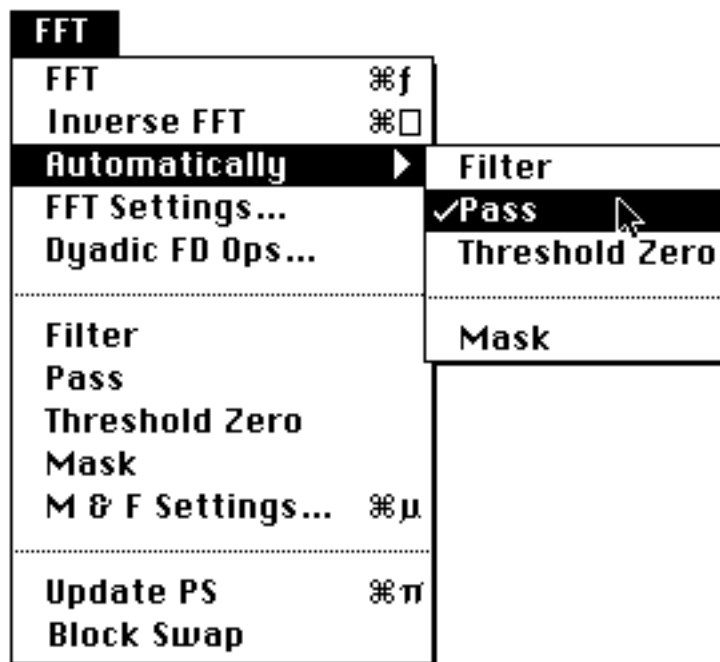


FFT EXTENSIONS TO *IMAGE* 1.25

Described herein are the extensions to *Image* 1.25 allowing frequency domain (power spectrum) display and editing of gray scale images. A description of each menu item and dialog box is followed below by a Quick Reference. If you have experience with earlier versions of this program or are simply ambitious, you may wish to skip ahead to the Quick Reference section which summarizes important key-combinations necessary to make full use of *Image*'s new capability.

• • •

The addition of FFT capability to *Image* has added one menu to the program:



FFT

The FFT menu item is enabled if the current selection is a valid FFT-able selection. Currently only square regions which are an integer power of 2 in size are FFT-able (the FFT Settings dialog box has a 'continuous selection size' setting, but that is not yet supported). Selections of this size are made most easily by holding down the command key while making a selection with the rectangle selection tool.

When the FFT is computed, the power spectrum is displayed in a new window (titled 'FFT #' where # increases serially from 1) with a slashed title bar. This title bar indicates that the window is a frequency domain window and that different editing rules apply. No cutting or pasting is allowed in a frequency domain window, but its contents may be copied and pasted into a space domain window. Frequency domain editing is limited to the operations Filter, Pass and Threshold Zero (described below), and selections may only

be made with the rectangle, oval and rounded rectangle tools. Furthermore, these tools behave differently than in the space domain as is described below under ‘Making Selections In the Frequency Domain’.

During the FFT computation, the portion of the image being transformed is displayed in the newly created frequency domain window. If automatic masking (described below) is enabled, a masked version of the original image will appear in the frequency domain window while the computation is taking place.

Once the FFT computation is complete and the power spectrum is displayed in the new frequency domain window, the results window shows the mouse location in a new format. If the mouse is over an active frequency domain window, its location is displayed in polar coordinates. The angle is expressed in degrees ($^{\circ}$), while the radius is expressed in either pixels per cycle (p/c) or, if a spatial scale calibration has been made using Set Scale..., in [units] per cycle (e.g. mm/c).

The contents of frequency domain windows can be saved as normal space domain images; the ability to save inverse-transformable power spectra is not yet supported.

The key-equivalent for FFT is Command-Option-F.

An Important Aside: Making Selections in the Frequency Domain

Selections in the frequency domain can be made with the rectangle, oval and rounded rectangle selection tools (the polygon and freehand selection tools will eventually be supported as well). These selection tools behave differently in the frequency domain because the frequency domain is radially symmetric and therefore redundant. When a selection is made with one of these tools, the effective selection is the union of the original selection with its reflection through the origin. The ‘marching ants’ marquee portrays this graphically as you make the selection.

