

Advanced Image Enhancement Part 3

From Chapters 17 and 18

We are about to explore a very powerful method of image enhancement. Until now, we have treated images as an array of numbers, a “spatial” function, and we have used various techniques to alter those numbers to enhance whatever portion of the image we are interested in. But there is another way to think about images that is not always intuitive. We can think of images as the superposition of a large number of frequencies. Every image can be decomposed into a unique set, or spectrum, of spatial frequencies. The mathematical technique that actually does this decomposition is called “the Fourier Transform”. The Fourier transform transforms images into an array of frequencies in frequency space. We can alter the array of frequencies, and then transform the altered frequency array back into image space. Depending on how we alter the frequencies, we can use Fourier transforms to enhance image detail, to alter the brightness distribution of an image, to remove noise, and accomplish many other tasks.

The word “frequency” implies a signal that varies with time, such as a sound wave or a radio wave. If you take one of our images and map the intensity of light across it, you will see that the intensity of light varies in space, just as a sound wave varies with time. Transitions that happen rapidly (like sharp edges) contain high frequencies, and transitions that happen slowly (like large scale galactic structure) contain low frequencies. The measurement units for a sound wave are cycles per second, and the measurement unit for a spatial frequency is cycles per pixel.

Think of an image of a white picket fence. You can take a ruler and count how many posts per given interval the fence has. This is an example of a spatial frequency. In a Fourier transform, simple, single spatial frequencies like our picket fence will appear as strips or bars. Now think of an image of a tiger with its complex strips. Fourier transforms of irregular objects (like galaxies and nebula) can be very complex.

The basic idea behind Fourier analysis is that it is possible to break any periodic function into a series of simple sine waves. The complementary idea is that it is possible to recreate any periodic function by adding together enough sine waves. This is the basis of the Fourier Series and the Fourier Integral (see Figure 17.3 in the book). The Fourier series consists of a series of sine and cosine functions (with their assorted coefficients) added together. The Fourier Series can be used to express any periodic mathematical function as a sum of any number of sine and cosine functions. In the real world, there are no strictly periodic functions. The Fourier Integral is a reformulation of the Fourier series using continuous functions rather than a discrete series of sine and cosine functions with their associated coefficients. One basis of the Fourier Integral is the fact that sine and cosine functions can be rewritten in exponential form (including both real and imaginary components, See section 17.2.2).

It is highly recommended that you read Chapter 17 before proceeding with this exercise. Frequency space can be a difficult concept to grasp, but it is not necessary to have a complete understanding of the mathematics. When we applied the Crispening Filter to the moon image in a previous exercise, we were enhancing the high frequencies (sharp transitions) to bring out detail. The Crispening filter is rather limited in what it can do and tends to be too rough. FFTs allow more flexibility and control. As you work through this exercise, you will develop a feel for Fourier transforms and how they work.

The type of Fourier transform most often used is the Fast Fourier Transform. The Fast Fourier Transform (FFT) is a computationally accelerated technique for determining the frequency spectrum of a signal. The basic steps in filtering an image using FFT techniques are:

- 1) Convert the image to suitable dimensions (the number of pixels in the x and y axis must be a factor of 2)
- 2) Perform a FFT on the image to transform it to frequency space.
- 3) Create a filter mask that does what you want it to do.
- 4) Multiply the image by the filter mask
- 5) Transform the result back to the spatial domain.

Step 1: The FFT technique does have an inherent limitation. Any image that you apply it to must have an equal number of pixels on the x and y axes. The dimensions must also be a factor of 2 (2x2, 4x4, 8x8, 16x16, 256x256, 512x512, 1024x1024, etc.). Load the image “triangle.fts”. This is a 128x128 image of a white triangle on a black background.

Step 2: Perform an FFT on this image. Click Enhance → Manual FFT → Forward FFT. An image (a two dimensional Fourier transform or frequency plot) appears that consists of various diagonal lines. Each line represents a spatial frequency found in the triangle image. The brightness of each line represents the amplitude of the spatial frequency. The frequency plot has a frequency of 0 at its center, and positive frequencies to the right and negative frequencies to the left. This is the “real” axis. The “imaginary” axis runs vertically through the center of the plot. Low frequencies are clustered around the center of the plot. Frequency increases as you move outward. So information about large scale structure is near the center of the plot, and information about details like edges can be found nearer the edges. Include a screen shot of your image in your report.

Step 3: Click Enhance → Manual FFT → Butterworth Filter Mask. The Butterworth Filter Generator window will pop up. Don’t be intimidated by the name. This is simply a method of generating a filter mask that will be applied to the image and used to “mask” some of the frequencies present in the image. The filter mask will be multiplied by the FFT image to change the distribution of frequencies. If the filter mask has small values around its center, it will reduce the large scale structure and enhance details. If the mask is dark around the edges, it will reduce details and enhance the large scale structure. With an FFT, you can also generate intermediate masks that remove low and high frequency and leave those in between, and vice versa.

