File Menu	- control file input and output.
Edit Menu	- allows image editing and simple image arithmetic.
View Menu	- controls visual parameters and set view characteristics.
Processing Menu	- routines for processing of the image.
Options Menu	- controls options for IPS system.
Tools Menu	- selects tools for image processing.
Help Menu	- on-line context sensitive help.

Other menus:
Section Menu
Pixel Menu

Open Image File	- opens an image file.
Save Image File	- saves the current image.
Save as...	- saves the current image under a new name.
Print Image	- outputs the current image to the printer.
Setup Printer	- sets up active printer options.
Exit IPS	- exit back to Windows.

UNDO	- undo last operation on image.
Copy	- copy image or region to clipboard
Paste	- paste image or region from clipboard
Region of Interest	- select interest area.
New Title	- change image title from default.
Copy Image	- copy the current image to another image.
Create Image	- create another image for data storage.
Close Image	- close a previously opened image.
Image Arithmetic	- perform simple arithmetic with images and constants.

Source	- change source image.
Palette	- change the color or gray scale palette.
Zoom	- zoom in on scrolled image.
Status Bar	- to enable or disable display of status bar.

New Palette - create a new table of colors (palette).

Convert to grey - conversion of a colored image into a black & white one by calculating brightness equivalence.

Read - read palette from file.

Save - save palette in current file.

Save As - save palette in a new file.

Brightness equivalence - level of brightness that corresponding to a given color. It is calculated according to the model of human eye susceptibility.Brightness is calculated as follows:

$$I = 0,3*R + 0,6*G + 0,1*B,$$

where R, G, B are the red, green and blue components.

Contrast Menu	- image contrast processing.
Look-up-table Menu	- change the image values selectively.
Filters Menu	- perform filter operations on an image.
Spatial Operators Menu	- processing operations based on local image regions.
Objects Menu	- perform object analysis of the image.
Morphology Menu	- use a window to search for objects in a binary image and adjust pixel values in object.
Transformation Menu	- perform spatial transformations on the image.

Image Negative	- perform negative of image.
Global Equalization	- change the image values to equalize the probability evenly over the full specified range.
Image Thresholding	- an upper and lower threshold are specified, and all values falling between these two thresholds are given the fill value. All image values NOT between the thresholds are given the value zero.
Global Normalization	- change the scaling of the image.

Median Filter	- compute median of the region around each pixel, and replace that as the value of the pixel.
Average Filter	- replace pixel values by the average of the region around the pixel.
Spots Filter	- remove spots from image by replacing all pixels in within the spot by the preceding values.
R-Filter	- a filter based on histogram analysis in a sliding window.

Smooth	- replace pixel values with average of window around pixel.
Laplacian Operator	- perform a Laplacian calculation of image.
Local Equalization	- perform local equalization over the defined window.
Local Normalization	- normalize local region to full range.
Variance	- computes variance of between neighbooring pixels.
Gradient Operator	- shadowing operation.
Convolution Operator	- perform convolution of image by matrix multiplication with a 3x3 kernel.

Connect Objects	- locate objects in image.
Object Statistics	- Computes statistics for objects in image.
Isoline	- find single equal value line on image and measure its length.
Contour Map	- convert image to contour map.

Dilating	- expand objects by window size.
Eroding	- contract objects by window size.
Opening	- dilate image and then erode. This causes near objects to be connected together.
Closing	- erode image and then dilate. This causes near objects to be separated.

Move Image	- Move the image data in any direction.
Rotate	- Rotate the image by defined angle.
Mirror	- Use one of several axes to mirror the data.
Smooth Transformation	- perform a geometric transformation on the image.

X-Axis	- Mirror the image along the X axis (left moves to right).
Y-AXIS	- Mirror the image along the Y axis (bottom moves to top).
Diagonal (Main)	- Mirror the image over the main axis (top right corner moves to bottom left corner).
Diagonal (Secondary)	- Mirror the image over the secondary axis (top left corner moves to bottom right corner).

Exac Transformation	- Transform the image using the points defined.
Set Points	- Set up the transformation points.
Show Set Points	- Show the transformation points already set.
Read	- Read a set of stored transformation points.
Save	- Save the set of defined transformation points.
Save As	- Save the set of defined transformation points under a new name.

Make LUT	- define the look-up-table matrix (LUT), and perform image lookup.
LUT Open	- open a previously defined LUT.
LUT Save	- save LUT.
Save as...	- save LUT under new file name.

Calibrate	- calibrate the resolution for length measurements.
Configuration	- configure the default system startup values.

Histogram Menu - histogram analysis of the image.

Histogram Analysis	- perform histogram analysis.
Histogram Open	- open a previously saved histogram.
Histogram Save	- save histogram.
Save as...	- save current histogram under new name.
Test Images	- create test images for processing evaluation.

IPS Help	- this help file.
About IPS	- vendor and copyright information.

Vertical Wedge	- vertical brightness gradient.
Horizontal Wedge	- horizontal brightness gradient.
Fractal Image	- image of fractals.
Polynomial Image	Polynomial distribution of brightnesses.

Section Data	- See cross section of image.
Set Points	- Set up the transformation points.
Move Image	- Move the image data in any direction.

Open

Opens a previously saved image file from disk. This opens the standard Open dialog box.

Choose the correct file type from the file type list.

Pressing the [format] button allows the format for the selected file type to be specified.

Parameters:

File Name	- is inserted in the corresponding field or the file to be opened is selected from the list of file names. The list gives the file names with extensions, as indicated in the field List of File Types. To look through the file list by extension, type *. followed by extensions (not more than three symbols) and press [open].
List of File Types	- the list of possible extensions: All Files (*.*) format None; BMP Files (*.BMP) format BMP; GIF Files (*.GIF) format GIF; TIFF Files (*.TIF) format TIFF.
Drives	- select a disc containing the file to be opened.
Directories	- select a directory containing the file to be opened.
[Format]	- install read parameters for different formats. For the format None the dialogue Format Details is called. For the formats TIFF, GIF and BMP the dialogue Read Details is called.
[Open]	- read the file.
[Cancel]	- exit without reading.

File Types

There are four currently supported file types for reading or writing:

TIFF - Tagged Information File Format.

GIF - Graphic Information File.

BMP - MS-Windows standard bitmap format.

None - arbitrary file format.

TIFF

To read the file saved in the TIFF format, one has to select extension *.TIF in the dialog Open Image. Then, press [Format] (in the dialogue Open Image), enter the dialog Read Details and install parameters for read.

When the extension *.TIF is selected the dialog Save Image As the file will be saved using a TIFF file format. The parameters for write are specified in the dialog Write Details, which is activated by pressing [Format] (in the dialog Save Image As).

GIF

To read the file saved in the GIF format, one has to select extension *.GIF in the dialog Open Image. Then, press [Format] (in the dialogue Open Image), enter the dialog Read Details and install parameters for read.

When the extension *.GIF is selected the dialog Save Image As the file will be saved using a GIF file format. The parameters for write are specified in the dialog Write Details, which is activated by pressing [Format] (in the dialog Save Image As).

BMP

To read the file saved in the BMP format, one has to select extension *.BMP in the dialog Open Image. Then, press [Format] (in the dialogue Open Image), enter the dialog Read Details and install parameters for read.

When the extension *.BMP is selected the dialog Save Image As the file will be saved using a BMP file format. The parameters for write are specified in the dialog Write Details, which is activated by pressing [Format] (in the dialog Save Image As).

None

The format in which the image is written in an arbitrary manner, line by line. Each pixel (see Image) is represented by one byte. If the image is multichannel, pixels in different channels are written successively:

[channel 1 pixel 1] [channel 1 pixel 2]...

[channel n pixel 1] [channel 1 pixel 2]...

The image can have any fixed line length. At the beginning of the file the header of an arbitrary length can also be specified.

To read the file of type None in the dialog Open Image, it is necessary to specify extension *.*. Then, when pressing [Format] (in the dialogue Open Image), you set the necessary parameters for writing in the Format Details dialog.

To write the image in the format None, it is necessary to specify extension *.*, in the dialog Save Image As. Then press [Format] (in the dialog Save Image As) and set the necessary parameters for writing in the Format Details dialog.

Saves the currently displayed image in a disk file. Uses the last specified file name.

Save As Command

Saves the currently displayed image file under a new file name. This opens the standard Save dialog box.

Choose the correct file type from the file type list.

Pressing the [format] button allows the format for the selected file type to be specified.

Parameters:

File Name - is inserted in the corresponding field or the file to be opened is selected from the list of file names. The list gives the file names with extensions, as indicated in the field List of File Types. To look through the file list by extension, type *. followed by extensions (not more than three symbols) and press [open].

