

Image Processing Toolbox™

Reference

R2012b

MATLAB®

How to Contact MathWorks



www.mathworks.com Web
comp.soft-sys.matlab Newsgroup
www.mathworks.com/contact_TS.html 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.
3 Apple Hill Drive
Natick, MA 01760-2098

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

Image Processing Toolbox™ Reference

© COPYRIGHT 1993–2012 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.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

Revision 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 13.SP1)
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)
March 2005	Online only	Revised for Version 5.0.2 (Release 14SP2)
September 2005	Online only	Revised for Version 5.1 (Release 14SP3)
March 2006	Online only	Revised for Version 5.2 (Release 2006a)
September 2006	Online only	Revised for Version 5.3 (Release 2006b)
March 2007	Online only	Revised for Version 5.4 (Release 2007a)
September 2007	Online only	Revised for Version 6.0 (Release 2007b)
March 2008	Online only	Revised for Version 6.1 (Release 2008a)
October 2008	Online only	Revised for Version 6.2 (Release 2008b)
March 2009	Online only	Revised for Version 6.3 (Release 2009a)
September 2009	Online only	Revised for Version 6.4 (Release 2009b)
March 2010	Online only	Revised for Version 7.0 (Release 2010a)
September 2010	Online only	Revised for Version 7.1 (Release 2010b)
April 2011	Online only	Revised for Version 7.2 (Release 2011a)
September 2011	Online only	Revised for Version 7.3 (Release 2011b)
March 2012	Online only	Revised for Version 8.0 (Release 2012a)
September 2012	Online only	Revised for Version 8.1 (Release 2012b)

Function Reference

Image Display and Exploration	1-2
Image Display and Exploration	1-2
Image File I/O	1-2
Image Types and Type Conversions	1-3
GUI Tools	1-5
Modular Interactive Tools	1-5
Navigational Tools for Image Scroll Panel	1-5
Utilities for Interactive Tools	1-6
Spatial Transformation and Image Registration	1-8
Spatial Transformations	1-8
Image Registration	1-9
Image Analysis and Statistics	1-10
Image Analysis	1-10
Texture Analysis	1-11
Pixel Values and Statistics	1-11
Image Arithmetic	1-12
Image Enhancement and Restoration	1-13
Image Enhancement	1-13
Image Restoration (Deblurring)	1-13
Linear Filtering and Transforms	1-15
Linear Filtering	1-15
Linear 2-D Filter Design	1-15
Image Transforms	1-15
Morphological Operations	1-17
Intensity and Binary Images	1-17
Binary Images	1-18

Structuring Element Creation and Manipulation	1-19
ROI-Based, Neighborhood, and Block Processing	1-20
ROI-Based Processing	1-20
Neighborhood and Block Processing	1-20
Colormaps and Color Space	1-22
Color Space Conversions	1-22
Utilities	1-24
Preferences	1-24
Validation	1-24
Mouse	1-25
Array Operations	1-25
Examples	1-25
Performance	1-25

Class Reference

2

Image Input and Output	2-2
ImageAdapter	2-2
Image Registration	2-3
registration.metric.MattesMutualInformation	2-3
registration.metric.MeanSquares,	2-3
registration.optimizer.RegularStepGradientDescent	2-3
registration.optimizer.OnePlusOneEvolutionary	2-3

Functions — Alphabetical List

3

Index

Function Reference

Image Display and Exploration (p. 1-2)	Display, import, and export images
GUI Tools (p. 1-5)	Modular interactive tools and associated utility functions.
Spatial Transformation and Image Registration (p. 1-8)	Spatial transformation and image registration
Image Analysis and Statistics (p. 1-10)	Image analysis, texture analysis, view pixel values, and calculate image statistics
Image Arithmetic (p. 1-12)	Add, subtract, multiply, and divide images
Image Enhancement and Restoration (p. 1-13)	Enhance and restore images
Linear Filtering and Transforms (p. 1-15)	Linear filters, filter design, and image transforms
Morphological Operations (p. 1-17)	Morphological image processing
ROI-Based, Neighborhood, and Block Processing (p. 1-20)	ROI-based, neighborhood, and block operations
Colormaps and Color Space (p. 1-22)	Manipulate image color
Utilities (p. 1-24)	Array operations, examples, preferences and other toolbox utility functions

Image Display and Exploration

Image Display and Exploration (p. 1-2)	Display and explore images
Image File I/O (p. 1-2)	Import and export images
Image Types and Type Conversions (p. 1-3)	Convert between the various image types

Image Display and Exploration

immovie	Make movie from multiframe image
implay	Play movies, videos, or image sequences
imshow	Display image
imtool	Image Tool
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

analyze75info	Read metadata from header file of Analyze 7.5 data set
analyze75read	Read image data from image file of Analyze 7.5 data set
dicomanon	Anonymize DICOM file
dicomdict	Get or set active DICOM data dictionary

<code>dicominfo</code>	Read metadata from DICOM message
<code>dicomlookup</code>	Find attribute in DICOM data dictionary
<code>dicomread</code>	Read DICOM image
<code>dicomuid</code>	Generate DICOM unique identifier
<code>dicomwrite</code>	Write images as DICOM files
<code>hdrread</code>	Read high dynamic range (HDR) image
<code>hdrwrite</code>	Write Radiance high dynamic range (HDR) image file
<code>interfileinfo</code>	Read metadata from Interfile file
<code>interfileread</code>	Read images in Interfile format
<code>isrset</code>	Check if file is R-Set
<code>makehdr</code>	Create high dynamic range image
<code>nitfinfo</code>	Read metadata from National Imagery Transmission Format (NITF) file
<code>nitfread</code>	Read image from NITF file
<code>openrset</code>	Open R-Set file
<code>rsetwrite</code>	Create reduced resolution data set from image file
<code>tonemap</code>	Render high dynamic range image for viewing

Image Types and Type Conversions

<code>demosaic</code>	Convert Bayer pattern encoded image to truecolor image
<code>gray2ind</code>	Convert grayscale or binary image to indexed image

<code>grayslice</code>	Convert grayscale image to indexed image using multilevel thresholding
<code>graythresh</code>	Global image threshold using Otsu's method
<code>im2bw</code>	Convert image to binary image, based on threshold
<code>im2double</code>	Convert image to double precision
<code>im2int16</code>	Convert image to 16-bit signed integers
<code>im2java2d</code>	Convert image to Java buffered image
<code>im2single</code>	Convert image to single precision
<code>im2uint16</code>	Convert image to 16-bit unsigned integers
<code>im2uint8</code>	Convert image to 8-bit unsigned integers
<code>ind2gray</code>	Convert indexed image to grayscale image
<code>ind2rgb</code>	Convert indexed image to RGB image
<code>label2rgb</code>	Convert label matrix into RGB image
<code>mat2gray</code>	Convert matrix to grayscale image
<code>rgb2gray</code>	Convert RGB image or colormap to grayscale

GUI Tools

Modular Interactive Tools (p. 1-5)	Modular interactive tool creation functions
Navigational Tools for Image Scroll Panel (p. 1-5)	Modular interactive navigational tools
Utilities for Interactive Tools (p. 1-6)	Modular interactive tool utility functions

Modular Interactive Tools

imageinfo	Image Information tool
imcontrast	Adjust Contrast tool
imdisplayrange	Display Range tool
imdistline	Distance tool
impixelinfo	Pixel Information tool
impixelinfoval	Pixel Information tool without text label
impixelregion	Pixel Region tool
impixelregionpanel	Pixel Region tool panel

Navigational Tools for Image Scroll Panel

immagbox	Magnification box for scroll panel
imoverview	Overview tool for image displayed in scroll panel
imoverviewpanel	Overview tool panel for image displayed in scroll panel
imscrollpanel	Scroll panel for interactive image navigation

Utilities for Interactive Tools

<code>axes2pix</code>	Convert axes coordinates to pixel coordinates
<code>getimage</code>	Image data from axes
<code>getimagemodel</code>	Image model object from image object
<code>imattributes</code>	Information about image attributes
<code>imellipse</code>	Create draggable ellipse
<code>imfreehand</code>	Create draggable freehand region
<code>imgca</code>	Get handle to current axes containing image
<code>imgcf</code>	Get handle to current figure containing image
<code>imgetfile</code>	Open Image dialog box
<code>imhandles</code>	Get all image handles
<code>inline</code>	Create draggable, resizable line
<code>impoint</code>	Create draggable point
<code>impoly</code>	Create draggable, resizable polygon
<code>imrect</code>	Create draggable rectangle
<code>imroi</code>	Region-of-interest (ROI) base class
<code>iptaddcallback</code>	Add function handle to callback list
<code>iptcheckhandle</code>	Check validity of handle
<code>iptgetapi</code>	Get Application Programmer Interface (API) for handle
<code>iptGetPointerBehavior</code>	Retrieve pointer behavior from HG object
<code>ipticondir</code>	Directories containing IPT and MATLAB® icons
<code>iptPointerManager</code>	Create pointer manager in figure

<code>iptremovecallback</code>	Delete function handle from callback list
<code>iptSetPointerBehavior</code>	Store pointer behavior structure in Handle Graphics object
<code>iptwindowalign</code>	Align figure windows
<code>makeConstrainToRectFcn</code>	Create rectangularly bounded drag constraint function
<code>trueSize</code>	Adjust display size of image

Spatial Transformation and Image Registration

Spatial Transformations (p. 1-8)

Spatial transformation of images and multidimensional arrays

Image Registration (p. 1-9)

Align two images using control points

Spatial Transformations

checkerboard

Create checkerboard image

findbounds

Find output bounds for spatial transformation

fliptform

Flip input and output roles of TFORM structure

imcrop

Crop image

impyramid

Image pyramid reduction and expansion

imresize

Resize image

imrotate

Rotate image

imtransform

Apply 2-D spatial transformation to image

makeresampler

Create resampling structure

maketform

Create spatial transformation structure (TFORM)

tformarray

Apply spatial transformation to N-D array

tformfwd

Apply forward spatial transformation

tforminv

Apply inverse spatial transformation

Image Registration

<code>cp2tform</code>	Infer spatial transformation from control point pairs
<code>cpcorr</code>	Tune control-point locations using cross correlation
<code>cpselect</code>	Control Point Selection Tool
<code>cpstruct2pairs</code>	Convert CPSTRUCT to valid pairs of control points
<code>imfuse</code>	Composite of two images
<code>imregconfig</code>	Configurations for intensity-based registration
<code>imregister</code>	Intensity-based image registration
<code>imshowpair</code>	Compare differences between images
<code>normxcorr2</code>	Normalized 2-D cross-correlation
<code>registration.metric.MattesMutualInformation</code>	Mattes mutual information metric configuration object
<code>registration.metric.MeanSquares</code>	Mean square error metric configuration object
<code>registration.optimizer.OnePlusOneEvolutionary</code>	One-plus-one evolutionary optimizer configuration object
<code>registration.optimizer.RegularStepGradientDescent</code>	Regular step gradient descent optimizer configuration object

Image Analysis and Statistics

Image Analysis (p. 1-10)	Trace boundaries, detect edges, and perform quadtree decomposition
Texture Analysis (p. 1-11)	Entropy, range, and standard deviation filtering; gray-level co-occurrence matrix
Pixel Values and Statistics (p. 1-11)	Create histograms, contour plots, and get statistics on image regions

Image Analysis

bwboundaries	Trace region boundaries in binary image
bwtraceboundary	Trace object in binary image
corner	Find corner points in image
cornermetric	Create corner metric matrix from image
edge	Find edges in grayscale image
hough	Hough transform
houghlines	Extract line segments based on Hough transform
houghpeaks	Identify peaks in Hough transform
imfindcircles	Find circles using circular Hough transform
qtdecomp	Quadtree decomposition
qtgetblk	Block values in quadtree decomposition
qtsetblk	Set block values in quadtree decomposition
viscircles	Create circle

Texture Analysis

entropy	Entropy of grayscale image
entropyfilt	Local entropy of grayscale image
graycomatrix	Create gray-level co-occurrence matrix from image
graycoprops	Properties of gray-level co-occurrence matrix
rangefilt	Local range of image
stdfilt	Local standard deviation of image

Pixel Values and Statistics

corr2	2-D correlation coefficient
imcontour	Create contour plot of image data
imhist	Display histogram of image data
imhistmatch	Adjust histogram of image to match N-bin histogram of reference image
impixel	Pixel color values
improfile	Pixel-value cross-sections along line segments
mean2	Average or mean of matrix elements
regionprops	Measure properties of image regions
std2	Standard deviation of matrix elements

Image Arithmetic

<code>imabsdiff</code>	Absolute difference of two images
<code>imadd</code>	Add two images or add constant to image
<code>imapplymatrix</code>	Linear combination of color channels
<code>imcomplement</code>	Complement image
<code>imdivide</code>	Divide one image into another or divide image by constant
<code>imlincomb</code>	Linear combination of images
<code>immultiply</code>	Multiply two images or multiply image by constant
<code>imsubtract</code>	Subtract one image from another or subtract constant from image

Image Enhancement and Restoration

Image Enhancement (p. 1-13)	Histogram equalization, decorrelation stretching, and 2-D filtering
Image Restoration (Deblurring) (p. 1-13)	Deconvolution for deblurring

Image Enhancement

adapthisteq	Contrast-limited adaptive histogram equalization (CLAHE)
decorrstretch	Apply decorrelation stretch to multichannel image
histeq	Enhance contrast using histogram equalization
imadjust	Adjust image intensity values or colormap
imnoise	Add noise to image
intlut	Convert integer values using lookup table
medfilt2	2-D median filtering
ordfilt2	2-D order-statistic filtering
stretchlim	Find limits to contrast stretch image
wiener2	2-D adaptive noise-removal filtering

Image Restoration (Deblurring)

deconvblind	Deblur image using blind deconvolution
deconvlucy	Deblur image using Lucy-Richardson method

deconvreg	Deblur image using regularized filter
deconvwnr	Deblur image using Wiener filter
edgetaper	Taper discontinuities along image edges
otf2psf	Convert optical transfer function to point-spread function
psf2otf	Convert point-spread function to optical transfer function

Linear Filtering and Transforms

Linear Filtering (p. 1-15)

Linear 2-D Filter Design (p. 1-15)

Image Transforms (p. 1-15)

Convolution, N-D filtering, and predefined 2-D filters

2-D FIR filters

Fourier, Discrete Cosine, Radon, and Fan-beam transforms

Linear Filtering

convmtx2

fspecial

imfilter

2-D convolution matrix

Create predefined 2-D filter

N-D filtering of multidimensional images

Linear 2-D Filter Design

freqz2

fsamp2

ftrans2

fwind1

fwind2

2-D frequency response

2-D FIR filter using frequency sampling

2-D FIR filter using frequency transformation

2-D FIR filter using 1-D window method

2-D FIR filter using 2-D window method

Image Transforms

dct2

dctmtx

2-D discrete cosine transform

Discrete cosine transform matrix

fan2para	Convert fan-beam projections to parallel-beam
fanbeam	Fan-beam transform
idct2	2-D inverse discrete cosine transform
ifanbeam	Inverse fan-beam transform
iradon	Inverse Radon transform
para2fan	Convert parallel-beam projections to fan-beam
phantom	Create head phantom image
radon	Radon transform

Morphological Operations

Intensity and Binary Images (p. 1-17)	Dilate, erode, reconstruct, and perform other morphological operations
Binary Images (p. 1-18)	Label, pack, and perform morphological operations on binary images
Structuring Element Creation and Manipulation (p. 1-19)	Create and manipulate structuring elements for morphological operations

Intensity and Binary Images

conndef	Create connectivity array
imbothat	Bottom-hat filtering
imclearborder	Suppress light structures connected to image border
imclose	Morphologically close image
imdilate	Dilate image
imerode	Erode image
imextendedmax	Extended-maxima transform
imextendedmin	Extended-minima transform
imfill	Fill image regions and holes
imhmax	H-maxima transform
imhmin	H-minima transform
imimposemin	Impose minima
imopen	Morphologically open image
imreconstruct	Morphological reconstruction
imregionalmax	Regional maxima

imregionalmin

Regional minima

imtophat

Top-hat filtering

watershed

Watershed transform

Binary Images

applylut

Neighborhood operations on binary images using lookup tables

bwarea

Area of objects in binary image

bwareaopen

Remove small objects from binary image

bwconncomp

Find connected components in binary image

bwconvhull

Generate convex hull image from binary image

bwdist

Distance transform of binary image

bwdistgeodesic

Geodesic distance transform of binary image

bweuler

Euler number of binary image

bwhitmiss

Binary hit-miss operation

bwlabel

Label connected components in 2-D binary image

bwlabeln

Label connected components in binary image

bwmorph

Morphological operations on binary images

bwpack

Pack binary image

bwperim

Find perimeter of objects in binary image

bwselect

Select objects in binary image

bwulterode	Ultimate erosion
bwunpack	Unpack binary image
graydist	Gray-weighted distance transform of grayscale image
imtophat	Top-hat filtering
makelut	Create lookup table for use with bwlookup

Structuring Element Creation and Manipulation

getheight	Height of structuring element
getneighbors	Structuring element neighbor locations and heights
getnhood	Structuring element neighborhood
getsequence	Sequence of decomposed structuring elements
isflat	True for flat structuring element
reflect	Reflect structuring element
strel	Create morphological structuring element (STREL)
translate	Translate structuring element (STREL)

ROI-Based, Neighborhood, and Block Processing

ROI-Based Processing (p. 1-20)	Define regions of interest (ROI) and perform operations on them
Neighborhood and Block Processing (p. 1-20)	Define neighborhoods and blocks and process them

ROI-Based Processing

poly2mask	Convert region of interest (ROI) polygon to region mask
roicolor	Select region of interest (ROI) based on color
roifill	Fill in specified region of interest (ROI) polygon in grayscale image
roifilt2	Filter region of interest (ROI) in image
roipoly	Specify polygonal region of interest (ROI)

Neighborhood and Block Processing

bestblk	Determine optimal block size for block processing
blockproc	Distinct block processing for image
close (ImageAdapter)	Close ImageAdapter object
col2im	Rearrange matrix columns into blocks
colfilt	Columnwise neighborhood operations
im2col	Rearrange image blocks into columns
ImageAdapter	Interface for image I/O

<code>nlfilter</code>	General sliding-neighborhood operations
<code>readRegion (ImageAdapter)</code>	Read region of image
<code>writeRegion (ImageAdapter)</code>	Write block of data to region of image

Colormaps and Color Space

Color Space Conversions (p. 1-22)

ICC profile-based device independent color space conversions and device-dependent color space conversions

Color Space Conversions

applycform	Apply device-independent color space transformation
iccfind	Search for ICC profiles
iccread	Read ICC profile
iccroot	Find system default ICC profile repository
iccwrite	Write ICC color profile to disk file
isicc	True for valid ICC color profile
lab2double	Convert $L^*a^*b^*$ data to double
lab2uint16	Convert $L^*a^*b^*$ data to uint16
lab2uint8	Convert $L^*a^*b^*$ data to uint8
makecform	Create color transformation structure
ntsc2rgb	Convert NTSC values to RGB color space
rgb2ntsc	Convert RGB color values to NTSC color space
rgb2ycbcr	Convert RGB color values to YCbCr color space
whitepoint	XYZ color values of standard illuminants
xyz2double	Convert XYZ color values to double

xyz2uint16

Convert *XYZ* color values to uint16

ycbcr2rgb

Convert YCbCr color values to RGB
color space

Utilities

Preferences (p. 1-24)	Set and determine value of toolbox preferences
Validation (p. 1-24)	Check input arguments and perform other common programming tasks
Mouse (p. 1-25)	Retrieve values of lines, points, and rectangles defined interactively using mouse
Array Operations (p. 1-25)	Circularly shift pixel values and pad arrays
Examples (p. 1-25)	Launch Image Processing Toolbox™ examples
Performance (p. 1-25)	Check for presence of Intel® Integrated Performance Primitives (Intel IPP) library

Preferences

<code>iptgetpref</code>	Get values of Image Processing Toolbox preferences
<code>iptprefs</code>	Display Image Processing Preferences dialog box
<code>iptsetpref</code>	Set Image Processing Toolbox preferences or display valid values

Validation

<code>getrangefromclass</code>	Default display range of image based on its class
<code>iptcheckconn</code>	Check validity of connectivity argument
<code>iptcheckinput</code>	Check validity of array

<code>iptcheckmap</code>	Check validity of colormap
<code>iptchecknargin</code>	Check number of input arguments
<code>iptcheckstrs</code>	Check validity of option string
<code>iptnum2ordinal</code>	Convert positive integer to ordinal string

Mouse

<code>getline</code>	Select polyline with mouse
<code>getpts</code>	Specify points with mouse
<code>getrect</code>	Specify rectangle with mouse

Array Operations

<code>padarray</code>	Pad array
-----------------------	-----------

Examples

<code>iptdemos</code>	Index of Image Processing Toolbox examples
-----------------------	--

Performance

<code>ippl</code>	Check for presence of Intel Integrated Performance Primitives (Intel IPP) library
-------------------	---

Class Reference

- “Image Input and Output” on page 2-2
- “Image Registration” on page 2-3

Image Input and Output

ImageAdapter

close

Close ImageAdapter object

ImageAdapter

Interface for image I/O

readRegion

Read region of image

writeRegion

Write block of data to region of image

Image Registration

registration.metric.MattesMutualInformation

registration.metric.MattesMutualInformation **Mattes** mutual information metric configuration object

registration.metric.MeanSquares,

registration.metric.MeanSquares **Mean square error metric** configuration object

registration.optimizer.RegularStepGradientDescent

registration.optimizer.RegularStepGradientDescent **Regular step** gradient descent optimizer configuration object

registration.optimizer.OnePlusOneEvolutionary

registration.optimizer.OnePlusOneEvolutionary **One plus** one evolutionary optimizer configuration object

Functions — Alphabetical List

adapthisteq

Purpose 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 grayscale 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'	<p>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.</p> <p>Default: 256</p>
'Range'	<p>String specifying the range of the output image data.</p> <p>'original' — Range is limited to the range of the original image, $[\min(I(:)) \quad \max(I(:))]$.</p> <p>'full' — Full range of the output image class is used. For example, for uint8 data, range is [0 255].</p> <p>Default: 'full'</p>
'Distribution'	<p>String specifying the desired histogram shape for the image tiles.</p> <p>'uniform' — Flat histogram</p> <p>'rayleigh' — Bell-shaped histogram</p> <p>'exponential' — Curved histogram</p> <p>Default: 'uniform'</p>
'Alpha'	<p>Nonnegative real scalar specifying a distribution parameter.</p> <p>Default: 0.4</p> <hr/> <p>Note Only used when 'Distribution' is set to either 'rayleigh' or 'exponential'.</p> <hr/>

Tips

- 'NumTiles' specifies the number of rectangular contextual regions (tiles) into which adapthisteq divides the image. adapthisteq calculates the contrast transform function for each of these regions individually. The optimal number of tiles depends on the type of the input image, and it is best determined through experimentation.
- 'ClipLimit' is a contrast factor that prevents over-saturation of the image specifically in homogeneous areas. These areas are characterized by a high peak in the histogram of the particular image tile due to many pixels falling inside the same gray level range. Without the clip limit, the adaptive histogram equalization technique could produce results that, in some cases, are worse than the original image.
- 'Distribution' specifies the distribution that adapthisteq uses as the basis for creating the contrast transform function. The distribution you select should depend on the type of the input image. For example, underwater imagery appears to look more natural when the Rayleigh distribution is used.

Class Support

Grayscale image I can be of class uint8, uint16, int16, single, or double. The output image J has the same class as I.

Examples

Apply 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);
```

Apply CLAHE to a color image.

```
[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(IMG, 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(IMG);
figure, imshow(J);
```

References

[1] Zuiderveld, Karel. “Contrast Limited Adaptive Histogram Equalization.” *Graphic Gems IV*. San Diego: Academic Press Professional, 1994. 474–485.

See Also

histeq

analyze75info

Purpose Read metadata from header file of Analyze 7.5 data set

Syntax
`info = analyze75info(filename)`
`info = analyze75info(filename, 'ByteOrder', endian)`

Description `info = analyze75info(filename)` reads the header file of the Analyze 7.5 data set specified by the string `filename`. The function returns `info`, a structure whose fields contain information about the data set.

Analyze 7.5 is a 3-D biomedical image visualization and analysis product developed by the Biomedical Imaging Resource of the Mayo Clinic. An Analyze 7.5 data set is made of two files, a header file and an image file. The files have the same name with different file extensions. The header file has the file extension `.hdr` and the image file has the file extension `.img`.

`info = analyze75info(filename, 'ByteOrder', endian)` reads the Analyze 7.5 header file using the byte ordering specified by *endian*, where *endian* can have either of the following values:

Value	Meaning
'ieee-le'	Byte ordering is Little Endian
'ieee-be'	Byte ordering is Big Endian

If the specified *endian* value results in a read error, `analyze75info` issues a warning message and attempts to read the header file with the opposite `ByteOrder` format.

Examples Read an Analyze 7.5 header file.

```
info = analyze75info('brainMRI.hdr');
```

Specify the byte ordering of the data set.

```
info = analyze75info('brainMRI', 'ByteOrder', 'ieee-le');
```

See Also `analyze75read`

Purpose Read image data from image file of Analyze 7.5 data set

Syntax
`X = analyze75read(filename)`
`X = analyze75read(info)`

Description `X = analyze75read(filename)` reads the image data from the image file of an Analyze 7.5 format data set specified by the string `filename`. The function returns the image data in `X`. For single-frame, grayscale images, `X` is an m -by- n array. `analyze75read` uses a data type for `X` that is consistent with the data type specified in the data set header file.

Analyze 7.5 is a 3-D biomedical image visualization and analysis product developed by the Biomedical Imaging Resource of the Mayo Clinic. An Analyze 7.5 data set is made of two files, a header file and an image file. The files have the same name with different file extensions. The header file has the file extension `.hdr` and the image file has the file extension `.img`.

`X = analyze75read(info)` reads the image data from the image file specified in the metadata structure `info`. `info` must be a valid metadata structure returned by the `analyze75info` function.

Note `analyze75read` returns image data in radiological orientation (LAS). This is the default used by the Analyze 7.5 format.

Examples

Example 1

Read image data from an Analyze 7.5 image file.

```
X = analyze75read('brainMRI');
```

Because Analyze 7.5 format uses radiological orientation (LAS), flip the data for correct image display in MATLAB.

```
X = flipdim(X,1);
```

Select frames 12 to 17 and use `reshape` to create an array for montage.

analyze75read

```
Y = reshape(X(:,:,12:17),[size(X,1) size(X,2) 1 6]);  
montage(Y);
```

Example 2

Call `analyze75read` with the metadata obtained from the header file using `analyze75info`.

```
info = analyze75info('brainMRI.hdr');  
X = analyze75read(info);
```

Class Support

X can be `logical`, `uint8`, `int16`, `int32`, `single`, or `double`. Complex and RGB data types are not supported.

See Also

`analyze75info`

Purpose Apply device-independent color space transformation

Syntax `B = applycform(A,C)`

Description `B = applycform(A,C)` converts the color values in `A` 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 `A` is two-dimensional, each row is interpreted as a color unless the color transformation structure contains a grayscale ICC profile. (See Note for this case.) `A` can have 1 or more columns, depending on the input color space. `B` has the same number of rows and 1 or more columns, depending on the output color space. (The ICC spec currently supports up to 15-channel device spaces.)

If `A` is three-dimensional, each row-column location is interpreted as a color, and `size(A,3)` is typically 1 or more, depending on the input color space. `B` has the same number of rows and columns as `A`, and `size(B,3)` is 1 or more, depending on the output color space.

Class Support `A` is a real, nonsparse array of class `uint8`, `uint16`, or `double` or a string. `A` is only a string if `C` was created with the following syntax:

```
C = makecform('named', profile, space)
```

The output array `B` has the same class as `A`, unless the output 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.

Note If the color transformation structure `C` contains a grayscale ICC profile, `applycform` interprets each pixel in `A` as a color. `A` can have any number of columns. `B` has the same size as `A`.

applycform

Examples

Read in a color image that uses the sRGB color space.

```
rgb = imread('peppers.png');
```

Create a color transformation structure that defines an sRGB to $L^*a^*b^*$ conversion.

```
C = makecform('srgb2lab');
```

Perform the transformation with `applycform`.

```
lab = applycform(rgb,C);
```

See Also

[lab2double](#) | [lab2uint8](#) | [lab2uint16](#) | [makecform](#) | [whitepoint](#) | [xyz2double](#) | [xyz2uint16](#)

How To

- “Converting Color Data Between Color Spaces”

Purpose 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`.

Algorithms `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.

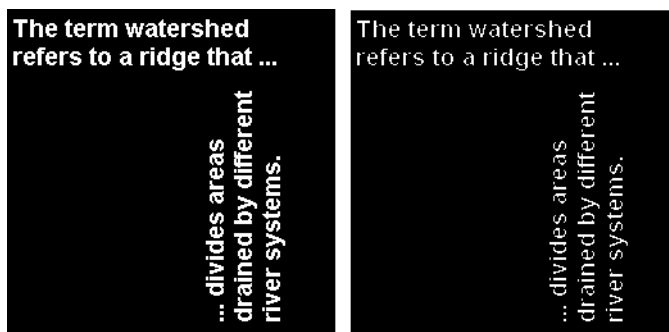
256	32	4
128	16	2
64	8	1

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.

Examples

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)
```



See Also

makelut

Purpose Convert axes coordinates to pixel coordinates

Syntax `pixelx = axes2pix(dim, XDATA, AXESX)`

Description `pixelx = axes2pix(dim, XDATA, AXESX)` converts an axes coordinate into an intrinsic (“pixel”) coordinate. For example, if `pt = get(gca, 'CurrentPoint')` then `AXESX` could be `pt(1,1)` or `pt(1,2)`. `AXESX` must be in intrinsic coordinates. `XDATA` is a two-element vector returned by `get(image_handle, 'XData')` or `get(image_handle, 'YData')`. `dim` is the number of image columns for the `x` coordinate, or the number of image rows for the `y` coordinate.

Class Support `dim`, `XDATA`, and `AXESX` can be double. The output is double.

Note `axes2pix` performs minimal checking on the validity of `AXESX`, `DIM`, or `XDATA`. For example, `axes2pix` returns a negative coordinate if `AXESX` is less than `XDATA(1)`. The function calling `axes2pix` bears responsibility for error checking.

Examples Example with default `XData` and `YData`.

```
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)
```

Example with non-default `XData` and `YData`.

```
xdata = [10 100]
ydata = [20 90]
px = axes2pix(ncols,xdata,30)
py = axes2pix(nrows,ydata,30)
```

See Also `impixelinfo` | `bwselect` | `imfill` | `impixel` | `improfile` | `roipoly`

Purpose	Determine optimal block size for block processing
Syntax	<pre>siz = bestblk([m n],k) [mb,nb] = bestblk([m n],k)</pre>
Description	<p><code>siz = bestblk([m n],k)</code> returns, for an m-by-n image, the optimal block size for block processing. <code>k</code> 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 <code>siz</code> is a 1-by-2 vector containing the row and column dimensions for the block.</p> <p><code>[mb,nb] = bestblk([m n],k)</code> returns the row and column dimensions for the block in <code>mb</code> and <code>nb</code>, respectively.</p>
Algorithms	<p><code>bestblk</code> returns the optimal block size given <code>m</code>, <code>n</code>, and <code>k</code>. The algorithm for determining <code>siz</code> is</p> <ul style="list-style-type: none">• If <code>m</code> is less than or equal to <code>k</code>, return <code>m</code>.• If <code>m</code> is greater than <code>k</code>, consider all values between $\min(m/10, k/2)$ and <code>k</code>. Return the value that minimizes the padding required. <p>The same algorithm is then repeated for <code>n</code>.</p>
Examples	<pre>siz = bestblk([640 800],72) siz = 64 50</pre>
See Also	<code>blockproc</code>

blockproc

Purpose Distinct block processing for image

Syntax

```
B = blockproc(A,[M N],fun)
B = blockproc(src_filename,[M N],fun)
B = blockproc(adapter,[M N],fun)
blockproc(...,Name,Value,...)
```

Description `B = blockproc(A,[M N],fun)` processes the image `A` by applying the function `fun` to each distinct `M`-by-`N` block of `A` and concatenating the results into `B`, the output matrix. `fun` is a function handle to a function that accepts a *block struct* as input and returns a matrix, vector, or scalar `Y`. For example, `Y = fun(block_struct)`. (For more information about a *block struct*, see the Definition section below.) For each block of data in the input image, `A`, `blockproc` passes the block in a *block struct* to the user function, `fun`, to produce `Y`, the corresponding block in the output image. If `Y` is empty, `blockproc` does not generate any output and returns empty after processing all blocks. Choosing an appropriate block size can significantly improve performance. For more information, see “Choosing Block Size” in the Image Processing Toolbox documentation.

`B = blockproc(src_filename,[M N],fun)` processes the image `src_filename`, reading and processing one block at a time. This syntax is useful for processing very large images since only one block of the image is read into memory at a time. If the output matrix `B` is too large to fit into memory, omit the output argument and instead use the 'Destination' parameter/value pair to write the output to a file.

`B = blockproc(adapter,[M N],fun)` processes the source image specified by `adapter`, an `ImageAdapter` object. An `ImageAdapter` is a user-defined class that provides `blockproc` with a common API for reading and writing to a particular image file format. For more information, see “Working with Data in Unsupported Formats” in the Image Processing Toolbox documentation.

`blockproc(...,Name,Value,...)` processes the input image, specifying parameters and corresponding values that control various aspects of the block behavior. Parameter names are not case sensitive.

Input Arguments

A

Input image.

[M N]

Block size of A.

fun

Function handle to a function that accepts a *block struct* as input and returns a matrix, vector, or scalar.

src_filename

Input image.

adapter

A user-defined class that provides `blockproc` with a common API for reading and writing to a particular image file format.

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as `Name1, Value1, ..., NameN, ValueN`.

BorderSize

A two-element vector, `[V H]`, specifying the amount of border pixels to add to each block. The function adds `V` rows above and below each block and `H` columns left and right of each block. The size of each resulting block will be:

$$[M + 2*V, N + 2*H]$$

By default, the function automatically removes the border from the result of `fun`. See the `'TrimBorder'` parameter for more information.

The function pads blocks with borders extending beyond the image edges with zeros.

Default: [0 0] (no border)

Destination

The destination for the output of `blockproc`. When you specify the 'Destination' parameter, `blockproc` does not return the processed image as an output argument, but instead writes the output to the 'Destination'. (You cannot request an output argument when the 'Destination' parameter is specified.)

Valid 'Destination' parameters are:

- TIFF filename: A string filename ending with '.tif'. If a file with this name already exists, it will be overwritten.
- ImageAdapter object: An instance of an ImageAdapter class. ImageAdapters provide an interface for reading and writing to arbitrary image file formats.

The 'Destination' parameter is useful when you expect your output to be too large to practically fit into memory. It provides a workflow for file-to-file image processing for arbitrarily large images.

PadPartialBlocks

A logical scalar. When set to `true`, `blockproc` pads partial blocks to make them full-sized (M-by-N) blocks. Partial blocks arise when the image size is not exactly divisible by the block size. If they exist, partial blocks lie along the right and bottom edge of the image. The default is `false`, meaning that the function does not pad the partial blocks, but processes them as-is. `blockproc` uses zeros to pad partial blocks when necessary.

Default: `false`

PadMethod

The 'PadMethod' determines how blockproc will pad the image boundary. Options are:

- X: Pads the image with a scalar (X) pad value. By default $X == 0$.
- 'replicate': Repeats border elements of image A.
- 'symmetric': Pads image A with mirror reflections of itself.

TrimBorder

A logical scalar. When set to true, the blockproc function trims off border pixels from the output of the user function, fun. The function removes V rows from the top and bottom of the output of fun, and H columns from the left and right edges. The 'BorderSize' parameter defines V and H. The default is true, meaning that the blockproc function automatically removes borders from the fun output.

Default: true

UseParallel

A logical scalar. Enabling this mode of image processing requires the Parallel Computing Toolbox™. When set to true, blockproc will attempt to run in parallel mode, distributing the processing across multiple workers (MATLAB sessions) in an open MATLAB pool. In parallel mode, the input image cannot be an ImageAdapter object. See the documentation for matlabpool for information on configuring your parallel environment.

Default: false

File Format Support: Input and output files for blockproc (as specified by src_filename and the 'Destination' parameter) must have one of the following file types and must be named with one of the listed file extensions:

- Read/Write File Formats: TIFF (*.tif, *.tiff), JPEG2000 (*.jp2, *.j2c, *.j2k)

- Read-Only File Formats: JPEG2000 (*.jpf, *.jpx)

Output Arguments

B
Output matrix.

Definitions

A *block struct* is a MATLAB structure that contains the block data as well as other information about the block. Fields in the *block struct* are:

- `block_struct.border`: A two-element vector, `[V H]`, that specifies the size of the vertical and horizontal padding around the block of data. (See the 'BorderSize' parameter in the Inputs section.)
- `block_struct.blockSize`: A two-element vector, `[rows cols]`, that specifies the size of the block data. If a border has been specified, the size does not include the border pixels.
- `block_struct.data`: M-by-N or M-by-N-by-P matrix of block data
- `block_struct.imageSize`: A two-element vector, `[rows cols]`, that specifies the full size of the input image.
- `block_struct.location`: A two-element vector, `[row col]`, that specifies the position of the first pixel (minimum-row, minimum-column) of the block data in the input image. If a border has been specified, the location refers to the first pixel of the discrete block data, not the added border pixels.

Examples

Generate an image thumbnail:

```
fun = @(block_struct) imresize(block_struct.data,0.15);  
I = imread('pears.png');  
I2 = blockproc(I,[100 100],fun);  
figure;  
imshow(I);  
figure;  
imshow(I2);
```



Set the pixels in each 32-by-32 block to the standard deviation of the elements in that block:

```
fun = @(block_struct) ...
    std2(block_struct.data) * ones(size(block_struct.data));
I2 = blockproc('moon.tif',[32 32],fun);
figure;
imshow('moon.tif');
figure;
imshow(I2,[]);
```



Original Image



Standard Deviation Image

Switch the red and green bands of an RGB image and write the results to a new TIFF file:

blockproc

```
I = imread('peppers.png');  
fun = @(block_struct) block_struct.data(:,:, [2 1 3]);  
blockproc(I,[200 200],fun,'Destination','grb_peppers.tif');  
figure;  
imshow('peppers.png');  
figure;  
imshow('grb_peppers.tif');
```



Original Image of Peppers



Recolored Image of Peppers

Convert a TIFF image into a new JPEG2000 image. Replace 'largeImage.tif' in the example below with the name of your file:

```
fun = @(block_struct) block_struct.data;  
blockproc('largeImage.tif',[1024 1024],fun,...  
    'Destination','New.jp2');
```

See Also


[colfilt](#) | [function_handle](#) | “ImageAdapter” on page 2-2 | [nlfilter](#)

How To

- “Anonymous Functions”
- “Parameterizing Functions”
- “Performing Distinct Block Operations”

bwarea

Purpose	Area of objects in 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> .
Algorithms	<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>
Examples	<p>Compute 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`

References

[1] Pratt, William K., *Digital Image Processing*, New York, John Wiley & Sons, Inc., 1991, p. 634.

See Also

`bweuler` | `bwperim`

bwareaopen

Purpose Remove small objects from binary image

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. This operation is known as an area opening. 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 0s and 1s. The 1-valued elements define neighborhood locations relative to the central element of conn. Note that conn must be symmetric about its central 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`.

Algorithms

The basic steps are

- 1 Determine the connected components:

```
CC = bwconncomp(BW, conn);
```

- 2 Compute the area of each component:

```
S = regionprops(CC, 'Area');
```


- 3 Remove small objects:

```
L = labelmatrix(CC);  
BW2 = ismember(L, find([S.Area] >= P));
```

Examples

Remove all objects in the image `text.png` containing fewer than 50 pixels:

```
BW = imread('text.png');  
BW2 = bwareaopen(BW, 50);  
imshow(BW);
```



```
figure, imshow(BW2)
```

bwareaopen



See Also


[bwconncomp](#) | [conndef](#)

Purpose Trace region boundaries in 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.

`B = bwboundaries(BW,conn)` specifies the connectivity to use when tracing parent and child boundaries. `conn` can have either of the following scalar values.

bwboundaries

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, two-dimensional, and nonsparse. `L` and `N` are double. `A` is sparse logical.

Examples

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);
figure, 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
figure; spy(A);
```

Example 3

Display object boundaries in red and hole boundaries in green.

```
BW = imread('blobs.png');
[B,L,N] = bwboundaries(BW);
figure; 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);
figure; 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
```

```
end  
end
```

Algorithms

The `bwboundaries` function implements the Moore-Neighbor tracing algorithm modified by Jacob's stopping criteria. This function is based on the `boundaries` function presented in the first edition of *Digital Image Processing Using MATLAB*[1].

References

[1] Gonzalez, R. C., R. E. Woods, and S. L. Eddins, *Digital Image Processing Using MATLAB*, New Jersey, Pearson Prentice Hall, 2004.

See Also

`bwlabel` | `bwlabeln` | `bwperim` | `bwtraceboundary`

bwconncomp

Purpose Find connected components in binary image

Syntax
CC = bwconncomp(BW)
CC = bwconncomp(BW,conn)

Description CC = bwconncomp(BW) returns the connected components CC found in BW. The binary image BW can have any dimension. CC is a structure with four fields.

Field	Description
Connectivity	Connectivity of the connected components (objects)
ImageSize	Size of BW
NumObjects	Number of connected components (objects) in BW
PixelIdxList	1-by-NumObjects cell array where the <i>k</i> th element in the cell array is a vector containing the linear indices of the pixels in the <i>k</i> th object.

bwconncomp uses a default connectivity of 8 for two dimensions, 26 for three dimensions, and `conndef(ndims(BW), 'maximal')` for higher dimensions.

CC = bwconncomp(BW,conn) specifies the desired connectivity for the connected components. conn can have 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 using a 3-by-3-by- ... -by-3 matrix of 0s and 1s. `conn` must be symmetric about its center element. The 1-valued elements define neighborhood locations relative to `conn`.

The functions `bwlabel`, `bwlabeln`, and `bwconncomp` all compute connected components for binary images. `bwconncomp` replaces the use of `bwlabel` and `bwlabeln`. It uses significantly less memory and is sometimes faster than the other functions.

Function	Input Dimension	Output Form	Memory Use	Connectivity
<code>bwlabel</code>	2-D	Label matrix with double-precision	High	4 or 8
<code>bwlabeln</code>	N-D	Double-precision label matrix	High	Any
<code>bwconncomp</code>	N-D	CC struct	Low	Any

Tips

To extract features from a binary image using `regionprops` with default connectivity, just pass `BW` directly into `regionprops` (i.e., `regionprops(BW)`).

To compute a label matrix having more memory-efficient data type (e.g., `uint8` versus `double`), use the `labelmatrix` function on the output of `bwconncomp`. See the documentation for each function for more information.

Class Support

`BW` can be a logical or numeric array of any dimension, and it must be real and nonsparse. `CC` is a structure.

Examples

Example 1

Calculate the centroids of the 3-D objects.

```
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]);
```

```
CC = bwconncomp(BW);  
S = regionprops(CC, 'Centroid');
```

Example 2

Erase the largest letter from the image.

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



```
CC = bwconncomp(BW);  
numPixels = cellfun(@numel, CC.PixelIdxList);  
[biggest, idx] = max(numPixels);  
BW(CC.PixelIdxList{idx}) = 0;  
  
figure, imshow(BW);
```

The term watershed
refers to a ridge that ...

... divides areas
drained by different
river systems.

Algorithms

The basic steps in finding the connected components are:

- 1 Search for the next unlabeled pixel, p .
- 2 Use a flood-fill algorithm to label all the pixels in the connected component containing p .
- 3 Repeat steps 1 and 2 until all the pixels are labelled.

See Also

`bwlabel` | `bwlabeln` | `labelmatrix` | `regionprops`

bwconvhull

Purpose Generate convex hull image from binary image

Syntax
CH = bwconvhull(BW)
CH = bwconvhull(BW, method)
CH = bwconvhull(BW, 'objects', conn)

Description CH = bwconvhull(BW) computes the convex hull of all objects in BW and returns CH, a binary convex hull image.
CH = bwconvhull(BW, method) specifies the desired method for computing the convex hull image.
CH = bwconvhull(BW, 'objects', conn) specifies the desired connectivity used when defining individual foreground objects. The conn parameter is only valid when the method is 'objects'.

Input Arguments
BW
A logical 2-D image

method

A string that can have the following values:

- 'union': Compute the convex hull of all foreground objects, treating them as a single object.
- 'objects': Compute the convex hull of each connected component of BW individually. CH contains the convex hulls of each connected component.

Default: 'union'

conn

Connectivity. Can have the following scalar values:

- 4: Two-dimensional, four-connected neighborhood.
- 8: Two-dimensional, eight-connected neighborhood.

Additionally, `conn` may be defined in a more general way, using a 3-by-3 matrix of 0s and 1s. The 1-valued elements define neighborhood locations relative to `conn`'s center element. `conn` must be symmetric about its center element.

Default: 8

Output Arguments

CH

A logical, convex hull image, containing the binary mask of the convex hull of all foreground objects in `BW`.

Examples

Display the binary convex hull of an image:


```
subplot(2,2,1);  
I = imread('coins.png');  
imshow(I);  
title('Original');  
  
subplot(2,2,2);  
BW = I > 100;  
imshow(BW);  
title('Binary');  
  
subplot(2,2,3);  
CH = bwconvhull(BW);  
imshow(CH);  
title('Union Convex Hull');  
  
subplot(2,2,4);  
CH_objects = bwconvhull(BW, 'objects');  
imshow(CH_objects);  
title('Objects Convex Hull');
```

bwconvhull

Original



Binary



Union Convex Hull



Objects Convex Hull



See Also

[bwconncomp](#) | [bwlabel](#) | [labelmatrix](#) | [regionprops](#)

Purpose Distance transform of binary image

Syntax

```
D = bwdist(BW)
[D,IDX] = bwdist(BW)
[D,IDX] = 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,IDX] = bwdist(BW)` also computes the closest-pixel map in the form of an index array, `IDX`. (The closest-pixel map is also called the feature map, feature transform, or nearest-neighbor transform.) `IDX` has the same size as `BW` and `D`. Each element of `IDX` contains the linear index of the nearest nonzero pixel of `BW`.

`[D,IDX] = bwdist(BW,method)` computes the distance transform, where `method` specifies an alternate distance metric. `method` can take any of the following values. The `method` string can be abbreviated.

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'	<p>In 2-D, the Euclidean distance between (x_1, y_1) and (x_2, y_2) is</p> $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}.$ <p>This is the default method.</p>
'quasi-euclidean'	<p>In 2-D, the quasi-Euclidean distance between (x_1, y_1) and (x_2, y_2) is</p> $ x_1 - x_2 + (\sqrt{2} - 1) y_1 - y_2 , \quad x_1 - x_2 > y_1 - y_2 $ $(\sqrt{2} - 1) x_1 - x_2 + y_1 - y_2 , \quad \text{otherwise.}$

Tips

`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.

The function `bwdist` changed in version 6.4 (R2009b). Previous versions of the Image Processing Toolbox used different algorithms for computing the Euclidean distance transform and the associated label matrix. If you need the same results produced by the previous implementation, use the function `bwdist_old`.

Class Support

BW can be numeric or logical, and it must be nonsparse. D is a single matrix with the same size as BW. The class of IDX depends on the number of elements in the input image, and is determined using the following table.

Class	Range
'uint32'	numel(BW) $\leq 2^{32} - 1$
'uint64'	numel(BW) $\geq 2^{32}$

Examples

Compute 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,IDX] = 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
```

```
IDX =
```

```


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 `IDX` 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`.

Compare 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)
```




Compare 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')

```



Algorithms

For Euclidean distance transforms, `bwdist` uses the fast algorithm described in

[1] Maurer, Calvin, Rensheng Qi, and Vijay Raghavan, "A Linear Time Algorithm for Computing Exact Euclidean Distance Transforms

of Binary Images in Arbitrary Dimensions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 25, No. 2, February 2003, pp. 265-270.

For cityblock, chessboard, and quasi-Euclidean distance transforms, `bwdist` uses the two-pass, sequential scanning algorithm described in

[2] Rosenfeld, Azriel and John 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

[3] 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

`bwulterode` | `watershed`

How To

- "Distance Transform"

Purpose	Geodesic distance transform of binary image
Syntax	<pre>D = bwdistgeodesic(BW,mask) D = bwdistgeodesic(BW,C,R) D = bwdistgeodesic(BW,ind) D = bwdistgeodesic(...,method)</pre>
Description	<p>D = bwdistgeodesic(BW,mask) computes the geodesic distance transform, given the binary image BW and the seed locations specified by mask. Regions where BW is true represent valid regions that can be traversed in the computation of the distance transform. Regions where BW is false represent constrained regions that cannot be traversed in the distance computation. For each true pixel in BW, the geodesic distance transform assigns a number that is the constrained distance between that pixel and the nearest true pixel in mask. Output matrix D contains geodesic distances.</p> <p>D = bwdistgeodesic(BW,C,R) computes the geodesic distance transform of the binary image BW. Vectors C and R contain the column and row coordinates of the seed locations.</p> <p>D = bwdistgeodesic(BW,ind) computes the geodesic distance transform of the binary image BW. ind is a vector of linear indices of seed locations.</p> <p>D = bwdistgeodesic(...,method) specifies an alternate distance metric.</p>
Input Arguments	<p>BW Binary image.</p> <p>mask Logical image the same size as BW that specifies seed locations.</p> <p>C,R</p>

bwdistgeodesic

Numeric vectors that contain the positive integer column and row coordinates of the seed locations. Coordinate values are valid C,R subscripts in BW.

ind

Numeric vector of positive integer, linear indices of seed locations.

method

Type of distance metric. method can have any of these values.

Method	Description
'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 $.
'chessboard'	The chessboard distance is $\max(x_1 - x_2 , y_1 - y_2)$.
'quasi-euclidean'	The quasi-Euclidean distance 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 $, otherwise.

Default: 'chessboard'

Output Arguments

D

Numeric array of class `single`, with the same size as input BW, that contains geodesic distances.

Class Support

BW is a logical matrix. C, R, and ind are numeric vectors that contain positive integer values. D is a numeric array of class `single` that has the same size as the input BW.

Examples

Compute the geodesic distance transformation of `BW` based on the seed locations specified by vectors `C` and `R`. Output pixels for which `BW` is false have undefined geodesic distance and contain NaN values. Because there is no connected path from the seed locations to element `BW(10,5)`, the output `D(10,5)` has a value of `Inf`.

```
BW = [1 1 1 1 1 1 1 1 1 1;...
      1 1 1 1 1 1 0 0 1 1;...
      1 1 1 1 1 1 0 0 1 1;...
      1 1 1 1 1 1 0 0 1 1;...
      0 0 0 0 0 1 0 0 1 0;...
      0 0 0 0 1 1 0 1 1 0;...
      0 1 0 0 1 1 0 0 0 0;...
      0 1 1 1 1 1 1 0 1 0;...
      0 1 1 0 0 0 1 1 1 0;...
      0 0 0 0 1 0 0 0 0 0];
```

```
BW = logical(BW);
C = [1 2 3 3 3];
R = [3 3 3 1 2];
```

```
D = bwdistgeodesic(BW,C,R);
```

Algorithms

`bwdistgeodesic` uses the geodesic distance algorithm described in Soille, P., *Morphological Image Analysis: Principles and Applications*, 2nd Edition, Secaucus, NJ, Springer-Verlag, 2003, pp. 219–221.

See Also

`bwdist` | `graydist`

bweuler

Purpose Euler number of 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`.

Examples

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



```
bweuler(BW)
```

```
ans =
```

```
-3
```


Algorithms

`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.

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.

See Also

`bwmorph` | `bwperim`

bwhitmiss

Purpose Binary hit-miss operation

Syntax
BW2 = bwhitmiss(BW1,SE1,SE2)
BW2 = bwhitmiss(BW1,INTERVAL)

Description BW2 = bwhitmiss(BW1,SE1,SE2) performs the hit-miss operation defined by the structuring elements SE1 and SE2. The hit-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-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(BW1,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 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.

Examples Perform the hit-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];
```

```
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](#)

bwlabel

Purpose Label connected components in 2-D binary image

Syntax
`L = bwlabel(BW, n)`
`[L, num] = bwlabel(BW, n)`

Description `L = bwlabel(BW, n)` returns a matrix `L`, of the same size as `BW`, containing labels for the connected objects in `BW`. The variable `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.

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.

`[L, num] = bwlabel(BW, n)` returns in `num` the number of connected objects found in `BW`.

The functions `bwlabel`, `bwlabeln`, and `bwconncomp` all compute connected components for binary images. `bwconncomp` replaces the use of `bwlabel` and `bwlabeln`. It uses significantly less memory and is sometimes faster than the other functions.

	Input Dimension	Output Form	Memory Use	Connectivity
<code>bwlabel</code>	2-D	Double-precision label matrix	High	4 or 8
<code>bwlabeln</code>	N-D	Double-precision label matrix	High	Any
<code>bwconncomp</code>	N-D	CC struct	Low	Any

Tips **Using find with bwlabel**

You can use the MATLAB `find` function in conjunction with `bwlabel` 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, enter the following:

```
[r, c] = find(bwlabel(BW)==2)
```

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. For more information, see `label2rgb`.

Using `labelmatrix` and `regionprops`

To compute a label matrix having a more memory-efficient data type (e.g., `uint8` versus `double`), use the `labelmatrix` function on the output of `bwconncomp`. For more information, see the reference page for each function.

To extract features from a binary image using `regionprops` with default connectivity, just pass `BW` directly into `regionprops`, i.e., `regionprops(BW)`.

Class Support

`BW` can be logical or numeric, and it must be real, two-dimensional, and nonsparse. `L` is of class `double`.

Examples

Label components using 4-connected objects. Notice objects 2 and 3; with 8-connected labeling, `bwlabel` would consider these a single object rather than two separate objects.

```
BW = logical ([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
```

bwlabel

```
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
```

Algorithms

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.

References

[1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision, Volume I*, Addison-Wesley, 1992, pp. 28-48.

See Also

bwconncomp | bwlabeln | bwselect | labelmatrix | label2rgb | regionprops

Purpose Label connected components in 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. The input image 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 0s and 1s. The 1-valued elements define neighborhood locations relative to the

bwlabeln

central element of `conn`. Note that `conn` must be symmetric about its central element.

The functions `bwlabel`, `bwlabeln`, and `bwconncomp` all compute connected components for binary images. `bwconncomp` replaces the use of `bwlabel` and `bwlabeln`. It uses significantly less memory and is sometimes faster than the older functions.

Function	Input Dimension	Output Form	Memory Use	Connectivity
<code>bwlabel</code>	2-D	Double-precision label matrix	High	4 or 8
<code>bwlabeln</code>	N-D	Double-precision label matrix	High	Any
<code>bwconncomp</code>	N-D	CC struct	Low	Any

Tips

To extract features from a binary image using `regionprops` with default connectivity, just pass `BW` directly into `regionprops`, i.e., `regionprops(BW)`.

To compute a label matrix having a more memory-efficient data type (e.g., `uint8` versus `double`), use the `labelmatrix` function on the output of `bwconncomp`:

```
C = bwconncomp(BW);
L = labelmatrix(CC);

CC = bwconncomp(BW, conn);
S = regionprops(CC);
```

Class Support

`BW` can be numeric or logical, and it must be real and nonsparse. `L` is of class `double`.

Examples

Calculate the centroids of the 3-D objects.

```
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])
```

```
bwlabeln(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
```

Algorithms

bwlabeln 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.
- 2 Resolve the equivalence classes using the union-find algorithm [1].
- 3 Relabel the pixels based on the resolved equivalence classes.

References

[1] Sedgewick, Robert, *Algorithms in C*, 3rd Ed., Addison-Wesley, 1998, pp. 11-20.

bwlabeln

See Also

[bwconncomp](#) | [bwlabel](#) | [labelmatrix](#) | [label2rgb](#) | [regionprops](#)

Purpose	Nonlinear filtering using lookup tables
Syntax	<code>A = bwlookup(BW, lut)</code>
Description	<p><code>A = bwlookup(BW, lut)</code> performs a 2-by-2 or 3-by-3 nonlinear neighborhood filtering operation on binary or grayscale image <code>BW</code> and returns the results in output image <code>A</code>. The neighborhood processing determines an integer index value used to access values in lookup table <code>lut</code>. The fetched <code>lut</code> value becomes the pixel value in output image <code>A</code> at the targeted position.</p> <ul style="list-style-type: none">• <code>A</code> is the same size as <code>BW</code>• <code>A</code> is the same data type as <code>lut</code> <p><code>bwlookup</code> supports the generation of efficient, production-quality C/C++ code from MATLAB. For best results, specify an input image of class <code>logical</code>. To see a complete list of toolbox functions that support code generation, see “Supported Functions”.</p>
Input Arguments	<p>BW - Input image binary image grayscale image</p> <p>Input image transformed by nonlinear neighborhood filtering operation, specified as either a grayscale or binary (logical) image. In the case of numeric values, non-zero pixels are considered true which is equivalent to logical 1.</p> <p>Data Types single double int8 int16 int32 int64 uint8 uint16 uint32 uint64 logical</p> <p>lut - Lookup table of output pixel values 16- or 256-element vector</p> <p>Lookup table of output pixel values, specified as a 16- or 256-element vector. The size of <code>lut</code> determines which of the two neighborhood operations is performed.</p>

- If `lut` contains 16 data elements, then the neighborhood matrix is 2-by-2.
- If `lut` contains 256 data elements, then the neighborhood matrix is 3-by-3.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32
| uint64 | logical

Output Arguments

A - Output image

binary image | grayscale image

Output image, returned as a grayscale or binary image whose size matches `BW`, and whose distribution of pixel values are determined by the content of `lut`.

- A is the same size as `BW`
- A is the same data type as `lut`

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32
| uint64 | logical

Examples

2-by-2 Neighborhood Erosion of Binary Image

This example performs an erosion along the edges of a binary image. Vector `lut` is constructed such that the filtering operation places a 1 at the targeted pixel location in image `A` only when all 4 pixels in the 2-by-2 neighborhood of `BW` are set to 1. For all other 2-by-2 neighborhood combinations in `BW`, the targeted pixel location in image `A` receives a 0.

```
% Construct lut so it is 'true' only
% when all 4 2-by2 locations equal 1
lutfun = @(x)(sum(x(:))==4);
lut     = makelut(lutfun,2);

% load binary image
BW1     = imread('text.png');
```

```
% Perform 2x2 Neighborhood processing
% with 16 element vector lut
BW2    = bwlookup(BW1,lut);


% show Zoomed before and after images
figure;
h1 = subplot(1,2,1); imshow(BW1), axis off; title('BW1')
h2 = subplot(1,2,2); imshow(BW2); axis off; title('BW2')

% do a 16X Zoom to see effects of erosion on text
set(h1,'Ylim',[.5 64.5]); set(h1,'Xlim',[.5 64.5]);
set(h2,'Ylim',[.5 64.5]); set(h2,'Xlim',[.5 64.5]);


lut =

    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    0
    1
```

BW1



BW2



Algorithm

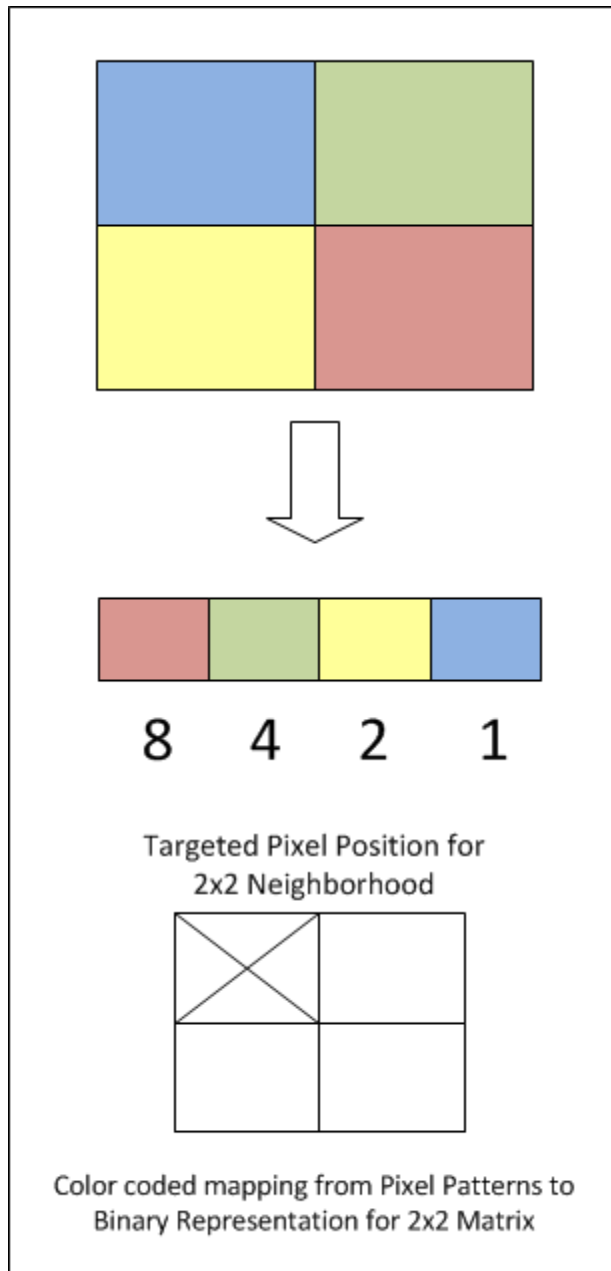
The first step in each iteration of the filtering operation performed by `bwlookup` entails computing the `index` into vector `lut` based on the binary pixel pattern of the neighborhood matrix on image `BW`. The value in `lut` accessed at `index`, `lut(index)`, is inserted into output image `A` at the targeted pixel location. This results in image `A` being the same data type as vector `lut`.

Since there is a 1-to-1 correspondence in targeted pixel locations, image `A` is the same size as image `BW`. If the targeted pixel location is on an edge of image `BW` and if any part of the 2x2 or 3x3 neighborhood matrix extends beyond the image edge, then these non-image locations are padded with 0 in order to perform the filtering operation.

Mapping from Neighborhood Matrices to Binary Notation

The following figures show the mapping from binary 0 and 1 patterns in the neighborhood matrices to its binary representation. Adding 1 to the binary representation yields `index` which is used to access `lut`.

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 $2^4 = 16$.



Example: This example demonstrates how the pixel pattern in a 2x2 matrix determines which entry in lut is placed in the targeted pixel location.

Create random 16 element lut vector containing uint8 data

```
scurr = rng; % save current seed state
rng('default') % always generate same set of random numbers
lut = uint8( round( 255*rand(16,1) ) ) % generate 16 uint8
rng(scurr); % restore
```

```
lut =
```

```
208
231
 32
233
161
 25
 71
139
244
246
 40
248
244
124
204
 36
```

Create a small 2x2 image; targeted pixel location will be at location BW(1,1).

```
BW = [1 0; 0 1]
```

```
BW =
```

```
1    0
```



```
0    1
```

By referring to the color coded mapping figure above, the binary representation for this 2x2 neighborhood can be computed as shown in the code snippet below. The logical 1 at BW(1,1) corresponds to blue in the figure which maps to the Least Significant Bit (LSB) at position 0 in the 4-bit binary representation ($2^0 = 1$). The logical 1 at BW(2,2) is red which maps to the Most Significant Bit (MSB) at position 3 in the 4-bit binary representation ($2^3 = 8$).

```
% BW(1,1): blue square; sets bit position 0 on right
% BW(2,2): red square; sets bit position 3 on left
binNot = '1 0 0 1'; % binary representation of 2x2 neighborhood matrix
```

```
X = bin2dec( binNot ); % convert from binary to decimal
index = X + 1 % add 1 to compute index value for uint8 vector lut
A11 = lut(index) % value at A(1,1)
```

```
index =
```

```
10
```

```
A11 =
```

```
246
```

The above calculation predicts that output image A should contain the value 246 at targeted position A(1,1).

```
A = bwlookup(BW,lut) % perform filtering
```

```
A =
```

```
246    32
161    231
```

A(1,1) does in fact equal 246.

Note For a more robust way to perform image erosion, see function `imerode`.

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 $2^9 = 512$.

The process for computing the binary representation of 3x3 neighborhood processing is the same as shown above for 2x2 neighborhoods.

