

Smoothing & Sharpening Functions

Introduction to Smoothing & Sharpening Functions

Smoothing and sharpening function use the pixels in an $N \times N$ neighborhood about each pixel to modify an image. For both smoothing and sharpening filters the larger the $N \times N$ neighborhood the stronger the smoothing or sharpening effect. Smoothing and sharpening function can be either non-adaptive or adaptive to local statistics found in each $N \times N$ neighborhood of an image.

Non-adaptive functions use the same calculation in each $N \times N$ neighborhood to modify a pixel. For example, a box smoothing filter smooths an image by calculating the average of all pixels in the $N \times N$ neighborhood of a pixel and replaces the pixel at the center of the $N \times N$ neighborhood with the average value. Gaussian blur and unsharp mask are also examples of non-adaptive filters. Non-adaptive filters are represented as an $N \times N$ convolution kernel. The dimensions of the $N \times N$ convolution kernel are typically 3×3 , 5×5 , 7×7 , 9×9 , or larger. A 3×3 convolution kernel is denoted by

$$\begin{vmatrix} a1 & a2 & a3 \\ a4 & a5 & a6 \\ a7 & a8 & a9 \end{vmatrix}$$

A convolution is a one-to-one linear function F that maps an $M \times N$ image Z and a $N \times N$ convolution kernel C onto a new $M \times N$ image W . The function F has the following properties:

- 1) A pixel from Z is mapped to the same position in W .
- 2) If the convolution kernel C is given by

$$\begin{vmatrix} c1 & c2 & c3 \\ c4 & c5 & c6 \\ c7 & c8 & c9 \end{vmatrix}$$

and

$$\begin{vmatrix} z1 & z2 & z3 \\ z4 & p & z5 \\ z6 & z7 & z8 \end{vmatrix}$$

is a neighborhood about a pixel p in Z , then p is mapped by the convolution function to a pixel q in W where

$$q = c1*z1 + c2*z2 + c3*z3 + c4*z4 + c5*p + c6*z5 + c7*z6 + c8*z7 + c9*z8$$

The above definition which is stated for a 3×3 image neighborhood and 3×3 convolution kernel can easily be generalized to a $N \times N$ image neighborhood with an $N \times N$ convolution kernel.

If the coefficients of a convolution kernel sum to one then the kernel represents smoothing or low pass filter. If the coefficients of a convolution kernel sum to zero then the kernel represents a sharpening or high pass filter.

Adaptive functions calculate statistics in the $N \times N$ neighborhood about each pixel and use the local statistics to modify the calculation that is used to modify the pixel at the center of the $N \times N$ neighborhood. Examples of adaptive smoothing and sharpening filter in ImagesPlus include the adaptive edge preserving selected neighborhood smoothing filter and adaptive unsharp mask, respectively.

Custom filters in ImagesPlus allow the user to enter their own non-adaptive filters by specifying the convolution kernel. User defined filters include gradient, emboss, outlier, smoothing, sharpening, and edge filters.

Multi-scale resolution tools in ImagesPlus allow the user to select which detail levels of an image are effected by smoothing and sharpening. Image detail levels can range from small and very fine detail to large scale detail.


The frequency based tools in ImagesPlus use the Fourier transform to convert a color or monochrome image from its usual spatial domain into the frequency domain representation. In the frequency domain the amplitudes of low and high frequencies can be modified to smooth and sharpen an image. The Frequency Filter tool in ImagesPlus has low pass smoothing, high pass sharpening, band pass, and band stop options. Light and dark line patterns can be suppressed or enhanced using the Line Suppression & Enhancement tool. Wider color bands can be removed from an image using the Band Suppression tool.

All of the smoothing and sharpening filters can be used with layers, blend modes, opacity, and masks to increase their usefulness. The dark background of an image can be smoothed using an inverse luminance masks without effecting the brighter areas of the image. A luminance mask can be used to sharpen the bright areas of an image without effecting the dark areas of the image. The luminance of an image can be sharpened and used as the new luminance component of an image. The [Making & Using Masks](#) and [Layers, Blend Modes, Opacity, & Masks](#) chapters of the ImagesPlus user manual contain detailed information creating masks and using them with layers, blend modes, and opacity.

Noise in an image can take several forms. Impulse noise appears as salt and pepper noise. Additive noise is independent of image intensity. Multiplicative noise is a function of image intensity.

Smoothing & Noise Reduction

The Smoothing & Noise Reduction commands are used to blur an image. The non-adaptive filters apply the same convolution kernel to each pixel of an image. The adaptive filters adapt to the local statistics of an image. The Filter Strength control can be used to vary the effect of the filter from no change to maximum strength for a given blur window kernel size.



A smoothing filter can be applied to just the background, midtones, highlights, or any other range. The Min Apply and Max Apply sliders and check boxes are used to select a range of brightness values where the smoothing filter is applied. The Feather Range slider is used to blend the smoothed data defined by the min and max apply levels into the image data.

The Set View toolbar button is used between each successive application of a smoothing filter with different parameters. For example, after the background is smoothed using one set of parameters the Set View toolbar button is pressed to save the smoothed result before smoothing with different filter parameters or apply range.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Non-adaptive Filters

Box: Selects a box or average filter with a low-pass convolution kernel of identical elements. For example, a 3x3 box filter has a convolution kernel of

$$1/9 \begin{vmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{vmatrix}$$

The box filter is a strong smoothing filter which is stronger than the Gaussian or binomial filter.

Binomial: Selects a binomial low-pass convolution kernel with rows and columns formed from the product of binomial coefficients. A 5 x 5 binomial convolution kernel would have coefficients as shown by

$$(1 / (\text{sum of coefficients})) \begin{vmatrix} 1 & 4 & 6 & 4 & 1 \\ 4 & 16 & 24 & 16 & 4 \\ 6 & 24 & 36 & 24 & 6 \\ 4 & 16 & 24 & 16 & 4 \\ 1 & 4 & 6 & 4 & 1 \end{vmatrix}$$

The binomial filter is a weaker smoothing filter compared to the Gaussian or box filter.

Gaussian: Selects a Gaussian low-pass convolution kernel with elements that form a Gaussian distribution. A global Gaussian filter can be applied using the Smoothing & Noise Reduction command. A local Gaussian filter can be applied using the mouse and the Mask Tools command. Produces a medium smoothing effect.

Median: Selects a standard rank-order median filter. The intensity of a pixel P is replaced with the median intensity of all pixels in the N x N neighborhood about P.

Adaptive Detail Preserving Filters

Selected Neighborhood: The selected neighborhood filter is an adaptive filter that performs a variable weighted average from pixel to pixel. The weights depend upon the difference between the intensities in the neighborhood of a pixel and the intensity at the pixel. This filter tries to suppress noise without losing edge sharpness. Excellent filter for removing hot and cold pixels while preserving edge sharpness when used with a 3 x 3 or 5 x 5 neighborhood or blur window.

Absolute Difference A median filter variant that selects the absolute difference trimmed mean filter. An adaptive rank-order filter with improved preservation of edges and excellent suppression of impulse and additive noise.

