements of L are integer values greater than or equal to 0. The elements labeled 0 do not belong to a unique watershed region. These are called *watershed pixels*. The elements labeled 1 belong to the first watershed region, the elements labeled 2 belong to the second watershed region, and so on.

By default, watershed uses 8-connected neighborhoods for 2-D inputs and 26-connected neighborhoods for 3-D inputs. For higher dimensions, watershed uses the connectivity given by `conndef(ndims(A), 'maximal')`.

L = watershed(A,CONN) specifies the connectivity to be used in the watershed computation. CONN can have any of the following scalar values.

Value	Meaning
Two-dimensional connectivities	
4	4-connected neighborhood
8	8-connected neighborhood
Three-dimensional connectivities	
6	6-connected neighborhood
18	18-connected neighborhood
26	26-connected neighborhood

Connectivity can be defined in a more general way for any dimension by using for CONN a 3-by-3-by-...-by-3 matrix of 0's and 1's. The 1-valued elements define neighborhood locations relative to the center element of CONN. Note that CONN must be symmetric about its center element.

Class Support

A can be a numeric or logical array of any dimension, and it must be nonsparse. The output array L is of class double.

Example**2-D Example**

- 1 Make a binary image containing two overlapping circular objects.

```
center1 = -10;
center2 = -center1;
dist = sqrt(2*(2*center1)^2);
radius = dist/2 * 1.4;
lims = [floor(center1-1.2*radius) ceil(center2+1.2*radius)];
[x,y] = meshgrid(lims(1):lims(2));
bw1 = sqrt((x-center1).^2 + (y-center1).^2) <= radius;
bw2 = sqrt((x-center2).^2 + (y-center2).^2) <= radius;
bw = bw1 | bw2;
figure, imshow(bw,'InitialMagnification','fit'), title('bw')
```

- 2 Compute the distance transform of the complement of the binary image.

```
D = bwdist(~bw);
figure, imshow(D,[],'InitialMagnification','fit')
title('Distance transform of ~bw')
```

- 3 Complement the distance transform, and force pixels that don't belong to the objects to be at $-\infty$.

```
D = -D;
D(~bw) = -Inf;
```

- 4 Compute the watershed transform and display it as an indexed image.

```
L = watershed(D);
rgb = label2rgb(L,'jet',[.5 .5 .5]);
figure, imshow(rgb,'InitialMagnification','fit')
title('Watershed transform of D')
```

3-D Example

- 1 Make a 3-D binary image containing two overlapping spheres.

```
center1 = -10;
center2 = -center1;
dist = sqrt(3*(2*center1)^2);
radius = dist/2 * 1.4;
```

watershed

```
lims = [floor(center1-1.2*radius) ceil(center2+1.2*radius)];
[x,y,z] = meshgrid(lims(1):lims(2));
bw1 = sqrt((x-center1).^2 + (y-center1).^2 + ...
           (z-center1).^2) <= radius;
bw2 = sqrt((x-center2).^2 + (y-center2).^2 + ...
           (z-center2).^2) <= radius;
bw = bw1 | bw2;
figure, isosurface(x,y,z,bw,0.5), axis equal, title('BW')
xlabel x, ylabel y, zlabel z
xlim(lims), ylim(lims), zlim(lims)
view(3), camlight, lighting gouraud
```

2 Compute the distance transform.

```
D = bwdist(-bw);
figure, isosurface(x,y,z,D,radius/2), axis equal
title('Isosurface of distance transform')
xlabel x, ylabel y, zlabel z
xlim(lims), ylim(lims), zlim(lims)
view(3), camlight, lighting gouraud
```

3 Complement the distance transform, force nonobject pixels to be -Inf, and then compute the watershed transform.

```
D = -D;
D(~bw) = -Inf;
L = watershed(D);
figure, isosurface(x,y,z,L==2,0.5), axis equal
title('Segmented object')
xlabel x, ylabel y, zlabel z
xlim(lims), ylim(lims), zlim(lims)
view(3), camlight, lighting gouraud
figure, isosurface(x,y,z,L==3,0.5), axis equal
title('Segmented object')
xlabel x, ylabel y, zlabel z
xlim(lims), ylim(lims), zlim(lims)
view(3), camlight, lighting gouraud
```

Algorithm

watershed uses a variation of the Vincent and Soille algorithm [1]. For details of the variation, see `toolbox/images/images/private/watershed_vs.h`.

See Also

`bwlabel`, `bwlabeln`, `bwdist`, `regionprops`

Reference

[1] Vincent, Luc, and Pierre Soille, "Watersheds in Digital Spaces: An Efficient Algorithm Based on Immersion Simulations," *IEEE Transactions of Pattern Analysis and Machine Intelligence*, Vol. 13, No. 6, June 1991, pp. 583-598.

whitepoint

Purpose Return *XYZ* values of standard illuminants

Syntax
`xyz = whitepoint`
`xyz = whitepoint(string)`

Description `xyz = whitepoint(string)` returns `xyz`, a three-element row vector of *XYZ* values scaled so that $Y = 1$. `string` specifies the white reference illuminant. The following table lists all the possible values for `string`. The default value is enclosed in braces (`{}`).

Value	Description
'a'	CIE standard illuminant A
'c'	CIE standard illuminant C
'd50'	CIE standard illuminant D50
'd55'	CIE standard illuminant D55
'd65'	CIE standard illuminant D65
{'icc'}	ICC standard profile connection space illuminant; a 16-bit fractional approximation of D50

`xyz = whitepoint` is the same as `xyz = whitepoint('icc')`.

Class Support `string` is a character array. `xyz` is of class `double`.