It is often desirable to center the selection on the origin and this can be done by pressing the command key while making the selection. As in the space domain, pressing the shift key during selection constrains the selection to have equal width and height. Finally, if you want to make complex selections, the active selection, or ‘region of interest’, can be added to and subtracted from using the same technique as in the space domain: To subtract a region from the currently active region, hold down the option key while making a new selection; when you’re done it will be subtracted from the current region. Do the same while holding down the control key if you want to add to the current selection. It might seem complicated at first using e.g. option-command-shift to subtract a centered square region from the current region of interest, but after a while you’ll get the hang of it. All the key combinations are summarized in the Quick Reference below, so you might want to print that page out and keep it around while you’re learning.

Inverse FFT

The inverse FFT menu item is enabled when a frequency domain window is front-most. When the inverse FFT is computed, the inverse transform is displayed in a new space domain window with title ‘I FFT #’ where # increases serially from 1. During the inverse FFT computation, a copy of the power spectrum of the original transform is displayed in this window. If automatic Filtering, Passing or Threshold Zeroing (described

below) are enabled, the copy of the power spectrum displayed will reflect these operations.

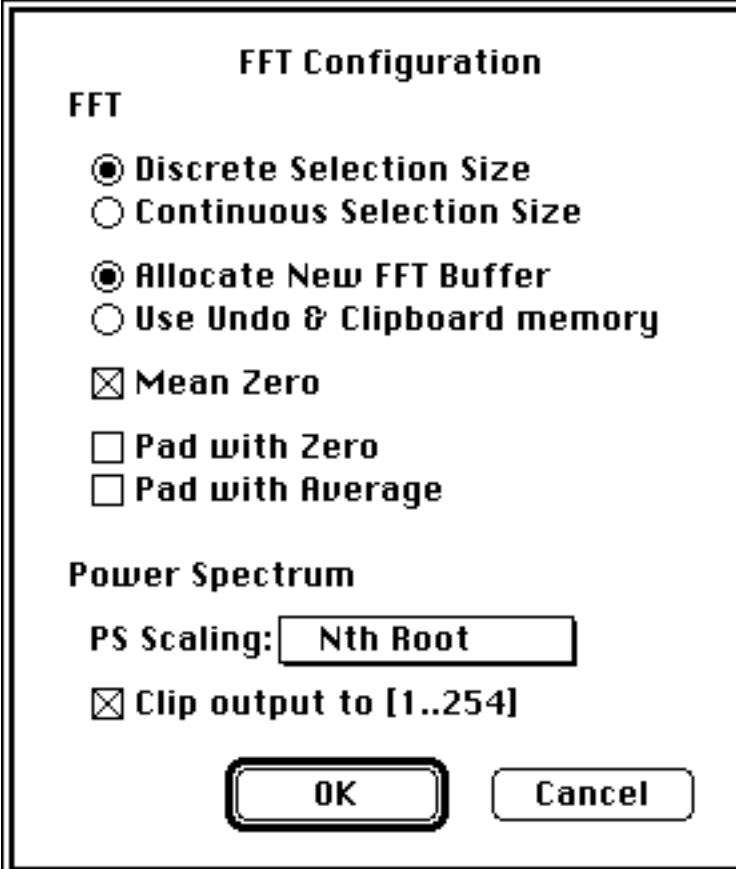
The key-equivalent for Inverse FFT is Command-Option-Shift F.

Automatically

The 'Automatically' sub-menu allows one to turn automatic Filtering, Passing, Threshold Zeroing and Masking on and off. The first three operations are used in Inverse FFTs, while Masking is used in forward FFTs. All four of these operations can either be done in place using the menu items in the main FFT menu or they can be done automatically to the copy generated during the transform process. By doing these operations automatically, the original copy of the image or power spectrum remains unchanged, eliminating the need to restore or re-compute it.

FFT Settings...

Choosing FFT Settings... brings up the following dialog box:



The image shows a dialog box titled "FFT Configuration". It contains two main sections: "FFT" and "Power Spectrum".

FFT

- Discrete Selection Size
- Continuous Selection Size
- Allocate New FFT Buffer
- Use Undo & Clipboard memory
- Mean Zero
- Pad with Zero
- Pad with Average

Power Spectrum

PS Scaling:

Clip output to [1..254]

At the bottom of the dialog box are two buttons: "OK" and "Cancel".