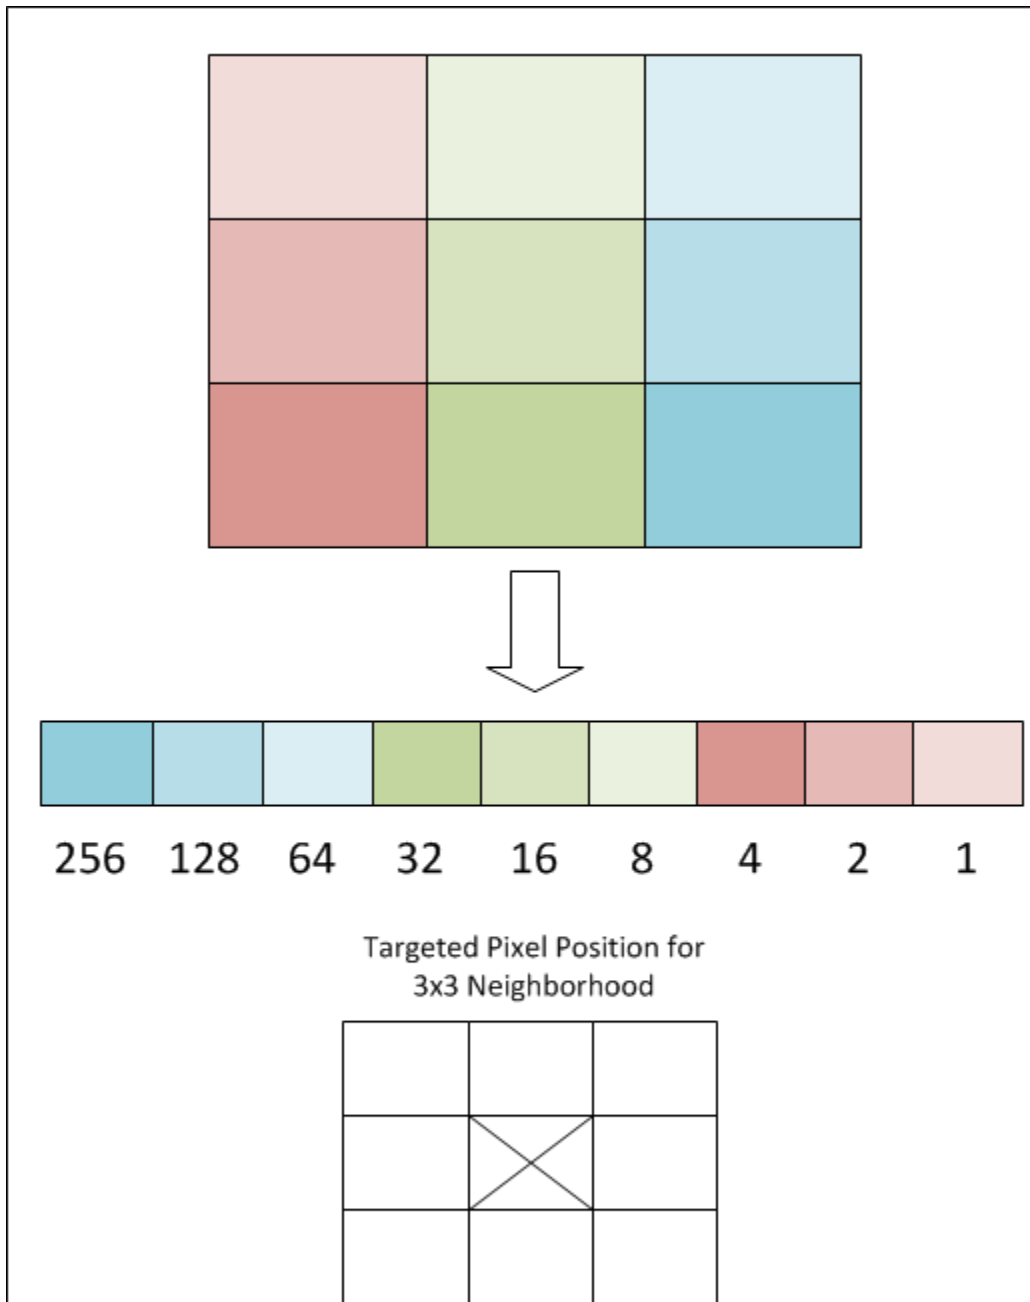


Figure 1: Color coded mapping from Pixel Patterns to Binary Representation for 3x3 Matrix

bwmorph

Purpose 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, returning the image minus the morphological closing of the image (dilation followed by erosion).
'branchpoints'	Find branch points of skeleton. For example: <pre>0 0 1 0 0 0 0 0 0 0 0 0 1 0 0 becomes 0 0 0 0 0 1 1 1 1 1 0 0 1 0 0 0 0 1 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0</pre>
'bridge'	Bridges unconnected pixels, that is, sets 0-valued pixels to 1 if they have two nonzero neighbors that are not connected. For example: <pre>1 0 0 1 1 0 1 0 1 becomes 1 1 1 0 0 1 0 1 1</pre>

Operation	Description
'clean'	Removes isolated pixels (individual 1s that are surrounded by 0s), such as the center pixel in this pattern. <pre>0 0 0 0 1 0 0 0 0</pre>
'close'	Performs morphological closing (dilation followed by erosion).
'diag'	Uses diagonal fill to eliminate 8-connectivity of the background. For example: <pre>0 1 0 0 1 0 1 0 0 becomes 1 1 0 0 0 0 0 0 0</pre>
'dilate'	Performs dilation using the structuring element ones(3).
'endpoints'	Finds end points of skeleton. For example: <pre>1 0 0 0 1 0 0 0 0 1 0 0 becomes 0 0 0 0 0 0 1 0 0 0 1 0 0 0 0 0 0 0 0 0</pre>
'erode'	Performs erosion using the structuring element ones(3).

Operation	Description
'fill'	<p>Fills isolated interior pixels (individual 0s that are surrounded by 1s), such as the center pixel in this pattern.</p> <pre> 1 1 1 1 0 1 1 1 1 </pre>
'hbreak'	<p>Removes H-connected pixels. For example:</p> <pre> 1 1 1 1 1 1 0 1 0 becomes 0 0 0 1 1 1 1 1 1 </pre>
'majority'	<p>Sets a pixel to 1 if five or more pixels in its 3-by-3 neighborhood are 1s; otherwise, it sets the pixel to 0.</p>
'open'	<p>Performs morphological opening (erosion followed by dilation).</p>
'remove'	<p>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.</p>
'shrink'	<p>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.</p>
'skel'	<p>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.</p>

Operation	Description
'spur'	Removes spur pixels. For example: <pre> 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 </pre>
'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. See “Algorithms” on page 3-75 for more detail.
'tophat'	Performs morphological "top hat" operation, returning the image minus the morphological opening of the image (erosion followed by dilation).

bwmorph supports the generation of efficient, production-quality C/C++ code from MATLAB. The text string specifying the operation must be a constant and, for best results, specify an input image of class `logical`. To see a complete list of toolbox functions that support code generation, see “Supported Functions”.

Class Support

The input image `BW` can be numeric or logical. It must be 2-D, real and nonsparse. The output image `BW2` is of class `logical`.


Examples

```


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

```


bwmorph



```
BW2 = bwmorph(BW,'remove');  
figure, imshow(BW2)
```



```
BW3 = bwmorph(BW,'skel',Inf);  
figure, imshow(BW3)
```

References

- [1] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Vol. 1, Addison-Wesley, 1992.
- [2] Kong, T. Yung and Azriel Rosenfeld, *Topological Algorithms for Digital Image Processing*, Elsevier Science, Inc., 1996.
- [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.
- [4] Pratt, William K., *Digital Image Processing*, John Wiley & Sons, Inc., 1991.

Algorithms

When used with the 'thin' option, bwmorph uses the following algorithm (References [3]):

- 1** Divide the image into two distinct subfields in a checkerboard pattern.
- 2** In the first subiteration, delete pixel p from the first subfield if and only if the conditions G_1 , G_2 , and G_3 are all satisfied.
- 3** In the second subiteration, delete pixel p from the second subfield if and only if the conditions G_1 , G_2 , and G_3' are all satisfied.

Condition G1:

$$X_H(p) = 1$$

where

$$X_H(p) = \sum_{i=1}^4 b_i$$

$$b_i = \begin{cases} 1, & \text{if } x_{2i-1} = 0 \text{ and } (x_{2i} = 1 \text{ or } x_{2i+1} = 1) \\ 0, & \text{otherwise} \end{cases}$$

x_1, x_2, \dots, x_8 are the values of the eight neighbors of p , starting with the east neighbor and numbered in counter-clockwise order.

Condition G2:

$$2 \leq \min\{n_1(p), n_2(p)\} \leq 3$$

where

$$n_1(p) = \sum_{k=1}^4 x_{2k-1} \vee x_{2k}$$

$$n_2(p) = \sum_{k=1}^4 x_{2k} \vee x_{2k+1}$$

Condition G3:

$$(x_2 \vee x_3 \vee \bar{x}_8) \wedge x_1 = 0$$

Condition G3':

$$(x_6 \vee x_7 \vee \bar{x}_4) \wedge x_5 = 0$$

The two subiterations together make up one iteration of the thinning algorithm. When the user specifies an infinite number of iterations (`n=Inf`), the iterations are repeated until the image stops changing. The conditions are all tested using `applylut` with precomputed lookup tables.

See Also

`bweuler` | `bwperim` | `imdilate` | `imerode`


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`.

Examples

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));
```

bwpack

See Also

bwunpack | imdilate | imerode

Purpose Find perimeter of objects 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.

Examples

Find the perimeter of objects in an image mask.

```
BW1 = imread('circbw.tif');  
BW2 = bwperim(BW1,8);  
imshow(BW1)  
figure, imshow(BW2)
```



See Also

[bwarea](#) | [bwboundaries](#) | [bweuler](#) | [bwtraceboundary](#) | [conndef](#) | [imfill](#)

Purpose Select objects in 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.



bwselect

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.

Examples

```
BW1 = imread('text.png');  
c = [43 185 212];  
r = [38 68 181];  
BW2 = bwselect(BW1,c,r,4);  
imshow(BW1), figure, imshow(BW2)
```



See Also

[bwlabel](#) | [imfill](#) | [impixel](#) | [roipoly](#) | [roifill](#)


Purpose Trace object in 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.

bwtraceboundary

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.

`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.

Examples

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)
```

```
        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

bwboundaries, bwperim
```

bwulterode

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.

Examples

```
originalBW = imread('circles.png');  
imshow(originalBW)  
ultimateErosion = bwulterode(originalBW);  
figure, imshow(ultimateErosion)
```

See Also

[bwdist](#) | [conndef](#) | [imregionalmax](#)

bwunpack

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`.

Examples 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`

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.

Examples Create a checkerboard where the side of every square is 20 pixels in length.


```
I = checkerboard(20);imshow(I)
```

checkerboard




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](#)

Purpose	Rearrange matrix columns into blocks
Syntax	<pre>A = col2im(B,[m n],[mm nn],'distinct')</pre> <pre>A = col2im(B,[m n],[mm nn],'sliding')</pre>
Description	<p><code>A = col2im(B,[m n],[mm nn],'distinct')</code> rearranges each column of B into a distinct m-by-n block to create the matrix A of size mm-by-nn. If $B = [A_{11}(:) \ A_{21}(:) \ A_{12}(:) \ A_{22}(:)]$, where each column has length $m \times n$, then $A = [A_{11} \ A_{12}; \ A_{21} \ A_{22}]$ where each A_{ij} is m-by-n.</p> <p><code>A = col2im(B,[m n],[mm nn],'sliding')</code> 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 <code>im2col(...,'sliding')</code> using a column compression function (such as <code>sum</code>).</p> <p><code>col2im(B,[m n],[mm nn])</code> is the same as <code>col2im(B, [m n], [mm nn],'sliding')</code>.</p>
Class Support	B can be logical or numeric. The return value A is of the same class as B.
Examples	<pre>B = reshape(uint8(1:25),[5 5])'</pre> <pre>C = im2col(B,[1 5])</pre> <pre>A = col2im(C,[1 5],[5 5],'distinct')</pre>
See Also	<code>blockproc</code> <code>colfilt</code> <code>im2col</code> <code>nlfilter</code>

colfilt

Purpose Columnwise neighborhood operations

Syntax

```
B = colfilt(A,[m n],block_type,fun)
B = colfilt(A,[m n],[mblock nblock],block_type,fun)
B = colfilt(A,'indexed',...)
```

Description `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. Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`. The function `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.

Note `colfilt` can perform operations similar to `blockproc` and `nlfilter`, but often executes much faster.

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

Value	Description
	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`.

Examples

Set 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

`blockproc` | `col2im` | `function_handle` | `im2col` | `nlfilter`

How To

- “Anonymous Functions”
- “Parameterizing Functions”

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
```

conndef

```
      0   1   0
      0   0   0

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	2-D convolution matrix
Syntax	<code>T = convmtx2(H,m,n)</code> <code>T = convmtx2(H,[m n])</code>
Description	<p><code>T = convmtx2(H,m,n)</code> returns the convolution matrix <code>T</code> for the matrix <code>H</code>. If <code>X</code> is an <code>m</code>-by-<code>n</code> matrix, then <code>reshape(T*X(:),size(H)+[m n]-1)</code> is the same as <code>conv2(X,H)</code>.</p> <p><code>T = convmtx2(H,[m n])</code> returns the convolution matrix, where the dimensions <code>m</code> and <code>n</code> are a two-element vector.</p>
Class Support	The inputs are all of class <code>double</code> . The output matrix <code>T</code> is of class <code>sparse</code> . The number of nonzero elements in <code>T</code> is no larger than <code>prod(size(H))*m*n</code> .
See Also	<code>conv2</code> <code>convmtx</code>

corner

Purpose

Find corner points in image

Syntax

```
C = corner(I)
C = corner(I, method)
C = corner(I, N)
C = corner(I, method, N)
C = corner(..., Name, Value)
```

Description

`C = corner(I)` detects corners in image `I` and returns them in matrix `C`.

`C = corner(I, method)` detects corners in image `I` using the specified method.

`C = corner(I, N)` detects corners in image `I` and returns a maximum of `N` corners.

`C = corner(I, method, N)` detects corners using the specified method and maximum number of corners.

`C = corner(..., Name, Value)` specifies parameters and corresponding values that control various aspects of the corner detection algorithm.

Tips

The `corner` and `cornermetric` functions both detect corners in images. For most applications, use the streamlined `corner` function to find corners in one step. If you want greater control over corner selection, use the `cornermetric` function to compute a corner metric matrix and then write your own algorithm to find peak values.

Input Arguments

I

A grayscale or binary image.

method

The algorithm used to detect corners. Supported methods are:

- 'Harris': The Harris corner detector.
- 'MinimumEigenvalue': Shi & Tomasi's minimum eigenvalue method.

Default: 'Harris'

N

The maximum number of corners the corner function can return.

Default: 200

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

FilterCoefficients

A vector, `V`, of filter coefficients for the separable smoothing filter. The outer product, `V*V'`, gives the full filter kernel.

Default: `fspecial('gaussian',[5 1],1.5)`

QualityLevel

A scalar value, `Q`, where $0 < Q < 1$, specifying the minimum accepted quality of corners. When candidate corners have corner metric values less than $Q * \max(\text{corner metric})$, the toolbox rejects them. Use larger values of `Q` to remove erroneous corners.

Default: 0.01

SensitivityFactor

A scalar value, `K`, where $0 < K < 0.25$, specifying the sensitivity factor used in the Harris detection algorithm. The smaller the value of `K`, the more likely the algorithm is to detect sharp corners. Use this parameter with the 'Harris' method only.

Default: 0.04

corner

Output Arguments

C

An M-by-2 matrix containing the X and Y coordinates of the corner points detected in I.


Class Support

I is a nonsparse numeric array. C is a matrix of class double.

Examples

Find and plot corner points in a checkerboard image.

```
I = checkerboard(50,2,2);  
C = corner(I);  
imshow(I);  
hold on  
plot(C(:,1), C(:,2), 'r*');
```



Algorithms

The corner function performs nonmaxima suppression on candidate corners, and corners are at least two pixels apart.

See Also

`cornermetric`

How To

- “Detecting Corners Using the `corner` Function”

cornermetric

Purpose

Create corner metric matrix from image

Description

`CM = cornermetric(I)` generates a corner metric matrix for the grayscale or logical image `I`. The corner metric, `CM`, is used to detect corner features in `I` and is the same size as `I`. Larger values in `CM` correspond to pixels in `I` with a higher likelihood of being a corner feature.

`CM = cornermetric(I, method)` generates a corner metric matrix for the grayscale or logical image `I` using the specified method. Valid values for `method` are:

Value	Description
'Harris'	The Harris corner detector. This is the default method.
'MinimumEigenvalue'	Shi and Tomasi's minimum eigenvalue method.

`CM = cornermetric(..., param1, val1, param2, val2, ...)` generates a corner metric matrix for `I`, specifying parameters and corresponding values that control various aspects of the corner metric calculation algorithm. Parameters include:

Parameter	Description
'FilterCoefficients'	A vector, <code>V</code> , of filter coefficients for the separable smoothing filter. This parameter is valid with the 'Harris' and 'MinimumEigenvalue' methods. The outer product, <code>V*V'</code> , gives the full filter kernel. The default is <code>fspecial('gaussian',[5 1],1.5)</code> .
'SensitivityFactor'	A scalar <code>k</code> , $0 < k < 0.25$, specifying the sensitivity factor used in the Harris detection algorithm. For smaller values of <code>k</code> , the algorithm is more likely to

Parameter	Description
	detect sharper corners. This parameter is only valid with the 'Harris' method. Default value: 0.04

Tips

The `corner` and `cornermetric` functions both detect corners in images. For most applications, use the streamlined `corner` function to find corners in one step. If you want greater control over corner selection, use the `cornermetric` function to compute a corner metric matrix and then write your own algorithm to find peak values.

Class Support

`I` is a nonsparse numeric array. `CM` is a matrix of class `double`.

Examples

Find corner features in grayscale image.

First generate a corner metric matrix.

```
I = imread('pout.tif');
I = I(1:150,1:120);
subplot(1,3,1);
imshow(I);
title('Original Image');
CM = cornermetric(I);
```

Adjust corner metric for viewing.

```
CM_adjusted = imadjust(CM);
subplot(1,3,2);
imshow(CM_adjusted);
title('Corner Metric');
```


Find and display some corner features.

```
corner_peaks = imregionalmax(CM);
corner_idx = find(corner_peaks == true);
[r g b] = deal(I);
```

cornermetric

```
r(corner_idx) = 255;  
g(corner_idx) = 255;  
b(corner_idx) = 0;  
RGB = cat(3,r,g,b);  
subplot(1,3,3);  
imshow(RGB);  
title('Corner Points');
```

The following figure shows the output of these operations.



See Also

[corner](#) | [edge](#) | [immovie](#) | [imshow](#)

Purpose	2-D correlation coefficient
Syntax	<code>r = corr2(A,B)</code>
Description	<code>r = corr2(A,B)</code> 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.

Examples

Compute the correlation coefficient between an image and the same image processed with a median filter.

```
I = imread('pout.tif');
J = medfilt2(I);
R = corr2(I,J)
```

```
R =
```

```
0.9959
```

Algorithms

`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`

Purpose

Infer spatial transformation from control point pairs

Syntax

```
TFORM = cp2tform(input_points, base_points, transformtype)
TFORM = cp2tform(CPSTRUCT, transformtype)
[TFORM, input_points, base_points] =
cp2tform(CPSTRUCT, ...)
TFORM = cp2tform(input_points, base_points, 'polynomial',
    order)
TFORM = cp2tform(CPSTRUCT, 'polynomial', order)
TFORM = cp2tform(input_points, base_points,
    'piecewise linear')
TFORM = cp2tform(CPSTRUCT, 'piecewise linear')
TFORM = cp2tform(input_points, base_points, 'lwm', N)
TFORM = cp2tform(CPSTRUCT, 'lwm', N)
[TFORM, input_points, base_points, input_points_bad,
    base_points_bad] = cp2tform(input_points, base_points,
    'piecewise linear')
[TFORM, input_points, base_points, input_points_bad,
    base_points_bad] = cp2tform(CPSTRUCT,
    'piecewise linear')
```

Description

`TFORM = cp2tform(input_points, base_points, transformtype)` infers a spatial transformation from control point pairs and returns this transformation as a TFORM structure.

`TFORM = cp2tform(CPSTRUCT, transformtype)` works on a CPSTRUCT structure that contains the control point matrices for the input and base images. The Control Point Selection Tool, `cpselect`, creates the CPSTRUCT.

`[TFORM, input_points, base_points] = cp2tform(CPSTRUCT, ...)` returns the control points that were used in `input_points` and `base_points`. Unmatched and predicted points are not used. See `cpstruct2pairs`.

`TFORM = cp2tform(input_points, base_points, 'polynomial', order)` lets you specify the order of the polynomials to use.

`TFORM = cp2tform(CPSTRUCT, 'polynomial', order)` works on a `CPSTRUCT` structure.

`TFORM = cp2tform(input_points, base_points, 'piecewise linear')` creates a Delaunay triangulation of the base control points, and maps corresponding input control points to the base control points. The mapping is linear (affine) for each triangle and continuous across the control points but not continuously differentiable as each triangle has its own mapping.

`TFORM = cp2tform(CPSTRUCT, 'piecewise linear')` works on a `CPSTRUCT` structure.

`TFORM = cp2tform(input_points, base_points, 'lwm', N)` creates a mapping by inferring a polynomial at each control point using neighboring control points. The mapping at any location depends on a weighted average of these polynomials. You can optionally specify the number of points, `N`, used to infer each 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)` works on a `CPSTRUCT` structure.

`[TFORM, input_points, base_points, input_points_bad, base_points_bad] = cp2tform(input_points, base_points, 'piecewise linear')` returns the control points that were used in `input_points` and `base_points` and 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')` works on a `CPSTRUCT` structure.

Tips

- When `transformtype` is `'nonreflective similarity'`, `'similarity'`, `'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` have more than the minimum number of control points, a least-squares solution is found. See `mldivide`.
- When either `input_points` or `base_points` has a large offset with respect to their origin (relative to range of values that it spans), `cp2tform` shifts the points to center their bounding box on the origin before fitting a `TFORM` structure. This enhances numerical stability and is handled transparently by wrapping the origin-centered `TFORM` within a custom `TFORM` that automatically applies and undoes the coordinate shift as needed. As a result, `fields(T)` can give different results for different coordinate inputs, even for the same transformation type.

Input Arguments

input_points

m-by-2, double matrix containing the *x*- and *y*-coordinates of control points in the image you want to transform.

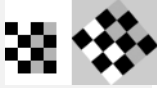
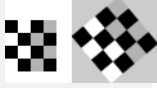



base_points



m-by-2, double matrix containing the *x*- and *y*-coordinates of control points in the base image.

transformtype

Specifies the type of spatial transformation to infer. The `cp2tform` function requires a minimum number of control point pairs to infer a structure of each transform type. The following 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.

Transformation Types

Transformation Type	Description	Minimum Number of Control Point Pairs	Example
'nonreflective similarity'	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	
'similarity'	Same as 'nonreflective similarity' with the addition of optional reflection.	3	
'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	
'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	
'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 (order 2) 10 (order 3) 15 (order 4)	

Transformation Type	Description	Minimum Number of Control Point Pairs	Example
'piecewise linear'	Use this transformation when parts of the image appear distorted differently.	4	
'lwm'	Use this transformation (local weighted mean), when the distortion varies locally and piecewise linear is not sufficient.	6 (12 recommended)	

CPSTRUCT

Structure containing control point matrices for the input and base images. Use the Control Point Selection Tool (cpselect) to create the CPSTRUCT.

'polynomial',order

Specifies the order of polynomials to use. order can be 2, 3, or 4.

Default: 3

'piecewise linear'

Linear for each piece and continuous, not continuously differentiable.

'lwm'

Local weighted mean.

N

Number of points.

Output Arguments

TFORM

Structure containing the spatial transformation.

input_points

Input control points that were used to infer the spatial transformation. Unmatched and predicted points are not used.

base_points

Base control points that were used to infer the spatial transformation. Unmatched and predicted points are not used.

input_points_bad

Input control points that were eliminated because they were determined to be outliers.

base_points_bad

Base control points that were eliminated because they were determined to be outliers.

Examples

Transform an image, use the `cp2tform` function to return the transformation, and compare the angle and scale of the TFORM to the angle and scale of the original transformation:

```
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,'nonreflective similarity');

% Recover angle and scale by checking how a unit vector
% parallel to the x-axis is rotated and stretched.
u = [0 1];
v = [0 0];
[x, y] = tformfwd(t, u, v);
dx = x(2) - x(1);
dy = y(2) - y(1);
```

```
angle = (180/pi) * atan2(dy, dx)
scale = 1 / sqrt(dx^2 + dy^2)
```

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 (x,y) according to transformtype.
- 2 Return the TFORM structure containing spatial transformation.

The procedure varies depending on the transformtype.

Nonreflective Similarity

Nonreflective similarity 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, and ty.

Similarity

Similarity transformations can include rotation, scaling, translation, and reflection. Shapes and angles are preserved. Parallel lines remain parallel. Straight lines remain straight.

Let

```
sc = s*cos(theta)
ss = s*sin(theta)
```


$$[u \ v] = [x \ y \ 1] * \begin{bmatrix} sc & -a*ss \\ ss & a*sc \\ tx & ty \end{bmatrix}$$

Solve for sc , ss , tx , ty , and a . If $a = -1$, reflection is included in the transformation. If $a = 1$, reflection is not included in the transformation.

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. Nonreflective similarity transformations are a subset of affine transformations.

For an affine transformation,

$$[u \ v] = [x \ y \ 1] * T_{inv}$$

T_{inv} is a 3-by-2 matrix. Solve for the six elements of T_{inv} :

```
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,

$$[up \ vp \ wp] = [x \ y \ w] * T_{inv}$$

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.

Note An affine or projective transformation can also be expressed like this, for a 3-by-2 Tinv:

$$[u \ v]^T = Tinv^T * [x \ y \ 1]^T$$

Or, like this, for a 3-by-3 Tinv:

$$[u \ v \ 1]^T = Tinv^T * [x \ y \ 1]^T$$

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:

$$[u \ v] = [1 \ x \ y \ x*y \ x^2 \ y^2 \ y*x^2 \ x*y^2 \ x^3 \ y^3] * Tinv$$

Both u and v are third-order polynomials of x and y . Each third-order polynomial has 10 terms. To specify all coefficients, $Tinv$ 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:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} 1 & x & y & x*y & x^2 & y^2 & y*x^2 & x*y^2 & x^3 & y^3 & x^3*y & x^2*y^2 & x*y^3 & x^4 \\ y^4 \end{bmatrix} * T_{inv}$$

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.

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.

See Also

`cpcorr` | `cpselect` | `cpstruct2pairs` | `imtransform` | `tformfwd`
| `tforminv`

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`.

Algorithms

`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.

Examples

Use `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,...)
cpselect(...,param1, val1,...)
```

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 grayscale, truecolor, or binary images, or strings that identify files containing these images. The Control Point Selection Tool returns the control points in a `CPSTRUCT` structure. (For more information, see “Specifying Control Points Using the Control Point Selection Tool”.)

`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)` syntax to close the tool from the command line.

`cpselect(...,param1, val1,...)` starts `cpselect`, specifying parameters and corresponding values that control various aspects of the tool. Parameter names can be abbreviated, and case does not matter. Parameters include:

Parameter	Description
'Wait'	<p>Logical scalar that controls whether cpselect waits for the user to finish the task of selecting points. If set to <code>false</code> (the default), you can run cpselect at the same time as you run other programs in MATLAB. If set to <code>true</code>, you must finish the task of selecting points before doing anything else in MATLAB.</p> <hr/> <p>Note When 'Wait' is set to true, cpselect returns the selected pairs of points, not a handle to the tool:</p> <pre>[xyinput_out, xybase_out] = cpselect(...,'Wait', true) xyinput_out and xybase_out are <i>P</i>-by-2 matrices that store the input and base image coordinates, respectively.</pre>

Class Support

The images can be grayscale, truecolor, or binary. A grayscale image can be `uint8`, `uint16`, `int16`, `single`, or `double`. A truecolor image can be `uint8`, `uint16`, `single`, or `double`. A binary image is of class `logical`.

Algorithms

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	Nonreflective similarity
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.

Examples

Start Control Point Selection tool with saved images.

```
cpselect('westconcordaerial.png','westconcordorthophoto.png')
```

Start Control Point Selection tool with images and control points stored in variables in the workspace.

```
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);
```

Use `cpselect` in a script, specifying the `'wait'` parameter to block until control point selection is complete.

```
aerial = imread('westconcordaerial.png');  
figure, imshow(aerial)  
figure, imshow('westconcordorthophoto.png')  
load westconcordpoints % load preselected control points  
  
% Block the tool until you pick some more control points  
[aerial_points,ortho_points] = ...  
    cpselect(aerial,'westconcordorthophoto.png',...  
            input_points,base_points,...  
            'Wait',true);
```

```
t_concord = cp2tform(aerial_points,ortho_points,'projective');
info = imfinfo('westconcordorthophoto.png');
aerial_registered = imtransform(aerial, t_concord,...
                               'XData',[1 info.Width],...
                               'YData',[1 info.Height]);
figure, imshow(aerial_registered)
```

See Also

[cpcorr](#) | [cp2tform](#) | [cpstruct2pairs](#) | [imtransform](#)

How To

- “Geometric Transformation and Image Registration”

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.

Examples Start the Control Point Selection Tool, `cpselect`.

```
aerial = imread('westconcordaerial.png');  
cpselect(aerial(:,:,1), 'westconcordorthophoto.png')
```

Using `cpselect`, pick control points in the images. Select **Export Points to Workspace** from the **File** menu to save the points to the workspace. On the **Export Points to Workspace** 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

2-D discrete cosine transform

Syntax

B = dct2(A)
 B = dct2(A,m,n)
 B = dct2(A,[m n])

Description

B = dct2(A) returns the two-dimensional discrete cosine transform of A. The matrix B is the same size as A and contains the discrete cosine transform coefficients $B(k_1, k_2)$.

B = dct2(A,m,n) pads the matrix A with 0's to size m-by-n before transforming. If m or n is smaller than the corresponding dimension of A, dct2 truncates A.

B = dct2(A,[m n]) same as above.

Class Support

A can be numeric or logical. The returned matrix B is of class double.

Algorithms

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 A and output image B is

$$B_{pq} = \alpha_p \alpha_q \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} A_{mn} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad \begin{matrix} 0 \leq p \leq M-1 \\ 0 \leq q \leq N-1 \end{matrix}$$

where

$$\alpha_p = \begin{cases} \frac{1}{\sqrt{M}}, & p = 0 \\ \sqrt{\frac{2}{M}}, & 1 \leq p \leq M-1 \end{cases}$$

and

$$\alpha_q = \begin{cases} \frac{1}{\sqrt{N}}, & q = 0 \\ \sqrt{\frac{2}{N}}, & 1 \leq q \leq N-1 \end{cases}$$


M and N are the row and column size of A , 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.

This transform can be inverted using `idct2`.

Examples

The commands below compute the discrete cosine transform for the autumn image. Notice that most of the energy is in the upper left corner.

```
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])
```



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.

See Also

`fft2` | `idct2` | `ifft2`

dctmtx

Purpose Discrete cosine transform matrix

Syntax `D = dctmtx(n)`

Description `D = dctmtx(n)` 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 `double`. D is returned as a matrix of class `double`.

Tips If A is square, the two-dimensional DCT of A can be computed as $D*A*D'$. This computation is sometimes faster than using `dct2`, especially if you are computing a large number of small DCTs, because D needs to be determined only once.

For example, in JPEG compression, the DCT of each 8-by-8 block is computed. To perform this computation, use `dctmtx` 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 `dct2` for each individual block.

Examples

```
A = im2double(imread('rice.png'));
D = dctmtx(size(A,1));
dct = D*A*D';
figure, imshow(dct)
```

See Also `dct2`

Purpose Deblur image using blind deconvolution

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, P1, P2,...,PN)
```

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 restored PSF is a positive array that is the same size as INITPSF, normalized so its sum adds up to 1. 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.

I can be a N-dimensional array.

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).

[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.

Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`.

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.

`I` and `INITPSF` can be `uint8`, `uint16`, `int16`, `single`, or `double`. `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`.

Examples

```
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));
[J P] = deconvblind(BlurredNoisy,INITPSF,20,10*sqrt(V),WT);
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`

How To

- “Anonymous Functions”
- “Parameterizing Functions”

deconvlucy

Purpose Deblur image using Lucy-Richardson method

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.

`I` can be a `N`-dimensional array.

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`.

Resuming Deconvolution

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.

`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`.

Examples

```
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');
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)');
```

References

- [1] Biggs, D.S.C. “Acceleration of Iterative Image Restoration Algorithms.” *Applied Optics*. Vol. 36. Number 8, 1997, pp. 1766–1775.
- [2] Hanisch, R.J., R.L. White, and R.L. Gilliland. “Deconvolution of Hubble Space Telescope Images and Spectra.” *Deconvolution of Images and Spectra* (P.A. Jansson, ed.). Boston, MA: Academic Press, 1997, pp. 310–356.

See Also

deconvblind | deconvreg | deconvwnr | otf2psf | padarray |
psf2otf

Purpose

Deblur image using regularized filter

Syntax

```
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, ...)
```

Description

`J = deconvreg(I, PSF)` deconvolves image `I` using the regularized filter algorithm, returning deblurred image `J`. The assumption is that the image `I` was created by convolving a true image with a point-spread function `PSF` and possibly by adding noise. The algorithm is a constrained optimum in the sense of least square error between the estimated and the true images under requirement of preserving image smoothness.

`I` can be a `N`-dimensional array.

`J = deconvreg(I, PSF, NOISEPOWER)` where `NOISEPOWER` is the additive noise power. The default value is 0.

`J = deconvreg(I, PSF, NOISEPOWER, 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]`.

`J = deconvreg(I, PSF, NOISEPOWER, LRANGE, 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`.

`[J, LAGRA] = deconvreg(I, PSF, ...)` outputs the value of the Lagrange multiplier `LAGRA` in addition to the restored image `J`.

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 = edgetaper(I,PSF)`.

Class Support

`I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. Other inputs have to be of class `double`. `J` has the same class as `I`.

Examples

```
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 Deblur image using Wiener filter

Syntax
`J = deconvwnr(I,PSF,NSR)`
`J = deconvwnr(I,PSF,NCORR,ICORR)`

Description `J = deconvwnr(I,PSF,NSR)` deconvolves image `I` using the Wiener filter algorithm, returning deblurred image `J`. Image `I` can be an `N`-dimensional array. `PSF` is the point-spread function with which `I` was convolved. `NSR` is the noise-to-signal power ratio of the additive noise. `NSR` can be a scalar or a spectral-domain array of the same size as `I`. Specifying 0 for the `NSR` is equivalent to creating an ideal inverse filter.

The algorithm is optimal in a sense of least mean square error between the estimated and the true images.

`J = deconvwnr(I,PSF,NCORR,ICORR)` deconvolves image `I`, where `NCORR` is the autocorrelation function of the noise and `ICORR` is the autocorrelation function of the original image. `NCORR` and `ICORR` can be of any size or dimension, not exceeding the original image. If `NCORR` or `ICORR` are `N`-dimensional arrays, the values correspond to the autocorrelation within each dimension. If `NCORR` or `ICORR` are vectors, and `PSF` is also a vector, the values represent the autocorrelation function in the first dimension. If `PSF` is an array, the 1-D autocorrelation function is extrapolated by symmetry to all non-singleton dimensions of `PSF`. If `NCORR` or `ICORR` is a scalar, this value represents the power of the noise of the image.

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)`

prior to calling `deconvwnr`.

Class Support

I can be of class uint8, uint16, int16, single, or double. Other inputs have to be of class double. J has the same class as I.

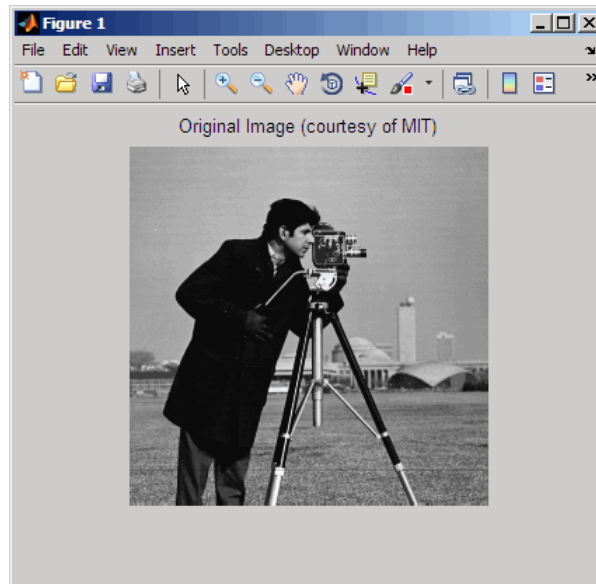
Examples

Use deconvwnr to Restore an Image

Create a noisy, blurry image and then apply the deconvwnr filter to deblur it.

Display the original image.

```
I = im2double(imread('cameraman.tif'));  
imshow(I);  
title('Original Image (courtesy of MIT)');
```




Simulate a motion blur.

```
LEN = 21;  
THETA = 11;  
PSF = fspecial('motion', LEN, THETA);
```

```
blurred = imfilter(I, PSF, 'conv', 'circular');  
figure, imshow(blurred)
```


Simulate additive noise.

```
noise_mean = 0;  
noise_var = 0.0001;  
blurred_noisy = imnoise(blurred, 'gaussian', ...  
                        noise_mean, noise_var);  
figure, imshow(blurred_noisy)  
title('Simulate Blur and Noise')
```



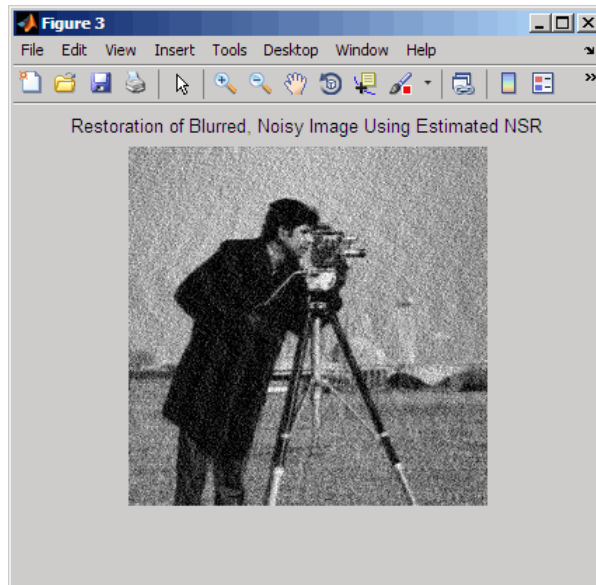
Try restoration assuming no noise.

```
estimated_nsr = 0;  
wnr2 = deconvwnr(blurred_noisy, PSF, estimated_nsr);  
figure, imshow(wnr2)  
title('Restoration of Blurred, Noisy Image Using NSR = 0')
```



Try restoration using a better estimate of the noise-to-signal-power ratio.

```
estimated_nsr = noise_var / var(I(:));  
wnr3 = deconvwnr(blurred_noisy, PSF, estimated_nsr);  
figure, imshow(wnr3)  
title('Restoration of Blurred, Noisy Image Using Estimated NSR');
```



References

"Digital Image Processing", R. C. Gonzalez & R. E. Woods, Addison-Wesley Publishing Company, Inc., 1992.

See Also

deconvblind | deconvlucy | deconvreg | edgetaper | otf2psf | padarray | psf2otf

decorrstretch

Purpose Apply decorrelation stretch to 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.

Examples

```
[X, map] = imread('forest.tif');  
S = decorrstretch(ind2rgb(X, map), 'tol', 0.01);  
figure, imshow(X, map)  
figure, imshow(S)
```

See Also imadjust | stretchlim

Purpose Convert Bayer pattern encoded image to truecolor image

Syntax `RGB = demosaic(I, sensorAlignment)`

Description `RGB = demosaic(I, sensorAlignment)` converts a Bayer pattern encoded image to a truecolor image using gradient-corrected linear interpolation. `I` is an M-by-N array of intensity values that are Bayer pattern encoded. `I` must have at least 5 rows and 5 columns.

A Bayer filter mosaic, or color filter array, refers to the arrangement of color filters that let each sensor in a single-sensor digital camera record only red, green, or blue data. The patterns emphasize the number of green sensors to mimic the human eye's greater sensitivity to green light. The `demosaic` function uses interpolation to convert the two-dimensional Bayer-encoded image into the truecolor image, `RGB`, which is an M-by-N-by-3 array.

`sensorAlignment` is one of the following text strings that specifies the Bayer pattern. Each string represents the order of the red, green, and blue sensors by describing the four pixels in the upper-left corner of the image (left-to-right, top-to-bottom).

demosaic

Value	2-by-2 Sensor Alignment				
'gbrg'	<table border="1"><tr><td>Green</td><td>Blue</td></tr><tr><td>Red</td><td>Green</td></tr></table>	Green	Blue	Red	Green
Green	Blue				
Red	Green				
'grbg'	<table border="1"><tr><td>Green</td><td>Red</td></tr><tr><td>Blue</td><td>Green</td></tr></table>	Green	Red	Blue	Green
Green	Red				
Blue	Green				
'bggr'	<table border="1"><tr><td>Blue</td><td>Green</td></tr><tr><td>Green</td><td>Red</td></tr></table>	Blue	Green	Green	Red
Blue	Green				
Green	Red				
'rggb'	<table border="1"><tr><td>Red</td><td>Green</td></tr><tr><td>Green</td><td>Blue</td></tr></table>	Red	Green	Green	Blue
Red	Green				
Green	Blue				

**Class
Support**

I can be uint8 or uint16, and it must be real. RGB has the same class as I.

Examples

Convert a Bayer pattern encoded image that was photographed by a camera with a sensor alignment of 'bggr'.

```
I = imread('mandi.tif');  
J = demosaic(I,'bggr');  
imshow(I);  
figure, imshow(J);
```



Bayer Pattern Encoded Image

demosaic



Truecolor Image

Purpose

Anonymize 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 whose fields are the names of the attributes to preserve. 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', ...
```

dicomanon

```
{'PatientAge', 'PatientSex', 'StudyDescription'})
```

Anonymize a series of images, keeping the hierarchy.

```
values.StudyInstanceUID = dicomuid;  
values.SeriesInstanceUID = 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](#)

Purpose Get or set active DICOM data dictionary

Syntax

```
dicomdict('set',dictionary)
dictionary = dicomdict('get')
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.

`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.

Examples

```
dictionary = dicomdict('get')

dictionary =

dicom-dict.mat
```

See Also `dicominfo` | `dicomread` | `dicomwrite`

dicominfo

Purpose Read metadata from DICOM message

Syntax
`info = dicominfo(filename)`
`info = dicominfo(filename, 'dictionary', D)`

Description `info = dicominfo(filename)` reads the metadata from the compliant Digital Imaging and Communications in Medicine (DICOM) file specified in the string `filename`.

`info = dicominfo(filename, 'dictionary', D)` uses the data dictionary file given in the string `D` to read the DICOM message. The file in `D` must be on the MATLAB search path. The default file is `dicom-dict.mat`.

Examples `info = dicominfo('CT-MONO2-16-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'  
    .  
    .  
    .
```

See Also

`dicomdict` | `dicomread` | `dicomwrite` | `dicomuid`

dicomlookup

Purpose Find attribute in DICOM data dictionary

Syntax
`name = dicomlookup(group, element)`
`[group, element] = dicomlookup(name)`

Description `name = dicomlookup(group, element)` looks into the current DICOM data dictionary for the attribute with the specified `group` and `element` tag and returns a string containing the name of the attribute. `group` and `element` can contain either a decimal value or hexadecimal string.

`[group, element] = dicomlookup(name)` looks into the current DICOM data dictionary for the attribute specified by `name` and returns the `group` and `element` tags associated with the attribute. The values are returned as decimal values.

Examples Find the names of DICOM attributes using their tags.

```
name1 = dicomlookup('7FE0', '0010')
name2 = dicomlookup(40, 4)
```

Look up a DICOM attribute's tag (GROUP and ELEMENT) using its name.

```
[group, element] = dicomlookup('TransferSyntaxUID')
```

Examine the metadata of a DICOM file. This returns the same value even if the data dictionary changes.

```
metadata = dicominfo('CT-MON02-16-ankle.dcm');
metadata.(dicomlookup('0028', '0004'))
```

See Also `dicomdict` | `dicominfo`

Purpose

Read 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'`.

dicomread

Tips

The `dicomread` function supports both reversible (lossless) and irreversible (lossy) JPEG-2000 compression in DICOM files.

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

Use `dicomread` to retrieve the data array, `X`, and colormap matrix, `map`, needed to create a montage.

```
[X, map] = dicomread('US-PAL-8-10x-echo.dcm');  
montage(X, map, 'Size', [2 5]);
```

Call `dicomread` with the information retrieved from the DICOM file using `dicominfo` and display the image using `imshow`. Adjust the contrast of the image using `imcontrast`.

```
info = dicominfo('CT-MONO2-16-ankle.dcm');  
Y = dicomread(info);  
figure, imshow(Y);  
imcontrast;
```

See Also

`dicomdict` | `dicominfo` | `dicomwrite`

Purpose	Generate DICOM unique identifier
Syntax	UID = dicomuid
Description	<p>UID = dicomuid creates a string UID containing a new DICOM unique identifier.</p> <p>Multiple calls to dicomuid produce globally unique values. Two calls to dicomuid always return different values.</p>
See Also	dicominfo dicomwrite

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 optional 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 software. The following table lists the options that you can specify, in alphabetical order. Default values are enclosed in braces (`{}`).

Option Name	Description
'CompressionMode'	<p>String specifying the type of compression to use when storing the image. Possible values:</p> <pre>{ 'None' } 'JPEG lossless' 'JPEG lossy' 'JPEG2000 lossy' 'JPEG2000 lossless' 'RLE'</pre>
'CreateMode'	<p>Specifies the method used for creating the data to put in the new file. Possible values:</p> <pre>{ 'Create' } — Verify input values and generate missing data values. 'Copy' — Copy all values from the input and do not generate missing values.</pre>
'Dictionary'	<p>String specifying the name of a DICOM data dictionary.</p>
'Endian'	<p>String specifying the byte ordering of the file.</p> <pre>'Big' {'Little'}</pre> <hr/> <p>Note If VR is set to 'Explicit', 'Endian' must be 'Big'. dicomwrite ignores this value if 'CompressionMode' or 'TransferSyntax' is set.</p> <hr/>

Option Name	Description
'MultiframeSingleFile'	Logical value indicating whether multiframe imagery should be written to one file. When true (default), one file is created regardless of how many frames X contains. When false, one file is written for each frame in the image.
'TransferSyntax'	A DICOM UID specifying the 'Endian', 'VR', and 'CompressionMode' options. Note If specified, dicomwrite ignores any values specified for the 'Endian', 'VR', and 'CompressionMode' options. The TransferSyntax value encodes values for these options.
'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 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.

Field	Description
'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.

Tips

The DICOM format specification lists several Information Object Definitions (IODs) that can be created. These IODs correspond to images and metadata produced by different real-world modalities (e.g., MR, X-ray, Ultrasound, etc.). For each type of IOD, the DICOM specification defines the set of metadata that must be present and possible values for other metadata.

`dicomwrite` fully implements a limited number of these IODs, listed above in the `ObjectType` syntax. For these IODs, `dicomwrite` verifies that all required metadata attributes are present, creates missing attributes if necessary, and specifies default values where possible. Using these supported IODs is the best way to ensure that the files you create conform to the DICOM specification. This is `dicomwrite` default behavior and corresponds to the `CreateMode` option value of `'Create'`.

To write DICOM files for IODs that `dicomwrite` doesn't implement, use the `'Copy'` value for the `CreateMode` option. In this mode, `dicomwrite` writes the image data to a file including the metadata that you specify as a parameter, shown above in the `info` syntax. The purpose of this option is to take metadata from an existing file of the same modality or IOD and use it to create a new DICOM file with different image pixel data.

Note Because `dicomwrite` copies metadata to the file without verification in 'copy' mode, it is possible to create a DICOM file that does not conform to the DICOM standard. For example, the file may be missing required metadata, contain superfluous metadata, or the metadata may no longer correspond to the modality settings used to generate the original image. When using 'Copy' mode, make sure that the metadata you use is from the same modality and IOD. If the copy you make is unrelated to the original image, use `dicomuid` to create new unique identifiers for series and study metadata. See the IOD descriptions in Part 3 of the DICOM specification for more information on appropriate IOD values.

Examples

Read a CT image from the sample DICOM file included with the toolbox and then write the CT image to a file, creating a secondary capture image.

```
X = dicomread('CT-MONO2-16-ankle.dcm');  
dicomwrite(X, 'sc_file.dcm');
```

Write the CT image, `X`, to a DICOM file along with its metadata. 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. In this mode, `dicomwrite` does not verify the metadata written to the file.

```
dicomwrite(X, 'ct_copy.dcm', metadata, 'CreateMode', 'copy');
```

See Also

`dicomdict` | `dicominfo` | `dicomread` | `dicomuid`

edge

Purpose

Find edges in grayscale image

Note

- The syntax `BW = edge(...,K)` has been removed. Use the `BW = edge(...,direction)` syntax instead.
 - The syntax `edge(I,'marr-hildreth',...)` has been removed. Use the `edge(I,'log',...)` syntax instead.
-

Syntax

```
BW = edge(I)
```

```
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

`BW = edge(I)` takes a grayscale or a binary 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.

By default, `edge` uses the Sobel method to detect edges but the following provides a complete list of all the edge-finding methods supported by this function:

- 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.

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' (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 = edge(I, 'roberts', ..., options)` where `options` can be the text string `'thinning'` or `'nothinning'`. When you specify `'thinning'`, or don't specify a value, the algorithm applies edge thinning. Specifying the `'nothinning'` option can speed up the operation of the algorithm by skipping the additional edge thinning stage.

`[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. If you specify a threshold of 0, the output image has closed contours, because it includes all the zero crossings in the input image.

`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. If you specify a threshold of 0, the output image has closed contours, because it includes all the zero crossings in the input image.

`[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 scalar 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. The value for `thresh` is relative to the highest value of the gradient magnitude of the image.

`BW = edge(I, 'canny', thresh, sigma)` specifies the Canny method, using `sigma` as the standard deviation of the Gaussian filter. The default `sigma` is `sqrt(2)`; 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`.

Tips

For the gradient-magnitude methods (Sobel, Prewitt, Roberts), `thresh` is used to threshold the calculated gradient magnitude. For the zero-crossing methods, including Lap, `thresh` is used as a threshold for the zero-crossings; in other words, a large jump across zero is an edge, while a small jump isn't.

The Canny method applies two thresholds to the gradient: a high threshold for low edge sensitivity and a low threshold for high edge sensitivity. `edge` starts with the low sensitivity result and then grows it to include connected edge pixels from the high sensitivity result. This helps fill in gaps in the detected edges.

In all cases, the default threshold is chosen heuristically in a way that depends on the input data. The best way to vary the threshold is to run `edge` once, capturing the calculated threshold as the second output argument. Then, starting from the value calculated by `edge`, adjust the threshold higher (fewer edge pixels) or lower (more edge pixels).


The function `edge` changed in Version 7.2 (R2011a). Previous versions of the Image Processing Toolbox used a different algorithm for computing the Canny method. If you need the same results produced by the previous implementation, use the following syntax:

```
BW = edge(I, 'canny_old', ...)
```

Examples


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);
```



Prewitt Method

```
figure, imshow(BW2)
```



Canny Method

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.

See Also

`fspecial` | `imgradient` | `imgradientxy`

Purpose	Taper discontinuities along image edges
Syntax	<code>J = edgetaper(I,PSF)</code>
Description	<p><code>J = edgetaper(I,PSF)</code> blurs the edges of the input image <code>I</code> using the point spread function <code>PSF</code>. The size of the <code>PSF</code> cannot exceed half of the image size in any dimension.</p> <p>The output image <code>J</code> is the weighted sum of the original image <code>I</code> and its blurred version. The weighting array, determined by the autocorrelation function of <code>PSF</code>, makes <code>J</code> equal to <code>I</code> in its central region, and equal to the blurred version of <code>I</code> near the edges.</p> <p>The <code>edgetaper</code> function reduces the ringing effect in image deblurring methods that use the discrete Fourier transform, such as <code>deconvwnr</code>, <code>deconvreg</code>, and <code>deconvlucy</code>.</p>
Class Support	<code>I</code> and <code>PSF</code> can be of class <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> . <code>J</code> is of the same class as <code>I</code> .
Examples	<pre>original = imread('cameraman.tif'); PSF = fspecial('gaussian',60,10); edgesTapered = edgetaper(original,PSF); figure, imshow(original,[]); figure, imshow(edgesTapered,[]);</pre>
See Also	<code>deconvlucy</code> <code>deconvreg</code> <code>deconvwnr</code> <code>otf2psf</code> <code>padarray</code> <code>psf2otf</code>

entropy

Purpose	Entropy of grayscale image
Syntax	$E = \text{entropy}(I)$
Description	<p>$E = \text{entropy}(I)$ returns E, a scalar value representing the entropy of grayscale 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</p> $-\sum(p.\log_2(p))$ <p>where p contains the histogram counts returned from <code>imhist</code>. By default, <code>entropy</code> uses two bins for logical arrays and 256 bins for <code>uint8</code>, <code>uint16</code>, or <code>double</code> arrays.</p> <p>I can be a multidimensional image. If I has more than two dimensions, the <code>entropy</code> function treats it as a multidimensional grayscale image and not as an RGB image.</p>
Class Support	I can be <code>logical</code> , <code>uint8</code> , <code>uint16</code> , or <code>double</code> and must be real, nonempty, and nonsparse. E is <code>double</code> .
Notes	<code>entropy</code> converts any class other than <code>logical</code> to <code>uint8</code> for the histogram count calculation so that the pixel values are discrete and directly correspond to a bin value.
Examples	<pre>I = imread('circuit.tif'); J = entropy(I)</pre>
References	[1] Gonzalez, R.C., R.E. Woods, S.L. Eddins, <i>Digital Image Processing Using MATLAB</i> , New Jersey, Prentice Hall, 2003, Chapter 11.
See Also	<code>imhist</code> <code>entropyfilt</code>

Purpose	Local entropy of grayscale image
Syntax	<pre>J = entropyfilt(I) J = entropyfilt(I,NHOOD)</pre>
Description	<p><code>J = entropyfilt(I)</code> returns the array <code>J</code>, where each output pixel contains the entropy value of the 9-by-9 neighborhood around the corresponding pixel in the input image <code>I</code>. <code>I</code> can have any dimension. If <code>I</code> has more than two dimensions, <code>entropyfilt</code> treats it as a multidimensional grayscale image and not as a truecolor (RGB) image. The output image <code>J</code> is the same size as the input image <code>I</code>.</p> <p>For pixels on the borders of <code>I</code>, <code>entropyfilt</code> uses symmetric padding. In symmetric padding, the values of padding pixels are a mirror reflection of the border pixels in <code>I</code>.</p> <p><code>J = entropyfilt(I,NHOOD)</code> performs entropy 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 neighbors. <code>NHOOD</code>'s size must be odd in each dimension.</p> <p>By default, <code>entropyfilt</code> uses the neighborhood <code>true(9)</code>. <code>entropyfilt</code> determines the center element of the neighborhood by <code>floor((size(NHOOD) + 1)/2)</code>. 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>
Class Support	<p><code>I</code> can be <code>logical</code>, <code>uint8</code>, <code>uint16</code>, or <code>double</code>, and must be real and nonsparse. <code>NHOOD</code> can be <code>logical</code> or numeric and must contain zeros or ones. The output array <code>J</code> is of class <code>double</code>.</p> <p><code>entropyfilt</code> converts any class other than <code>logical</code> to <code>uint8</code> for the histogram count calculation so that the pixel values are discrete and directly correspond to a bin value.</p>
Examples	<pre>I = imread('circuit.tif'); J = entropyfilt(I);</pre>

entropyfilt

```
imshow(I), figure, imshow(J,[]);
```

References

[1] Gonzalez, R.C., R.E. Woods, S.L. Eddins, *Digital Image Processing Using MATLAB*, New Jersey, Prentice Hall, 2003, Chapter 11.

See Also

entropy | imhist | rangefilt | stdfilt

Purpose Convert fan-beam projections to parallel-beam

Syntax

```
P = fan2para(F,D)
P = fan2para(..., param1, val1, param2, val2,...)
[P ,parallel_locations,
  parallel_rotation_angles] = fan2para(...)
```

Description `P = fan2para(F,D)` converts the fan-beam data `F` to the parallel-beam data `P`. `D` is the distance in pixels from the fan-beam vertex to the center of rotation that was used to obtain the projections.

`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.

Parameter	Description
'FanCoverage'	String specifying the range through which the beams are rotated. 'cycle' — Rotate through the full range [0,360). This is the default. 'minimal' — 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. Default value is 1.

Parameter	Description
'FanSensorGeometry'	<p>String specifying how sensors are positioned.</p> <p>'arc' — Sensors are spaced equally along a circular arc at distance D from the center of rotation. Default value is 'arc'</p> <p>'line' — Sensors are spaced equally along a line, the closest point of which is distance D from the center of rotation.</p> <p>See fanbeam for details.</p>
'FanSensorSpacing'	<p>Positive real scalar specifying the spacing of the fan-beam sensors. Interpretation of the value depends on the setting of 'FanSensorGeometry'.</p> <p>If 'FanSensorGeometry' is set to 'arc' (the default), the value defines the angular spacing in degrees. Default value is 1.</p> <p>If 'FanSensorGeometry' is 'line', the value specifies the linear spacing. Default value is 1. See fanbeam for details.</p> <hr/> <p>Note This linear spacing is measured on the x' axis. The x' axis for each column, col, of F is oriented at fan_rotation_angles(col) degrees counterclockwise from the x-axis. The origin of both axes is the center pixel of the image.</p> <hr/>

Parameter	Description
'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>'cycle' — Parallel data covers 360 degrees</p> <p>{'halfcycle'} — Parallel data covers 180 degrees</p>
'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 PAR_ROT_INC, where PAR_ROT_INC is the value of 'ParallelRotationIncrement'. $180/PAR_ROT_INC$ 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> $[D \cdot \tan(\min(FAN_ANGLES)), \dots, D \cdot \tan(\max(FAN_ANGLES))]$

Parameter	Description
	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 ,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 nonsparse. All other numeric inputs are double. The output P is double.

Examples

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(Psynthetic,[],...  
        'XData',theta,'YData',xp,'InitialMagnification','fit')  
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(Precovered,[],'XData',Pangles,...
```



```
        'YData',Ploc,'InitialMagnification','fit')
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


fanbeam

Purpose Fan-beam transform

Syntax
`F = fanbeam(I,D)`
`F = fanbeam(..., param1, val1, param1, val2,...)`
`[F, fan_sensor_positions, fan_rotation_angles] = fanbeam(...)`

Description `F = fanbeam(I,D)` computes the fan-beam data (sinogram) `F` from the image `I`. A sinogram is a special x-ray procedure that is done with contrast media (x-ray dye) to visualize any abnormal opening (sinus) in the body.

`D` is the distance in pixels from the fan-beam vertex to the center of rotation. The center of rotation is the center pixel of the image, defined as `floor((size(I)+1)/2)`. `D` must be large enough to ensure that the fan-beam vertex is outside of the image at all rotation angles. See “Tips” on page 3-183 for guidelines on specifying `D`. The following figure illustrates `D` in relation to the fan-beam vertex for one fan-beam geometry. See the `FanSensorGeometry` parameter for more information.



Each column of `F` contains the fan-beam sensor samples at one rotation angle. The number of columns in `F` is determined by the fan rotation increment. By default, the fan rotation increment is 1 degree so `F` has 360 columns.


The number of rows in F is determined by the number of sensors. `fanbeam` determines the number of sensors by calculating how many beams are required to cover the entire image for any rotation angle.

For information about how to specify the rotation increment and sensor spacing, see the documentation for the `FanRotationIncrement` and `FanSensorSpacing` parameters, below.


$F = \text{fanbeam}(\dots, \text{param1}, \text{val1}, \text{param1}, \text{val2}, \dots)$ specifies parameters, listed below, that control various aspects of the fan-beam projections. Parameter names can be abbreviated, and case does not matter.

'`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. Valid values are 'arc' or 'line'. In the 'arc' geometry, sensors are spaced equally along a circular arc, as shown below. This is the default value.



In 'line' geometry, sensors are spaced equally along a line, as shown below.



'FanSensorSpacing' -- Positive real scalar specifying the spacing of the fan-beam sensors. Interpretation of the value depends on the setting of 'FanSensorGeometry'. If 'FanSensorGeometry' is set to 'arc' (the default), the value defines the angular spacing in degrees. Default value is 1. If 'FanSensorGeometry' is 'line', the value specifies the linear spacing. Default value is 1.

Note This linear spacing is measured on the x' axis. The x' axis for each column, `col`, of `F` is oriented at `fan_rotation_angles(col)` degrees counterclockwise from the x -axis. The origin of both axes is the center pixel of the image.

`[F, fan_sensor_positions, fan_rotation_angles] = fanbeam(...)` returns the location of fan-beam sensors in `fan_sensor_positions` and the rotation angles where the fan-beam projections are calculated in `fan_rotation_angles`.

If 'FanSensorGeometry' is 'arc' (the default), `fan_sensor_positions` contains the fan-beam spread angles. If 'FanSensorGeometry' is 'line', `fan_sensor_positions` contains the fan-beam sensor positions along the x' axis. See 'FanSensorSpacing' for more information.

Class Support

I can be logical or numeric. All other numeric inputs and outputs can be double. None of the inputs can be sparse.

Tips

As a guideline, try making D a few pixels larger than half the image diagonal dimension, calculated as follows

```
sqrt(size(I,1)^2 + size(I,2)^2)
```

The values returned in F are a numerical approximation of the fan-beam projections. The algorithm depends on the Radon transform, interpolated to the fan-beam geometry. The results vary depending on the parameters used. You can expect more accurate results when the image is larger, D is larger, and for points closer to the middle of the image, away from the edges.

Examples

The following example computes the fan-beam projections for rotation angles that cover the entire image.

```
iptsetpref('ImshowAxesVisible','on')
ph = phantom(128);
imshow(ph)
[F,Fpos,Fangles] = fanbeam(ph,250);
figure
imshow(F,[],'XData',Fangles,'YData',Fpos,...
       'InitialMagnification','fit')
axis normal
xlabel('Rotation Angles (degrees)')
ylabel('Sensor Positions (degrees)')
colormap(hot), colorbar
```

The following example computes the Radon and fan-beam projections and compares the results at a particular rotation angle.