Example This example returns the *XYZ* color space representation of the default white reference illuminant `'icc'`.

```
wp_icc = whitepoint
wp_icc =
    0.9642    1.0000    0.8249
```

Purpose	Perform two-dimensional adaptive noise-removal filtering
Syntax	<pre>J = wiener2(I,[m n],noise) [J,noise] = wiener2(I,[m n])</pre>
Description	<p><code>wiener2</code> lowpass-filters an intensity image that has been degraded by constant power additive noise. <code>wiener2</code> uses a pixelwise adaptive Wiener method based on statistics estimated from a local neighborhood of each pixel.</p> <p><code>J = wiener2(I,[m n],noise)</code> filters the image <code>I</code> using pixelwise adaptive Wiener filtering, using neighborhoods of size <code>m</code>-by-<code>n</code> to estimate the local image mean and standard deviation. If you omit the <code>[m n]</code> argument, <code>m</code> and <code>n</code> default to 3. The additive noise (Gaussian white noise) power is assumed to be <code>noise</code>.</p> <p><code>[J,noise] = wiener2(I,[m n])</code> also estimates the additive noise power before doing the filtering. <code>wiener2</code> returns this estimate in <code>noise</code>.</p>
Class Support	The input image <code>I</code> is a two-dimensional image of class <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> . The output image <code>J</code> is of the same size and class as <code>I</code> .
Example	For an example, see “Using Adaptive Filtering” on page 10-52.
Algorithm	<code>wiener2</code> estimates the local mean and variance around each pixel,

$$\mu = \frac{1}{NM} \sum_{n_1, n_2 \in \eta} a(n_1, n_2)$$

$$\sigma^2 = \frac{1}{NM} \sum_{n_1, n_2 \in \eta} a^2(n_1, n_2) - \mu^2$$

where η is the N -by- M local neighborhood of each pixel in the image `A`. `wiener2` then creates a pixelwise Wiener filter using these estimates,

$$b(n_1, n_2) = \mu + \frac{\sigma^2 - v^2}{\sigma^2} (a(n_1, n_2) - \mu)$$

where v^2 is the noise variance. If the noise variance is not given, `wiener2` uses the average of all the local estimated variances.

wiener2

See Also

`filter2`, `medfilt2`

Reference

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 536-540.

Purpose Convert color data from *XYZ* representation to double

Syntax `xyzd = xyz2double(XYZ)`

Description `xyzd = xyz2double(XYZ)` converts an M-by-3 or M-by-N-by-3 array of *XYZ* color values to double. `xyzd` has the same size as *XYZ*.

The Image Processing Toolbox follows the convention that double-precision *XYZ* arrays contain 1931 CIE *XYZ* values. *XYZ* arrays that are `uint16` follow the convention in the ICC profile specification (ICC.1:2001-4, www.color.org) for representing *XYZ* values as unsigned 16-bit integers. There is no standard representation of *XYZ* values as unsigned 8-bit integers. The ICC encoding convention is illustrated by this table.

Value (X, Y, or Z)	uint16 Value
0.0	0
1.0	32768
$1.0 + (32767/32768)$	65535

Class Support `xyz` is a `uint16` or `double` array that must be real and nonsparse. `xyzd` is of class `double`.

Example Convert `uint16`-encoded *XYZ* values to double.

```
xyz2double(uint16([100 32768 65535]))
ans =

    3.0518e-003    1.0000e+000    2.0000e+000
```

See Also `applycform`, `lab2double`, `lab2uint16`, `lab2uint8`, `makecform`, `whitepoint`, `xyz2uint16`

xyz2uint16

Purpose Convert color data from *XYZ* representation to uint16

Syntax `xyz16 = xyz2uint16(xyz)`

Description `xyz16 = xyz2uint16(xyz)` converts an M-by-3 or M-by-N-by-3 array of *XYZ* color values to uint16. `xyz16` has the same size as `xyz`.

The Image Processing Toolbox follows the convention that double-precision *XYZ* arrays contain 1931 CIE *XYZ* values. *XYZ* arrays that are uint16 follow the convention in the ICC profile specification (ICC.1:2001-4, www.color.org) for representing *XYZ* values as unsigned 16-bit integers. There is no standard representation of *XYZ* values as unsigned 8-bit integers. The ICC encoding convention is illustrated by this table.

Value (X, Y, or Z)	uint16 Value
0.0	0
1.0	32768
$1.0 + (32767/32768)$	65535

Class Support `xyz` is a uint16 or double array that must be real and nonsparse. `xyz16` is uint8.

Example Convert *XYZ* values to uint16 encoding.

```
xyz2uint16([0.1 0.5 1.0])
ans =
```

```
3277 16384 32768
```

See Also `applycform`, `lab2double`, `lab2uint16`, `lab2uint8`, `makecform`, `whitepoint`, `xyz2double`

Purpose Convert YCbCr values to RGB color space

Syntax
`rgbmap = ycbcr2rgb(ycbcrmap)`
`RGB = ycbcr2rgb(YCBCR)`

Description `rgbmap = ycbcr2rgb(ycbcrmap)` converts the YCbCr values in the colormap `ycbcrmap` to the RGB color space. If `ycbcrmap` is *m*-by-3 and contains the YCbCr luminance (*Y*) and chrominance (*C_b* and *C_r*) color values as columns, then `rgbmap` is returned as an *m*-by-3 matrix that contains the red, green, and blue values equivalent to those colors.

`RGB = ycbcr2rgb(YCBCR)` converts the YCbCr image `YCBCR` to the equivalent truecolor image `RGB`.

Class Support If the input is a YCbCr image, it can be of class `uint8`, `uint16`, or `double`; the output image is of the same class as the input image. If the input is a colormap, the input and output colormaps are both of class `double`.

Example Convert image from RGB space to YCbCr space and back.

```
rgb = imread('board.tif');  
ycbcr = rgb2ycbcr(rgb);  
rgb2 = ycbcr2rgb(ycbcr);
```

See Also `ntsc2rgb`, `rgb2ntsc`, `rgb2ycbcr`

zoom

Purpose

Zoom in and out on an image

zoom is a MATLAB function. To get help for this function, select **MATLAB Help** from the Help menu and view the online function reference page.