Quantile: The quantile filter is an adaptive filter that performs weak smoothing on dark areas that are rich in contrast. In bright areas the smoothing is stronger. The quantile filter is an appropriate filter for smoothing multiplicative noise.

Adaptive Median Selects the K-nearest neighborhood mean filter.

Strength: Enable only when the adaptive median filter is selected. Sets the number of pixels, from the local N x N window, that are used to calculate the resulting median pixel.

Min and Max apply can be set using the Set Min Max Apply button on the Histogram command or by a left-click on an image. Changes made to Min and Max apply are updated on the Histogram command if the Histogram command is open.

Min Apply: Specifies the minimum pixel brightness value where the smoothing filter is applied. The default value is 0 or black. Check the Min Apply box to the left of the slider then left-click on the minimum brightness in the image where the smoothing filtered is to be applied.

Max Apply: Specifies the maximum pixel brightness value where the smoothing filter is applied. The default value is 65535 or white. Check the Max Apply box to the left of the slider then left-click on the maximum brightness pixel in the image where the smoothing filtered is to be applied.

Feather Range: Specifies a feather range greater than the max apply level and less than the min apply level where filtered data is blended or feathered into the unaffected image data.

Filter Strength: Varies the effect of the blurring filter from no change at 0% to the maximum filter effect at 100%.


Blur Window Size: Selects the size of the N x N convolution kernel. The size can range from 3 x 3 to 999 x 999 when using the box, binomial, or Gaussian filters. The median, selected neighborhood, absolute difference, quantile, and adaptive median filters have convolution kernels that range from 3 x 3 to 51 x 51. The larger the blur window size the stronger the blurring effect.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Smoothing & Noise Reduction dialog box.

Multiresolution Smoothing/Sharpening



Wavelet based smoothing and sharpening tool. This command can be used to independently smoothing or sharpening different detail levels of a color or monochrome image. For example, fine detail can be sharpened or smoothed without effecting larger scale detail. Larger scale detail can be smoothed or sharpened without changing smaller scale detail. The detail levels correspond to a local neighborhood window size as follows:

Finest	3x3
Fine	5x5
Medium	7x7
Large	9x9
Largest	11x11

Smoothing or sharpening can be applied to just the background, midtones, highlights, or any other range. The Min Apply and Max Apply sliders and check boxes are used to select a range of brightness values where the smoothing or sharpening filter is applied. The Feather Range slider is used to blend the smoothed or sharpened data defined by the min and max apply levels into the unaffected image data.

The Set View toolbar button is used between each successive application of a smoothing or sharpening filter with different parameters. For example, after the background is smoothed or sharpened using one set of parameters the Set View toolbar button is pressed to save the result before smoothing or sharpening with different parameters or apply range.

Min and Max apply can be set using the Set Min Max Apply button on the Histogram command or by a left-click on an image. Changes made to Min and Max apply are updated on the Histogram command if the Histogram command is open.


Control Summary

Min Apply: Specifies the minimum pixel brightness value where the smoothing or sharpening filter is applied. The default value is 0 or black. Check the Min Apply box to the left of the slider then left-click on the minimum brightness in the image where the smoothing or sharpening filtered is to be applied.

Max Apply: Specifies the maximum pixel brightness value where the smoothing or sharpening filter is applied. The default value is 65535 or white. Check the Max Apply box to the left of the slider then left-click on the maximum brightness pixel in the image where the smoothing or sharpening filtered is to be applied.

Feather Range: Specifies a feather range greater than the max apply level and less than the min apply level where filtered data is blended or feathered into the unaffected image data.

Reduce Artifacts: Eliminates haloes around bright objects.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Sharpen Slider Strength

Finest: Finest sharpening with control range [1, 4]

Very Fine: Very fine sharpening with control range [1, 5].

Fine: Fine sharpening with control range [1, 7].

Medium: Medium sharpening with control range [1, 10].

Strong: Strong sharpening with aggressive control range [1, 20].

Enable Slider If checked, moving a slider will automatically apply the filter to an image.

To All Images: The operation is applied to each image if this box is checked.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Reset All: All slider positions are set to 0.0% or no change.

Done: Choose this button to close the dialog box.

Reduce High Low Bloom Pixel



Problem high, low, and bloom pixels can be removed using this tool. After the filter is applied to an image test the correction using a brightness and contrast increase function like digital development, levels, or curves. Undo and the Process History window can be used to reverse the brightness increase and allow parameter adjustment. After the parameters are set for a given image type use the Save Current Parameters As button to save the parameters for future use with this command or the extra filters option of the auto image set processing light frame parameters.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Reduce

High Pixels: Only high pixels are removed.

Low Pixels: Only low pixels are removed.

High and Low: Both high and low pixels are removed.

Bloom: Monochrome CCD bloom spikes are removed.

Reduction Strength

Low: Reduced high, low, and bloom pixel detection and removal.

Normal: Normal high, low, and bloom pixel detection and removal.

Window Size: The convolution reduction kernel. Larger window size reduces more of a problem pixel.

Effect Size: Size of the area in pixels that is reduced around each problem pixel.

Background to High Low Pixel Difference

Defines the difference between the image background and problem pixels. A smaller difference removes more problem pixels. Adjust this value to remove problem pixels and leave stars unaffected.

Saved Parameter List

Save Current Parameters As: The name used to save the current parameters. The parameter name is entered into the saved parameter list box of this command window and also the Bloom, High, Low Pixel Reduction list box in the Extra Filters section of the Light Frames Processing Parameters window of Image Set Operations | Auto Image Set Processing.

Saved Parameter List: List of all saved parameter settings.

Save: Saves the current parameters using the name entered in the Save Current Parameters As box. The name is entered into the saved parameter list. Saved parameters can be selected and set on this command or used by auto image set processing by selecting the saved parameter name on the Light Frame Processing Parameters window.

Update: Edit the parameters saved with the current name selected in the saved parameter list.

Delete: Deletes the saved parameter entry displayed in the saved parameter list.

To All Images: The operation is applied to each image if this box is checked.


Enable Slider If checked, moving a slider will automatically apply the filter to an image.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the dialog box.

Background/Highlight Smoothing



Iterative wavelet based smoothing tool for reducing fine grain noise from background areas. A preview image is used to adjust parameters before they are applied to a large full size image. To set noise parameters from the image check the **Select From Image** box then left-click on an area in the image.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Smooth Resolution Level

Finest: Smooths the finest level of detail.

Fine: Smooths fine detail.

Medium: Medium size detail is smoothed.

For each resolution level 0% is no smoothing and 100% is full smoothing.

Luminance and Noise Level

Select From Image: Check this box then left-click on a **Background** or **Midtone-Highlight** area in the image to set **Min Luminance**, **Max Luminance**, and **Noise Aggressive** smoothing parameters.

Min Luminance: Defines the lower luminance level to smooth.

Max Luminance: Defines the upper luminance level to smooth.

Noise Aggressive: Defines the maximum noise level within min and max luminance.