List of File Types - the list of possible extensions: All Files (*.*) format None; BMP Files (*.BMP) format BMP; GIF Files (*.GIF) format GIF; TIFF Files (*.TIF) format TIFF.

Drives - select a disc containing the file to be opened.

Directories - select a directory containing the file to be opened.

[Format] - install read parameters for different formats. For the format None the dialogue Format Details is called. For the formats TIFF, GIF and BMP the dialogue Write Details is called.

[Open] - read the file.

[Cancel] - exit without reading.

Prints the currently displayed image on the active Windows printer.

Sets up the active printer using the standard printer setup dialog box. Use this option to change the printer configuration.

File Formats

Depending on the type of operation (Open or Save) and the type of file specified this command allows access to file type specific formats available:

Arbitrary Format Parameters

Standard Formats for File Open

Standard Formats for File Save

Read/Write Format Command

This command allows custom specification of image file format. The exact, technical image file format must be known before these options are changed. After reading the file, reset the format to the standard IPS for easier data retrieval, or use the same format to reinsert the processed data back into the input file.

The compression coefficient (n) can take on values higher or lower than unity. If n>1, the image from the file is read with a thinout, i.e. each nth image point and each nth line are read. If n<1, the image is stretched during reading. In this case, the image will be 1/n times larger than the origin, i.e. will contain 1/n times more points in the line and/or in the column. On stretching, a linear interpolation from the neighboring points in lines and columns will occur into the points that don't have a value in the initial image. The compression coefficient can be real or integer.

Parameters:

Header size - number of bytes in the header information of the image file.

Line size - number of pixels per line in image file.

Number of Bands - If channel data is interleaved (multiple channel image file) enter the total number of bands (channels).

Band's Number - the number of the band (channel) you wish to retrieve.

X compression - enter a value for data compression along the X axis. Numbers less than one will cause the image to be enlarged.

Y compression - enter a value for data compression along the Y axis. Numbers less than one will cause the image to be enlarged.

X shift in File - enter the number of pixels to shift the image in the X direction. This allows a subset of the image to be read from the image file.

Y shift in File - enter the number of pixels to shift the image in the Y direction.

Press 

to store parameters, or

 to abort.

Read Details

In this dialog parameters for reading images saved in the formats TIFF, GIF, BMP are specified.

Parameters:

Read fullsize - when this is selected, the whole image is read from the file using compression so that it fully fits into the image size.

Xshift - Xcoordinate of the image point in the file, from which read begins. In the mode Read Full Size X is automatically set equal to 0.

Yshift - Ycoordinate of the image point in the file, from which read begins. In the mode Read Full Size, Y is automatically set equal to 0.

Read Component - the name of the component for read if a colored image is saved in the file: Rred, Ggreen, Bblue. When parameters for read from the GIF format are installed, this field is absent in the dialogue.

Press 

to store parameters, or


to abort.

Write Details

In this dialog, parameters for image write are installed in TIFF, GIF, and BMP formats.

Parameters:

Type - One of the standard types for write: RGB, Palette, or Scale is specified.

Compression - if this parameter is selected, the TIFF file is written with compression. For GIF and BMP this parameter is not applied.

Press 

to store parameters, or


to abort.

Exits application without saving.

Reverses the last processing operation, showing the original image. Executing this command causes the menu item to change to REDO, which redoes the processing operation and changes the unprocessed image back to the processed one.

To save intermediate steps in a long chain of operations, use Save and Save Image as.

A dialog box requests the new image title. Enter any character string of up to 18 characters and press 

to accept new title. Changed titles are only valid for current session and are not saved with image.

Copy

Copying a current image to the clipboard to provide exchange with other image processing programs of the IPS series (see Paste).

Paste

Copying an image from the clipboard (provided it has been placed on the clipboard via programs of the IPS series) to the current image. If the image was not placed on the on the clipboard (see Copy), the operation is impossible.

Copies the currently displayed image to another image. A list box is shown containing all available destination images. Choose the desired image and press 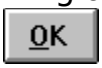

to execute. This will copy the currently displayed image to the chosen target image, and display the copied image.
The copy has the same title as the old image with addition of an asterisk. For example, if the starting image had the title TEST.DAT, the new one will be TEST.DAT*.

Creates a new image for use in image processing. This image is initially blank, but can be filled by a copy operation, a file open command or when used as the result of image arithmetic.

The IPS system initially creates two untitled images for image processing use. If more images are needed for simultaneous processing or comparison, create as many as needed, up to a maximum number of seven available images, but note that creating more images reduces system memory and slows down performance.

Closes the currently displayed image. Frees memory and does not save the image.

Image Arithmetic Command

Use this command to perform simple image arithmetic on any image. The procedure for doing this begins by specifying the first image operand, the image on which the operation is to be performed. Next, specify the second image operand, or choose ***constant*** to specify a constant for the second operand. The constant value is entered in the constant entry box immediately beneath the list boxes. Finally, the result image is chosen, the image where the output from the arithmetic operation will be stored. After having specified all required operand information, press any of the operation buttons to perform that particular operation, or choose 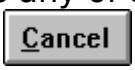

to exit without performing operation.
Arithmetic operations on images. Standard arithmetic operations: addition, subtraction, multiplication and division, and logic operations: OR, AND, XOR, NOT, are performed. The operations can be performed both with two images or an image plus constant. For this purpose, the constant is entered in the Constant box and the name of the second operand Constant is installed.

Remember that the range of brightness variations in the is restricted to 0255. For this reason the brightness of the points for which the operation result is less than zero, will become zero, while the brightness of the points for which the operation result is more than 255, will become 255. Consequently, for optimal operation it is necessary to transform initial images into the brightness ranges so that the final result, if possible, won't be beyond the dynamic range (see Image). This can be achieved using, e.g., a conversion table Lookuptable.

New Palette

Creates a new RGB palette by successively setting transfer functions for each of the three colors: Red, Green and Blue.

By specifying the color in the color box one can edit the transfer function for a chosen color.

In the upper left corner of the dialog there is a window which shows the transfer functions for each color. On entering the dialog, the current state of the RGB palette is drawn. The cursor can be used to change independent points by setting the necessary values. Current cursor positions are reflected in the fields: Value image brightness, Brightness current component value. Selected values can be saved by pressing [Set color] or by double clicking the mouse.

It is possible to construct a new transfer function for any color by resetting all its values. To do this, current transfer function should be set to zero using [Clear] and then new values should be set with the cursor or scroll bars. To construct a continuous transfer function, it is convenient to use [Joint Points], which connects independent points. This function works only if current transfer function has been zeroed before setting the new values.

A transfer function can be set with the help of one of the standard functions (Gaussian Function, Maxwellian Function, Parabola Function). For this purpose one selects the name of the desired standard function in the window Function Type and installs its parameters: X position of the function maximum, D function half width.

The transfer function is set by pressing [Set Function].

The transfer function can be smoothed by pressing [Smooth], or can be treated logarithmically by pressing [Log].

It is possible to set a linear transfer function, which changes linearly from 0 to 255, by pressing [Linear color].

By pressing [Gray Palette] one can set a gray-scale color table.

To have a look at an image in the selected palette without quitting the dialogue, press [Show].

To apply the new palette, press [OK].

To quit the dialogue without changing the palette, press Cancel].

Transfer function is the correspondence between image brightness and the value of an independent palette color. The value of each color can vary within the 0255 range.

Read

Palette read. The file for reading is selected using a standard dialogue Open File.

Save

Save palette with current name.

Save As

Save the palette. Use of a standard dialogue Save As, in which the palette name is selected and saved on disc.

Zoom Command

The mode is intended for a detailed exploded view of the image. To enter this mode, one puts a check-marker on line Zoom in the menu VIEW. The Zoom dialog appears on the screen and a selected image fragment with a given zooming coefficient is displayed. In the main window in this mode a rectangle limiting a viewed area is drawn by solid lines. The rectangle can be moved on the screen with mouse or keyboard (see Functions of Mouse and Keyboard). To move the rectangle with the mouse, click the right key. On releasing the key, a selected image fragment is zoomed. The same procedure can be performed using the keyboard: the rectangle is moved by arrows as the key Shift is pressed. To choose a zooming area, press Shift + Space. Switch between the Main and Zoom windows is performed by pressing Alt + F6 on the keyboard.

In the upper part of this dialog there is a field where a selected image fragment appears. Zooming coefficients can be changed using scroll bars located below and on the right of the image subsection. A lower scroll bar changes the zooming coefficient along axis X (provided the mode nonUniform is set) or simultaneously zooming coefficients along both axes (provided the mode nonUniform is not set). A vertical scroll bar changes the zooming coefficient along axis Y (for the nonUniform mode). As the zooming coefficient changes, the size of the viewed area and, correspondingly, the size of the rectangle that outlines the zoomed area in the main window also change. The current zooming coefficients along the horizontal and vertical lines are specified in the fields Horizontal Rate and Vertical Rate, respectively.