A

- adapthisteq 15-21
 - example 10-43
- adaptive filters 10-52
 - definition 10-2
 - two-dimensional adaptive noise filtering 15-527
- affine transformations
 - definition 6-12
 - using `imtransform` 5-11
- aliasing
 - definition 5-2, 5-7
- alpha channel 13-4
- analyzing images
 - edge detection 15-132
 - intensity profiles 15-330
 - pixel values 15-316
 - quadtree decomposition 15-459
- antialiasing 5-7
 - definition 5-2
- `applycform` 15-24
- `applylut` 15-25
 - example 9-45
 - modifying text in an image 9-44
- approximation
 - definition 13-2
 - in indexed images 13-12
 - of an image background 1-11
 - of colors 13-7
- area
 - of binary images 9-42
 - of image regions 10-12
- arrays
 - storing images 2-2

- averaging filter
 - creating 15-159
 - example 7-7
- `axes2pix` 15-27

B

- background
 - of a binary image 9-2
 - of an intensity image 1-11
- `basicimageinfo` 15-245
- `bestblk` 15-28
- bicubic interpolation 5-3
- bilinear interpolation 5-3
- binary images 2-4
 - applying a lookup table 15-25
 - calculating the Euler number 9-43
 - changing the display colors of 3-48
 - constructing lookup tables 15-418
 - displaying 3-47
 - feature measurement 9-40
 - finding area 15-33
 - flood fill operation 9-26
 - image area 9-42
 - labeling connected components 15-51
 - labeling connected components labeling 9-40
 - lookup table operations 9-44
 - morphological operations 9-4
 - object selection 15-66
 - packing 9-3
 - perimeter determination 9-17
 - selecting objects in 9-42
 - using neighborhood operations 15-25
- binary masks

- creating 11-3
 - definition 11-2
 - bit depth
 - 1-bit images 2-17
 - screen bit depth 13-3
 - blind deconvolution algorithm
 - used for deblurring 12-16
 - blkproc 15-29
 - specifying overlap 14-10
 - using 14-9
 - block processing 14-2
 - block size 15-28
 - column processing 14-11
 - distinct blocks 14-8
 - padding borders in block processing 14-5
 - sliding neighborhoods 14-4
 - using `nlfilt` 15-433
 - border padding
 - definition 14-2
 - in block processing 14-5
 - border replication
 - in image processing 7-12
 - boundary padding
 - See* border padding
 - boundary ringing
 - in image deblurring 12-23
 - boundary tracing 10-15
 - bounding box
 - finding for a region 10-12
 - brightness adjustment 10-38
 - bwarea 15-27
 - bwarea 15-33
 - using 9-42
 - bwareaopen 15-35
 - bwboundaries 15-37
 - bwdist 15-41
 - using 9-37
 - bweuler 15-47
 - bwhitmiss 15-49
 - bwlabel 15-51
 - bwlabeln 15-54
 - bwmorph 15-57
 - skeletonization example 9-17
 - bwpack 15-62
 - bwperim 15-64
 - bwselect 15-66
 - bwtraceboundary 15-68
 - using 10-16
 - bwulterode 15-71
 - bwunpack 15-73
- ## C
- camera read-out noise 12-12
 - Canny edge detector 15-133
 - using the edge function 10-13
 - center of mass
 - calculating for region 10-12
 - center pixel
 - calculating 14-4
 - definition 14-2
 - checkerboard 15-74
 - chrominance
 - in CIE color spaces 13-16
 - in NTSC color space 13-23
 - in YCbCr color space 13-24
 - CIE color spaces 13-15
 - CIELAB color space 13-15
 - closing 15-57
 - morphology 9-14
 - cmpermute 15-76
 - cmunique 15-77
 - col2im 15-79
 - colfilt 15-80

- example 14-12, 14-14
- using 14-11
- color
 - approximation 13-7
 - dithering 13-13
 - quantization 13-7
 - quantization performed by `rgb2ind` 15-483
 - reducing number of colors 13-6
- color approximation
 - performed by `rgb2ind` 15-483
- color cube
 - a description of 13-7
 - quantization of 13-8
- color planes 13-9, 13-25
 - of an HSV image 13-25
 - of an RGB image 2-9
- color reduction 13-6
- color spaces
 - converting among 13-15
 - converting between 15-436, 15-485
 - data encodings 13-18
 - device-independent color spaces 13-16
 - HSV 13-24
 - NTSC 13-23, 15-436, 15-485
 - YCbCr 13-24
 - YIQ 13-23
- colorbar 3-51
- colorcube 13-12
- colormap mapping 13-11
- colormaps
 - brightening 15-32
 - creating a colormap using `colorcube` 13-12
 - darkening 15-32
 - rearranging colors in 15-76
 - removing duplicate entries in 15-77
- column processing 15-80
 - definition 14-2
 - in neighborhood operations 14-11
 - reshaping columns into blocks 15-79
- composite transformations 5-11
- conformal transformations 6-12
- `conndef` 15-83
- connected components
 - definition 9-2
 - labeling 9-40
 - using `bwlabel` 15-51
- connectivity
 - definition 9-2
 - overview 9-23
 - specifying custom 9-25
- constant component
 - See* zero-frequency component
- contour plots
 - text labels 10-10
- contours
 - definition 10-2
- contrast adjustment
 - decreasing contrast 10-38
 - increasing contrast 10-36
 - specifying limits automatically 10-39
- contrast stretching
 - with decorrelation stretching 10-47
 - See also* contrast adjustment
- contrast-limited adaptive histogram equalization (CLAHE) 10-43
- Control Point Selection Tool
 - appearance of control point symbols 6-26
 - changing view of images 6-17
 - saving a session 6-30
 - saving control points 6-29
 - specifying control points 6-22
 - starting 6-15
 - using 6-14
 - using point prediction 6-24