Iteration Control

Iterations: Sets the number of iterations to be performed. The range can be from 1 to 10,000 iterations.

Iterations Done: Displays the number of iterations performed.

Step Through: Check this box if you want to stop after each iteration. Use the Next button to perform the next iteration or the Cancel button to stop the iterative restoration.

Cancel: This button stops the iterative processing after a complete iteration. The Iterations Done text displays the number of iterations performed.

Next: Click this button to perform the next iteration.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Reset Image: Resets the image to its state before smoothing so that different parameters can be set using the mouse.

Apply: This button starts the iterative restoration process. The frame around the image will turn yellow indicating that a restoration process is in progress. The frame will return to blue when the process finishes or it is canceled.

Done: Choose this button to close the adaptive Richardson-Lucy iterative restoration dialog box.

Star Size Reduction/Smoothing



The Star Size Reduction and Smoothing option is an iterative process that reduces star size and also removes high and low pixels. Erosion and dilation variations are also included.

A preview image is used to adjust parameters before they are applied to a large full size image.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Star Size Reduction: Stars and problem pixels are reduced. The strength of the reduction is controlled by the **Windows Size** and **Strength** parameters. A larger window size produces greater size reduction for each iteration. For a given window size the reduction increases as the strength parameter increases.

Erosion: Image objects are reduced in size.

Dilation: Image objects are increased in size.

Iteration Controls

Iterations: Sets the number of iterations to be performed. The range can be from 1 to 10,000 iterations.

Iterations Done: Displays the number of iterations performed.

Step Through: Check this box if you want to stop after each iteration. Use the Next button to perform the next iteration or the Cancel button to stop the iterative restoration.

Cancel: This button stops the iterative processing after a complete iteration. The Iterations Done text displays the number of iterations performed.


Next: Click this button to perform the next iteration.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button starts the iterative restoration process. The frame around the image will turn yellow indicating that a restoration process is in progress. The frame will return to blue when the process finishes or it is canceled.

Done: Choose this button to close the dialog box.

Local Mouse Effects



The local mouse effect tools are used to apply local changes to an image. A local blur, sharpen, intensity increase, intensity decrease, or constant gray level can be applied to an image using the mouse. Local effects are applied using a left-click and drag of the mouse to define a path along which the effect is applied. The width of the path is defined by its effect size.

Each local path change to an image is tracked by the unlimited undo and redo built into the Process History command. The commands tracked by the process history list can be copied and applied to other images, saved in a project file for later use, or applied to a set of images using Filter Files.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Local Effects

Blur Use this effect to apply a localized Gaussian blur to the image along the path selected by the mouse. The Window Size list box controls the amount of blur. The Effect Size list box defines the width of the blurred path.

Sharpen Use this effect to apply a localized unsharp mask to the image along the path selected by the mouse. The Window Size list box controls the amount of sharpening. The Effect Size list box defines the width of the sharpened path.

Darker This effect is used to darken a mask along a path selected by the mouse. The Effect Size list box defines the width of the path. The Intensity Value list box controls the amount that is subtracted from the image.

Brighter This effect is used to brighten a mask along a path selected by the mouse. The Effect Size list box defines the width of the path. The Intensity Value list box controls the amount that is added to the image.

Fill With Intensity Value This effect is used to assign a constant gray value to a mask along a path selected by the mouse. The Effect Size list box defines the width of the path. The Intensity Value list box defines the constant gray level. A value of 0 defines black, 32500 defines medium gray, and 65535 defines white. Any value between 0 and 65535 can be used to define a shade of gray.

Clone Use this effect to clone pixels from one area of an image along the path selected by the mouse. The Window Size list box controls the distance from the path to the copied pixels. The Effect Size list box defines the number of pixels changed or width of the cloned path. After Effect Size is set right-click on the image to select the pixels to be cloned. Next set blend mode and feather radius for the clone option then left click and drag a path to clone the selected pixels.

Column Fix Use this effect to repair defective columns. Left-click to repair a column. The Pixel Distance value controls the distance from the column to the copied pixels.

Horizontal Line Fill Similar to the fill with intensity value option except the path is forced to be a horizontal line. Most often used to paint on the spectrum of an image to set frequency amplitudes equal to 0 to remove pattern noise.

Vertical Line Fill Similar to the fill with intensity value option except the path is forced to be a vertical line. Most often used to paint on the spectrum of an image to set frequency amplitudes equal to 0 to remove pattern noise.

None This option disables mouse input to this command. If the command is open while not in use then choose this option to disable all of the mouse controlled commands.

Tool Parameters

Window Size Defines the convolution kernel size or strength of the local blur and local sharpen effects. The weakest effect is achieved with a window size of 3x3. The effect increases as the window size is increased.

Effect Size Defines the size of the local effect. If effect size is 1x1 then the local effect is applied only to the pixel that is selected using the cursor. If the effect size is 15x15 then all pixels with a 15x15 neighborhood of the selected pixel also have the effect applied. Effect size defines the size of the paintbrush when using a left-click and drag of the mouse to paint a local effect.

Intensity Value This edit box is used to define the gray level used by the local and global fill effects and also local darker and brighter. A value of 0 defines black, 32500 defines medium gray, and 65535 defines white. Any value between 0 and 65535 can be used to define a shade of gray.

Blend Mode Sets the blend mode that is used to blend each path with the image. Normal, average, lighten (max), and darken (min) blend modes are supported.


Feather Pixel Radius Sets the width to be smoothed along the edge of a path.

Pixel Distance Enabled only when the column fix option is selected. The pixel distance value controls the distance from the column to the copied pixels.

Opacity Sets opacity that is used along with blend mode to blend the path into the image.

Done Use this button to close the dialog box.

Sharpening & Texture Enhancement - Adaptive Contrast



This operator provides locally adaptive contrast modification. The amount of contrast enhancement is locally adaptively controlled by an edge intensity L . The value L is a weighted average of all edge pixels in a local window about a pixel. The Prewitt edge operator is used to compute the edge pixels. The local contrast can be increased or decreased.

The result of this operator has higher or lower contrast at edges in comparison to the original image. The Filter Strength control on the RANGE TAB can be used to vary the effect of the filter from no change to maximum strength for a given kernel size.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Contrast

Increase Use the slider to increase contrast.

Value: Increase contrast parameter set by the slider.

Decrease Use the slider to decrease contrast.

Value: Decrease contrast parameter set by the slider.

Reduce Artifacts: Eliminates haloes around bright objects.

Window Size: The convolution kernel size. The range is from 5x5 to 65x65.

To All Images: The operation is applied to each image if this box is checked.


Enable Slider If checked, moving a slider will automatically apply the filter to an image.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Adaptive Contrast dialog box.

Statistical Difference



The Statistical Difference command is used to sharpen or smooth an image. The initial slider positions will generally produce a sharper image. Shifting the Deviation Factor slider to the right will sharpen the image while shifting the same control to the left will smooth the image. Image brightness and contrast can be controlled using the Edge/Background slider. The Mean Factor control is used to adjust the range of the contrast and brightness adjustment produced by the Edge/Background control. The strength of the smoothing or sharpening operation can be adjusted using the Gain Factor slider and Window Size control.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Edge / Background: Controls the ratio of edge to background.