Discrete & Continuous Selection Size

The first radio button pair allows one to choose between discrete and continuous selection sizes. Currently, only discrete selection sizes are supported, but a new version due out 'any time now' will allow square selections of any size to be transformed.

Allocate New FFT Buffer / Use Undo & Clipboard memory

The second radio button pair is only to be used by those with severe memory constraints (those with only 2 Megabytes of memory, for example). Normally, *Image* will try to allocate a new power spectrum buffer for every new frequency domain window it creates, but if you choose 'Use Undo & Clipboard memory', the power spectrum will be computed in the memory normally reserved for the Clipboard and Undo buffer. This allows those with only 2M to compute larger transforms, but the power spectra generated in this way cannot be Inverse transformed, filtered or updated (re-computed).

Mean Zero

Mean Zero before FFT, when enabled, will cause the mean of the image being transformed to be computed and subtracted from the copy of the image being transformed. This effectively removes the DC component (the central pixel) of the transform and, because this component can otherwise dominate the power spectrum, improve the overall scaling.

Pad with Zero / Pad with Average

When computing the convolution or correlation of two images by performing dyadic frequency domain operations (described below), one may choose to pad the data to be transformed. When this option is enabled, the transform window will be twice the size (four times the area) of the original transform selection; during the transform computation, the selection will appear in the upper left quadrant of the new frequency domain window. If zero padding is enabled, the other three quadrants will be filled with zeros, while if average padding is enabled, they will be filled with the average value of the selection. Zero padding is necessary if the convolution of two images is to be subsequently sought; without zero padding, the cyclic, or circular, convolution will result. Likewise, padding is necessary for the correlation and autocorrelation operations, however, in this case more detail is brought out if the data is padded with its average value.

Power Spectrum Scaling

Power spectra are normally log-scaled, but this scaling sometimes brings out more spectral detail at low spatial frequencies than is desired. For this reason, Nth Root scaling, for N in [2..9] and Linear scaling are also supported. After the scaling has been changed using this pop-up menu, the power spectrum displayed in the front-most frequency domain window can be updated to reflect this scaling with the Update PS menu item.

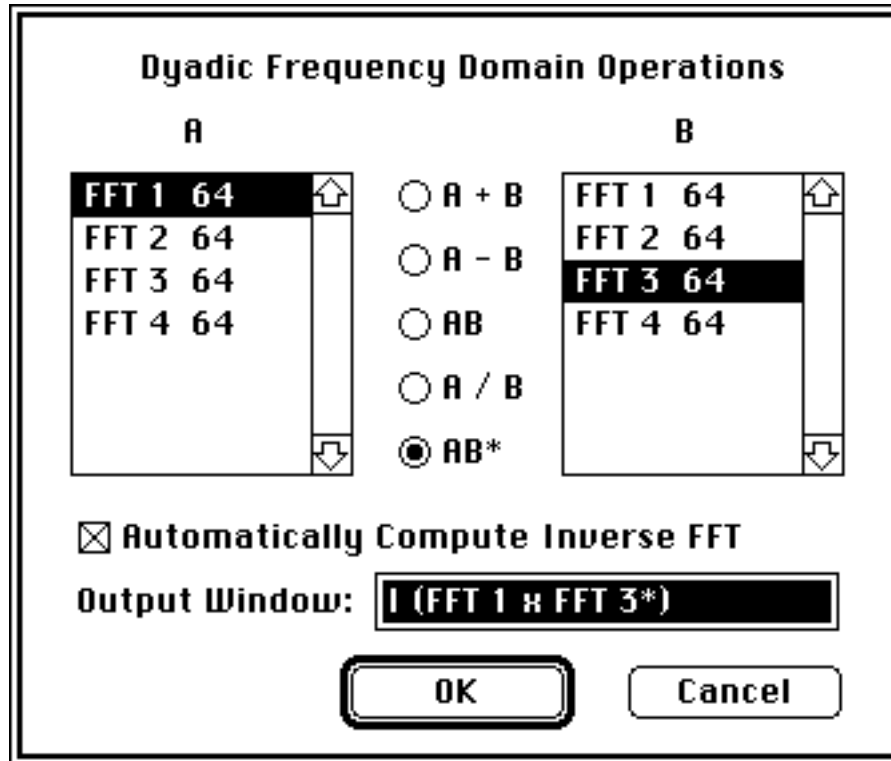
Clip Output to [1..254]