- control points
 - appearance of 6-26
 - prediction 6-24
 - saving 6-29
 - selecting 6-14
 - specifying 6-22
 - conv2
 - compared to `imfilter` 7-14
 - conversions between image types 2-11
 - convmtx2 15-86
 - convn
 - compared to `imfilter` 7-14
 - convolution
 - convolution matrix 15-86
 - definition 7-4
 - Fourier transform and 8-12
 - two-dimensional 7-4
 - with `imfilter` 7-9
 - convolution kernel
 - definition 7-4
 - coordinate systems
 - pixel coordinates 2-30
 - spatial coordinates 2-31
 - corr2 15-88
 - calculating summary statistics 10-12
 - correlation
 - definition 7-6
 - Fourier transform 8-13
 - with `imfilter` 7-9
 - correlation coefficient 15-88
 - correlation kernel
 - definition 7-6
 - cp2tform 15-89
 - using 6-12
 - cpcorr 15-98
 - example 6-31
 - cpselect 15-100
 - using 6-7
 - cpstruct2pairs 15-102
 - cropping an image 5-10
 - cross-correlation
 - improving control point selection 6-31
- D**
- damping
 - for noise reduction 12-12
 - data types
 - converting between 2-19
 - double-precision (`double`) 2-19
 - in image filtering 7-7
 - DC component
 - See* zero-frequency component
 - dct2 15-103
 - equations 8-17
 - dctmtx 15-106
 - using 8-18
 - deblurring
 - avoiding boundary ringing 12-23
 - conceptual model 12-3
 - overview 12-3
 - overview of functions 12-6
 - use of frequency domain 12-22
 - using the blind deconvolution algorithm 12-16
 - using the Lucy-Richardson algorithm 12-11
 - with regularized filter 12-9
 - with the Wiener filter 12-7
 - decomposition of structuring elements
 - example 9-9
 - getting sequence 15-185
 - deconvblind 15-107
 - example 12-16
 - deconvlucy 15-110
 - example 12-11

- deconvreg 15-113
 - example 12-9
 - deconvwnr 15-115
 - example 12-7
 - decorrelation stretching 10-44
 - See also* contrast adjustment
 - decorrstretch 15-117
 - detail rectangle
 - in Control Point Selection Tool 6-16
 - DICOM files
 - reading and writing 2-21
 - DICOM unique identifier
 - generating 2-27
 - dicomanon 15-118
 - dicomdict 15-120
 - dicominfo 15-121
 - dicomread 15-122
 - dicomuid 15-124
 - dicomwrite 15-125
 - Digital Imaging and Communications in Medicine
 - See* DICOM files
 - dilation 9-4
 - grayscale 15-437
 - discrete cosine transform 8-17
 - image compression 8-19
 - transform matrix 8-18, 15-106
 - discrete Fourier transform 8-8
 - discrete transform
 - definition 8-2
 - display depth 13-3
 - See* screen color resolution
 - display techniques
 - displaying images at true size 15-516
 - multiple images 15-506
 - displaying images
 - adding a colorbar 3-51
 - binary 3-47
 - binary images with different colors 3-48
 - comparison of functions 3-2
 - indexed images 3-45
 - initial size 3-5
 - intensity images 3-46
 - multiple images in the same figure window 3-42
 - texture mapping 3-55
 - toolbox preferences for 3-58
 - unconventional ranges of data 3-46
 - using imshow 3-4
 - using the Image Tool 3-8, 3-11
 - distance
 - between pixels 10-3
 - Euclidean 10-3
 - distance transform 9-37
 - distinct block operations 14-8
 - definition 14-2
 - overlap 14-9
 - zero padding 14-8
 - dither 15-129
 - dithering 13-13
 - example 13-13
- E**
- edge 15-132
 - example 10-14
 - edge detection 10-13
 - Canny method 10-13
 - example 10-14
 - methods 15-132
 - Sobel method 10-13
 - edges
 - definition 10-2
 - edgetaper 15-137

- avoiding boundary ringing 12-23
- 8-bit image files
 - creating 2-18
 - reading 2-15
- enhancing images
 - decorrelation stretching 10-44
 - intensity adjustment 10-37
 - noise removal 10-49
- entropy 15-138
- entropyfilt 15-139
- erosion 9-4
 - grayscale 15-437
- Euclidean distance 10-3, 15-451
- Euler number
 - calculating 9-43
- exploring images
 - using Image Tool 3-16
- F**
- fan2para 15-141
- fanbeam 15-145
 - using 8-36
- fan-beam projection data
 - arc geometry 8-37
 - computing 8-36
 - line geometry 8-38
 - reconstructing image from 8-39
- fast Fourier transform 8-8
 - zero padding 8-10
 - See also* Fourier transform
- feature measurement 1-19
 - area 10-12
 - binary images 9-40
 - bounding box 10-12
 - center of mass 10-12
- fft 8-8
 - fft2 8-8
 - example 8-9
 - using 8-11
 - fftn 8-8
 - fftshift
 - example 7-18
 - using 8-11
- files
 - displaying images from disk 3-4
- filling a region 11-9
 - definition 11-2
- filling holes in images 9-28
- filter design 7-17
 - frequency sampling method 7-19, 15-156
 - frequency transformation method 7-18, 15-164
 - windowing method 7-20, 15-167, 15-171
- filter2
 - compared to imfilter 7-14
 - example 3-47, 10-51
- filtering
 - a region 11-6
 - masked filtering 11-6
 - regions 11-2
- filters
 - adaptive 10-52, 15-527
 - averaging 15-159
 - binary masks 11-6
 - computing frequency response 8-11
 - designing 7-17
 - finite impulse response (FIR) 7-17
 - frequency response 7-23
 - imfilter 7-7
 - Infinite Impulse Response (IIR) 7-18
 - Laplacian of Gaussian 15-159
 - linear 7-4
 - median 10-49, 15-429
 - multidimensional 7-13