Value The ratio set by the slider.

Gain Factor Prevents large output values when the deviation is small.

Value The gain factor set by the slider.

Mean Factor The desired mean factor.

Value The mean factor set by the slider.

Deviation Factor The desired standard deviation factor.

Value The deviation factor set by the slider.

Reduce Artifacts: Eliminates haloes around bright objects.

Window Size: The convolution kernel size. The range is from 3x3 to 65x65.

Enable Slider If checked, moving a slider will automatically apply the filter to an image.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

To All Images: The operation is applied to each image if this box is checked.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Statistical Difference dialog box.

Unsharp Mask



The Unsharp Mask command is used to sharpen an image. An image I can be modeled as the sum of a detailed image D and an unsharp image U ,

$$I = D + U$$

The unsharp component can be approximated by applying a low-pass or unsharp filter L_p to I ,

$$I * L_p = (D + U) * L_p \approx U$$

where $*$ is the convolution operator. The detail component can now be found,

$$D = I - U \approx I - I * L_p$$

The image sharpness is increased by a linear high-pass filter that results from the image difference between the original image and a low-pass filtered image. Several variations of this idea are possible.

The Unsharp Mask command uses a background weight factor and the detailed image is expressed as

$$D = (1 - a) * Lp / (1 - a)$$

where $0 \leq a < 1$.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

The custom filter is not tracked by Process History and must be selected before use using the Custom Filter command.

Control Summary

Unsharp Filter

Box The unsharp image is created using a box or average filter. Produces the strongest sharpening effect.

Median The unsharp image is created using a median filter. Produces a sharpening effect that is a little weaker than the box filter.

Gauss The unsharp image is created using a Gaussian filter. Produces a medium strength sharpening effect.

Binomial The unsharp image is created using a binomial filter. Produces a sharpening effect that is a little weaker than the Gauss filter.

Custom The unsharp image is created using a custom filter defined by the Custom Filter command. Allows the user to design their own unsharp or blurring filter.

Background Weight Factor A weight factor applied to the unsharp image.

Variant

Binarization A binary image is produced from the sharpened image using the threshold defined below.

Threshold All pixels above this value are set to white, otherwise a pixel is set to black.

Image Sharpening Image sharpening is performed.

Reduce Artifacts: Eliminates haloes around bright objects.

Window Size: The convolution kernel size. The lower cutoff frequency decreases by increasing the window size. The range is from 3x3 to 65x65.

Enable Slider If checked, moving a slider will automatically apply the filter to an image.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

To All Images: The operation is applied to each image if this box is checked.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Unsharp Mask dialog box.

Adaptive Unsharp Mask



The Adaptive Unsharp Mask command sharpens an image using an adaptive local window function. The Adaptive Unsharp Mask is calculated the same way as described for the unsharp mask except an adaptive low-pass filter for L_p is used.

An adaptive unsharp mask provides greater control with less noise amplification as compared to a non-adaptive unsharp mask.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Adaptive Unsharp Filter

Selected Neighborhood The unsharp image is created using a selected neighborhood technique.

Quantile The unsharp image is created using a locally adaptive quantile filter.

Background Weight Factor A weight factor applied to the unsharp image.

Adaptive Filter Weight A weight factor applied to the adaptive low-pass filter.

Reduce Artifacts: Eliminates haloes around bright objects.

Window Size: The convolution kernel size. The lower cutoff frequency decreases by increasing the window size. The range is from 3x3 to 65x65.

Enable Slider If checked, moving a slider will automatically apply the filter to an image.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

To All Images: The operation is applied to each image if this box is checked.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Unsharp Mask dialog box.

Range



The range controls are used to specify a feathered luminance range where the current filter is applied.

Control Summary

Min and Max apply can be set using the Set Min Max Apply button on the Histogram command or by a left-click on an image. Changes made to Min and Max apply are updated on the Histogram command if the Histogram command is open.

Min Apply: Sets the minimum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Min Apply then left click on the area of an image that you want to set as the minimum apply value.

Max Apply: Sets the maximum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Max Apply then left click on the area of an image that you want to set as the maximum apply value.

Feather Range: Specifies a feather range greater than the max apply level and less than the min apply level where filtered data is blended or feathered into the unaffected image data.

Filter Strength: Varies the effect of the filter from no change at 0% to the maximum filter effect at 100%.

Enable Slider If checked, moving a slider will automatically apply the filter to an image.

Reset Sliders: Sets the min, max, feather range, and filter strength sliders to their default position.

Apply: This button applies the selected filter to the current active image.

Introduction to Restoration

A star or other point source of light should theoretically correspond to a single pixel on the CCD array after passing through a camera. This situation never occurs in practice since the atmosphere and optics induce a blurring or spreading of a point into adjacent pixels. We can imagine a CCD image or photograph as a collection of point sources each of which is blurred or spread out by a point spread function. If we know the point spreading function that the ideal image was convolved with to produce the observed blurred image, then we can use the known point spread function to deconvolve the observed image thereby restoring the ideal sharp image.

If the image contains a star, then its blurred image can be used as the point-spread function or PSF. Otherwise we can assume that the image was blurred using a standard point-spreading function such as a box, Gaussian, binomial, or custom user defined function. The point spread function can now be used to restore the observed image to the ideal image.

This process can be stated as:

$$I(x,y) = O(x,y) * P(x,y) + N(x,y) \quad (1)$$

Where $I(x,y)$ is the intensity distribution or the blurred image captured by the CCD, $O(x,y)$ is the sharp object image before it passed through the atmosphere and optics, $P(x,y)$ is the point-spread function, $N(x,y)$ represents additive noise, and $*$ is the convolution operator.

The problem at hand is to solve the above equation (1) for $O(x,y)$ given the captured image $I(x,y)$, an estimate of $P(x,y)$ taken from either the observed image or given an assumed shape, and an unknown noise function $N(x,y)$. This problem does not have a unique solution because of the estimated point-spread function and unknown noise. So additional assumptions are made to regularize the problem. The solution can then be found using either a direct or iterative procedure.

Direct restoration methods try to restore the image in a single pass of the image pixels. Examples of this would be inverse, Weiner, and blind deconvolution restoration. This type of restoration works well on images with a high signal-to-noise ratio since in this case it is reasonable to assume that $N(x,y) \approx 0$.

Iterative restoration methods start with the observed image and create a sequence of improved images that converge to the ideal image. Iterative restoration methods restore the image in a much softer way than the direct methods. The iterative process is stopped when we judge that noise amplification is inhibiting. The maximum entropy, Van Cittert, one-step gradient, and Richardson-Lucy algorithms are examples of iterative restoration methods.