When the power spectrum is computed, it normally is scaled to fill the output range [0..255]. Thresholding, however, allows the user to threshold only pixels in the range [1..254]. Therefore, if the user wants to use Threshold Zero (described below), he will be able to threshold any portion of the power spectrum when 'Clip Output to [1..254]' was enabled prior to the last power spectrum update.

All FFT Settings can be recorded permanently by choosing the 'Record Preferences' menu item in the file menu.

Dyadic FD Ops...

New in this version of image is the support of Dyadic Frequency Domain Operations. With them, *Image* is now capable computing convolutions, deconvolutions, correlations and autocorrelations.



When a frequency domain window is open, this menu item is enabled and choosing it will bring up the dialog box above. The titles of the currently open frequency domain windows appear in two lists, followed by their width (=height) in pixels (Dyadic operations are allowed only on images of equal dimensions). Between these two lists, five radio buttons allow one to select the frequency domain operation desired. In addition to the four common operations, conjugate multiplication is also supported for correlation computations.

The result of the dyadic frequency domain operation appears in a new window, the title of which may be entered into the Output Window field. A default title generated from the selections is provided until one types or clicks in this field.

Automatically Compute Inverse FFT

Often one is more interested in the inverse transform of a dyadic frequency domain operation than in the transform itself. When this is the case, the inverse transform can be automatically computed from the result of the frequency domain operation by enabling this item. The new window created by the operation is then a space domain window which, during the inverse transform computation, displays the power spectrum of the dyadic operation's result. Not only does this proceed automatically, but it spares the memory otherwise used by the intermediate frequency domain window and its buffer.

Addition and Subtraction

The addition and subtraction operations in the frequency domain are equivalent to those in the space domain. They therefore do not hold much interest except to verify the linearity of the transform process. Still, they are included here for completeness.

Multiplication (Convolution)

Multiplication in the frequency domain translates to convolution in the space domain; two images can be convolved by transforming each of them, multiplying their transforms and inverse transforming the result. While the same computation can be done directly in the space domain using *Image's* convolution feature, the space domain kernels must be typed in as text files and are limited in size to 63x63. Using the frequency domain operation removes the size restriction and allows one to draw the convolution kernels using *Image's* drawing tools. While the set of practical applications that demand this new flexibility may be small, the process is a great educational tool.

If data padding (see FFT Settings... above) is disabled, the inverse transform of the product of two transforms will produce the *cyclic* convolution of the two original images. This arises from the sampling-induced periodicity of the image plane. Since the Fourier transform of a sampled image is equivalent to the transform of an infinite periodic replication of the image, the shifting inherent in the convolution operation results in the image overlapping its periodic extension. This creates a circular convolution and the wrap around affect can be easily seen by experimenting with this operation.

In order to produce normal, not circular, convolution, one can avoid the overlap artifact by padding the image to be convolved with zeros. Unlike the autocorrelation function, however, convolution is sensitive to the placement of the extra zeros. By placing them to the lower right (see description under FFT Settings...), no artificial translation is induced.

Division (Deconvolution)

Division in the frequency domain corresponds to deconvolution in the space domain and serves as the inverse of the multiplication (convolution) operation. For example, if one first multiplies the transforms FFT1 and FFT2 together, leaving the result in a new frequency domain window, FFT1 x FFT2, and then divides this spectrum by either FFT1 or FFT2, the complementary FFT results. Unfortunately, however, the nature of the division operation creates output with a wide dynamic range (not to mention some singularities) that cannot be fully accommodated by the 16 bit integer arithmetic used

throughout these routines. Consequently, often large portions of the quotient transform are lost and the resulting inverse transform suffers accordingly. While this limits the practical use of this deconvolution operation, it is nonetheless impressive to see how well the original object can be recovered from its image when one knows the point spread function with which it was convolved.