- order-statistic 15-437
 - predefined types 15-159
 - Prewitt 15-159
 - Sobel 15-159
 - unsharp 15-159
 - unsharp masking 7-15
 - findbounds 15-151
 - FIR filters 7-17
 - transforming from one-dimensional to two-dimensional 7-18
 - flat-field correction 12-12
 - fliptform 15-152
 - flood-fill operation 9-26
 - foreground
 - of a binary image 9-2
 - 4-bit image files 2-17
 - Fourier transform 8-3
 - applications of the Fourier transform 8-11
 - centering the zero-frequency coefficient 8-11
 - computing frequency response 8-11
 - convolution and 8-12
 - correlation 8-13
 - DFT coefficients 8-9
 - examples of transform on simple shapes 8-7
 - fast convolution with 8-12
 - for performing correlation 8-13
 - frequency domain 8-3
 - increasing resolution 8-10
 - padding before computation 8-10
 - two-dimensional 8-3
 - zero-frequency component 8-3
 - freqspace 7-22
 - example 7-19
 - frequency domain 8-3
 - definition 8-2
 - frequency response
 - computing 7-23, 8-11, 15-154
 - desired response matrix 7-22
 - frequency sampling method (filter design) 7-19
 - using fsamp2 15-156
 - frequency transformation method (filter design) 7-18
 - using ftrans2 15-164
 - freqz
 - example 7-18
 - freqz2 15-154
 - computing frequency response of linear filters 8-11
 - example 7-20
 - fsamp2 15-156
 - example 7-19
 - fspecial 15-159
 - creating predefined filters 7-15
 - ftrans2 15-164
 - example 7-18
 - function functions
 - using 14-9
 - fwind1 15-167
 - example 7-21
 - fwind2 15-171
 - example 7-20
- ## G
- gamma correction 10-39
 - Gaussian convolution kernel
 - frequency response of 8-12
 - Gaussian filter 15-159
 - Gaussian noise 10-52
 - adding 15-309
 - geocoded images 6-6
 - geometric operations
 - cropping 5-10
 - definition 5-2

- interpolation 5-3
 - resizing 5-5
 - georegistered images 6-6
 - getheight 15-175
 - getimage 15-176
 - example 3-5, 3-14
 - getimagemodel 15-178
 - getline 15-179
 - getneighbors 15-180
 - getnhood 15-181
 - getpts 15-182
 - getrangefromclass 15-183
 - getrect 15-184
 - getsequence 15-185
 - graphics card 13-4
 - graphics file formats
 - converting from one format to another 2-19
 - writing data 2-16
 - gray2ind 15-186
 - grayscale morphological operations 15-437
 - grayscale 15-193
 - graythresh 15-194
 - thresholding image values 1-15
- ## H
- histeq 15-195
 - example 10-42
 - increase contrast example 10-41, 10-44
 - histogram equalization 10-41, 15-195
 - histograms 10-10
 - definition 10-2
 - holes
 - filling 9-26
 - tracing boundaries 10-15
 - hough 15-198
 - using 10-21
 - houghlines 15-200
 - using 10-22
 - houghpeaks 15-203
 - using 10-22
 - HSV color space 13-24
 - color planes of 13-25
 - hsv2rgb 13-24
 - http
 - [//www.mathworks.com/products/image/related.html](http://www.mathworks.com/products/image/related.html) 1-3
 - hue
 - in HSV color space 13-24
 - in NTSC color space 13-23
- ## I
- ICC profiles
 - processing 13-19
 - iccread 15-206
 - iccwrite 15-209
 - idct2 15-210
 - ifanbeam 15-212
 - using 8-39
 - ifft 8-8
 - ifft2 8-8
 - ifftn 8-8
 - IIR filters 7-18
 - im2bw 15-217
 - im2col 15-219
 - im2double 15-221
 - example 14-14
 - im2int16 15-222, 15-225
 - im2java2d 15-224
 - im2single 15-225
 - im2uint16 15-226
 - im2uint8 15-227
 - imabsdiff 15-228

- imadd 15-230
- imadjust 15-232
 - brightening example 10-38
 - example 10-37
 - gamma correction and 10-39
 - gamma correction example 10-40
 - increase contrast example 10-37
 - setting limits automatically 10-39
- image analysis
 - contour plots 10-9
 - edge detection 10-13
 - histograms 10-10
 - intensity profiles 10-5
 - overview 10-13
 - pixel values 10-3
 - quadtree decomposition 10-24
 - summary statistics 10-12
- image area (binary images) 15-33
- image arithmetic
 - combining functions 2-29
 - overview 2-28
 - truncation rules 2-28
- image editing 11-9
- image filtering
 - data types 7-7
 - unsharp masking 7-15
 - with `imfilter` 7-7
- Image Information tool
 - using 3-29
- image profiles
 - definition 10-2
- image properties
 - definition 10-2
 - set by `imshow` 3-2
- image registration
 - fine-tuning point placement 6-31
 - overview 6-4
 - procedure 6-4
 - selecting control points 6-14
 - specifying control point pairs 6-22
 - types of transformations 6-12
 - using control point prediction 6-24
- image rotation 5-8
- Image Tool
 - compared to `imshow` 3-2
 - controlling initial magnification 3-11
 - exploring images 3-16
 - overview 3-8
- image transformations
 - affine 6-12
 - custom 5-11
 - local weighted mean 6-13
 - piecewise linear 6-13
 - polynomial 6-13
 - projective 6-12
 - supported by `cp2tform` 6-12
 - types of 6-12
 - using `imtransform` 5-11
- image types 2-3
 - binary 2-4
 - converting 2-11
 - indexed 2-4
 - intensity 2-5
 - interpolation and 5-4
 - multiframe images 2-14
 - RGB 2-6
 - supported by the toolbox 2-3
 - See also* indexed, intensity, binary, RGB, multiframe
- Image Viewer
 - extended example 1-10
 - panning images 3-19
 - specifying image magnification 3-20
 - using the Pixel Region tool 3-22