```
I = ones(100);
D = 200;
dtheta = 45;
```

```
% Compute fan-beam projections for 'arc' geometry
[Farc,FposArcDeg,Fangles] = fanbeam(I,D,...
    'FanSensorGeometry','arc',...
    'FanRotationIncrement',dtheta);
% Convert angular positions to linear distance
% along x-prime axis
FposArc = D*tan(FposArcDeg*pi/180);

% Compute fan-beam projections for 'line' geometry
[Fline,FposLine] = fanbeam(I,D,...
    'FanSensorGeometry','line',...
    'FanRotationIncrement',dtheta);

% Compute the corresponding Radon transform
[R,Rpos]=radon(I,Fangles);

% Display the three projections at one particular rotation
% angle. Note the three are very similar. Differences are
% due to the geometry of the sampling, and the numerical
% approximations used in the calculations.
figure
idx = find(Fangles==45);
plot(Rpos,R(:,idx),...
    FposArc,Farc(:,idx),...
    FposLine,Fline(:,idx))
legend('Radon','Arc','Line')
```

References

[1] Kak, A.C., & Slaney, M., *Principles of Computerized Tomographic Imaging*, IEEE Press, NY, 1988, pp. 92-93.


See Also

fan2para | ifanbeam | iradon | para2fan | phantom | radon

Purpose	Find output bounds for spatial transformation
Syntax	<code>outbounds = findbounds(TFORM,inbounds)</code>
Description	<p><code>outbounds = findbounds(TFORM,inbounds)</code> estimates the output bounds corresponding to a given spatial transformation and a set of input bounds. <code>TFORM</code>, as returned by <code>maketform</code>, is a spatial transformation structure. <code>inbounds</code> is a 2-by-<code>num_dims</code> matrix. The first row of <code>inbounds</code> specifies the lower bounds for each dimension, and the second row specifies the upper bounds. <code>num_dims</code> has to be consistent with the <code>ndims_in</code> field of <code>TFORM</code>.</p> <p><code>outbounds</code> has the same form as <code>inbounds</code>. It is an estimate of the smallest rectangular region completely containing the transformed rectangle represented by the input bounds. Since <code>outbounds</code> is only an estimate, it might not completely contain the transformed input rectangle.</p>
Tips	<code>findbounds</code> gets called by <code>imtransform</code> if 'XData' and 'YData', the parameters that control the output-space bounding box in <code>imtransform</code> , are not specified. <code>findbounds</code> computes the output-space bounding box.
Algorithms	<ol style="list-style-type: none"><code>findbounds</code> first creates a grid of input-space points. These points are located at the center, corners, and middle of each edge in the image. <pre>I = imread('rice.png'); h = imshow(I); set(h,'AlphaData',0.3); axis on, grid on in_points = [... 0.5000 0.5000 0.5000 256.5000 256.5000 0.5000 256.5000 256.5000 0.5000 128.5000 128.5000 0.5000</pre>

findbounds

```
128.5000 128.5000
128.5000 256.5000
256.5000 128.5000];
hold on
plot(in_points(:,1),in_points(:,2),'.','MarkerSize',18)
hold off
```



Grid of Input-Space Points

- Next, `findbounds` transforms the grid of input-space points to output space. If `tform` contains a forward transformation (a nonempty `forward_fcn` field), then `findbounds` transforms the input-space points using `tformfwd`. For example:

```
tform = maketform('affine', ...
    [1.1067 -0.2341 0; 0.5872 1.1769 0; 1000 -300 1]);
out_points = tformfwd(tform, in_points)
```

The output appears below:

```
out_points =
```

```
1.0e+003 *
```


1.0008	-0.2995
1.1512	0.0018
1.2842	-0.3595
1.4345	-0.0582
1.0760	-0.1489
1.1425	-0.3295
1.2177	-0.1789
1.2928	-0.0282
1.3593	-0.2088

If TFORM does not contain a forward transformation, then `findbounds` estimates the output bounds using the Nelder-Mead optimization function `fminsearch`.

- 3 Finally, `findbounds` computes the bounding box of the transformed grid of points.

See Also

`cp2tform` | `imtransform` | `maketform` | `tformarray` | `tformfwd` | `tforminv`

fliptform

Purpose Flip input and output roles of 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.

Examples

```
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

2-D 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.

`freqz2(...)` produces a mesh plot of the two-dimensional magnitude frequency response when no output arguments are specified.

freqz2


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`.

Examples

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`

Purpose	2-D FIR filter using frequency sampling
Syntax	<pre>h = fsamp2(Hd) h = fsamp2(f1, f2, Hd, [m n])</pre>
Description	<p><code>h = fsamp2(Hd)</code> designs a two-dimensional FIR filter with frequency response <code>Hd</code>, and returns the filter coefficients in matrix <code>h</code>. (<code>fsamp2</code> returns <code>h</code> as a computational molecule, which is the appropriate form to use with <code>filter2</code>.) The filter <code>h</code> has a frequency response that passes through points in <code>Hd</code>. If <code>Hd</code> is <code>m</code>-by-<code>n</code>, then <code>h</code> is also <code>m</code>-by-<code>n</code>.</p> <p><code>fsamp2</code> designs two-dimensional FIR filters based on a desired two-dimensional frequency response sampled at points on the Cartesian plane. <code>Hd</code> 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.</p> $H_d(f_1, f_2) = H_d(\omega_1, \omega_2) \Big _{\omega_1 = \pi f_1, \omega_2 = \pi f_2}$ <p>For accurate results, use frequency points returned by <code>freqspace</code> to create <code>Hd</code>.</p> <p><code>h = fsamp2(f1, f2, Hd, [m n])</code> produces an <code>m</code>-by-<code>n</code> FIR filter by matching the filter response at the points in the vectors <code>f1</code> and <code>f2</code>. The frequency vectors <code>f1</code> and <code>f2</code> 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 <code>m*n</code> desired frequency points. <code>fsamp2</code> issues a warning if you specify fewer than <code>m*n</code> points.</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>fsamp2</code> must be of class <code>double</code> . All outputs are of class <code>double</code> .
Examples	Use <code>fsamp2</code> 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 H_d that contains the desired bandpass response. Use `freqspace` to create the frequency range vectors f_1 and f_2 .

```
[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)
```



Algorithms

`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.

References

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 213-217.

See Also

`conv2` | `filter2` | `freqspace` | `ftrans2` | `fwind1` | `fwind2`

fspecial

Purpose Create predefined 2-D filter

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
average	Averaging filter
disk	Circular averaging filter (pillbox)
gaussian	Gaussian lowpass filter
laplacian	Approximates the two-dimensional Laplacian operator
log	Laplacian of Gaussian filter
motion	Approximates the linear motion of a camera
prewitt	Prewitt horizontal edge-emphasizing filter
sobel	Sobel horizontal edge-emphasizing filter
unsharp	Unsharp contrast enhancement filter

`h = fspecial(type, parameters)` accepts the filter specified by `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 $2 \times \text{radius} + 1$. 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.

To compute the filter coefficients, `h`, for 'motion':

- 1 Construct an ideal line segment with the desired length and angle, centered at the center coefficient of `h`.
- 2 For each coefficient location (i, j) , compute the nearest distance between that location and the ideal line segment.

```
3 h = max(1 - nearest_distance, 0);
```

```
4 Normalize h:h = h/(sum(h(:)))
```

- `h = fspecial('prewitt')` returns the 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. `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.

Note Do not be confused by the name of this filter: an unsharp filter is an operator used to sharpen images. The name comes from a publishing industry process in which an image is sharpened by subtracting a blurred (unsharp) version of the image from itself.

fspecial supports the generation of efficient, production-quality C/C++ code from MATLAB. For best results, all inputs must be constants at compilation time. Expressions or variables are allowed if their values do not change. To see a complete list of toolbox functions that support code generation, see “Supported Functions”.

Class Support

h is of class double.

Examples

```
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

$$h_g(n_1, n_2) = e^{-\frac{(n_1^2 + n_2^2)}{2\sigma^2}}$$

$$h(n_1, n_2) = \frac{h_g(n_1, n_2)}{\sum_{n_1} \sum_{n_2} h_g}$$

`fspecial` creates Laplacian filters using

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$$

$$\nabla^2 = \frac{4}{(\alpha + 1)} \begin{bmatrix} \frac{\alpha}{4} & \frac{1-\alpha}{4} & \frac{\alpha}{4} \\ \frac{1-\alpha}{4} & -1 & \frac{1-\alpha}{4} \\ \frac{\alpha}{4} & \frac{1-\alpha}{4} & \frac{\alpha}{4} \end{bmatrix}$$

`fspecial` creates Laplacian of Gaussian (LoG) filters using

$$h_g(n_1, n_2) = e^{-\frac{(n_1^2 + n_2^2)}{2\sigma^2}}$$

$$h(n_1, n_2) = \frac{(n_1^2 + n_2^2 - 2\sigma^2)h_g(n_1, n_2)}{2\pi\sigma^6 \sum_{n_1} \sum_{n_2} h_g}$$

`fspecial` creates averaging filters using

`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` | `del2`

ftrans2

Purpose 2-D 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, Type I (even symmetric, odd-length) filter such as can be returned by `fir1`, `fir2`, or `remez` in the Signal Processing Toolbox software. 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;
```

All inputs and outputs should be of class `double`.

Tips 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 $H(\omega_1, \omega_2)$:


$$h(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2.$$

Examples

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)))
```



ftrans2

- 2 Use `ftrans2` with the default McClellan transformation to create the desired approximately circularly symmetric filter.

```
h = ftrans2(b);  
freqz2(h)
```



References

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 218-237.

See Also

`conv2` | `filter2` | `fsamp2` | `fwind1` | `fwind2`

Purpose 2-D FIR filter using 1-D window method

Syntax

```
h = fwind1(Hd, win)
h = fwind1(Hd, win1, win2)
h = fwind1(f1, f2, Hd, ...)
```

Description `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.

`h = fwind1(Hd, win)` designs a two-dimensional FIR filter `h` with frequency response `Hd`. (`fwind1` returns `h` as a computational molecule, which is the appropriate form to use with `filter2`.) `fwind1` uses the one-dimensional window `win` to form an approximately circularly symmetric two-dimensional window using Huang's method. You can specify `win` using windows from the Signal Processing Toolbox software, such as `boxcar`, `hamming`, `hanning`, `bartlett`, `blackman`, `kaiser`, or `chebwin`. If `length(win)` is `n`, then `h` is `n`-by-`n`.

`Hd` 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 `freqspace` to create `Hd`. (See the entry for `freqspace` for more information.)

`h = fwind1(Hd, win1, win2)` uses the two one-dimensional windows `win1` and `win2` to create a separable two-dimensional window. If `length(win1)` is `n` and `length(win2)` is `m`, then `h` is `m`-by-`n`.

`h = fwind1(f1, f2, Hd, ...)` lets you specify the desired frequency response `Hd` at arbitrary frequencies (`f1` and `f2`) along the x - and y -axes. The frequency vectors `f1` and `f2` 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.

Class Support


The input matrix `Hd` can be of class `double` or of any integer class. All other inputs to `fwind1` must be of class `double`. All outputs are of class `double`.

Examples

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)
```



Algorithms

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

fwind1 calls fwind2 with Hd and the two-dimensional window. fwind2 computes h using an inverse Fourier transform and multiplication by the two-dimensional window:

$$h_d(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H_d(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2$$

$$h(n_1, n_2) = h_d(n_1, n_2)w(n_2, n_2).$$

References


[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990.

See Also

conv2 | filter2 | fsamp2 | freqspace | ftrans2 | fwind2


Purpose	2-D FIR filter using 2-D 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> .
Examples	<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>.

```
[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 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)
```



Algorithms

fwind2 computes h using an inverse Fourier transform and multiplication by the two-dimensional window win .

$$h_d(n_1, n_2) = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} H_d(\omega_1, \omega_2) e^{j\omega_1 n_1} e^{j\omega_2 n_2} d\omega_1 d\omega_2$$

$$h(n_1, n_2) = h_d(n_1, n_2)w(n_1, n_2)$$

References

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 202-213.

See Also

conv2 | filter2 | fsamp2 | freqspace | ftrans2 | fwind1

Purpose	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 <code>double</code> .
Examples	<pre>se = strel(ones(3,3),magic(3)); getheight(se)</pre>
See Also	<code>strel</code> <code>getnhood</code>

getimage

Purpose 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; <code>A</code> 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 <code>uint8</code> arrays, [0,65535] for <code>uint16</code> 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.

Examples

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;
```

getimage

```
imtool cameraman.tif  
I = getimage(imgca);
```

See Also

[imshow](#) | [imtool](#)

Purpose 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.

Examples

```
h = imshow('bag.png');  
imgmodel = getimagemodel(h);
```

See Also `imagemodel`

getline

Purpose Select polyline with mouse

Syntax

```
[x, y] = getline(fig)
[x, y] = getline(ax)
[x, y] = getline
[x, y] = getline(..., 'closed')
```

Description `[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.

`[x, y] = getline(ax)` lets you select a polyline in the axes specified by the handle `ax`.

`[x, y] = getline` is the same as `[x,y] = getline(gcf)`.

`[x, y] = getline(..., 'closed')` animates and returns a closed polygon.

See Also `getpts` | `getrect`

Purpose 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.

Examples

```
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`

getnhood

Purpose	Structuring element neighborhood
Syntax	<code>NHOOD = getnhood(SE)</code>
Description	<code>NHOOD = getnhood(SE)</code> returns the neighborhood associated with the structuring element <code>SE</code> .
Class Support	<code>SE</code> is a STREL object. <code>NHOOD</code> is a logical array.
Examples	<pre>se = strel(eye(5)); NHOOD = getnhood(se)</pre>
See Also	<code>strel</code> <code>getneighbors</code>

Purpose Specify 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`

getrangefromclass

Purpose Default display range of image based on its class

Syntax `range = getrangefromclass(I)`

Description `range = getrangefromclass(I)` returns the default display range of the image `I`, based on its class type. The function returns `range`, a two-element vector specifying the display range in the form `[min max]`.

Class Support `I` can be `uint8`, `uint16`, `int16`, `logical`, `single`, or `double`. `range` is of class `double`.

Note For `single` and `double` data, `getrangefromclass` returns the range `[0 1]` to be consistent with the way `double` and `single` images are interpreted in MATLAB. For integer data, `getrangefromclass` returns the default display range of the class. For example, if the class is `uint8`, the dynamic range is `[0 255]`.

Examples Read in the 16-bit DICOM image and get the default display range.

```
CT = dicomread('CT-MONO2-16-ankle.dcm');  
r = getrangefromclass(CT)  
r =  
  
        -32768         32767
```

See Also `intmin` | `intmax`

Purpose	Specify rectangle with mouse
Syntax	<pre>rect = getrect rect = getrect(fig) rect = getrect(ax)</pre>
Description	<p><code>rect = getrect</code> lets you select a rectangle in the current axes using the mouse. Use the mouse to click and drag the desired rectangle. <code>rect</code> is a four-element vector with the form <code>[xmin ymin width height]</code>. To constrain the rectangle to be a square, use a shift- or right-click to begin the drag.</p> <p><code>rect = getrect(fig)</code> lets you select a rectangle in the current axes of figure <code>fig</code> using the mouse.</p> <p><code>rect = getrect(ax)</code> lets you select a rectangle in the axes specified by the handle <code>ax</code>.</p>
Examples	<p>Select a rectangle in an image of the moon:</p> <pre>imshow('moon.tif') rect = getrect</pre>
See Also	<code>getline</code> <code>getpts</code>

getsequence

Purpose Sequence of decomposed structuring elements

Syntax SEQ = getsequence(SE)

Description SEQ = getsequence(SE) returns the array of structuring elements SEQ, containing the individual structuring elements that form the decomposition of SE. SE can be an array of structuring elements. SEQ is equivalent to SE, but the elements of SEQ have no decomposition.

Class Support SE and SEQ are arrays of STREL objects.

Examples 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)
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 = getsequence(se)
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 grayscale or binary image to indexed image

Syntax `[X, map] = gray2ind(I,n)`
`[X, map] = gray2ind(BW,n)`

Description `[X, map] = gray2ind(I,n)` converts the grayscale image `I` to an indexed image `X`. `n` specifies the size of the colormap, `gray(n)`. `n` must be an integer between 1 and 65536. If `n` is omitted, it defaults to 64.

`[X, map] = gray2ind(BW,n)` converts the binary image `BW` to an indexed image `X`. `n` specifies the size of the colormap, `gray(n)`. If `n` is omitted, it defaults to 2.

`gray2ind` scales and then rounds the intensity image to produce an equivalent indexed image.

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`.

Examples Convert a grayscale 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 Create gray-level co-occurrence matrix from image

Syntax

```
glcm = graycomatrix(I)
glcms = graycomatrix(I, param1, val1, param2, val2,...)
[glcm, SI] = graycomatrix(...)
```


Description

`glcm = 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 .

`graycomatrix` calculates the GLCM from a scaled version of the image. By default, if `I` is a binary image, `graycomatrix` scales the image to two gray-levels. If `I` is an intensity image, `graycomatrix` scales the image to eight gray-levels. You can specify the number of gray-levels `graycomatrix` uses to scale the image by 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 several values in the GLCM of the 4-by-5 image `I`. Element (1,1) in the GLCM contains the value 1 because there is only one instance in the image where two, horizontally adjacent pixels have the values 1 and 1. Element (1,2) in the GLCM contains the value 2 because there are two instances in the image where two, horizontally adjacent pixels have the values 1 and 2. `graycomatrix` continues this processing to fill in all the values in the GLCM.

graycomatrix



`glcms = graycomatrix(I, param1, val1, param2, val2,...)`
 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'	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]
'NumLevels'	Integer specifying the number of gray-levels to use when scaling the grayscale values in I. For example, if NumLevels is 8, graycomatrix scales	8 (numeric) 2 (binary)

Parameter	Description	Default										
	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).											
'Offset'	<p>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]											
'Symmetric'	Boolean that creates a GLCM where the ordering of values in the pixel pairs is not considered. For example, when 'Symmetric' is set to true,	false										

graycomatrix

Parameter	Description	Default
	graycomatrix counts both 1,2 and 2,1 pairings when calculating the number of times the value 1 is adjacent to the value 2. When 'Symmetric' is set to false, graycomatrix only counts 1,2 or 2,1, depending on the value of 'offset'. See "Notes" on page 3-228.	

[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. Also, the word co-occurrence is frequently used in the literature without a hyphen, cooccurrence.

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.

The GLCM created when 'Symmetric' is set to true is symmetric across its diagonal, and is equivalent to the GLCM described by Haralick (1973). The GLCM produced by the following syntax, with 'Symmetric' set to true

```
graycomatrix(I, 'offset', [0 1], 'Symmetric', true)
```

is equivalent to the sum of the two GLCMs produced by the following statements where 'Symmetric' is set to false.

```
graycomatrix(I, 'offset', [0 1], 'Symmetric', false)
graycomatrix(I, 'offset', [0 -1], 'Symmetric', false)
```

Examples

Calculate the gray-level co-occurrence matrix for a grayscale image.

```
I = imread('circuit.tif');
glcm = graycomatrix(I,'Offset',[2 0]);
```

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',[])
```

References

Haralick, R.M., K. Shanmugan, and I. Dinstein, "Textural Features for Image Classification", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, 1973, pp. 610-621.

Haralick, R.M., and L.G. Shapiro. Computer and Robot Vision: Vol. 1, Addison-Wesley, 1992, p. 459.

See Also

graycoprops

Purpose Properties of 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.	$\sum_{i,j} \frac{(i-\mu_i)(j-\mu_j)p(i,j)}{\sigma_i\sigma_j}$

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.

Examples

```
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'})
```

graycoprops

See Also

`graycomatrix`

Purpose	Gray-weighted distance transform of grayscale image
Syntax	<code>T = graydist(A,mask)</code> <code>T = graydist(A,C,R)</code> <code>T = graydist(A,ind)</code> <code>T = graydist(...,method)</code>
Description	<p><code>T = graydist(A,mask)</code> computes the gray-weighted distance transform of the grayscale image <code>A</code>. Locations where <code>mask</code> is true are seed locations.</p> <p><code>T = graydist(A,C,R)</code> uses vectors <code>C</code> and <code>R</code> to specify the row and column coordinates of seed locations.</p> <p><code>T = graydist(A,ind)</code> specifies the linear indices of seed locations using the vector <code>ind</code>.</p> <p><code>T = graydist(...,method)</code> specifies an alternate distance metric. <code>method</code> determines the chamfer weights that are assigned to the local neighborhood during outward propagation. Each pixel's contribution to the geodesic time is based on the chamfer weight in a particular direction multiplied by the pixel intensity.</p>
Input Arguments	<p>A Grayscale image.</p> <p>mask Logical image the same size as <code>A</code> that specifies seed locations.</p> <p>C,R Numeric vectors that contain the positive integer row and column coordinates of the seed locations. Coordinate values are valid <code>C,R</code> subscripts in <code>A</code>.</p> <p>ind Numeric vector of positive integer, linear indices of seed locations.</p>

method

Type of distance metric. `method` can have any of these values.

Method	Description
'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 $.
'chessboard'	The chessboard distance is $\max(x_1 - x_2 , y_1 - y_2)$.
'quasi-euclidean'	The quasi-Euclidean distance 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 $, otherwise.

Default: 'chessboard'

Output Arguments

T

Array the same size as A that specifies the gray-weighted distance transform. If the input numeric type of A is `double`, the output numeric type of T is `double`. If the input is any other numeric type, the output T is `single`.

Class Support

A can be numeric or logical, and it must be nonsparse. `mask` is a logical array of the same size as A. `C`, `R`, and `ind` are numeric vectors that contain positive integer values.

The output T is an array of the same size as A. If the input numeric type of A is `double`, the output T is `double`. If the input is any other numeric type, the output T is `single`.

Examples

Matrices generated by the `magic` function have equal row, column and diagonal sums. The minimum path between the upper left and lower

right corner is along the diagonal. The following example demonstrates how the `graydist` function computes this path:

```
A = magic(3)
T1 = graydist(A,1,1);
T2 = graydist(A,3,3);
T = T1 + T2
```

A =

8	1	6
3	5	7
4	9	2

T =

10	11	17
13	10	13
17	17	10

As expected, there is a constant-value minimum path along the diagonal.

Algorithms

`graydist` uses the geodesic time algorithm described in Soille, P., *Generalized geodesy via geodesic time*, Pattern Recognition Letters, vol.15, December 1994; pp. 1235–1240

The basic equation for geodesic time along a path is:

$$\tau_f(P) = \frac{f(p_o)}{2} + \frac{f(p_l)}{2} + \sum_{i=1}^{l-1} f(p_i)$$

graydist

See Also

`bwdist` | `bwdistgeodesic` | `watershed`

Purpose Convert grayscale image to indexed image using multilevel thresholding

Syntax `X = grayscale(I, n)`

Description `X = grayscale(I, n)` thresholds the intensity image `I` returning an indexed image in `X`. `grayscale` uses the threshold values:

$$\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}$$

`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`.

Examples

```
I = imread('snowflakes.png');
X = grayscale(I,16);
imshow(I)
figure, imshow(X, jet(16))
```

grayscale



See Also

[gray2ind](#)

Purpose	Global image threshold using Otsu's method
Syntax	<pre>level = graythresh(I) [level EM] = graythresh(I)</pre>
Description	<p><code>level = graythresh(I)</code> computes a global threshold (<code>level</code>) that can be used to convert an intensity image to a binary image with <code>im2bw</code>. <code>level</code> is a normalized intensity value that lies in the range <code>[0, 1]</code>.</p> <p>The <code>graythresh</code> function uses Otsu's method, which chooses the threshold to minimize the intraclass variance of the black and white pixels.</p> <p>Multidimensional arrays are converted automatically to 2-D arrays using <code>reshape</code>. The <code>graythresh</code> function ignores any nonzero imaginary part of <code>I</code>.</p> <p><code>[level EM] = graythresh(I)</code> returns the effectiveness metric, <code>EM</code>, as the second output argument. The effectiveness metric is a value in the range <code>[0 1]</code> 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.</p>
Class Support	The input image <code>I</code> can be of class <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , or <code>double</code> and it must be nonsparse. The return value <code>level</code> is a double scalar. The effectiveness metric <code>EM</code> is a double scalar.
Examples	<pre>I = imread('coins.png'); level = graythresh(I); BW = im2bw(I,level); imshow(BW)</pre>
References	[1] Otsu, N., "A Threshold Selection Method from Gray-Level Histograms," <i>IEEE Transactions on Systems, Man, and Cybernetics</i> , Vol. 9, No. 1, 1979, pp. 62-66.
See Also	<code>im2bw</code>

hdrread

Purpose	Read high dynamic range (HDR) image
Syntax	<code>hdr = hdrread(filename)</code>
Description	<code>hdr = hdrread(filename)</code> reads the high dynamic range (HDR) image from the file specified by <code>filename</code> . <code>hdr</code> is an m-by-n-by-3 RGB array in the range <code>[0,inf)</code> of type <code>single</code> . For scene-referred data sets, these values usually are scene illumination in radiance units. To display these images, use an appropriate tone-mapping operator.
Class Support	The output image <code>hdr</code> is an m-by-n-by-3 image of type <code>single</code> .
Examples	<pre>hdr = hdrread('office.hdr'); rgb = tonemap(hdr); imshow(rgb);</pre>
References	[1] Larson, Greg W. "Radiance File Formats" http://radsite.lbl.gov/radiance/refer/filefmts.pdf
See Also	<code>hdrwrite</code> <code>makehdr</code> <code>tonemap</code>

- Purpose** Write Radiance high dynamic range (HDR) image file
- Syntax** `hdrwrite(hdr, filename)`
- Description** `hdrwrite(hdr, filename)` creates a Radiance high dynamic range (HDR) image file from HDR, a single- or double-precision high dynamic range RGB image. The HDR file with the name `filename` uses run-length encoding to minimize file size.
- See Also** `hdrread` | `makehdr` | `tonemap`

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 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)`.

`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 syntax 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 syntax 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`.

Examples

Enhance the contrast of an intensity image using histogram equalization.


```
I = imread('tire.tif');  
J = histeq(I);  
imshow(I)  
figure, imshow(J)
```



Display a histogram of the original image.


```
figure; imhist(I,64)
```

histeq



Compare it to a histogram of the processed image.

```
figure; imhist(J,64)
```



Algorithms

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 the transformation $b = T(a)$ to map the gray levels in X (or the `colormap`) to their new values.

If you do not specify `hgram`, `histeq` creates a flat `hgram`,

```
hgram = ones(1,n)*prod(size(A))/n;
```

and then applies the previous algorithm.

histeq

See Also

brighten | imadjust | imhist

Purpose

Hough transform

Syntax

```
[H, theta, rho] = hough(BW)
[H, theta, rho] = hough(BW, ParameterName, ParameterValue)
```

Description

[H, theta, rho] = hough(BW) computes the Standard Hough Transform (SHT) of the binary image BW. 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 hough generates the Hough transform matrix. BW can be logical or numeric, and it must be real, 2-D, and nonsparse.

[H, theta, rho] = hough(BW, ParameterName, ParameterValue) computes the SHT using parameter name/value pairs. When ParameterName is 'RhoResolution', specify a real scalar value between 0 and norm(size(BW)), exclusive, to determine the spacing of the Hough transform bins along the *rho* axis. The default value is 1.

When ParameterName is 'Theta', specify a vector of Hough transform *theta* values. Each element of the vector determines the *theta* value for the corresponding column of the output matrix H. The acceptable range of *theta* values is $-90^\circ \leq \theta < 90^\circ$, and the default is -90:89.

Examples

Compute and display the Hough transform of a gantrycrane image.

```
RGB = imread('gantrycrane.png');

% Convert to intensity.
I = rgb2gray(RGB);


% Extract edges.
BW = edge(I, 'canny');
[H,T,R] = hough(BW, 'RhoResolution', 0.5, 'Theta', -90:0.5:89.5);

% Display the original image.
subplot(2,1,1);
imshow(RGB);
```

hough

```
title('Gantrycrane Image');

% Display the Hough matrix.
subplot(2,1,2);
imshow(imadjust(mat2gray(H)), 'XData',T, 'YData',R,...
       'InitialMagnification','fit');
title('Hough Transform of Gantrycrane Image');
xlabel('\theta'), ylabel('\rho');
axis on, axis normal, hold on;
colormap(hot);
```




Compute the Hough transform over a limited θ range for the gantrycrane image.

```
RGB = imread('gantrycrane.png');
I = rgb2gray(RGB);
BW = edge(I, 'canny');
[H,T,R] = hough(BW, 'Theta', 44:0.5:46);
figure
```

```

imshow(imadjust(mat2gray(H)), 'XData', T, 'YData', R, ...
       'InitialMagnification', 'fit');
title('Limited Theta Range Hough Transform of Gantrycrane Image');
xlabel('\theta'), ylabel('\rho');
axis on, axis normal;
colormap(hot)

```




Algorithms

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 ρ is the distance from the origin to the line along a vector perpendicular to the line. θ is the angle of the perpendicular projection from the origin to the line measured in degrees clockwise from the positive x -axis. The range of θ is $-90^\circ \leq \theta < 90^\circ$. The angle of the line itself is $\theta + 90^\circ$, also measured clockwise with respect to the positive x -axis.

hough



The SHT is a parameter space matrix whose rows and columns correspond to ρ and θ values respectively. The elements in the SHT represent accumulator cells. Initially, the value in each cell is zero. Then, for every non-background point in the image, ρ is calculated for every θ . ρ is rounded off to the nearest allowed row in SHT. That accumulator cell is incremented. At the end of this procedure, a value of Q in $SHT(r,c)$ means that Q points in the xy -plane lie on the line specified by $\theta(c)$ and $\rho(r)$. Peak values in the SHT represent potential lines in the input image.

The Hough transform matrix, H , is $n\rho$ -by- $n\theta$.

Values of $n\rho$ and $n\theta$

```
nrho = 2*(ceil(D/RhoResolution)) + 1, where  
D = sqrt((numRowsInBW - 1)^2 + (numColsInBW - 1)^2).  
rho values range from -diagonal to diagonal, where  
diagonal = RhoResolution*ceil(D/RhoResolution).
```

```
ntheta = length(theta)
```

See Also

[houghlines](#) | [houghpeaks](#)

How To

- “Detecting Lines Using the Hough Transform”

houghlines

Purpose Extract line segments based on Hough transform

Syntax
`lines = houghlines(BW, theta, rho, peaks)`
`lines = houghlines(..., param1, val1, param2, val2)`

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, val1, param2, val2)` 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.

Examples

Search for line segments in an image and highlight 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];
    plot(xy(:,1),xy(:,2),'LineWidth',2,'Color','green');
```

houghlines

```
% 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','blue');
```

See Also

[hough](#) | [houghpeaks](#)

Purpose Identify peaks in Hough transform

Syntax
`peaks = houghpeaks(H, numpeaks)`
`peaks = houghpeaks(..., param1, val1, param2, val2)`

Description `peaks = houghpeaks(H, numpeaks)` locates peaks in the Hough transform matrix, H, generated by the `hough` function. `numpeaks` is a scalar value that specifies the maximum number of peaks to identify. If you omit `numpeaks`, it defaults to 1.

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.

`peaks = houghpeaks(..., 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
'Threshold'	Nonnegative scalar value that specifies the threshold at which values of H are considered to be peaks. Threshold can vary from 0 to Inf. Default is $0.5 * \max(H(:))$.
'NHoodSize'	Two-element vector of positive odd integers: [M N]. '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 $\text{size}(H) / 50$.

Class Support H is the output of the `hough` function. `numpeaks` is a positive integer scalar.

Examples Locate and display two peaks in the Hough transform of a rotated image.

```
I = imread('circuit.tif');
BW = edge(imrotate(I,50,'crop'),'canny');
```

houghpeaks

```
[H,T,R] = hough(BW);  
P = houghpeaks(H,2);  
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	Search for ICC profiles
Syntax	<pre>P = iccfind(directory) [P, descriptions] = iccfind(directory) [...] = iccfind(directory, pattern)</pre>
Description	<p><code>P = iccfind(directory)</code> searches for all of the ICC profiles in the directory specified by <code>directory</code>. The function returns <code>P</code>, a cell array of structures containing profile information.</p> <p><code>[P, descriptions] = iccfind(directory)</code> searches for all of the ICC profiles in the specified directory and returns <code>P</code>, a cell array of structures containing profile information, and <code>descriptions</code>, a cell array of text strings, where each string describes the corresponding profile in <code>P</code>. Each text string is the value of the <code>Description.String</code> field in the profile information structure.</p> <p><code>[...] = iccfind(directory, pattern)</code> returns all of the ICC profiles in the specified directory with the given <code>pattern</code> in their <code>Description.String</code> fields. <code>iccfind</code> performs case-insensitive pattern matching.</p>

Note To improve performance, `iccfind` caches copies of the ICC profiles in memory. Adding or modifying profiles might not change the results of `iccfind`. To clear the cache, use the `clear functions` command.

Examples	<p>Get all the ICC profiles in the default system directory where profiles are stored.</p> <pre>profiles = iccfind(iccroot);</pre> <p>Get a listing of all the ICC profiles with text strings that describe each profile.</p> <pre>[profiles, descriptions] = iccfind(iccroot);</pre>
-----------------	--

iccfind

Find the profiles whose descriptions contain the text string RGB.

```
[profiles, descriptions] = iccfind(iccroot, 'rgb');
```

See Also

[iccread](#) | [iccroot](#) | [iccwrite](#)

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` looks for the file in the current directory, a directory on the MATLAB path, or in the directory returned by `iccrroot`, in that order.

`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.

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 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
Description	1-by-1 struct array	The <code>String</code> field in this structure contains a text string describing the profile.
MediaWhitepoint	double array	<i>XYZ</i> 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 <code>iccread</code> 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
AToB1	BToA1	Preview1
AToB2	BToA2	Preview2
AToB3	BToA3	Gamut

Examples

The example reads the ICC profile that describes a typical PC computer monitor.

```
P = iccread('sRGB.icm')
```

```
P =
```

```

    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 =
```

iccread

RGB

See Also

`applycform` | `iccfind` | `iccroot` | `iccwrite` | `isicc` | `makecform`

Purpose Find system default ICC profile repository

Syntax `rootdir = iccroot`

Description `rootdir = iccroot` returns the system directory containing ICC profiles. Additional profiles can be stored in other directories, but this is the default location used by the color management system.

Note Only Windows and Mac OS X platforms are supported.

Examples Return information on all the profiles in the root directory.

```
iccfind(iccroot)
```

See Also `iccfind` | `iccread` | `iccwrite`

iccwrite

Purpose Write ICC color profile to disk file

Syntax `P_new = iccwrite(P, filename)`

Description `P_new = iccwrite(P, filename)` writes the International Color Consortium (ICC) color profile data in structure `P` to the file specified by `filename`.

`P` is a structure representing an ICC profile in the data format returned by `iccread` and used by `makecform` and `applycform` to compute color-space transformations. `P` 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. `iccwrite` makes any necessary corrections to the profile structure before writing it to the file and returns this corrected structure in `P_new`.

Note Because some applications use the profile description string in the ICC profile to present choices to users, the ICC recommends modifying the profile description string in the ICC profile data before writing the data to a file. Each profile should have a unique description string. For more information, see the example.

`iccwrite` 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 `Version` field in the file profile header. If any required fields are missing, `iccwrite` errors. For more information about ICC profiles, visit the ICC web site, www.color.org.

Note `iccwrite` 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 `isicc` to validate a profile.

Examples

Read a profile into the MATLAB workspace and export the profile data to a new file. The example changes the profile description string in the profile data before writing the data to a file.

```
P = iccread('monitor.icm');  
  
P.Description.String  
  
ans =  
  
sgC4_050102_d50.pf  
  
P.Description.String = 'my new description';  
  
pmon = iccwrite(P, 'monitor2.icm');
```

See Also

[applycform](#) | [iccread](#) | [isicc](#) | [makecform](#)

idct2

Purpose 2-D 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) pads A with 0's to size m-by-n before transforming. If [m n] < size(A), idct2 crops A before transforming.

B = idct2(A,[m n]) same as above.

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.

Algorithms idct2 computes the two-dimensional inverse DCT using:

$$A_{mn} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \alpha_p \alpha_q B_{pq} \cos \frac{\pi(2m+1)p}{2M} \cos \frac{\pi(2n+1)q}{2N}, \quad 0 \leq m \leq M-1, \quad 0 \leq n \leq N-1,$$

where

$$\alpha_p = \begin{cases} \frac{1}{\sqrt{M}}, & p = 0 \\ \sqrt{\frac{2}{M}}, & 1 \leq p \leq M-1 \end{cases}$$

and

$$\alpha_q = \begin{cases} \frac{1}{\sqrt{N}}, & q = 0 \\ \sqrt{\frac{2}{N}}, & 1 \leq q \leq N-1 \end{cases}.$$

Examples

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])
```

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.

See Also

`dct2` | `dctmtx` | `fft2` | `ifft2`

ifanbeam

Purpose 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. <code>{'cycle'}</code> — Rotate through the full range <code>[0,360)</code> . <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.

Parameter	Description
'FanSensorGeometry'	<p>String specifying how sensors are positioned.</p> <p>'arc' — Sensors are spaced equally along a circular arc at distance D from the center of rotation. Default value is 'arc'</p> <p>'line' — Sensors are spaced equally along a line, the closest point of which is distance D from the center of rotation.</p> <p>See <code>fanbeam</code> for details.</p>
'FanSensorSpacing'	<p>Positive real scalar specifying the spacing of the fan-beam sensors. Interpretation of the value depends on the setting of 'FanSensorGeometry'.</p> <p>If 'FanSensorGeometry' is set to 'arc' (the default), the value defines the angular spacing in degrees. Default value is 1.</p> <p>If 'FanSensorGeometry' is 'line', the value specifies the linear spacing. Default value is 1. See <code>fanbeam</code> for details.</p>
'Filter'	<p>String specifying the name of a filter. See <code>iradon</code> for details.</p>
'FrequencyScaling'	<p>Scalar in the range (0,1] that modifies the filter by rescaling its frequency axis. See <code>iradon</code> for details.</p>

Parameter	Description
'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>
'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> <hr/> <p>Note If the projections were calculated with the fanbeam function, the reconstructed image might not be the same size as the original image.</p> <hr/>

[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.

Examples

Example 1

This example creates a fan-beam transformation of the phantom head image and then calls the ifanbeam function to recreate the phantom image from the fan-beam transformation.

```
ph = phantom(128);
d = 100;
F = fanbeam(ph,d);
I = ifanbeam(F,d);
imshow(ph), figure, imshow(I);
```

Example 2

This example illustrates use of the ifanbeam function with the 'fancoverage' option set to 'minimal' .

```
ph = phantom(128);
P = radon(ph);
[F,obeta,otheta] = para2fan(P,100,...
    'FanSensorSpacing',0.5,...
    'FanCoverage','minimal',...
    'FanRotationIncrement',1);
phReconstructed = ifanbeam(F,100,...
    'FanSensorSpacing',0.5,...
    'Filter','Shepp-Logan',...
    'OutputSize',128,...
    'FanCoverage','minimal',...);
```

ifanbeam

```
                                'FanRotationIncrement',1);  
imshow(ph), figure, imshow(phReconstructed)
```

References

[1] Kak, A. C., and M. Slaney, *Principles of Computerized Tomographic Imaging*, New York, NY, IEEE Press, 1988.

See Also

fan2para | fanbeam | iradon | para2fan | phantom | radon

Purpose	Convert image to binary image, based on threshold
Syntax	<pre>BW = im2bw(I, level) BW = im2bw(X, map, level) BW = im2bw(RGB, level)</pre>
Description	<p><code>BW = im2bw(I, level)</code> converts the grayscale image <code>I</code> to a binary image. The output image <code>BW</code> replaces all pixels in the input image with luminance greater than <code>level</code> with the value 1 (white) and replaces all other pixels with the value 0 (black). Specify <code>level</code> in the range [0,1]. This range is relative to the signal levels possible for the image's class. Therefore, a <code>level</code> value of 0.5 is midway between black and white, regardless of class. To compute the <code>level</code> argument, you can use the function <code>graythresh</code>. If you do not specify <code>level</code>, <code>im2bw</code> uses the value 0.5.</p> <p><code>BW = im2bw(X, map, level)</code> converts the indexed image <code>X</code> with colormap <code>map</code> to a binary image.</p> <p><code>BW = im2bw(RGB, level)</code> converts the truecolor image <code>RGB</code> to a binary image.</p> <p>If the input image is not a grayscale image, <code>im2bw</code> converts the input image to grayscale, and then converts this grayscale image to binary by thresholding.</p>
Class Support	The input image can be of class <code>uint8</code> , <code>uint16</code> , <code>single</code> , <code>int16</code> , or <code>double</code> , and must be nonsparse. The output image <code>BW</code> is of class <code>logical</code> . <code>I</code> and <code>X</code> must be 2-D. RGB images are M-by-N-by-3.
Examples	<pre>load trees BW = im2bw(X,map,0.4); imshow(X,map), figure, imshow(BW)</pre>




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,'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 <i>distinct</i> 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 $A = [A11 A12; A21 A22]$, where each A_{ij} is m -by- n , then $B = [A11(:) A12(:) A21(:) A22(:)]$.
{'sliding'}	Converts each <i>sliding</i> 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 Examples, 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`.

im2col

Class Support

The input image A can be numeric or logical. The output matrix B is of the same class as the input image.

Examples

Calculate the local mean using a [2 2] neighborhood with zero padding:

```
A = reshape(linspace(0,1,16),[4 4])'  
B = im2col(A,[2 2])  
M = mean(B)  
newA = col2im(M,[1 1],[3 3])
```

The output appears like this:

```
newA =  
  
    0.1667    0.2333    0.3000  
    0.4333    0.5000    0.5667  
    0.7000    0.7667    0.8333
```

See Also

[blockproc](#) | [col2im](#) | [colfilt](#) | [nlfilter](#)

Purpose	Convert image to double precision
Syntax	<pre>I2 = im2double(I) RGB2 = im2double(RGB) I = im2double(BW) X2 = im2double(X, 'indexed')</pre>
Description	<p><code>I2 = im2double(I)</code> converts the intensity image <code>I</code> to double precision, rescaling the data if necessary.</p> <p>If the input image is of class <code>double</code>, the output image is identical.</p> <p><code>RGB2 = im2double(RGB)</code> converts the truecolor image <code>RGB</code> to double precision, rescaling the data if necessary.</p> <p><code>I = im2double(BW)</code> converts the binary image <code>BW</code> to a double-precision intensity image.</p> <p><code>X2 = im2double(X, 'indexed')</code> converts the indexed image <code>X</code> to double precision, offsetting the data if necessary.</p>
Class Support	Intensity and truecolor images can be <code>uint8</code> , <code>uint16</code> , <code>double</code> , <code>logical</code> , <code>single</code> , or <code>int16</code> . Indexed images can be <code>uint8</code> , <code>uint16</code> , <code>double</code> or <code>logical</code> . Binary input images must be <code>logical</code> . The output image is <code>double</code> .
Examples	<pre>I1 = reshape(uint8(linspace(1,255,25)), [5 5]) I2 = im2double(I1)</pre>
See Also	<code>double</code> <code>im2single</code> <code>im2int16</code> <code>im2uint8</code> <code>im2uint16</code>

im2int16

Purpose Convert image to 16-bit signed integers

Syntax

```
I2 = im2int16(I)
RGB2 = im2int16(RGB)
I = im2int16(BW)
```

Description

`I2 = im2int16(I)` converts the intensity image `I` to `int16`, rescaling the data if necessary. If the input image is of class `int16`, the output image is identical to it.

`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`.

Examples

```
I = reshape(linspace(0,1,20),[5 4])
I2 = im2int16(I)
```

See Also `im2double` | `im2single` | `im2uint8` | `im2uint16` | `int16`

Purpose	Convert image to Java buffered image
Syntax	<pre>jimage = im2java2d(I) jimage = im2java2d(X,MAP)</pre>
Description	<p><code>jimage = im2java2d(I)</code> converts the image <code>I</code> to an instance of the Java image class <code>java.awt.image.BufferedImage</code>. The image <code>I</code> can be an intensity (grayscale), RGB, or binary image.</p> <p><code>jimage = im2java2d(X,MAP)</code> converts the indexed image <code>X</code> with colormap <code>MAP</code> to an instance of the Java class <code>java.awt.image.BufferedImage</code>.</p> <hr/> <p>Note The <code>im2java2d</code> function works with the Java 2D API. The <code>im2java</code> function works with the Java Abstract Windowing Toolkit (AWT).</p> <hr/>
Class Support	Intensity, indexed, and RGB input images can be of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . Binary input images must be of class <code>logical</code> .
Examples	<p>Read an image into the MATLAB workspace and then use <code>im2java2d</code> to convert it into an instance of the Java class <code>java.awt.image.BufferedImage</code>.</p> <pre>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</pre>

im2single

Purpose Convert image to single precision

Syntax

```
I2 = im2single(I)
RGB2 = im2single(RGB)
I = im2single(BW)
X2 = im2single(X, 'indexed')
```

Description

`I2 = im2single(I)` converts the intensity image `I` to single, rescaling the data if necessary. If the input image is of class `single`, the output image is identical to it.

`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`.

Examples

```
I = reshape(uint8(linspace(1,255,25)),[5 5])
I2 = im2single(I)
```

See Also [im2double](#) | [im2int16](#) | [im2uint8](#) | [im2uint16](#) | [single](#)

Purpose	Convert image to 16-bit unsigned integers
Syntax	<pre>I2 = im2uint16(I) RGB2 = im2uint16(RGB) I = im2uint16(BW) X2 = im2uint16(X,'indexed')</pre>
Description	<p><code>I2 = im2uint16(I)</code> converts the intensity image <code>I</code> to <code>uint16</code>, rescaling the data if necessary. If the input image is of class <code>uint16</code>, the output image is identical to it.</p> <p><code>RGB2 = im2uint16(RGB)</code> converts the truecolor image <code>RGB</code> to <code>uint16</code>, rescaling the data if necessary.</p> <p><code>I = im2uint16(BW)</code> converts the binary image <code>BW</code> to a <code>uint16</code> intensity image, changing 1-valued elements to 65535.</p> <p><code>X2 = im2uint16(X,'indexed')</code> converts the indexed image <code>X</code> to <code>uint16</code>, offsetting the data if necessary. If <code>X</code> is of class <code>double</code>, <code>max(X(:))</code> must be 65536 or less.</p>
Class Support	Intensity and truecolor images can be <code>uint8</code> , <code>uint16</code> , <code>double</code> , <code>logical</code> , <code>single</code> , or <code>int16</code> . Indexed images can be <code>uint8</code> , <code>uint16</code> , <code>double</code> , or <code>logical</code> . Binary input images must be <code>logical</code> . The output image is <code>uint16</code> .
Examples	<pre>I = reshape(linspace(0,1,20),[5 4]) I2 = im2uint16(I)</pre>
See Also	<code>im2uint8</code> <code>double</code> <code>im2double</code> <code>uint8</code> <code>uint16</code> <code>imapprox</code>

im2uint8

Purpose Convert image to 8-bit unsigned integers

Syntax

```
I2 = im2uint8(I1)
RGB2 = im2uint8(RGB1)
I = im2uint8(BW)
X2 = im2uint8(X1,'indexed')
```

Description `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 the input image. If the input image is not `uint8`, `im2uint8` returns the equivalent image of class `uint8`, rescaling or offsetting the data as necessary.

`I2 = im2uint8(I1)` converts the grayscale image `I1` to `uint8`, rescaling the data if necessary.

`RGB2 = im2uint8(RGB1)` converts the truecolor image `RGB1` to `uint8`, rescaling the data if necessary.

`I = im2uint8(BW)` converts the binary image `BW` to a `uint8` grayscale image, changing 1-valued elements to 255.

`X2 = im2uint8(X1,'indexed')` converts the indexed image `X1` to `uint8`, offsetting the data if necessary. Note that it is not always possible to convert an indexed image to `uint8`. If `X1` is of class `double`, the maximum value of `X1` must be 256 or less; if `X1` is of class `uint16`, the maximum value of `X1` must be 255 or less.

Class Support Grayscale 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 `uint8`.

Examples

```
I1 = reshape(uint16(linspace(0,65535,25)),[5 5])
I2 = im2uint8(I1)
```

See Also `im2double` | `im2int16` | `im2single` | `im2uint16` | `uint8`

Purpose 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 This function may take advantage of hardware optimization for data types `uint8`, `int16`, and `single` to run faster. Hardware optimization requires that arrays `X` and `Y` are of the same size and class.

Examples

Calculate 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);
```

imabsdiff

`imshow(K,[]) % [] = scale data automatically`

See Also

`imadd | imcomplement | imdivide | imlincomb | immultiply |
imsubtract | ippl`

Purpose Add two images or add constant to 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.

Note On Intel architecture processors, `imadd` can take advantage of the Intel Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The Intel IPP library is activated if arrays `X`, `Y`, and `Z` are of class `logical`, `uint8`, or `single` and are of the same class. The Intel IPP library 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');
```

imadd

```
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](#) | [ipl](#)

Purpose

Adjust image intensity values or colormap

Syntax

```
J = imadjust(I)
J = imadjust(I,[low_in; high_in],[low_out; high_out])
J = imadjust(I,[low_in; high_in],[low_out; high_out],gamma)
newmap = imadjust(map,[low_in; high_in],[low_out; high_out],
    gamma)
RGB2 = imadjust(RGB1,...)
```

Description

`J = imadjust(I)` maps the intensity values in grayscale 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 `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 for `low_in`, `high_in`, `low_out`, and `high_out` must be between 0 and 1. 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 `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`.

Examples

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');  
RGB2 = imadjust(RGB1,[.2 .3 0; .6 .7 1],[,]);  
imshow(RGB1), figure, imshow(RGB2)
```



See Also

brighten | histeq | stretchlim

ImageAdapter

Purpose Interface for image I/O

Description ImageAdapter, an abstract class, specifies the Image Processing Toolbox interface for region-based reading and writing of image files. You can use classes that inherit from the ImageAdapter interface with the `blockproc` function for file-based processing of arbitrary image file formats.

Construction `adapter = ClassName(...)` handles object initialization, manages file opening or creation, and sets the initial values of class properties. The class constructor can take any number of arguments.

Properties

Colormap

Specifies a colormap. Use the `Colormap` property when working with indexed images.

Data Type: 2-D, real, M-by-3 matrix

Default: Empty (`[]`), indicating either a grayscale, logical, or truecolor image

ImageSize

Holds the size of the entire image. When you construct a new class that inherits from `ImageAdapter`, set the `ImageSize` property in your class constructor.

Data Type: 2- or 3-element vector specified as `[rows cols]` or `[rows cols bands]`, where `rows` indicates height and `cols` indicates width

Methods

Classes that inherit from `ImageAdapter` must implement the `readRegion` and `close` methods to support basic region-based reading of images. The `writeRegion` method allows for incremental, region-based writing of images and is optional. Image Adapter classes that do not implement the `writeRegion` method are read-only.

close	Close ImageAdapter object
readRegion	Read region of image
writeRegion	Write block of data to region of image

See Also

blockproc

Tutorials

- [Computing Statistics for Large Images](#)

How To

- [“Defining Abstract Classes”](#)
- [“Working with Data in Unsupported Formats”](#)

ImageAdapter.close

Purpose Close ImageAdapter object

Syntax `adapter.close`

Description `adapter.close` closes the ImageAdapter object and performs any necessary clean-up, such as closing file handles. When you construct a class that inherits from the ImageAdapter class, implement this method.

Purpose Read region of image

Syntax `data = adapter.readRegion(region_start, region_size)`

Description `data = adapter.readRegion(region_start, region_size)` reads a region of the image. `region_start` and `region_size` define a rectangular region in the image. `region_start`, a two-element vector, specifies the [row col] of the first pixel (minimum-row, minimum-column) of the region. `region_size`, a two-element vector, specifies the size of the requested region in [rows cols]. When you construct a class that inherits from the ImageAdapter class, implement this method.

ImageAdapter.writeRegion

Purpose Write block of data to region of image

Syntax `adapter.writeRegion(region_start, region_data)`

Description `adapter.writeRegion(region_start, region_data)` writes a contiguous block of data to a region of the image. The method writes the block of data specified by the `region_data` argument. The two-element vector, `region_start`, specifies the [row col] location of the first pixel (minimum-row, minimum-column) of the target region in the image. When you construct a class that inherits from the `ImageAdapter` class, implement this method if you need to write data.

Purpose 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 information about the basic attributes of the target image in a separate figure. `imageinfo` gets information about 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> . Note For single or <code>int16</code> images, <code>imageinfo</code> returns a class value of <code>double</code> , because image objects convert the <code>CData</code> of these images to <code>double</code> .
Image type	One of the image types identified by the Image Processing Toolbox software: 'intensity', 'truecolor', 'binary', or 'indexed'.

imageinfo

Attribute Name	Value
Minimum intensity or index	For grayscale images, this value represents the lowest intensity value of any pixel. For indexed images, this value represents the lowest index value into a color map. Not included for 'binary' or 'truecolor' images.
Maximum intensity or index	For grayscale images, this value represents the highest intensity value of any pixel. For indexed images, this value represents the highest index value into a color map. 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, NITF, Interfile, or Analyze file.

`imageinfo(info)` creates an Image Information tool containing the image metadata in the structure `info`. `info` is a structure returned by the functions `imfinfo`, `dicominfo`, `nitfinfo`, `interfileinfo`, or `analyze75info`. `info` can also 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.

Examples

```
imageinfo('peppers.png')
```

```
h = imshow('bag.png');  
info = imfinfo('bag.png');  
imageinfo(h,info);
```

```
imshow('canoe.tif');  
imageinfo;
```

See Also

[analyze75info](#) | [dicominfo](#) | [imatattributes](#) | [imfinfo](#) | [imformats](#) | [imtool](#) | [interfileinfo](#) | [nitfinfo](#)

imagemodel

Purpose Image Model object

Syntax `imgmodel = imagemodel(himage)`

Description `imgmodel = imagemodel(himage)` creates an image model object associated with a target image. The target image `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.

`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'.

API Functions

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. Brief descriptions of these methods follow.

Methods

`imagemodel` supports the following methods. Type methods `imagemodel` to see a list of methods, or type `help imagemodel/methodname` for more information about a specific method.

getClassType — Return class of image

`str = getClassType(imgmodel)` returns a string indicating the class of the image, where `imgmodel` is a valid image model and `str` is a text string, such as 'uint8'.

getDisplayRange — Return display range of intensity image

`disp_range = getDisplayRange(imgmodel)`, where `imgmodel` is a valid image model and `disp_range` is an array of doubles such as `[0 255]`, 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.

getImageHeight — Return number of rows

`height = getImageHeight(imgmodel)`, where `imgmodel` is a valid image model and `height` is a double scalar, returns a double scalar containing the number of rows.

getImageType — Return image type

`str = getImageType(imgmodel)`, where `imgmodel` is a valid image model and `str` is one of the text strings ('intensity', 'truecolor', 'binary', or 'indexed'), returns a text string indicating the image type.

getImageWidth — Return number of columns

`width = getImageWidth(imgmodel)`, where `imgmodel` is a valid image model and `width` is a double scalar, returns a double scalar containing the number of columns.

getMinIntensity — Return minimum value in image

`minval = getMinIntensity(imgmodel)`, where `imgmodel` is a valid image model and `minval` is a numeric value, returns the minimum value in the image calculated as `min(Image(:))`. 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. The class of `minval` depends on the class of the target image.

getMaxIntensity — Return maximum value in image

`maxval = getMaxIntensity(imgmodel)`, where `imgmodel` is a valid image model and `maxval` is a numeric value, returns the maximum value in the image calculated as `max(Image(:))`. 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. The class of `maxval` depends on the class of the target image.

getNumberFormatFcn — Return handle to function that converts numeric value into string

`fun = getNumberFormatFcn(imgmodel)`, where `imgmodel` is a valid image model, returns the handle to a function that converts a numeric value into a string.

For example, `str = fun(getPixelValue(imgmodel, 100, 100))` converts the numeric return value of the `getPixelValue` method into a text string.

getPixelInfoString — Return value of specific pixel as text string

`str = getPixelInfoString(imgmodel, row, column)`, where `str` is a character array, `imgmodel` is a valid image model, and `row` and `column` are numeric scalar values, returns a text string containing the value of the pixel at the location specified by `row` and `column`. For example, for an RGB image, the method returns a text string such as `'[66 35 60]'`.

getPixelRegionFormatFcn — Return handle to function that formats value of pixel into text string

`fun = getPixelRegionFormatFcn(imgmodel)`, takes a valid image model, `imgmodel`, and returns `fun`, a handle to a function that accepts the location (`row`, `column`) 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.

```
str = fun(100,100)
```

getPixelValue — Return value of specific pixel as numeric array

`val = getPixelValue(imgmodel, row, column)`, where `imgmodel` is a valid image model and `row` and `column` are numeric scalar values, returns the value of the pixel at the location specified by `row` and `column` as a numeric array. The class of `val` depends on the class of the target image.

getDefaultPixelInfoString — Return pixel information type as text string

`str = getDefaultPixelInfoString(imgmodel)`, where `imgmodel` is a valid image model, returns a text string indicating the pixel information type. This string can be used in place of actual pixel information

values. Depending on the image type, `str` can be the text string 'Intensity', '[R G B]', 'BW', or '<Index> [R G B]'.

getDefaultPixelRegionString — Return type of information displayed in Pixel Region tool

`str = getDefaultPixelRegionString(imgmodel)`, where `imgmodel` is a valid image model, 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. Depending on the image type, `str` can be the text string '000', 'R:000 G:000 B:000', '0', or '<000> R:0.00 G:0.00 B:0.00'.

getScreenPixelRGBValue — Return screen display value of specific pixel

`val = getScreenPixelRGBValue(imgmodel, row, col)` returns the screen display value of the pixel at the location specified by `row` and `col` as a double array. `imgmodel` is a valid image model, `row` and `col` are numeric scalar values, and `val` is an array of doubles, such as [0.2 0.5 0.3].

Examples

Create 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`

imapplymatrix

Purpose Linear combination of color channels

Syntax
`Y = imapplymatrix(M,X)`
`Y = imapplymatrix(M,X,C)`
`Y = imapplymatrix(..., output_type)`

Description `Y = imapplymatrix(M,X)` computes the linear combination of the rows of `M` with the color channels of `X`. The output data type is the same as the type of `X`.

`Y = imapplymatrix(M,X,C)` computes the linear combination of the rows of `M` with the color channels of `X`, adding the corresponding constant value from `C` to each combination. The output data type is the same as the type of `X`.

`Y = imapplymatrix(..., output_type)` returns the result of the linear combination in an array of type `output_type`.

Input Arguments

M

Matrix that contains the coefficients of the weighting matrix. If `X` is m -by- n -by- p , `M` must be q -by- p , where q is in the range $[1,p]$.

X

An image.

C

A vector with the same number of elements as the number of rows in `M`, or a scalar applied to every channel.

output_type

A string that describes the output type. Possible values include `uint8`, `int8`, `uint16`, `int16`, `uint32`, `int32`, `single`, or `double`.

Output Arguments

Y

Array that contains the linear combination of the rows of M with the color channels of X.

Examples


Convert RGB values to grayscale:

```
RGB = imread('peppers.png');  
M = [0.30, 0.59, 0.11];  
gray = imapplymatrix(M, RGB);  
figure  
subplot(1,2,1), imshow(RGB), title('Original RGB')  
subplot(1,2,2), imshow(gray), title('Grayscale Conversion')
```

Original RGB



Grayscale Conversion



See Also

[imlincomb](#) | [immultiply](#)

imattributes

Purpose 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> .
	<hr/> Note For single 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> . <hr/>

Attribute Name	Value
Image type	One of the image types identified by the Image Processing Toolbox software: 'intensity', 'truecolor', 'binary', or 'indexed'.
Minimum intensity	For intensity images, this value represents the lowest intensity value of any pixel. For indexed images, this value represents the lowest index value into a color map. Not included for 'binary' or 'truecolor' images.
Maximum intensity	For intensity images, this value represents the highest intensity value of any pixel. For indexed images, this value represents the highest index value into a color map. Not included for 'binary' or 'truecolor' images.

`attrs = imattributes(imgmodel)` returns information about the image represented by the image model object, `imgmodel`.

Examples

Retrieve the attributes of a grayscale 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'
```

Retrieve the attributes of a truecolor image.

imattributes

```
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`

Purpose	Bottom-hat filtering
Syntax	<pre>IM2 = imbothat(IM,SE) IM2 = imbothat(IM,NHOOD)</pre>
Description	<p><code>IM2 = imbothat(IM,SE)</code> performs morphological bottom-hat filtering on the grayscale or binary input image, <code>IM</code>, returning the filtered image, <code>IM2</code>. The argument <code>SE</code> is a structuring element returned by the <code>strel</code> function. <code>SE</code> must be a single structuring element object, not an array containing multiple structuring element objects.</p> <p><code>IM2 = imbothat(IM,NHOOD)</code> performs morphological bottom-hat filtering where <code>NHOOD</code> is an array of 0's and 1's that specifies the size and shape of the structuring element. This is equivalent to <code>imbothat(IM,strel(NHOOD))</code>.</p>
Class Support	<code>IM</code> 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.
Examples	<p>Top-hat filtering and bottom-hat filtering can be used together to enhance contrast in an image.</p> <ol style="list-style-type: none">1 Read the image into the MATLAB workspace. <pre>I = imread('pout.tif'); imshow(I)</pre>



- 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. (In other words, use this function to clear the image border.) IM can be a grayscale or binary image. The output image, IM2, is grayscale 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 grayscale 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
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.

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.

Examples

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
```

imclearborder

```
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
```

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
```

Algorithms

`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.

References


[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer, 1999, pp. 164-165.

See Also

`conndef`

Purpose	Morphologically close image
Syntax	<code>IM2 = imclose(IM,SE)</code> <code>IM2 = imclose(IM,NHOOD)</code>
Description	<p><code>IM2 = imclose(IM,SE)</code> performs morphological closing on the grayscale or binary image <code>IM</code>, returning the closed image, <code>IM2</code>. The structuring element, <code>SE</code>, must be a single structuring element object, as opposed to an array of objects. The morphological close operation is a dilation followed by an erosion, using the same structuring element for both operations.</p> <p><code>IM2 = imclose(IM,NHOOD)</code> performs closing with the structuring element <code>strel(NHOOD)</code>, where <code>NHOOD</code> is an array of 0's and 1's that specifies the structuring element neighborhood.</p>
Class Support	<code>IM</code> can be any numeric or logical class and any dimension, and must be nonsparse. If <code>IM</code> is logical, then <code>SE</code> must be flat. <code>IM2</code> has the same class as <code>IM</code> .
Examples	<p>Use <code>imclose</code> to join the circles in the image together by filling in the gaps between them and by smoothing their outer edges.</p> <ol style="list-style-type: none">1 Read the image into the MATLAB workspace and view it. <pre>originalBW = imread('circles.png'); imshow(originalBW);</pre>

imclose



- 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);
```

- 3 Perform a morphological close operation on the image.

```
closeBW = imclose(originalBW,se);  
figure, imshow(closeBW)
```



See Also

[imdilate](#) | [imerode](#) | [imopen](#) | [strel](#)

Purpose Choose Colormap tool

Syntax

```
imcolormaptool  
imcolormaptool(hclientfig)  
hfig = imcolormaptool(...)
```

Description

The Choose Colormap tool is an interactive colormap selection tool that allows you to change the colormap of the target (current) figure by selecting a colormap from a list of MATLAB colormap functions or workspace variables, or by entering a custom MATLAB expression.

`imcolormaptool` launches the Choose Colormap tool in a separate figure, which is associated with the target figure.

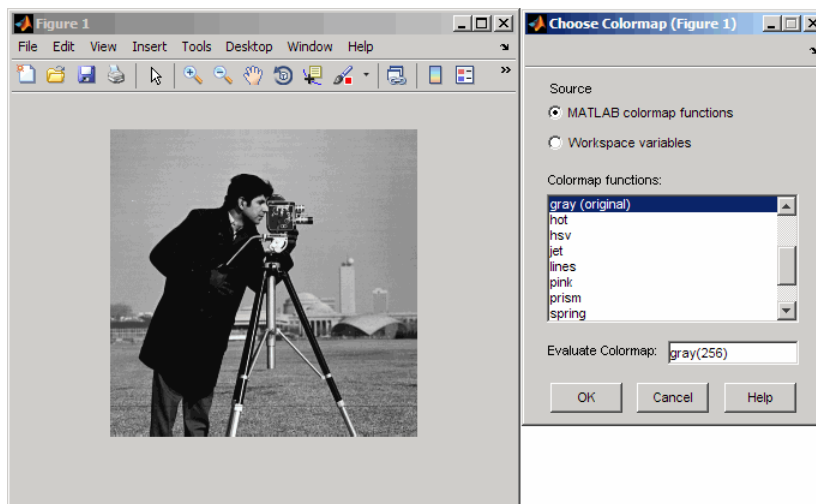
`imcolormaptool(hclientfig)` launches the Choose Colormap tool using `hclientfig` as the target figure. `hclientfig` must contain either a grayscale or an indexed image.

`hfig = imcolormaptool(...)` returns a handle to the Choose Colormap tool, `hfig`.

Examples

```
h = figure;  
imshow('cameraman.tif');  
imcolormaptool(h);
```

imcolormaptool



Choose Colormap Tool

See Also

[colormap](#) | [imshow](#) | [imtool](#)

Purpose Complement image

Syntax `IM2 = imcomplement(IM)`

Description `IM2 = imcomplement(IM)` computes the complement of the image `IM`. `IM` can be a binary, grayscale, 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 grayscale 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);
```

imcomplement

`imshow(I), figure, imshow(J)`



Original Image




Complement Image

See Also

`imabsdiff` | `imadd` | `imdivide` | `imlincomb` | `immultiply` | `imsubtract`

Purpose	Create contour plot of image data
Syntax	<pre>imcontour(I) imcontour(I,n) imcontour(I,v) imcontour(x,y,...) imcontour(...,LineStyle) [C,handle] = imcontour(...)</pre>
Description	<p><code>imcontour(I)</code> draws a contour plot of the grayscale image <code>I</code>, automatically setting up the axes so their orientation and aspect ratio match the image.</p> <p><code>imcontour(I,n)</code> draws a contour plot of the grayscale image <code>I</code>, automatically setting up the axes so their orientation and aspect ratio match the image. <code>n</code> 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.</p> <p><code>imcontour(I,v)</code> draws a contour plot of <code>I</code> with contour lines at the data values specified in vector <code>v</code>. The number of contour levels is equal to <code>length(v)</code>.</p> <p><code>imcontour(x,y,...)</code> uses the vectors <code>x</code> and <code>y</code> to specify the <code>x</code>- and <code>y</code>-axis limits.</p> <p><code>imcontour(...,LineStyle)</code> draws the contours using the line type and color specified by <code>LineStyle</code>. Marker symbols are ignored.</p> <p><code>[C,handle] = imcontour(...)</code> returns the contour matrix <code>C</code> and a handle to an <code>hggroup</code> object.</p>
Class Support	The input image can be of class <code>uint8</code> , <code>uint16</code> , <code>int16</code> , <code>single</code> , <code>double</code> , or <code>logical</code> .
Examples	<pre>I = imread('circuit.tif'); imcontour(I,3)</pre>

imcontour




See Also

[clabel](#) | [contour](#) | [LineSpec](#)

Purpose Adjust Contrast tool

Syntax
`imcontrast`
`imcontrast(h)`
`hfigure = imcontrast(...)`

Description `imcontrast` creates an Adjust Contrast tool in a separate figure that is associated with the grayscale image in the current figure, called the target image. The Adjust Contrast tool is an interactive contrast and brightness adjustment tool, shown in the following figure, that you can use to adjust the black-to-white mapping used to display the image. When you use the tool, `imcontrast` adjusts the contrast of the displayed image by modifying the axes `CLim` property. To modify the actual pixel values in the target image, click the **Adjust Data** button. (This button is unavailable until you make a change to the contrast of the image.) For more information about using the tool, see “Tips” on page 3-322.