Conjugate Multiplication (Correlation)

Conjugate multiplication in the frequency domain translates to correlation in the space domain. The wide use of correlation and autocorrelation operations in imaging science makes this perhaps the most important dyadic frequency domain operation supported by *Image*.

The best way to compute the autocorrelation of an image is to select 'Pad with Average' in FFT Settings before Fourier transforming. This accomplishes two things. First, data padding avoids the characteristic wrap around affect to which correlation, like convolution, is subject. Second, by padding the data with its average value instead of with zeros, more detail is brought out in the autocorrelation function. By padding with zeros, the extended data looks like a modulated rectangle function. In the autocorrelation operation, the signature of the rect function dominates over that of the modulation (your data). By padding with the average value of the data, this rect function autocorrelation signature is removed and the autocorrelation features of the data stand out.

Finally, the result of the autocorrelation operation produces an image that is inherently 'block swapped' or in its 'wrap around' state. If you prefer the more common 'un-swapped' representation of the data, it can be rearranged using the Block Swap menu item described below.

Like the FFT Settings, the preferred Dyadic Operation settings can be recorded permanently using 'Record Preferences.'

Filter & Pass

Once a region is selected in the frequency domain, these menu items become enabled and the selection can be either Filtered or Passed. 'Filter' removes or suppresses the spatial frequencies in the current selection while 'Pass' performs the complementary operation. The details of these operations are described under 'Mask & Filter Settings...' below.

Threshold Zero

Another means of frequency domain editing is provided by the Threshold Zero operation. When thresholding in the frequency domain, this menu item is enabled and allows one to select, by using the LUT tool to manipulate the threshold 'density slice', which part of the power spectrum to filter out. Choosing Threshold Zero will zero out all parts of the power spectrum within the thresholding interval (colored red by default). A typical application of this would be to remove all components of the power spectrum except the most dominant peaks.

In order to make the entire power spectrum fall within the threshold-able interval [1..254], the 'Clip Output to [1..254]' option in the FFT Settings dialog box should be enabled.

Like the Filter and Pass operations, Threshold Zeroing can be performed automatically on inverse transformation to avoid permanent (non-undoable) alteration of the original power spectrum.

Mask

Masking is the space domain analog of Passing: when Masked, an image selection is given a smooth transition to zero at its edges. In contrast to Passing, however, masking is less flexible because it is only enabled for rectangular selections in the space domain and regions inside the mask transition region are passed completely while those outside are zeroed completely.

Why use masking? The discrete two dimensional Fourier transform of an image is effectively the transform of the periodic extension of the image. Tiling the image in the plane can give rise to sharp edges where the right (top) edge of the image meets the left (bottom) edge of its periodic extension. This in turn increases the image's power spectral density in the horizontal and vertical directions; in extreme cases the power spectrum of an unmasked image can be dominated by this cross shaped signature. By masking, however, one guarantees that an image's periodic extension has no sharp edge discontinuities, effectively removing the cross-shaped artifact from the power spectrum. Be aware, however: masking (multiplying) in the space domain is equivalent to convolving in the frequency domain; by masking an image, you are convolving its transform with the transform of the mask, which can also have undesirable effects. In general, masking is used to remove the cross from a power spectrum; if the power spectrum is to be modified and inverse transformed, masking is not used.

Like Filtering, Passing and Threshold Zeroing, Masking can be done automatically during the forward FFT, eliminating the need modify the original image. Unlike the other operations, however, masking *is* undoable.

Mask & Filter Settings...

Choosing 'M & F Settings...' brings up the following dialog box:

Mask and Filter Configuration

Transition Type

Gaussian **Linear**
 Cosine **Parabolic**

Transition Size & Range

FFT

Percent Width: %

Inverse FFT


Pixel Width: pixels

Minimum Level: %

Maximum Level: %

Transition Type

Both the space domain operation Mask as well as the frequency domain operations Filter and Pass support the four different transition types shown below:



These are the transitions that Masks or Filters take from full scale to zero. Both the space and frequency domain images are altered by such Masks and Filters using a multiply operation; each pixel in the transition region is reduced in proportion to the transition function at that point.

Percent Width


For Masking, the width of the transition region is specified as a proportion of the minimum dimension of the selection rectangle, denoted Percent Width. A Percent Width of 100% will set the width of the transition region to half of the minimum dimension of the selection rectangle and the entire region inside the rectangle will be altered by the mask. Any percent width in [0..100%] can be entered in the M & F Settings dialog box.