- zoom tools 3-19
- imageinfo 15-235
- imagemodel 15-238
- images
 - analyzing 10-3
 - arithmetic operations 2-28
 - causes of blurring 12-3
 - creating movies 15-306
 - data types 2-19
 - displaying multiple images 3-41, 15-506
 - displaying multiple images in the same figure window 3-42
 - feature measurement 1-19
 - filling holes in 9-28
 - finding image minima and maxima 9-29
 - getting data from axes 15-176
 - how MATLAB stores 2-2
 - image types 2-3
 - improving contrast 1-14
 - reducing number of colors 13-6
 - registering 6-4
 - restoring blurred images 12-3
 - returning information about 2-19
 - statistical analysis of 1-20
 - using imshow 3-4
 - viewing as a surface plot 1-12
- imapprox 15-243
 - example 13-12
- imbothat 15-248
- imclearborder 15-250
- imclose 15-253
 - using 9-14
- imcomplement 15-255
- imcontour 15-257
 - example 10-10
- imcontrast 15-259
- imcrop 15-261
 - example 5-10
- imdilate 15-264
- imdisplayrange 15-268
- imdivide 15-270
- imerode 15-272
 - closure example 9-15
- imextendedmax 15-275
 - example 9-31
- imextendedmin 15-277
- imfill 15-279
 - example 9-28
- imfilter 15-283
 - compared to other filtering functions 7-14
 - convolution option 7-9
 - correlation option 7-9
 - padding options 7-10
 - using 7-7
- imfinfo 2-19
 - example 2-17
- imgca 15-287
- imgcf 15-288
- imgetfile 15-289
- imhandles 15-290
- imhist 15-291
 - example 10-11, 10-36, 10-37
- imhmax 15-293
- imhmin 15-295
- imimposemin 15-297
- imlincomb 15-301
 - example 2-29
- immagbox 15-304
- immovie 15-306
 - example 3-54
- immultiply 15-307
- imnoise 15-309
 - example 10-52
 - salt & pepper example 10-50

- imopen 15-312
 - using 9-14
- imoverview 15-314
- imoverviewpanel 15-315
- impixel 15-316
 - example 10-4
- impixelinfo 15-319
- impixelinfoval 15-321
- impixelregion 15-323
- impixelregionpanel 15-325
- impositionrect 15-326
- improfile 15-330
 - example 10-7
 - grayscale example 10-6
- imread 2-15
 - example for multiframe image 2-16
- imreconstruct 15-334
 - example 9-20
- imregionalmax 15-336
- imregionalmin 15-338
- imresize 15-340
 - using 5-5
- imrotate 15-342
 - example 5-8, 5-9
 - using 5-8
- imscrollpanel 15-344
- imshow 15-349
 - compared to Image Tool 3-2
 - displaying images 3-4
 - displaying unconventional range data 3-46
 - example for binary images 3-48
 - example for intensity images 3-46
 - example for RGB images 3-50
 - preferences for 3-58
 - used with indexed images 3-45
- imsubtract 15-353
- imtool 15-355
 - compared to imshow 3-2
 - overview 3-8
- imtophat 15-359
- imtransform 15-361
 - using 5-11
- imview 15-367
 - displaying unconventional range data 3-46
 - used with indexed images 3-45
- imwrite
 - example 2-17
- ind2gray 15-369
- ind2rgb 15-370
- indexed images 2-4
 - converting from intensity 15-186
 - converting from RGB 15-482
 - reducing number of colors in 13-6, 13-12
- infinite impulse response (IIR) filter 7-18
- inline 14-9
- Intel Performance Primitives Library 15-372
- intensity adjustment 10-37
 - gamma correction 10-39
 - histogram equalization 10-41
 - specifying limits automatically 10-39
 - See also* contrast adjustment
- intensity images 2-5
 - converting from matrices 15-427
 - converting from RGB 15-480
 - converting to indexed 15-186
 - displaying 3-46
 - flood-fill operation 9-26
- intensity profiles 10-5, 15-330
- interpolation 5-3
 - bicubic 5-3
 - bilinear 5-3
 - default 5-4
 - definition 5-2
 - intensity profiles 10-5

- nearest-neighbor 5-3
 - of binary images 5-4
 - of indexed images 5-4
 - of RGB images 5-4
 - tradeoffs between methods 5-3
 - within a region of interest 11-9
- intlut 15-371
- inverse Radon transform 8-29
 - example 8-34
 - filtered backprojection algorithm 8-31
- inverse transform
 - definition 8-2
- ipp1 15-372
- iptaddcallback 15-373
- iptcheckconn 15-374
- iptcheckhandle 15-375
- iptcheckinput 15-377
- iptcheckmap 15-379
- iptcheckmargin 15-380
- iptcheckstrs 15-381
- iptcondir 15-386
- iptdemos 15-383
- iptgetapi 15-384
- iptgetpref 3-59, 15-385
- iptnum2ordinal 15-387
- iptremovecallback 15-388
- iptsetpref 15-389
- iptwindowalign 15-393
- iradon 8-29, 15-395
 - example 8-29
- isbw 15-398
- isflat 15-399
- isgray 15-400
- isicc 15-401
- isind 15-402
- isrgb 15-403