Note The Adjust Contrast tool can handle grayscale images of class `double` and `single` with data ranges beyond the default display range, which is `[0 1]`. For these images, `imcontrast` sets the histogram limits to fit the image data range, with padding at the upper and lower bounds.

`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 an axes or figure handle, `imcontrast` uses the first image returned by `findobj(H, 'Type', 'image')`.

`hfigure = imcontrast(...)` returns a handle to the Adjust Contrast tool figure.

Tips

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 (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.

For more information about using the Adjust Contrast tool, see "Adjusting Image Contrast Using the Adjust Contrast Tool".

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 before clicking and dragging the mouse accelerates the rate of change; holding down the **Shift** key before clicking and dragging the mouse slows the rate of change. Keys must be pressed before clicking and dragging.

Examples

```
imshow('pout.tif')
imcontrast(gca)
```

See Also

```
imadjust | imtool | stretchlim
```


Purpose


Crop image

Syntax


```
I = imcrop
I2 = imcrop(I)
X2 = imcrop(X, map)
I = imcrop(h)
I2 = imcrop(I, rect)
X2 = imcrop(X, map, rect)
[...] = imcrop(x, y,...)
[I2 rect] = imcrop( )
[X,Y,I2,rect] = imcrop( )
```

Description

`I = imcrop` creates an interactive Crop Image tool associated with the image displayed in the current figure, called the target image. The Crop Image tool is a moveable, resizable rectangle that you can position interactively using the mouse. When the Crop Image tool is active, the pointer changes to cross hairs  when you move it over the target image. Using the mouse, you specify the crop rectangle by clicking and dragging the mouse. You can move or resize the crop rectangle using the mouse. When you are finished sizing and positioning the crop rectangle, create the cropped image by double-clicking the left mouse button or by choosing **Crop Image** from the context menu. `imcrop` returns the cropped image, `I`. The following figure illustrates the Crop Image tool with the context menu displayed. For more information about the interactive capabilities of the tool, see the table that follows.



Interactive Behavior	Description
Deleting the Crop Image tool.	Press Backspace , Escape or Delete , or right-click inside the crop rectangle and select Cancel from the context menu. Note: If you delete the ROI, the function returns empty values.
Resizing the Crop Image tool.	Select any of the resize handles on the crop rectangle. The pointer changes to a double-headed arrow \leftrightarrow . Click and drag the mouse to resize the crop rectangle.

Interactive Behavior	Description
Moving the Crop Image tool.	Move the pointer inside the boundary of the crop rectangle. The pointer changes to a fleur shape  . Click and drag the mouse to move the rectangle over the image.
Changing the color used to display the crop rectangle.	Right-click inside the boundary of the crop rectangle and select Set Color from the context menu.
Cropping the image.	Double-click the left mouse button or right-click inside the boundary of the crop rectangle and select Crop Image from the context menu.
Retrieving the coordinates of the crop rectangle.	Right-click inside the boundary of the crop rectangle and select Copy Position from the context menu. <code>imcrop</code> copies a four-element position vector (<code>[xmin ymin width height]</code>) to the clipboard.

`I2 = imcrop(I)` displays the image `I` in a figure window and creates a cropping tool associated with that image. `I` can be a grayscale image, a truecolor image, or a logical array. The cropped image returned, `I2`, is of the same type as `I`.

`X2 = imcrop(X, map)` displays the indexed image `X` in a figure using the colormap `map`, and creates a cropping tool associated with that image.

`I = imcrop(h)` creates a cropping tool associated with the image specified by handle `h`. `h` may be an image, axes, uipanel, or figure handle. If `h` is an axes, uipanel, or figure handle, the cropping tool acts on the first image found in the container object.

Note With these interactive syntaxes, the cropping tool blocks the MATLAB command line until you complete the operation.

`I2 = imcrop(I, rect)` crops the image `I`. `rect` is a four-element position vector [`xmin ymin width height`] that specifies the size and position of the crop rectangle.

`X2 = imcrop(X, map, rect)` crops the indexed image `X`. `map` specifies the colormap used with `X`. `rect` is a four-element position vector [`xmin ymin width height`] that specifies the size and position of the cropping rectangle.

`[...] = imcrop(x, y,...)` specifies a non-default spatial coordinate system for the target image. `x` and `y` are two-element vectors specifying `XData` and `YData`.

`[I2 rect] = imcrop()` returns the cropping rectangle in `rect`, a four-element position vector.

`[X,Y,I2,rect] = imcrop()` returns `x` and `y`, two-element vectors that specify the `XData` and `YData` of the target 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`. A truecolor image 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 or a handle, the output image has the same class as the input image except for `int16` or `single`. If the input image is `int16` or `single`, the output image is `double`.

Tips

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.

Examples

```
I = imread('circuit.tif');  
I2 = imcrop(I,[75 68 130 112]);  
imshow(I), figure, imshow(I2)
```



See Also

`imrect` | `zoom`

Purpose Dilate image

Syntax

```
IM2 = imdilate(IM,SE)
IM2 = imdilate(IM,NHOOD)
IM2 = imdilate(IM,SE,PACKOPT)
IM2 = imdilate(...,SHAPE)
```

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> . <code>IM</code> 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.

imdilate

`IM2 = imdilate(...,SHAPE)` specifies the size of the output image. `SHAPE` can have either of the following values. Default value is enclosed in braces (`{}`).

Value	Description
{ 'same' }	Make the output image the same size as the input image. If the value of <code>PACKOPT</code> is 'ispacked', <code>PADOPT</code> must be 'same'.
'full'	Compute the full dilation.

Definitions

The *binary dilation* of A by B , denoted $A \oplus B$, is defined as the set operation:

$$A \oplus B = \left\{ z \mid \left(\hat{B} \right)_z \cap A \neq \emptyset \right\},$$

where \hat{B} is the reflection of the structuring element B . In other words, it is the set of pixel locations z , where the reflected structuring element overlaps with foreground pixels in A when translated to z . Note that some people use a definition of dilation in which the structuring element is not reflected.

In the general form of *gray-scale dilation*, the structuring element has a height. The gray-scale dilation of $A(x,y)$ by $B(x,y)$ is defined as:

$$(A \oplus B)(x,y) = \max \left\{ A(x-x',y-y') + B(x',y') \mid (x',y') \in D_B \right\},$$

where D_B is the domain of the structuring element B and $A(x,y)$ is assumed to be $-\infty$ outside the domain of the image. To create a structuring element with nonzero height values, use the syntax `strel(nhood,height)`, where `height` gives the height values and `nhood` corresponds to the structuring element domain, D_B .

Most commonly, gray-scale dilation is performed with a flat structuring element ($B(x,y) = 0$). Gray-scale dilation using such a structuring element is equivalent to a local-maximum operator:

$$(A \oplus B)(x, y) = \max\{A(x - x', y - y') \mid (x', y') \in D_B\}.$$

All of the `strel` syntaxes except for `strel(nhood,height)`, `strel('arbitrary',nhood,height)`, and `strel('ball', ...)` produce flat structuring elements.

For more information about binary dilation, see [1].



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

Dilate 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')
```



Dilate a grayscale image with a rolling ball structuring element.

```
I = imread('cameraman.tif');
se = strel('ball',5,5);
I2 = imdilate(I,se);
imshow(I), title('Original')
```

imdilate

```
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

Algorithms

`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 [3].

References

[1] Gonzalez, R. C., R. E. Woods, and S. L. Eddins, *Digital Image Processing Using MATLAB*, Gatesmark Publishing, 2009.

[2] Haralick, R. M., and L. G. Shapiro, *Computer and Robot Vision*, Vol. I, Addison-Wesley, 1992, pp. 158-205.

[3] van den Boomgard, R, and R. van Balen, "Methods for Fast Morphological Image Transforms Using Bitmapped Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, Number 3, pp. 254-258, May 1992.

See Also

`bwpack` | `bwunpack` | `conv2` | `filter2` | `imclose` | `imerode` | `imopen` | `strel`

imshow_range


Purpose Display Range tool

Syntax

```
imshow_range  
imshow_range(h)  
imshow_range(hparent, himage)  
hpanel = imshow_range(...)
```

Description `imshow_range` 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_range(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_range(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_range(...)` returns a handle to the Display Range tool uipanel.

Note The Display Range tool can work with multiple images in a figure. When the pointer 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_range;
```

Import a 16-bit DICOM image and display it with its default range and scaled range in the same figure.

```
dcm = dicomread('CT-MONO2-16-ankle.dcm');  
subplot(1,2,1), imshow(dcm);  
subplot(1,2,2), imshow(dcm, []);  
imshow_range;
```

See also

`imshow`

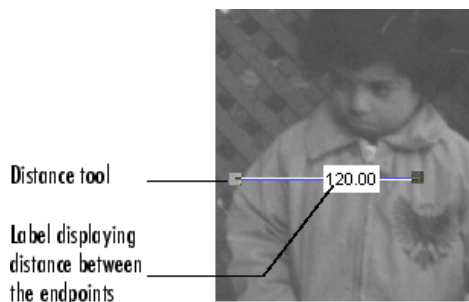
imdistline



Purpose Distance tool

Syntax
`h = imdistline`
`h = imdistline(hparent)`
`h = imdistline(..., x, y)`

Description `h = imdistline` creates a Distance tool on the current axes. The function returns `h`, a handle to an `imdistline` object.

The Distance tool is a draggable, resizable line, superimposed on an axes, that measures the distance between the two endpoints of the line. The Distance tool displays the distance in a text label superimposed over the line. The tools specifies the distance in data units determined by the `XData` and `YData` properties, which is pixels, by default. The following figure shows a Distance tool on an axes.



To move the Distance tool, position the pointer over the line, the shape changes to the fleur, . Click and drag the line using the mouse. To resize the Distance tool, move the pointer over either of the endpoints of the line, the shape changes to the pointing finger, . Click and drag the endpoint of the line using the mouse. The line also supports a context menu that allows you to control various aspects of its functioning and appearance. See Context Menu for more information. Right-click the line to access the context menu.

`h = imdistline(hparent)` creates a draggable Distance tool on the object specified by `hparent`. `hparent` specifies the Distance tool's parent, which is typically an axes object, but can also be any other object that can be the parent of an `hggroup` object.

`h = imdistline(..., x, y)` creates a Distance tool with endpoints located at the locations specified by the vectors `x` and `y`, where `x = [x1 x2]` and `y = [y1 y2]`.

Context Menu

Distance Tool Behavior	Context Menu Item
Export endpoint and distance data to the workspace	Select Export to Workspace from the context menu.
Toggle the distance label on/off.	Select Show Distance Label from the context menu.
Specify horizontal and vertical drag constraints	Select Constrain Drag from the context menu.
Change the color used to display the line.	Select Set Color from the context menu.
Delete the Distance tool object	Select Delete from the context menu.

API Functions

The Distance tool contains a structure of function handles, called an API, that can be used to retrieve distance information and control other aspects of Distance tool behavior. To retrieve this structure from the Distance tool, use the `iptgetapi` function, where `h` is a handle to the Distance tool.

`api = iptgetapi(h)`

The following table lists the functions in the API, with their syntax and brief descriptions.

Function	Description
<code>getDistance</code>	Returns the distance between the endpoints of the Distance tool. <code>dist = getDistance()</code> The value returned is in data units determined by the <code>XData</code> and <code>YData</code> properties, which is pixels, by default.
<code>getAngleFromHorizontal</code>	Returns the angle in degrees between the line defined by the Distance tool and the horizontal axis. The angle returned is between 0 and 180 degrees. (For information about how this angle is calculated, see “Tips” on page 3-341.) <code>angle = getAngleFromHorizontal()</code>
<code>setLabelTextFormatter</code>	Sets the format string used in displaying the distance label. <code>setLabelTextFormatter(str)</code> <code>str</code> is a character array specifying a format string in the form expected by <code>sprintf</code> .
<code>getLabelTextFormatter</code>	Returns a character array specifying the format string used to display the distance label. <code>str = getLabelTextFormatter()</code> <code>str</code> is a character array specifying a format string in the form expected by <code>sprintf</code> .

Function	Description
setLabelVisible	<p>Sets visibility of Distance tool text label.</p> <pre>setLabelVisible(h,TF)</pre> <p>h is the Distance tool. TF is a logical scalar. When the distance label is visible, TF is true. When the distance label is invisible, TF is false.</p>
getLabelVisible	<p>Gets visibility of Distance tool text label.</p> <pre>TF = getLabelVisible(h)</pre> <p>h is the Distance tool. TF is a logical scalar. When TF is true, the distance label is visible. When TF is false, the distance label is invisible.</p>
setPosition	<p>Sets the endpoint positions of the Distance tool.</p> <pre>setPosition(X,Y) setPosition([X1 Y1; X2 Y2])</pre>
getPosition	<p>Returns the endpoint positions of the Distance tool.</p> <pre>pos = getPosition()</pre> <p>pos is a 2-by-2 array [X1 Y1; X2 Y2].</p>
delete	<p>Deletes the Distance tool associated with the API.</p> <pre>delete()</pre>

imdistline

Function	Description
setColor	<p>Sets the color used to draw the Distance tool.</p> <pre>setColor(new_color)</pre> <p><code>new_color</code> can be a three-element vector specifying an RGB triplet, or a text string specifying the long or short names of a predefined color, such as 'white' or 'w'. For a complete list of these predefined colors and their short names, see <code>ColorSpec</code>.</p>
getColor	<p>Gets the color used to draw the ROI object <code>h</code>.</p> <pre>color = getColor(h)</pre> <p><code>color</code> is a three-element vector that specifies an RGB triplet.</p>
addNewPositionCallback	<p>Adds the function handle <code>fcn</code> to the list of new-position callback functions.</p> <pre>id = addNewPositionCallback(fcn)</pre> <p>Whenever the Distance tool changes its position, each function in the list is called with the following syntax.</p> <pre>fcn(pos)</pre> <p><code>pos</code> is a 2-by-2 array [X1 Y1; X2 Y2].</p> <p>The return value, <code>id</code>, is used only with <code>removeNewPositionCallback</code>.</p>

Function	Description
removeNewPositionCallback	<p>Removes the corresponding function from the new-position callback list.</p> <pre>removeNewPositionCallback(id)</pre> <p>id is the identifier returned by addNewPositionCallback.</p>
setPositionConstraintFcn	<p>Sets the position constraint function to be the specified function handle, fcn. Use this function to control where the Distance tool can be moved and resized.</p> <pre>setPositionConstraintFcn(fcn)</pre> <p>Whenever the Distance tool is moved or resized because of a mouse drag, the constraint function is called using the following syntax.</p> <pre>constrained_position = fcn(new_position)</pre> <p>new_position is a 2-by-2 array [X1 Y1; X2 Y2].</p>
getPositionConstraintFcn	<p>Returns the function handle of the current drag constraint function.</p> <pre>fcn = getDragConstraintFcn()</pre>

Tips

If you use imdistline with an axes that contains an image object, and do not specify a drag constraint function, users can drag the line outside the extent of the image. When used with an axes created by the plot function, the axes limits automatically expand to accommodate the movement of the line.

To understand how imdistline calculates the angle returned by getAngleToHorizontal, draw an imaginary horizontal vector from the bottom endpoint of the distance line, extending to the right. The value

returned by `getAngleToHorizontal` is the angle from this horizontal vector to the distance line, which can range from 0 to 180 degrees.

Examples

Example 1

Insert a Distance tool into an image. Use `makeConstrainToRectFcn` to specify a drag constraint function that prevents the Distance tool from being dragged outside the extent of the image. Right-click the Distance tool and explore the context menu options.

```
figure, imshow('pout.tif');
h = imdistline(gca);
api = iptgetapi(h);
fcn = makeConstrainToRectFcn('imline',...
                             get(gca,'XLim'),get(gca,'YLim'));
api.setDragConstraintFcn(fcn);
```

Example 2

Position endpoints of the Distance tool at the specified locations.

```
close all, imshow('pout.tif');
h = imdistline(gca,[10 100],[10 100]);
```

Delete the Distance tool.

```
api = iptgetapi(h);
api.delete();
```

Example 3

Use the Distance tool with `XData` and `YData` of associated image in non-pixel units. This example requires the `boston.tif` image from the Mapping Toolbox software, which includes material copyrighted by GeoEye™, all rights reserved.

```
start_row = 1478;
end_row = 2246;
meters_per_pixel = 1;
rows = [start_row meters_per_pixel end_row];
```

```
start_col = 349;
end_col = 1117;
cols = [start_col meters_per_pixel end_col];
img = imread('boston.tif','PixelRegion',{rows,cols});
figure;
hImg = imshow(img);
title('1 meter per pixel');

% Specify initial position of distance tool on Harvard Bridge.
hline = imdistline(gca,[271 471],[108 650]);
api = iptgetapi(hline);
api.setLabelTextFormatter('%02.0f meters');

% Repeat process but work with a 2 meter per pixel sampled image. Verify
% that the same distance is obtained.
meters_per_pixel = 2;
rows = [start_row meters_per_pixel end_row];
cols = [start_col meters_per_pixel end_col];
img = imread('boston.tif','PixelRegion',{rows,cols});
figure;
hImg = imshow(img);
title('2 meters per pixel');

% Convert XData and YData to meters using conversion factor.
XDataInMeters = get(hImg,'XData')*meters_per_pixel;
YDataInMeters = get(hImg,'YData')*meters_per_pixel;

% Set XData and YData of image to reflect desired units.
set(hImg,'XData',XDataInMeters,'YData',YDataInMeters);
set(gca,'XLim',XDataInMeters,'YLim',YDataInMeters);

% Specify initial position of distance tool on Harvard Bridge.
hline = imdistline(gca,[271 471],[108 650]);
api = iptgetapi(hline);
api.setLabelTextFormatter('%02.0f meters');
```

See Also

[iptgetapi](#) | [makeConstrainToRectFcn](#)

imdivide

Purpose Divide one image into another or divide image by 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.

Note On Intel architecture processors, `imdivide` can take advantage of the Intel Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The Intel IPP library is activated only if arrays `X` and `Y` are of class `uint8`, `int16`, or `single` and are of the same size and class.

Examples

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](#)

imellipse

Purpose Create draggable ellipse

Syntax

```
h = imellipse
h = imellipse(hparent)
h = imellipse(hparent, position)
H = imellipse(...,param1, val1, ...)
```

Description `h = imellipse` begins interactive placement of an ellipse on the current axes. The function returns `h`, a handle to an `imellipse` object. The ellipse has a context menu associated with it that controls aspects of its appearance and behavior—see “Interactive Behavior” on page 3-347. Right-click on the line to access this context menu.

`h = imellipse(hparent)` begins interactive placement of an ellipse on the object specified by `hparent`. `hparent` specifies the HG parent of the ellipse graphics, which is typically an axes but can also be any other object that can be the parent of an `hggroup`.


`h = imellipse(hparent, position)` creates a draggable ellipse on the object specified by `hparent`. `position` is a four-element vector that specifies the initial location of the ellipse in terms of a bounding rectangle. `position` has the form `[xmin ymin width height]`.

`H = imellipse(...,param1, val1, ...)` creates a draggable ellipse, specifying parameters and corresponding values that control the behavior of the ellipse. The following table lists the parameter available. Parameter names can be abbreviated, and case does not matter.

Parameter	Value
'PositionConstraintFcn'	Function handle that is called whenever the mouse is dragged. You can use this to control where the ellipse may be dragged. See the help for the <code>setPositionConstraintFcn</code> method for information about valid function handles.

Interactive Behavior

When you call `imellipse` with an interactive syntax, the pointer changes to a cross hairs $+$ when over an image. Click and drag the mouse to specify the size and position of the ellipse. The ellipse also supports a context menu that you can use to control aspects of its appearance and behavior. The following figure illustrates the ellipse with its context menu.



The following table lists the interactive behavior supported by `imellipse`.

Interactive Behavior	Description
Moving the entire ellipse.	Move the pointer inside the ellipse. The pointer changes to a fleur shape \blacklozenge . Click and drag the mouse to move the ellipse.
Resizing the ellipse.	Move the pointer over a resizing handle on the ellipse. The pointer changes to a double-ended arrow shape \longleftrightarrow . Click and drag the mouse to resize the ellipse.

Interactive Behavior	Description
Changing the color used to display the ellipse.	Move the pointer inside the ellipse. Right-click and select Set Color from the context menu.
Retrieving the current position of the ellipse.	Move the pointer inside the ellipse. Right-click and select Copy Position from the context menu. <code>imellipse</code> copies a four-element position vector [xmin ymin width height] to the clipboard.
Preserving the current aspect ratio of the ellipse during resizing.	Move the pointer inside the ellipse. Right-click and select Fix Aspect Ratio from the context menu.

Methods

Each `imellipse` object supports a number of methods. Type methods `imellipse` to see a complete list.

addNewPositionCallback – Add new-position callback to ROI object

See `imroi` for information.

createMask – Create mask within image

See `imroi` for information.

delete – Delete ROI object

See `imroi` for information.

getColor – Get color used to draw ROI object

See `imroi` for information.

getPosition – Return current position of ellipse

See `imrect` for information.

getPositionConstraintFcn – Return function handle to current position constraint function

See `imroi` for information.

getVertices – Return vertices on perimeter of ellipse

`vert = getVertices(h)` returns a set of vertices which lie along the perimeter of the ellipse `h`. `vert` is a N-by-2 array.

removeNewPositionCallback – Remove new-position callback from ROI object.

See `imroi` for information.

resume – Resume execution of MATLAB command line

See `imroi` for information.

setColor – Set color used to draw ROI object

See `imroi` for information.

setConstrainedPosition – Set ROI object to new position

See `imroi` for information.

setFixedAspectRatioMode – Control whether aspect ratio preserved during resize

See `imrect` for information.

setPosition – Set ellipse to new position

See `imrect` for information.

setPositionConstraintFcn – Set position constraint function of ROI object.

See `imroi` for information.

setResizable – Set resize behavior of ellipse

See `imrect` for information.

wait – Block MATLAB command line until ROI creation is finished

`vert = wait(h)` blocks execution of the MATLAB command line until you finish positioning the ROI object `h`. You indicate completion by double-clicking on the ROI object. The returned vertices, `vert`, is of the form returned by the `getVertices` method.

Tips

If you use `imellipse` with an axes that contains an image object, and do not specify a position constraint function, users can drag the ellipse outside the extent of the image and lose the ellipse. When used with an axes created by the `plot` function, the axes limits automatically expand to accommodate the movement of the ellipse.

Examples

Example 1

Create an ellipse, using callbacks to display the updated position in the title of the figure. The example illustrates using the `makeConstrainToRectFcn` to keep the ellipse inside the original `xlim` and `ylim` ranges.

```
figure, imshow('cameraman.tif');  
h = imellipse(gca, [10 10 100 100]);  
addNewPositionCallback(h,@(p) title(mat2str(p,3)));  
fcn = makeConstrainToRectFcn('imellipse',get(gca,'XLim'),get(gca,'YLim'));  
setPositionConstraintFcn(h,fcn);
```

Example 2

Interactively place an ellipse by clicking and dragging. Use `wait` to block the MATLAB command line. Double-click on the ellipse to resume execution of the MATLAB command line.

```
figure, imshow('pout.tif');  
h = imellipse;  
position = wait(h);
```

See Also

`imfreehand` | `imline` | `impoint` | `impoly` | `imrect` | `imroi` |
`iptgetapi` | `makeConstrainToRectFcn`

Purpose Erode image

Syntax

```
IM2 = imerode(IM,SE)
IM2 = imerode(IM,NHOOD)
IM2 = imerode(...,PACKOPT,M)
IM2 = imerode(...,SHAPE)
```

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 erosion; 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(...,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(..., SHAPE)` specifies the size of the output image. `SHAPE` 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>SHAPE</code> must be <code>'same'</code> .
<code>'full'</code>	Compute the full erosion.

Definitions

The *binary erosion* of A by B , denoted $A \ominus B$, is defined as the set operation $A \ominus B = \{z \mid (B_z \subseteq A)\}$. In other words, it is the set of pixel locations z , where the structuring element translated to location z overlaps only with foreground pixels in A .

In the general form of *gray-scale erosion*, the structuring element has a height. The gray-scale erosion of $A(x, y)$ by $B(x, y)$ is defined as:

$$(A \ominus B)(x, y) = \min \{A(x + x', y + y') - B(x', y') \mid (x', y') \in D_B\},$$

where D_B is the domain of the structuring element B and $A(x, y)$ is assumed to be $+\infty$ outside the domain of the image. To create a structuring element with nonzero height values, use the syntax `strel(nhood, height)`, where `height` gives the height values and `nhood` corresponds to the structuring element domain, D_B .

Most commonly, gray-scale erosion is performed with a flat structuring element ($B(x, y) = 0$). Gray-scale erosion using such a structuring element is equivalent to a local-minimum operator:

$$(A \ominus B)(x, y) = \min \{A(x + x', y + y') \mid (x', y') \in D_B\}.$$

All of the `strel` syntaxes except for `strel(nhood, height)`, `strel('arbitrary', nhood, height)`, and `strel('ball', ...)` produce flat structuring elements.

For more information on binary erosion, see [1].


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

Erode 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)
```



Erode 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')
```



Algorithms

`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 [3].

References

[1] Gonzalez, R. C., R. E. Woods, and S. L. Eddins, *Digital Image Processing Using MATLAB*, Gatesmark Publishing, 2009.

[2] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Vol. I, Addison-Wesley, 1992, pp. 158-205.

[3] van den Boomgard, R, and R. van Balen, "Methods for Fast Morphological Image Transforms Using Bitmapped Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, Number 3, pp. 254-258, May 1992.

See Also

`bwpack` | `bwunpack` | `conv2` | `filter2` | `imclose` | `imdilate` | `imopen`
| `strel`

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 a constant intensity value, and whose external boundary pixels all have a lower value.

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.


imextendedmax

Class Support

I can be of any nonsparse numeric class and any dimension. BW has the same size as I and is always logical.

Examples

```
I = imread('glass.png');  
BW = imextendedmax(I,80);  
imshow(I), figure, imshow(BW)
```



Original Image



Extended Maxima Image

References

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

See Also

`conndef` | `imextendedmin` | `imhmax` | `imreconstruct` | `imregionalmax`

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 a constant intensity value, and whose external boundary pixels all have a higher value.

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.


imextendedmin

Class Support

I can be of any nonsparse numeric class and any dimension. BW has the same size as I and is always logical.

Examples

```
I = imread('glass.png');  
BW = imextendedmin(I,50);  
imshow(I), figure, imshow(BW)
```



Original Image



Extended Minima Image

References

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

See Also

`conndef` | `imextendedmax` | `imhmin` | `imreconstruct` | `imregionalmin`

Purpose

Fill image regions and holes

Syntax

```
BW2 = imfill(BW)
[BW2,locations] = imfill(BW)
BW2 = imfill(BW,locations)
BW2 = imfill(BW,'holes')
I2 = imfill(I)
BW2 = imfill(BW,locations,conn)
```

Description

`BW2 = imfill(BW)` displays the binary image `BW` on the screen and lets you define the region to fill by selecting points interactively by using the mouse. To use this interactive syntax, `BW` must be a 2-D image. Press **Backspace** or **Delete** to remove the previously selected point. A shift-click, right-click, or double-click selects a final point and starts the fill operation. Pressing **Return** finishes the selection without adding a point.

`[BW2,locations] = imfill(BW)` returns the locations of points selected interactively in `locations`. `locations` is a vector of linear indices into the input image. To use this interactive syntax, `BW` must be a 2-D image.

`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 grayscale image `I`. In this syntax, a hole is defined as an area of dark pixels surrounded by lighter pixels.

`BW2 = imfill(BW,locations,conn)` fills the area defined by `locations`, 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.

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)
```

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 a grayscale image.

```
I = imread('tire.tif');
I2 = imfill(I,'holes');
figure, imshow(I), figure, imshow(I2)
```

imfill



Original Image



Filled Image

Algorithms

`imfill` uses an algorithm based on morphological reconstruction [1].

References

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 173-174.

See Also

`bwselect` | `imreconstruct` | `roifill`

Purpose N-D filtering of multidimensional images

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 `logical` or 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 or `logical` 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.

Tips

On Intel architecture processors, `imfilter` can take advantage of the Intel Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The Intel IPP library is activated only if `A` and `H` are both two-dimensional and `A` is of class `uint8`, `uint16`, `int16`, `single`, or `double`.

When the Intel IPP library is used, `imfilter` has different rounding behavior on some processors. Normally, when `A` is an integer class, filter outputs such as 1.5, 4.5, etc are rounded away from zero. However, when the Intel IPP library is used, these values are rounded toward zero.

Also, a special case exists if `A` is of class `single` or `double`, your input image contains NaN values, and your filtering kernel contains zero values. In this situation, you will receive different results if the Intel IPP library is enabled versus if it is not enabled. If you write code that depends on output values in the same neighborhood as NaN values in an input image, be aware that your code may behave differently on different machines, depending on the availability of the Intel IPP library.

To disable the Intel IPP library, use this command:

```
iptsetpref('UseIPPL', false)
```

Example

This example shows the different results obtained with the Intel IPP library enabled versus not enabled in the special case when you are filtering an array containing NaN elements with a filter containing zero values.

First, consider the case in which the Intel IPP library is enabled (the default case).

```
A = [2 2 2 2; 2 NaN 2 2; 2 2 2 2; 2 2 2 2];
h = [1 0; 0 1];
imfilter(A, h)
```

```
ans =
```

NaN	NaN	4	2
NaN	NaN	4	2
4	4	4	2
2	2	2	2

Now, compare this result to that obtained when you disable the Intel IPP library.

```
iptsetpref('UseIPPL', false);
imfilter(A, h)
```

imfilter

```
ans =  
  
    NaN     4     4     2  
     4    NaN     4     2  
     4     4     4     2  
     2     2     2     2
```

As you can see, your output is not the same.

Examples

Read a color image into the workspace and view it.

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



Original Image


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 `originalRGB` to create a new image, `filteredRGB`.

```
filteredRGB = imfilter(originalRGB, h);
```

```
figure, imshow(filteredRGB)
```



Filtered Image

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
```

Name	Size	Bytes	Class	Attributes
filteredRGB	384x512x3	589824	uint8	
h	37x37	10952	double	
originalRGB	384x512x3	589824	uint8	

Specify the `replicate` boundary option.

```
boundaryReplicateRGB = imfilter(originalRGB, h, 'replicate');  
figure, imshow(boundaryReplicateRGB)
```

imfilter




Image with Replicate Boundary

See Also

`conv2` | `convn` | `filter2` | `fspecial` | `ippl`

Purpose

Find circles using circular Hough transform

Syntax

```
centers = imfindcircles(A,radius)
[centers,radii] = imfindcircles(A,radiusRange)
[centers,radii,metric] = imfindcircles(A,radiusRange)
[centers,radii,metric] = imfindcircles(___,Name,Value)
```

Description

`centers = imfindcircles(A,radius)` finds the circles in image `A` whose radii are approximately equal to `radius`. The output, `centers`, is a two-column matrix containing the x,y coordinates of the circle centers in the image.

`[centers,radii] = imfindcircles(A,radiusRange)` finds circles with radii in the range specified by `radiusRange`. The additional output argument, `radii`, contains the estimated radii corresponding to each circle center in `centers`.

`[centers,radii,metric] = imfindcircles(A,radiusRange)` also returns a column vector, `metric`, containing the magnitudes of the accumulator array peaks for each circle (in descending order). The rows of `centers` and `radii` correspond to the rows of `metric`.

`[centers,radii,metric] = imfindcircles(___,Name,Value)` specifies additional options with one or more `Name,Value` pair arguments, using any of the previous syntaxes.

Input Arguments

A - Input image

grayscale image | truecolor image | binary image

Input image is the image in which to detect circular objects, specified as a grayscale, truecolor, or binary image.

Data Types

single | double | int16 | uint8 | uint16 | logical

radius - Circle radius

scalar integer

Circle radius is the approximate radius of the circular objects you want to detect, specified as a scalar integer of any numeric type.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32 | uint64

radiusRange - Range of radii

two-element vector of integers

Range of radii for the circular objects you want to detect, specified as a two-element vector, [*rmin* *rmax*], of integers of any numeric type.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32 | uint64

Name-Value Pair Arguments

Specify optional comma-separated pairs of **Name**, **Value** arguments, where **Name** is the argument name and **Value** is the corresponding value. **Name** must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as **Name1**, **Value1**, ..., **NameN**, **ValueN**.

Example: 'ObjectPolarity', 'bright' specifies bright circular objects on a dark background.

ObjectPolarity - Object polarity

'bright' (default) | 'dark'

Object polarity indicates whether the circular objects are brighter or darker than the background, specified as the comma-separated pair consisting of 'ObjectPolarity' and either of the text strings in the following table.

'bright'	The circular objects are brighter than the background.
'dark'	The circular objects are darker than the background.

Example: 'ObjectPolarity', 'bright' specifies bright circular objects on a dark background.

Method - Computation method

'PhaseCode' (default) | 'TwoStage'

Computation method is the technique used to compute the accumulator array, specified as the comma-separated pair consisting of 'Method' and either of the text strings in the following table.

'PhaseCode'	Atherton and Kerbyson's [1] phase coding method . This is the default.
'TwoStage'	The method used in two-stage circular Hough transform [2], [3].

Example: 'Method', 'PhaseCode' specifies the Atherton and Kerbyson's phase coding method.

Sensitivity - Sensitivity factor

nonnegative scalar between 0 and 1 | 0.85 (default)

Sensitivity factor is the sensitivity for the circular Hough transform accumulator array, specified as the comma-separated pair consisting of 'Sensitivity' and a nonnegative scalar value in the range [0, 1]. As you increase the sensitivity factor, `imfindcircles` detects more circular objects, including weak and partially obscured circles. Higher sensitivity values also increase the risk of false detection.

Example: 'Sensitivity', 0.3 sets the sensitivity factor to 0.3.

EdgeThreshold - Edge gradient threshold

nonnegative scalar between 0 and 1

imfindcircles

Edge gradient threshold sets the gradient threshold for determining edge pixels in the image, specified as the comma-separated pair consisting of 'EdgeThreshold' and a nonnegative scalar value in the range [0, 1]. Specify 0 to set the threshold to zero-gradient magnitude. Specify 1 to set the threshold to the maximum gradient magnitude. `imfindcircles` detects more circular objects (with both weak and strong edges) when you set the threshold to a lower value. It detects fewer circles with weak edges as you increase the value of the threshold. By default, `imfindcircles` chooses the edge gradient threshold automatically using the function `graythresh`.

Example: 'EdgeThreshold', 0.5 sets the edge gradient threshold to 0.5.

Output Arguments

centers - Coordinates of circle centers

two-column matrix

Coordinates of the circle centers, returned as a P-by-2 matrix containing the *x*-coordinates of the circle centers in the first column and the *y*-coordinates in the second column. The number of rows, P, is the number of circles detected. `centers` is sorted based on the strength of the circles.

radii - Estimated radii

column vector

The estimated radii for the circle centers, returned as a column vector. The radius value at `radii(j)` corresponds to the circle centered at `centers(j, :)`.

metric - Circle strengths

column vector

Circle strengths is the relative strengths for the circle centers, returned as a vector. The value at `metric(j)` corresponds to the circle with radius `radii(j)` centered at `centers(j, :)`.

Tips

- Specify a relatively small `radiusRange` for better accuracy. A good rule of thumb is to choose `radiusRange` such that $r_{\max} < 3 \cdot r_{\min}$ and $(r_{\max} - r_{\min}) < 100$.
- The accuracy of `imfindcircles` is limited when the value of `radius` (or `rmin`) is less than 10.
- The radius estimation step is typically faster if you use the (default) `'PhaseCode'` method instead of `'TwoStage'`.
- Both computation methods, `'PhaseCode'` and `'TwoStage'` are limited in their ability to detect concentric circles. The results for concentric circles can vary depending on the input image.
- `imfindcircles` does not find circles with centers outside the domain of the image.
- `imfindcircles` preprocesses binary (logical) images to improve the accuracy of the result. It converts truecolor images to grayscale using the function `rgb2gray` before processing them.

Examples

Detection of Five Strongest Circles in an Image

Read the image into the workspace and display it.

```
A = imread('coins.png');  
imshow(A)
```

imfindcircles



Find all the circles with radius r such that $15 \leq r \leq 30$.


```
[centers, radii, metric] = imfindcircles(A,[15 30]);
```

Retain the five strongest circles according to the metric values.

```
centersStrong5 = centers(1:5,:);  
radiiStrong5 = radii(1:5);  
metricStrong5 = metric(1:5);
```

Draw the five strongest circle perimeters.

```
viscircles(centersStrong5, radiiStrong5,'EdgeColor','b');
```




Detection of Bright and Dark Circles in an Image

Read the image into the workspace and display it.

```
A = imread('circlesBrightDark.png');  
imshow(A)
```

imfindcircles



Define the radius range.

```
Rmin = 30;  
Rmax = 65;
```

Find all the bright circles in the image within the radius range.


```
[centersBright, radiiBright] = imfindcircles(A,[Rmin Rmax], 'ObjectPolarity', 'bright');
```

Find all the dark circles in the image within the radius range.

```
[centersDark, radiiDark] = imfindcircles(A,[Rmin Rmax], 'ObjectPolarity', 'dark');
```


Plot bright circles in blue.

```
viscircles(centersBright, radiiBright,'EdgeColor','b');
```



Plot dark circles in dashed red boundaries.

```
viscircles(centersDark, radiiDark,'LineStyle','--');
```




Algorithm

Function `imfindcircles` uses a Circular Hough Transform (CHT) based algorithm for finding circles in images. This approach is used because of its robustness in the presence of noise, occlusion and varying illumination.

The CHT is not a rigorously specified algorithm, rather there are a number of different approaches that can be taken in its implementation. However, by and large, there are three essential steps which are common to all.

- 1 Accumulator Array Computation.

Foreground pixels of high gradient are designated as being candidate pixels and are allowed to cast 'votes' in the accumulator array. In a classical CHT implementation, the candidate pixels vote in pattern around them that forms a full circle of a fixed radius. Figure 1a shows an example of a candidate pixel lying on an actual circle (solid circle) and the classical CHT voting pattern (dashed circles) for the candidate pixel.



(a)

Figure 1: classical CHT voting pattern

2 Center Estimation

The votes of candidate pixels belonging to an image circle tend to accumulate at the accumulator array bin corresponding to the circle's center. Therefore, the circle centers are estimated by detecting the peaks in the accumulator array. Figure 1b shows an example of the candidate pixels (solid dots) lying on an actual circle (solid circle), and their voting patterns (dashed circles) which coincide at the center of the actual circle.

3 Radius Estimation

If the same accumulator array is used for more than one radius value, as is commonly done in CHT algorithms, radii of the detected circles have to be estimated as a separate step.

Function `imfindcircles` provides two algorithms for finding circles in images: Phase-Coding (default) and Two-Stage. Both share some common computational steps, but each has its own unique aspects as well.

The common computational features shared by both algorithms are as follow:

- Use of 2-D Accumulator Array:

The classical Hough Transform requires a 3-D array for storing votes for multiple radii, which results in large storage requirements and long processing times. Both the Phase-Coding and Two-Stage methods solve this problem by using a single 2-D accumulator array for all the radii. Although this approach requires an additional step of radius estimation, the overall computational load is typically lower, especially when working over large radius range. This is a widely adopted practice in modern CHT implementations.

- Use of Edge Pixels

Overall memory requirements and speed is strongly governed by the number of candidate pixels. To limit their number, the gradient magnitude of the input image is threshold so that only pixels of high gradient are included in tallying votes.

- Use of Edge Orientation Information:

Another way to optimize performance is to restrict the number of bins available to candidate pixels. This is accomplished by utilizing locally available edge information to only permit voting in a limited interval along direction of the gradient (Figure 2).




Figure 2: Voting mode: multiple radii, along direction of the gradient

r_{\min}	Minimum search radius
r_{\max}	Maximum search radius
r_{actual}	Radius of the circle that the candidate pixel belongs to
c_{\min}	Center of the circle of radius r_{\min}
c_{actual}	Center of the circle of radius r_{actual}

The two CHT methods employed by function `imfindcircles` fundamentally differ in the manner by which the circle radii are computed.

- Two-Stage

Radii are explicitly estimated utilizing the estimated circle centers along with image information. The technique is based on computing radial histograms; see references 1 & 2 for a detailed explanation.

- Phase-Coding

The key idea in Phase Coding is the use of complex values in the accumulator array with the radius information encoded in the phase of the array entries. The votes cast by the edge pixels contain information not only about the possible center locations but also about the radius of the circle associated with the center location. Unlike the Two-Stage method where radius has to be estimated explicitly using radial histograms, in Phase Coding the radius can be estimated by simply decoding the phase information from the estimated center location in the accumulator array. (see reference 3).

References

- [1] T.J Atherton, D.J. Kerbyson. "Size invariant circle detection." *Image and Vision Computing*. Volume 17, Number 11, 1999, pp. 795-803.
- [2] H.K Yuen, .J. Princen, J. Illingworth, and J. Kittler. "Comparative study of Hough transform methods for circle finding." *Image and Vision Computing*. Volume 8, Number 1, 1990, pp. 71-77.

imfindcircles

[3] E.R. Davies, *Machine Vision: Theory, Algorithms, Practicalities*.
Chapter 10. 3rd Edition. Morgan Kaufman Publishers, 2005,

See Also

hough | houghpeaks | houghlines | viscircles

Purpose Create draggable freehand region

Syntax

```
h = imfreehand
h = imfreehand(hparent)
h = imfreehand(...,param1, val1,...)
```

Description `h = imfreehand` begins interactive placement of a freehand region of interest on the current axes. The function returns `h`, a handle to an `imfreehand` object. A freehand region of interest can be dragged interactively using the mouse and supports a context menu that controls aspects of its appearance and behavior. See “Interactive Behavior” on page 3-386.

`h = imfreehand(hparent)` begins interactive placement of a freehand region of interest on the object specified by `hparent`. `hparent` specifies the HG parent of the freehand region graphics, which is typically an axes, but can also be any other object that can be the parent of an `hgroup`.

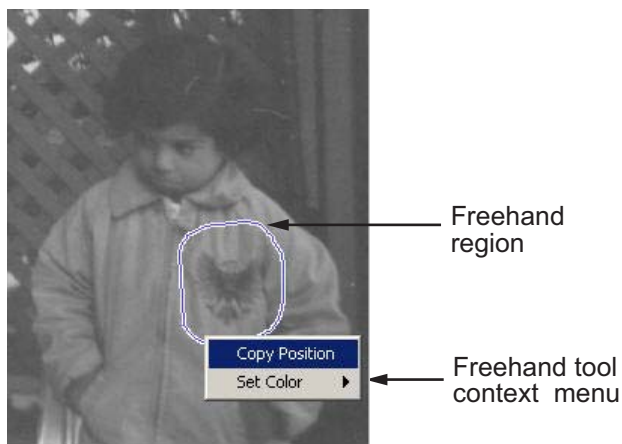
`h = imfreehand(...,param1, val1,...)` creates a freehand ROI, specifying parameters and corresponding values that control the behavior of the tool. The following table lists the parameters available. Parameter names can be abbreviated, and case does not matter.

Parameter	Description
'Closed'	Scalar logical that controls whether the freehand region is closed. When set to <code>true</code> (the default), <code>imfreehand</code> draws a straight line to connect the endpoints of the freehand line to create a closed region. If set to <code>false</code> , <code>imfreehand</code> leaves the region open.
'PositionConstraintFcn'	Function handle specifying the function that is called whenever the freehand region is dragged using the mouse. Use this parameter to control where the freehand region can be dragged. See the help for


Parameter	Description
	the setPositionConstraintFcn method for information about valid function handles.

Interactive Behavior

When you call `imfreehand` with an interactive syntax, the pointer changes to a cross hairs \oplus when positioned over an image. Click and drag the mouse to draw the freehand region. By default, `imfreehand` draws a straight line connecting the last point you drew with the first point, but you can control this behavior using the 'Closed' parameter. The following figure illustrates a freehand region with its context menu.



The following table lists the interactive features supported by `imfreehand`.

Interactive Behavior	Description
Moving the region.	Move the pointer inside the freehand region. The pointer changes to a fleur shape  . Click and hold the left mouse button to move the region.
Changing the color used to draw the region.	Move the pointer inside the freehand region. Right-click and select Set Color from the context menu.
Retrieving the current position of the freehand region.	Move the pointer inside the freehand region. Right-click and select Copy Position from the context menu. <code>imfreehand</code> copies an <i>n</i> -by-2 array of coordinates on the boundary of the ROI to the clipboard..

Methods

The `imfreehand` object supports the following methods. Type methods `imfreehand` to see a complete list of all methods.

addNewPositionCallback – Add new-position callback to ROI object

See `imroi` for information.

createMask – Create mask within image

See `imroi` for information.

delete – Delete ROI object

See `imroi` for information.

getColor – Get color used to draw ROI object

See `imroi` for information.

getPosition – Return current position of freehand region

`pos = getPosition(h)` returns the current position of the freehand region `h`. The returned position, `pos`, is an *N*-by-2 array `[X1 Y1; ... ; XN YN]`.

getPositionConstraintFcn – Return function handle to current position constraint function

See imroi for information.

removeNewPositionCallback – Remove new-position callback from ROI object.

See imroi for information.

resume – Resume execution of MATLAB command line

See imroi for information.

setClosed – Set geometry of freehand region

`setClosed(h, TF)` sets the geometry of the freehand region `h`. `TF` is a logical scalar. True means that the freehand region is closed. False means that the freehand region is open.

setColor – Set color used to draw ROI object

See imroi for information.

setConstrainedPosition – Set ROI object to new position

See imroi for information.

setPositionConstraintFcn – Set position constraint function of ROI object.

See imroi for information.

wait – Block MATLAB command line until ROI creation is finished

See imroi for information.

Tips

If you use `imfreehand` with an axes that contains an image object, and do not specify a position constraint function, users can drag the freehand region outside the extent of the image and lose the freehand region. When used with an axes created by the `plot` function, the axes limits automatically expand to accommodate the movement of the freehand region.

Examples

Interactively place a closed freehand region of interest by clicking and dragging over an image.

```
figure, imshow('pout.tif');  
h = imfreehand(gca);
```

Interactively place a freehand region by clicking and dragging. Use the `wait` method to block the MATLAB command line. Double-click on the freehand region to resume execution of the MATLAB command line.

```
figure, imshow('pout.tif');  
h = imfreehand;  
position = wait(h);
```

See Also

[imellipse](#) | [imline](#) | [impoint](#) | [impoly](#) | [imrect](#) | [iptgetapi](#) | [makeConstrainToRectFcn](#)

imfuse

Purpose

Composite of two images

Syntax

```
C = imfuse(A,B)
C = imfuse(A,B,method)
C = imfuse( __ ,Name,Value)
```

Description

`C = imfuse(A,B)` creates a composite image from two images, A and B. If A and B are different sizes, the smaller dimensions are padded with zeros so that both images are the same size before creating the composite.

`C = imfuse(A,B,method)` uses the visualization method specified by `method`.

`C = imfuse(__ ,Name,Value)` specifies additional options with one or more `Name,Value` pair arguments, using any of the previous syntaxes.

Input Arguments

A,B - Input images

grayscale images | truecolor images | binary images

Input images to be combined into a composite image, specified as any combination of grayscale, truecolor, or binary images.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32
| uint64 | logical

method - Visualization method

'falsecolor' (default) | 'blend' | 'diff' | 'montage'

Visualization method for creating the composite image, specified as one of the text strings in the following table.

Method	Description
'falsecolor'	Creates a composite RGB image showing A and B overlaid in different color bands. Gray regions in the composite image show where the two images have the same intensities. Magenta and green regions show where the intensities are different. This is the default method.
'blend'	Overlays A and B using alpha blending.
'diff'	Creates a difference image from A and B.
'montage'	Places A and B next to each other in the same image.

Example: `C = imfuse(A,B,'montage')` places A and B next to each other in the output image.

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

Example: `'Scaling','joint'` scales the intensity values of A and B together as a single data set.

Scaling - Image data scaling option

`'independent'` (default) | `'joint'` | `'none'`

Image data scaling option that controls the intensity scaling of the pixels representing the input images in the output image, specified as a comma-separated pair consisting of `'Scaling'` and `'independent'`, `'joint'`, or `'none'`. These options are described in the table below.

'independent'	Scales the intensity values of A and B independently when C is created.
'joint'	Scales the intensity values in the images jointly as if they were together in the same image. This option is useful when you want to visualize registrations of monomodal images, where one image contains fill values that are outside the dynamic range of the other image.
'none'	No additional scaling.

ColorChannels - Output color channel for each input image

[R G B] | 'red-cyan' | 'green-magenta' (default)

Output color channel for each input image, specified as the comma-separated pair consisting of 'ColorChannels' and one of the text strings in the following table. 'ColorChannels' can only be used with the 'falsecolor' visualization method.

[R G B]	A three element vector that specifies which image to assign to the red, green, and blue channels. The R, G, and B values must be 1 (for the first input image), 2 (for the second input image), and 0 (for neither image).
'red-cyan'	A shortcut for the vector [1 2 2], which is suitable for red/cyan stereo anaglyphs.
'green-magenta'	A shortcut for the vector [1 2 2], which is a high contrast option, ideal for people with many kinds of color blindness.

Output Arguments**C - Fused image**

grayscale image | truecolor image | binary image

Fused image that is a composite of the input images, returned as a grayscale, truecolor, or binary image.

Tips

- Use `imfuse` to create composite visualizations that you can save to a file. Use `imshowpair` to display composite visualizations to the screen.

Examples**Create Overlay Image of Two Images**

Load an image into the workspace. Create a copy and apply a rotation and offset.

```
A = imread('cameraman.tif');  
B = imrotate(A,5,'bicubic','crop');
```

Create blended overlay image, scaling the intensities of A and B jointly as a single data set.

```
C = imfuse(A,B,'blend','Scaling','joint');
```

Save the resulting image as a `.png` file.

```
imwrite(C,'my_blend_overlay.png');
```

See Also

`imregister` | `imshowpair` | `imtransform`

imgca

Purpose Get handle to current axes containing image

Syntax
`h = imgca`
`h = imgca(hfig)`

Description `h = imgca` returns the handle of the current axes that contains an image. The current axes can be in a regular figure window or in an Image Tool window.

If no figure contains an axes that contains an image, `imgca` creates a new axes.

`h = imgca(hfig)` returns the handle to the current axes that contains an image in the specified figure. (It need not be the current figure.)

Note `imgca` can be useful in returning the handle to the Image Tool axes. Because the Image Tool turns graphics object handle visibility off, you cannot retrieve a handle to the tool figure using `gca`.

Examples Compute the centroid of each coin, and superimpose its location on the image. View the results using `imtool` and `imgca`:

```
I = imread('coins.png');  
figure, imshow(I)
```




Original Image

```
bw = im2bw(I, graythresh(getimage));
```



```
figure, imshow(bw)
```




Binary Image

```
bw2 = imfill(bw,'holes');
s = regionprops(bw2, 'centroid');
centroids = cat(1, s.Centroid);
```

Display original image I and superimpose centroids:

```
imshow(I)
hold(imgca,'on')
plot(imgca,centroids(:,1), centroids(:,2), 'r*')
hold(imgca,'off')
```



Centroids of Coins

See also

[gca](#), [gcf](#), [imgcf](#), [imhandles](#)

imgcf

Purpose Get handle to current figure containing image

Syntax `hfig = imgcf`

Description `hfig = imgcf` returns the handle of the 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`.

Examples

```
imtool rice.png
cmap = copper(256);
set(imgcf, 'Colormap', cmap)
```

See also `gca`, `gcf`, `imgca`, `imhandles`

Purpose Open Image dialog box

Syntax [filename, user_canceled] = imgetfile

Description [filename, user_canceled] = imgetfile displays the Open Image dialog box. You can use this dialog box in imaging applications to get the name of the image file a user wants to open. The Open Image dialog box includes only files using supported image file formats (listed in `imformats`) and DICOM files. When the user selects a file and clicks **Open**, `imgetfile` returns the full path of the file in the return value `filename` and sets the `user_canceled` return value to `FALSE`. If the user clicks **Cancel**, `imgetfile` returns an empty string in `filename` and sets the `user_canceled` return value to `TRUE`.

Note The Open Image dialog box is modal; it blocks the MATLAB command line until the user responds.

See Also `imformats` | `imtool` | `uigetfile`

imgradient

Purpose Gradient magnitude and direction of an image

Syntax

```
[Gmag,Gdir] = imgradient(I)
[Gmag,Gdir] = imgradient(I,method)
[Gmag,Gdir] = imgradient(Gx,Gy)
```

Description [Gmag,Gdir] = imgradient(I) returns the gradient magnitude, Gmag, and the gradient direction, Gdir, for input image I.

- Gmag and Gdir are the same size as input image I
- Gdir contains angles in degrees within the range [-180 180]

[Gmag,Gdir] = imgradient(I,method) returns the gradient magnitude and direction using specified method.

[Gmag,Gdir] = imgradient(Gx,Gy) returns the gradient magnitude and direction using directional gradients.

Input Arguments

I - Input image

grayscale image | binary image

Input image, specified as a grayscale or binary image.

Data Types

single | double | int8 | int32 | uint8 | uint16 | uint32 | logical

method - Gradient operator

`Sobel' (default) | `Prewitt' | `CentralDifference' | `IntermediateDifference' | `Roberts'

Gradient operator, specified as one of the text strings `Sobel', `Prewitt', `CentralDifference', `IntermediateDifference', or `Roberts'.

- Central difference gradient ($dI/dx = (I(x+1) - I(x-1))/ 2$)
- Intermediate difference gradient ($dI/dx = I(x+1) - I(x)$)

Data Types

char

Gx, Gy - Directional gradients along x-axis (horizontal) and y-axis (vertical)

matrix

Directional gradients along x -axis (horizontal) and y -axis (vertical), specified as non-sparse matrices equal in size to image **I**. Typically they are obtained from function `imgradientxy` which returns them as output arguments.

Data Types

single | double | int8 | int32 | uint8 | uint16 | uint32

Output Arguments**Gmag - Gradient magnitude**

matrix

Gradient magnitude, returned as a non-sparse matrix the same size as image **I**. **Gmag** is of class `double`, unless the input image **I** to function `imgradientxy` is of class `single`, in which case it is of class `single`.

Data Types

double | single

Gdir - Gradient direction

matrix

Gradient direction, returned as a non-sparse matrix the same size as image **I**. **Gdir** contains angles in degrees within the range $[-180\ 180]$ measured counterclockwise from the positive x -axis (horizontal). **Gdir** is of class `double`, unless the input image **I** to function `gradientxy` is of class `single`, in which case it is of class `single`.

Data Types

double | single

imgradient


Examples

Display gradient magnitude and gradient direction

Read image and compute gradient magnitude and gradient direction using Prewitt's gradient operator.

```
I = imread('coins.png');  
[Gmag, Gdir] = imgradient(I,'prewitt');  
figure  
imshowpair(Gmag, Gdir, 'montage');  
title('Gradient Magnitude, Gmag (left), and Gradient Direction, Gdir (right)');  
axis off;
```

Gradient Magnitude, Gmag (left), and Gradient Direction, Gdir (right), using Prewitt's gradient operator.




Display directional gradients in addition to gradient magnitude and direction


Read in image and return directional gradients, G_x and G_y , as well as gradient magnitude and direction, G_{mag} and G_{dir} , utilizing default Sobel gradient operator.

```
I = imread('coins.png');  
[Gx, Gy] = imgradientxy(I);  
[Gmag, Gdir] = imgradient(Gx, Gy);  
figure; imshowpair(Gmag, Gdir, 'montage'); axis off;  
title('Gradient Magnitude (Gmag) and Gradient Direction (Gdir) using Sobel method');  
figure; imshowpair(Gx, Gy, 'montage'); axis off;  
title('Directional Gradients, Gx and Gy, using Sobel method');
```

Gradient Magnitude, G_{mag} (left), and Gradient Direction, G_{dir} (right)



Directional Gradients, G_x and G_y , using Sobel method



Tips

- When applying the gradient operator at the boundaries of the image, values outside the bounds of the image are assumed to equal the nearest image border value.
-

Algorithms

The algorithmic approach taken in `imgradient` for each of the listed gradient methods is to first compute directional gradients, G_x and G_y , with respect to the x-axis and y-axis. The x-axis is defined along the columns going right and the y-axis is defined along the rows going down. The gradient magnitude and direction are then computed from their orthogonal components G_x and G_y .

See Also

`imgradientxy` | `edge` | `fspecial`

Purpose

Directional gradients of an image

Syntax

```
[Gx,Gy] = imgradientxy(I)  
[Gx,Gy] = imgradientxy(I,method)
```

Description

[Gx,Gy] = imgradientxy(I) returns the directional gradients, Gx and Gy.

- Gx and Gy. are the same size as input image I
- When applying the gradient operator at the boundaries of the image, values outside the bounds of the image are assumed to equal the nearest image border value.

[Gx,Gy] = imgradientxy(I,method) returns the directional gradients using specified method

Input Arguments

I - Input image

grayscale image | binary image

Input image, specified as a grayscale or binary image.

Data Types

single | double | int8 | int32 | uint8 | uint16 | uint32 | logical

method - Gradient operator

'Sobel' (default) | 'Prewitt' | 'CentralDifference' | 'IntermediateDifference'

Gradient operator, specified as one of the text strings 'Sobel', 'Prewitt', 'CentralDifference', or 'IntermediateDifference'.

- Central difference gradient ($dI/dx = (I(x+1) - I(x-1))/ 2$)
- Intermediate difference gradient ($dI/dx = I(x+1) - I(x)$)

Data Types

char

imgradientxy

Output Arguments

Gx, Gy - Directional gradients along **x-axis (horizontal)** and **y-axis (vertical)**

matrix

Directional gradients along *x*-axis (horizontal) and *y*-axis (vertical), returned as non-sparse matrices equal in size to image *I*. The output matrices are of class `double`, unless the input image is of class `single`, in which case they are of class `single`.

Data Types

`single` | `double`


Examples

Display directional gradients

Read image and compute gradient magnitude and gradient direction using Prewitt's gradient operator.

```
I = imread('coins.png');
[Gx, Gy] = imgradientxy(I,'prewitt');
figure
imshowpair(Gx, Gy, 'montage');
title('Directional Gradients: x-direction, Gx (left), y-direction, Gy (right)');
axis off;
```

Directional Gradients: x-direction, G_x (left), y-direction, G_y (right),




Display gradient magnitude and direction in addition to directional gradients


Read in image and return directional gradients, G_x and G_y, as well as gradient magnitude and direction, G_{mag} and G_{dir}, utilizing default Sobel gradient operator.

```
I = imread('coins.png');
[Gx, Gy] = imgradientxy(I);
[Gmag, Gdir] = imgradient(Gx, Gy);
figure; imshowpair(Gmag, Gdir, 'montage'); axis off;
title('Gradient Magnitude, Gmag (left), and Gradient Direction, Gdir (right)');
figure; imshowpair(Gx, Gy, 'montage'); axis off;
title('Directional Gradients, Gx and Gy, using Sobel method')
```

Gradient Magnitude, G_{mag} (left), and Gradient Direction, G_{dir} (right), us



Directional Gradients, G_x and G_y , using Sobel me



Tips

- When applying the gradient operator at the boundaries of the image, values outside the bounds of the image are assumed to equal the nearest image border value.

Algorithms

The algorithmic approach is to compute directional gradients, G_x and G_y , with respect to the x-axis and y-axis. The x-axis is defined along the columns going right and the y-axis is defined along the rows going down.

See Also

`edge` | `fspecial` | `imgradient`

imhandles

Purpose	Get all image handles
Syntax	<code>himage = imhandles(h)</code>
Description	<p><code>himage = imhandles(h)</code> takes a graphics handle <code>h</code> as an input and returns all of the image handles whose ancestor is <code>h</code>. <code>h</code> can be an array of valid figure, axes, image, or uipanel handles.</p> <p><code>himage</code> is an array of image handles.</p> <p><code>imhandles</code> ignores colorbars in <code>h</code> and does not include its handle in <code>himage</code>.</p>
Note	<code>imhandles</code> errors if the image objects in <code>himage</code> do not have the same figure as their parent.
Examples	<p>Return the handle to the image object in the current axes.</p> <pre>figure, imshow('moon.tif'); himage = imhandles(gca)</pre> <p>Display two images in a figure and uses <code>imhandles</code> to get handles to both of the image objects in the figure.</p> <pre>subplot(1,2,1), imshow('autumn.tif'); subplot(1,2,2), imshow('glass.png'); himages = imhandles(gcf)</pre>
See Also	<code>imgca</code> <code>imgcf</code>

Purpose Display histogram of image data

Syntax

```
imhist(I)
imhist(I, n)
imhist(X, map)
[counts,x] = imhist(...)
```

Description `imhist(I)` displays a histogram for the 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 two 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, `imhist` returns the histogram counts for each colormap entry; the length of `counts` is the same as the length of the colormap.

Note The maximum value on the y-axis may be automatically reduced, so outlier spikes do not dominate. To show the full range of y-axis values, call `imhist` with the following syntax:

```
[counts,x] = imhist(...)
```

Then call `stem`:

```
stem(x,counts)
```

Tips

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

$$\frac{A(p-1.5)}{(n-1)} \leq x < \frac{A(p-0.5)}{(n-1)},$$

where x is the intensity value. For `int16` intensity images, the p th bin is the half-open interval

$$\frac{A(p-1.5)}{(n-1)} - 32768 \leq x < \frac{A(p-0.5)}{(n-1)} - 32768,$$


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`.

Class Support

An input intensity image can be of class `uint8`, `int8`, `uint16`, `int16`, `uint32`, `int32`, `single`, `double`, or `logical`. An input indexed image can be of class `uint8`, `uint16`, `single`, `double`, or `logical`.

Examples

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

See Also `histeq` | `hist`

imhistmatch

Purpose Adjust histogram of image to match N-bin histogram of reference image

Syntax

```
B = imhistmatch(A,Ref)
B = imhistmatch(A,Ref,N)
[B,hgram] = imhistmatch( ___ )
```

Description `B = imhistmatch(A,Ref)` image `A` is transformed so that the histogram of the returned image `B` approximately matches the histogram of reference image `Ref` built with 64 (default value) equally spaced histogram bins. The returned image `B` will have no more than 64 discrete levels.

- Images `A` and `Ref` can be any of the permissible data types.
- If both `A` and `Ref` are truecolor RGB images, then each color channel of `A` is matched independently to the corresponding color channel of `Ref`.
- If `A` is a truecolor RGB image and `Ref` is a grayscale image, then each channel of `A` is matched against the single histogram derived from `Ref`.
- If `A` is a grayscale image, then `Ref` must also be a grayscale image.
- Images `A` and `Ref` need not be equal in size.

`B = imhistmatch(A,Ref,N)` image `A` is transformed using a histogram derived from image `Ref` consisting of `N` equally spaced bins within the appropriate range for the given image data type. The returned image `B` will have no more than `N` discrete levels. The default value for `N` is 64.

- If the data type of the image is either single or double, then the histogram range is `[0, 1]`.
- If the data type of the image is `uint8`, then the histogram range is `[0, 255]`
- If the data type of the image is `uint16`, then the histogram range is `[0, 65535]`

- If the data type of the image is int16, then the histogram range is [-32768, 32767]

[B,hgram] = imhistmatch(___) returns hgram, a vector of length N, containing the histogram counts derived from reference image Ref, using either of the previous two syntaxes.

Input Arguments

A - Input image

truecolor image | grayscale image

Input image to be transformed, specified as a truecolor or grayscale image. The returned image will take the data type class of the input image.

Data Types

single | double | int16 | uint8 | uint16

Ref - Reference image whose histogram is the reference histogram

truecolor image | grayscale image

Reference image whose histogram is the reference histogram, specified as truecolor or grayscale image. The reference image provides the equally spaced N bin reference histogram which output image B is trying to match.

Data Types

single | double | int16 | uint8 | uint16

N - Number of equally spaced bins in reference histogram

64 (default) | scalar

Number of equally spaced bins in reference histogram, specified as a scalar value. In addition to specifying the number of equally spaced bins in the histogram for image Ref, N also represents the upper limit of the number of discrete data levels present in output image B.

Data Types

double

Output Arguments

B - Output image

truecolor RGB image | grayscale image

Output image, returned as a truecolor or grayscale image. The output image is derived from image A whose histogram is an approximate match to the histogram of input image Ref built with N equally spaced bins. Image B is of the same size and data type as input image A. Input argument N represents the upper limit of the number of discrete levels contained in image B.

Data Types

single | double | int16 | uint8 | uint16

hgram - Histogram counts derived from reference image Ref

vector | matrix

Histogram counts derived from reference image Ref, specified as a vector or matrix. When Ref is a truecolor image, hgram is a 3xN matrix. When Ref is a grayscale image, hgram is a 1xN vector.

Data Types

double

Examples

Histogram matching of aerial images

These aerial images, taken at different times, represent overlapping views of the same terrain in Concord, Massachusetts. This example demonstrates that input images A and Ref can be of different sizes and image types.

Load both images.