Iterative restoration algorithms can be classified as adaptive or non-adaptive. Non-adaptive restoration algorithms use a space-invariant point-spread function and apply the same number of iterations to each pixel of the image to be restored. Adaptive restoration algorithms use a space-variant point-spread function or apply a variable number of iterations depending on the local statistics at a given iteration. Good results are usually obtained using the non-adaptive methods after 10 to 15 iterations. The main drawback with non-adaptive restoration algorithms is that they diverge if the image contains noise or they try to over-fit the observed data to the restoration model producing a *speckled* appearance in the restored image. Adaptive restoration algorithms can greatly reduce or eliminate noise amplification and

the over-fitting of observed data to the restoration model. Good results are usually obtained using the adaptive methods after 10 to 100 iterations.

The Van Cittert method is very good at restoring high signal-to-noise ratio images. Images of the sun, moon, planets, and daytime terrestrial scenes would all be appropriate candidates for a Van Cittert restoration. If we assume that $N(x,y) = 0$, then the Van Cittert iterative method can be expressed as:

$$O(n+1)(x,y) = O(n)(x,y) + \alpha (I(x,y) - P(x,y) * O(n)(x,y))$$

Where $O(n)(x,y)$ is the restored image after n iterations and α is a relaxation or convergence factor generally taken to be 1. This method diverges in the presence of noise or if $N(x,y) \neq 0$. The term $I(x,y) - P(x,y) * O(n)(x,y)$ is called the residual error and represents remaining detail to be added to the restored image during the remaining iterations. A non-adaptive implementation of the Van Cittert method is accessed in ImagesPlus by using the non-adaptive iterative restoration dialog box.

The one-step gradient method is similar to the Van Cittert algorithm but produces a smoother restored image. The smoothing is caused by the additional convolution with $P(-x,-y)$. The one-step gradient algorithm is expressed as:

$$O(n+1)(x,y) = O(n)(x,y) + \alpha P(-x,-y) * (I(x,y) - P(x,y) * O(n)(x,y))$$

A non-adaptive implementation of the one-step gradient method is accessed in ImagesPlus by using the non-adaptive iterative restoration dialog box.

The Richardson-Lucy (R-L) method has been widely used for restoring HST images. The classical Richardson-Lucy algorithm produces restored images with several desirable properties:

- 1) The Richardson-Lucy iteration converges to the maximum likelihood image based on Poisson statistics. This is appropriate for optical images and noise from counting.
- 2) The restored image is not affected by small errors in the point-spread function.
- 3) The algorithm can be implemented efficiently on a computer.

The Richardson-Lucy algorithm is expressed as:


$$O(n+1)(x,y) = O(n)(x,y) [(I(x,y) / (P(x,y) * O(n)(x,y))) * P(-x,-y)]$$

The classical non-adaptive version of the Richardson-Lucy algorithm is limited by the over fitting of the observed image data to the restoration model. If several iterations are applied to an image a *speckled* appearance will begin to appear and this limits the number of iterations. An implementation of this version of the Richardson-Lucy method is accessed in ImagesPlus by using the non-adaptive iterative restoration dialog box.

The adaptive Richardson-Lucy algorithm applies more iterations in regions of the image where the restored image and restoration model differ greatly when compared to the noise. The iteration process is stopped in regions where the restored image and restoration model differ little when compared to the noise. In this way the *speckled* appearance caused by over fitting of the observed data to the restoration model can be greatly reduced or eliminated allowing a greater number of iterations to be applied. An implementation of this version of the Richardson-Lucy method is accessed in ImagesPlus by using the adaptive Richardson-Lucy restoration dialog box.

Photometry information is lost after a Van Cittert, one-step gradient, Richardson-Lucy, or maximum entropy restoration since these methods are nonlinear.

Non-Adaptive Iterative Restoration



This dialog box is used to select and control non-adaptive iterative restoration. The restoration is started using the Apply button and can be canceled using the Cancel Button. A preview image is used to adjust parameters before they are applied to a large full size image. Other images can be processed while this restoration method is being executed.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Restoration Method

Selects the type of restoration method to be performed. Choose the Van Cittert, one-step gradient, or Richardson-Lucy restoration method.

Van Cittert is an iterative restoration method that is well suited for enhancing the images of bright objects. This method can produce very fine detail but has low tolerance for noise. The Van Cittert method is similar to applying several unsharp masks of decreasing strength to an image.

The one-step gradient method is an iterative restoration method that is related to the Van Cittert method. The one-step gradient and Van Cittert methods both calculate the residual detail at a given iteration in the same way. At a given iteration the one-step gradient method smooths the residual detail before adding it to the restored image. The Van Cittert method just adds the residual detail to the restored image. Thus the restored image produced by the one-step gradient method is smoother than the restored image produced by the Van Cittert method.

Point Spread Function Type

Selects the type of point spread function to be used with the selected restoration method. The preset types are Box, Gaussian, and Binomial filters. The box, Gaussian, and binomial type produce a strong, medium, and weaker restoration effect respectively for a given PSF size and number of iterations.

The Custom type uses the current system PSF that is set using the Custom Point Spread Function dialog box. Custom PSF functions can be selected from an image by choosing an isolated star or entered using the keyboard. The custom point spread filter is not tracked by Process History and must be selected before use using the Custom Point Spread command.

Restoration Controls

Relaxation: Controls the strength of the restoration method. The setting depends upon the image to be restored, number of iterations, and PSF size.

The Van Cittert method generally works well when relaxation is set near 0.20. A value of 0 produces no restoration effect and a value of 2.50 produces the maximum restoration effect.

The one-step gradient method is similar to the Van Cittert method but produces a smoother restored image. The relaxation value is usually set to 1 when using the one-step gradient method.

The Richardson-Lucy method usually works best when relaxation is set to 1.0. A value less than 1 will darken the image and a value greater than 1 will lighten the image.

PSF Size: Controls the window size of the point spread function. A smaller PSF size like 3x3

generally produces finer results than a larger PSF size.

Min and Max apply can be set using the Set Min Max Apply button on the Histogram command or by a left-click on an image. Changes made to Min and Max apply are updated on the Histogram command if the Histogram command is open.

Min Apply: Sets the minimum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Min Apply then left click on the area of an image that you want to set as the minimum apply value.

Max Apply: Sets the maximum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Max Apply then left click on the area of an image that you want to set as the maximum apply value.

Feather Range: Specifies a feather range greater than the max apply level and less than the min apply level where filtered data is blended or feathered into the unaffected image data.

A good minimum apply value for restoring an image of a planet can be set by selecting the brightness a little inside the edge of the planet. For a deep sky object set minimum apply equal to the dark area of the object and above the background level. Artifacts that form in and around bright stars during restoration can be controlled by setting the maximum apply value between 50000 and 60000.

Reduce Artifacts: Eliminates haloes around bright objects.

Iteration Controls

Iterations: Sets the number of iterations to be performed. The range can be from 1 to 10,000 iterations.

Iterations Done: Displays the number of iterations performed.

Step Through: Check this box if you want to stop after each iteration. Use the Next button to perform the next iteration or the Cancel button to stop the iterative restoration.

Cancel: This button stops the iterative processing after a complete iteration. The Iterations Done text displays the number of iterations performed.