J

- JPEG compression
 - discrete cosine transform and 8-19

L

- lab2double 15-404
- lab2uint16 15-406
- lab2uint8 15-408
- label matrix
 - creating 9-40
 - viewing as pseudocolor image 1-18, 9-41
- label2rgb 15-410
- labeling
 - connected components 9-40
 - levels of contours 10-10
- Laplacian of Gaussian edge detector 15-133
- Laplacian of Gaussian filter 15-159
- line detection 8-27
- line segment
 - pixel values along 10-5
- linear conformal transformations 6-12
- linear filtering 7-4, 14-5
 - convolution 7-4
 - filter design 7-17
 - FIR filters 7-17
 - IIR filters 7-18
 - noise removal and 10-49
- local weighted mean transformations 6-13
- lookup tables
 - constructing 15-418
 - using 9-44
- Lucy-Richardson algorithm
 - used for deblurring 12-11
- luminance
 - in NTSC color space 13-23
 - in YCbCr color space 13-24

M

- magnification
 - specifying in Image Tool 3-11
- magnifying
 - See* resizing images
- makecform 15-412
- makelut 15-418
 - example 9-44
- makeresampler 15-420
- maketform 15-424
- marker image
 - creating 9-34
 - definition 9-19
- mask image
 - definition 9-19
- masked filtering 11-6
 - definition 11-2
- mat2gray 15-427
- matrices
 - converting to intensity images 15-427
 - storing images in 2-2
- maxima
 - finding in images 9-29
 - imposing 9-33
 - suppressing 9-31
- McClellan transform 15-164
- mean2 10-12, 15-428
- medfilt2 15-429
 - example 10-51
 - using 10-50
- median filtering 10-49, 15-429
- minima
 - finding in images 9-29
 - imposing 9-33
 - suppressing 9-31
- minimum variance quantization 13-9
- moiré patterns 5-7
- montage 15-431
 - example 3-52, 3-53
- morphological operations 9-4
 - closing 15-57
 - diagonal fill 15-58
 - dilation 9-4
 - erosion 9-4
 - grayscale 15-437
 - opening 9-14
 - overview 9-1
 - predefined operations 9-16
 - shrinking objects 15-59
 - skeletonization 9-17
- morphological reconstruction
 - finding peaks and valleys 9-29
 - overview 9-19
- morphology
 - definition 9-2
 - See also* morphological reconstruction
- movies
 - creating from images 3-54, 15-306
 - playing 3-54
- multidimensional filters 7-13
- multiframe images
 - about 2-14
 - displaying 15-431
 - limitations 2-14
- multilevel thresholding 15-193

N

- nearest-neighbor interpolation 5-3
- neighborhood operations
 - definition 14-2
- neighborhoods
 - binary image operations 15-25
 - definition 9-2

- neighborhood operations 14-2
- `nlfilter` 14-6, 15-433
 - example 14-6
- noise
 - adding to an image 15-309
 - definition 10-2
- noise amplification
 - reducing 12-12
- noise removal 10-49
 - adaptive filtering (Weiner) and 10-52
 - Gaussian noise 10-52
 - grain noise 10-49
 - linear filtering and 10-49
 - median filter and 10-50
 - salt and pepper noise 10-50
- nonlinear filtering 14-5
- normalized cross-correlation 6-31
- `normxcorr2` 15-434
- NTSC color space 13-23, 15-436, 15-485
- `ntsc2rgb` 13-23, 15-436

O

- object selection 15-66
- objects
 - tracing boundaries 10-15
- observed image
 - in image registration 6-14
- 1-bit image files 2-17
- opening
 - morphology 9-14
- optical transfer function (OTF)
 - definition 12-2
- order-statistic filtering 15-437
- `ordfilt2` 15-437
- orthonormal matrix 8-19
- orthophoto

- defined 6-6
- orthorectified image 6-6
- OTF
 - See* optical transfer function
- `otf2psf` 15-439
 - use of 12-22
- overlap
 - in block operations 14-2
 - in distinct block operations 14-8
- Overview window
 - using 3-16

P

- packed binary image
 - definition 9-3
- `padarray` 15-440
- padding borders
 - in block processing 14-5
 - options with `imfilter` 7-10
- panning images
 - using the Image Viewer 3-19
- `para2fan` 15-444
- parallel beam projections 8-30
- perimeter determination
 - in binary images 9-17
- phantom 8-32, 15-448
- piecewise linear transformations 6-13
- Pixel Region tool
 - using 3-22
- pixel values
 - along a line segment 10-5
 - returning using a mouse 10-3
 - using `impixel` 15-316
 - using `pixval` 15-451
 - using the Pixel Region tool 3-22
- pixels

- correcting for bad pixels 12-12
 - defining connectivity 9-23
 - definition 2-2
 - displaying coordinates of 10-3
 - Euclidean distance between 10-3
 - returning coordinates of 10-4
 - selecting 10-3
- pixval 15-451
- using 10-3
- PNG (Portable Network Graphics) files
- writing as 16-bit 2-17
- point mapping
- for image registration 6-4
- point spread function (PSF)
- definition 12-2
 - importance of in deblurring 12-4
- Poisson noise
- adding 15-309
- poly2mask 15-453
- polygon
- pixels inside 11-3
 - selecting a polygonal region of interest 11-3
- polynomial transformations 6-13
- predicting control point locations
- in image registration 6-24
- preferences
- Image Processing Toolbox display preferences 3-58
 - ImshowAxesVisible 3-58
 - ImshowBorder 3-58
 - ImshowTruesize 3-58
- Prewitt edge detector 15-133
- Prewitt filter 15-159
- profiles
- reading ICC color profiles 13-19
- projections
- parallel beam 8-30
 - projective transformations 5-11, 6-12
- PSF
- See* point spread function
- psf2otf 15-458
- use of 12-22
- ## Q
- qtdecomp 15-459
- example 10-24
- qtgetblk 15-462
- qtsetblk 15-463
- quadtree decomposition 10-24, 15-459
- definition 10-2
 - getting block values 15-462
 - setting block values 15-463
- quantization 13-7
- minimum variance 13-9
 - performed by rgb2ind 15-483
 - tradeoffs between using minimum variance and uniform quantization methods 13-11
- ## R
- radon 15-464
- example 8-24
- Radon transform
- center pixel 8-23
 - description of 8-21
 - example 8-32
 - inverse Radon transform 8-29
 - line detection example 8-27
 - of the Shepp-Logan Head phantom 8-33
 - relationship to Hough transform 8-27
- rangefilt 15-469
- rank filtering 10-50
- See also* order-statistic filtering