```
A = imread('concordaerial.png');  
Ref = imread('concordorthophoto.png');
```

Querying for the size of the images reveals that both are different in size and type; image A is a truecolor RGB image, while image Ref is a grayscale image. Both images are of data type uint8.

```
size(A)
```

```
ans =  
      2036      3060      3
```

```
size(Ref)
```

```
ans =  
      2215      2956
```

Generate the histogram matched output image; each channel of **A** is matched against the single histogram of **Ref** built with 64 (default value) equally spaced bins. Output image **B** takes on the characteristics of image **A**; it is an RGB image whose size and data type is the same as image **A**. The number of distinct levels present in each RGB channel of image **B** is determined by the number of bins in the single aim histogram built from grayscale image **Ref** which in this case is 64.

```
B = imhistmatch(A,Ref);
```

RGB image with color cast



Reference grayscale image



Multiple N values applied to RGB Images


In this example, you will see the effect on output image B of varying N, the number of equally spaced bins in the aim histogram of image Ref, from its default value 64 to the maximum value of 256 for uint8 pixel data.

The following images were taken with a digital camera and represent two different exposures of the same scene.

```
A = imread('office_2.jpg'); % Dark Image
Ref = imread('office_4.jpg'); % Reference image
```

Image A, being the darker image, has a preponderance of its pixels in the lower bins. The reference image, Ref, is a properly exposed image and fully populates all of the available bins values in all three RGB channels: as shown in the table below, all three channels have 256 unique levels for 8-bit pixel values.


A: Dark Image








Ref: Reference Image



RGB Histograms: 256 bins



RGB Histograms: 256 bins



imhistmatch

Unique 8-bit level values in each RGB Channel		
	Image A	Image Ref
Red Channel	205	256
Green Channel	193	256
Blue Channel	224	256

Output image B is generated using three different values of N: 64, 128 and 256. The objective of function imhistmatch is to transform image A such that the histogram of output image B is a match to the histogram of Ref built with N equally spaced bins. As a result, N represents the upper limit of the number of discrete data levels present in image B.

```
[B64, hgram] = imhistmatch(A, Ref, 64);  
[B128, hgram] = imhistmatch(A, Ref, 128);  
[B256, hgram] = imhistmatch(A, Ref, 256);
```



Output Image B64








Output Image B128



RGB Histograms: 64 bins



RGB Histograms: 128 bins



Unique 8-bit level values in each RGB Channel for N = [64 128 256]			
	output Image B64	output image B128	output image B256
Red Channel	57	101	134
Green Channel	60	104	135
Blue Channel	58	104	136

Note that as N increases, the number of levels in each RGB channel of output image B also increases.

Histogram matching a 16-bit grayscale MRI image

In this example, a 16 bit grayscale MRI image is loaded, darkened and histogram matched at two values of N.

A 16-bit DICOM image of a knee imaged via MRI is loaded into the workspace.

```
K = dicomread('knee1.dcm'); % read in original image
K = dicomread('knee1.dcm'); % read in original 16-bit image
LevelsK = unique(K(:)); % determine number of unique code values
disp(['image K: # levels: ' num2str(length(LevelsK))]);
disp(['max level = ' num2str( max(LevelsK) )]);
disp(['min level = ' num2str( min(LevelsK) )]);
```

```
image K: # levels = 448
max level = 473
min level = 0
```

Since it appears that all 448 discrete values are at low code values (darker), the image data will be scaled to span the entire 16-bit range of [0 65535]

```
% Scale it to full 16-bit range
```

```
Kdouble = double(K); % cast uint16 to double
kmult = 65535/(max(max(Kdouble(:)))); % full range multiplier
Ref = uint16(kmult*Kdouble); % full range 16-bit reference image
```


The reference image is darkened to yield the input A image argument to imhistmatch

```
% build concave bow-shaped curve for darkening Reference image
ramp = [0:65535]/65535;
ppconcave = spline([0 .1 .50 .72 .87 1],[0 .025 .25 .5 .75 1]);
Ybuf = ppval( ppconcave, ramp);
Lut16bit = uint16( round( 65535*Ybuf ) );
```


```
% pass image Ref through LUT to darken image
A = intlut(Ref,Lut16bit);
```

The two images contain the same number of discrete code values, but differ in overall brightness.

A: Darkened Image



Ref: Reference Image




imhistmatch


Next, the histogram matched output image, **B**, is generated at two values of **N**. The first is the default value of 64, the second is the number of values present in image **A** of 448.

```
B16bit64 = imhistmatch(A(:,:,1),Ref(:,:,1)); % default # bins: N = 64  
  
N = length(LevelsA); % number of unique 16-bit code values in image  
B16bitUniq = imhistmatch(A(:,:,1),Ref(:,:,1),N);
```

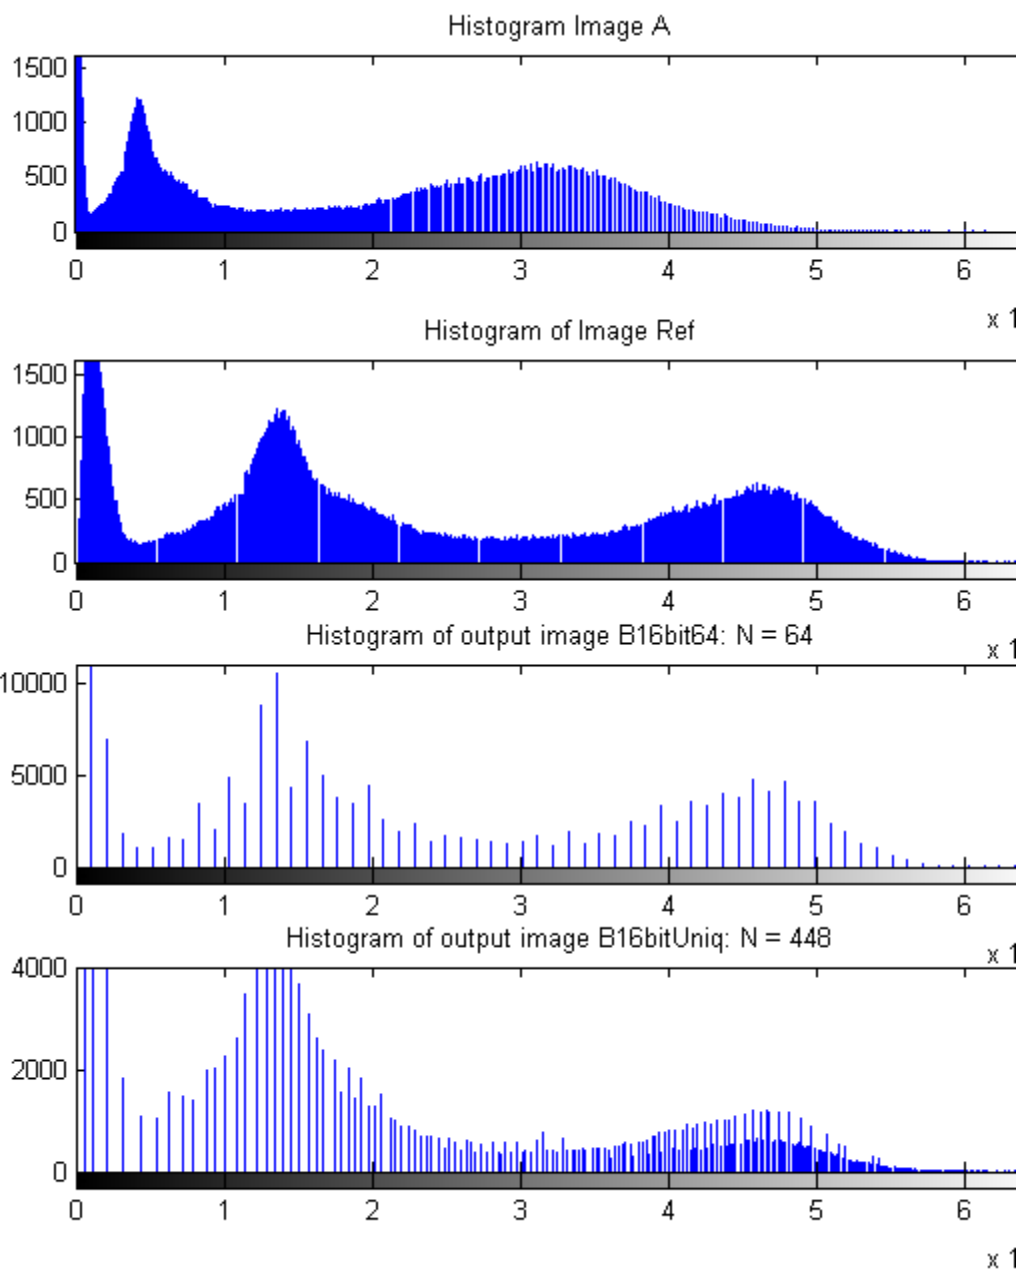
B16bit64: N = 64



B16bitUniq: N = 448



The following figure shows the 16 bit histograms of all four images; the y-axis scaling is the same for plots.




Unique 16-bit code values in output B Images		
	Levels	N
B16bit64	63	64
B16bitUniq	222	448

N also represents the upper limit of discrete levels in the output image which is shown above; the number of levels increases from 63 to 222 when the number of histogram bins increases from 64 to 448. But note, in the above histogram plots, there are rapid fluctuations in adjacent histogram bins for the B image containing 222 levels, especially in the upper portion of the histogram range. By comparison, the 63 level B histogram has a relatively smooth and continuous progression of peaks in this region.

Algorithms

The objective of `imhistmatch` is to transform image A such the the histogram of image B matches the histogram derived from image Ref. It consists of N equally spaced bins which span the full range of the image data type. A consequence of matching histograms in this way is that N also represents the upper limit of the number of discrete data levels present in image B.

An important behavioral aspect of this algorithm to note is that as N increases in value, the degree of rapid fluctuations between adjacent populated peaks in the histogram of image B tends to increase. This can be seen in the following histogram plots taken from the 16-bit grayscale MRI example.



An optimal value for N represents a trade-off between more output levels (larger values of N) while minimizing peak fluctuations in the histogram (smaller values of N).

See Also

[histeq](#) | [imadjust](#) | [imhist](#)

imhmax

Purpose H-maxima transform

Syntax
`I2 = imhmax(I,h)`
`I2 = imhmax(I,h,conn)`

Description `I2 = 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 a constant intensity value, and whose external boundary pixels all have a lower value.

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 = imhmax(I,h,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.

Examples

```
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.
```

References

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

See Also

conndef | imextendedmax | imhmin | imreconstruct | imregionalmax

imhmin

Purpose H-minima transform

Syntax `I2 = imhmin(I,h)`
`I2 = imhmin(I,h,conn)`

Description `I2 = imhmin(I,h)` suppresses all minima in `I` whose depth is less than `h`. `I` is a grayscale image and `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.

Examples

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

imhmin

10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10

References

[1] Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer-Verlag, 1999, pp. 170-171.

See Also

`conndef` | `imextendedmin` | `imhmax` | `imreconstruct` | `imregionalmin`

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`.

Examples

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);
```

imimposemin

```
title('Regional Minima in Original Image');  
BW2 = imregionalmin(K);  
figure, imshow(BW2);  
title('Regional Minima After Processing');
```



Regional Minima in Original Image



Regional Minima After Processing

Algorithms

`imimposemin` uses a technique based on morphological reconstruction.

See Also

`conndef` | `imreconstruct` | `imregionalmin`

Purpose

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 Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The Intel IPP library is activated only in the following cases:

```
Z = imlincomb( 1.0, A1, 1.0, A2)
Z = imlincomb( 1.0, A1, -1.0, A2)
Z = imlincomb(-1.0, A1, 1.0, A2)
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`

imline

Purpose Create draggable, resizable line

Syntax

```
h = imline
h = imline(hparent)
h = imline(hparent, position)
h = imline(hparent, x, y)
H = imline(..., param1, val1,...)
```

Description `h = imline` begins interactive placement of a line on the current axes. The function returns `h`, a handle to an `imline` object. The line has a context menu associated with it that controls aspects of its appearance and behavior—see “Interactive Behavior” on page 3-439. Right-click on the line to access this context menu.

`h = imline(hparent)` begins interactive placement of a line on the object specified by `hparent`. `hparent` specifies the HG parent of the line graphics, which is typically an axes but can also be any other object that can be the parent of an `hggroup`


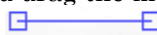
`h = imline(hparent, position)` creates a draggable, resizable line on the object specified by `hparent`. `position` is a 2-by-2 array that specifies the initial endpoint positions of the line in the form `[X1 Y1; X2 Y2]`.



`h = imline(hparent, x, y)` creates a line on the object specified by `hparent`. `x` and `y` are two-element vectors that specify the initial endpoint positions of the line in the form `x = [X1 X2]`, `y = [Y1 Y2]`.

`H = imline(..., param1, val1,...)` creates a draggable, resizable line, specifying parameters and corresponding values that control the behavior of the line. The following tables lists the parameter available. Parameter names can be abbreviated, and case does not matter.

Parameter	Description
'PositionConstraintFcn	Function handle fcn that is called whenever the object is dragged using the mouse. You can use this function to control where the line can be dragged. See the help for the setPositionConstraintFcn method for information about valid function handles.

Interactive Behavior

When you call `imline` with an interactive syntax, the pointer changes to a cross hairs  when over the image. Click and drag the mouse to specify the position and length of the line, such as . The line supports a context menu that you can use to control aspects of its appearance and behavior. For more information about these interactive features, see the following table.

Interactive Behavior	Description
Moving the line.	Move the pointer over the line. The pointer changes to a fleur shape  . Click and drag the mouse to move the line.
Moving the endpoints of the line.	Move the pointer over either end of the line. The pointer changes to the pointing finger,  . Click and drag the mouse to resize the line.

Interactive Behavior	Description
Changing the color used to display the line.	Move the pointer over the line. Right-click and select Set Color from the context menu.
Retrieving the coordinates of the endpoints of the line.	Move the pointer over the line. Right-click and select Copy Position from the context menu. <code>imline</code> copies a 2-by-2 array to the clipboard specifying the coordinates of the endpoints of the line in the form [X1 Y1; X2 Y2].

Methods

Each `imline` object supports a number of methods. Type methods `imline` to see a list of the methods.

addNewPositionCallback – Add new-position callback to ROI object

See `imroi` for information.

delete – Delete ROI object

See `imroi` for information.

getColor – Get color used to draw ROI object.

See `imroi` for information.

getPosition – Return current position of line

Returns the endpoint positions of the line.

```
pos = api.getPosition()
```

`pos` is a 2-by-2 array [X1 Y1; X2 Y2].

getPositionConstraintFcn – Return function handle to current position constraint function

See `imroi` for information.

removeNewPositionCallback – Remove new-position callback from ROI object.

See `imroi` for information.

resume – Resume execution of MATLAB command line

See `imroi` for information.

setColor – Set color used to draw ROI object

See `imroi` for information.

setConstrainedPosition – Set ROI object to new position

See `imroi` for information.

setPosition – Set line to new position

`setPosition(h,pos)` sets the line `h` to a new position. The new position, `pos`, has the form, `[X1 Y1; X2 Y2]`.

`setPosition(h,x,y)` sets the line `h` to a new position. `x` and `y` specify the endpoint positions of the line in the form `x = [x1 x2]`, `y = [y1 y2]`.

setPositionConstraintFcn – Set position constraint function of ROI object.

See `imroi` for information.

wait – Block MATLAB command line until ROI creation is finished

See `imroi` for information.

Tips

If you use `imline` with an axes that contains an image object, and do not specify a position constraint function, users can drag the line outside the extent of the image and lose the line. When used with an axes created by the `plot` function, the axis limits automatically expand to accommodate the movement of the line.

Examples**Example 1**

Use a custom color for displaying the line. Use `addNewPositionCallback` method. Move the line, note that the 2-by-2 position vector of the line is displayed in the title above the image. Explore the context menu of the line by right clicking on the line.

```
figure, imshow('pout.tif');
h = imline(gca,[10 100], [100 100]);
setColor(h,[0 1 0]);
id = addNewPositionCallback(h,@(pos) title(mat2str(pos,3)));
```

imline

```
% After observing the callback behavior, remove the callback.  
% using the removeNewPositionCallback API function.  
removeNewPositionCallback(h,id);
```

Example 2

Interactively place a line by clicking and dragging. Use `wait` to block the MATLAB command line. Double-click on the line to resume execution of the MATLAB command line

```
figure, imshow('pout.tif');  
h = imline;  
position = wait(h);
```

See Also

[imellipse](#) | [imfreehand](#) | [impoint](#) | [impoly](#) | [imrect](#) | [imroi](#) | [makeConstrainToRectFcn](#)

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.

Function	Description
<code>setMagnification</code>	<p>Sets the magnification in units of screen pixels per image pixel.</p> <p><code>setMagnification(new_mag)</code></p> <p>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>setMagnification(2)</code>, the magnification box will show the string '200%'.</p>

Examples Add 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`

Purpose

Make movie from multiframe image

Note `immovie(D,size)` is an obsolete syntax and is no longer supported. Use `immovie(X,map)` instead.

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`. For details about the movie structure array, see the reference page for `getframe`. To play the movie, call `implay`.

`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.

Tips

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.

Examples

```
load mri
mov = immovie(D,map);
implay(mov)
```

See Also

`avifile` | `getframe` | `montage` | `movie` | `movie2avi`

immultiply

Purpose Multiply two images or multiply image by 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.

Note On Intel architecture processors, `immultiply` can take advantage of the Intel Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The Intel IPP library is activated only if arrays X, Y, and Z are of class `logical`, `uint8`, or `single`, and are of the same class.

Examples 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');  
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
```

imnoise

Purpose Add noise to image

Syntax

```
J = imnoise(I,type)
J = imnoise(I,type,parameters)
J = imnoise(I,'gaussian',m,v)
J = imnoise(I,'localvar',V)
J = imnoise(I,'localvar',image_intensity,var)
J = imnoise(I,'poisson')
J = imnoise(I,'salt & pepper',d)
J = imnoise(I,'speckle',v)
```

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

`J = imnoise(I,type,parameters)` Depending on `type`, you can specify additional parameters to `imnoise`. 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. If `I` is double precision, then input pixel values are interpreted as means of Poisson distributions scaled up by $1e12$. For example, if an input pixel has the value $5.5e-12$, then the corresponding output pixel will be generated from a Poisson distribution with mean of 5.5 and then scaled back down by $1e12$. If `I` is single precision, the scale factor used is $1e6$. If `I` is `uint8` or `uint16`, then input pixel values are used directly without scaling. For example, if a pixel in a `uint8` input has the value 10, then the corresponding output pixel will be generated from a Poisson distribution with mean 10.

`J = imnoise(I, 'salt & pepper', d)` adds salt and pepper noise to the image `I`, where `d` is the noise density. This affects approximately $d * \text{numel}(I)$ pixels. The default for `d` is 0.05.

`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.

imnoise

Class Support

For most noise types, I can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. For Poisson noise, `int16` is not allowed. 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.

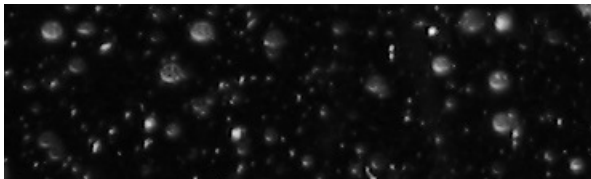
Examples

```
I = imread('eight.tif');  
J = imnoise(I, 'salt & pepper', 0.02);  
figure, imshow(I)  
figure, imshow(J)
```



See Also


[rand](#) | [randn](#)

Purpose	Morphologically open image
Syntax	<pre>IM2 = imopen(IM,SE) IM2 = imopen(IM,NHOOD)</pre>
Description	<p><code>IM2 = imopen(IM,SE)</code> performs morphological opening on the grayscale or binary image <code>IM</code> with the structuring element <code>SE</code>. The argument <code>SE</code> must be a single structuring element object, as opposed to an array of objects. The morphological open operation is an erosion followed by a dilation, using the same structuring element for both operations.</p> <p><code>IM2 = imopen(IM,NHOOD)</code> performs opening with the structuring element <code>strel(NHOOD)</code>, where <code>NHOOD</code> is an array of 0's and 1's that specifies the structuring element neighborhood.</p>
Class Support	<code>IM</code> can be any numeric or logical class and any dimension, and must be nonsparse. If <code>IM</code> is logical, then <code>SE</code> must be flat. <code>IM2</code> has the same class as <code>IM</code> .
Examples	<p>Remove the smaller objects in an image.</p> <ol style="list-style-type: none">1 Read the image into the MATLAB workspace and display it. <pre>I = imread('snowflakes.png'); imshow(I)</pre>2 Create a disk-shaped structuring element with a radius of 5 pixels. <pre>se = strel('disk',5);</pre>

imopen

- 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](#)

Purpose Overview tool for image displayed in scroll panel

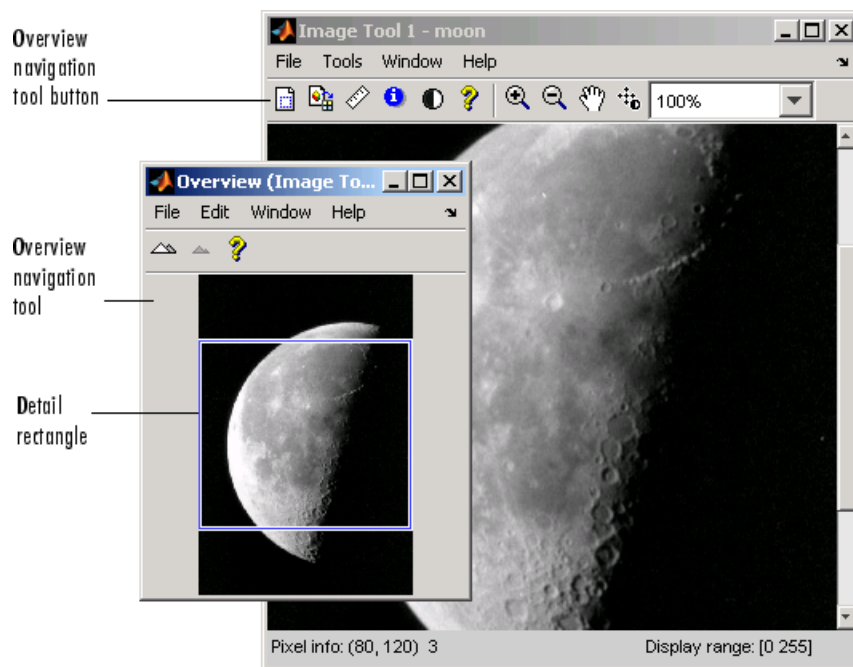
Syntax
`imoverview(himage)`
`hfig = imoverview(...)`

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.

The following figure shows the Image Tool with the Overview tool.

imoverview



`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`.

Examples

Create a figure, disabling the toolbar and menubar, because the toolbox navigation tools are not compatible with the standard MATLAB zoom and pan tools. Then create a scroll panel in the figure and use 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](#)

imoverviewpanel

Purpose Overview tool panel for image displayed in scroll panel

Syntax `hpanel = imoverviewpanel(hparent,himage)`

Description `hpanel = imoverviewpanel(hparent,himage)` creates an Overview tool panel associated with the image specified by the handle `himage`, called the target image. `himage` must be contained in a scroll panel created by `imscrollpanel`. `hparent` is a handle to the figure or `uipanel` object that will contain the Overview tool panel. `imoverviewpanel` returns `hpanel`, a handle to the Overview tool `uipanel` object.

The Overview tool is a navigation aid for images displayed in a scroll panel. `imoverviewpanel` 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.

Note To create an Overview tool in a separate figure, use `imoverview`. When created using `imoverview`, the Overview tool includes zoom-in and zoom-out buttons.

Examples Create an Overview tool that is embedded in the same figure that contains the target image.

```
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])
```

See Also `imoverview` | `imscrollpanel`

Purpose

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 syntax 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 syntax.

```
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 syntax.

```
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`.

Tips

`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.

Examples

```
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`

impixelinfo

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 pointer is positioned over. The tool can display pixel information for all the images in a figure.

The Pixel Information tool is a uipanel object, positioned in the lower-left corner of the figure. The tool contains the text string `Pixel info:` followed by the pixel information. Before you move the pointer over the image, the tool contains the default pixel information text string `(X,Y) Pixel Value`. Once you move the pointer over the image, the information displayed varies by image type, as shown in the following table. If you move the pointer off the image, the pixel information tool displays the default pixel information string for that image type.

Image Type	Pixel Information	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]
Floating point image with <code>CDataMapping</code> property set to <code>direct</code>	(X,Y) value <index> [R G B]	(19,10) 82 <4> [15 255 10]

For example, for grayscale (intensity) images, the pixel information tool displays the x and y coordinates of the pixel and its value, as shown in the following figure.

X and Y coordinates	Pixel Value
Pixel info: (418, 261) 143	

If you want to display the pixel information without the “Pixel Info” label, use the `impixelinfoval` function.

`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 pointer is positioned over a pixel. In the context menu displayed, choose **Copy pixel info**.

Examples

Display an image and add 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',[5 1 300 20]);
```

Use the Pixel Information tool in a figure containing multiple images of different types.

imshowinfo

```
figure
subplot(1,2,1), imshow('liftingbody.png');
subplot(1,2,2), imshow('autumn.tif');
imshowinfo;
```

See Also

[imshowinfoval](#) | [imtool](#)

Purpose Pixel Information tool without 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 pointer 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 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 pointer is positioned over a pixel. In the context menu displayed, choose **Copy pixel info**.

Examples

Add a Pixel Information 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-MONO2-16-ankle.dcm');
h = imshow(ankle,[]);
```

impixelinfoval

```
hText = impixelinfoval(gcf,h);  
set(hText,'FontWeight','bold')  
set(hText,'FontSize',10)
```

See also

`impixelinfo`

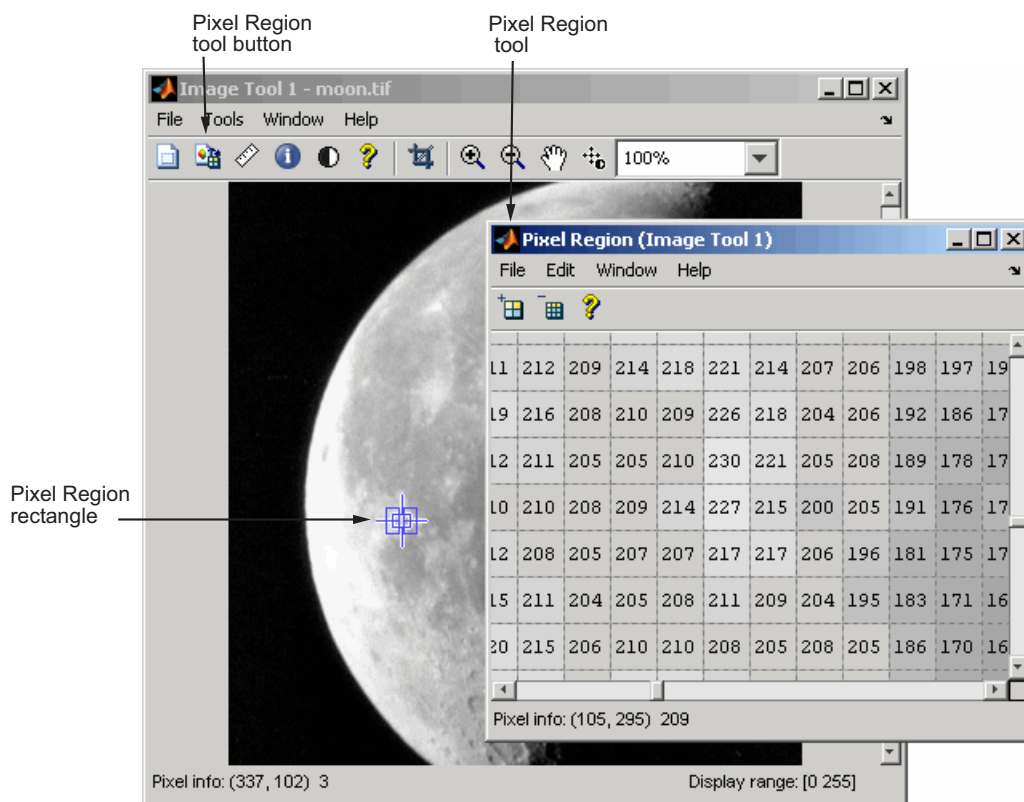
Purpose Pixel Region tool

Syntax

```
impixelregion  
impixelregion(h)  
hfig = impixelregion(...)
```

Description `impixelregion` creates a Pixel Region 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, as shown in the following figure.

impixelregion



The Pixel Region rectangle defines the area of the target image that is displayed in the Pixel Region tool. You can move this rectangle over the target image using the mouse to view different regions. To get a closer view of the pixels displayed in the tool, use the zoom buttons on the Pixel Region tool toolbar or change the size of the Pixel Region rectangle using the mouse. You can also resize the Pixel Region tool itself to view more or fewer pixels. If the size of the pixels allows, the tool superimposes the numeric value of the pixel over each pixel.

To get the current position of the Pixel Region rectangle, right-click on the rectangle and select **Copy Position** from the context menu. The Pixel Region tool copies a four-element position vector to the clipboard.

To change the color of the Pixel Region rectangle, right-click and select **Set Color**.

`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`.

Examples

Display an image and then create a Pixel Region tool associated with the image.

```
imshow peppers.png  
impixelregion
```

See Also

`impixelinfo` | `impixelregionpanel` | `imtool`

impixelregionpanel

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. If the size of the pixels allows, 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. See `impixelregion` for more information.

Note To create a Pixel Region tool in a separate figure window, use `impixelregion`.

Examples

```
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` | `imrect` | `imscrollpanel`

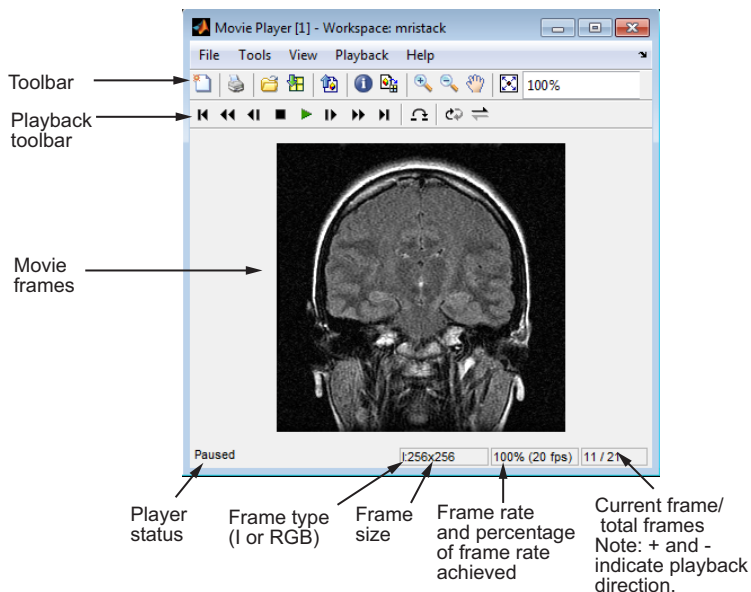
Purpose Play movies, videos, or image sequences

Syntax

```
implay
implay(filename)
implay(I)
implay(..., FPS)
```

Description `implay` opens a Movie Player for showing MATLAB movies, videos, or image sequences (also called image stacks). Use the `implay` File menu to select the movie or image sequence that you want to play. Use `implay` toolbar buttons or menu options to play the movie, jump to a specific frame in the sequence, change the frame rate of the display, or perform other exploration activities. You can open multiple `implay` movie players to view different movies simultaneously.

The following figure shows the Movie Player containing an image sequence.



imshow

`imshow(filename)` opens the `imshow` movie player, displaying the content of the file specified by `filename`. The file can be an Audio Video Interleaved (AVI) file. `imshow` reads one frame at a time, conserving memory during playback. `imshow` does not play audio tracks.

`imshow(I)` opens the `imshow` movie player, displaying the first frame in the multiframe image array specified by `I`. `I` can be a MATLAB movie structure, or a sequence of binary, grayscale, or truecolor images. A binary or grayscale image sequence can be an `M-by-N-by-1-by-K` array or an `M-by-N-by-K` array. A truecolor image sequence must be an `M-by-N-by-3-by-K` array.

`imshow(..., FPS)` specifies the rate at which you want to view the movie or image sequence. The frame rate is specified as frames-per-second. If omitted, `imshow` uses the frame rate specified in the file or the default value 20.

Class Support


`I` can be numeric but `uint8` is preferred. The actual data type used to display pixels may differ from the source data type.

Tips

- You can access `imshow` through the Plot Selector workspace tool,

which is represented by this icon:  Select data to plot

In your workspace, select the data you want to display. The Plot

Selector icon updates like this:  `imshow(mristack)`

. To open `imshow` with your chosen data, either click on the Plot Selector icon or select `imshow` from the Plot Selector drop-down menu.

Examples

Animate a sequence of images.

```
load cellsequence
imshow(cellsequence,10);
```

Visually explore a stack of MRI images.

```
load mristack
imshow(mristack);
```

Play an AVI file.

```
imshow('rhinos.avi');
```

impoint

Purpose Create draggable point

Syntax

```
h = impoint
h = impoint(hparent)
p = impoint(hparent, position)
p = impoint(hparent, x, y)
p = impoint(..., param, val)
```

Description `h = impoint` begins interactive placement of a draggable point on the current axes. The function returns `h`, a handle to an `impoint` object. The point has a context menu associated with it that controls aspects of its appearance and behavior—see “Interactive Behavior” on page 3-473. Right-click on the point to access this context menu.

`h = impoint(hparent)` begins interactive placement of a point on the object specified by `hparent`. `hparent` specifies the HG parent of the point graphics, which is typically an axes but can also be any other object that can be the parent of an `hggroup`.



`p = impoint(hparent, position)` creates a draggable point on the object specified by `hparent`. `position` is a 1-by-2 array of the form `[x y]` that specifies the initial position of the point.


`p = impoint(hparent, x, y)` creates a draggable point on the object specified by `hparent`. `x` and `y` are both scalars that together specify the initial position of the point.

`p = impoint(..., param, val)` creates a draggable point, specifying parameters and corresponding values that control the behavior of the point. The following table lists the parameter available. Parameter names can be abbreviated, and case does not matter.

Parameter	Description
'PositionConstraintFcn	Function handle fcn that is called whenever the point is dragged using the mouse. You can use this function to control where the point can be dragged. See the help for the setPositionConstraintFcn method for information about valid function handles.

Interactive Behavior

When you call `impoint` with an interactive syntax, the pointer changes to a cross hairs  when over the image. Click and drag the mouse to specify the position of the point, such as . The following table describes the interactive behavior of the point, including the right-click context menu options.

Interactive Behavior	Description
Moving the point.	Move the mouse pointer over the point. The mouse pointer changes to a fleur shape  . Click and drag the mouse to move the point.
Changing the color used to display the point.	Move the mouse pointer over the point. Right-click and select Set Color from the context menu and specify the color you want to use.
Retrieving the coordinates of the point.	Move the mouse pointer over the point. Right-click and select Copy Position from the context menu to copy a 1-by-2 array to the clipboard specifying the coordinates of the point [X Y].

Methods

The following lists the methods supported by the `impoint` object. Type `methods impoint` to see a complete list of all the methods.

addNewPositionCallback – Add new-position callback to ROI object

See `imroi` for information.

delete – Delete ROI object

See `imroi` for information.

getColor – Get color used to draw ROI object

See `imroi` for information.

getPosition – Return current position of point

`pos = getPosition(h)` returns the current position of the point `h`. The returned position, `pos`, is a two-element vector `[x y]`.

getPositionConstraintFcn – Return function handle to current position constraint function

See `imroi` for information.

removeNewPositionCallback – Remove new-position callback from ROI object.

See `imroi` for information.

resume – Resume execution of MATLAB command line

See `imroi` for information.

setColor – Set color used to draw ROI object

See `imroi` for information.

setConstrainedPosition – Set ROI object to new position

See `imroi` for information.

setPosition – Set point to new position

`setPosition(h,pos)` sets the point `h` to a new position. The new position, `pos`, has the form, `[x y]`.

`setPosition(h,new_x,new_y)` sets the point `h` to a new position. `new_x` and `new_y` are both scalars that together specify the position of the point.

setPositionConstraintFcn — Set position constraint function of ROI object.

See imroi for information.

setString — Set text label for point

setString(h,s) sets a text label for the point h. The string, s, is placed to the lower right of the point.

wait — Block MATLAB command line until ROI creation is finished

See imroi for information.

Tips

If you use impoint with an axes that contains an image object, and do not specify a drag constraint function, users can drag the point outside the extent of the image and lose the point. When used with an axes created by the plot function, the axes limits automatically expand to accommodate the movement of the point.

Examples**Example 1**

Use impoint methods to set custom color, set a label, enforce a boundary constraint, and update position in title as point moves.

```
figure, imshow rice.png
h = impoint(gca,100,200);
% Update position in title using newPositionCallback
addNewPositionCallback(h,@(h) title(sprintf('%1.0f,%1.0f',h(1),h(2))));
% Construct boundary constraint function
fcn = makeConstrainToRectFcn('impoint',get(gca,'XLim'),get(gca,'YLim'));
% Enforce boundary constraint function using setPositionConstraintFcn
setPositionConstraintFcn(h,fcn);
setColor(h,'r');
setString(h,'Point label');
```

Example 2

Interactively place a point. Use wait to block the MATLAB command line. Double-click on the point to resume execution of the MATLAB command line.

```
figure, imshow('pout.tif');
```

impoint

```
p = impoint(gca,[]);  
p = wait(p);
```

See Also

[imellipse](#) | [imfreehand](#) | [imline](#) | [impoly](#) | [imrect](#) | [imroi](#) | [makeConstrainToRectFcn](#)

Purpose

Create draggable, resizable polygon

Syntax

```
h = impoly
h = impoly(hparent)
h = impoly(hparent, position)
h = impoly(..., param1, val1, ...)
```

Description

`h = impoly` begins interactive placement of a polygon on the current axes. The function returns `h`, a handle to an `impoly` object. The polygon has a context menu associated with it that controls aspects of its appearance and behavior—see “Interactive Behavior” on page 3-478. Right-click on the polygon to access this context menu.


`h = impoly(hparent)` begins interactive placement of a polygon on the object specified by `hparent`. `hparent` specifies the HG parent of the polygon graphics, which is typically an axes but can also be any other object that can be the parent of an `hggroup`.


`h = impoly(hparent, position)` creates a draggable, resizable polygon on the object specified by `hparent`. `position` is an n -by-2 array that specifies the initial position of the vertices of the polygon. `position` has the form `[X1,Y1;...;XN,YN]`.

`h = impoly(..., param1, val1, ...)` creates a draggable, resizable polygon, specifying parameters and corresponding values that control the behavior of the polygon. The following table lists available parameters. Parameter names can be abbreviated, and case does not matter.

Parameter	Description
'Closed'	Scalar logical that controls whether the polygon is closed. When set to true (the default), <code>impoly</code> creates a closed polygon, that is, it draws a straight line between the last vertex specified and the first vertex specified to create a closed region. When <code>Closed</code> is false, <code>impoly</code> does not connect the last vertex with the first vertex, creating an open polygon (or polyline).
'PositionConstraintFcn'	Function handle <code>fcn</code> that is called whenever the object is dragged using the mouse. You can use this function to control where the line can be dragged. See the help for the <code>setPositionConstraintFcn</code> method for information about valid function handles.





Interactive Behavior

When you call `impoly` with an interactive syntax, the pointer changes to a cross hairs  when over the image. Click and drag the mouse to define the vertices of the polygon and adjust the size, shape, and position of the polygon. The polygon also supports a context menu that you can use to control aspects of its appearance and behavior. The choices in the context menu vary whether you position the pointer on an edge of the polygon (or anywhere inside the region) or on one of the vertices. The following figure shows a polygon being created.



The following table lists the interactive behaviors supported by impoly.

Interactive Behavior	Description
Closing the polygon.	<p>Use any of the following mechanisms:</p> <ul style="list-style-type: none"> • Move the pointer over the initial vertex of the polygon that you selected. The pointer changes to a circle ○. Click either mouse button. • Double-click the left mouse button. This action creates a vertex at the point under the mouse and draws a straight line connecting this vertex with the initial vertex. • Click the right mouse button. This action draws a line connecting the last vertex

Interactive Behavior	Description
	selected with the initial vertex; it does not create a new vertex.
Adding a new vertex.	Move the pointer over an edge of the polygon and press the A key. The shape of the pointer changes  . Click the left mouse button to create a new vertex at that position on the line.
Moving a vertex. (Reshaping the polygon.)	Move the pointer over a vertex. The pointer changes to a circle  . Click and drag the vertex to its new position.
Deleting a vertex.	Move the pointer over a vertex. The shape changes to a circle  . Right-click and select Delete Vertex from the vertex context menu. This action deletes the vertex and adjusts the shape of the polygon, drawing a new straight line between the two vertices that were neighbors of the deleted vertex.
Moving the polygon.	Move the pointer inside the polygon. The pointer changes to a fleur shape  . Click and drag the mouse to move the polygon.
Changing the color of the polygon	Move the pointer inside the polygon. Right-click and select Set Color from the context menu.
Retrieving the coordinates of the vertices	Move the pointer inside the polygon. Right-click and select Copy Position from the context menu. <code>impoly</code> copies an n -by-2 array containing the x - and y -coordinates of each vertex to the clipboard. n is the number of vertices you specified.

Methods

Each `impoly` object supports a number of methods, listed below. Methods inherited from the base class are links to that class.

addNewPositionCallback — Add new-position callback to ROI object

See `imroi` for information.

createMask — Create mask within image

See `imroi` for information.

delete — Delete ROI object

See `imroi` for information.

getColor — Get color used to draw ROI object

See `imroi` for information.

getPosition — Return current position of polygon

`pos = getPosition(h)` returns the current position of the polygon `h`. The returned position, `pos`, is an N-by-2 array `[X1 Y1; ...; XN YN]`.

getPositionConstraintFcn — Return function handle to current position constraint function

See `imroi` for information.

removeNewPositionCallback — Remove new-position callback from ROI object.

See `imroi` for information.

resume — Resume execution of MATLAB command line

See `imroi` for information.

setClosed — Set geometry of polygon

`setClosed(TF)` sets the geometry of the polygon. `TF` is a logical scalar. `true` means that the polygon is closed. `false` means that the polygon is an open polyline.

setColor — Set color used to draw ROI object

See `imroi` for information.

setConstrainedPosition — Set ROI object to new position

See `imroi` for information.

setPosition – Set polygon to new position

`setPosition(h,pos)` sets the polygon `h` to a new position. The new position, `pos`, is an n -by-2 array, `[x1 y1; ..; xn yn]` where each row specifies the position of a vertex of the polygon.

setPositionConstraintFcn – Set position constraint function of ROI object.

See `imroi` for information.

setVerticesDraggable – Control whether vertices may be dragged

`setVerticesDraggable(h,TF)` sets the interactive behavior of the vertices of the polygon `h`. `TF` is a logical scalar. `True` means that the vertices of the polygon are draggable. `False` means that the vertices of the polygon are not draggable.

wait – Block MATLAB command line until ROI creation is finished

See `imroi` for information.

Tips

If you use `impoly` with an axes that contains an image object, and do not specify a position constraint function, users can drag the polygon outside the extent of the image and lose the polygon. When used with an axes created by the `plot` function, the axes limits automatically expand when the polygon is dragged outside the extent of the axes.

Examples

Example 1

Display updated position in the title. Specify a position constraint function using `makeConstrainToRectFcn` to keep the polygon inside the original `xlim` and `ylim` ranges.

```
figure, imshow('gantrycrane.png');  
h = impoly(gca, [188,30; 189,142; 93,141; 13,41; 14,29]);  
setColor(h,'yellow');  
addNewPositionCallback(h,@(p) title(mat2str(p,3)));  
fcn = makeConstrainToRectFcn('impoly',get(gca,'XLim'),...  
    get(gca,'YLim'));  
setPositionConstraintFcn(h,fcn);
```


Example 2

Interactively place a polygon by clicking to specify vertex locations. Double-click or right-click to finish positioning the polygon. Use `wait` to block the MATLAB command line. Double-click on the polygon to resume execution of the MATLAB command line.

```
figure, imshow('pout.tif');  
h = impoly;  
position = wait(h);
```

See also

`imellipse`, `imfreehand`, `imline`, `impoint`, `imrect`, `imroi`,
`makeConstraintToRectFcn`

impositionrect

Purpose Create draggable position rectangle

Syntax `H = impositionrect(hparent,position)`

Note This function is obsolete and may be removed in future versions. Use `imrect` instead.

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.

Tips 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> can be a three-element vector specifying an RGB triplet, or a text string specifying the long or short names of a predefined color, such as 'white' or 'w'. For a complete list of these predefined colors and their short names, see <code>ColorSpec</code>.</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>

impositionrect

Function	Description
	The return value, <code>id</code> , is used only with <code>removeNewPositionCallback</code> .
<code>removeNewPositionCallback</code>	Removes the corresponding function from the new-position callback list. <code>api.removeNewPositionCallback(id)</code> where <code>id</code> is the identifier returned by <code>api.addNewPositionCallback</code>
<code>setDragConstraintCallback</code>	Sets the drag constraint function to be the specified function handle, <code>fcn</code> . <code>api.setDragConstraintCallback(fcn)</code> Whenever the position rectangle is moved because of a mouse drag, the constraint function is called using the syntax: <code>constrained_position = fcn(new_position)</code> 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.

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 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 images and RGB images.

If you call `improfile` with one of these syntax, 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 non-interactively, using these syntax.

```
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 syntax 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 non-default spatial coordinate system for the input image, use these syntax.

```
[...] = 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 (`{}`).

improfile


Value	Description
{'nearest'}	Nearest-neighbor interpolation
'bilinear'	Bilinear interpolation
'bicubic'	Bicubic interpolation

Class Support

The input image can be `uint8`, `uint16`, `int16`, `single`, `double`, or `logical`. All other inputs and outputs must be `double`.

Examples

```
I = imread('liftingbody.png');  
x = [19 427 416 77];  
y = [96 462 37 33];  
improfile(I,x,y),grid on;
```




See Also

`impixel` | `interp2`

Purpose Display Save Image dialog box

Syntax [filename, ext, user_canceled] = imputfile

Description [filename, ext, user_canceled] = imputfile displays the Save Image dialog box (shown below) which you can use to specify the full path and format of a file. Using the dialog box, you can navigate to folders in a file system and select a particular file or specify the name of a new file. `imputfile` limits the types of files displayed in the dialog box to the image file format selected in the Files of Type menu.



When you click **Save**, `imputfile` returns the full path to the file in `filename` and the file type selected from the Files of Type menu in `ext`. `imputfile` does not automatically add a file name extension (such as `.jpg`) to the file name.

imputfile

If the user clicks **Cancel** or closes the Save Image dialog box, `imputfile` returns, setting `user_canceled` to `True` (1), and `settingfilename` and `ext` to empty strings; otherwise, `user_canceled` is `False` (0).

Note The Save Image dialog box is modal; it blocks the MATLAB command line until you click **Save** or cancel the operation.

See Also

`imformats` | `imtool` | `imgetfile` | `imsave`

Purpose Image pyramid reduction and expansion

Syntax `B = impyramid(A, direction)`

Description `B = impyramid(A, direction)` computes a Gaussian pyramid reduction or expansion of `A` by one level. `direction` can be 'reduce' or 'expand'.

If `A` is `m`-by-`n` and `direction` is 'reduce', then the size of `B` is `ceil(M/2)`-by-`ceil(N/2)`. If `direction` is 'expand', then the size of `B` is `(2*M-1)`-by-`(2*N-1)`.

Reduction and expansion take place only in the first two dimensions. For example, if `A` is 100-by-100-by-3 and `direction` is 'reduce', then `B` is 50-by-50-by-3.

Note that `impyramid` uses the kernel specified on page 533 of the Burt and Adelson paper cited below:

$$w = \left[\frac{1}{4} - \frac{a}{2}, \frac{1}{4}, a, \frac{1}{4}, \frac{1}{4} - \frac{a}{2} \right], \text{ where } a = 0.375.$$

The parameter `a` is chosen to be 0.375 so that the equivalent weighting function is close to a Gaussian shape and the weights can be readily applied using fixed-point arithmetic.

Class support `A` can be any numeric class except `uint64` or `int64`, or it can be logical. The class of `B` is the same as the class of `A`.

Examples Compute a four-level multiresolution pyramid of the cameraman image.

```
I0 = imread('cameraman.tif');
I1 = impyramid(I0, 'reduce');
I2 = impyramid(I1, 'reduce');
I3 = impyramid(I2, 'reduce');

imshow(I0)
```

impyramid

```
figure, imshow(I1)  
figure, imshow(I2)  
figure, imshow(I3)
```

References

[1] Burt and Adelson, "The Laplacian Pyramid as a Compact Image Code," IEEE Transactions on Communications, vol. COM-31, no. 4, April 1983, pp. 532-540.

[2] Burt, "Fast Filter Transforms for Image Processing," Computer Graphics and Image Processing, vol. 16, 1981, pp. 20-51

See Also

`imresize`

Purpose Quantize image using specified quantization levels and output values

Syntax

```
quant_A = imquantize(A,levels)
quant_A = imquantize( ___,values)
[quant_A,index] = imquantize( ___ )
```

Description `quant_A = imquantize(A,levels)` quantizes image `A` using specified quantization values contained in the `N` element vector `levels`. Output image `quant_A` is the same size as `A` and contains `N + 1` discrete integer values in the range 1 to `N + 1` which are determined by the following criteria:

- If $A(k) \leq levels(1)$, then $quant_A(k) = 1$.
- If $levels(m-1) < A(k) \leq levels(m)$, then $quant_A(k) = m$.
- If $A(k) > levels(N)$, then $quant_A(k) = N + 1$.

Note that `imquantize` assigns values to the two implicitly defined end intervals:

- $A(k) \leq levels(1)$
- $A(k) > levels(N)$

`quant_A = imquantize(___,values)` adds the `N + 1` element vector `values` where `N = length(levels)`. Each of the `N + 1` elements of `values` specify the quantization value for one of the `N + 1` discrete pixel values in `quant_A`.

- If $A(k) \leq levels(1)$, then $quant_A(k) = values(1)$.
- If $levels(m-1) < A(k) \leq levels(m)$, then $quant_A(k) = values(m)$.
- If $A(k) > levels(N)$, then $quant_A(k) = values(N + 1)$.

`[quant_A,index] = imquantize(___)` returns matrix *index* such that:

```
quant_A = values(index)
```

imquantize

Input Arguments

A - Input image

image

Input image, specified as a numeric array of any dimension.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32 | uint64

levels - Quantization levels

vector

Quantization levels, specified as an N element vector. Values of the discrete quantization levels must be in monotonically increasing order.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32 | uint64

values - Quantization values

vector

Quantization values, specified as an N + 1 element vector.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32 | uint64

Output Arguments

quant_A - Quantized output image

image

Quantized output image, returned as a numeric array the same size as A. If input argument *values* is specified, then *quant_A* is the same data type as *values*. If *values* is not specified, then *quant_A* is of class double.

index - Mapping matrix

matrix

Mapping matrix, returned as a matrix the same size as input image A. It contains integer indices which access *values* to construct the

output image: `quant_A = values(index)`. If input argument *values* is not defined, then `index = quant_A`.

Data Types

double

Examples**Image Threshold**

Compute multiple thresholds for an image using `multithresh` and apply those thresholds to the image using `imquantize` to get segment labels.

```
I = imread('circlesBrightDark.png');
```

Computing two thresholds will quantize the image into three discrete levels

```
thresh = multithresh(I,2);
```

```
seg_I = imquantize(I,thresh); % apply the thresholds to obtain segments
```

```
RGB = label2rgb(seg_I); % convert to color image
```


```
figure;
```

```
imshowpair(I,RGB,'montage'); % display images side-by-side
```

```
axis off;
```

```
title('Original Image
```

```
RGB Segmented Image')
```



Two Schemes for RGB Image Quantization

Quantize truecolor RGB image to 8 levels using multi-level thresholding function `multithresh`. Compare results between thresholding the entire RGB image versus plane-by-plane thresholding.

Read in truecolor RGB image.


```
I = imread('peppers.png');  
imshow(I); axis off;  
title('RGB Image');
```

RGB Image



imquantize

Generate thresholds for seven levels from the entire RGB image. then repeat process for each plane in the RGB image.

```
threshRGB = multithresh(I,7) % 7 thresholds from entire RGB image
```

```
threshPlanes = zeros(3,7); % initialize
% Compute thresholds for each R, G and B plane
for i = 1:3
    threshPlane(i,:) = multithresh(I(:,:,i),7);
end
threshPlane
```

```
threshRGB =
```

```
    24    51    79   109   141   177   221
```

```
threshPlane =
```

```
    40    69    92   125   159   195   231
    27    49    74   100   128   164   209
    18    38    60    86   127   165   222
```

Process entire image with set of threshold values computed from entire image. Similarly, process each RGB plane separately using threshold vector computed from the given plane. Add black (0) and white (255) to value vector which assigns values to output image.

```
value = [0 threshRGB(2:end) 255];
% Quantize entire image using one threshold vector
quantRGB = imquantize(I, threshRGB, value);
```

```
quantPlane = zeros( size(I) );
% Quantize each RGB plane using threshold vector generated for that plane
for i = 1:3
    value = [0 threshPlane(i,2:end) 255] % output value to assign
    quantPlane(:,:,i) = imquantize(I(:,:,i),threshPlane(i,:),value);
end
```

```
quantPlane = uint8(quantPlane); % convert from double to uint8
```

Display both posterized images and note the visual differences in the two thresholding schemes.

Full RGB Image Quantization



The following code snippet computes the number of unique RGB pixel vectors in each output image. Note that the plane-by-plane thresholding scheme yields about 23% more colors than the full RGB image scheme.

imquantize

```
% convert images to mx3 matrices
dim = size( quantRGB );
quantRGBmx3 = reshape(quantRGB, prod(dim(1:2)), 3);
quantPlanemx3 = reshape(quantPlane, prod(dim(1:2)), 3);

% extract only unique 3 element RGB pixel vectors from each matrix
colorsRGB = unique( quantRGBmx3, 'rows' );
colorsPlane = unique( quantPlanemx3, 'rows' );

disp(['Number of unique colors in RGB image          : ' int2str(length(colorsRGB))])
disp(['Number of unique colors in Plane-by-Plane image : ' int2str(length(colorsPlane))])

Number of unique colors in RGB image          : 188
Number of unique colors in Plane-by-Plane image : 231
```

Threshold a Grayscale Image from 256 to 8 levels

This example reduces the number of discrete levels in an image from 256 to 8. It then demonstrates two different schemes for assigning values to each of the eight output levels.

```
I = imread('coins.png');
imshow(I); axis off;
title('Grayscale Image');
```

Grayscale Image



Obtain seven thresholds from `multithresh` to split the image into eight levels.

```
thresh = multithresh(I,7); % 7 thresholds result in 8 image levels
```

Construct the `valuesMax` vector such that the maximum value in each quantization interval is assigned to the eight levels of the output image.

```
valuesMax = [thresh max(I(:))]  
[quant8_I_max,index] = imquantize(I,thresh,valuesMax);
```

```
valuesMax =
```

```
65  88  119  149  169  189  215  255
```

imquantize

Similarly, construct `valuesMin` such that the minimum value in each quantization interval is assigned to the eight levels of the output image

```
valuesMin = [min(I(:)) thresh]
```


```
valuesMin =
```

```
    23    65    88   119   149   169   189   215
```

Instead of calling `imquantize` again with the vector `valuesMin`, use the output argument `index` to assign those values to the output image.

```
quant8_I_min = valuesMin(index);
```

Display both eight-level output images side by side.



See Also

`multithresh` | `label2rgb` | `rgb2ind`

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 (including `uint64` or `int64`) or logical arrays with the same class and any dimension. `IM` is of the same class as `marker` and `mask`.

Performance Note

This function may take advantage of hardware optimization for data types `logical`, `uint8`, and `single` to run faster. Hardware optimization requires `marker` and `mask` to be 2-D images and `conn` to be either 4 or 8.

Algorithms

`imreconstruct` uses the fast hybrid grayscale reconstruction algorithm described in [1].

References

[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.

See Also

`imclearborder` | `imextendedmax` | `imextendedmin` | `imfill` | `imhmax` | `imhmin` | `imimposemin`

imrect

Purpose Create draggable rectangle

Syntax

```
H = imrect
h = imrect(hparent)
h = imrect(hparent, position)
H = imrect(..., param1, val1,...)
```

Description `H = imrect` begins interactive placement of a rectangle on the current axes. The function returns `h`, a handle to an `imrect` object. The rectangle has a context menu associated with it that controls aspects of its appearance and behavior—see “Interactive Behavior” on page 3-509. Right-click on the rectangle to access this context menu.


`h = imrect(hparent)` begins interactive placement of a rectangle on the object specified by `hparent`. `hparent` specifies the HG parent of the rectangle graphics, which is typically an axes but can also be any other object that can be the parent of an `hggroup`.

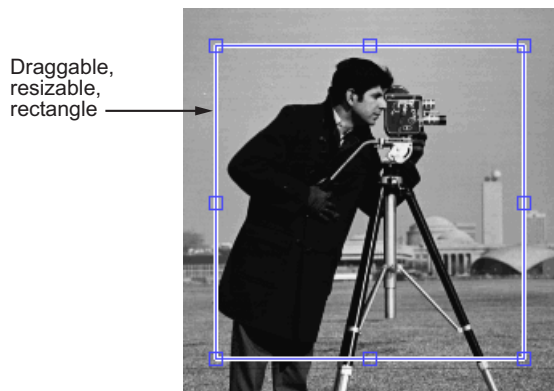
`h = imrect(hparent, position)` creates a draggable rectangle on the object specified by `hparent`. `position` is a four-element vector that specifies the initial size and location of the rectangle. `position` has the form `[xmin ymin width height]`.

`H = imrect(..., param1, val1,...)` creates a draggable rectangle, specifying parameters and corresponding values that control the behavior of the rectangle. The following table lists the parameter available. Parameter names can be abbreviated, and case does not matter.



Parameter	Description
'PositionConstraintFcn	Function handle <code>fcn</code> that is called whenever the mouse is dragged. You can use this function to control where the rectangle can be dragged. See the help for the <code>setPositionConstraintFcn</code> method for information about valid function handles.

Interactive Behavior

When you call `imrect` with an interactive syntax, the pointer changes to a cross hairs  when over the image. You can create the rectangle and adjust its size and position using the mouse. The rectangle also supports a context menu that you can use to control aspects of its appearance and behavior. The following figure shows the rectangle.



The following table lists the interactive behaviors supported by `imrect`.

Interactive Behavior	Description
Moving the rectangle.	Move the pointer inside the rectangle. The pointer changes to a fleur shape  . Click and drag the mouse to move the rectangle.
Resizing the rectangle.	Move the pointer over any of the edges or corners of the rectangle, the shape changes to a double-ended arrow,  . Click and drag the edge or corner using the mouse.
Changing the color of the rectangle.	Move the pointer inside the rectangle. Right-click and select Set Color from the context menu.

Interactive Behavior	Description
Retrieving the coordinates of the current position	Move the pointer inside the polygon. Right-click and select Copy Position from the context menu. <code>imrect</code> copies a four-element position vector to the clipboard.
Preserve the current aspect ratio of the rectangle during interactive resizing.	Move the pointer inside the rectangle. Right-click and select Fix Aspect Ratio from the context menu.

Methods

Each `imrect` object supports a number of methods, listed below. Type `methods imrect` to see a list of the methods.

addNewPositionCallback – Add new-position callback to ROI object

See `imroi` for information.

createMask – Create mask within image

See `imroi` for information.

delete – Delete ROI object

See `imroi` for information.

getColor – Get color used to draw ROI object

See `imroi` for information.

getPosition – Return current position of rectangle

`pos = getPosition(h)` returns the current position of the rectangle `h`. The returned position, `pos`, is a 1-by-4 array [`xmin` `ymin` `width` `height`].

getPositionConstraintFcn – Return function handle to current position constraint function

See `imroi` for information.

removeNewPositionCallback – Remove new-position callback from ROI object.

See `imroi` for information.

resume – Resume execution of MATLAB command line

See `imroi` for information.

setColor – Set color used to draw ROI object

See `imroi` for information.

setConstrainedPosition – Set ROI object to new position

See `imroi` for information.

setFixedAspectRatioMode – Control whether aspect ratio preserved during resize

`setFixedAspectRatioMode(h,TF)` sets the interactive resize behavior of the rectangle `h`. `TF` is a logical scalar. True means that the current aspect ratio is preserved during interactive resizing. False means that interactive resizing is not constrained.

setPosition – Set rectangle to new position

`setPosition(h,pos)` sets the rectangle `h` to a new position. The new position, `pos`, has the form `[xmin ymin width height]`.

setPositionConstraintFcn – Set position constraint function of ROI object

See `imroi` for information.

setResizable – Set resize behavior of rectangle

`setResizable(h,TF)` sets whether the rectangle `h` may be resized interactively. `TF` is a logical scalar.

wait – Block MATLAB command line until ROI creation is finished

See `imroi` for information.

Tips

If you use `imrect` with an axes that contains an image object, and do not specify a position constraint function, users can drag the rectangle outside the extent of the image. When used with an axes created by the `plot` function, the axes limits automatically expand to accommodate the movement of the rectangle.

When the API function `setResizable` is used to make the rectangle non-resizable, the **Fix Aspect Ratio** context menu item is not provided.

Examples

Example 1

Display updated position in the title. Specify a position constraint function using `makeConstrainToRectFcn` to keep the rectangle inside the original `Xlim` and `Ylim` ranges.

```
figure, imshow('cameraman.tif');  
h = imrect(gca, [10 10 100 100]);  
addNewPositionCallback(h,@(p) title(mat2str(p,3)));  
fcn = makeConstrainToRectFcn('imrect',get(gca,'XLim'),get(gca,'YLim'));  
setPositionConstraintFcn(h,fcn);
```

Now drag the rectangle using the mouse.