A zoomed image fragment can be normalized differently depending on the mode set in Normalize Interval:

None	- no normalization.
Auto	- normalization is performed so that a zero brightness corresponds to a minimum brightness of the chosen subsection and a 255 brightness to a maximum brightness of the subsection.
Manual	- a maximal and minimum brightness can be set using the corresponding scroll bars located in the lower part of the dialog. If a maximum is set to less than the minimum, a negative image is produced.

In the lower part of the dialog there is a box for Brightness specification, in the Min and Max fields of which there appear a minimum and a maximum brightness of a zoomed fraction.

To switch off the mode, press [Close] or the key Esc. Earlier modes and parameters set are preserved.

Scroll between the images using this option. A list box is shown containing all available images. Choose the desired image and press 

to accept new image and continue. The IPS window caption displays the source image title.

Converts the image to a negative of the original, making dark pixels light and light pixels dark.

Global Equalization

This contrast operation adjusts pixel values so that the full dynamic range of pixel values for the image is used. The image is passed through a Look-Up-Table constructed on the basis of a normalized integral of the histogram. The operation works well when the image has a single-modal histogram. The operation does not produce new brightness levels, but distributes the current brightness levels equally across the whole dynamic range. This can lead to merging of different brightness levels for the output image, thereby decreasing the number of levels after this operation.

The equalization is done on the histogram of the whole image. See Local Equalization for slide window equalization (equalizing to each local area's dynamic range).

Parameters:

Maximum brightness - set the maximum brightness of the output image. The output image brightness will be adjusted evenly to the new minimum-to maximum range.

Minimum brightness - set the minimum brightness for the output image.

Recalc. Histogram - check this box to recalculate the histogram used for equalizing, or uncheck it to use the currently defined histogram. This allows standard equalization across several images, or selection of the histogram of a region of interest for equalizing the whole image.

Press 

to execute equalize, or


to abort.

Image Thresholding

This converts a range of pixel values to a single output value, thereby producing a binarized image. All pixels whose brightness is outside the range specified are set to zero, and the pixels inside the range are given the specified value. Setting the lower level greater than the upper level causes the range in between the lower and upper thresholds to be set to zero, and all other values to the fill value.

This operation is usually used to create various masks. Note that if one needs to cut out a brightness range without changing the brightness values themselves, the function [Cut] in the menu Histogram Analysis is to be used.

Parameters:

Lower level - minimum of threshold range.

Upper level - maximum of threshold range.

Fill value - value to be filled in the threshold range.

Press 

to execute threshold calculation, or

to abort.

Global Normalization

Normalizing converts an old range of pixel values into a new range of pixel values, either linearly or logarithmically. It is also possible to make the dynamic range broader or narrower, decrease the number of brightness values and restrict them to given limits.

These manipulations should be carried out before operations (e.g., Image Arithmetic) which can change the brightness values beyond the dynamic range.

Parameters:

Old min. brightness - the minimum of the old range of pixel values.

New min. brightness - the minimum of the new range of pixel values.

Old max. brightness - the maximum of the old range of pixel values.

New max. brightness - the maximum of the new range of pixel values.

Recalculate Max/Min - check this box to recalculate the old minimum and maximum of the image values for use in the conversion to the new range, or uncheck it to allow use of a the entered minimum and maximum values (see above).

Lin. - check this box to make the conversion of old range to new range linearly.

Log. - check this box to make the conversion of old range to new range logarithmically.

Press 

to execute, or


to abort.

Median Filter

The median filter reduces image noise by computing replacing the pixel value with the median of the surrounding region. This reduces the noise of an image, with less blurring than is caused by the Average filter.

Parameters:

Space Size of Filter - specifies how many pixels in each direction to use as the region for finding the median.

Shift from Median - causes the pixel value to be shifted from the median.

Press 

to execute, or

to abort.

Global Average

Replaces pixel values by the average of the surrounding region. This reduces noise by producing a smoothing effect.

The window sizes in terms of coordinates X and Y can be different and are set as operation parameters. This filter is a smoothing one (see also Median Filter) unless the condition Exclude Center Point has been selected.

If the latter condition is selected, the filter is differential. In this case a difference between a weighting value of brightness in the point of starting image and an average value in the window is placed in every image point. Thus, the filter can be applied to stress changes in brightness.

Parameters:

X Size of Filter	- the size of the averaging region in the X direction.
Y Size of Filter	- the size of the averaging region in the Y direction.
Shift from Average	- enter a value to add to the calculated average.
Weight of Center	- enter the weight to be given the center pixel (old value) in computing the average. A weighted average causes the central pixel to be multiplied the specified number of times when computing the average.
Real Brightness	- check this box if you desire the resulting average to be assigned the value of the nearest real pixel value of the pixels in the averaging region.
Exclude from Average	- check this box if you wish the center pixel excluded from the average calculation. This causes the old pixel value to subtracted from the average. The resulting image is the absolute value of these calculations.

Press 

to execute, or


to abort.

Spots Filter

This a digital noise filter designed to eliminate digital noise from the image.

The spot filter is used to change the brightness value in overshoot points, i.e. where it exceeds by a given threshold the brightness value of neighboring points. Different space sizes of can be set. The brightness at point of overshoot is replaced by the brightness of the nearest nonfailure pixel. This filter is usually used when an image has spacelocalized noises.

Parameters:

Space Size of Filter - sets the minimum pixel size of the digital noise to be removed.

Break Value - sets the minimum difference in values between the digital noise and the image.

Filter Direction - sets the direction the filter will search for spots (X, Y, or BOTH).

Press 

to execute**,** or


to abort.

R-Filter Command

This filter is based on the histogram calculated in the window around a filtered point. For each point, a new value of brightness is selected depending on where the largest number of histogram points are concentrated. In one case the value of brightness is replaced by the minimum value found in the histogram, in another case by the maximum value. If the parameter Use Binary mode is selected, point brightness is replaced, respectively, by zero or a value given in Object Brightness rather than by a minimum or a maximum value.

A threshold value is selected depending on the mode used. In the mode Internal it equals an average (between maximum and minimum) brightness value in the histogram. In the mode External it equals the value set in External Threshold Level, in the mode General a maximum value is selected from External and Internal for each point. The R-Filter is convenient to use when one needs to outline homogeneous areas and sharpen the boundaries between them.

Parameters:

Space Size of Filter - space size of filter in X and Y directions.

Type Filter - select filter type:
Internal
External
General

External Threshold Level - a threshold value used for External Type is selected.

Use Binary Mode - is selected to make the output image binary.

Object Brightness - brightness of the objects after the image has been made binary.

Press 

to execute, or


to abort.

Spatial Operators Options

Use this command to set the window size for spatial operations requiring a slide window.

Operations in the sliding window: the image is successively looked through and brightness of each point changes depending on the point environment (a window around it). By varying the window size, one can stress or smooth the image details with different characteristic sizes. Window Size d is determined by a window half width. The full size of the window around the point equals 2d + 1.

Smooth Command

Replaces the pixel value with the average of the slide window defined around the point. This operation is similar to that of Average Filter; however, it permits averaging not only for the full image but also for any fragment, provided the mode Mask Used is selected. Operation will occur for only those image points where the mask value is other than zero.

Parameters:

X Size of Window - size of slide window on either side of central pixel in the X direction.

Y Size of Window - size of slide window on either side of central pixel in the Y direction.

Mask Used - is selected if one has to carry out processing only for points marked by mask.

Mask Image - selects the title of the mask image

Press 

to accept entered values, or


to abort.

Laplacian Operator

The laplacian spatial operation allows enhancement of features that are close to the size of the region defined by the parameters.

Laplacian is the difference between an average value calculated in the window (2X+1)*(6Y+1), where X and Y are given in the dialogue box parameters.

This operation is differential and is usually applied to emphasize the areas with sharp drops of brightness and to outline the boundaries (see Equalization and Normalization).

Parameters:

X Size of Window - size of slide window on either side of central pixel in the X direction.

Y Size of Window - size of slide window on either side of central pixel in the Y direction.

Press 

to accept entered values, or

to abort.

Local Equalization

Performs histogram equalization over the defined window. See also Global Equalization.

The features of this operation are similar to those of Local Normalization. Application of the operation is determined by the properties of a particular image.

Parameters:

X Size of Window - size of slide window on either side of central pixel in the X direction.

Y Size of Window - size of slide window on either side of central pixel in the Y direction.

Press 

to accept entered values, or


to abort.

Local Normalization

Normalizes each pixel to full scale brightness relative to the surrounding region. See also Global Normalization.