- read-out noise
 - correcting 12-12
- real orthonormal matrix 8-19
- reconstruction
 - morphological 9-19
- reference image
 - in image registration 6-14
- reflect 15-471
- region labeling 9-40
- region of interest
 - based on color or intensity 11-5
 - binary masks 11-3
 - definition 11-2
 - filling 11-9
 - filtering 11-6
 - polygonal 11-3
 - selecting 11-3
 - using arbitrary binary masks 11-4
- region property measurement 10-12
- regional maxima
 - definition 9-3, 9-29
 - imposing 9-33
 - suppressing 9-31
- regional minima
 - definition 9-3, 9-29
 - imposing 9-33
 - suppressing 9-31
- regionprops 15-472
 - using 1-19, 10-12
- registering an image 6-4
- regularized filter
 - used for deblurring 12-9
- replication
 - to avoid border effect 7-12
- resizing images 5-5
 - antialiasing 5-7
- resolution
 - screen color resolution 13-3
 - See also* bit depth 13-3
- RGB color cube
 - description of 13-7
 - quantization of 13-8
- RGB images 2-6
 - converting to indexed 15-482
 - converting to intensity 15-480
 - displaying 3-49
 - measuring the intensities of each color plane 10-7
 - reducing number of colors 13-6
- rgb2gray 15-480
- rgb2hsv 13-24
 - converting RGB to HSV color space 13-25
 - example 13-25
- rgb2ind 15-482
 - colormap mapping example 13-12
 - example 13-12, 13-13
 - in minimum variance quantization 13-9
 - minimum variance quantization example 13-10
 - specifying a colormap to use 13-11
 - uniform quantization example 13-9
 - used in dithering 13-13
 - using 13-7
- rgb2ntsc 15-485
 - example 13-23
- rgb2ycbcr 15-486
 - example 13-24
- Richardson-Lucy algorithm
 - See* Lucy-Richardson
- ringing
 - in image deblurring 12-23
- Roberts edge detector 15-133
- roicolor 11-5, 15-488
- roifill 15-489
 - example 11-9

- roifilt2 15-491
 - contrast example 11-6
 - inline example 11-8
- roipoly 15-493
 - example 11-3
- rotation
 - of images 5-8

S

- salt and pepper noise 10-50
 - adding 15-309
- sampling
 - handling undersampled images 12-13
- saturation
 - in HSV color space 13-24
 - in NTSC color space 13-23
- screen bit depth 13-3
 - definition 13-2
- screen color resolution 13-3
 - definition 13-2
- ScreenDepth 13-3
- Shepp-Logan head phantom 8-32
- shrinking
 - See* resizing images
- Signal Processing Toolbox
 - hamming function 7-21
- 16-bit image files
 - creating 2-18
 - reading 2-15
- skeletonization 9-17
- sliding neighborhood operations 14-4
 - center pixel in 14-4
 - padding in 14-5
- Sobel edge detector 15-132
- Sobel filter 15-159
- spatial coordinates 2-31

- spatial domain
 - definition 8-2
- speckle noise
 - adding 15-309
- statistical properties
 - mean 15-428
 - of image objects 1-20
 - standard deviation 15-495
- std2 10-12, 15-495
- stdfilt 15-496
- storage classes
 - converting between 2-13
- strel 15-497
- stretchlim 15-504
 - adjusting image contrast 1-14
 - using 10-39
- structuring elements 9-7
 - creating 9-8
 - decomposition of 9-9
 - decomposition sequence 15-185
 - definition 9-3
 - determining composition 15-266
- subimage 3-43, 15-506
- subplot 3-42
- subtraction
 - of one image from another 1-13
- surf
 - viewing images 1-12

T

- template matching 8-13
- texture mapping 3-55
- tformarray 15-507
- tformfwd 15-511
- tforminv 15-513
- thresholding

- to create a binary image 1-15
- to create indexed image from intensity image 15-193
- tomography 8-29
- tracing boundaries 10-15
- transformation matrix 7-18
- transforms 8-1
 - definition 8-2
 - discrete cosine 8-17, 15-103
 - discrete Fourier transform 8-8
 - Fourier 8-3
 - inverse Radon 8-29
 - Radon 8-21
 - two-dimensional Fourier transform 8-3
- translate 15-515
- transparency 13-4
- truesize 15-516
- truncation rules
 - for image arithmetic operators 2-28
- 24-bit image files 2-6

U

- uint16
 - storing images in 2-15
- uint8
 - storing images in 2-15
- uintlut 15-519
- undersampling
 - correcting 12-13
- uniform quantization
 - See quantization
- unsharp filter 15-159
- unsharp masking 7-15

W

- warp 3-55, 15-520
 - example 3-55
- watershed 15-522
- weight array
 - in deblurring 12-12
- whitepoint 15-526
- Wiener filter
 - deblurring with 12-7
- wiener2 15-527
 - adaptive filtering 10-52
 - using 10-52
- windowing method (filter design) 7-20, 15-167, 15-171

X

- X-ray absorption tomography 8-29
- XYZ color space 13-15
- xyz2double 15-529
- xyz2uint16 15-530

Y

- YCbCr color space 13-24
- ycbcr2rgb 15-531
 - using 13-24
- YIQ color space 13-23

Z

- zero padding 8-13
 - and the fast Fourier transform 8-10
 - image boundaries 7-10
- zero-cross edge detector 15-133
- zero-frequency component 8-3
- zooming

Control Point Selection Tool 6-19
with Image Viewer 3-19