Next: Click this button to perform the next iteration.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button starts the iterative restoration process. The frame around the image will turn yellow indicating that a restoration process is in progress. The frame will return to blue when the process finishes or it is canceled.

Done: Choose this button to close the Iterative Restoration dialog box.

Adaptive Richardson-Lucy Restoration



The adaptive Richardson-Lucy iterative restoration method is used to control noise amplification in the restored image. The characteristic *speckled* appearance in restored images produced by the classical non-adaptive Richardson-Lucy algorithm can be greatly reduced or eliminated by using this adaptive method. By adjusting the noise threshold value and PSF size 40 to 100 iterations of this method can be applied without noticeable noise amplification. The minimum and maximum apply values are also used to control noise and artifacts in the restored image by controlling where the restoration is applied. The minimum and maximum apply values are used to define a brightness range that receives most of the restoration For example, if minimum apply is set to 0 and maximum apply is set to 65535 then restoration is applied equally to all areas of the image. Set minimum apply between 20000 and 30000 to apply the restoration to only the midtones and highlights.

Please see Introduction to Restoration for a general discussion of iterative restoration methods.

The restoration is started using the Apply button and can be canceled using the Cancel Button. A preview image is used to adjust parameters before they are applied to a large full size image. Other images can be processed while this restoration method is being executed.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Point Spread Function Parameters

PSF Size: Controls the window size of the point spread function. A smaller PSF size like 3x3 generally produces finer results than a larger PSF size.

Select the type of point spread function to be used with the adaptive Richardson-Lucy restoration. The preset types are Box, Gaussian, and Binomial filters. The box, Gaussian, and binomial produce strong, medium, and weaker restoration effect respectively for a given PSF size and number of iterations.

The Custom type uses the current system PSF that is set using the Custom Point Spread Function dialog box. Custom PSF functions can be selected from an image by choosing an isolated star or entered using the keyboard. The custom point spread filter is not tracked by Process History and must be selected before use using the Custom Point Spread command.

Application Parameters

Noise Threshold:

The noise threshold value controls the amount of noise amplification in the restored image. If noise threshold is set to 0.0, then the adaptive Richardson-Lucy algorithm reduces to the classical non-adaptive Richardson-Lucy algorithm. As the noise threshold is increased from 1 to 10 the amount of noise amplification decreases. The characteristic *speckled* appearance in restored images produced by non-adaptive Richardson-Lucy can be greatly reduced or eliminated by adjusting the noise threshold value.

Min and Max apply can be set using the Set Min Max Apply button on the Histogram command or by a left-click on an image. Changes made to Min and Max apply are updated on the Histogram command if the Histogram command is open.

Min Apply: Sets the minimum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Min Apply then left click on the area of an image that you want to set as the minimum apply value.

Max Apply: Sets the maximum brightness value that receives most of the restoration. The brightness can be measured using the Cross-Hair tool then set manually. For automatic selection check the Select box to the left of Max Apply then left click on the area of an image that you want to set as the maximum apply value.

Feather Range: Specifies a feather range greater than the max apply level and less than the min apply level where filtered data is blended or feathered into the unaffected image data.

A good minimum apply value for restoring an image of a planet can be set by selecting the brightness a little inside the edge of the planet. For a deep sky object set minimum apply equal

to the dark area of the object and above the background level. Artifacts that form in and around bright stars during restoration can be controlled by setting the maximum apply value between 50000 and 60000.

Reduce Artifacts: Eliminates haloes around bright objects.

Iteration Controls

Iterations: Sets the number of iterations to be performed. The range can be from 1 to 10,000 iterations.

Iterations Done: Displays the number of iterations performed.

Step Through: Check this box if you want to stop after each iteration. Use the Next button to perform the next iteration or the Cancel button to stop the iterative restoration.

Cancel: This button stops the iterative processing after a complete iteration. The Iterations Done text displays the number of iterations performed.


Next: Click this button to perform the next iteration.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button starts the iterative restoration process. The frame around the image will turn yellow indicating that a restoration process is in progress. The frame will return to blue when the process finishes or it is canceled.

Done: Choose this button to close the adaptive Richardson-Lucy iterative restoration dialog box.

Custom Restoration Point Spread Function



This dialog box is used to define custom point spread functions that are applied using the non-adaptive and adaptive restoration commands. Custom point spread functions allow the user to tailor a function to meet their needs. This dialog box is used to define and edit custom point spread functions. A custom PSF can be created by editing one of the PSF template files stored on disk, selected from an image using the mouse, or entered using the keyboard.

Control Summary

Point Spread Function Values The values of the convolution kernel. Values can be set using the mouse, reading a template PSF file from disk, or by entering a value.

PSF Size

Selects the size of the PSF convolution kernel. New values are automatically set to the constant defined by Value:.

PSF Name: Enter the custom point spread function name. If the function is saved to disk, then .psf is appended to the end of the name.

Persistence

Open: Opens a PSF file from disk and sets it as the current system PSF. Opens files with an extension of .psf.

Save: Saves the current PSF to disk using the name <PSF NAME>.psf and sets the current system PSF.

Edit PSF Values

The edit commands are used to perform operations upon the PSF grid elements. If a change is made to the PSF grid you must either save the PSF to disk or click Set PSF to update the current system PSF.

Set With Mouse: Check this box to enable setting the PSF grid elements using the mouse. Click the mouse on a grid element to set it to the constant defined by Value:.

Value: Defines a constant value that is used to fill the PSF grid elements by clicking them with the mouse or automatically when you increase the size of the grid.



Adds the constant Value: to each pixel of the PSF.



Subtracts the constant Value: from each pixel of the PSF.



Multiplies each pixel of the PSF by the constant Value:.



Divides each pixel of the PSF by the constant Value:.

Unitize: Divides each element of the PSF by the maximum value of the PSF. All elements will be changed to a value ≤ 1.0 .

Normalize: Changes the elements so that they all have a value ≤ 1.0 and sum to 1.0.

Positive: If the value of a grid element < 0.0 the element is replaced with 0.0.

Reset: Sets every PSF grid element to the constant defined by Value:.

Set PSF: Sets the currently displayed PSF grid as the system PSF.

Get PSF: Retrieves the system PSF and display its name and grid elements.

Select From Image: Check this box to choose a PSF by right clicking the mouse on an image point. The PSF Size is fixed at 9x9 and the Move View arrows are enabled.

Use the arrows to position the selected image area within the PSF grid. The maximum values are usually positioned at the center of the grid. If the maximums are offset from the image center the image will appear to flow in the opposite direction of the offset. Uncheck this box to save the selected image area in the PSF grid. Click the Set PSF button to set the current system PSF.

The grid is set to the pixel intensities in a 9x9 neighborhood about the mouse click point. For a color image the grid is set to the average intensity of the color channels.


Move View



The arrows are enabled when you are selecting a PSF from an image using the mouse. Use the arrows to move the selected area from the image in the direction shown.

Done: Select this button to close the Custom Point Spread Function dialog box.