The objective of this operation is to replace the brightness value for each image point by the value point vicinitynormalized value. For small window sizes this operation emphasizes different boundaries. With increasing window size, the details of the image are emphasized.

For Local Normalization, minimum and maximum brightness values are calculated in the vicinity of the point. These values are then used for normalization (see Global Normalization) within the range 0255. After calculation of normalization coefficients the brightness values in the central point are transformed.

Parameters:

X Size of Window - size of slide window on either side of central pixel in the X direction.

Y Size of Window - size of slide window on either side of central pixel in the Y direction.

Press 

to accept entered values, or


to abort.

Variance Command

As a result of this operation, the brightness value in each image point is replaced by brightness disparity value in the vicinity of the pixel. It is used to stress the areas with sharp drops in brightness (see Equalization and Normalization), and also to distinguish homogenous image fragments with the same statistic parameters.

Parameters:

X Size of Window - size of slide window on either side of central pixel in the X direction.

Y Size of Window - size of slide window on either side of central pixel in the Y direction.

Press 

to accept entered values, or


to abort.

Gradient Operator

This operation produces a shadowing effect on the image. The shadowing is produced by computing the gradient of the pixel to the window in the specified direction.

The operator calculates the brightness gradient in the sliding window in a given direction. The image becomes pseudothreedimensional: "shadows" appear as if the image was lighted from a given direction. It is thus possible to emphasize particular features.

The direction of illumination is determined by the horizontaltovertical window size ratio and by the quadrant chosen.

Gradient brightness is equal to the difference between brightness in the processed pixel and the average value in the window, for which this pixel is one of the corners. The corner chosen is given by the Quadrant Number, in which the window is located relative to the processed point.

Parameters:

X Size of Window - size of window (shadow) in the X direction.

Y Size of Window - size of window (shadow) in the Y direction.

Quadrant Number - mathematical quadrant of the gradient window (i.e. direction of shadowing).

Press 

to execute, or

to abort

Convolution Operator

This is a general purpose, custom filter operator. It can be used for high or low pass filtering, edge detection in any direction, custom laplacian operations or other enhancement methods. The filter is defined by a 3x3 kernel whose values define the behavior of the filter.

The brightness value at each image point is replaced by a new value equal to brightness sum in neighboring points multiplied by the corresponding kernel values and divided by the total weight.

There are two ways to specify a kernel for convolution use. One is by direct specification using the weight and kernel boxes. The other is by reading a previously stored kernel from a disk file.

Parameters:

Weight - Set the real number by which to divide the kernel output value. This value should be set to the sum of the kernel values in order to reduce the summation effect of the filter.

Kernel - Enter the kernel values in the appropriate boxes.

Kernel File Name - sets the file name for *Read Kernel* and *Write Kernel* operations.

Read Kernel - Use the directory and filename boxes to obtain a saved kernel name, and press this button to read that kernel from the disk file. The saved kernel will be displayed.

Write Kernel - If you wish to use this kernel in a later session, you can save it to disk by using the ***Kernel File Name*** and directory choices to obtain the path and filename, and pressing ***Write Kernel*** to save it to a disk file.

Press 

to execute the convolution operation, or


to abort.

Kernels for noise suppression:

NOIS1.KRN:

1 1 1

1 1 1

1 1 1

with weight 9.

NOIS2.KRN:

1 1 1

1 2 1

1 1 1

with weight 10.

NOIS3.KRN:

1 2 1

2 4 2

1 2 1

with weight 16.

Kernel for restoring image focus:

FOCUS.KRN:

-1 -2 -1

-2 19 -2

-1 -2 -1

with weight 7.

Kernels for outlining and distinguishing boundaries:

BOUND1.KRN:

0 -1 0

-1 5 -1

0 -1 0

with weight 1.

BOUND2.KRN:

-1 -1 -1

-1 9 -1

-1 -1 -1

with weight 1.

BOUND3.KRN:

-1 -2 -1

-2 13 -2

-1 -2 -1

with weight 1.

Kernels providing twodimensional differentiation are used to distinguish drops in brightness in different directions:

NORTH.KRN

1 1 1

1 -2 1

-1 -1 -1

with weight 1.

NORTHW.KRN

1 1 1

-1 -2 1

-1 -1 1

with weight 1.

WEST.KRN:

-1 1 1

-1 -2 1

-1 1 1

with weight 1.

SOUTHW.KRN:

1 1 1

1 2 1

1 1 1

with weight 1.

SOUTH.KRN:

1 1 1

1 2 1

1 1 1

with weight 1.

SOUTHE.KRN:

1 1 1

1 2 1

1 1 1

with weight 1.

EAST.KRN:

1 1 1

1 2 1

1 1 1

with weight 1.

NORTHE.KRN:

1 1 1

1 2 1

1 1 1

with weight 1.

Kernels for calculating the Laplacian operator are used to distinguish drops in brightness without taking into account orientation:

LAP1.KRN:

0 1 0

1 4 1

0 1 0

with weight 1.

LAP2.KRN:

1 1 1

1 8 2

1 1 1

with weight 1.

LAP3.KRN:

1 2 1

2 4 2

1 2 1

with weight 1.

Kernels used to distinguish lines following different directions:

LIN1.KRN:

1 2 1

1 2 1

1 2 1

with weight 1.

LIN2.KRN:

1 1 1

2 2 2

1 1 1

with weight 1.

LIN3.KRN:

1 1 2

1 2 1

2 1 1

with weight 1.

LIN4.KRN:

2 1 1

1 2 1

1 1 2

with weight 1.

Connect Objects Command

This operation finds connected objects based on a range of pixel values and a minimum area criterion. The resulting image displays the objects found where each object's pixels have a unique value depending on the objects size. The largest object's pixels will have the maximum value (255) and the smallest the minimum (1). The non objects in the resulting image will have a value of 0. Therefore, this operation can find a maximum of 254 discrete objects.

This operation can be used as a preliminary if one has to analyze the shape statistics of various objects or shapes of particular objects. The operation is also useful when independent objects and their boundaries with a high level of noise have to be outlined.

Parameters:

Lower Level - enter the lower level of the pixel value range defining the objects.

Upper Level - enter the upper level of the pixel value range defining the objects.

Minimum Area - enter the minimum area of the objects.

Press 

to execute, or


to abort.

Object Statistics

Computes standard object statistics, for each object number found using Connect Objects.

The operation is used to analyze the shape of different objects. An object is a set of points of the same brightness, therefore, prior to the operation one has to pre-process an image by "coloring" the necessary objects (see Connected Objects). In the upper left corner of the dialogue is a field, where the number of pixel of different objects is plotted as a function of their brightness, i.e. actually a histogram of the processed image.

A vertical marker can be moved all over the histogram with the aid of the lower scroll bar. In the corresponding fields on the right the parameters of the object where the marker is currently set are displayed.

A vertical scroll bar helps vary the scale of histogram representation in order to select the scale which is most convenient for analysis.

Note that if prior to statistics calculation, the image was not processed in accordance with the operation Connected Objects, on pressing [Calculate] you get a corresponding warning. If the object shape statistics have been calculated, shape parameters of particular objects can be obtained without entering the dialogue. For this purpose bring the cursor to any point of a desired object and press the left key of the mouse. The Object Shape can be selected from the menu that will appear.

Parameters:

Brightness - brightness of object points.

Area - number of pixels in object.

Mass Center Coordinates - (Xc,YC) mass center coordinates for the object. For this calculation all points are assumed to have the same weight.

Nearest Circle Parameters - parameters that specify the object difference from circle.

Radius - radius of circle that has the same area as the object.

Difference - ratio of the number of points which are not included in the circle (of the above radius centered at the object mass center) to the total number of object points.

Shape Parameters - object shape parameters.

Angle - object orientation angle.

Perimeter - perimeter of object shape in pixel units.

Formfactor - the maximum to minimum object cross-section ratio

Press **{bmc calc.bmp)** to perform computation of statistics, and then use the horizontal scroll bar to move the histogram cursor to the desired object number. The statistics will appear in the boxes to the left.

To view previously computed statistics, just press **{bmc show.bmp)**.

Object Shapes

The dialogue Object Shapes shows shape parameters of the object indicated by the cursor. A menu item for Object Shape parameters appears on clicking the mouse once only if the object shape statistics have already been calculated for a given image.

Isoline Command

This operation will find an equal valued line of pixels in the image. Reentering the dialog box after the operation displays the length of the resulting line (in calibrated units).

Parameters:

Isoline Level - enter the pixel value of the line to find.

Brightness Level - enter the brightness to assign the pixels found in the line.

Isoline Length - re-enter the Isoline routine to see the line length.

Press 

to execute, or

to abort.

Contour Map Command

This operation generates a contour map of the image.