Example 2

Interactively place a rectangle by clicking and dragging. Use `wait` to block the MATLAB command line. Double-click on the rectangle to resume execution of the MATLAB command line.

```
figure, imshow('pout.tif');  
h = imrect;  
position = wait(h);
```

See Also

`imellipse` | `imfreehand` | `imline` | `impoint` | `impoly` | `imroi` | `makeConstrainToRectFcn`

Purpose 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 a constant intensity value, and whose external boundary pixels all have a lower value.

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

imregionalmax

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.

Examples

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](#) | [imextendedmax](#) | [imhmax](#) | [imreconstruct](#) | [imregionalmin](#)

imregionalmin

Purpose 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 a constant intensity value, and whose external boundary pixels all have a higher value.

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.

Examples

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
```

imregionalmin

0 0 0 0 0 0 0 0 0 0

See Also

conndef | imextendedmin | imhmin | imimposemin | imreconstruct |
imregionalmax

Purpose	Configurations for intensity-based registration
Syntax	<code>[optimizer,metric] = imregconfig(modality)</code>
Description	<code>[optimizer,metric] = imregconfig(modality)</code> creates <code>optimizer</code> and <code>metric</code> configurations that you pass to <code>imregister</code> to perform intensity-based image registration. <code>imregconfig</code> returns <code>optimizer</code> and <code>metric</code> with default settings to provide a basic registration configuration.
Input Arguments	<p>modality - Image capture modality 'monomodal' 'multimodal'</p> <p>Image capture modality describes how your images have been captured, specified as either 'monomodal' (captured on the same device) or 'multimodal' (captured on different devices).</p>
Output Arguments	<p>optimizer - Optimization configuration <code>optimizer</code> object</p> <p>Optimization configuration describes the method for optimizing the similarity metric, returned as one of the <code>optimizer</code> objects, <code>registration.optimizer.RegularStepGradientDescent</code> or <code>registration.optimizer.OnePlusOneEvolutionary</code></p> <p>metric - Metric configuration <code>metric</code> object</p> <p>Metric configuration describes the image similarity metric to be optimized during registration, returned as one of the <code>metric</code> objects, <code>registration.metric.MeanSquares</code> or <code>registration.metric.MattesMutualInformation</code>.</p>
Tips	<ul style="list-style-type: none">• Your registration results can improve if you adjust the <code>optimizer</code> or <code>metric</code> settings. For example, if you increase the number of iterations in the optimizer, reduce the optimizer step size, or change

imregconfig

the number of samples in a stochastic metric, the registration improves to a point, at the expense of performance.

Definitions

Monomodal

Images captured on the same device. Monomodal images have similar brightness ranges.

Multimodal

Images captured on different devices. Multimodal images usually have different brightness ranges.

Examples

Create Optimizer and Metric Configurations to Register Images Captured on the Same Device

Load the images into the workspace and display them.

```
fixed = imread('pout.tif');  
moving = imrotate(fixed, 5, 'bilinear', 'crop');  
imshowpair(fixed, moving, 'Scaling', 'joint');
```



Create the optimizer and metric. The two images in this example were captured on the same device, so we'll set the modality to 'monomodal'.

```
[optimizer, metric] = imregconfig('monomodal')

optimizer =

    registration.optimizer.RegularStepGradientDescent

Properties:
    GradientMagnitudeTolerance: 1.000000e-04
    MinimumStepLength: 1.000000e-05
    MaximumStepLength: 6.250000e-02
    MaximumIterations: 100
    RelaxationFactor: 5.000000e-01

metric =
```

imregconfig

```
registration.metric.MeanSquares
```

This class has no properties.

Pass optimizer and metric to `imregister` to perform the registration.

```
movingRegistered = imregister(moving, fixed, 'rigid', optimizer, metric);
```

View the registered images

```
figure  
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



See Also

```
imshowpair |  
imregisterregistration.metric.MattesMutualInformation  
| registration.metric.MeanSquares |  
registration.optimizer.RegularStepGradientDescent |  
registration.optimizer.OnePlusOneEvolutionary |
```


Concepts

- “Intensity-Based Automatic Image Registration”

imregister

Purpose

Intensity-based image registration

Syntax

```
movingRegistered = imregister(moving, fixed, transformType,  
    optimizer, metric)  
movingRegistered = imregister(moving, fixed, transformType,  
    optimizer, metric, Name, Value)
```

Description

`movingRegistered = imregister(moving, fixed, transformType, optimizer, metric)` transforms a 2-D or 3-D moving image so that it is registered with the fixed image.

- Both `moving` and `fixed` images must be of the same dimensionality, either 2-D or 3-D

`movingRegistered = imregister(moving, fixed, transformType, optimizer, metric, Name, Value)` specifies additional options with one or more `Name, Value` pair arguments.

Input Arguments

moving - Moving image

grayscale image

Moving image is the image to be transformed into alignment with a reference image, specified as a 2-D or 3-D grayscale image.

Data Types

single | double | int8 | int16 | int32 | uint8 | uint16 | uint32

fixed - Fixed image

grayscale image

Fixed image is the reference image in the target orientation, specified as a 2-D or 3-D grayscale image.

Data Types

single | double | int8 | int16 | int32 | uint8 | uint16 | uint32

transformType - Transform type`'translation' | 'rigid' | 'similarity' | 'affine'`

Transform type is the geometric transformation to be applied to the moving image, specified as one of the text strings listed in this table.

Transform Type	Description
'translation'	(x,y) translation.
'rigid'	Rigid transformation consisting of translation and rotation.
'similarity'	Nonreflective similarity transformation consisting of translation, rotation, and scale.
'affine'	Affine transformation consisting of translation, rotation, scale, and shear.

optimizer - Optimization configuration

optimizer object

Optimizer configuration describes the method for optimizing the similarity metric, specified as an optimizer object returned by the `imregconfig` function: `registration.optimizer.RegularStepGradientDescent` or `registration.optimizer.OnePlusOneEvolutionary`.

metric - Metric configuration

metric object

Metric configuration describes the image similarity metric to be optimized during registration, specified as a metric object returned by the `imregconfig` function: `registration.metric.MeanSquares` or `registration.metric.MattesMutualInformation`.

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can

imregister

specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

Example: `'DisplayOptimization',1` enables the verbose optimization mode.

DisplayOptimization - Verbose optimization flag

`false (default) | true`

Verbose optimization flag controls whether optimization information is displayed in the command window during the registration process, specified as the comma-separated pair consisting of `'DisplayOptimization'` and either `true` or `false`.

Data Types

logical

PyramidLevels - Number of pyramid levels

`3 (default) | positive integer`

Number of pyramid levels that `imregister` uses during the registration process, specified as the comma-separated pair consisting of `'PyramidLevels'` and a positive integer.

Example: `'PyramidLevels',4` sets the number of pyramid levels to 4.

Data Types

double

Output Arguments

movingRegistered - Transformed moving image

matrix

Transformed moving image, returned as a matrix. Any fill pixels introduced that do not correspond to locations in the original image are 0.

Tips

- Create `optimizer` and `metric` with the `imregconfig` function before calling `imregister`.
- Your registration results can improve if you adjust the `optimizer` or `metric` settings. For example, if you increase the number of

iterations in the optimizer, reduce the optimizer step size, or change the number of samples in a stochastic metric, the registration improves to a point, at the expense of performance.

- If the spatial scaling of your images differs by more than 10%, you should resize them with `imresize` before registering them.
- Use `imshowpair` or `imfuse` to visualize the results of registration.
- Use `imregister` to register several images in an automated workflow.

Definitions

Monomodal

Images captured on the same device. Monomodal images have similar brightness ranges.

Multimodal

Images captured on different devices. Multimodal images usually have different brightness ranges.

Examples

Register 2-D Images from Same Capture Device

Read the reference image and create an unregistered copy.

```
fixed = imread('pout.tif');  
moving = imrotate(fixed, 5, 'bilinear', 'crop');
```

View the misaligned images.

```
imshowpair(fixed, moving, 'Scaling', 'joint');
```

imregister



Create a configuration suitable for registering images from the same sensor.

```
[optimizer, metric] = imregconfig('monomodal')
```

```
optimizer =
```

```
    registration.optimizer.RegularStepGradientDescent
```

```
Properties:
```

```
    GradientMagnitudeTolerance: 1.000000e-04
```

```
    MinimumStepLength: 1.000000e-05
```

```
    MaximumStepLength: 6.250000e-02
```

```
    MaximumIterations: 100
```

```
    RelaxationFactor: 5.000000e-01
```

```
metric =
```

```
registration.metric.MeanSquares
```

This class has no properties.

Now modify the optimizer configuration to get more precision.

```
optimizer.MaximumIterations = 300;  
optimizer.MinimumStepLength = 5e-4;
```

Perform the registration.

```
movingRegistered = imregister(moving, fixed, 'rigid', optimizer, metric);
```

View registered images.

```
figure  
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



imregister

See Also

`imregconfig` | `imshowpair` | `imfuse` |
`imtransformregistration.metric.MattesMutualInformation`
| `registration.metric.MeanSquares` |
`registration.optimizer.RegularStepGradientDescent` |
`registration.optimizer.OnePlusOneEvolutionary` |

Related Examples

- “Registering Multimodal MRI Images”

Concepts

- “Intensity-Based Automatic Image Registration”

Purpose

Resize image

Syntax

```
B = imresize(A, scale)
B = imresize(A, [numrows numcols])
[Y newmap] = imresize(X, map, scale)
[...] = imresize(..., method)
[...] = imresize(..., parameter, value, ...)
```

Description

`B = imresize(A, scale)` returns image `B` that is `scale` times the size of `A`. The input image `A` can be a grayscale, RGB, or binary image. If `scale` is between 0 and 1.0, `B` is smaller than `A`. If `scale` is greater than 1.0, `B` is larger than `A`.

`B = imresize(A, [numrows numcols])` returns image `B` that has the number of rows and columns specified by `[numrows numcols]`. Either `numrows` or `numcols` may be NaN, in which case `imresize` computes the number of rows or columns automatically to preserve the image aspect ratio.

`[Y newmap] = imresize(X, map, scale)` resizes the indexed image `X`. `scale` can either be a numeric scale factor or a vector that specifies the size of the output image (`[numrows numcols]`). By default, `imresize` returns a new, optimized colormap (`newmap`) with the resized image. To return a colormap that is the same as the original colormap, use the 'Colormap' parameter (see below).

`[...] = imresize(..., method)` resizes the indexed image. *method* can be (1) a text string that specifies a general interpolation method, (2) a text string that specifies an interpolation kernel, or (3) a two-element cell array that specifies an interpolation kernel.

(1) Text String Specifying Interpolation Method

Method Name	Description
'nearest'	Nearest-neighbor interpolation; the output pixel is assigned the value of the pixel that the point falls within. No other pixels are considered.
'bilinear'	Bilinear interpolation; the output pixel value is a weighted average of pixels in the nearest 2-by-2 neighborhood
'bicubic'	Bicubic interpolation (the default); the output pixel value is a weighted average of pixels in the nearest 4-by-4 neighborhood

(2) Text String Specifying Interpolation Kernel

Kernel Name	Description
'box'	Box-shaped kernel
'triangle'	Triangular kernel (equivalent to 'bilinear')
'cubic'	Cubic kernel (equivalent to 'bicubic')
'lanczos2'	Lanczos-2 kernel
'lanczos3'	Lanczos-3 kernel

(3) Two-element Cell Array Specifying Interpolation Kernel

Form	Description
{f,w}	f is a function handle for a custom interpolation kernel and w is the custom kernel's width. $f(x)$ must be zero outside the interval $-w/2 \leq x < w/2$. Your function handle f may be called with a scalar or a vector input.

[...] = imresize(..., parameter, value, ...) you can control various aspects of the resizing operation by specifying parameter/value pairs with any of the previous syntaxes. The following table lists these parameters.

Parameter	Value
'Antialiasing'	A Boolean value that specifies whether to perform antialiasing when shrinking an image. The default value depends on the interpolation method. If the method is nearest-neighbor ('nearest'), the default is false; for all other interpolation methods, the default is true.
'Colormap'	A text string that specifies whether imresize returns an optimized colormap or the original colormap (Indexed images only). If set to 'original', the output colormap (newmap) is the same as the input colormap (map). If set to 'optimized', imresize returns a new optimized colormap. The default value is 'optimized'.
'Dither'	A Boolean value that specifies whether to perform color dithering (Indexed images only). The default value is true.

Parameter	Value
'Method'	As described above
'OutputSize'	A two-element vector, [numrows numcols], that specifies the size of the output image. If you specify NaN for one of the values, <code>imresize</code> computes the value of the dimension to preserve the aspect ratio of the original image.
'Scale'	A scalar or two-element vector that specifies the resize scale factors. If you specify a scalar, <code>imresize</code> uses the value as the scale factor for each dimension. If you specify a vector, <code>imresize</code> uses the individual values as the scale factors for the row and column dimensions, respectively.

Tips

The function `imresize` changed in version 5.4 (R2007a). Previous versions of the Image Processing Toolbox used a somewhat different algorithm by default. If you need the same results produced by the previous implementation, use the function `imresize_old`.

For bicubic interpolation, the output image may have some values slightly outside the range of pixel values in the input image. This may also occur for user-specified interpolation kernels.

Class Support

The input image can be numeric or logical and it must be nonsparse. The output image is of the same class as the input image. An input image that is an indexed image can be `uint8`, `uint16`, or `double`.

Examples

Shrink by factor of two using the defaults of bicubic interpolation and antialiasing.

```
I = imread('rice.png');  
J = imresize(I, 0.5);  
figure, imshow(I), figure, imshow(J)
```

Shrink by factor of two using nearest-neighbor interpolation. (This is the fastest method, but it has the lowest quality.)

```
J2 = imresize(I, 0.5, 'nearest');
```

Resize an indexed image

```
[X, map] = imread('trees.tif');  
[Y, newmap] = imresize(X, map, 0.5);  
imshow(Y, newmap)
```

Resize an RGB image to have 64 rows. The number of columns is computed automatically.

```
RGB = imread('peppers.png');  
RGB2 = imresize(RGB, [64 NaN]);
```

See Also

[imrotate](#) | [imtransform](#) | [tformarray](#) | [interp2](#)

Purpose Region-of-interest (ROI) base class

Description Because the `imroi` class is abstract, creating an instance of the `imroi` class is not allowed.

Methods

`imroi` supports the following methods. Type `methods imroi` to see a complete list.

addNewPositionCallback – Add new-position callback to ROI object

`id = addNewPositionCallback(h, fcn)` adds the function handle `fcn` to the list of new-position callback functions of the ROI object `h`. Whenever the ROI object changes its position each function in the list is called with the syntax:

```
fcn(pos)
```

where `pos` is of the form returned by the object's `getPosition` method.

The return value, `id`, is used only with `removeNewPositionCallback`.

createMask – Create mask within image

`BW = createMask(h)` returns a mask, or binary image, that is the same size as the input image with 1s inside the ROI object `h` and 0s everywhere else. The input image must be contained within the same axes as the ROI.

`BW = createMask(h, h_im)` returns a mask the same size as the image `h_im` with 1s inside the ROI object `h` and 0s outside. This syntax is required when the axes that contain the ROI hold more than one image.

delete – Delete ROI object

`delete(h)` deletes the ROI object `h`

getColor – Get color used to draw ROI object.

`color = getColor(h)` gets the color used to draw the ROI object `h`. The three-element vector `color` specifies an RGB triplet.

getPosition – Return current position of ROI object

`pos = getPosition(h)` returns current position of the ROI object `h`.

getPositionConstraintFcn – Return function handle to current position constraint function

`fcn = getPositionConstraintFcn(h)` returns a function handle `fcn` to the current position constraint function of the ROI object `h`.

removeNewPositionCallback – Remove new-position callback from ROI object

`removeNewPositionCallback(h,id)` removes the corresponding function from the new-position callback list of the ROI object `h`. `id` is the identifier returned by the `addNewPositionCallback` method.

resume – Resume execution of MATLAB command line

`resume(h)` resumes execution of the MATLAB command line. When called after a call to `wait`, `resume` causes `wait` to return an accepted position. The `resume` method is useful when you need to exit `wait` from a callback function.

setColor – Set color used to draw ROI object.

`setColor(h,new_color)` sets the color used to draw the ROI object `h`. `new_color` can be a three-element vector specifying an RGB triplet, or a text string specifying the long or short name of a predefined color, such as 'white' or 'w'. See `ColorSpec` for a list of predefined colors.

setConstrainedPosition – Set ROI object to new position

`setConstrainedPosition(h,candidate_position)` sets the ROI object `h` to a new position. The candidate position is subject to the position constraint function. `candidate_position` is of the form expected by the `setPosition` method.

setPositionConstraintFcn – Set position constraint function of ROI object

`setPositionConstraintFcn(h,fcn)` sets the position constraint function of the ROI object `h` to be the specified function handle, `fcn`. Whenever the object is moved because of a mouse drag, the constraint function is called using the syntax:

```
constrained_position = fcn(new_position)
```

where `new_position` is of the form returned by the `getPosition` method. You can use the `makeConstrainToRectFcn` to create this function.

wait – Block MATLAB command line until ROI creation is finished

`accepted_pos = wait(h)` blocks execution of the MATLAB command line until you finish positioning the ROI object `h`. You indicate completion by double-clicking on the ROI object. The returned position, `accepted_pos`, is of the form returned by the object's `getPosition` method.

See Also

`makeConstrainToRectFcn`

Purpose Rotate image

Syntax

```
B = imrotate(A,angle)
B = imrotate(A,angle,method)
B = imrotate(A,angle,method,bbox)
```

Description `B = imrotate(A,angle)` rotates image `A` by `angle` degrees in a counterclockwise direction around its center point. To rotate the image clockwise, specify a negative value for `angle`. `imrotate` makes the output image `B` large enough to contain the entire rotated image. `imrotate` uses nearest neighbor interpolation, setting the values of pixels in `B` that are outside the rotated image to 0 (zero).

`B = imrotate(A,angle,method)` rotates image `A`, using the interpolation method specified by `method`. `method` is a text 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 image `A`, where `bbox` specifies the size of the returned image. `bbox` is a text string that can have one of the following values. The default value is enclosed in braces (`{}`).

imrotate

Value	Description
'crop'	Make output image B the same size as the input image A, cropping the rotated image to fit
{'loose'}	Make output image B large enough to contain the entire rotated image. B is generally larger than A.

Class Support

The input image can be numeric or logical. The output image is of the same class as the input image.


Performance Note

This function may take advantage of hardware optimization for data types `uint8`, `uint16`, and `single` to run faster.


Examples

Read a solar spectra image, stored in FITS format, and rotate 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');  
figure, imshow(I)  
figure, imshow(J)
```



Original Image



Rotated Image

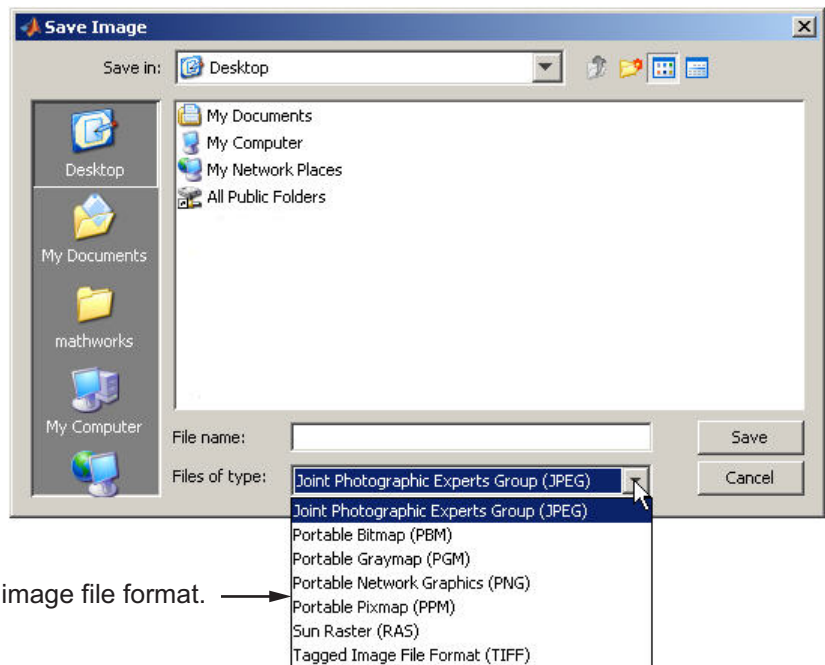
See Also

[imcrop](#) | [imresize](#) | [imtransform](#) | [tformarray](#)

Purpose Save Image Tool

Syntax
 imsave
 imsave(h)
 [filename, user_canceled] = imsave()

Description imsave creates a Save Image tool in a separate figure that is associated with the image in the current figure, called the target image. The Save Image tool displays an interactive file chooser dialog box (shown below) in which you can specify a path and filename. When you click **Save**, the Save Image tool writes the target image to a file using the image file format you select in the Files of Type menu. imsave uses imwrite to save the image, using default options.



imsave

If you specify a filename that already exists, `imsave` displays a warning message. Select **Yes** to use the filename or **No** to return to the dialog to select another filename. If you select **Yes**, the Save Image tool attempts to overwrite the target file.

Note The Save Image tool is modal; it blocks the MATLAB command line until you respond.

`imsave(h)` creates a Save Image tool associated with the image specified by the handle `h`. `h` can be an image, axes, uipanel, or figure handle. If `h` is an axes or figure handle, `imsave` uses the first image returned by `findobj(h, 'Type', 'image')`.

`[filename, user_canceled] = imsave()` returns the full path to the file selected in `filename`. If you press the **Cancel** button, `imsave` sets `user_canceled` to true (1); otherwise, false (0).

Remarks

In contrast to the **Save as** option in the figure **File** menu, the Save Image tool saves only the image displayed in the figure. The **Save as** option in the figure window File menu saves the entire figure window, not just the image.

Examples

```
imshow peppers.png
imsave
```

See Also

`imformats` | `imgetfile` | `imputfile` | `imwrite`

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. The function returns `hpanel`, a 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.

imscrollpanel

Function	Description
setMagnification	<p>Sets the magnification of the target image in units of screen pixels per image pixel.</p> <pre>mag = api.setMagnification(new_mag)</pre> <p>where <code>new_mag</code> is a scalar magnification factor.</p>
getMagnification	<p>Returns the current magnification factor of the target image in units of screen pixels per image pixel.</p> <pre>mag = api.getMagnification()</pre> <p>Multiply <code>mag</code> by 100 to convert to percentage. For example if <code>mag=2</code>, the magnification is 200%.</p>
setMagnificationAndCenter	<p>Changes the magnification and makes the point <code>cx,cy</code> in the target image 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 target 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) api.setVisibleLocation([xmin ymin])</pre>

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>

imscrollpanel

Function	Description
<code>addNewLocationCallback</code>	<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>
<code>removeNewLocationCallback</code>	<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>
<code>replaceImage</code>	<pre>api.replaceImage(...,PARAM1,VAL1,PARAM2,VAL2,...)</pre> <p>replaces the image displayed in the scroll panel.</p> <pre>api.replaceImage(I) api.replaceImage(BW) api.replaceImage(RGB) api.replaceImage(I,MAP) api.replaceImage(filename)</pre> <p>By default, the new image data is displayed centered, at 100% magnification. The image handle is unchanged.</p>

Function	Description
	<p>The parameters you can specify include many of the parameters supported by <code>imshow</code>, including <code>'Colormap'</code>, <code>'DisplayRange'</code>, and <code>'InitialMagnification'</code>. In addition, you can use the <code>'PreserveView'</code> parameter to preserve the current magnification and centering of the image during replacement. Specify the logical scalar <code>True</code> to preserve current centering and magnification. Parameter names can be abbreviated, and case does not matter.</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`.

Examples

Create a scroll panel with a Magnification Box and an Overview tool.

1 Create a scroll panel.

```
hFig = figure('Toolbar','none',...
             'Menubar','none');
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.

imscrollpanel

```
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](#)

How To

- “Adding Navigation Aids to a GUI”

Purpose Display image

Note

- The syntax `imshow(x,y,...)` has been removed. Use the `imshow(...,'XData',x,'YData',y)` syntax instead.
 - The syntax `imshow(I,N)` has been removed. Your grayscale image will be displayed using 256 shades of gray.
 - The syntax `imshow(...,'truesize')` has been removed. Use the `imshow(...,'InitialMagnification',100)` syntax instead.
 - The syntax `imshow(...,'notruesize')` has been removed. Use the `imshow(...,'InitialMagnification','fit')` syntax instead.
-

Syntax

```
imshow(I)
imshow(I,[low high])
imshow(RGB)
imshow(BW)
imshow(X,map)
imshow(filename)
himage = imshow(...)
imshow(..., param1, val1, param2, val2,...)
```

Description

`imshow(I)` displays the grayscale image `I`.

`imshow(I,[low high])` displays the grayscale image `I`, specifying the display range for `I` in `[low high]`. 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

`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` displays the first image in the file. 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 in alphabetical order. Parameter names can be abbreviated, and case does not matter.

Parameter	Value
'Border'	<p>Text string that controls whether <code>imshow</code> includes a border around the image displayed in the figure window. Valid strings are 'tight' and 'loose'.</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> <p>By default, the border is set to the value returned by <code>iptgetpref('ImshowBorder')</code>.</p>
'Colormap'	<p>2-D, real, m-by-3 matrix specifying a colormap. <code>imshow</code> uses this to set the figure's <code>colormap</code> property. Use this parameter to view grayscale images in false color. If</p>

Parameter	Value
	you specify an empty colormap (<code>[]</code>), <code>imshow</code> ignores this parameter.
'DisplayRange'	<p data-bbox="590 413 1329 508">Two-element vector [<code>LOW HIGH</code>] that controls the display range of a grayscale image. See the <code>imshow(I,[low high])</code> syntax for more details about how to set this parameter.</p> <hr data-bbox="590 564 1329 567"/> <p data-bbox="590 574 1243 800">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 <code>'DisplayRange'</code> parameter must be specified when calling <code>imshow</code> with a filename, for example <code>imshow(filename,'DisplayRange'[LOW HIGH])</code>.</p>
'InitialMagnification'	<p data-bbox="590 847 1322 1038">A numeric scalar value, or the text string <code>'fit'</code>, 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 <code>'fit'</code>, <code>imshow</code> scales the entire image to fit in the window.</p> <p data-bbox="590 1055 1329 1211">On initial display, <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.</p> <p data-bbox="590 1229 1191 1324">By default, the initial magnification parameter is set to the value returned by <code>iptgetpref('ImshowInitialMagnification')</code>.</p> <p data-bbox="590 1341 1329 1437">If the image is displayed in a figure with its <code>'WindowStyle'</code> property set to <code>'docked'</code>, <code>imshow</code> warns and displays the image at the largest magnification that fits in the figure.</p>

Parameter	Value
	<p>Note: If you specify the axes position (using <code>subplot</code> or <code>axes</code>), <code>imshow</code> ignores any initial magnification you might have specified and defaults to the 'fit' behavior.</p> <p>When used with the 'Reduce' parameter, only 'fit' is allowed as an initial magnification.</p>
'Parent'	Handle of an axes that specifies the parent of the image object that will be created by <code>imshow</code> .
'Reduce'	Logical value that specifies whether <code>imshow</code> subsamples the image in <code>filename</code> . The 'Reduce' parameter is only valid for TIFF images and you must specify a file name. Use this parameter to display overviews of very large images.
'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.

Note If you are building a GUI where you want to control the figure and axes properties, be sure to use the `imshow(..., 'Parent', ax)` syntax.

Class Support

A truecolor image can be `uint8`, `uint16`, `single`, or `double`. An indexed image can be `logical`, `uint8`, `single`, or `double`. A grayscale image can be `logical`, `uint8`, `int16`, `uint16`, `single`, or `double`. A binary image must be of class `logical`.

For grayscale images of class `single` or `double`, the default display range is `[0 1]`. If your image's data range is much larger or smaller than the default display range, you might need to experiment with setting the display range to see features in the image that would not be visible using the default display range. For all grayscale images having integer types, the default display range is `[intmin(class(I)) intmax(class(I))]`.

If your image is `int8`, `int16`, `uint32`, `int32` 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 `iptprefs`.

Tips

`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.

The `imshow` function is not supported when MATLAB is started with the `-nojvm` option.

imshow

You can access `imshow` through the Plot Selector workspace tool, which is represented by this icon:  Select data to plot. In your workspace, select the data you want to display. The Plot Selector icon changes to look like this:  `plot(x,y)`. Scroll down to **Image Processing Toolbox Plots**. Select **imshow(I)**.

Examples

Display an image from a file.

```
imshow('board.tif')
```

Display an indexed image.

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

Display a grayscale image.

```
I = imread('cameraman.tif');  
imshow(I)
```

Display the same grayscale image, adjusting the display range.

```
h = imshow(I,[0 80]);
```

See Also

`imread` | `imtool` | `iptprefs` | `subimage` | `trueimage` | `warp` | `image`
| `imagesc`

Purpose

Compare differences between images

Syntax

```
imshowpair(A,B)  
imshowpair(A,B,method)  
imshowpair( ___,Name,Value)
```

Description

`imshowpair(A,B)` displays the differences between images A and B. If A and B are different sizes, the smaller dimensions are padded with zeros on the bottom and right edges so that the two image are the same size in the display.

`imshowpair(A,B,method)` uses the visualization method specified by `method`.

`imshowpair(___,Name,Value)` specifies additional options with one or more `Name,Value` pair arguments, using any of the previous syntaxes.

Input Arguments

A,B - Input images

grayscale images | truecolor images | binary images

Input images to be combined into a composite image, specified as any combination of grayscale, truecolor, or binary images.

Data Types

single | double | int8 | int16 | int32 | int64 | uint16 | uint32 | uint64
| logical

method - Visualization method

'falsecolor' (default) | 'blend' | 'diff' | 'montage'

Visualization method for creating the composite image, specified as one of the text strings in the following table.

Method	Description
'falsecolor'	Creates a composite RGB image showing A and B overlaid in different color bands. Gray regions in the composite image show where the two images have the same intensities. Magenta and green regions show where the intensities are different. This is the default method.
'blend'	Overlays A and B using alpha blending.
'diff'	Creates a difference image from A and B.
'montage'	Places A and B next to each other in the same image.

Example: `imshowpair(A,B,'montage')` displays A and B next to each other.

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

Example: `'Scaling','joint'` scales the intensity values of A and B together as a single data set.

Scaling - Image data scaling option

`'independent'` (default) | `'joint'` | `'none'`

Image data scaling option that controls the intensity scaling of the pixels representing the input images in the output image, specified as the comma-separated pair consisting of `'Scaling'` and one of the text strings in the following table.

'independent'	Scales the intensity values of A and B independently from each other.
'joint'	Scales the intensity values in the images jointly as if they were together in the same image. This option is useful when you want to visualize registrations of monomodal images, where one image contains fill values that are outside the dynamic range of the other image.
'none'	No additional scaling.

ColorChannels - Output color channel for each input image

[R G B] | 'red-cyan' | 'green-magenta' (default)

Output color channel for each input image, specified as the comma-separated pair consisting of 'ColorChannels' and one of the text strings from the following table. 'ColorChannels' can only be used with the 'falsecolor' visualization method.

[R G B]	A three element vector that specifies which image to assign to the red, green, and blue channels. The R, G, and B values must be 1 (for the first input image), 2 (for the second input image), and 0 (for neither image).
'red-cyan'	A shortcut for the vector [1 2 2], which is suitable for red/cyan stereo anaglyphs.
'green-magenta'	A shortcut for the vector [1 2 2], which is a high contrast option, ideal for people with many kinds of color blindness.

Tips

- Use `imfuse` to create composite visualizations that you can save to a file. Use `imshowpair` to display composite visualizations to the screen.

Examples

Display Two Images That Differ by Rotation and Offset

Display a pair of grayscale images with two different visualization methods, 'diff' and 'blend'.

Load an image into the workspace. Create a copy with a rotation and offset applied.

```
A = imread('cameraman.tif');  
B = imrotate(A,5,'bicubic','crop');
```

Display the difference of A and B.

```
imshowpair(A,B,'diff');
```



Display a blended overlay A and B.

```
figure;  
imshowpair(A,B,'blend','Scaling','joint');
```



Display Two Images with Different Brightness Ranges


Load an image into the workspace. Create a copy and apply rotation, offset, and a brightness adjustment.

```
A = dicomread('CT-MONO2-16-ankle.dcm');  
B = double(imrotate(A,10,'bicubic','crop'));  
B = B * 0.2;
```

Display the images with the default, 'falsecolor', method and apply brightness scaling independently to each image.

```
figure;  
imshowpair(A,B,'Scaling','independent');
```

imshowpair



See Also

`imfuse` | `imregister` | `imshow` | `imtransform`

Purpose Subtract one image from another or subtract constant from 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, elements of the output that exceed the range of the integer type are truncated, and fractional values are rounded.

Note On Intel architecture processors, `imsubtract` can take advantage of the Intel Integrated Performance Primitives (Intel IPP) library, thus accelerating its execution time. The (Intel IPP) library 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 = imsubtract(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,[])
```

imsubtract

Subtract a constant value from an image:

```
I = imread('rice.png');  
Iq = imsubtract(I,50);  
figure, imshow(I), figure, imshow(Iq)
```

See Also

[imabsdiff](#) | [imadd](#) | [imcomplement](#) | [imdivide](#) | [imlincomb](#) |
[immultiply](#) | [ippl](#)

Purpose

Image Tool

Syntax

```
imtool
imtool(I)
imtool(I,[low high])
imtool(RGB)
imtool(BW)
imtool(X,map)
imtool(filename)
hfigure = 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 grayscale image `I`.

`imtool(I,[low high])` displays the grayscale image `I`, specifying the display range for `I` in the vector `[low high]`. 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`. Pixel values of 0 display as black; pixel values of 1 display 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` or a reduced resolution dataset (R-Set) created by `rsetwrite`. If the file contains multiple images, the first one is

displayed. The file must be in the current directory or on the MATLAB path.

`hfigure = imtool(...)` returns `hfigure`, a handle to the figure created by `imtool`. `close(Hfigure)` 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
'Colormap'	2-D, real, m -by-3 matrix specifying the colormap to use for the figure's colormap property. Use this parameter to view grayscale images in false color. If you specify an empty colormap (<code>[]</code>), <code>imtool</code> ignores this parameter.
'DisplayRange'	Two-element vector [<code>LOW HIGH</code>] that controls the display range of a grayscale image. See the <code>imtool(I,[low high])</code> syntax for more details about how to set this parameter. 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.

Parameter	Value
	<p>When set to 'adaptive', the entire image is visible on initial display. If the image is too large to display on the screen, imtool displays the image at the largest magnification that fits on the screen.</p> <p>When set to 'fit', imtool scales the entire image to fit in the window.</p> <p>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).</p> <hr/> <p>Note When the image aspect ratio is such that less than one pixel would be displayed in either dimension at the requested magnification, imtool issues a warning and displays the image at 100%.</p> <hr/> <p>By default, the initial magnification parameter is set to the value returned by <code>iptgetpref('ImtoolInitialMagnification')</code>.</p>

Class Support

A truecolor image can be `uint8`, `uint16`, `single`, or `double`. An indexed image can be `logical`, `uint8`, `single`, or `double`. A grayscale image can be `uint8`, `uint16`, `int16`, `single`, or `double`. A binary image must be `logical`. A binary image is of class `logical`.

For all grayscale images having integer types, the default display range is `[intmin(class(I)) intmax(class(I))]`.

For grayscale images of class `single` or `double`, the default display range is `[0 1]`. If the data range of a `single` or `double` image is much larger or smaller than the default display range, you might need to

experiment with setting the display range to see features in the image that would not be visible using the default display range.

Large Data Support

To view very large TIFF or NITF images that will not fit into memory, you can use `rsetwrite` to create a reduced resolution dataset (R-Set) viewable in `imtool`. R-Sets can also improve performance of `imtool` for large images that fit in memory.

The following tools can be used with an R-Set: Overview, Zoom, Pan, Image Information, and Distance. Other tools, however, will not work with an R-Set. You cannot use the Pixel Region, Adjust Contrast, Crop Image, and Window/Level tools. Please note that the Pixel Information tool displays only the x and y coordinates of a pixel and not the associated intensity, index, or [R G B] values.

Related Toolbox Preferences

You can use the Image Processing Preferences dialog box to set toolbox preferences that modify the behavior of `imtool`. To access the dialog, select **File > Preferences** in the MATLAB desktop or Image Tool menu. Also, you can set preferences programmatically with `iptsetpref`. The `imtool` preferences include:

- 'ImtoolInitialMagnification' controls the initial magnification for image display. To override this toolbox preference, specify the 'InitialMagnification' parameter when you call `imtool`, as follows:



```
imtool(...,'InitialMagnification',initial_mag).
```

- 'ImtoolStartWithOverview' controls whether the Overview tool opens automatically when you open an image using the Image Tool. Possible values:
 - `true`— Overview tool opens when you open an image.
 - `{false}`— Overview tool does not open when you open an image. This is the default behavior.

For more information about these preferences, see `iptprefs`.

Tips

`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.

You can access `imtool` through the Plot Selector workspace tool, which is represented by this icon:  Select data to plot. In your workspace, select the data you want to display. The Plot Selector icon changes to look like this:  `plot(x,y)`. Scroll down to **Image Processing Toolbox Plots**. Select **`imtool(I)`**.

Examples

Display an image from a file.

```
imtool('board.tif')
```

Display an indexed image.

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

Display a grayscale image.

```
I = imread('cameraman.tif');  
imtool(I)
```

Display a grayscale image, adjusting the display range.

```
h = imtool(I,[0 80]);  
close(h)
```

See Also

`imageinfo` | `imcontrast` | `imoverview` | `impixelregion` | `imread` | `imshow` | `iptprefs` | `rsetwrite`

imtophat

Purpose 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`. Top-hat filtering computes the morphological opening of the image (using `imopen`) and then subtracts the result from the original image. `imtophat` uses 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 0s and 1s 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.

Examples 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)
```




- 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(...,Name,Value)
```

Description

`B = imtransform(A,tform)` transforms the image `A` according to the 2-D spatial transformation defined by `tform`. If `ndims(A) > 2`, such as for an RGB image, then `imtransform` applies the same 2-D transformation to all 2-D planes along the higher dimensions.

`B = imtransform(A,tform,interp)` specifies the form of interpolation to use.

`[B,xdata,ydata] = imtransform(...)` returns the location of the output image `B` in the output X-Y space. By default, `imtransform` calculates `xdata` and `ydata` automatically so that `B` contains the entire transformed image `A`. However, you can override this automatic calculation by specifying values for the 'XData' and 'YData' arguments.

`[B,xdata,ydata] = imtransform(...,Name,Value)` transforms the image with additional options for controlling various aspects of the spatial transformation specified by one or more `Name,Value` pair arguments.

Tips

- **Image Registration.** The `imtransform` function automatically shifts the origin of your output image to make as much of the transformed image visible as possible. If you use `imtransform` to do image registration, the syntax `B = imtransform(A,tform)` can produce unexpected results. To control the spatial location of the output image, set 'XData' and 'YData' explicitly.
- **Pure Translation.** Calling the `imtransform` function with a purely translational transformation, results in an output image that is exactly like the input image unless you specify 'XData' and 'YData' values in your call to `imtransform`. For example, if you want the

output to be the same size as the input revealing the translation relative to the input image, call `imtransform` as shown in the following syntax:

```
B = imtransform(A,T,'XData',[1 size(A,2)],...  
              'YData',[1 size(A,1)])
```

For more information about this topic, see “Translate Image Using `maketform` and `imtransform`” in the User’s Guide, especially the section “Understanding the Coordinates of the Transformed Image”.

- **Transformation Speed.** When you do not specify the output-space location for `B` using `'XData'` and `'YData'`, `imtransform` estimates the location automatically using the function `findbounds`. You can use `findbounds` as a quick forward-mapping option for some commonly used transformations, such as affine or projective. For transformations that do not have a forward mapping, such as the polynomial ones computed by `cp2tform`, `findbounds` can take much longer. If you can specify `'XData'` and `'YData'` directly for such transformations, `imtransform` may run noticeably faster.
- **Clipping.** The automatic estimate of `'XData'` and `'YData'` using `findbounds` sometimes clips the output image. To avoid clipping, set `'XData'` and `'YData'` directly.
- **Arbitrary Dimensional Transformations.** Use a 2-D transformation for `tform` when using `imtransform`. For arbitrary-dimensional array transformations, see `tformarray`.

Input Arguments

A

An image of any nonsparse numeric class (real or complex) or of class `logical`.

tform

A spatial transformation structure returned by `maketform` or `cp2tform`. `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. This convention is the reverse of the array subscripting convention in MATLAB.

interp

A string that specifies the form of interpolation to use. *interp* can be one of the following strings: 'bicubic', 'bilinear', or 'nearest' (nearest-neighbor). Alternatively, *interp* can be a resampler structure returned by *makeresampler*. This option allows more control over how *imtransform* performs resampling.

Default: 'bilinear'

Name-Value Pair Arguments

Optional comma-separated pairs of *Name*, *Value* arguments, where *Name* is the argument name and *Value* is the corresponding value. *Name* must appear within single quotes (' ') and is not case sensitive. You can specify several name and value pair arguments in any order as *Name1*, *Value1*, ..., *NameN*, *ValueN*.

UData

A two-element, real vector that, when combined with 'VData', specifies the spatial location of image *A* in the 2-D input space U-V. The two elements of 'UData' give the u -coordinates (horizontal) of the first and last columns of *A*, respectively.

Default: [1 size(A,2)]

VData

A two-element, real vector that, when combined with 'UData', specifies the spatial location of image *A* in the 2-D input space U-V. The two elements of 'VData' give the v -coordinates (vertical) of the first and last rows of *A*, respectively.

Default: [1 size(A,1)]

XData

A two-element, real vector that, when combined with 'YData', specifies the spatial location of the output image B in the 2-D output space X-Y. The two elements of 'XData' give the *x*-coordinates (horizontal) of the first and last columns of B, respectively.

Default: If you do not specify 'XData' and 'YData', `imtransform` estimates values that contain the entire transformed output image. To determine these values, `imtransform` uses the `findbounds` function.

YData

A two-element real vector that, when combined with 'XData', specifies the spatial location of the output image B in the 2-D output space X-Y. The two elements of 'YData' give the *y*-coordinates (vertical) of the first and last rows of B, respectively.

Default: If you do not specify 'XData' and 'YData', `imtransform` estimates values that contain the entire transformed output image. To determine these values, `imtransform` uses the `findbounds` function.

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 specifies both width and height.

Default: If you do not specify 'XYScale' but you do specify 'Size', then `imtransform` calculates 'XYScale' from 'Size', 'XData', and 'YData'. If you do not provide 'XYScale' or 'Size', then `imtransform` uses the scale of the input pixels for 'XYScale', except in cases where an excessively large output image would result.

Note In cases where preserving the scale of the input image would result in an excessively large output image, the `imtransform` function automatically increases the `'XYScale'`. To ensure that the output pixel scale matches the input pixel scale, specify the `'XYScale'` parameter. For example, call `imtransform` as shown in the following syntax:

```
B = imtransform(A,T,'XYScale',1)
```

Size

A two-element vector of nonnegative integers that specifies the number of rows and columns of the output image `B`. For higher dimensions, `imtransform` takes the size of `B` directly from the size of `A`. Thus, `size(B,k)` equals `size(A,k)` for `k > 2`.

Default: If you do not specify `'Size'`, `imtransform` derives this value from `'XData'`, `'YData'`, and `'XYScale'`.

FillValues

An array containing one or several fill values. The `imtransform` function uses fill values 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'` requires a scalar. However, if `A`'s dimension is greater than two, then you can specify `'FillValues'` as an array whose size satisfies the following constraint: `size(fill_values,k)` must equal either `size(A,k+2)` or 1.

For example, if `A` is a `uint8` RGB image that is 200-by-200-by-3, then possibilities for `'FillValues'` include the following values.

Value	Fill
0	Fill with black
[0;0;0]	Fill with black

Value	Fill
255	Fill with white
[255;255;255]	Fill with white
[0;0;255]	Fill with blue
[255;255;0]	Fill with yellow

If A is 4-D with size 200-by-200-by-3-by-10, then you can specify 'FillValues' as a scalar, 1-by-10, 3-by-1, or 3-by-10.

Output Arguments

B

Output image of any nonsparse numeric class (real or complex) or of class logical.

xdata

Two-element vector that specifies the *x*-coordinates of the first and last columns of B.

Note Sometimes the output values *xdata* and *ydata* do not exactly equal the input 'XData' and 'YData' arguments. The values differ either because of the need for an integer number of rows and columns, or because 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* can be different.

ydata

Two-element vector that specifies the *y*-coordinates of the first and last rows of B.

Examples

Simple Transformation. 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)
```



Horizontal Shear

Projective Transformation. Map a square to a quadrilateral with a projective transformation:

```
% Set up an input coordinate system so that the input image
% fills the unit square with vertices (0 0),(1 0),(1 1),(0 1).
I = imread('cameraman.tif');
udata = [0 1]; vdata = [0 1];

% Transform to a quadrilateral with vertices (-4 2),(-8 3),
% (-3 -5),(6 3).
tform = maketform('projective',[ 0 0; 1 0; 1 1; 0 1],...
                  [-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.
```

imtransform

```
[B,xdata,ydata] = imtransform(I, tform, 'bicubic', ...
                              'udata', udata,...
                              'vdata', vdata,...
                              'size', size(I),...
                              'fill', 128);
subplot(1,2,1), imshow(I, 'XData', udata, 'YData', vdata), ...
    axis on
subplot(1,2,2), imshow(B, 'XData', xdata, 'YData', ydata), ...
    axis on
```



Projective Transformation

Image Registration. Register an aerial photo to an orthophoto:

```
% Read in the aerial photo.
unregistered = imread('westconcordaerial.png');
figure, imshow(unregistered)
```




Aerial Photo

```
% Read in the orthophoto.  
figure, imshow('westconcordorthophoto.png')
```



Orthophoto

```
% Load control points that were previously picked.  
load westconcordpoints
```

imtransform

```
% 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)
```



Transformed Image

See Also

[checkerboard](#) | [cp2tform](#) | [imresize](#) | [imrotate](#) | [maketform](#) | [makesampler](#) | [tformarray](#)

Tutorials

- “Translate Image Using maketform and imtransform”
- Exploring Slices from a 3-Dimensional MRI Data Set
- Padding and Shearing an Image Simultaneously

Purpose

Display image in image tool

Note `imview` has been removed. Use `imtool` instead.

ind2gray

Purpose Convert indexed image to grayscale image

Syntax `I = ind2gray(X,map)`

Description `I = ind2gray(X,map)` converts the image `X` with colormap `map` to a grayscale image `I`. `ind2gray` removes the hue and saturation information from the input image while retaining the luminance.

Note A grayscale image is also called a gray-scale, gray scale, or gray-level image.

Class Support `X` can be of class `uint8`, `uint16`, `single`, or `double`. `map` is `double`. `I` is of the same class as `X`.

Examples

```
load trees
I = ind2gray(X,map);
imshow(X,map)
figure,imshow(I)
```



Image Courtesy of Susan Cohen

Algorithms `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` | `mat2gray` | `rgb2gray` | `rgb2ntsc`

ind2rgb

Purpose	Convert indexed image to 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 Read metadata from Interfile file

Syntax `info = interfileinfo(filename)`

Description `info = interfileinfo(filename)` returns a structure whose fields contain information about an image in a Interfile file. `filename` is a string that specifies the name of the file. The file must be in the current directory or in a directory on the MATLAB path.

The Interfile file format was developed for the exchange of nuclear medicine data. In Interfile 3.3, metadata is stored in a header file, separate from the image data. The two files have the same name with different file extensions. The header file has the file extension `.hdr` and the image file has the file extension `.img`.

Examples Read metadata from an Interfile file.

```
info = interfileinfo('MyFile.hdr');
```

For more information about this file format, visit the Interfile Archive, maintained by the Department of Medical Physics and Bioengineering, University College, London, UK.

See Also `interfileread`

interfileread

Purpose Read images in Interfile format

Syntax
`A = interfileread(filename)`
`A = interfileread(filename, window)`

Description `A = interfileread(filename)` reads the images in the first energy window of `filename` into `A`, where `A` is an M-by-N array for a single image and an M-by-N-by-P array for multiple images. The file must be in the current directory or in a directory on the MATLAB path.

`A = interfileread(filename, window)` reads the images in the energy window specified by `window` of `filename` into `A`.

The images in the energy window must be of the same size.

Examples Read image data from an Interfile file.

```
img = interfileread('MyFile.hdr');
```

For more information about this file format, visit the [Interfile Archive](#), maintained by the Department of Medical Physics and Bioengineering, University College, London, UK.

See also `interfileinfo`

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> .
Examples	<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>

Purpose Check for presence of Intel Integrated Performance Primitives (Intel IPP) library

Syntax TF = ippl
[TF B] = ippl

Description The Intel Integrated Performance Primitives (Intel IPP) library provides a collection of basic functions used in signal and image processing. The Intel IPP library 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. The Intel IPP library is used by some of the Image Processing Toolbox functions to accelerate their execution time.

TF = ippl returns true (1) if the Intel IPP library is available and false (0) otherwise.

[TF B] = ippl returns an additional column cell array B. The array B contains a string describing the Intel IPP library module.

When the Intel IPP library is available, some of the Image Processing Toolbox arithmetic functions (imabsdiff, imdivide, and immultiply) and the imfilter function take advantage of it. Toolbox functions that use these functions also benefit.

Notes The Intel IPP library is used only for some data types and only under specific conditions. See the help sections of the functions listed above for detailed information on when the Intel IPP library is activated.

You can enable or disable the Intel IPP library by using the Image Processing Preferences dialog box, part of the MATLAB Preferences dialog box. You can open this dialog box by selecting **Preferences** from the MATLAB Desktop **File** menu and then choosing **Image Processing**.

You can also enable or disable the Intel IPP library at the command line. To disable the Intel IPP library, use this command:

```
iptsetpref('UseIPPL', false)
```

To enable IPPL, use this command:

```
iptsetpref('UseIPPL', true)
```

Note that enabling or disabling the Intel IPP library has the effect of clearing all loaded MEX-files.

See Also

`imabsdiff` | `imdivide` | `imfilter` | `immultiply` | `iptprefs`

iptaddcallback

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.

Examples Create a figure and register 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

Purpose	Check validity of connectivity argument
Syntax	<code>iptcheckconn(conn, func_name, var_name, arg_pos)</code>
Description	<p><code>iptcheckconn(conn, func_name, var_name, arg_pos)</code> checks whether <code>conn</code> is a valid connectivity argument. If it is invalid, the function issues a formatted error message.</p> <p>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.</p> <p><code>func_name</code> is a string that specifies the name used in the formatted error message to identify the function checking the connectivity argument.</p> <p><code>var_name</code> is a string that specifies the name used in the formatted error message to identify the argument being checked.</p> <p><code>arg_pos</code> is a positive integer that indicates the position of the argument being checked in the function argument list. <code>iptcheckconn</code> includes this information in the formatted error message.</p>
Class Support	<code>conn</code> must be of class <code>double</code> or <code>logical</code> and must be real and nonsparse.
Examples	<p>Create a 4-by-4 array and pass it as the connectivity argument.</p> <pre>iptcheckconn(eye(4), 'func_name', 'var_name', 2)</pre>

iptcheckhandle

Purpose	Check validity of handle
Syntax	<code>iptcheckhandle(H, valid_types, func_name, var_name, arg_pos)</code>
Description	<p><code>iptcheckhandle(H, valid_types, func_name, var_name, arg_pos)</code> checks the validity of the handle <code>H</code> and issues a formatted error message if the handle is invalid. <code>H</code> must be a handle to a single figure, <code>uipanel</code>, <code>hggroup</code>, <code>axes</code>, or <code>image</code> object.</p> <p><code>valid_types</code> is a cell array of strings specifying the set of Handle Graphics object types to which <code>H</code> is expected to belong. For example, if you specify <code>valid_types</code> as <code>{'uipanel', 'figure'}</code>, <code>H</code> can be a handle to either a <code>uipanel</code> object or a <code>figure</code> object.</p> <p><code>func_name</code> is a string that specifies the name used in the formatted error message to identify the function checking the handle.</p> <p><code>var_name</code> is a string that specifies the name used in the formatted error message to identify the argument being checked.</p> <p><code>arg_pos</code> is a positive integer that indicates the position of the argument being checked in the function argument list. <code>iptcheckhandle</code> converts this value to an ordinal number and includes this information in the formatted error message.</p>

Examples

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.

```
func_name      arg_pos      var_name
  |            |            |
  v            v            v
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](#)

iptcheckinput

Purpose Check validity of array

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>nonempty</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.

Examples

To trigger this error message, create a three-dimensional array and then check for the attribute '2d'.

```
A = [ 1 2 3; 4 5 6 ];
B = [ 7 8 9; 10 11 12];
C = cat(3,A,B);
iptcheckinput(C,{'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
      |              |              |
      v              v              v
Function FUNC_NAME expected its second input, var_name, to be
two-dimensional.
      ^
      |
attributes
```

See Also

```
iptcheckhandle | iptcheckmap | iptchecknargin | iptcheckstrs
| iptnum2ordinal
```

iptcheckmap

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 a formatted 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` includes this information in the formatted error message.

Examples

```
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 input number 2, 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`

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.

Examples

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
```

iptcheckstrs

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.

Examples

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.

See Also

[iptcheckhandle](#) | [iptcheckinput](#) | [iptcheckmap](#) | [iptchecknargin](#)
| [iptnum2ordinal](#)

iptdemos

Purpose	Index of Image Processing Toolbox examples
Syntax	<code>iptdemos</code>
Description	<code>iptdemos</code> displays the HTML page that lists all the Image Processing examples. <code>iptdemos</code> displays the page in the MATLAB Help browser.

Purpose	Get Application Programmer Interface (API) for handle
Syntax	<code>API = iptgetapi(h)</code>
Description	<p><code>API = iptgetapi(h)</code> returns the API structure associated with handle <code>h</code> if there is one. Otherwise, <code>iptgetapi</code> returns an empty array.</p> <p>For more information about handle APIs, see the help for <code>immagbox</code>, <code>impositionrect</code>, or <code>imscrollpanel</code>.</p>
Examples	<pre>hFig = figure('Toolbar','none',... 'Menubar','none'); hIm = imshow('tape.png'); hSP = imscrollpanel(hFig,hIm); api = iptgetapi(hSP); api.setMagnification(2) % 2X = 200%</pre>
See Also	<code>immagbox</code> <code>imrect</code> <code>imscrollpanel</code>

iptGetPointerBehavior

Purpose Retrieve pointer behavior from HG object

Syntax `pointerBehavior = iptGetPointerBehavior(h)`

Description `pointerBehavior = iptGetPointerBehavior(h)` returns the pointer behavior structure associated with the Handle Graphics object `h`. A pointer behavior structure contains function handles that interact with a figure's pointer manager (see `iptPointerManager`) to control what happens when the figure's mouse pointer moves over and then exits the object. See `iptSetPointerBehavior` for details.

If `h` does not contain a pointer behavior structure, `iptGetPointerBehavior` returns `[]`.

See Also `iptPointerManager` | `iptSetPointerBehavior`

Purpose Get values of Image Processing Toolbox preferences

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.

`value = iptgetpref(prefname)` returns the value of the Image Processing Toolbox preference specified by the string `prefname`. See `iptprefs` for a complete list of valid preference names or access the Image Processing preferences dialog box from the **File** menu in the MATLAB desktop. Preference names are not case sensitive and can be abbreviated.

Examples

```
value = iptgetpref('ImshowAxesVisible')  
  
value =  
  
off
```

See Also `imshow` | `iptprefs` | `iptsetpref`

ipticondir

Purpose Directories containing IPT and MATLAB icons

Syntax `[D1, D2] = ipticondir`

Description `[D1, D2] = ipticondir` returns the names of the directories containing the Image Processing Toolbox icons (D1) and the MATLAB icons (D2).

Examples `[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`.

Examples The following example returns the string 'fourth'.

```
str = iptnum2ordinal(4)
```

The following example returns the string '23rd'.

```
str = iptnum2ordinal(23)
```

iptPointerManager

Purpose Create pointer manager in figure

Syntax

```
iptPointerManager(hFigure)
iptPointerManager(hFigure, 'disable')
iptPointerManager(hFigure, 'enable')
```

Description `iptPointerManager(hFigure)` creates a pointer manager in the specified figure. The pointer manager controls pointer behavior for any Handle Graphics objects in the figure that contain pointer behavior structures. Use `iptSetPointerBehavior` to associate a pointer behavior structure with a particular object to define specific actions that occur when the mouse pointer moves over and then leaves the object. See `iptSetPointerBehavior` for more information.

`iptPointerManager(hFigure, 'disable')` disables the figure's pointer manager.

`iptPointerManager(hFigure, 'enable')` enables and updates the figure's pointer manager.

Note If the figure already contains a pointer manager, `iptPointerManager(hFigure)` does not create a new one. It has the same effect as `iptPointerManager(hFigure, 'enable')`.

Tips `iptPointerManager` considers not just the object the pointer is over, but all objects in the figure. `iptPointerManager` searches the HG hierarchy to find the first object that contains a pointer behavior structure. The `iptPointerManager` then executes that object's pointer behavior function. For example, you could set the pointer to be a fleur and associate that pointer with the axes. Then, when you slide the pointer into the figure window, it will initially be the default pointer, then change to a fleur when you cross into the axes, and remain a fleur when you slide over the objects parented to the axes.

Examples

Plot a line. Create a pointer manager in the figure. Then, associate a pointer behavior structure with the line object in the figure that changes the mouse pointer into a fleur whenever the pointer is over it.

```
h = plot(1:10);  
iptPointerManager(gcf);  
enterFcn = @(hFigure, currentPoint)...  
           set(hFigure, 'Pointer', 'fleur');  
iptSetPointerBehavior(h, enterFcn);
```

See Also

[iptGetPointerBehavior](#) | [iptSetPointerBehavior](#)

iptprefs

Purpose Display Image Processing Preferences dialog box

Syntax `iptprefs`

Description `iptprefs` opens the Image Processing Preferences dialog box, part of the MATLAB Preferences dialog box. You can also open this dialog box by selecting **Preferences** from the MATLAB Desktop **File** menu, and then choosing **Image Processing**.

The Image Processing Preferences dialog box contains display preferences for `imtool` and `imshow` and an option for enabling the Intel Integrated Performance Primitives (Intel IPP) library. Also, you can set all preferences at the command line with the `iptsetpref` function. To see how the options in the dialog box correspond to the commands, view the following picture.




Image Processing Preferences Dialog Box

The following table details the available preferences and their syntaxes. Note that preference names are case insensitive and you can abbreviate them. The default value appears enclosed in braces ({}).

Available Image Processing Preferences

Preference Name	Description
'ImtoolStartWithOverview'	<p>Controls whether the Overview tool opens automatically when you open an image using the Image Tool (imtool). Possible values:</p> <p><code>true</code>— Overview tool opens when you open an image.</p> <p><code>{false}</code>— Overview tool does not open when you open an image. This is the default behavior.</p>
'ImtoolInitialMagnification'	<p>Controls the initial magnification of the image displayed by imtool. Possible values:</p> <p><code>{'adaptive'}</code> — 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. A magnification of 100% means that there should be one screen pixel for every image pixel.</p> <p><code>'fit'</code> — 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 imtool.</p>

Available Image Processing Preferences (Continued)

Preference Name	Description
'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>
'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> <hr/> <p>Note There still can 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> <hr/> <p>You can override this preference by specifying the 'Border' parameter when you call <code>imshow</code>.</p>

Available Image Processing Preferences (Continued)

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. A magnification of 100% 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>
'UseIPP'	<p>Controls whether some toolbox functions use the Intel IPP library or not. Possible values:</p> <p>{true} — Enable use of the Intel IPP library.</p> <p>false — Disable use of the Intel IPP library.</p> <p>Note: Setting this preference value clears all loaded MEX-files.</p>

See Also

`imshow` | `imtool` | `iptgetpref` | `iptsetpref`

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.

Examples Register three 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`

iptSetPointerBehavior

Purpose Store pointer behavior structure in Handle Graphics object

Syntax `iptSetPointerBehavior(h, pointerBehavior)`
`iptSetPointerBehavior(h, [])`
`iptSetPointerBehavior(h, enterFcn)`

Description `iptSetPointerBehavior(h, pointerBehavior)` stores the specified pointer behavior structure in the specified Handle Graphics object, `h`. If `h` is an array of objects, `iptSetPointerBehavior` stores the same structure in each object.

When used with a figure's pointer manager (see `iptPointerManager`), a pointer behavior structure controls what happens when the figure's mouse pointer moves over and then exits an object in the figure. For details about this structure, see "Pointer Behavior Structure" on page 3-614.

`iptSetPointerBehavior(h, [])` clears the pointer behavior from the Handle Graphics object or objects.

`iptSetPointerBehavior(h, enterFcn)` creates a pointer behavior structure, setting the `enterFcn` field to the function handle specified, and setting the `traverseFcn` and `exitFcn` fields to `[]`. See "Pointer Behavior Structure" on page 3-614 for details about these fields. This syntax is provided as a convenience because, for most common uses, only the `enterFcn` is necessary.

Pointer Behavior Structure

A pointer behavior structure contains three fields: `enterFcn`, `traverseFcn`, and `exitFcn`. You set the value of these fields to function handles and use the `iptSetPointerBehavior` function to associate this structure with an HG object in a figure. If the figure has a pointer manager installed, the pointer manager calls these functions when the following events occur. If you set a field to `[]`, no action is taken.

Function Handle	When Called
enterFcn	Called when the mouse pointer moves over the object.
traverseFcn	Called once when the mouse pointer moves over the object, and called again each time the mouse moves within the object.
exitFcn	Called when the mouse pointer leaves the object.

When the pointer manager calls the functions you create, it passes two arguments: a handle to the figure and the current position of the pointer.

Examples

Example 1

Change the mouse pointer to a fleur whenever it is over a specific object and restore the original pointer when the mouse pointer moves off the object. The example creates a patch object and associates a pointer behavior structure with the object. Because this scenario requires only an `enterFcn`, the example uses the `iptSetPointerBehavior(n, enterFcn)` syntax. The example then creates a pointer manager in the figure. Note that the pointer manager takes care of restoring the original figure pointer.

```
hPatch = patch([.25 .75 .75 .25 .25],...
               [.25 .25 .75 .75 .25], 'r');
xlim([0 1]);
ylim([0 1]);

enterFcn = @(figHandle, currentPoint)...
    set(figHandle, 'Pointer', 'fleur');
iptSetPointerBehavior(hPatch, enterFcn);
iptPointerManager(gcf);
```

Example 2

Change the appearance of the mouse pointer, depending on where it is within the object. This example sets up the pointer behavior structure, setting the `enterFcn` and `exitFcn` fields to `[]`, and setting `traverseFcn` to a function named `overMe` that handles the position-specific behavior. `overMe` is an example function (in `\toolbox\images\imdemos`) that varies the mouse pointer depending on the location of the mouse within the object. For more information, edit `overMe`.

```
hPatch = patch([.25 .75 .75 .25 .25],...
               [.25 .25 .75 .75 .25], 'r');
xlim([0 1])
ylim([0 1])

pointerBehavior.enterFcn = [];
pointerBehavior.exitFcn = [];
pointerBehavior.traverseFcn = @overMe;

iptSetPointerBehavior(hPatch, pointerBehavior);
iptPointerManager(gcf);
```

Example 3

Change the figure's title when the mouse pointer is over the object. In this scenario, `enterFcn` and `exitFcn` are used to achieve the desired side effect, and `traverseFcn` is `[]`.

```
hPatch = patch([.25 .75 .75 .25 .25],...
               [.25 .25 .75 .75 .25], 'r');
xlim([0 1])
ylim([0 1])

pointerBehavior.enterFcn = ...
    @(figHandle, currentPoint)...
    set(figHandle, 'Name', 'Over patch');
pointerBehavior.exitFcn = ...
    @(figHandle, currentPoint) set(figHandle, 'Name', '');
pointerBehavior.traverseFcn = [];
```

```
iptSetPointerBehavior(hPatch, pointerBehavior);  
iptPointerManager(gcf)
```

See Also

[iptGetPointerBehavior](#) | [iptPointerManager](#)

iptsetpref

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 you change it. You can also use the Image Processing Preferences dialog box to set the preferences. To access the dialog, select **Preferences** from the **File** menu in the MATLAB desktop. For more information about available preferences, see `iptprefs`.