Edge Extraction



The edge extraction operators generate an edge map. An edge map contains all contour lines of objects or image regions. Edge maps are used for image segmentation and can also be used to sharpen an image.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Standard Edge Extractors

One Pixel: Selects the one pixel edge operator. The edge map produced by this operator is defined by the difference between the current intensity and the minimum intensity inside a 3x3 window. The edges have a width of one pixel.

Sobel: Selects the Sobel edge operator. The edge map produced by the Sobel operator is the sum of a horizontal and vertical edge map. Some smoothing in the edge direction is performed to suppress noise.

Kirsch: Selects the Kirsch edge operator. The edge map produced by the Kirsch operator is produced from eight directional edge maps. Some smoothing in the edge direction is performed to suppress noise.

Laplace: Selects the Laplace edge operator. The Laplace edge operator does not suppress

noise. It follows that the Laplace operator is very sensitive to noise.

Morphological: Selects the morphological edge operator. The edge map is based on the fact that erosion and dilation shift the intensity values in opposite directions. The point-to-point difference between an image and the erosion or dilation of the image produces unsharp edge maps. The point-to-point combination of these edge maps by a minimum operation yields a sharp edge map.

Better Edge Extractors

Deriche: Selects the Deriche edge operator. The Deriche operator produces accurately placed edges, has good resistance to noise, and is sensitive to true edges. The filter has both high-pass and low-pass characteristics. The cutoff frequencies can be adjusted.

Contra-harmonic: Selects the contra-harmonic edge operator. The edge map is the result of the weighted difference between the local maximum and local minimum. This operator has good robustness against noise and outliers. Low contrast edges can be suppressed.

Edge Width: Adjusts the edge map contour line width.

Value: Edge width parameter value. Set by the slider or entered using the keyboard.

Window Size: Selects the size of the convolution kernel. The size ranges from 3x3 to 7x7.

Binarization: Enables binarization of the edge map.

Threshold: Sets the binarization threshold value.

Value: Pixel intensity less than or equal to this value is set to black, otherwise it is set to white.

Brightness Scale: Sets the brightness scale factor.

Scale: The scale factor applied to the edge map.

Enable Slider If checked, moving a slider will automatically apply the filter to an image.


To All Images: The operation is applied to each image if this box is checked.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.


Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Edge Extraction dialog box.

Custom Filters



Custom filters allow the user to tailor a filter to meet their needs. This dialog box is used to define, edit, and set the system custom filter. A custom filter can be created by editing one of the filter template files stored on disk or entered using the keyboard.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Filter Values The values of the convolution kernel. Values can be set using the mouse, reading a filter file from disk, or by entering a value.

Filter Size

Selects the size of the convolution kernel. New values are automatically set to the constant defined by Value.

Divide By Sum (Low Pass) The custom filter will be used as a low-pass filter. The sum of the filter values cannot be zero.

Filter Name: Enter the name of the custom filter. If you save the function to disk .ftr is appended to the end of the name.

Persistence

Open: Opens a filter file from disk and sets it as the current system filter. Opens files with an extension of ftr.

Save: Saves the current filter to disk using the name <Filter Name>.ftr and sets the current system filter.

Filter Values

The Filter Value commands are used to perform operations upon the filter grid elements. If a change is made to the filter grid you must either save the filter to disk or click Set Filter to update the current custom system filter.

Set With Mouse: Check this box to enable setting the filter grid elements using the mouse. Sets grid element that the mouse is clicked on to the constant defined by Value.

Value: Defines a constant value that is used to fill the filter grid elements by clicking them with the mouse or automatically when you increase the size of the grid.

Reset: Sets every filter grid element to the constant defined by Value.

Set Filter: Sets the current displayed filter grid as the custom system filter.

Get Filter: Retrieves the current system filter and display its name and grid elements.


To All Images: The operation is applied to each image if this box is checked.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the current custom system filter to the active image.

Done: Select this button to close the Custom Filter dialog box.

Line Suppression & Enhancement - Suppress



The Suppress operator is used to eliminate horizontal or vertical line patterns from an image. Vertical Height and Horizontal Width should be chosen to be a multiple of the width and height of the line to be suppressed. The Feature Mask tool can be used to split an image into high and low frequency images so that line suppression does not change the high frequency image data such as stars and edge detail. Bright areas of the image are also removed from the low frequency image using the Feature Mask tool. Line suppression is applied to the low frequency background image that contains the line pattern before it is merged with the high frequency image data to build the line suppressed image.

The Smoothing & Noise Reduction section of the ImagesPlus [image processing tutorial page](#) has several example of line suppression using the Line Suppression tool with the Feature Mask.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Eliminate

Horizontal Lines: The filter will eliminate horizontal lines.

Vertical Lines: The filter will eliminate vertical lines.

Line Detection Window Size in Pixels

Vertical Height: Sets the window height of the filter.

Horizontal Width: Sets the window width of the filter.


To All Images: The operation is applied to each image if this box is checked.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Line Suppression & Enhancement dialog box.

Enhance



The Enhance operator is used to enhance line-like patterns. Bands on Jupiter and Saturn can be enhanced using the Dark option.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Line Type

Light: Enhance light line patterns.

Dark: Enhance dark line patterns.

Light and Dark: Enhance light and dark line patterns.

Window Size: The window size ranges from 3x3 to 65x65.


To All Images: The operation is applied to each image if this box is checked.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Line Suppression & Enhancement dialog box.

Band Suppression



The Band Suppress operator is used to remove wide horizontal or vertical bands from an image. The Feature Mask tool is used to split an image into high and low frequency images so that band suppression does not change the high frequency image data such as stars and edge detail. Bright areas of the image are also removed from the low frequency image using the Feature Mask tool. Band suppression is applied to the low frequency background image that contains the band pattern before it is merged with the high frequency image data to build the band suppressed image.

The Smoothing & Noise Reduction section of the ImagesPlus [image processing tutorial page](#) has several examples of line and band suppression using the Band Suppression tool with the Feature Mask.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Suppress Vertical: Controls the width of vertical band suppression.

Suppress Horizontal: Controls the width of horizontal band suppression.

Suppression Strength: Controls the aggressiveness of band suppression.

Enable Vertical Band Suppression: Enables or disables vertical band suppression.

Enable Horizontal Band Suppression: Enables or disables horizontal band suppression.


Enable Slider If checked, moving a slider will automatically apply the filter to an image.

Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Done: Choose this button to close the Band Suppression dialog box.

Frequency Filter



The Frequency Filter transforms an image from the spatial domain to the frequency domain using a fast Fourier transform. The image in the frequency domain can be smoothed or sharpened by adjusting frequency amplitudes. If low frequency amplitudes are reduced or high frequency amplitudes are increased then the image is sharpened. If low frequency amplitudes are increased or high frequency amplitudes are decreased then the image is smoothed.

The Frequency Tools section of the ImagesPlus [image processing tutorial page](#) has several example of using the Frequency Filter tool. Start with the [Frequency Domain Concepts and Practical Applications of the Frequency Filter](#) tutorial for a detail description of frequency based processing.



The color channel toolbar buttons are enabled on a color image and have no effect on a monochrome image.