A number of isolines are performed on the image, whose brightness level is within a given range. The isolines are constructed with a uniform step in brightness.

Parameters:

Lower Level - the lower level of pixel value range for contour map.

Upper Level - the upper level of pixel value range for contour map.

Number of Lines - defines the number of contour lines to produce in the pixel value range.

Press 

to execute, or

to abort.

Morphology Operations

Morphology operations are used to change the shape of objects in the image by either increasing (dilating) or decreasing (eroding) the object shapes. There are also combination processes for combining or differentiating between objects (closing and opening).

Morphologic operations are used to transform brightness on the image depending upon brightness of neighboring points. The neighboring points are defined by an ellipse inscribed in the window with a given size around the processed point. Morphologic operations, like the RFilter, are usually used to distinguish homogeneous objects of given sizes, smooth boundaries and eliminate various noise - like object shape distortions. These operations are applied mainly for analysis of binarized objects.

These routines define the window size parameters that determine what size and shape objects will be increased to or reduced by.

Dilating Command

Dilating an object increases the size and rounds the edges (makes the edges convex) of the object. The amount each object is increased and rounded is determined by the Window Size parameters.

Morphologic vicinity is an ellipse inscribed into a rectangular window around the processed point.

Parameters:

X Size of Window - defines the X axis window size.

Y Size of Window - defines the Y axis window size.

Press 

to accept parameters, or


to abort.

Eroding Command

Eroding an object decreases the size and makes the edges of the object more pointed (makes the edges concave). The amount each object is decreased and made concave is determined by the Window Size parameters:

Morphologic vicinity is an ellipse inscribed into a rectangular window around the processed point.

Parameters:

X Size of Window - defines the X axis window size.

Y Size of Window - defines the Y axis window size.

Press 

to accept parameters, or


to abort.

Opening Command

This causes near objects to be connected together. This command is a combination of first dilating and then eroding the image. Enter the Window Size parameters to determine the basic object shapes to be connected:

Morphologic vicinity is an ellipse inscribed into a rectangular window around the processed point.

A morphologic operation is called stable when its repeated application does not affect the image.

Parameters:

X Size of Window - defines the X axis window size.

Y Size of Window - defines the Y axis window size.

Press 

to accept parameters, or


to abort.

Closing Command

This causes near objects to be separated. This command is a combination of first eroding and then dilating the image. Enter the Window Size parameters to determine the basic object shapes to be separated.

Morphologic vicinity is an ellipse inscribed into a rectangular window around the processed point.

A morphologic operation is called stable when its repeated application does not affect the image.

Parameters:

X Size of Window - defines the X axis window size.

Y Size of Window - defines the Y axis window size.

Press 

to accept parameters, or


to abort.

Rotate Image

Rotation of an image through an arbitrary angle with reference to an arbitrary image point. One can use for this transformation different image approximation degrees to the points of a new image, into which the initial image points failed to get during the rotation.

Coordinates on the image are calculated from the upper left corner of the image. The coordinate X increases from left to right, the coordinate Y from top downwards. When passing to a neighboring point, the coordinate values increase by unity, i.e. the coordinates change with a step 1.

Parameters:

Pivot Point - point around which to rotate the image (i.e. center of rotation).

Angle - enter the angle to rotate the image (positive for clockwise and negative for counter-clockwise).

# of Smooth Points - Number of points to smooth in determining the new rotated positions of the pixels. A smaller number means a finer output image.
1 - only the nearest point is used for calculating rotation.
2 - two nearest points are used;
4 - four nearest points are used.

Press 

to execute, or


to abort.

X-axis Mirror

Mirrors the image along the X axis. This swaps between the left and right side pixels for each line. Image reflection with respect to a horizontal axis passing through its center.

Y-axis Mirror

Mirrors the image along the Y axis. This swaps between the bottom and top pixels for each line. Image reflection with respect to a vertical axis passing through its center.

Diagonal (Main) Mirror

Mirrors the image along the main diagonal axis. This swaps between the pixels in the top right corner with those in the bottom left corner. The main diagonal remains the same.

Diagonal (Secondary) Mirror

Mirrors the image along the secondary diagonal axis. This swaps between the pixels in the top left corner with those in the bottom right corner. The secondary diagonal remains the same.

Move Image

Shifts the image the specified number of pixels in any direction.

This produces a parallel shift of all image points in a given direction. The shift direction can be indicated on setting shifts along axes X and Y in the corresponding fields of the dialogue. The dialogue can be called either from the main menu or from the menu that appears near the cursor on drawing the line on the image (see Functions of Mouse and Keyboard). In this case the drawn line is assumed to be a direction and in the fields Shift on X, Shift on Y the corresponding values are set, which can be corrected whenever necessary.

Parameters:

Shift on X - defines the X axis shift.

Shift on Y - defines the Y axis shift.

Press 

to accept parameters, or

to abort.

Exec Transformation

Performs smooth geometric transformations according to given shift points. During the process, the given points are shifted into new positions, while their vicinities are smoothly transformed, i.e. any straight lines on the image are transformed without breaks and inflections. Prior to the operation, it is necessary to give starting and end coordinates of all shift points. This can be done by two methods: using the Set Points menu item or the cursor menu. When the line is drawn by the cursor (see Functions of Mouse and Keyboard), releasing the left mouse key leads to appearance of the section menu from which Set Points can be selected.

A rectangular grid is drawn on the full image, and smooth transformation functions are constructed in the grid cells. The given points are part of the grid knots. Shifts of other knots are calculated automatically depending on input parameters. The number of lines in the grid can be increased to the Number of Grid Points. In this case the grid will be uniform. If the grid knots are closely spaced, a use of the parameter Operating Radius determines the area around the shift points that will be transformed. If these parameters are correctly selected, one can transform only vicinities of shift points of a given radius rather than the full image. Don't forget to set the radius to greater than the maximum shift value.

Parameters:

Number of Grid Points - number of lines in the grid.

Operating Radius - selection of the influence radius of shift points.

Press 

to accept parameters, or

to abort.

Set Points

Sets shift points for geometric transformations. In the center of the dialogue there is a window which has all set points and their shifts. When entering the dialog from the section menu (see Functions of Mouse and Keyboard), parameters of a new point (number, start and end coordinates) are displayed in the corresponding windows. The parameters can be corrected or a point can be added to the list immediately by pressing [Add]. On entering the dialog from Smooth Transformation Menu, only a current free number is filled. Starting and end coordinates are introduced into the corresponding fields.

Parameters:

Number	- ordinal number of shift.
Starting Point	- shift starting coordinates.
Ending Point	- shift ending coordinates.
[Add]	- add a new point to the list.
[Delete]	- delete a point marked by the left mouse key from the list.
[Sort]	- sort points according to increasing ordinal number and eliminate jumps between the numbers.
[Clear]	- remove all points from the list.
[Exit]	- quit the dialog.

Show Set Points

All introduced shift points are displayed by arrows showing the shift direction and number.

Read (Point Coordinate)

Selects a file to read selected shift points using a standard Open File dialog.

Save As (Point Coordinate)

Save point coordinates. A standard Save As dialog is used to select the file name and save the shift points on the disk.

Save (Point Coordinate)

Save point coordinates using the current file name.

Calibration

Sets the resolution of the pixel size, and the measurement units.

The units can be set directly by giving the size of a pixel (image point) and measurement unit in the fields Value and Units. It is also possible to set the mode to Section and give starting and end coordinates of any segment on the image whose length is known. Then one has to indicate this length in the field Value. Press [OK] and the size of an independent pixel will be calculated proceeding from the above data.

For a convenient setting of segment starting and end coordinates, one can draw the necessary line on the image (see Functions of Mouse and Keyboard) and then enter this menu item. This can be done using the menu that appears on manipulating the cursor, whereupon the fields with the segment starting and end coordinates are filled automatically.

Parameters:

Value	- enter the distance between pixels (resolution) or the length of the section defined by the endpoints (if ***to Section*** box is checked).
Units	- enter the calibration units (i.e. km, mm, µm, etc.).
to Section	- check this box to calibrate the value entered to the last defined section, or to a section entered using the ***Endpoint*** parameters.
Endpoint X	- enter the X coordinates for the endpoints.
Endpoint Y	- enter the Y coordinates for the endpoints.

Press 

to execute, or


to abort.

Configuration

Sets the default parameters for system startup, and the directories for data files.

Load

Read all dialog box parameters. The read file is selected using a standard Open File dialog. The starting state of the program is kept in the IPS.INI file in the operating directory.

Save

Saves all parameters using a specified file name. If a set of parameters is saved in the IPS.INI file in the operating directory, on a subsequent initiation of the IPS program the saved state will be loaded automatically.