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

See Also `imshow` | `imtool` | `iptgetpref` | `iptprefs`

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.



iptwindowalign

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` constrains the position adjustment of `moving_fig` to keep it entirely visible on the screen.

`iptwindowalign` has no effect if either figure window is docked.


Examples

To illustrate some possible figure window alignments, first create two figures: `fig1` and `fig2`. Initially, `fig2` overlays `fig1` on the screen.

```
fig1 = figure;  
fig2 = figure;
```

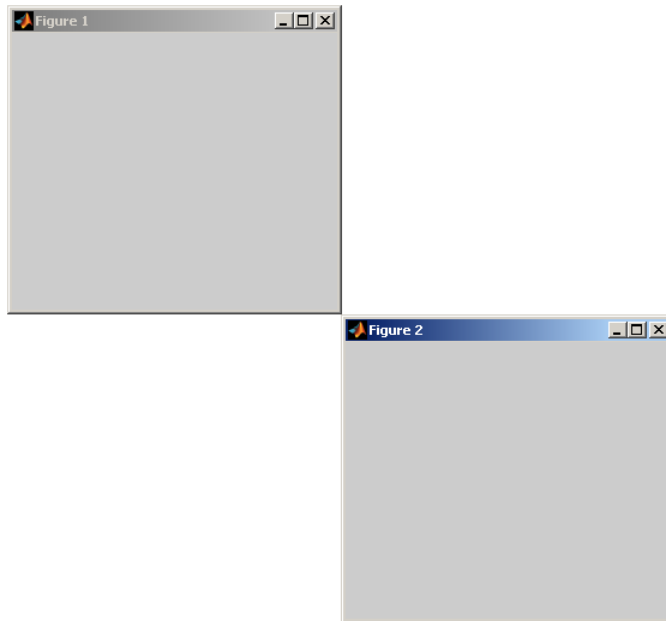
Use `iptwindowalign` to move `fig2` so its left edge is aligned with the right edge of `fig1`.

```
iptwindowalign(fig1, 'right', fig2, 'left');
```



Now move `fig2` so its top edge is aligned with the bottom edge of `fig1`.


```
iptwindowalign(fig1, 'bottom', fig2, 'top');
```



Now move `fig2` so the two figures are centered horizontally.

```
iptwindowalign(fig1, 'hcenter', fig2, 'hcenter');
```

iptwindowalign



See Also

`imtool`

Purpose Inverse Radon transform

Syntax

```
I = iradon(R, theta)
I = iradon(P, 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.

Value	Description
'nearest'	Nearest-neighbor interpolation
'linear'	Linear interpolation (the default)

Value	Description
'spline'	Spline interpolation
'pchip'	Shape-preserving piecewise cubic interpolation
'cubic'	Same as 'pchip'
'v5cubic'	Cubic interpolation from MATLAB 5. This method does not extrapolate, and it issues a warning and uses 'spline' if X is not equally spaced.

`filter` specifies the filter to use for frequency domain filtering. `filter` can be any of the strings that specify standard filters.

Value	Description
'Ram-Lak'	Cropped Ram-Lak or ramp filter. This is the default. 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
'None'	No filtering. When you specify this value, <code>iradon</code> returns unfiltered backprojection data.

`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, `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.

```
output_size = 2*floor(size(R,1)/(2*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`.

Examples

Compare filtered and unfiltered backprojection.

```
P = phantom(128);
R = radon(P,0:179);
I1 = iradon(R,0:179);
I2 = iradon(R,0:179,'linear','none');
subplot(1,3,1), imshow(P), title('Original')
subplot(1,3,2), imshow(I1), title('Filtered backprojection')
subplot(1,3,3), imshow(I2,[]), title('Unfiltered backprojection')
```



Compute the backprojection of a single projection vector. The `iradon` syntax does not allow you to do this directly, because if `theta` is a scalar it is treated as an increment. You can accomplish the task by passing in two copies of the projection vector and then dividing the result by 2.

```
P = phantom(128);  
R = radon(P,0:179);  
r45 = R(:,46);  
I = iradon([r45 r45], [45 45])/2;  
imshow(I, [])  
title('Backprojection from the 45-degree projection')
```

Algorithms

`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.

References

[1] Kak, A. C., and M. Slaney, *Principles of Computerized Tomographic Imaging*, New York, NY, IEEE Press, 1988.

See Also

[fan2para](#) | [fanbeam](#) | [ifanbeam](#) | [para2fan](#) | [phantom](#) | [radon](#)

Purpose True for binary image

Syntax `flag = isbw(A)`

`isbw` has been removed.

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`

isflat

Purpose	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

Purpose True for grayscale image

Syntax `flag = isgray(A)`

`isgray` has been removed.

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 grayscale 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`

isicc

Purpose 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.

Examples Read in an ICC profile and isicc returns True.

```
P = iccread('sRGB.icm');
```

```
TF = isicc(P)
```

```
TF =
```

```
1
```

This example creates a MATLAB structure and uses isicc to test if it's a valid ICC profile. isicc returns False.

```
S.name = 'Any Student';
```

```
S.score = 83;
```

```
S.grade = 'B+'
```

```
TF = isicc(S)
```

```
TF =
```

```
0
```

See Also

`applycform` | `iccread` | `icwrite` | `makecform`

isind

Purpose True for indexed image

Syntax `flag = isind(A)`

`isind` has been removed.

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 Check if file is National Imagery Transmission Format (NITF) file

Syntax `[tf, NITF_version] = isnif(filename)`

Description `[tf, NITF_version] = isnif(filename)` returns True (1) if the file specified by `filename` is a National Imagery Transmission Format (NITF) file, otherwise False (0). If the file is a NITF file, `isnif` returns a text string identifying the NITF version in `NITF_version`, such as '2.1'. If the file is not a NITF file, `NITF_version` contains the text string 'UNK'.

See Also `nitfinfo` | `nitfread`

isrgb

Purpose True for RGB image

Syntax `flag = isrgb(A)`

`isrgb` has been removed.

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`

Purpose Check if file is R-Set

Syntax [tf, supported] = isrset(filename)

Description [tf, supported] = isrset(filename) sets tf to true if the file filename is a reduced resolution dataset (R-Set) created by rsetwrite and false if it is not. The value of supported is true if the R-Set file is compatible with the R-Set tools (such as imtool) in the version of the Image Processing Toolbox you are using. If supported is false, the R-Set file was probably created by a newer version of rsetwrite than the one in the version of the Image Processing Toolbox you are using.

See Also rsetwrite

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 software 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`.

Examples 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 software 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.

Examples 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 software 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.

Examples 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 label matrix into 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 `labelmatrix`, `bwlabel`, `bwlabeln`, 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

Value	Color
'r'	Red
'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.

`label2rgb` supports the generation of efficient, production-quality C/C++ code from MATLAB. For best results, when using the standard syntax:

```
RGB = label2rgb(L, map, zerocolor, order)
```

- Submit at least two input arguments: the label matrix, `L`, and the colormap matrix, `map`.
- `map` must be an `n-by-3`, `double`, colormap matrix. You cannot use a string containing the name of a MATLAB colormap function or a function handle of a colormap function.
- If you set the boundary color `zerocolor` to the same color as one of the regions, `label2rgb` will not issue a warning.
- If you supply a value for `order`, it must be `'noshuffle'`.

To see a complete list of toolbox functions that support code generation, see “Supported Functions”.

Class Support


The input label matrix `L` can have any numeric class. It must contain finite, nonnegative integers. The output of `label2rgb` is of class `uint8`.

label2rgb

Examples

Use `label2rgb` to customize display of `labelmatrix`.

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



Original

```
BW = im2bw(I, graythresh(I));  
CC = bwconncomp(BW);  
L = labelmatrix(CC);  
RGB = label2rgb(L);  
figure, imshow(RGB)
```



RGB

```
RGB2 = label2rgb(L, 'spring', 'c', 'shuffle');  
figure, imshow(RGB2)
```



RGB2

See Also

`bwconncomp` | `bwlabel` | `colormap` | `ismember` | `labelmatrix` | `watershed`

labelmatrix

Purpose Create label matrix from bwconncomp structure

Syntax `L = labelmatrix(CC)`

Description `L = labelmatrix(CC)` creates a label matrix from the connected components structure `CC` returned by `bwconncomp`. The size of `L` is `CC.ImageSize`. 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 class of `L` depends on `CC.NumObjects`, as shown in the following table.

Class	Range
'uint8'	$CC.NumObjects \leq 255$
'uint16'	$256 \leq CC.NumObjects \leq 65535$
'uint32'	$65536 \leq CC.NumObjects \leq 2^{32} - 1$
'double'	$CC.NumObjects \geq 2^{32}$

`labelmatrix` is more memory efficient than `bwlabel` and `bwlabeln` because it returns its label matrix in the smallest numeric class necessary for the number of objects.

Class Support `CC` is a structure returned by `bwconncomp`. The label matrix `L` is `uint8`, `uint16`, `uint32`, or `double`.

Examples Calculate the connected components and display results:

```
BW = imread('text.png');  
CC = bwconncomp(BW);  
L = labelmatrix(CC);  
L2 = bwlabel(BW);
```

```
whos L L2
```

The output for `whos` appears as follows:

Name	Size	Bytes	Class	Attributes
L	256x256	65536	uint8	
L2	256x256	524288	double	

```
figure, imshow(label2rgb(L));
```

The output for `imshow` appears as follows:

**The term watershed
refers to a ridge that ...**

**... divides areas
drained by different
river systems.**

See Also

[bwconncomp](#) | [bwlabel](#) | [bwlabeln](#) | [label2rgb](#) | [regionprops](#)

makecform

Purpose

Create color transformation structure

Syntax

```
C = makecform(type)
C = makecform(type, 'WhitePoint', WP)
C = makecform(type, 'AdaptedWhitePoint', WP)
C = makecform('srgb2cmyk', 'RenderingIntent', intent)
C = makecform('cmyk2srgb', 'RenderingIntent', intent)
C = makecform('adapt', 'WhiteStart', WPS, 'WhiteEnd', WPE,
    'AdaptModel', modelName)
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)
C = makecform('mattrc', profile, 'Direction', direction)
C = makecform('mattrc', profile, 'Direction', direction,
    'RenderingIntent', intent)
C = makecform('graytrc', profile, 'Direction', direction)
C = makecform('graytrc', profile, 'Direction', direction,
    'RenderingIntent', intent)
C = makecform('named', profile, space)
```

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` also supports conversions to and from the *sRGB* and *CMYK* color spaces. For a list of the abbreviations used by the Image Processing Toolbox software for each color space, see the Remarks section of this reference page.

Type	Description
'cmyk2srgb'	Convert from the <i>CMYK</i> color space to the <i>sRGB</i> color space.
'lab2lch'	Convert from the $L^*a^*b^*$ to the L^*ch color space.
'lab2srgb'	Convert from the $L^*a^*b^*$ to the <i>sRGB</i> color space.
'lab2xyz'	Convert from the $L^*a^*b^*$ to the <i>XYZ</i> color space.
'lch2lab'	Convert from the L^*ch to the $L^*a^*b^*$ color space.
'srgb2cmyk'	Convert from the <i>sRGB</i> to the <i>CMYK</i> color space.
'srgb2lab'	Convert from the <i>sRGB</i> to the $L^*a^*b^*$ color space.
'srgb2xyz'	Convert from the <i>sRGB</i> to the <i>XYZ</i> color space.
'upvpl2xyz'	Convert from the $u'v'L$ to the <i>XYZ</i> color space.
'uvl2xyz'	Convert from the uvL to the <i>XYZ</i> color space.
'xyl2xyz'	Convert from the xyY to the <i>XYZ</i> color space.
'xyz2lab'	Convert from the <i>XYZ</i> to the $L^*a^*b^*$ color space.
'xyz2srgb'	Convert from the <i>XYZ</i> to the <i>sRGB</i> color space.
'xyz2upvpl'	Convert from the <i>XYZ</i> to the $u'v'L$ color space.
'xyz2uvl'	Convert from the <i>XYZ</i> to the uvL color space.
'xyz2xyl'	Convert from the <i>XYZ</i> to the xyY color space.

`C = makecform(type, 'WhitePoint', WP)` specifies the value of the reference white point. *type* can be either 'xyz2lab' or 'lab2xyz'. *WP* is a 1-by-3 vector of *XYZ* values scaled so that $Y = 1$. The default is `whitepoint('ICC')`. Use the `whitepoint` function to create the *WP* vector.

`C = makecform(type, 'AdaptedWhitePoint', WP)` specifies the adapted white point. *type* can be either 'srgb2lab', 'lab2srgb', 'srgb2xyz', or 'xyz2srgb'. As above, *WP* is a row vector of *XYZ* values scaled so that $Y = 1$. If not specified, the default adapted white point is `whitepoint('ICC')`. To get answers consistent with some published *sRGB* equations, specify `whitepoint('D65')` for the adapted white point.

`C = makecform('srgb2cmyk', 'RenderingIntent', intent)` and `C = makecform('cmyk2srgb', 'RenderingIntent', intent)` specify the rendering intent for transforms of type `srgb2cmyk` and `cmyk2srgb`. These transforms convert data between *sRGB* IEC61966-2.1 and "Specifications for Web Offset Publications" (SWOP) *CMYK*. *intent* must be one of these strings: 'AbsoluteColorimetric', 'Perceptual', 'RelativeColorimetric', or 'Saturation'. For more information, see the table Rendering Intent on page 3-651.

`C = makecform('adapt', 'WhiteStart', WPS, 'WhiteEnd', WPE, 'AdaptModel', modelName)` creates a linear chromatic-adaptation transform. *WPS* and *WPE* are row vectors of *XYZ* values, scaled so that $Y = 1$, specifying the starting and ending white points. *modelName* is either 'vonKries' or 'Bradford' and specifies the type of chromatic-adaptation model to be employed. If 'AdaptModel' is not specified, it defaults to 'Bradford'.

`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 intents are listed below. Each rendering intent has distinct aesthetic and color-accuracy trade-offs.

Rendering Intent

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.

Rendering Intent (Continued)

Value	Description
'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. Each `LUTtype` listed in the table below contains the components of an 8-bit or 16-bit LUTtag that performs a transformation between device colors and PCS colors using a particular rendering. For more information about 'clut' transformations, see Section 6.5.7 of the International Color Consortium specification ICC.1:2001-04 (Version 2) or Section 6.5.9 of ICC.1:2001-12 (Version 4), available at www.color.org.

LUT Type	Description
'AToB0' (default)	Device to PCS: perceptual rendering intent
'AToB1'	Device to PCS: media-relative colorimetric rendering intent
'AToB2'	Device to PCS: saturation rendering intent
'AToB3'	Device to PCS: ICC-absolute rendering intent

LUT Type	Description
'BToA0'	PCS to device: perceptual rendering intent
'BToA1'	PCS to device: media-relative colorimetric rendering intent
'BToA2'	PCS to device: saturation rendering intent
'BToA3'	PCS to device: ICC-absolute rendering intent
'Gamut'	Determines which PCS colors are out of gamut for a given device
'Preview0'	PCS colors to the PCS colors available for soft proofing using the perceptual rendering
'Preview1'	PCS colors available for soft proofing using the media-relative colorimetric rendering.
'Preview2'	PCS colors to the PCS colors available for soft proofing using the saturation rendering.

`C = makeform('mattrc', MatTrc, 'Direction', direction)` creates the color transformation structure `C` based on a Matrix/Tone Reproduction Curve (`MatTRC`) model, containing an *RGB-to-XYZ* matrix and *RGB* Tone Reproduction Curves. `MatTRC` is typically the '`MatTRC`' field of an ICC profile structure returned by `iccread`, based on tags 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 (*RGB to XYZ*) or inverse (*XYZ to RGB*) direction. For more information, see section 6.3.1.2 of the International Color Consortium specification ICC.1:2001-04 or ICC.1:2001-12, available at www.color.org.

`C = makeform('mattrc', profile, 'Direction', direction)` creates a color transform based on the `MatTRC` field of the given ICC profile structure `profile`. `direction` is either '`forward`' or '`inverse`' and specifies whether the `MatTRC` is applied in the forward (*RGB to XYZ*) or inverse (*XYZ to RGB*) direction.

`C = makecform('mattrc', profile, 'Direction', direction, 'RenderingIntent', intent)` is similar, but adds the option of specifying the rendering intent. *intent* must be either 'RelativeColorimetric' (the default) or 'AbsoluteColorimetric'. When 'AbsoluteColorimetric' is specified, the colorimetry is referenced to a perfect diffuser, rather than to the Media White Point of the profile.

`C = makecform('graytrc', profile, 'Direction', direction)` creates the color transformation structure `C` that specifies a monochrome transform based on a single-channel Tone Reproduction Curve (GrayTRC) contained in an ICC color profile. *direction* can be either 'forward' or 'inverse' and specifies whether the grayTRC transform is to be applied in the forward (device to PCS) or inverse (PCS to device) direction. ("Device" here refers to the grayscale signal communicating with the monochrome device. "PCS" is the Profile Connection Space of the ICC profile and can be either XYZ or $L^*a^*b^*$, depending on the 'ConnectionSpace' field in profile.Header.)

`C = makecform('graytrc', profile, 'Direction', direction, 'RenderingIntent', intent)` is similar but adds the option of specifying the rendering intent. *intent* must be either 'RelativeColorimetric' (the default) or 'AbsoluteColorimetric'. When 'AbsoluteColorimetric' is specified, the colorimetry is referenced to a perfect diffuser, rather than to the Media White Point of the profile.

`C = makecform('named', profile, space)` creates the color transformation structure `C` that specifies the transformation from color names to color-space coordinates. *profile* must be a profile structure for a Named Color profile (with a NamedColor2 field). *space* is either 'PCS' or 'Device'. The 'PCS' option is always available and will return $L^*a^*b^*$ or XYZ coordinates, depending on the 'ConnectionSpace' field in profile.Header, in 'double' format. The 'Device' option, when active, returns device coordinates, the dimension depending on the 'ColorSpace' field in profile.Header, also in 'double' format.

Tips

The Image Processing Toolbox software 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), where x and y refer to the xy -coordinates of the associated CIE chromaticity diagram, and 1 refers to Y (luminance).
uv1	1960 CIE uvY values, where u and v refer to the uv -coordinates, and 1 refers to Y (luminance).
upvp1	1976 CIE $u'v'Y$ values, where up and vp refer to the $u'v'$ -coordinates and 1 refers to Y (luminance).
lab	1976 CIE $L^*a^*b^*$ values Note 1 or L refers to L^* (CIE 1976 psychometric lightness) rather than luminance (Y).
lch	Polar transformation of CIE $L^*a^*b^*$ values, where c = chroma and h = hue
cmyk	Standard values used by printers
srgb	Standard computer monitor RGB values, (IEC 61966-2-1)

Examples

Convert RGB image to $L^*a^*b^*$, assuming input image is *sRGB*.

```
rgb = imread('peppers.png');
cform = makecform('srgb2lab');
lab = applycform(rgb,cform);
```

makecform

Convert from a non-standard RGB color profile to the device-independent XYZ profile connection space. Note that the ICC input profile must include a MatTRC value.

```
InputProfile = iccread('myRGB.icc');  
C = makecform('matttrc',InputProfile.MatTRC, ...  
             'direction', 'forward');
```

See Also

[applycform](#) | [iccread](#) | [iccwrite](#) | [isicc](#) | [lab2double](#) | [lab2uint8](#)
| [lab2uint16](#) | [whitepoint](#) | [xyz2double](#) | [xyz2uint16](#)

How To

- “Converting Color Data Between Color Spaces”

Purpose

Create rectangularly bounded drag constraint function

Syntax

```
fcn = makeConstrainToRectFcn(type, xlim, ylim)
```

Description

`fcn = makeConstrainToRectFcn(type, xlim, ylim)` creates a position constraint function for draggable tools of a given type, where `type` is one of the following strings: 'imellipse', 'imfreehand', 'imline', 'impoint', 'impoly', or 'imrect'. The rectangular boundaries of the position constraint function are described by the vectors `xlim` and `ylim` where `xlim = [xmin xmax]` and `ylim = [ymin ymax]`.

Examples

Constrain drag of `impoint` within axes limits.

```
figure, plot(1:10);  
h = impoint(gca,2,6);  
api = iptgetapi(h);  
fcn = makeConstrainToRectFcn('impoint',get(gca,'XLim'),...  
    get(gca,'YLim'));  
api.setPositionConstraintFcn(fcn);
```

See Also

`imdistanline` | `imellipse` | `imfreehand` | `imline` | `impoint` | `impoly` | `imrect`

makehdr

Purpose Create high dynamic range image

Syntax
HDR = makehdr(files)
HDR = makehdr(files, param1, val1,...)

Description HDR = makehdr(files) creates the single-precision, high dynamic range image from the set of spatially registered low dynamic range images listed in the files cell array. makehdr uses the middle exposure between the brightest and darkest images as the base exposure for the high dynamic range calculations. (This value does not need to appear in any particular file.)

Note When you call makehdr with this syntax, the low dynamic range image files must contain exif exposure metadata.

HDR = makehdr(files, param1, val1,...) creates a high dynamic range image from the low dynamic range images in files, specifying parameters and corresponding values that control various aspects of the image creation. Parameter names can be abbreviated and case does not matter.

Note Only one of the BaseFile, ExposureValues, and RelativeExposure parameters may be used at a time.

Parameter	Description
'BaseFile'	Character array containing the name of the file to use as the base exposure.
'ExposureValues'	Vector of exposure values, with one element for each low dynamic range image in the cell array <code>files</code> . An increase in one exposure value (EV) corresponds to a doubling of exposure, while a decrease in one EV corresponds to a halving of exposure. Any positive value is allowed. This parameter overrides EXIF exposure metadata.
'RelativeExposure'	Vector of relative exposure values, with one element for each low dynamic range image in the cell array <code>files</code> . An image with a relative exposure (RE) of 0.5 has half as much exposure as an image with an RE of 1. An RE value of 3 has three times the exposure of an image with an RE of 1. This parameter overrides EXIF exposure metadata.
'MinimumLimit'	Numeric scalar value in the range [0 255] that specifies the minimum correctly exposed value. For each low dynamic range image, pixels with smaller values are considered underexposed and will not contribute to the final high dynamic range image.
'MaximumLimit'	Numeric scalar value in the range [0 255] that specifies the maximum correctly exposed value. For each low dynamic range image, pixels with larger values are considered overexposed and will not contribute to the final high dynamic range image.

Examples

Make a high dynamic range image from a series of six low dynamic range images that share the same *f*/stop number and have different exposure times. Use `tonemap` to visualize the HDR image.

makehdr

```
files = {'office_1.jpg', 'office_2.jpg', 'office_3.jpg', ...
         'office_4.jpg', 'office_5.jpg', 'office_6.jpg'};
expTimes = [0.0333, 0.1000, 0.3333, 0.6250, 1.3000, 4.0000];

hdr = makehdr(files, 'RelativeExposure', expTimes ./ expTimes(1));
rgb = tonemap(hdr);
figure; imshow(rgb)
```

References

[1] Reinhard, et al. "High Dynamic Range Imaging." 2006. Ch. 4.

See Also

hdrread | tonemap

Purpose

Create lookup table for use with `bwlookup`

Syntax

```
lut = makelut(fun,n)
```

Description

`lut = makelut(fun,n)` returns a lookup table for use with `bwlookup`. `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.

Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`.

Class Support

`lut` is returned as a vector of class `double`.

Examples

Construct a lookup table for 2-by-2 neighborhoods. In this example, the function passed to `makelut` returns `TRUE` if the number of 1's in the neighborhood is 2 or greater, and returns `FALSE` otherwise.

```
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

`bwlookup`

How To

- “Anonymous Functions”
- “Parameterizing Functions”

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}}
```

makeresampler

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.
- 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 software. 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.

makeresampler

Property	Description
'ResampleFcn'	<p>Handle to a function that performs the resampling. The function is called with the following interface.</p> $B = \text{resample_fcn}(A, M, \text{TDIMS_A}, \text{TDIMS_B}, \text{FSIZE_A}, \text{FSIZE_B}, F, R)$ <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 (<code>N = length(TDIMS_A)</code>) and <code>B</code> has <code>P</code> transform dimensions (<code>P = length(TDIMS_B)</code>), then <code>ndims(M) = P + 1</code>, if <code>N > 1</code> and <code>P</code> if <code>N == 1</code>, and <code>size(M, P + 1) = N</code>.</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 <code>double</code>. <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.

Examples

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 spatial transformation structure (TFORM)

Syntax

```
T = maketform(transformtype,...)
T = maketform('affine',A)
T = maketform('affine',U,X)
T = maketform('projective',A)
T = maketform('projective',U,X)
T = maketform('custom', NDIMS_IN, NDIMS_OUT, FORWARD_FCN,
    INVERSE_FCN, TDATA)
T = maketform('box',tsize,LOW,HIGH)
T = maketform('box',INBOUNDS, OUTBOUNDS)
T = maketform('composite',T1,T2,...,TL)
T = maketform('composite', [T1 T2 ... TL])
```

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 syntax.

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.

Note An affine or projective transformation can also be expressed like this, for a 3-by-2 A:

$$[X \ Y]' = A' * [U \ V \ 1]'$$

Or, like this, for a 3-by-3 A:

$$[X \ Y \ 1]' = A' * [U \ V \ 1]'$$

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 syntax:

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 Kth 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.

Note that the inputs T1, T2, ..., TL are ordered just as they would be when using the standard notation for function composition: $T = T1 \circ T2 \circ \dots \circ TL$ and note also that composition is associative, but not commutative. This means that in order to apply T to the input U, one must apply TL first and T1 last. Thus if L = 3, for example, 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.

Examples

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);
```

maketform

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

See Also

`tformfwd` | `tforminv` | `fliptform` | `imtransform` | `tformarray`

Purpose Convert matrix to grayscale 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.

Examples
`I = imread('rice.png');`
`J = filter2(fspecial('sobel'),I);`
`K = mat2gray(J);`
`imshow(I), figure, imshow(K)`



See Also [gray2ind](#) | [ind2gray](#) | [rgb2gray](#)

mean2

Purpose	Average or mean of matrix elements
Syntax	<code>B = mean2(A)</code>
Description	<code>B = mean2(A)</code> computes the mean of the values in A.
Class Support	The input image A can be <code>numeric</code> or <code>logical</code> . The output image B is a scalar of class <code>double</code> .
Algorithms	<code>mean2</code> computes the mean of an array A using <code>mean(A(:))</code> .
See Also	<code>std2</code> <code>mean</code> <code>std</code>

Purpose 2-D median filtering

Note The syntax `medfilt2(A,[M N],[Mb Nb],...)` has been removed.

Syntax

```
B = medfilt2(A, [m n])
B = medfilt2(A)
B = medfilt2(A, 'indexed', ...)
B = medfilt2(..., padopt)
```

Description Median filtering is a nonlinear operation often used in image processing to reduce "salt and pepper" noise. A median filter is more effective than convolution when the goal is to simultaneously reduce noise and preserve edges.

`B = medfilt2(A, [m n])` performs median filtering of the matrix `A` in two dimensions. Each output pixel contains the median value in the `m`-by-`n` neighborhood around the corresponding pixel in the input image. `medfilt2` pads the image with 0s on the edges, so the median values for the points within `[m n]/2` of the edges might appear distorted.

`B = medfilt2(A)` performs median filtering of the matrix `A` using the default 3-by-3 neighborhood.

`B = medfilt2(A, 'indexed', ...)` processes `A` as an indexed image, padding with 0s if the class of `A` is `uint8`, or 1s if the class of `A` is `double`.

`B = medfilt2(..., padopt)` controls how the matrix boundaries are padded. `padopt` may be `'zeros'` (the default), `'symmetric'`, or `'indexed'`. If `padopt` is `'symmetric'`, `A` is symmetrically extended at the boundaries. If `padopt` is `'indexed'`, `A` is padded with ones if it is `double`; otherwise it is padded with zeros.

Class Support The input image `A` can be of class `logical` or `numeric` (unless the `'indexed'` syntax is used, in which case `A` cannot be of class `uint16`). The output image `B` is of the same class as `A`.

Note For information about performance considerations, see `ordfilt2`.

Tips

If the input image `A` is of an integer class, all the output values are returned as integers. If the number of pixels in the neighborhood (i.e., $m \times n$) is even, some of the median values might not be integers. In these cases, the fractional parts are discarded. Logical input is treated similarly.

For example, suppose you call `medfilt2` using 2-by-2 neighborhoods, and the input image is a `uint8` array that includes this neighborhood.

```
1 5
4 8
```

`medfilt2` returns an output value of 4 for this neighborhood, although the true median is 4.5.

Examples

Add salt and pepper noise to an image and then restore the image using `medfilt2`.

```
I = imread('eight.tif');
J = imnoise(I,'salt & pepper',0.02);
K = medfilt2(J);
imshow(J), figure, imshow(K)
```



Algorithms

medfilt2 uses ordfilt2 to perform the filtering.

References

[1] Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, pp. 469-476.

See Also

filter2 | ordfilt2 | wiener2

montage

Purpose Display multiple image frames as rectangular montage

Syntax

```
montage(filenamees)
montage(I)
montage(X, map)
montage(..., param1, value1, param2, value2, ...)
h = montage(...)
```

Description `montage(filenamees)` displays a montage of the images specified in `filenamees`. `filenamees` is an N-by-1 or 1-by-N cell array of filenames. If the files are not in the current directory or in a directory on the MATLAB path, you must specify the full pathname. See the `imread` command for more information. If one or more of the image files contains an indexed image, `montage` uses the colormap from the first indexed image file. `montage` arranges the frames so that they roughly form a square.

`montage(I)` displays all the frames of a multiframe image array `I` in a single image object. `I` can be a sequence of binary, grayscale, or truecolor images. A binary or grayscale image sequence must be an M-by-N-by-1-by-K array. A truecolor image sequence must be an M-by-N-by-3-by-K array.

`montage(X, map)` displays all the frames of the indexed image array `X`, using the colormap `map` for all frames. `X` is an M-by-N-by-1-by-K array.

`montage(..., param1, value1, param2, value2, ...)` returns a customized display of an image montage, depending on the values of the optional parameter/value pairs. See “Parameters” on page 3-678 for a list of available parameters.

`h = montage(...)` returns the handle to the image object.

Parameters

The following table lists the parameters available, alphabetically by name. Parameter names can be abbreviated, and case does not matter.

Parameter	Value
'DisplayRange'	<p>A 1-by-2 vector, [LOW HIGH] that controls the display range of each image in a grayscale sequence. The value LOW (and any value less than LOW) displays as black; the value HIGH (and any value greater than HIGH) displays as white. If you specify an empty matrix ([]), <code>montage</code> uses the minimum and maximum values of the images to be displayed in the montage as specified by 'Indices'. For example, if 'Indices' is 1:K and the 'DisplayRange' is set to [], the minimum value in I ($\min(I(:))$) is displayed as black, and the maximum value ($\max(I(:))$) is displayed as white.</p> <p>Default: Range of the data type of I.</p>
'Indices'	<p>A numeric array specifying which frames to display in the montage. For example, to create a montage of the first four frames in I, enter <code>montage(I, 'Indices', 1:4);</code>. You can use this parameter to specify individual frames or skip frames. For example, the value 1:2:20 displays every other frame.</p> <p>Default: 1:K, where K is the total number of frames to display.</p>
'Size'	<p>A 2-element vector, [NROWS NCOLS], specifying the number of rows and number of columns in the montage. You can use NaNs in the size vector to indicate that <code>montage</code> should calculate size in a particular dimension in a way that includes all the images in the montage. For example, if 'Size' is [2 NaN], the montage will have 2 rows, and the number of columns will be computed automatically to include all of the images in the montage. The images are displayed horizontally across columns.</p> <p>Default: Calculated so the images in the montage roughly form a square. <code>montage</code> considers the aspect ratio when calculating the number of images to display horizontally and vertically.</p>

Class Support

A grayscale image array can be `logical`, `uint8`, `uint16`, `int16`, `single`, or `double`. An indexed image can be `logical`, `uint8`, `uint16`, `single`, or `double`. The colormap must be `double`. A truecolor image can be

montage

uint8, uint16, single, or double. The output is a handle to the image object produced by montage.

Examples

Create a montage from a series of files. By default, montage arranges the images into a square.


```
fileFolder = fullfile(matlabroot, 'toolbox', 'images', 'imdemos');  
dirOutput = dir(fullfile(fileFolder, 'AT3_1m4_*.tif'));  
fileNames = {dirOutput.name}'  
montage(fileNames);
```

Display the same set of images in two rows and five columns.

```
montage(fileNames, 'Size', [2 5]);
```

This example shows you how to customize the number of images in the montage and the contrast in the montage.

```
load mri  
montage(D, map)
```



Create a new montage of the first 9 images.

```
figure
montage(D, map, 'Indices', 1:9);
```

Maximize the contrast of D without using the colormap.

```
figure
montage(D, 'DisplayRange', []);
```

See Also

`immovie` | `imshow`

multithresh

Purpose Multilevel image thresholds using Otsu's method

Syntax

```
thresh = multithresh(A)
thresh = multithresh(A,N)
[thresh,metric] = multithresh(____)
```

Description `thresh = multithresh(A)` returns the single threshold value `thresh` computed for image `A` using Otsu's method.

- `thresh` is an input argument to `imquantize` that converts `A` into a binary image.

`thresh = multithresh(A,N)` returns `thresh` containing `N` threshold values using Otsu's method.

- `thresh` is a `1xN` vector which can be used to convert image `A` into an image with `N + 1` discrete levels using `imquantize`.

`[thresh,metric] = multithresh(____)` returns `metric`, a measure of the effectiveness of the computed thresholds. `metric` is in the range `[0 1]` and a higher value indicates greater effectiveness of the thresholds in separating the input image into `N + 1` classes based on Otsu's objective criterion.

Input Arguments

A - Input image

`image`

Input image, specified as a non-sparse numeric matrix of any dimension. `multithresh` finds the thresholds based on the aggregate histogram of the entire matrix. An RGB image is considered a 3-D numeric array and the thresholds are computed for the combined data from all three color planes.

`multithresh` uses the range of the input `A` `[min(A(:)) max(A(:))]` as the limits for computing the histogram which is used in subsequent computations. Any NaNs are ignored in computation. Any `Inf`s and `-Inf`s are counted in the first and last bin of the histogram, respectively.

For degenerate inputs where the number of unique values in A is less than or equal to N , there is no viable solution using Otsu's method. For such inputs, `thresh` contains all the unique values from A and possibly some extra values that are chosen arbitrarily.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32
| uint64

N - Number of threshold values

1 (default) | scalar

Number of threshold values, specified as a scalar value. For $N > 2$, `multithresh` uses search-based optimization of Otsu's criterion to find the thresholds. The search-based optimization guarantees only locally optimal results. Since the chance of converging to local optimum increases with N , it is preferable to use smaller values of N , typically $N < 10$. The maximum allowed value for N is 20.

Data Types

single | double | int8 | int16 | int32 | int64 | uint8 | uint16 | uint32
| uint64

Output Arguments**thresh - Set of values used to quantize an image**

1xN vector

Set of values used to quantize an image, returned as a 1xN vector, whose data type is the same as image A .

metric - Measure of the effectiveness of the thresholds

scalar

Measure of the effectiveness of the thresholds, returned as a scalar value. Higher values indicates greater effectiveness of the thresholds in separating the input image into $N+1$ classes based on Otsu's objective criterion. For degenerate inputs where the number of unique values in A is less than or equal to N , `metric` equals 0.

Data Types

double

multithresh

Examples

Image Threshold

Compute multiple thresholds for an image using `multithresh` and apply those thresholds to the image using `imquantize` to get segment labels.

```
I = imread('circlesBrightDark.png');
```

Computing two thresholds quantizes the image into three discrete levels.

```
thresh = multithresh(I,2);
```

```
seg_I = imquantize(I,thresh); % Apply the thresholds to obtain segmented
```

```
RGB = label2rgb(seg_I); % convert to color image
```


```
figure;
```

```
imshowpair(I,RGB,'montage'); % display images side-by-side
```

```
axis off;
```

```
title('Original Image
```

```
RGB Segmented Image')
```



Two Schemes for RGB Image Quantization

Quantize the truecolor RGB image to eight levels using multilevel thresholding function `multithresh`. Compare results between thresholding the entire RGB image versus plane-by-plane thresholding.

Read in a truecolor RGB image.

multithresh

```
I = imread('peppers.png');  
imshow(I); axis off;  
title('RGB Image');
```

RGB Image



Generate thresholds for seven levels from the entire RGB image; then repeat process for each plane in the RGB image

```
threshRGB = multithresh(I,7) % 7 thresholds from entire RGB image
```

```
threshPlanes = zeros(3,7); % initialize
% Compute thresholds for each R, G, and B plane
for i = 1:3
    threshPlane(i,:) = multithresh(I(:,:,i),7);
end
threshPlane
```

```
threshRGB =
```

```
    24    51    79   109   141   177   221
```

```
threshPlane =
```

```
    40    69    92   125   159   195   231
    27    49    74   100   128   164   209
    18    38    60    86   127   165   222
```

Process the entire image with the set of threshold values computed from the entire image. Similarly, process each RGB plane separately using the threshold vector computed from the given plane. Add black (0) and white (255) to value vector which assigns values to the output image.

```
value = [0 threshRGB(2:end) 255];
```

```
% Quantize entire image using one threshold vector
```

```
quantRGB = imquantize(I, threshRGB, value);
```

```
quantPlane = zeros( size(I) );
```

```
% Quantize each RGB plane using threshold vector generated for that p
```

```
for i = 1:3
```

```
    value = [0 threshPlane(i,2:end) 255] % output value to assign
```

```
    quantPlane(:,:,i) = imquantize(I(:,:,i), threshPlane(i,:), value);
```

```
end
```

multithresh

```
quantPlane = uint8( quantPlane ); % convert from double to uint8
```

Display both posterized images and note the visual differences in the two thresholding schemes.

Full RGB Image Quantization



The following code snippet computes the number of unique RGB pixel vectors in each output image. Note that the plane-by-plane thresholding scheme yields about 23% more colors than the full RGB image scheme.

```

% Convert images to mx3 matrices
dim = size( quantRGB );
quantRGBmx3 = reshape(quantRGB, prod(dim(1:2)), 3);
quantPlanemx3 = reshape(quantPlane, prod(dim(1:2)), 3);

% Extract only unique 3 element RGB pixel vectors from each matrix
colorsRGB = unique( quantRGBmx3, 'rows');
colorsPlane = unique( quantPlanemx3, 'rows');

disp(['Number of unique colors in RGB image          : ' int2str(len
disp(['Number of unique colors in Plane-by-Plane image : ' int2str(len

Number of unique colors in RGB image          : 188
Number of unique colors in Plane-by-Plane image : 231

```

Behavior of Output Argument metric

Compute various thresholds for different values of N and on different images as a way of showing the variation of `metric`.

```

I = imread('circlesBrightDark.png');

% find all unique grayscale values in image
uniqLevels = unique(I(:));

disp(['Number of unique levels = ' int2str( length(uniqLevels) )]);

Number of unique levels = 148

Compute a series of thresholds at monotonically increasing values of N

% Compute the thresholds
Nvals = [1 2 4 8];
for i = 1:length(Nvals)
    [thresh, metric] = multithresh(I, Nvals(i) );
    disp(['N = ' int2str(Nvals(i)) ' | metric = ' num2str(metric)]);
end

N = 1 | metric = 0.54767

```

multithresh

```
N = 2 | metric = 0.98715
N = 4 | metric = 0.99648
N = 8 | metric = 0.99902
```

Apply the N=8 set of thresholds to obtain a 9 level segmented image via `imquantize`.

```
seg_Neq8 = imquantize(I,thresh);
uniqLevels = unique( seg_Neq8(:) ) % Verify that image has only 9 levels
```

```
uniqLevels =
```

```
1
2
3
4
5
6
7
8
9
```

`seg_Neq8` is input to `multithresh` and `N` is set to equal to 8 which is 1 less than the number of levels in this segmented image.

```
[thresh, metric] = multithresh(seg_Neq8,8)
```

```
thresh =
```

```
Columns 1 through 7
```

```
1.8784    2.7882    3.6667    4.5451    5.4549    6.3333    7.2118
```

```
Column 8
```

```
8.1216
```



```
metric =
```

```
1
```

Note what happens when N is increased by 1 and now equals the number of levels in the image.

```
[thresh, metric] = multithresh(seg_Neq8,9)
```

```
Warning: No solution exists because the number of unique levels in image is
too few to find 9 thresholds. Returning an arbitrarily chosen solution.
> In multithresh at 171
```

```
thresh =
```

```
1 2 3 4 5 6 7 8 9
```

```
metric =
```

```
0
```

Here the input was degenerate because the number of levels in the image was too few for the number of requested thresholds. Hence metric is 0.

References

[1] 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.

See Also

[graythresh](#) | [imquantize](#) | [im2bw](#) | [rgb2ind](#)

nitfinfo

Purpose Read metadata from National Imagery Transmission Format (NITF) file

Syntax `metadata = nitfinfo(filename)`

Description `metadata = nitfinfo(filename)` returns a structure whose fields contain file-level metadata about the images, annotations, and graphics in a National Imagery Transmission Format (NITF) file. NITF is an image format used by the U.S. government and military for transmitting documents. A NITF file can contain multiple images and include text and graphic layers. *filename* is a character array that specifies the name of the NITF file, which must be in the current directory, in a directory on the MATLAB path, or contain the full path to the file.

`nitfinfo` supports version 2.0 and 2.1 NITF files, at all Joint Interoperability Test Command (JITC) compliance levels, as well as the NATO Secondary Image Format (NSIF) 1.0. `nitfinfo` does not support NITF 1.1 files.

See Also `isnitf` | `nitfread`

Purpose	Read image from NITF file
Syntax	<pre>X = nitfread(filename) X = nitfread(filename,idx) X = nitfread(...,'PixelRegion',{rows,cols})</pre>
Description	<p><code>X = nitfread(filename)</code> reads the first image from the National Imagery Transmission Format (NITF) file specified by the character array <code>filename</code>. The <code>filename</code> array must be in the current directory or in a directory on the MATLAB path, or it must contain the full path to the file.</p> <p><code>X = nitfread(filename,idx)</code> reads the image with index number <code>idx</code> from a NITF file that contains multiple images.</p> <p><code>X = nitfread(...,'PixelRegion',{rows,cols})</code> reads a region of pixels from a NITF image. <code>rows</code> and <code>cols</code> are two or three element vectors, where the first value is the starting location, and the last value is the ending location. In the three value syntax, the second value is the increment.</p> <p>This function supports version 2.0 and 2.1 NITF files, as well as NSIF 1.0. Compressed images, image submasks, and NITF 1.1 files are not supported.</p>
See Also	<code>isnitf</code> <code>nitfinfo</code>

nlfilter

Purpose

General sliding-neighborhood operations

Syntax

```
B = nlfilter(A, [m n], fun)
B = nlfilter(A, 'indexed',...)
```

Description

`B = nlfilter(A, [m n], fun)` applies the function `fun` to each `m`-by-`n` sliding block of the grayscale image `A`. `fun` is a function that accepts an `m`-by-`n` matrix as input and returns a scalar result.

```
c = fun(x)
```

`fun` must be a function handle.

Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`.

`c` is the output value for the center pixel in the `m`-by-`n` block `x`. `nlfilter` calls `fun` for each pixel in `A`. `nlfilter` zero-pads the `m`-by-`n` block at the edges, if necessary.

`B = nlfilter(A, 'indexed',...)` processes `A` as an indexed image, padding with 1's if `A` is of class `single` or `double` and 0's if `A` is of class `logical`, `uint8`, or `uint16`.

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`. When `A` is grayscale, it can be any numeric type or `logical`. When `A` is indexed, it can be `logical`, `uint8`, `uint16`, `single`, or `double`.

Tips

`nlfilter` can take a long time to process large images. In some cases, the `colfilt` function can perform the same operation much faster.

Examples

This example produces the same result as calling `medfilt2` with a 3-by-3 neighborhood.

```
A = imread('cameraman.tif');
A = im2double(A);
```

```
fun = @(x) median(x(:));  
B = nlfilter(A,[3 3],fun);  
imshow(A), figure, imshow(B)
```

See Also

[blockproc](#) | [colfilt](#) | [function_handle](#)

How To

- “Anonymous Functions”
- “Parameterizing Functions”

normxcorr2

Purpose Normalized 2-D 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 numeric. The output matrix `C` is double.

Tips Normalized cross-correlation is an undefined operation in regions where `A` has zero variance over the full extent of the template. In these regions, we assign correlation coefficients of zero to the output `C`.

Algorithms `normxcorr2` uses the following general procedure [1], [2]:

- 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.

The implementation closely follows following formula from [1]:

$$\gamma(u,v) = \frac{\sum_{x,y} [f(x,y) - \bar{f}_{u,v}] [t(x-u, y-v) - \bar{t}]}{\left\{ \sum_{x,y} [f(x,y) - \bar{f}_{u,v}]^2 \sum_{x,y} [t(x-u, y-v) - \bar{t}]^2 \right\}^{0.5}}$$

where

- f is the image.
- \bar{t} is the mean of the template
- $\bar{f}_{u,v}$ is the mean of $f(x,y)$ in the region under the template.

Examples

```

template = .2*ones(11); % Make light gray plus on dark gray background
template(6,3:9) = .6;
template(3:9,6) = .6;
BW = template > 0.5;      % Make white plus on black background
figure, imshow(BW), figure, imshow(template)
% Make new image that offsets the template
offsetTemplate = .2*ones(21);
offset = [3 5]; % Shift by 3 rows, 5 columns
offsetTemplate( (1:size(template,1))+offset(1),...
               (1:size(template,2))+offset(2) ) = template;
figure, imshow(offsetTemplate)

% Cross-correlate BW and offsetTemplate to recover offset
cc = normxcorr2(BW,offsetTemplate);
[max_cc, imax] = max(abs(cc(:)));
[ypeak, xpeak] = ind2sub(size(cc),imax(1));
corr_offset = [ (ypeak-size(template,1)) (xpeak-size(template,2)) ];
isequal(corr_offset,offset) % 1 means offset was recovered

```

References

- [1] Lewis, J. P., "Fast Normalized Cross-Correlation," *Industrial Light & Magic*
- [2] Haralick, Robert M., and Linda G. Shapiro, *Computer and Robot Vision*, Volume II, Addison-Wesley, 1992, pp. 316-317.

See Also

corrcoef

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`.

Examples Convert RGB image to NTSC and back.

```
RGB = imread('board.tif');  
NTSC = rgb2ntsc(RGB);  
RGB2 = ntsc2rgb(NTSC);
```

See Also `rgb2ntsc` | `rgb2ind` | `ind2rgb` | `ind2gray`

Purpose	Open R-Set file
Syntax	<code>openrset(<i>filename</i>)</code>
Description	<code>openrset(<i>filename</i>)</code> opens the reduced resolution dataset (R-Set) specified by <i>filename</i> for viewing.
See Also	<code>imtool</code> <code>rsetwrite</code>

ordfilt2

Purpose

2-D order-statistic filtering

Syntax

```
B = ordfilt2(A, order, domain)
B = ordfilt2(A, order, domain, S)
B = ordfilt2(..., padopt)
```

Description

`B = ordfilt2(A, order, domain)` replaces each element in `A` by the `orderth` element in the sorted set of neighbors specified by the nonzero elements in `domain`.

`B = ordfilt2(A, order, domain, S)` where `S` is the same size as `domain`, uses the values of `S` corresponding to the nonzero values of `domain` as additive offsets.

`B = ordfilt2(..., padopt)` controls how the matrix boundaries are padded. Set `padopt` to `'zeros'` (the default) or `'symmetric'`. If `padopt` is `'zeros'`, `A` is padded with 0's at the boundaries. If `padopt` is `'symmetric'`, `A` is symmetrically extended at the boundaries.

Class Support

The class of `A` can be `logical`, `uint8`, `uint16`, or `double`. The class of `B` is the same as the class of `A`, unless the additive offset form of `ordfilt2` is used, in which case the class of `B` is `double`.

Tips

`domain` 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.

For example, `B = ordfilt2(A,5,ones(3,3))` implements a 3-by-3 median filter; `B = ordfilt2(A,1,ones(3,3))` implements a 3-by-3 minimum filter; and `B = ordfilt2(A,9,ones(3,3))` implements a 3-by-3 maximum filter. `B = ordfilt2(A,1,[0 1 0; 1 0 1; 0 1 0])` replaces each element in `A` by the minimum of its north, east, south, and west neighbors.

The syntax that includes `S` (the matrix of additive offsets) can be used to implement grayscale morphological operations, including grayscale dilation and erosion.

Performance Considerations

When working with large domain matrices that do not contain any zero-valued elements, `ordfilt2` can achieve higher performance if `A` is in an integer data 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.

Examples

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)
```

References

- [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

See Also

`medfilt2`

otf2psf

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.

Examples

```
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

Purpose

Pad array

Syntax

```
B = padarray(A, padsize)
B = padarray(A, padsize, padval)
B = padarray(A, padsize, padval, direction)
```

Description

`B = padarray(A, padsize)` pads array `A` with 0's (zeros). `padsize` is a vector of nonnegative 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 `padsize` 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, padsize, 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 (`{}`).

padarray

Value	Meaning
{ 'both' }	Pads before the first element and after the last array element along each dimension. This is the default.
'post'	Pad after the last array element along each dimension.
'pre'	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.

Examples

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')  
    3  2  1  1  2  3  4  
b ==
```

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')  
    1  2  2  2  
    3  4  4  4  
    3  4  4  4  
    3  4  4  4  
    3  4  4  4  
B =
```

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])
```

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(:,:,1) ==
```

padarray

```
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](#)

Purpose Convert parallel-beam projections to fan-beam

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 ifanbeam 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.

para2fan

Parameter	Description
'FanSensorGeometry'	<p>Text string specifying how sensors are positioned.</p> <p>Possible values: {'arc'} or 'line'</p> <p>See <code>fanbeam</code> 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> <p>If 'FanSensorSpacing' is not specified, the default is the smallest value implied by 'ParallelSensorSpacing' such that</p> <p>If 'FanSensorGeometry' is 'arc', 'FanSensorSpacing' is</p> $180/\text{PI}*\text{ASIN}(\text{PSPACE}/D)$ <p>where PSPACE is the value of 'ParallelSensorSpacing'.</p> <p>If 'FanSensorGeometry' is 'line', 'FanSensorSpacing' is</p> $D*\text{ASIN}(\text{PSPACE}/D)$

Parameter	Description
'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>'cycle' -- Parallel data covers 360 degrees</p> <p>{ 'halfcycle' } — Parallel data covers 180 degrees</p>
'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> $[D \cdot \sin(\min(\text{FAN_ANGLES})), D \cdot \sin(\max(\text{FAN_ANGLES}))]$ <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>

`[F, fan_positions, fan_rotation_angles] = fan2para(...)`
returns the fan-beam sensor measurement *angles* in `fan_positions`, if 'FanSensorGeometry' is 'arc'. 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.

Examples

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](#)

Purpose Create 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] = phantom(...)` returns the matrix E used to generate the phantom.

Class Support

All inputs and all outputs must be of class `double`.

Tips

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 $[0,1]$.

Examples

```
P = phantom('Modified Shepp-Logan',200);  
imshow(P)
```



References

[1] Jain, Anil K., *Fundamentals of Digital Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1989, p. 439.

See Also

radon | iradon

poly2mask

Purpose Convert region of interest (ROI) polygon to region mask

Syntax `BW = poly2mask(x, y, m, n)`

Description `BW = poly2mask(x, y, m, n)` computes a binary region of interest (ROI) mask, `BW`, from an ROI polygon, represented by the vectors `x` and `y`. The size of `BW` is `m`-by-`n`. `poly2mask` sets pixels in `BW` that are inside the polygon (`X,Y`) to 1 and sets pixels outside the polygon to 0.

`poly2mask` closes the polygon automatically if it isn't already closed.

Note on Rectangular Polygons

When the input polygon goes through the middle of a pixel, sometimes the pixel is determined to be inside the polygon and sometimes it is determined to be outside (see Algorithm for details). To specify a polygon that includes a given rectangular set of pixels, make the edges of the polygon lie along the outside edges of the bounding pixels, instead of the center of the pixels.

For example, to include pixels in columns 4 through 10 and rows 4 through 10, you might specify the polygon vertices like this:

```
x = [4 10 10 4 4];  
y = [4 4 10 10 4];  
mask = poly2mask(x,y,12,12)
```

```
mask =
```

```
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 0 0 0 0 0 0 0 0 0  
0 0 0 0 1 1 1 1 1 1 0 0  
0 0 0 0 1 1 1 1 1 1 0 0  
0 0 0 0 1 1 1 1 1 1 0 0  
0 0 0 0 1 1 1 1 1 1 0 0  
0 0 0 0 1 1 1 1 1 1 0 0  
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 0 0 0 0

```

In this example, the polygon goes through the center of the bounding pixels, with the result that only some of the desired bounding pixels are determined to be inside the polygon (the pixels in row 4 and column 4 and not in the polygon). To include these elements in the polygon, use fractional values to specify the outside edge of the 4th row (3.5) and the 10th row (10.5), and the outside edge of the 4th column (3.5) and the outside edge of the 10th column (10.5) as vertices, as in the following example:

```

x = [3.5 10.5 10.5 3.5 3.5];
y = [3.5 3.5 10.5 10.5 3.5];
mask = poly2mask(x,y,12,12)

```

mask =

```

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 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 0 0
0 0 0 1 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 0 0 0 0

```

Class Support

The class of BW is logical

Examples

```

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
```

Algorithms

When creating a region of interest (ROI) 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.



Pixels on the Edge of an ROI Are Only Partially Covered by Border


To determine which pixels are in the region, poly2mask uses the following algorithm:

- 1** Divide each pixel (unit square) into a 5-by-5 grid. See “Dividing Pixels into a 5-by-5 Subpixel Grid” on page 3-718 for an illustration.
- 2** Adjust the position of the vertices to be on the intersections of the subpixel grid. See “Adjusting the Vertices to the Subpixel Grid” on page 3-718 for an illustration.
- 3** Draw a path from each adjusted vertex to the next, following the edges of the subpixel grid. See “Drawing a Path Between the Adjusted Vertices” on page 3-719 for an illustration.
- 4** Determine which border pixels are inside the polygon using this rule: if a pixel’s central subpixel is inside the boundaries defined by the

path between adjusted vertices, the pixel is considered inside the polygon. See “Determining Which Pixels Are in the Region” on page 3-720 for an illustration.


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.



See Also `roipoly`

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.

Examples

```
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

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`. Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`.

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`.

Tips

`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.

Examples

```
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](#)

How To

- “Anonymous Functions”
- “Parameterizing Functions”

Purpose 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`.

Tips 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.

Examples

```
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);
```

qtgetblk

```
[vals,r,c] = qtgetblk(I,S,4)
```

See Also

qtdecomp | qtsetblk

Purpose Set block values in quadtree decomposition

Syntax `J = qtsetblk(I, S, dim, vals)`

Description `J = qtsetblk(I, S, dim, vals)` replaces each `dim`-by-`dim` block in the quadtree decomposition of `I` with the corresponding `dim`-by-`dim` block in `vals`. `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.

Class Support `I` can be of class `logical`, `uint8`, `uint16`, `int16`, `single`, or `double`. `S` is of class `sparse`.

Tips The ordering of the blocks in `vals` must match the columnwise order of the blocks in `I`. For example, if `vals` is 4-by-4-by-2, `vals(:, :, 1)` contains the values used to replace the first 4-by-4 block in `I`, and `vals(:, :, 2)` contains the values for the second 4-by-4 block.

Examples

```
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)
```

See Also `qtdecomp` | `qtgetblk`

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

```
floor((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`.

Examples

```
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  
iptsetpref('ImshowAxesVisible','off')
```



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.

Algorithms

The Radon transform of an image is the sum of the Radon transforms of each individual pixel.

The algorithm first divides pixels in the image into four subpixels and projects each subpixel separately, as shown in the following figure.



Each subpixel's contribution is proportionally split into the two nearest bins, according to the distance between the projected location and the bin centers. If the subpixel projection hits the center point of a bin, the bin on the axes gets the full value of the subpixel, or one-fourth the value of the pixel. If the subpixel projection hits the border between two bins, the subpixel value is split evenly between the bins.

See Also

[fan2para](#) | [fanbeam](#) | [ifanbeam](#) | [iradon](#) | [para2fan](#) | [phantom](#)

Purpose	Local range of image
Syntax	<code>J = rangefilt(I)</code> <code>J = rangefilt(I, NHOOD)</code>
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>. The image <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</p>

then use the `getnhood` method to extract the neighborhood from the structuring element object.

Examples

(2-D) Identify the two flying objects by measuring the local range.

```
I = imread('liftingbody.png');
J = rangefilt(I);
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.
Examples	<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(BW, properties)  
STATS = regionprops(CC, properties)  
STATS = regionprops(L, properties)  
STATS = regionprops(..., I, properties)
```

Description

STATS = regionprops(BW, *properties*) measures a set of properties for each connected component (object) in the binary image, BW. The image BW is a logical array; it can have any dimension.

STATS = regionprops(CC, *properties*) measures a set of properties for each connected component (object) in CC, which is a structure returned by bwconncomp.

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.

STATS = regionprops(..., I, *properties*) measures a set of properties for each labeled region in the image I. The first input to regionprops—either BW, CC, or L—identifies the regions in I. The sizes must match: size(I) must equal size(BW), CC.ImageSize, or size(L).

STATS is a structure array with length equal to the number of objects in BW, CC.NumObjects, or max(L(:)). The fields of the structure array denote different properties for each region, as specified by *properties*.

Properties

properties can be a comma-separated list of strings, a cell array containing strings, the single string 'all', or the string 'basic'. If *properties* is the string 'all', regionprops computes all the shape measurements, listed in Shape Measurements on page 3-737. If called with a grayscale image, regionprops also returns the pixel value measurements, listed in Pixel Value Measurements on page 3-737. If *properties* is not specified or if it is the string 'basic', regionprops computes only the 'Area', 'Centroid', and 'BoundingBox'

measurements. You can calculate the following properties on N-D inputs: 'Area', 'BoundingBox', 'Centroid', 'FilledArea', 'FilledImage', 'Image', 'PixelIdxList', 'PixelList', and 'SubarrayIdx'.

Shape Measurements

'Area'	'EulerNumber'	'Orientation'
'BoundingBox'	'Extent'	'Perimeter'
'Centroid'	'Extrema'	'PixelIdxList'
'ConvexArea'	'FilledArea'	'PixelList'
'ConvexHull'	'FilledImage'	'Solidity'
'ConvexImage'	'Image'	'SubarrayIdx'
'Eccentricity'	'MajorAxisLength'	
'EquivDiameter'	'MinorAxisLength'	

Pixel Value Measurements

'MaxIntensity'	'MinIntensity'	'WeightedCentroid'
'MeanIntensity'	'PixelValues'	

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' — The smallest rectangle containing the region, a 1-by-Q *2 vector, where Q is the number of image dimensions: `ndims(L)`, `ndims(BW)`, or `numel(CC.ImageSize)`. `BoundingBox` is `[u1_corner width]`, where:


regionprops

`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-Q vector that specifies 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 for a discontinuous region. The region consists of the white pixels; the green box is the bounding box, and the red dot is the centroid.



'ConvexHull' — p-by-2 matrix that specifies 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) that specifies 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 that specifies the number of pixels in 'ConvexImage'. This property is supported only for 2-D input label matrices.

'Eccentricity' — Scalar that specifies 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.

'EquivDiameter' — Scalar that specifies the diameter of a circle with the same area as the region. Computed as $\sqrt{4 \cdot \text{Area} / \pi}$. This property is supported only for 2-D input label matrices.

'EulerNumber' — Scalar that specifies 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. `regionprops` uses 8-connectivity to compute the EulerNumber measurement. To learn more about connectivity, see “Pixel Connectivity”.

'Extent' — Scalar that specifies the ratio of pixels in the region to pixels in the total bounding box. Computed as the Area divided by the area of the bounding box. This property is supported only for 2-D input label matrices.

'Extrema' — 8-by-2 matrix that specifies 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.


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.

regionprops



'FilledArea' — Scalar specifying the number of on pixels in FilledImage.

'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.



Original Image, Containing a Single Region




Image Returned

'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.

'MajorAxisLength' — Scalar specifying 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.

'MaxIntensity' — Scalar specifying the value of the pixel with the greatest intensity in the region.

'MeanIntensity' — Scalar specifying the mean of all the intensity values in the region.

'MinIntensity' — Scalar specifying the value of the pixel with the lowest intensity in the region.

'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.

'Orientation' — Scalar; the angle (in degrees ranging from -90 to 90 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.


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.
- The orientation is the angle between the horizontal dotted line and the major axis.



'Perimeter' — Scalar; the distance around the boundary of the region. `regionprops` computes the perimeter by calculating the distance between each adjoining pair of pixels around the border of the region. If the image contains discontinuous regions, `regionprops` returns

unexpected results. The following figure shows the pixels included in the perimeter calculation for this object.



'PixelIdxList' — p -element vector containing the linear indices of the pixels in the region.

'PixelList' — p -by- Q matrix specifying the locations of 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.

'PixelValues' — p -by-1 vector, where p is the number of pixels in the region. Each element in the vector contains the value of a pixel in the region.

'Solidity' — Scalar specifying 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.

'SubarrayIdx' — Cell-array containing indices such that $L(\text{idx}\{\})$ extracts the elements of L inside the object bounding box.

'WeightedCentroid' — p -by- Q vector of coordinates specifying the center of the region based on location and intensity value. The first element of **WeightedCentroid** is the horizontal coordinate (or x -coordinate) of the weighted centroid. The second element

is the vertical coordinate (or y -coordinate). All other elements of `WeightedCentroid` are in order of dimension.

Class Support

If the first input is `BW`, `BW` must be a logical array and it can have any dimension. If the first input is `CC`, `CC` must be a structure returned by `bwconncomp`. If the first input is `L`, `L` must be real, nonsparse, and contain integers. `L` can have any numeric class and any dimension.

Tips

Note on Terminology

You can use `regionprops` on contiguous regions and discontinuous regions.

Contiguous regions are also called "objects," "connected components," and "blobs." A label matrix containing contiguous regions might look like this:

```
1 1 0 2 2 0 3 3
1 1 0 2 2 0 3 3
```

Elements of `L` equal to 1 belong to the first contiguous region or connected component; elements of `L` equal to 2 belong to the second connected component; etc.

Discontinuous regions are regions that might contain multiple connected components. A label matrix containing discontinuous regions might look like this:

```
1 1 0 1 1 0 2 2
1 1 0 1 1 0 2 2
```

Elements of `L` equal to 1 belong to the first region, which is discontinuous and contains two connected components. Elements of `L` equal to 2 belong to the second region, which is a single connected component.

Selecting Regions Based on Certain Criteria

The function `ismember` is useful in conjunction with `regionprops`, `bwconncomp`, and `labelmatrix` for creating a binary image containing only objects or regions that meet certain criteria. For example, these

commands create a binary image containing only the regions whose area is greater than 80.