Pixel Width, Minimum & Maximum Level

Since the Filter and Pass operations can be applied to much more complicated regions in the frequency domain, the transition width is specified in pixels instead. Furthermore, since it may be desirable to suppress, instead of completely eliminate, particular spatial frequencies, you can specify a maximum and minimum level of the transition function.

For example, if the user wants to suppress all spatial frequencies outside the selection by 50% while completely passing all spatial frequencies within the selection, he can set the transition pixel width to 0, the minimum level to 50%, the maximum level to 100% and use the Pass operation.

A graphical representation of a Pass and Filter mask cross section is shown below:



The key equivalent for Mask and Filter Settings is Command-Option-M.

All Mask & Filter Settings can be recorded permanently by choosing the 'Record Preferences' menu item in the file menu.

Update PS

At any time a frequency domain window is front-most, it's power spectrum can be re-computed using Update PS. If, however, the power spectrum was computed using the 'Use Undo & Clipboard memory' option enabled, the power spectrum cannot be re-computed because the memory used in the original power spectrum computation has likely been corrupted by Undo and Clipboard operations.

Typically, Update PS is used to re-compute the power spectrum after the default scaling or the 'Clip Output to [1..254]' variables have changed.

The key equivalent for Update Power Spectrum is Command-Option-P.

Block Swap

Block swapping is enabled for space domain windows whose width and height are equal and an integer power of two. This happens to include all windows produced by inverse Fourier transformation. Why Block Swap? What *is* Block Swap?

Block Swapping (I'm sure it has a more official name) is performed automatically every time a power spectrum is computed. If you think of an image's origin as being at its geometric center, block swapping simply swaps quadrant 1 of the image with quadrant 3; likewise for quadrants 2 and 4. In the normal, un-swapped, state, the power spectrum's central peak is distributed among the four corners of the image matrix. While this is the format used in all computations, it doesn't correspond to nature's FFT analog, the diffraction pattern, so most spectra are block swapped before display. This menu item allows block swapping for space domain windows only, however, because it is useful in two other contexts.

First, convolution kernels are typically arranged in a block swapped or 'wrap around' state. This puts the maximum of the kernel at matrix index 0,0 (upper left corner) so that convolving with the kernel will not result in a translation. Second, the cross and auto-correlation operations result in inherently un-swapped space domain images, yet like power spectra, correlation functions are often displayed in block swapped format. The Block Swap menu item lets you choose the format you desire.

FFT QUICK REFERENCE

Space Domain

- With the rectangle selection tool, hold down the command key to make a square selection an integer power of 2 in size.
- To make the largest FFT selection allowable in the current window, hold down the command key while double clicking the rectangle selection tool.
- Masking vignettes a rectangular image selection so that it has a smooth transition to zero at its edges.
- Automatic Masking, when enabled, will cause the copy of the image selection used in computing the FFT to be masked, leaving the original selection unchanged.

Frequency Domain

- Only the rectangle, oval and rounded rectangle selection tools allow selections to be made (Some other tools work too, like the cross-section plot tool).
- Holding down the command key will center these selections on the origin.
- Holding down the shift key will constrain these selections to have equal height and width (whether or not they are centered at the origin).
- Holding down the option key before making a new selection will cause it to be subtracted from the current selection
- Holding down the control key before making a new selection will cause it to be added to the current selection.
- Filtering removes or suppresses the spatial frequencies in the current selection.
- Passing removes or suppresses the spatial frequencies outside the current selection.
- Threshold Zeroing zeroes out all spatial frequencies within the current threshold density slice.
- Automatic Filtering, Passing and Threshold Zeroing performs these operations on the copy of the power spectrum used in computing the Inverse FFT without altering the original power spectrum.
- Update Power Spectrum re-computes the power spectrum to reflect the current scaling default (chosen in FFT Settings).

Key Equivalents

- | | |
|---------------------|------------------------|
| • FFT | Command-Option-F |
| • Inverse FFT | Command-Shift-Option-F |
| • M & F Settings... | Command-Option-M |
| • Update PS | Command-Option-P |