Histogram Analysis

This routine opens a histogram analysis dialog, with boxes to control and execute an analysis. The top left corner contains the histogram window, with a scaling bar to the left and two scroll bars underneath. The scroll bar marked ***L*** controls the movement of the left limit bar, and the bar marked ***R*** controls the right one. The left and right limit bars can also be controlled from the ***Left*** and ***Right Bar*** boxes in the bottom left corner by entering the respective positions. Changing the limit bars and recalculating forces all pixels outside the limit bars to be given the corresponding limit bar position value.

In the top right corner there is a histogram statistics box that shows the vital statistics for the histogram. The ***mean*** is the mean of all pixel values in the image, the ***variance*** is the average deviation of the pixel values from the mean (dispersity), the ***skewness*** is a measure of the asymmetry of the histogram, and the ***Kurtosis*** is a measure of its flatness (excess).

Below that is the accumulator box that allows the storing of histogram data for cumulative analysis of several histograms. The buttons allows the display of a new histogram from information stored in accumulator 

, adding to the accumulator


, subtracting from the accumulator


, and clearing the accumulator


. The accumulator functions can be used to combine histogram data from several images to allow group analysis.

In the bottom right corner is the Average window, which is used to perform an average of the image before the histogram is analyzed. The ***window size*** is the X and Y size of the averaging window. Pressing the 

button performs the average.

In the bottom center is the ***partial sum*** of the image which is the number of pixels between the horizontal scroll bars. When the parameters are ready press 

to compute the histogram.


allows the display of accumulator data, and
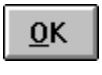

leaves the window.

When quitting the dialogue with the help of [Cut], all brightness values on the image lying within the range restricted by the scroll bars, will remain unchanged, while all the other values will be set to zero.

Opens a previously saved histogram file.

Saves the currently displayed histogram in a disk file. Uses the last specified file name.

Saves the currently displayed histogram file under a new file name.

Image with a vertical brightness wedge.

Image with a horizontal brightness wedge.

Fractal image.

Polynomial

In this dialogue, in the corresponding fields the coefficients of the following second order polynomial are given:

b(x,y)=aO + ax*x + ay*y + ax2*x**2 + ay2*y**2 + axy*x*y,

where b(x,y) is test image brightness in the point with coordinates x,y, aO, ax, ay, ax2, ay2; axy are coefficients. With the Auto normalization mode the test image is automatically normalized within the range 0255.

Look-up-table Creation

Defining and applying an arbitrary brightness transfer table (LUT) on the basis of an old or a new table.

The left side of the lookup dialog contains a window to view the lookup table, scroll bars to define lookup table points, a Clear button to clear defined points or current lookup table, and a Join Points button to connect defined points to interpolate the lookup table. If Join Points is shaded, press Clear to clear the lookup table and start specifying a new one.

On the right side is the Values box , which allows the defining of points by value and includes a Set button to enter a point (whether defined by scroll bars or values). Below that is the Operations window, which can be used to perform simple arithmetic operations on the LUT. After entering the function and constant (if needed) press Exec to perform the operation. When the LUT is defined to your approval press OK to convert image to new values based on the lookup table.

Opens a previously saved LUT file. To apply this table enter Make and press 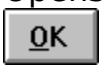

to convert image to new values based on the lookup table.

Saves the currently displayed LUT in a disk file. Uses the last specified file name.

Saves the currently displayed LUT file under a new file name.

Section Command

This dialog allows the user to analyze brightness distribution over the straight line drawn by the cursor. The main role in the dialog is played by the window which shows the section brightness graph.

If you press the left key or Space and update the cursor, a line connecting the starting point with the current cursor position will be drawn on the image. By releasing the mouse key or pressing Space one can access the menu Section.

After drawing a line, a menu item appears that enables the Section command. Choosing this item causes a cross section of the line to appear in the section dialog box. Values can be sampled along this line. Changing the section endpoints and pressing

allows fine tuning of the cross section. The cross section statistics are show in the dialog box.

Under the window there is a scroll bar that controls a marker, which can move along the graph. The marker can be adjusted at any section point. In this case, in the block Bar Position the coordinates of the current point and its brightness (X, Y, Value) are set. Block section statistics show a mean value (mean) and dispersity (variance) of brightness over the whole section, as well as the section length (length). In blocks Start Position and End Position the coordinates of the starting and end points of the section are set.

Press  for further explanation.

Point Menu

The cursor is manipulated as follows. It is moved with the help of the mouse or arrows on the keyboard. When selecting the necessary point, press the left mouse key and , without moving the cursor, release it (or press twice Space on the keyboard). Then the menu Set New Brightness, from which you can choose the dialog Set New Brightness, will appear near the cursor.

If the image, for which object shape parameters have been calculated, is current, the line Object Shape will appear in the menu. From Object Shape you can learn the shape parameters of the object whose point was marked by the cursor.

Set New Brightness

Within this menu, a new brightness value for independent image points can be set.

Parameters:

Brightness in this field the required brightness value for the selected point is introduced.

[OK]	operation performance
[Cancel]	quit without calculations

Warnings and Errors:

Errors in Image Allocation

User Input Errors

Errors in File Reading

Errors in Image Allocation

Warning "Unable to allocate <name> image!" may appear when you enter the program or attempt to open an additional image.

In the name field, the image title to be put down cannot be placed. The same error may arise when some other operations that require placement of additional buffer images are performed. The error is due to a lack of free memory in Windows. To continue work, strip off part of applications currently operating in windows and restart the program.

User Input Errors

When the dialogue fields are filled with wrong values (e.g., not with numbers but with symbols or numbers that are beyond the range of possible variations of installed parameters), on pressing [OK] to quit the dialogue a warning appears. The warning indicates the name of the parameter that has been erroneously installed, as well as the cause of the error and the limit of the range exceeded at parameter installation. The program remains in the current dialogue. To quit the dialogue, either install a correct value of the parameter or use the button [Cancel].

Errors in File Reading

Errors can be due either to wrong installation of parameters or to lesser than anticipated information in the file, or this information cannot be read.

In this case warning "Unable to read file" appears. Only the part of data read before the error has occurred is inserted in the current image.

The system menu appears on pressing the button in the left corner of the main window header. This menu is standard for Windows applications. Its description can be easily found in the operation manual or in Help for Windows.

Status Bar

This menu selection toggles the status bar, which appears below the image. The status bar displays the current pixel coordinates of the cursor and the current region of interest.

IPS Tutorial, by Bryon Gomberg

Getting Started

The basic concepts that the user needs to understand before entering the world of image processing are:

What is an Image?

Where do Images come from?

Why do you need Image Processing?

Starting an IPS session

Image Restoration

Image Enhancement

Advanced Processing

Further Reading

Books on Image Processing

Introduction

Image processing is basically divided into two major fields: Image Restoration and Image Enhancement. These types of processing are distinguished by the different objectives of each. Image Restoration involves making the image as much like the image originally observed as possible, and Image enhancement involves changing the image to make the features we wish to see more pronounced.

This tutorial is designed for the inexperienced user, and explains the basic terms associated with image processing, image processing does, and lead the user through several examples of image restoration and enhancement so that they may acquire the knowledge needed to successfully process images with the IPS software.

What is an Image?

An image is basically a picture that represents some objective. The image is a collection of pixels (short for picture elements), arranged in rows and columns. Each pixel has a certain value that is used to represent what the objective "looks like" at that particular point.

The image type is defined by the type of value the pixels have, and can one of the following: 1 bit, 2 bit, 4 bit, 8 bit, or true-color (24 bits). Image types of 1, 2, and 4 bits are not usually used in image processing, so they will not be considered here.

8 bit Images

An 8 bit image means that each pixel can have a value from 0 to 255, and can be either a gray scale image or an indexed image. A gray scale image uses each pixel value as the brightness of the image at that point, and an indexed image uses the value as a representation of a color (also called pseudo or false color images). The palette is a table of colors that shows what color each pixel value should be.

True Color Images

A true color image is an image that has 24 bits of data for each pixel, and the 24 bits are divided into three 8 bit values: one each for the red, green and blue brightness (referred to as RGB values). When this image is displayed, each pixel has a unique color depending on the RGB values.

Where do Images come from?

Images start off as some form of detectable output from the objective, and the process of obtaining an image starts with an input device that measures this output. The image data is then transferred to the computer through some form of communication link and stored on the computer (in an Image File), or sent to an application for initial processing.