```
cc = bwconncomp(BW);
stats = regionprops(cc, 'Area');
idx = find([stats.Area] > 80);
BW2 = ismember(labelmatrix(cc), idx);
```

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`.

Performance Considerations

Most of the measurements take very little time to compute. However, there are a few measurements, listed below, that 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`.

Using `bwlabel`, `bwlabeln`, `bwconncomp`, and `regionprops`

The functions `bwlabel`, `bwlabeln`, and `bwconncomp` all compute connected components for binary images. `bwconncomp` replaces the use of `bwlabel` and `bwlabeln`. It uses significantly less memory and is sometimes faster than the other functions.

Function	Input Dimension	Output Form	Memory Use	Connectivity
<code>bwlabel</code>	2-D	Double-precision label matrix	High	4 or 8
<code>bwlabeln</code>	N-D	Double-precision label matrix	High	Any
<code>bwconncomp</code>	N-D	CC struct	Low	Any

The output of `bwlabel` and `bwlabeln` is a double-precision label matrix. To compute a label matrix using a more memory-efficient data type, use the `labelmatrix` function on the output of `bwconncomp`:

```
CC = bwconncomp(BW);
L = labelmatrix(CC);
```

To extract features from a binary image using `regionprops` with the default connectivity, it is no longer necessary to call `bwlabel` or `bwlabeln` first. You can simply pass the binary image directly to `regionprops`, which then uses the memory-efficient `bwconncomp` to

regionprops


compute the connected components automatically for you. To extract features from a binary image using a nondefault connectivity, call `bwconncomp` first and then pass the result to `regionprops`:

```
CC = bwconncomp(BW, CONN);  
S = regionprops(CC);
```

Examples

Label the connected pixel components in the `text.png` image, compute their centroids, and superimpose the centroid locations on the image:

```
BW = imread('text.png');  
s = regionprops(BW, 'centroid');  
centroids = cat(1, s.Centroid);  
imshow(BW)  
hold on  
plot(centroids(:,1), centroids(:,2), 'b*')  
hold off
```



See Also

`bwconncomp` | `bwlabel` | `bwlabeln` | `ismember` | `labelmatrix` | `watershed`

registration.metric.MattesMutualInformation

Purpose	Mattes mutual information metric configuration object
Description	A <code>MattesMutualInformation</code> object describes a mutual information metric configuration that you pass to the function <code>imregister</code> to solve image registration problems.
Construction	<code>metric = registration.metric.MattesMutualInformation()</code> constructs a <code>MattesMutualInformation</code> object.
Properties	<p>NumberOfSpatialSamples</p> <p>Number of spatial samples used to compute the metric.</p> <p><code>NumberOfSpatialSamples</code> is a positive scalar integer value that defines the number of random pixels <code>imregister</code> uses to compute the metric. Your registration results are more reproducible (at the cost of performance) as you increase this value. <code>imregister</code> only uses <code>NumberOfSpatialSamples</code> when <code>UseAllPixels = 0</code> (false). The default value for <code>NumberOfSpatialSamples</code> is 500.</p> <p>NumberOfHistogramBins</p> <p>Number of histogram bins used to compute the metric.</p> <p><code>NumberOfHistogramBins</code> is a positive scalar integer value that defines the number of bins <code>imregister</code> uses to compute the joint distribution histogram. The default value is 50, and the minimum value is 5.</p> <p>UseAllPixels</p> <p>Logical scalar that specifies whether <code>imregister</code> should use all pixels in the overlap region of the images to compute the metric.</p> <p>You can achieve significantly better performance if you set this property to 0 (false). When <code>UseAllPixels = 0</code>, the <code>NumberOfSpatialSamples</code> property controls the number of random pixel locations that <code>imregister</code> uses to compute the metric. The results of your registration might not be reproducible when <code>UseAllPixels = 0</code>. This is because <code>imregister</code> selects a</p>

registration.metric.MattesMutualInformation

random subset of pixels from the images to compute the metric.
The default value for `UseAllPixels` is 1 (true).

Definitions

Mutual Information Metric

Metric used to maximize the number of coincident pixels with the same relative brightness value. This metric is best suited for images with different brightness ranges.

Copy Semantics

Value. To learn how value classes affect copy operations, see [Copying Objects in the MATLAB Programming Fundamentals documentation](#).

Tips

- Larger values of mutual information correspond to better registration results. You can examine the computed values of Mattes mutual information if you enable `'DisplayOptimization'` when you call `imregister`, for example:

```
movingRegistered = imregister(moving, fixed, 'rigid', optimizer, metric, 'DisplayOptimization',
```

Examples

Register MRI Images with MattesMutualInformation Metric


Register two MRI images of a knee that were obtained using different protocols.

Read the images into the workspace.

```
fixed = dicomread('knee1.dcm');  
moving = dicomread('knee2.dcm');
```

View the misaligned images.

```
imshowpair(fixed, moving, 'Scaling', 'joint');
```



Create the optimizer configuration object suitable for registering images from different sensors.

```
optimizer = registration.optimizer.OnePlusOneEvolutionary;
```

Create the MattesMutualInformation metric configuration object.

```
metric = registration.metric.MattesMutualInformation
```

```
metric =
```

```
    registration.metric.MattesMutualInformation
```

Properties:

registration.metric.MattesMutualInformation

```
NumberOfSpatialSamples: 500
NumberOfHistogramBins: 50
UseAllPixels: 1
```

Tune the properties of the optimizer so that the problem will converge on a global maxima. Increase the number of iterations the optimizer will use to solve the problem.


```
optimizer.InitialRadius = 0.009;
optimizer.Epsilon = 1.5e-4;
optimizer.GrowthFactor = 1.01;
optimizer.MaximumIterations = 300;
```

Register the moving and fixed images.

```
movingRegistered = imregister(moving, fixed, 'affine', optimizer, metric)
```

View the registered images.

```
figure;
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



Algorithms

The `imregister` function uses an iterative process to register images. The metric you pass to `imregister` defines the image similarity metric for evaluating the accuracy of the registration. An image similarity metric takes two images and returns a scalar value that describes how similar the images are. The optimizer you pass to `imregister` defines the methodology for minimizing or maximizing the similarity metric.

Mutual information metrics are information theoretic techniques for measuring how related two variables are. These algorithms use the joint probability distribution of a sampling of pixels from two images to measure the certainty that the values of one set of pixels map to similar values in the other image. This information is a quantitative measure of how similar the images are. High mutual information implies a large

registration.metric.MattesMutualInformation

reduction in the uncertainty (entropy) between the two distributions, signaling that the images are likely better aligned.

The Mattes mutual information algorithm uses a single set of pixel locations for the duration of the optimization, instead of drawing a new set at each iteration. The number of samples used to compute the probability density estimates and the number of bins used to compute the entropy are both user selectable. The marginal and joint probability density function is evaluated at the uniformly spaced bins using the samples. Entropy values are computed by summing over the bins. Zero-order and third-order B-spline kernels are used to compute the probability density functions of the fixed and moving images, respectively.[1]

References

[1] Rahunathan, Smriti, D. Stredney, P. Schmalbrock, and B.D. Clymer. Image Registration Using Rigid Registration and Maximization of Mutual Information. Poster presented at: MMVR13. The 13th Annual Medicine Meets Virtual Reality Conference; 2005 January 26–29; Long Beach, CA.

[2] D. Mattes, D.R. Haynor, H. Vesselle, T. Lewellen, and W. Eubank. "Non-rigid multimodality image registration." (Proceedings paper). *Medical Imaging 2001: Image Processing*. SPIE Publications, 3 July 2001. pp. 1609–1620.

Alternatives

Use `imregconfig` to construct a metric configuration for typical image registration scenarios.

See Also

`registration.metric.MeanSquares` | `imregister`

Concepts

- Class Attributes
- Property Attributes

Purpose	Mean square error metric configuration object
Description	A MeanSquares object describes a mean square error metric configuration that you pass to the function <code>imregister</code> to solve image registration problems.
Construction	<code>metric = registration.metric.MeanSquares()</code> constructs a MeanSquares object.
Copy Semantics	Value. To learn how value classes affect copy operations, see Copying Objects in the MATLAB Programming Fundamentals documentation.
Tips	<ul style="list-style-type: none">• This metric is an element-wise difference between two input images. The ideal value is zero. You can examine the computed values of mean square error if you enable 'DisplayOptimization' when you call <code>imregister</code>. For example, <code>movingRegistered = imregister(moving, fixed, 'rigid', optimizer, metric)</code>

Examples

Register Remote Sensing Images with MeanSquares Metric

Create a MeanSquares object and use it to register two images captured with different sensors.

In general, `imregister` doesn't support perspective transformations. However it returns good results for this problem, which uses a similarity transformation.

Read the images into the workspace.

```
fixed = imread('westconcordorthophoto.png');  
moving = rgb2gray(imread('westconcordaerial.png'));
```

View the misaligned images.

```
imshowpair(fixed, moving, 'Scaling', 'joint');
```

registration.metric.MeanSquares



Create the optimizer configuration object suitable for registering images from different sensors.

```
optimizer = registration.optimizer.OnePlusOneEvolutionary;
```

Create the MeanSquares metric configuration object. Even though the images came from different sensors, they have an intensity relationship similar enough to use mean square error as the similarity metric.

```
metric = registration.metric.MeanSquares
```

```
metric =
```



```
registration.metric.MeanSquares
```

This class has no properties.

Increase `MaximumIterations` property of the optimizer to allow for more iterations.

```
optimizer.MaximumIterations = 1000;
```

Register the moving and fixed images.

```
movingRegistered = imregister(moving, fixed, 'similarity', optimizer);
```

View the registered images.

```
figure;  
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



Algorithms

The `imregister` function uses an iterative process to register images. The metric you pass to `imregister` defines the image similarity metric for evaluating the accuracy of the registration. An image similarity metric takes two images and returns a scalar value that describes how similar the images are. The optimizer you pass to `imregister` defines the methodology for minimizing or maximizing the similarity metric.

The mean squares image similarity metric is computed by squaring the difference of corresponding pixels in each image and taking the mean of the those squared differences.

Use `imregconfig` to construct a metric configuration for typical image registration scenarios.

See Also

`registration.metric.MattesMutualInformation` | `imregister`

registration.optimizer.OnePlusOneEvolutionary

- Purpose** One-plus-one evolutionary optimizer configuration object
- Description** A `OnePlusOneEvolutionary` object describes a one-plus-one evolutionary optimization configuration that you pass to the function `imregister` to solve image registration problems.
- Construction** `optimizer = registration.optimizer.OnePlusOneEvolutionary()`
Constructs a `OnePlusOneEvolutionary` object.

Properties

GrowthFactor

Growth factor of the search radius.

`GrowthFactor` is a positive scalar value that the optimizer uses to control the rate at which the search radius grows in parameter space. If you set `GrowthFactor` to a large value, the optimization is fast, but it might result in finding only the metric's local extrema. If you set `GrowthFactor` to a small value, the optimization is slower, but it is likely to converge on a better solution. The default value of `GrowthFactor` is 1.05.

Epsilon

Minimum size of the search radius.

`Epsilon` is a positive scalar value that controls the accuracy of convergence by adjusting the minimum size of the search radius. If you set `Epsilon` to a small value, the optimization of the metric is more accurate, but the computation takes longer. If you set `Epsilon` to a large value, the computation time decreases at the expense of accuracy. The default value of `Epsilon` is $1.5e-6$.

InitialRadius

Initial size of search radius.

`InitialRadius` is a positive scalar value that controls the initial search radius of the optimizer. If you set `InitialRadius` to a large value, the computation time decreases. However, overly large

values of `InitialRadius` might result in an optimization that fails to converge. The default value of `InitialRadius` is $6.25e-3$.

MaximumIterations

Maximum number of optimizer iterations.

`MaximumIterations` is a positive scalar integer value that determines the maximum number of iterations the optimizer performs at any given pyramid level. The registration could converge before the optimizer reaches the maximum number of iterations. The default value of `MaximumIterations` is 100.

Copy Semantics

Value. To learn how value classes affect copy operations, see [Copying Objects in the MATLAB Programming Fundamentals documentation](#).

Examples

Register MRI Images with OnePlusOneEvolutionary Optimizer

Register two MRI images of a knee that were obtained using different protocols.


Read the images into the workspace.

```
fixed = dicomread('knee1.dcm');  
moving = dicomread('knee2.dcm');
```

View the misaligned images.

```
imshowpair(fixed, moving, 'Scaling', 'joint');
```

registration.optimizer.OnePlusOneEvolutionary



Create the optimizer configuration object suitable for registering images from different sensors.

```
optimizer = registration.optimizer.OnePlusOneEvolutionary
```

```
optimizer =
```

```
    registration.optimizer.OnePlusOneEvolutionary
```

```
Properties:
```

```
    GrowthFactor: 1.050000e+00
```

```
    Epsilon: 1.500000e-06
```

```
    InitialRadius: 6.250000e-03
```

```
    MaximumIterations: 100
```

Create the MattesMutualInformation metric configuration object.

```
metric = registration.metric.MattesMutualInformation;
```

Tune the properties of the optimizer so that the problem will converge on a global maxima. Increase the number of iterations the optimizer will use to solve the problem.


```
optimizer.InitialRadius = 0.009;  
optimizer.Epsilon = 1.5e-4;  
optimizer.GrowthFactor = 1.01;  
optimizer.MaximumIterations = 300;
```

Register the moving and fixed images.

```
movingRegistered = imregister(moving, fixed, 'affine', optimizer, metric);
```

View the registered images.

```
figure  
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



Algorithms

The `imregister` function uses an iterative process to register images. The metric you pass to `imregister` defines the image similarity metric for evaluating the accuracy of the registration. An image similarity metric takes two images and returns a scalar value that describes how similar the images are. The optimizer you pass to `imregister` defines the methodology for minimizing or maximizing the similarity metric.

An evolutionary algorithm iterates to find a set of parameters that produce the best possible registration result. It does this by perturbing, or mutating, the parameters from the last iteration (the parent). If the new (child) parameters yield a better result, then the child becomes the new parent whose parameters are perturbed, perhaps more aggressively. If the parent yields a better result, it remains the parent and the next perturbation is less aggressive.

References

[1] Styner, M., C. Brechbuehler, G. Székely, and G. Gerig. "Parametric estimate of intensity inhomogeneities applied to MRI." *IEEE Transactions on Medical Imaging*. Vol. 19, Number 3, 2000, pp. 153-165.

Alternatives

Use `imregconfig` to construct an optimizer configuration for typical image registration scenarios.

See Also

`registration.optimizer.RegularStepGradientDescent` |
`imregister`

Concepts

- Class Attributes
- Property Attributes

registration.optimizer.RegularStepGradientDescent

Purpose Regular step gradient descent optimizer configuration object

Description A `RegularStepGradientDescent` object describes a regular step gradient descent optimization configuration that you pass to the function `imregister` to solve image registration problems.

Construction `optimizer = registration.optimizer.RegularStepGradientDescent()` constructs a `RegularStepGradientDescent` object.

Properties **GradientMagnitudeTolerance**

Gradient magnitude tolerance.

`GradientMagnitudeTolerance` is a positive scalar value that controls the optimization process. When the value of the gradient is smaller than `GradientMagnitudeTolerance`, it is an indication that the optimizer might have reached a plateau. The default value of `GradientMagnitudeTolerance` is $1e-4$.

MinimumStepLength

Tolerance for convergence.

`MinimumStepLength` is a positive scalar value that controls the accuracy of convergence. If you set `MinimumStepLength` to a small value, the optimization takes longer to compute, but it is likely to converge on a more accurate metric value. The default value of `MinimumStepLength` is $1e-5$.

MaximumStepLength

Initial step length.

`MaximumStepLength` is a positive scalar value that controls the initial step length used in optimization. If you set `MaximumStepLength` to a large value, the computation time decreases. However, the optimizer might fail to converge if you set `MaximumStepLength` to an overly large value. The default value of `MaximumStepLength` is 0.0625 .

MaximumIterations

Maximum number of iterations.

MaximumIterations is a positive scalar integer value that determines the maximum number of iterations the optimizer performs at any given pyramid level. The registration could converge before the optimizer reaches the maximum number of iterations. The default value of MaximumIterations is 100.

RelaxationFactor

Step length reduction factor.

RelaxationFactor is a scalar value between 0 and 1 that defines the rate at which the optimizer reduces step size during convergence. Whenever the optimizer determines that the direction of the gradient changed, it reduces the size of the step length. If your metric is noisy, you can set RelaxationFactor to a larger value. This leads to a more stable convergence at the expense of computation time. The default value of RelaxationFactor is 0.5.

Copy Semantics

Value. To learn how value classes affect copy operations, see Copying Objects in the MATLAB Programming Fundamentals documentation.

Examples

Register Images with RegularStepGradientDescent Optimizer

Create a RegularStepGradientDescent object and use it to register two images captured with the same device.

Read the reference image and create an unregistered copy.

```
fixed = imread('pout.tif');  
moving = imrotate(fixed, 5, 'bilinear', 'crop');
```

View the misaligned images.

```
imshowpair(fixed, moving, 'Scaling', 'joint');
```

registration.optimizer.RegularStepGradientDescent



Create the optimizer configuration object suitable for registering images from the same device.

```
optimizer = registration.optimizer.RegularStepGradientDescent
```

Create the metric configuration object.

```
metric = registration.metric.MeanSquares;
```

Now modify the optimizer configuration to get more precision.

```
optimizer.MaximumIterations = 300;  
optimizer.MinimumStepLength = 5e-4;
```

Perform the registration.

```
movingRegistered = imregister(moving, fixed, 'rigid', optimizer, metric);
```

View registered images.

```
figure  
imshowpair(fixed, movingRegistered, 'Scaling', 'joint');
```



Algorithms

The `imregister` function uses an iterative process to register images. The metric you pass to `imregister` defines the image similarity metric for evaluating the accuracy of the registration. An image similarity metric takes two images and returns a scalar value that describes how similar the images are. The optimizer you pass to `imregister` defines the methodology for minimizing or maximizing the similarity metric.

The regular step gradient descent optimization adjusts the transformation parameters so that the optimization follows the gradient of the image similarity metric in the direction of the extrema. It uses constant length steps along the gradient between computations until the gradient changes direction, at which point the step length is halved.

Use `imregconfig` to construct an optimizer configuration for typical image registration scenarios.

registration.optimizer.RegularStepGradientDescent

See Also

`registration.optimizer.OnePlusOneEvolutionary` | `imregister`

Concepts

- Class Attributes
- Property Attributes

Purpose Convert RGB image or colormap to grayscale

Syntax `I = rgb2gray(RGB)`
`newmap = rgb2gray(map)`

Description `I = rgb2gray(RGB)` converts the truecolor image `RGB` to the grayscale intensity image `I`. `rgb2gray` converts RGB images to grayscale by eliminating the hue and saturation information while retaining the luminance.

`newmap = rgb2gray(map)` returns a grayscale colormap equivalent to `map`.

Note A grayscale image is also called a gray-scale, gray scale, or gray-level image.

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`.

Examples 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);
```

Algorithms `rgb2gray` converts RGB values to grayscale values by forming a weighted sum of the *R*, *G*, and *B* components:

rgb2gray

$$0.2989 * R + 0.5870 * G + 0.1140 * B$$

Note that these are the same weights used by the `rgb2ntsc` function to compute the *Y* component.

See Also

`ind2gray` | `mat2gray` | `ntsc2rgb` | `rgb2ind` | `rgb2ntsc`

Purpose Convert RGB color values to NTSC color space

Syntax `yiqmap = rgb2ntsc(rgbmap)`
`YIQ = rgb2ntsc(RGB)`

Description `yiqmap = rgb2ntsc(rgbmap)` converts the m-by-3 RGB values in `rgbmap` to NTSC color space. `yiqmap` is an m-by-3 matrix that contains the NTSC luminance (*Y*) and chrominance (*I* and *Q*) color components as columns that are equivalent to the colors in the RGB colormap.

`YIQ = rgb2ntsc(RGB)` converts the truecolor image `RGB` to the equivalent NTSC image `YIQ`.

Tips 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.

`rgb2ntsc` defines the NTSC components using

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Class Support `RGB` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. `RGBMAP` can be `double`. The output is `double`.

See Also `ntsc2rgb` | `rgb2ind` | `ind2rgb` | `ind2gray`

rgb2ycbcr

Purpose	Convert RGB color values to YCbCr color space
Syntax	<pre>ycbcrmap = rgb2ycbcr(map) YCBCR = rgb2ycbcr(RGB)</pre>
Description	<p><code>ycbcrmap = rgb2ycbcr(map)</code> converts the RGB values in <code>map</code> to the YCbCr color space. <code>map</code> must be an M-by-3 array. <code>ycbcrmap</code> is an M-by-3 matrix that contains the YCbCr luminance (<i>Y</i>) and chrominance (<i>Cb</i> and <i>Cr</i>) color values as columns. Each row in <code>ycbcrmap</code> represents the equivalent color to the corresponding row in the RGB colormap, <code>map</code>.</p> <p><code>YCBCR = rgb2ycbcr(RGB)</code> converts the truecolor image <code>RGB</code> to the equivalent image in the YCbCr color space. <code>RGB</code> must be a M-by-N-by-3 array.</p> <p>If the input is <code>uint8</code>, <code>YCBCR</code> is <code>uint8</code>, where <i>Y</i> is in the range [16 235], and <i>Cb</i> and <i>Cr</i> are in the range [16 240]. If the input is a <code>double</code>, <i>Y</i> is in the range [16/255 235/255] and <i>Cb</i> and <i>Cr</i> are in the range [16/255 240/255]. If the input is <code>uint16</code>, <i>Y</i> is in the range [4112 60395] and <i>Cb</i> and <i>Cr</i> are in the range [4112 61680].</p>
Class Support	If the input is an RGB image, it can be of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . If the input is a colormap, it must be <code>double</code> . The output image is of the same class as the input image.
Examples	<p>Convert RGB image to YCbCr.</p> <pre>RGB = imread('board.tif'); YCBCR = rgb2ycbcr(RGB);</pre> <p>Convert RGB color space to YCbCr.</p> <pre>map = jet(256); newmap = rgb2ycbcr(map);</pre>
References	[1] Poynton, C. A. <i>A Technical Introduction to Digital Video</i> , John Wiley & Sons, Inc., 1996, p. 175.

[2] Rec. ITU-R BT.601-5, *Studio Encoding Parameters of Digital Television for Standard 4:3 and Wide-screen 16:9 Aspect Ratios*, (1982-1986-1990-1992-1994-1995), Section 3.5.

See Also

[ntsc2rgb](#) | [rgb2ntsc](#) | [ycbcr2rgb](#)

roicolor

Purpose Select region of interest (ROI) based on color

Syntax `BW = roicolor(A,low,high)`
`BW = roicolor(A,v)`

Description `roicolor` selects a region of interest (ROI) 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 an ROI selected as those pixels that lie within the colormap range `[low high]`.

`BW = (A >= low) & (A <= high)`

`BW` is a binary image with 0's outside the region of interest and 1's inside.

`BW = roicolor(A,v)` returns an ROI 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`.

Examples

```
load clown
BW = roicolor(X,10,20);
imshow(X,map)
figure,imshow(BW)
```



See Also

`roifilt2` | `roipoly`

Purpose Fill in specified region of interest (ROI) polygon in grayscale image

Syntax

```
J = roifill
J = roifill(I)
J = roifill(I, c, r)
J = roifill(I, BW)
[J,BW] = roifill(...)
J = roifill(x, y, I, xi, yi)
[x, y, J, BW, xi, yi] = roifill(...)
```

Description Use `roifill` to fill in a specified region of interest (ROI) polygon in a grayscale image. `roifill` smoothly interpolates inward from the pixel values on the boundary of the polygon by solving Laplace's equation. The boundary pixels are not modified. `roifill` can be used, for example, to erase objects in an image.

`J = roifill` creates an interactive polygon tool, associated with the image displayed in the current figure, called the target image. You use the mouse to define the ROI – see “Interactive Behavior” on page 3-777. When you are finished defining the ROI, fill in the area specified by the ROI by double-clicking inside the region or by right-clicking anywhere inside the region and selecting **Fill Area** from the context menu. `roifill` returns the image, `J`, which is the same size as `I` with the region filled in (see “Examples” on page 3-780).

Note If you do not specify an output argument, `roifill` displays the filled image in a new figure.

`J = roifill(I)` displays the image `I` and creates an interactive polygon tool associated with the 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, 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.




`[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.


Interactive Behavior

When you call `roifill` with an interactive syntax, the pointer changes to a cross hairs shape  when you move it over the target image. Using the mouse, you specify a region-of-interest by selecting vertices of a polygon. You can change the size or shape of the polygon using the mouse. The following figure illustrates a polygon defined by multiple vertices. For more information about all the interactive capabilities of `roifill`, see the table that follows.



Interactive Behavior	Description
Closing the polygon. (Completing the region-of-interest.)	Use any of the following mechanisms: <ul style="list-style-type: none">• Move the pointer over the initial vertex of the polygon that you selected. The shape changes to a circle ○. Click either mouse button.• Double-click the left mouse button. This action creates a vertex at the point under the mouse and draws a straight line connecting this vertex with the initial vertex.• Click the right mouse button. This action draws a line connecting the last vertex

Interactive Behavior	Description
	selected with the initial vertex; it does not create a new vertex.
Deleting the polygon	<p>Press Backspace, Escape or Delete, or right-click inside the region and select Cancel from the context menu.</p> <p>Note: If you delete the ROI, the function returns empty values.</p>
Moving the polygon	<p>Move the pointer inside the region. The pointer changes to a fleur . Click and drag the mouse to move the polygon.</p>
Changing the color of the polygon	<p>Move the pointer inside the region. Right-click and select Set color from the context menu.</p>
Adding a new vertex.	<p>Move the pointer over an edge of the polygon and press the A key. The shape of the pointer changes . Click the left mouse button to create a new vertex at that position on the line.</p>
Moving a vertex. (Reshaping the region-of-interest.)	<p>Move the pointer over a vertex. The pointer changes to a circle . Click and drag the vertex to its new position.</p>

Interactive Behavior	Description
Deleting a vertex.	Move the pointer over a vertex. The pointer changes to a circle  . Right-click and select Delete Vertex from the context menu. This action deletes the vertex and adjusts the shape of the polygon, drawing a new straight line between the two vertices that were neighbors of the deleted vertex.
Retrieving the coordinates of the vertices	Move the pointer inside the region. Right-click and select Copy position from the context menu to copy the current position to the Clipboard. Position is an n -by-2 array containing the x - and y -coordinates of each vertex, where n is the number of vertices you selected.

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`.

Examples

This example uses `roifill` to fill a region in the input image, I . For more examples, especially of the interactive syntaxes, see “Filling an ROI”.

```
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

[impoly](#) | [roifilt2](#) | [roipoly](#)

roifilt2

Purpose Filter region of interest (ROI) in image

Syntax
`J = roifilt2(h, I, BW)`
`J = roifilt2(I, BW, fun)`

Description `J = roifilt2(h, I, BW)` filters the data in `I` with the two-dimensional linear filter `h`. `BW` is a binary image the same size as `I` that defines an ROI used as a mask for filtering. `roifilt2` returns an image that consists of filtered values for pixels in locations where `BW` contains 1's, and unfiltered values for pixels in locations where `BW` contains 0's. For this syntax, `roifilt2` calls `filter2` to implement the filter.

`J = roifilt2(I, BW, fun)` processes the data in `I` using the function `fun`. The result `J` contains computed values for pixels in locations where `BW` contains 1's, and the actual values in `I` for pixels in locations where `BW` contains 0's. `fun` must be a function handle. Parameterizing Functions, in the MATLAB Mathematics documentation, explains how to provide additional parameters to the function `fun`.

Class Support For the syntax that includes a filter `h`, the input image can be logical or numeric, and the output array `J` has the same class as the input image. For the syntax that includes a function, `I` can be of any class supported by `fun`, and the class of `J` depends on the class of the output from `fun`.

Examples This example continues the `roipoly` example, filtering the region of the image `I` specified by the mask `BW`. The `roifilt2` function returns the filtered image `J`, shown in the following figure.

```
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);  
figure, imshow(I), figure, imshow(J)
```

**See Also**

`filter2` | `function_handle` | `imfilter` | `roipoly`

How To

- “Anonymous Functions”
- “Parameterizing Functions”

roipoly


Purpose Specify polygonal region of interest (ROI)

Syntax



```
BW = roipoly
BW = roipoly(I)
BW = roipoly(I, c, r)
BW = roipoly(x, y, I, xi, yi)
[BW, xi, yi] = roipoly(...)
[x, y, BW, xi, yi] = roipoly(...)
```




Description Use `roipoly` to specify a polygonal region of interest (ROI) within an image. `roipoly` returns a binary image that you can use as a mask for masked filtering.

`BW = roipoly` creates an interactive polygon tool, associated with the image displayed in the current figure, called the target image. With the polygon tool active, the pointer changes to cross hairs \oplus when you move the pointer over the image in the figure. Using the mouse, you specify the region by selecting vertices of the polygon. You can move or resize the polygon using the mouse. The following figure illustrates a polygon defined by multiple vertices. The following table describes all the interactive behavior of the polygon tool.



When you are finished positioning and sizing the polygon, create the mask by double-clicking, or by right-clicking inside the region and selecting **Create mask** from the context menu. roipoly returns the mask as a binary image, BW, the same size as I. In the mask image, roipoly sets pixels inside the region to 1 and pixels outside the region to 0.

Interactive Behavior	Description
Closing the polygon. (Completing the region-of-interest.)	Use any of the following mechanisms: <ul style="list-style-type: none"> • Move the pointer over the initial vertex of the polygon that you selected. The pointer changes to a circle . Click either mouse button. • Double-click the left mouse button. This action creates a vertex at the point under the mouse pointer and draws a straight line connecting this vertex with the initial vertex. • Right-click the mouse. This draws a line connecting the last vertex selected with the initial vertex; it does not create a new vertex at the point under the mouse.
Moving the entire polygon	Move the pointer inside the region. The pointer changes to a fleur shape  . Click and drag the polygon over the image.
Deleting the polygon	Press Backspace , Escape or Delete , or right-click inside the region and select Cancel from the context menu. Note: If you delete the ROI, the function returns empty values.

Interactive Behavior	Description
Moving a vertex. (Reshaping the region-of-interest.)	Move the pointer over a vertex. The pointer changes to a circle  . Click and drag the vertex to its new position.
Adding a new vertex.	Move the pointer over an edge of the polygon and press the A key. The pointer changes shape to  . Click the left mouse button to create a new vertex at that point on the edge.
Deleting a vertex. (Reshaping the region-of-interest.)	Move the pointer over the vertex. The pointer changes to a circle  . Right-click and select Delete vertex from the context menu. roipoly draws a new straight line between the two vertices that were neighbors of the deleted vertex.
Changing the color of the polygon	Move the pointer anywhere inside the boundary of the region and click the right mouse button. Select Set color from the context menu.
Retrieving the coordinates of the vertices	Move the pointer inside the region. Right-click and select Copy position from the context menu to copy the current position to the Clipboard. The position is an n -by-2 array containing the x - and y -coordinates of each vertex, where n is the number of vertices.

Note If you call roipoly without specifying any output arguments, roipoly displays the resulting mask image in a new figure window.

`BW = roipoly(I)` displays the image `I` and creates an interactive polygon tool associated with that image.

`BW = roipoly(I, c, r)` returns the ROI specified by the polygon described by vectors `c` and `r`, which specify the column and row indices of each vertex, respectively. `c` and `r` must be the same size.

`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 x - and y -coordinates of the polygon vertices in `xi` and `yi`.

Note `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`.

Class Support

The input image `I` can be of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image `BW` is of class `logical`. All other inputs and outputs are of class `double`.

Tips

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.

Examples

Use `roipoly` to create a mask image, `BW`, the same size as the input image, `I`. The example in `roifilt2` continues this example, filtering the specified region in the image. For more examples, especially of the interactive syntaxes, see “Specifying a Region of Interest (ROI)”.

```
I = imread('eight.tif');  
c = [222 272 300 270 221 194];  
r = [21 21 75 121 121 75];  
BW = roipoly(I,c,r);  
figure, imshow(I)  
figure, imshow(BW)
```



See Also

[impoly](#) | [poly2mask](#) | [roifilt2](#) | [roicolor](#) | [roifill](#)

Purpose

Create reduced resolution data set from image file

Syntax

```
rsetfile = rsetwrite(File_Name)
rsetfile = rsetwrite(File_Name, output_filename)
rsetfile = rsetwrite(adapter, output_filename)
```

Description

`rsetfile = rsetwrite(File_Name)`, where `File_Name` is a TIFF or NITF image file, creates a reduced resolution data set (R-Set) from the specified file. The R-Set file is written to the current working directory with a name based on the input file name. For example, if `File_Name` is `'VeryLargeImage.tiff'`, `rsetfile` will be `'VeryLargeImage.rset'`. If an image file contains multiple images, only the first one is used.

`rsetfile = rsetwrite(File_Name, output_filename)` creates an R-Set from the specified image file, using `output_filename` as the name of the new file. In this case, `rsetfile` and `output_filename` contain the same string.

`rsetfile = rsetwrite(adapter, output_filename)` creates an R-Set from the specified Image Adapter object, `adapter`. Image Adapters are user-defined classes that provide `rsetwrite` a common API for reading a particular image file format. See the documentation for `ImageAdapter` for more details.

Tips

`rsetwrite` creates an R-Set file by dividing an image into spatial tiles and resampling the image at different resolution levels. When you open the R-Set file in the Image Tool and zoom in, you view tiles at a higher resolution. When you zoom out, you view tiles at a lower resolution. In this way, clarity of the image and memory usage are balanced for optimal performance. The R-Set file contains a compressed copy of the full-resolution data.

Because R-Set creation can be time consuming, a progress bar shows the status of the operation. If you cancel the operation, processing stops, no file is written, and the `rsetfile` variable will be empty.

`rsetwrite` supports NITF image files that are uncompressed and Version 2.0 or higher. It does not support NITF files with more than

three bands or with floating point data. Images with more than one data band are OK if they contain unsigned integer data.

While it is possible to create an R-Set from an image where the dimensions are smaller than the size of a single R-Set tile, the resulting R-set file will likely be larger and take longer to load than the original file. The current size of an R-Set tile is 512 x 512 pixels.

Examples

Example 1: Create an R-Set File

Visualize a very large image by using an R-Set. Replace 'MyReallyBigImage.tif' in the example below with the name of your file:

```
big_file = 'MyReallyBigImage.tif';
rset_file = rsetwrite(big_file);
imtool(rset_file)
```

Example 2: Convert TIFF Files to R-Set Files

Create R-Set files for every TIFF in a directory containing very large images. Put the R-Set files into a temporary directory:

```
d = dir('*.tif*');
image_dir = pwd;
cd(tempdir)
for p = 1:numel(d)
    big_file = fullfile(image_dir, d(p).name);
    rsetwrite(big_file);
end
```

See Also

[imread](#) | [imtool](#)

Purpose	Standard deviation of matrix elements
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.
Algorithms	<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>

stdfilt

Purpose Local standard deviation of 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` method to extract the neighborhood from the structuring element object.

Examples

```
I = imread('circuit.tif');  
J = stdfilt(I);  
imshow(I);  
figure, imshow(J,[]);
```

See Also `entropyfilt` | `getnhood` | `rangefilt` | `std2` | `strel`

Purpose Create morphological structuring element (STREL)

Syntax

```
SE = strel(shape, parameters)
SE = strel('arbitrary', NHOOD)
SE = strel('arbitrary', NHOOD, HEIGHT)
SE = strel('ball', R, H, N)
SE = strel('diamond', R)
SE = strel('disk', R, N)
SE = strel('line', LEN, DEG)
SE = strel('octagon', R)
SE = strel('pair', OFFSET)
SE = strel('periodicline', P, V)
SE = strel('rectangle', MN)
SE = strel('square', W)
```

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

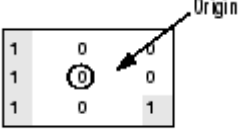
'arbitrary'	'pair'
'diamond'	'periodicline'
'disk'	'rectangle'
'line'	'square'
'octagon'	

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 = 
NHOOD = [1 0 0; 1 0 0; 1 0 1];

`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 that is symmetric with respect to the neighborhood center. 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 $1*V$, $-1*V$, $2*V$, $-2*V$, ..., $P*V$, $-P*V$.



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
<code>getheight</code>	Get height of structuring element
<code>getneighbors</code>	Get structuring element neighbor locations and heights
<code>getnhood</code>	Get structuring element neighborhood
<code>getsequence</code>	Extract sequence of decomposed structuring elements
<code>isflat</code>	Return true for flat structuring element
<code>reflect</code>	Reflect structuring element
<code>translate</code>	Translate structuring element

Examples

```
se1 = strel('square',11)           % 11-by-11 square
se2 = strel('line',10,45)         % length 10, angle 45 degrees
se3 = strel('disk',15)           % disk, radius 15
se4 = strel('ball',15,5)         % ball, radius 15, height 5
```

Algorithms

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 spheres" [2].

References

- [1] van den Boomgard, R, and R. van Balen, "Methods for Fast Morphological Image Transforms Using Bitmapped Images," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 54, Number 3, pp. 252–254, May 1992.
- [2] Adams, R., "Radial Decomposition of Discs and Spheres," *Computer Vision, Graphics, and Image Processing: Graphical Models and Image Processing*, Vol. 55, Number 5, pp. 325–332, September 1993.
- [3] Jones, R., and P. Soille, "Periodic lines: Definition, cascades, and application to granulometrie," *Pattern Recognition Letters*, Vol. 17, pp. 1057–1063, 1996.

See Also

`imdilate` | `imerode`

stretchlim

Purpose Find limits to contrast stretch image

Syntax

```
LOW_HIGH = stretchlim(I)
LOW_HIGH = stretchlim(I, TOL)
LOW_HIGH = stretchlim(RGB, TOL)
```

Description `LOW_HIGH = stretchlim(I)` returns `LOW_HIGH`, a two-element vector of pixel values that specify lower and upper limits that can be used for contrast stretching image `I`. By default, values in `LOW_HIGH` specify the bottom 1% and the top 1% of all pixel values. The gray values returned can be used by the `imadjust` function to increase the contrast of an image.

`LOW_HIGH = stretchlim(I, TOL)` where `TOL` is a two-element vector `[LOW_FRACT HIGH_FRACT]` that specifies the fraction of the image to saturate at low and high pixel values.

If `TOL` is a scalar, `LOW_FRACT = TOL`, and `HIGH_FRACT = 1 - LOW_FRACT`, which saturates equal fractions at low and high pixel values.

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.

Note If `TOL` is too big, such that no pixels would be left after saturating low and high pixel values, `stretchlim` returns `[0 1]`.

Class Support The input image can be of class `uint8`, `uint16`, `int16`, `double`, or `single`. The output limits returned, `LOW_HIGH`, are of class `double` and have values between 0 and 1.

Examples

```
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 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 truecolor 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`.

Examples

```
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`

Purpose Apply spatial transformation to N-D 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.

tformarray

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

$$[D1 \ D2 \ D3 \ \dots \ DN \ L]$$

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`:

- `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`.

Examples

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](#)

tformfwd

Purpose Apply forward spatial transformation

Syntax

```
[X, Y] = tformfwd(T, U, V)
[X1, X2, X3, ...] = tformfwd(T, U1, U2, U3, ...)
X = tformfwd(T, U)
X = tformfwd(T, U)
[X1, X2, X3, ...] = tformfwd(T, U)
X = tformfwd(T, U1, U2, U3, ...)
X = tformfwd(U, T)
```

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) ... UNDIMS_IN(k)]`

to the point

`[X1(k) X2(k) ... 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.

$X = \text{tformfwd}(T, U)$, where U is an $(N+1)$ -dimensional array, maps the point $U(k_1, k_2, \dots, k_N, :)$ to the point $X(k_1, k_2, \dots, k_N, :)$. $\text{size}(U, N+1)$ must equal $NDIMS_IN$. 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.

Examples

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`

tforminv

Purpose Apply inverse spatial transformation

Syntax

```
[U,V] = tforminv(T,X,Y)
[U1,U2,U3,...] = tforminv(T,X1,X2,X3,...)
U = tforminv(T,X)
[U1,U2,U3,...] = tforminv(T,X)
U = tforminv(T,X1,X2,X3,...)
```

Description `[U,V] = 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.ndims_in` and `T.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,...] = tforminv(T,X1,X2,X3,...)` applies the NDIMS_OUT-to-NDIMS_IN inverse transformation defined in TFORM structure T to the coordinate arrays `X1,X2,...,XNDIMS_OUT` (where `NDIMS_IN = T.ndims_in` and `NDIMS_OUT = T.ndims_out`). The number of output arguments must equal `NDIMS_IN`. The transformation maps the point

$$[X1(k) \ X2(k) \ \dots \ XNDIMS_OUT(k)]$$

to the point

$$[U1(k) \ U2(k) \ \dots \ UNDIMS_IN(k)].$$

`X1,X2,X3,...` can have any dimensionality, but must be the same size.

`U1,U2,U3,...` have this size also.

`U = 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 = tforminv(T,X)`, where `X` is an $(N+1)$ -dimensional array, maps the point `X(k1,k2,...,kN,:)` to the point `U(k1,k2,...,kN,:)`. `size(X,N+1)` must equal `NDIMS_OUT`. `U` is an $(N+1)$ -dimensional array, with `size(U,I)` equal to `size(X,I)` for $I = 1, \dots, N$ and `size(U,N+1)` equal to `NDIMS_IN`.

`[U1,U2,U3,...] = tforminv(T,X)` maps an $(N+1)$ -dimensional array to `NDIMS_IN` equally-sized N -dimensional arrays.

`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.

Examples

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` | `tformfwd` | `maketform` | `fliptform`

tonemap

Purpose Render high dynamic range image for viewing

Syntax
RGB = tonemap(HDR)
RGB = tonemap(HDR, param1, val1, ...)

Description RGB = tonemap(HDR) converts the high dynamic range image HDR to a lower dynamic range image, RGB, suitable for display, using a process called tone mapping. Tone mapping is a technique used to approximate the appearance of high dynamic range images on a display with a more limited dynamic range.

RGB = tonemap(HDR, param1, val1, ...) performs tone mapping where parameters control various aspects of the operation. The following table lists these parameters.

Parameter	Description
'AdjustLightness'	A two-element vector in the form [low high] that specifies the overall lightness of the rendered image. low and high are luminance values of the low dynamic range image, in the range [0, 1]. These values are passed to imadjust.
'AdjustSaturation'	A numeric value that specifies the saturation of colors in the rendered image. When the value is greater than 1, the colors are more saturated. When the value is in the range [0, 1) colors are less saturated.
'NumberOfTiles'	A two-element vector of the form [rows cols] that specifies the number of tiles used during the adaptive histogram equalization part of the tone mapping operation. rows and cols specify the number of tile rows and columns. Both rows and cols must be at least 2. The total number of image tiles is equal to rows * cols. A larger number of tiles results in an

Parameter	Description
	image with greater local contrast. The default for rows and cols is 4.

Class Support

The high dynamic range image HDR must be a m-by-n-by-3 single or double array. The output image RGB is an m-by-n-by-3 uint8 image.

Examples

Load a high dynamic range image, convert it to a low dynamic range image while deepening shadows and increasing saturation, and display the results.


```
hdr = hdrread('office.hdr');
imshow(hdr)
rgb = tonemap(hdr, 'AdjustLightness', [0.1 1], ...
               'AdjustSaturation', 1.5);
figure;
imshow(rgb)
```

See Also

[adapthisteq](#) | [hdrread](#) | [stretchlim](#)

translate

Purpose	Translate structuring element (STREL)
Syntax	<code>SE2 = translate(SE,V)</code>
Description	<p><code>SE2 = translate(SE,V)</code> translates the structuring element <code>SE</code> in N-D space. <code>SE</code> is an array of structuring elements, created using the <code>strel</code> function.</p> <p><code>V</code> is an N-element vector that specifies the offsets of the desired translation in each dimension, relative to the structuring element's origin. If you specify an array, <code>translate</code> reshapes the array into a vector.</p> <p><code>SE2</code> is an array of structuring elements the same size as <code>SE</code>. Each individual structuring element in <code>SE2</code> is the translation of the corresponding structuring element in <code>SE</code>.</p>
Class Support	<code>SE</code> and <code>SE2</code> are STREL objects; <code>V</code> is a vector of doubles that must contain only integers.
Examples	<p>Translate a 3-by-3 structuring element.</p> <pre>se = strel(ones(3)) se2 = translate(se,[-2 -2])</pre> <p>The following figure shows the original structuring element and the translated structuring element.</p>



Dilating with a translated version of `strel(1)` is a way to translate an input image in space by an integer number of pixels. 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 image

Syntax `truesize(fig,[mrows ncols])`

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 having one screen pixel for each image pixel.

If you do not specify a figure, `truesize` uses the current figure.

Examples Fit image to figure window.

```
imshow(checkerboard,'InitialMagnification','fit')
```

Resize image and figure to show image at its 80-by-80 pixel size.

```
truesize
```

See Also `imshow` | `iptsetpref` | `iptgetpref`

Purpose Compute new values of A based on lookup table (LUT)

Syntax `B = uintlut(A,LUT)`

uintlut has been removed. Use intlut instead.

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.

Examples

```
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`

viscircles

Purpose Create circle

Syntax
`viscircles(centers, radii)`
`viscircles(ax, centers, radii)`
`h = viscircles(ax, centers, radii)`
`h = viscircles(___, Name, Value)`

Description `viscircles(centers, radii)` draws circles with specified centers and radii onto the current axes.

`viscircles(ax, centers, radii)` draws circles onto the axes specified by `ax`.

`h = viscircles(ax, centers, radii)` returns a vector of `hggroup` handles to each circle. These handles are children of the axes object, `ax`.

`h = viscircles(___, Name, Value)` specifies additional options with one or more `Name, Value` pair arguments, using any of the previous syntaxes.

Input Arguments

centers - Coordinates of circle centers

two-column matrix

Coordinates of circle centers, specified as a P-by-2 matrix, such as that obtained from `imfindcircles`. The *x*-coordinates of the circle centers are in the first column and the *y*-coordinates are in the second column. The coordinates can be integers (of any numeric type) or floating-point values (of type `double` or `single`).

Data Types

`single` | `double` | `int8` | `int16` | `int32` | `int64` | `uint8` | `uint16` | `uint32` | `uint64`

radii - Circle radii

column vector

Circle radii, specified as a column vector such as that returned by `imfindcircles`. The radius value at `radii(j)` corresponds to the circle with center coordinates `centers(j,:)`. The values of `radii` can be nonnegative integers (of any numeric type) or floating-point values (of type `double` or `single`).

Data Types

`single` | `double` | `int8` | `int16` | `int32` | `int64` | `uint8` | `uint16` | `uint32` | `uint64`

ax - Axes in which to draw circles

`handle`

Axes in which to draw circles, specified as a handle object returned by `gca` or `axes`.

Data Types

`double`

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments, where `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (`' '`). You can specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

Example: `'EdgeColor','b'` specifies blue circle edges, using the short name for blue.

EdgeColor - Color of circle edge

`[R G B]` | short name | long name | `'red'` (default) | `'none'`

Color of circle edges, specified as the comma-separated pair consisting of `'EdgeColor'`, followed by any `ColorSpec` value or `'none'`.

Example: `'EdgeColor','b'` specifies blue circle edges.

LineStyle - Line style of circle edge

`'-'` (default) | `'--'` | `':'`

Line style of circle edge, specified as the comma-separated pair consisting of 'LineStyle' and any line specifier in the table below.

Specifier	Line Style
'_'	Solid line (default)
'--'	Dashed line
':'	Dotted line
'-.'	Dash-dot line
'none'	No line

Example: 'LineStyle', '--' specifies a dashed line at each circle edge.

LineWidth - Width of circle edge

double | 2 (default)

Width of circle edge expressed in points, where each point equals 1/72 of an inch, specified as the comma-separated pair consisting of 'LineWidth' followed by a positive floating-point double value.

Example: 'LineWidth', 4 specifies a 4-point line width.

Output Arguments

h - Handles of the circles

vector


Handles of the circles, returned as a vector of hggroup object handles. These handles are children of the axes object ax.

Examples

Draw Edge Lines Around Both Bright and Dark Circles in an Image

Read the image into the workspace and display it.

```
A = imread('circlesBrightDark.png');  
imshow(A)
```



Define the radius range.

```
Rmin = 30;  
Rmax = 65;
```

Find all the bright circles in the image within the radius range.

```
[centersBright, radiiBright] = imfindcircles(A,[Rmin Rmax], 'ObjectPolarity');
```


Find all the dark circles in the image within the radius range.

```
[centersDark, radiiDark] = imfindcircles(A,[Rmin Rmax], 'ObjectPolarity');
```

Draw blue lines at bright circle edges.


viscircles

```
viscircles(centersBright, radiiBright,'EdgeColor','b');
```



Draw red dashed lines at dark circle edges.

```
viscircles(centersDark, radiiDark,'LineStyle','--');
```



See Also

`imfindcircles` | `imdistrline` | `imtool`

warp

Purpose Display image as 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`.

Tips Texture-mapped surfaces are generally rendered more slowly than images.

See Also `imshow` | `image` | `imagesc` | `surf`

Purpose Watershed transform

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.

Remarks

The watershed transform algorithm used by this function changed in version 5.4 (R2007a) of the Image Processing Toolbox software. The previous algorithm occasionally produced labeled watershed basins that were not contiguous. If you need to obtain the same results as the previous algorithm, use the function `watershed_old`.

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`.

Examples

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 `-Inf`.


```
D = -D;
D(~bw) = -Inf;
```

- 4 Compute the watershed transform and display the resulting label matrix as an RGB images.

```
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;
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)
```

watershed

```
view(3), camlight, lighting gouraud
```

- 3 Complement the distance transform, force nonobject pixels to be $-\text{Inf}$, and then compute the watershed transform.

```
D = -D;  
D(~bw) = -Inf;  
L = watershed(D);  
figure  
isosurface(x,y,z,L==2,0.5)  
isosurface(x,y,z,L==3,0.5)  
axis equal  
title('Segmented objects')  
xlabel x, ylabel y, zlabel z  
xlim(lims), ylim(lims), zlim(lims)  
view(3), camlight, lighting gouraud
```

Algorithms

watershed uses the Fernand Meyer algorithm [1].

References

[1] Meyer, Fernand, "Topographic distance and watershed lines," *Signal Processing*, Vol. 38, July 1994, pp. 113-125.

See Also

[bwlabel](#) | [bwlabeln](#) | [bwdist](#) | [regionprops](#)

Purpose `XYZ` color values of standard illuminants

Syntax
`xyz = whitepoint(string)`
`xyz = whitepoint`

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`.

Examples Return the `XYZ` color space representation of the default white reference illuminant `'icc'`.

```
wp_icc = whitepoint
wp_icc =
    0.9642    1.0000    0.8249
```

whitepoint

See Also

`applycform` | `lab2double` | `lab2uint8` | `lab2uint16` | `makecform` |
`xyz2double` | `xyz2uint16`

Purpose 2-D adaptive noise-removal filtering

Note The syntax `wiener2(I,[m n],[mblock nblock],noise)` has been removed. Use the `wiener2(I,[m n],noise)` syntax instead.

Syntax

```
J = wiener2(I,[m n],noise)
[J,noise] = wiener2(I,[m n])
```

Description `wiener2` lowpass-filters a grayscale image that has been degraded by constant power additive noise. `wiener2` uses a pixelwise adaptive Wiener method based on statistics estimated from a local neighborhood of each pixel.

`J = wiener2(I,[m n],noise)` filters the image `I` using pixelwise adaptive Wiener filtering, using neighborhoods of size `m`-by-`n` to estimate the local image mean and standard deviation. If you omit the `[m n]` argument, `m` and `n` default to 3. The additive noise (Gaussian white noise) power is assumed to be `noise`.

`[J,noise] = wiener2(I,[m n])` also estimates the additive noise power before doing the filtering. `wiener2` returns this estimate in `noise`.

Class Support The input image `I` is a two-dimensional image of class `uint8`, `uint16`, `int16`, `single`, or `double`. The output image `J` is of the same size and class as `I`.

Examples For an example, see “Removing Noise By Adaptive Filtering”.

Algorithms `wiener2` estimates the local mean and variance around each pixel.

$$\mu = \frac{1}{NM} \sum_{n_1, n_2 \in \eta} a(n_1, n_2)$$

and

wiener2

$$\sigma^2 = \frac{1}{NM} \sum_{n_1, n_2 \in \eta} a^2(n_1, n_2) - \mu^2,$$

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.

References

Lim, Jae S., *Two-Dimensional Signal and Image Processing*, Englewood Cliffs, NJ, Prentice Hall, 1990, p. 548, equations 9.44 -- 9.46.

See Also

`filter2` | `medfilt2`

Purpose Convert XYZ color values to double

Syntax `xyxd = xyz2double(XYZ)`

Description `xyxd = 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 software 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`.

Examples Convert `uint16`-encoded XYZ values to double.

```
xyz2double(uint16([100 32768 65535]))
ans =

    0.0031    1.0000    2.0000
```

See Also `applycform` | `lab2double` | `lab2uint16` | `lab2uint8` | `makecform` | `whitepoint` | `xyz2uint16`

xyz2uint16

Purpose Convert *XYZ* color values 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 software 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.

Examples 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 color values to RGB color space
Syntax	<pre>rgbmap = ycbcr2rgb(ycbcrmap) RGB = ycbcr2rgb(YCBCR)</pre>
Description	<p><code>rgbmap = ycbcr2rgb(ycbcrmap)</code> converts the YCbCr values in the colormap <code>ycbcrmap</code> to the RGB color space. If <code>ycbcrmap</code> is M-by-3 and contains the YCbCr luminance (<i>Y</i>) and chrominance (<i>Cb</i> and <i>Cr</i>) color values as columns, <code>rgbmap</code> is returned as an M-by-3 matrix that contains the red, green, and blue values equivalent to those colors.</p> <p><code>RGB = ycbcr2rgb(YCBCR)</code> converts the YCbCr image <code>YCBCR</code> to the equivalent truecolor image <code>RGB</code>.</p>
Class Support	If the input is a YCbCr image, it can be of class <code>uint8</code> , <code>uint16</code> , or <code>double</code> . 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 <code>double</code> .
Examples	<p>Convert image from RGB space to YCbCr space and back.</p> <pre>rgb = imread('board.tif'); ycbcr = rgb2ycbcr(rgb); rgb2 = ycbcr2rgb(ycbcr);</pre>
References	<p>[1] Poynton, C. A. <i>A Technical Introduction to Digital Video</i>, John Wiley & Sons, Inc., 1996, p. 175.</p> <p>[2] Rec. ITU-R BT.601-5, <i>Studio Encoding Parameters of Digital Television for Standard 4:3 and Wide-screen 16:9 Aspect Ratios</i>, (1982-1986-1990-1992-1994-1995), Section 3.5.</p>
See Also	<code>ntsc2rgb</code> <code>rgb2ntsc</code> <code>rgb2ycbcr</code>

ycbcr2rgb

A

- adapthisteq 3-2
- affine transformations
 - definition 3-111
- analyze75info 3-6
- analyze75read 3-7
- analyzing images
 - edge detection 3-165
 - intensity profiles 3-488
 - pixel values 3-457
 - quadtree decomposition 3-723
- applycform 3-9
- applylut 3-11
- averaging filter
 - creating 3-194
- axes2pix 3-14

B

- basicimageinfo 3-304
- bestblk 3-15
- binary images
 - applying a lookup table 3-11
 - constructing lookup tables 3-661
 - labeling connected components 3-54
 - object selection 3-83
 - using neighborhood operations 3-11
- block processing
 - block size 3-15
 - using nlfilter 3-694
- bwarea 3-14 3-24
- bwareaopen 3-26
- bwboundaries 3-29
- bwdist 3-41
- bweuler 3-50
- bwhitmiss 3-52
- bwlabel 3-54
- bwlabeln 3-57
- bwmorph 3-70
- bwpack 3-78

- bwperim 3-81
- bwselect 3-83
- bwtraceboundary 3-85
- bwulterode 3-88
- bwunpack 3-90

C

- Canny edge detector 3-165
- checkerboard 3-91
- chessboard distance transform
 - using bwdist 3-41
- cityblock distance transform
 - using bwdist 3-41
- closing 3-71
- col2im 3-93
- colfilt 3-94
- color spaces
 - converting between 3-698 3-771
 - NTSC 3-698 3-771
- column processing 3-94
 - reshaping columns into blocks 3-93
- conndef 3-97
- connected components
 - using bwlabel 3-54
- convmtx2 3-99
- convolution
 - convolution matrix 3-99
- corr2 3-107
- correlation coefficient 3-107
- cp2tform 3-108
- cpcorr 3-120
- cpselect 3-122
- cpstruct2pairs 3-126

D

- dct2 3-127
- dctmtx 3-130
- decomposition of structuring elements

- getting sequence 3-222
- deconvblind 3-131
- deconvlucy 3-134
- deconvreg 3-137
- deconvwnr 3-139
- decorrstretch 3-144
- demosaic 3-145
- dicomanon 3-149
- dicomdict 3-151
- dicominfo 3-152
- dicomlookup 3-154
- dicomread 3-155
- dicomuid 3-157
- dicomwrite 3-158
- dilation
 - grayscale 3-700
- discrete cosine transform
 - transform matrix 3-130
- disk filter 3-194
- display techniques
 - displaying images at true size 3-816
 - multiple images 3-802

E

- edge 3-164
- edge detection
 - methods 3-165
- edgetaper 3-171
- entropy 3-172
- entropyfilt 3-173
- erosion
 - grayscale 3-700
- Euclidean distance transform
 - using bwdist 3-42

F

- fan2para 3-175
- fanbeam 3-180

- filter design
 - frequency sampling method 3-191
 - frequency transformation method 3-200
 - windowing method 3-207
- filters
 - averaging 3-194
 - disk 3-194
 - Gaussian 3-194
 - Laplacian 3-194
 - Log 3-194
 - median 3-675
 - motion 3-194
 - order-statistic 3-700
 - Prewitt 3-194
 - Sobel 3-194
 - unsharp 3-194
- findbounds 3-185
- fliptform 3-188
- frequency response
 - computing 3-189
- frequency sampling method (filter design)
 - using fsamp2 3-191
- frequency transformation method (filter design)
 - using ftrans2 3-200
- freqz2 3-189
- fsamp2 3-191
- fspecial 3-194
- ftrans2 3-200
- fwind1 3-203
- fwind2 3-207

G

- Gaussian filter 3-194
- Gaussian noise
 - adding 3-448
- getheight 3-211
- getimage 3-212
- getimagemodel 3-215
- getline 3-216

getneighbors 3-217
getnhood 3-218
getpts 3-219
getrangefromclass 3-220
getrect 3-221
getsequence 3-222
gray2ind 3-224
grayscale morphological operations 3-700
grayscale 3-237
graythresh 3-239

H

hdrread 3-240
hdrwrite 3-241
histeq 3-242
histogram equalization 3-242
hough 3-247
houghlines 3-252
houghpeaks 3-255

I

iccfind 3-257
iccread 3-259
iccroot 3-263
iccwrite 3-264
idct2 3-266
ifanbeam 3-268
im2bw 3-273
im2col 3-275
im2double 3-277
im2int16 3-278 3-280
im2java2d 3-279
im2single 3-280
im2uint16 3-281
im2uint8 3-282
imabsdiff 3-283
imadd 3-285
imadjust 3-287

image transformations
 affine 3-111
 local weighted mean 3-112
 nonreflective similarity 3-111
 piecewise linear 3-112
 polynomial 3-111
 projective 3-111
 similarity 3-111
imageinfo 3-295
imagemodel 3-298
images
 creating movies 3-445
 displaying multiple images 3-802
 getting data from axes 3-212
imbothat 3-307
imclearborder 3-310
imclose 3-313
imcolormaptool 3-315
imcomplement 3-317
imcontour 3-319
imcontrast 3-321
imcrop 3-324
imdilate 3-329
imdisplayrange 3-334
imdistline 3-336
imdivide 3-344
imellipse 3-346
imerode 3-351
imextendedmax 3-355
imextendedmin 3-357
imfill 3-359
imfilter 3-363
imfreehand 3-385
imgca 3-394
imgcf 3-396
imgetfile 3-397
imhandles 3-408
imhist 3-409
imhmax 3-426
imhmin 3-428

imimposemin 3-431
imlincomb 3-435
inline 3-438
immagbox 3-443
immovie 3-445
immultiply 3-446
imnoise 3-448
imopen 3-451
imoverview 3-453
imoverviewpanel 3-456
impixel 3-457
impixelinfo 3-460
impixelinfoval 3-463
impixelregion 3-465
impixelregionpanel 3-468
implay 3-469
impoint 3-472
impoly 3-477
impositionrect 3-484
improfile 3-488
imputfile 3-491
impyramid 3-493
imreconstruct 3-506
imrect 3-508
imregionalmax 3-513
imregionalmin 3-516
imresize 3-531
imroi 3-536
imrotate 3-539
imsave 3-541
imscrollpanel 3-543
imshow 3-549
imsubtract 3-561
imtool 3-563
imtophat 3-568
imtransform 3-571
imview 3-581
ind2gray 3-582
ind2rgb 3-584
indexed images
 converting from intensity 3-224
Intel® Integrated Performance Primitives
 Library 3-588
intensity images
 converting from matrices 3-673
 converting from RGB 3-769
 converting to indexed 3-224
intensity profiles 3-488
interfileinfo 3-585
interfileread 3-586
intlut 3-587
ippl 3-588
iptaddcallback 3-590
iptcheckconn 3-591
iptcheckhandle 3-592
iptcheckinput 3-594
iptcheckmap 3-596
iptcheckmargin 3-597
iptcheckstrs 3-598
iptcondir 3-604
iptdemos 3-600
iptgetapi 3-601
iptGetPointerBehavior 3-602
iptgetpref 3-603
iptnum2ordinal 3-605
iptPointerManager 3-606
iptremovecallback 3-613
iptSetPointerBehavior 3-614
iptsetpref 3-618
iptwindowalign 3-619
iradon 3-623
isbw 3-627
isflat 3-628
isgray 3-629
isicc 3-630
isind 3-632
isnitf 3-633
isrgb 3-634

L

- lab2double 3-636
- lab2uint16 3-638
- lab2uint8 3-640
- label2rgb 3-642
- Laplacian filter 3-194
- Laplacian of Gaussian edge detector 3-165
- local weighted mean transformations 3-112
- Log filters 3-194
- lookup tables
 - constructing 3-661

M

- makecform 3-648
- makeConstrainToRectFcn 3-657
- makehdr 3-658
- makelut 3-661
- makeresampler 3-663
- maketform 3-668
- mat2gray 3-673
- matrices
 - converting to intensity images 3-673
- McClellan transform 3-200
- mean2 3-674
- medfilt2 3-675
- median filtering 3-675
- montage 3-678
- morphological operations
 - closing 3-71
 - diagonal fill 3-71
 - grayscale 3-700
 - shrinking objects 3-73
- motion filters 3-194
- movies
 - creating from images 3-445
- multiframe images
 - displaying 3-678
- multilevel thresholding 3-237

N

- neighborhoods
 - binary image operations 3-11
- nitfinfo 3-692
- nitfread 3-693
- nlfilter 3-694
- noise
 - adding to an image 3-448
- nonreflective similarity transformations 3-111
- normxcorr2 3-696
- NTSC color space 3-698 3-771
- ntsc2rgb 3-698

O

- object selection 3-83
- order-statistic filtering 3-700
- ordfilt2 3-700
- otf2psf 3-702

P

- padarray 3-703
- para2fan 3-707
- phantom 3-711
- piecewise linear transformations 3-112
- pixel values
 - using impixel 3-457
- Poisson noise
 - adding 3-448
- poly2mask 3-714
- polynomial transformations 3-111
- Prewitt edge detector 3-165
- Prewitt filters 3-194
- projective transformations 3-111
- psf2otf 3-722

Q

- qtdecomp 3-723

- qtgetblk 3-727
- qtsetblk 3-729
- quadtree decomposition 3-723
 - getting block values 3-727
 - setting block values 3-729
- quasi-euclidean distance transform
 - using bwdist 3-42

R

- radon 3-730
- rangefilt 3-733
- reflect 3-735
- regionprops 3-736
- RGB images
 - converting to intensity 3-769
- rgb2gray 3-769
- rgb2ntsc 3-771
- rgb2ycbcr 3-772
- Roberts edge detector 3-165
- roicolor 3-774
- roifill 3-776
- roifilt2 3-782
- roipoly 3-784

S

- salt and pepper noise
 - adding 3-448
- similarity transformations 3-111
- Sobel edge detector 3-165
- Sobel filters 3-194
- speckle noise
 - adding 3-448
- statistical properties
 - mean 3-674
 - standard deviation 3-791
- std2 3-791
- stdfilt 3-792
- strel 3-793

- stretchlim 3-800
- structuring elements
 - decomposition sequence 3-222
 - determining composition 3-332
- subimage 3-802

T

- tformarray 3-803
- tformfwd 3-808
- tforminv 3-810
- thresholding
 - to create indexed image from intensity image 3-237
- tonemap 3-812
- transforms
 - discrete cosine 3-127
- translate 3-814
- trueSize 3-816

U

- uintlut 3-817
- unsharp filters 3-194

W

- warp 3-824
- watershed 3-825
- whitepoint 3-829
- wiener2 3-831
- windowing method (filter design) 3-207

X

- xyz2double 3-833
- xyz2uint16 3-834

Y

- ycbcr2rgb 3-835

Z

zero-cross edge detector 3-165