Control Summary

Frequency Amplitude Mask

The controls of the Frequency Amplitude Mask group are used to shape the frequency amplitude mask. The Low Frequency – High Frequency horizontal grayscale bar displays a radial section of the frequency spectrum mask from the low frequency center of the spectrum out through the high frequency area near the edge of the spectrum mask. To better understand the spectrum mask that is generated use the Show Spectrum and Show Spectrum Mask options as you change mask parameters.

Low Pass: Generates a round mask with center at the center of the frequency spectrum image so that low frequency amplitudes are unchanged while the outer high frequency amplitudes are reduce. The Low Pass option smooths an image since high frequency amplitudes are reduced. The [Smoothing an Image Using the Frequency Filter](#) tutorial is an example of using the Low Pass option.

High Pass: Generates a mask with a round dark center at the center of the frequency spectrum image so that low frequency amplitudes are reduced while the outer high frequency amplitudes are maintained. The High Pass option creates a high frequency edge map that can be used to sharpen an image if the Add Scaled Mask to Amplitude option is checked. The [Sharpening an Image Using the Frequency Filter](#) tutorial is an example of using the High Pass option.

Band Pass: Generates a white ring shaped mask centered about the low frequency center of the spectrum image. The Cutoff parameter defines the radius of the ring. The Width parameter defines the width of the ring. The Decay parameter defines how much the edges of the ring are smoothed out. Frequencies that are covered by the white ring mask are passed while other frequencies are blocked or reduced.

Band Stop: Generates a black ring shaped mask centered about the low frequency center of the spectrum image. The Cutoff parameter defines the radius of the ring. The Width parameter defines the width of the ring. The Decay parameter defines how much the edges of the ring are smoothed out. Frequencies that are covered by the black ring mask are blocked or reduced while other frequencies are passed.

Cutoff: Defines the cutoff radius measured from the center of the frequency spectrum.

Decay: Defines how much the edges of the ring or circular mask are smoothed out.

Width: Defines the width of the band pass or band stop frequency ring mask.

Scale Mask: Scales the mask by a factor and also by the 1x, 2x, 5x, 10x, 20x, or 100x multiplier.

Add Scaled Mask to Amplitude: If checked the scaled masked amplitudes are added to the frequency spectrum.

Spectrum

Show Spectrum: Displays the frequency spectrum of an image. Low frequencies are near the center of the spectrum image. Frequencies increase radially as you move from the center to the edge of the spectrum image. For a given frequency amplitude increase from black for 0 amplitude to white for maximum amplitude.

Show Masked Spectrum: Displays the spectrum mask generated by the controls in the Frequency Amplitude Mask group.

Show Spectrum Mask: Displays the round or ring shaped mask generated by the controls in the Frequency Amplitude Mask group.

Normalize:

Enable Slider If checked, moving a slider will automatically apply the filter to an image.


Lock Image Check this box to lock or couple an image with this command. Each time you press the Apply button the command is applied to the same image even if the image is not the current active image.

Apply: This button applies the selected filter to the current active image.

Reset: Sets the controls to their default position.


Done: Choose this button to close the Frequency Filter dialog box.

Split Frequency Spectrum



The Split Frequency Spectrum tool is used to split a color or monochrome image into grayscale, red, green, blue, or RGB spectra so that frequency amplitudes can be adjusted using the local mouse controlled paint tool which is located at Smooth Sharpen | Local Mouse Effects. A common use of painting on the frequency spectrum with the paint tool is to greatly reduce or remove fixed pattern noise such as horizontal and vertical lines and bands. The [Removing Structured Noise in the Frequency Domain](#) tutorial shows an example of removing vertical line noise by painting on the spectrum of the problem image. The [Detect & Remove Residual Structured Noise after Calibration and Stacking](#) tutorial is a more detailed example of painting on a frequency spectrum to remove pattern noise.

Note: As of ImagesPlus 6.5 x64 the paint tool is used to paint only black or 0 onto a frequency spectrum to set the painted frequency amplitudes to 0. Painting with any other color other than black has no effect.



The color channel toolbar buttons have no effect on the Split Frequency Spectrum tool.

Control Summary

Split Into Frequency Amplitude Spectrum

The controls in this group specify how the spectrum is split into component spectra.

Color: The spectrum is split into a color RGB spectrum so that the red, green, and blue spectra are accessed simultaneously. So painting on the RGB frequency spectrum changes the red, green, and blue spectra in the same way.

Red, Green, and Blue Spectra: Choose this option to create separate red, green, and blue spectra so that each color can receive separate painting.

Monochrome: This option appears only when the current active image is a monochrome grayscale image. A single grayscale frequency spectrum will be created and used for painting.

Link Spectra to Image: Check this box to automatically setup the split frequency spectra in the Combine Frequency Spectrum tool to rebuild the image from its paint modified spectra.

Apply: This button applies the frequency spectrum split to the current active image.

Done: Choose this button to close the Split Frequency Spectrum dialog box.

Combine Frequency Spectrum

Combine Spectrum

Input Spectra

Single Color R, G, B Spectra Single Monochrome

Select Input Spectra:

Single Color []

Red []

Green []

Blue []

Select Input Image

Single Color []

Combined Spectra Output Image Name:

CombineFrequencyTarget

Clear Selection Reset All Display Selection Display Combination

Apply Done

The Combine Frequency Spectrum tool is used to build a color or monochrome image from its frequency spectra and the image that was used to create the frequency spectra. The frequency spectra are created by applying the Split Frequency Spectrum tool to an image.



The color channel toolbar buttons have no effect on the Combine Spectrum tool.

Control Summary

Input Spectra This group of controls is used to specify the input spectra type.

Single Color: A single RGB frequency spectrum is used to build an image.

R, G, B Spectra: Separate red, green, and blue frequency spectra are used to build an image.

Single Monochrome: A single grayscale frequency spectrum is used to build an image.

Select Input Spectra

Single Color / Mono: Select this option then left click near the center of a color or monochrome grayscale frequency spectrum image to select it.

Red: Select this option then left click near the center of the red frequency spectrum image to select it.

Green: Select this option then left click near the center of the green frequency spectrum image to select it.

Blue: Select this option then left click near the center of the blue frequency spectrum image to select it.

Select Input Image The input image is the image that was used to split the frequency spectra images. The input image provides phase information while building the image from modified frequency spectra images.

Single Color / Mono: Select this option then left click near the center of a color or monochrome grayscale image to select it.

Combined Spectra Output Image Name: The name of the image that is built from the input image and the frequency spectra images.

Clear Selection: Selection a frequency spectra in the Select Input Spectra group or the input image in the Select Input Image group then press this button to clear the selection so another image can be selected.

Reset All: Press this button to clear all selected frequency spectra and input image.

Display Selection: Selection a frequency spectra in the Select Input Spectra then press this button to display the frequency spectrum image.

Display Combination: Displays the image created using the selected frequency spectra and input image.

Apply: Press this button to rebuild the image from its frequency components.

Done: Choose this button to close the Combine Spectrum dialog box.