The following picture visualizes the process:

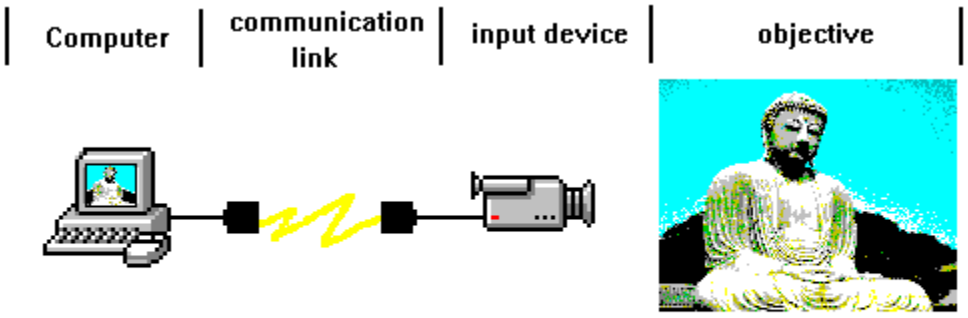

Detectable output is usually some form of electromagnetic (EM) output, over a certain frequency range such as: Infrared, Ultraviolet, X-rays, visible light, gamma-rays, etc.

Input Devices

Images are produced by a variety of input devices, and each input device typically produces a certain type of image. Input devices include: video cameras, scanners, electromagnetic sensors (a class of devices that collect various frequencies of light from the objective -IR, UV, visible, etc.). Each input device has a certain resolution, and image size. The input device usually converts the image from analog data to digital pixel values.

Communication Link

After the image is produced, it is transferred to the computer by some form of communication link. This link can be anything from a simple computer cable, to a complex satellite link, traveling from an space probe. Typical links used are: computer cables, modems, LANs, radio links (for remotely situated input devices), diskette (from standalone input systems), GPIB (IEEE standard), etc. Data transfer is usually in digital form (the input device converts the data to digital format) to ensure reliable data transfer.

Why do you need Image Processing?

Throughout the process of acquiring the image data, the pixel values go through numerous devices and interfaces along the path from the objective to the computer. Each device or interface can change the image data from how it should look under ideal circumstances. For example: there can be interference between the objective and the input device, such as haze or EM noise; faulty input devices can cause analog noise in the image data; if a radio link is used to transfer the data from a remote location, digital noise can be introduced; etc.

Image restoration is the means by which we can restore the image data to how it should look under ideal conditions (i.e. no noise).

Image enhancement is then used to change the image data to a more useful representation of the objective. For example: searching for objects in the image, increasing the visibility of low contrast artifacts, shadowing the image to make subtle details more pronounced.

Starting an IPS session

An IPS session involves four major procedures:

- Activating the IPS application
- Retrieving an image
- Processing the image
- Saving the image

When starting the IPS application a blank (black) image is displayed. This image is one of two blank images that IPS initializes, named "untitled1" and "untitled2". These images are 256x256 memory images (256 columns and 256 rows) for processing image data, and you can use the Create Image Command to create more images. Use the Source Command to switch between available memory images.

All functions that require cursor control can be carried out with the keyboard in addition to the mouse. Mouse and keyboard input can be used in parallel.

At any particular time the main window shows one image of the whole number of open images in the complex. This image is called current. For a detailed view of images the mode Zoom is used.

Closing these images frees up computer memory, and improves system performance.

Press  to find out how to load image data from a file.

Mouse and Keyboard Usage

To set the cursor in the desired position, first set it approximately by the mouse and then fix accurately using the keyboard. As a rule, control by arrows is similar to that with the mouse: click and release of the left key of the mouse corresponds to Space, double click of the left key to Ctr+Space. Pressing on the right key corresponds to Esc.

As the cursor appears in the field where the image is located it acquires the form of a crosshair. Using the cursor, one can select the desired point or construct a section.

To begin image processing, an image data file must be opened, using the Open File Command. Opening an image data file causes the image data stored in the file to be can be copied to the currently displayed memory image.

Use the Format Command from the Open File dialog to change the image file format to the desired format. This option can also be used to change the image files data size, shape, or load only a smaller part of the whole image if the image is larger than 256x256 pixels.

NOTE: The IPS system only supports 8 bit image data, both gray scale and indexed types. If the image was an indexed image, the palette colors are displayed, but in order to start processing, the image must be converted to a gray scale image. If the image was a gray scale image the current IPS palette will be used to visualize the data. To change the palette, use the Palette Command from the View Menu.

Press  to find out how to process the image.

Processing the Image

Image processing is performed by executing a series of operations on the image in discrete, consecutive steps. Each operation is performed on the output from the previous operation. IPS operations are almost always done on the currently displayed image data (except when specifying a different source image in the Image Arithmetic Command).

A very useful feature of the IPS is the UNDO Command, which allows the previous operation to be reversed for comparison. Re-executing the UNDO command redoes the aforementioned operation.

The IPS system performs processing operations only on the region of interest. On the image one can select any rectangular area as the region of interest. The current area of interest is outlined by a dotted rectangle. All operations on the image are carried out only within the area of interest, which makes it possible to process differently independent image fragments.

It is also useful to periodically save the image data in a temporary image file, when good intermediate results have been achieved. It is very bad practice to save the image data in the previously opened image file, because this obliterates the original image data! Use a new file name for the intermediate and final image data to avoid this.

Every image created by the program has its own title, which is displayed in the caption of the main window. When reading the image, its default title is the name of file from which the image has been read. The Save operation allows one to save the image in the file with its current title.

For your convenience, the IPS allows assigning new titles to memory images. This will help you keep track of the various images by using a relevant title describing the image. Use the Title Command to specify up to 20 characters for the image title.

Press  to find out how to save the image data.

To enter the region of interest area selection, doubleclick the left mouse key or Ctrl+Space on the keyboard. Besides, one can also enter the mode from the menu EDIT by selecting Set Edit Region. On entering the mode the cursor acquires the form of a rectangle with an arrow. To choose the area of interest that coincides with the entire image, doubleclick the left mouse key or press Ctrl+Space on the keyboard. To choose a rectangular area as an area of interest, first fix position of one of the corners by holding down the left mouse key or Space. Now stretch the rectangle by moving the cursor with a pressed mouse key or using the arrow keys (on the keyboard), a "rubber rectangle" is drawn. Its opposite corner is fixed by releasing the left mouse key or pressing Space. One can quit the mode at any moment and save the outlined area of interest by pressing the key Esc or the releasing the left mouse button.

When you have completed the image processing and are satisfied with the results, use the Save As Command to save the image data to a new image file name.

You can also print the processed image using the Print Command an image involves recording the image data to a piece of paper by using a gray scale mapping. When the printer is not capable of producing shades of gray directly, a method called dithering is sometimes used to produce the same effect as a gray scale by numeric manipulation of the image data.

Now Press here to find out about image restoration.

Overview

Image restoration is a type of image processing that is involved with changing the Image data to represent the original source (the objective being shown in the image) as closely as possible. Imperfections in the image data can usually be attributed to three major sources: analog noise - from the input device, or from the observational medium; digital noise - caused be faulty data transfer; and geometric distortions - caused from perspective of the input device.

For example a video picture on a remote location sent by radio signal to the control center and inputted into the computer goes through the video camera, A/D converter, transmitter, receiver, amplifier, and disk controller. Through any of these stages the Image can be degraded by noise, faulty equipment, or low resolution.

Analog Noise

Analog noise can usually be identified by the image looking splotchy or speckled.

Ways to reduce this type of noise are through the use of smoothing, averaging, and median filters.

Digital Noise

Digital noise produces rectangular areas or strips on the image, and the most prominent characteristic of these strips is that they have almost the same pixel values, regardless of the changes seen in the surrounding image.

Ways to reduce this type of noise are through the use of spot and median filters.

Geometric Distortions

Geometric distortions are harder to detect than noise, because you must have some kind of reference to compare the image data to. Typically, geometric distortion will cause straight lines to become curved, but in order to evaluate which parts of the image are distorted, and by how much, you must have several reference points visible in the image, that you can measure the distortions against.

Once you have the reference data available, you can use the Geometric Transformation Command to correct the images distortion. First you must draw lines that show the changes that must be made in the image, in order to correct the distortions, using the Set Points Option. The more points you define to transform the image data, the more accurate the transformation will be. Secondly you must perform the Geometric Transformation to reposition the pixels in their correct places.

Additional Methods

Other possible restoration abilities include correcting the contrast or dynamic range of the image using equalization or normalization (either locally or globally).

Overview

Image enhancement is primarily involved with modifying the image to bring out features that were not noticeable previously. Such operations include outlining, background elimination, edge detection, object distinction, contour mapping, contrast enhancement.

It is possible to enhance the image to such an extent that the original objective is no longer recognizable, to enable distinction of very subtle features of the image.

Contrast Enhancement

Sometimes it might be necessary to over enhance the images contrast, when the image area that we are interested in has very subtle features, and we do not require the area outside of our local region of interest to be meaningful.

Sharpening

When the image originally contained large amounts of noise, and we had to perform extensive filter processing, the image may become blurred. To improve the sharpness of the image we can use the Convolution Command and specify a High Pass Filter kernel, such as:

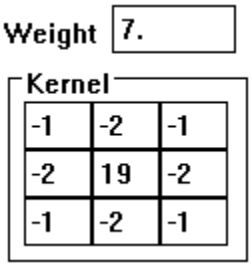

Shadowing

In some cases, a shadowing effect can enhance the details we wish to visualize, such as for images with uniform background coloration. The Gradient Command can be used to produce a shadowing effect on the image, and this command allows choice of the direction the shadow will be projected from, and the shadow length in both the X and Y directions.

Edge Detection

There are several processing operations that can be used to enhance edges in images. These include convolutions, local equalization, laplacian, and gradients

Convolutions using an edge detection kernel can enhance edges in any direction desired. For example:

Northwest

Kernel

1	1	1
1	-2	-1
1	-1	-1

North

Kernel

1	1	1
1	-2	1
-1	-1	-1

Northeast

Kernel

1	1	1
-1	-2	1
-1	-1	1

West

Kernel

1	1	-1
1	-2	-1
1	1	-1

Kernel

-1	1	1
-1	-2	1
-1	1	1

East

Kernel

-1	-1	1
-1	-2	1
1	1	1

Kernel

-1	-1	-1
1	-2	1
1	1	1

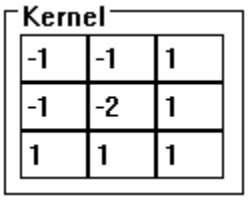

Southwest South Southeast

Local equalization using a small slide window will cause edges in any direction to be more pronounced.

Laplacian operations enhance and object in the image of the given window size, so if the window is set very small, this will enhance edges.

Gradient operations can show the edges by shadowing the image.

Overview

Advanced processing involves any type of processing that doesn't restoring or enhancing the image.

Some types of advanced processing are objects processing and producing visual effects, such as image combinations.

Object processing includes: locating objects; sorting them; defining their shape, size, and position; enlarging or reducing their spatial size; connecting or separating closely spaced objects.

Object Detection

Detecting objects can be performed using the Connect Objects Command. This command locates the objects according to the parameters entered, sorts them by size and gives each object's pixels a unique value (255 for the largest object found). The total number of objects that can found is 255. The pixel value of 0 is reserved for the background (non object area).

After the objects have been found, you can compute the object statistics using the Object Statistics Command. This will display the computed parameters for each object that was found using the connect object command.

Object Enhancement

After we have found the desired objects, it is possible to enhance their appearance in several ways

Objects that are very close to one another spatially or only slightly connecting can be connected together as one object, or the distance between them enlarged by using the morphology commands Closing and Opening.

Objects can also be enlarged or reduced with the morphology commands dilating and eroding.

Another procedure to enhance the objects can be achieved by overlaying the original image on the objects found so that only the object areas of the original image are shown, and the rest of the image is zero. This is done by thresholding the object image so that all object pixel values are 255 (i.e. min=1, max=255, fill=255) and then performing image arithmetic and ANDing the thresholded object image with the original image.

Visual Effects

Here are several types of visual effects used in image processing:

Image combinations: This can be achieved by saving the image you wish to combine into a file, and then choosing the region where you want to place the image by resizing and scrolling the window, and then opening the image file over the same source image. Set the read write format to produce the desired result (X and Y axis shifting and compression).

Books on Image Processing

U. Prett. Digital Processing of Images. 2 books, Moscow, Mir, 1982.

Fast algorithm in Digital Processing of images. Transformations and Median Filters. Ed. T.S. Huang, Moscow, Radio: Svyaz, 1984.

R. Beits, M. McDonnel, Image Reduction and Reconstruction. Moscow, Mir, 1989.

L. P. Yaroslavsky. Introduction into Digital Processing of Images. Moscow, Sovetskoe Radio, 1979.

Format description can be found in: S. Rimmer. BitMapped Graphics, Windcrest Books, 1990.

TIFF specification revision 6.0 Adus corporation, 1992.

Glossary

Image File
image
image enhancement
image restoration
histogram
Look-Up-Table(LUT)
byte
pixel
kernel
region of interest

Image File

Image data is usually stored in files on disk, and these files can have various formats. Some of the more popular formats are TIFF (Tagged Image File Format), GIF (Graphics Interchange Format), PCX (Picture Format), and BMP (Windows Bitmap Format). The image information stored in the data file is usually the size of the image (rows and columns), the type of image (gray scale, indexed, or true color), and the palette information if needed.

The image file is a file containing the image data, in a byte per pixel format (8 bits depth). The file starts from an optional header followed by rows of data. The term "interleaved data" is used when several channels have been recorded and the data of each pixel channel are in sequential bytes of the file. In files of this type, the channels are called "bands".

Image

An image for image processing is any number of rows and columns of byte per pixel data (8 bits depth). The IPS system treats images as being 256x256 or 512x512 in size, and all processing operations will be performed on this size image. If an image is smaller or larger than this standard size, IPS will fill the rest of the image with zeros or truncate the extraneous data. The Read/Write Format command provides the option of changing this behavior.

Image enhancement

Image enhancement comprises combinations of operations and commands that magnify certain features or characteristics of the original image. Processing Commands can be used to enhance images.

Image restoration

histogram

See Histogram Analysis

Look-Up-Table (LUT)

The Look Up Table is an array of pixel values that define how the current image pixel values will be mapped into new pixel values. For example, if you set the old (or input) value to 200 and the new (output) value to 255, all pixels of the image that have the value 200 will be changed to the value 255.

See Also LUT Command

byte

A byte is the basic unit of computer storage and operation. It is a element of memory that contains 8 binary bits and can have a value from 0 to 255.

pixel
Short for picture element, this is the basic unit of an image. Each pixel of an image has a value, and when these values are displayed in the proper order, the image is displayed.

Region of Interest

This is a subset of the whole image used to limit all operations to this area.
To select a region see Region Selection

kernel

This is a group of numbers that defines how the convolution command will change the image. The numbers are arranged in a 3x3 array, and each number is the multiplier for the corresponding adjacent pixel of the image. The result of this operation for one pixel is the sum of the product of each pixel times its multiplier, divided by the weight.

Median

The median is determined by the surrounding pixels such that half of the pixels for the given region are greater in value than the median and half are less.

IPS Main Index

Processing Commands

arithmetic
average, global
closing
connect objects
contour map
convolution
dilating
equalization, local
equalization, global
eroding
gradient
isoline
kernel
Laplacian Operator
look up table
median
mirror
morphology
move image
negative
normalization, local
normalization, global
object statistics
objects
opening
pixel brightness
R-filter
rotate
smooth
spots
thresholding
variance

Image Control Operations

Help Sections

Keyboard Commands

To avoid a long selection of a desired operation by going through the sections of the main menu, the operation can be called using the corresponding combination of buttons as shown below. Press them simultaneously, and a desired dialogue will be activated or a desired operation will be performed.

See Mouse and Keyboard Usage for more information.

Open	Ctrl + O
Save As	Ctrl + S
Print	Ctrl + P
Exit	Alt + F4
UNDO	Ctrl + U
Copy	Ctrl + C
Paste	Ctrl + A
New Title	Ctrl + T
Copy Image	Ctrl + I
Source	S
Palette	Ctrl + L
Zoom	Ctrl + Z
Image Negative	Shift + G
Global Equalization	Shift + E
Image Thresholding	Shift + T
Global Normalization	Shift + N
Make	Shift + L
Median Filter	Shift + M
Average Filter	Shift + A
Spots Filter	Shift + S
RFilter	Shift + R
Smooth	Shift + W
Laplacian	Shift + P
Equalization	Shift + Q
Normalization	Shift + Z
Variance	Shift + V
Gradient Operator	Shift + G
Convolution Operator	Shift + C
Connected Objects	Shift + J
Objects Statistic	Shift + O
Isoline	Shift + I
Contour Map	Shift + K
Dilating	Shift + Ctrl + D
Eroding	Shift + Ctrl + E
Opening	Shift + Ctrl + O
Closing	Shift + Ctrl + C
Move Image	Shift + X

Rotate	Shift + Y
Axis X (Mirror Menu)	Shift + Ctrl + X
Axis Y (Mirror Menu)	Shift + Ctrl + Y
Diagonal (Main) (Mirror Menu)	Shift + Ctrl + M
Diagonal (Second) (Mirror Menu)	Shift + Ctrl + S
Exec Transformation	Shift + H
Set Points	Shift + F
Show Set Points	Shift + U
Calibration	Ctrl + B
Histogram	Ctrl + H

Procedure Index