The Butterworth Filter Generator window has a display showing a plot of the frequency characteristics of the filter it will generate, along with controls to change the shape of the filter. The first thing to check is that the filter size is the same as your image. This image is 128x128, so change the filter size to match.

Click the Low Pass, High Pass, Band Pass, and Band Stop buttons to get an idea of what the various filter shapes look like. For example, the low pass filter has high amplitude at low frequencies (the left part of the graph), and very small amplitude at high frequencies (the right part of the graph). Therefore, only low frequency information will be allowed, high frequencies will be suppressed.

- 1) What types of structures do you think will be preserved by the low pass filter?
- 2) Record the shapes of the remaining filters.
- 3) What types of structures do you think will be preserved by each of the filters recorded in question 2?

To a certain extent, you can also control the shape of each filter. For example, the Low Pass filter has a Low-pass radius control. Slide the control around.

- 1) What happens to the filter plot?
- 2) What do you think happens to the frequencies suppressed by the filter?

The filter order gives you a rough idea of the number of sine waves used to generate the filter.

- 1) What happens to the filter plot when you increase the filter order?
- 2) Compare this with figure 17.3 in the book.

Set the filter mask for Low Pass filter with a Filter Size of 128x128. Set the Filter Order to 1 and a Low Pass Radius of 16. Click “Generate Filter” and a filter mask will be generated. This is the image that will be multiplied by the FFT image to create a new image of the triangle. The Filter is bright in the center, and gradually fades towards the edges. Remembering that low frequencies are near the center of the FFT image, this mask will allow low frequencies, but will suppress the high frequency information.

Step 4: Let’s apply the mask to the FFT. Click Enhance→Manual FFT→Mask Fourier Transform. The Mask window will appear. The window contains 4 controls which allow you to select the Fourier Image to use, the Mask to use, the Base Value and the Contrast. When the base value is 0, frequencies blocked by the filter remain blocked; when the base value is 1, all frequencies present in the original image are passed through the filter. The Contrast multiplies the transmission of the filter by a contrast factor. This allows the filter to suppress frequencies and enhance the frequencies of interest. Leave the Base Value set to 0, and the Contrast to 1.

Click “Apply”. A new FFT image appears, one that is bright in the center and darker towards the edges. This is because the original FFT was multiplied by the mask. Include a screen shot of this image in your report.

Step 5: Click on the masked version of the FFT that you just produced to make sure it is active. Now click Enhance→Manual FFT→Inverse FFT. A new image of the triangle will appear.

- 1) Compare this new triangle with the original
- 2) Explain any differences in terms of the FFT and the mask.
- 3) Include a screen shot of this image in your report.

Step 6: Repeat the above steps using the high pass filter, band pass, and band stop. Describe the results of each.

Step 7: We just went through the complex steps to perform FFT filtering manually. Fortunately, AIP4WIN has an automated tool that is much easier to use. Click Enhance→FFT Filter. The Butterworth FFT filter window will appear. This window is a combination of the windows we were working with before. The auto FFT will automatically square up your image, and change the dimensions to a power of 2 if necessary. It will then run the steps we went through above and you will see the individual windows appear as the steps are performed. Select the “triangle.fts” image, the low pass filter, a filter order of 4, a Low-Pass radius of 4, base value of 0 and contrast of 1. Click Apply.

- 1) What do you see happening?
- 2) What does the final image look like?
- 3) Include a screen shot of your image in your report.

Step 8: Now we need to try a real image. Open an averaged stacked image of an object. Nebulae or galaxies usually work better than images of stars. The File “m27stack.fts” is available for use in the General observations section of the class webpage. This is an image of the Dumbbell Nebula we took in class. Click Enhance→FFT Filter. Select the low pass filter, a filter order of 1, a low-pass radius of 32, a base value of 0, and a contrast of 1. Click Apply.

- 1) Describe the resulting image.
- 2) What happens if you rerun the FFT with a high pass filter?
- 3) Include screen shots of your low pass and high pass images.

Step 9: Open an image of Mars (available in the General observations section of the class web page). Click Enhance→FFT Filter. Select the high pass filter, a filter order of 1, a high-pass radius of 8, a base value of 0, and a contrast of 1. Click Apply.

- 1) Describe the resulting image.
- 2) What happens if you rerun the FFT with a low pass filter?
- 3) Which filter do you think is best for bringing out surface details on Mars?
- 4) Include a screen shot of your best Mars image.

Close all your images when you are done with this section.

Wavelet Spatial Filtering

Wavelet filtering is a new image processing technique. It is covered in detail in Chapter 18 of the book. Wavelets are a hybrid between the spatial methods of operators and kernels and the frequency-based methods of the Fourier Transform. The formal definition of a wavelet is a function with positive and negative values locally, zero everywhere else, and an integral of 0. There is a graph of a wavelet on page 477 of the book that gives you a better visual idea of what a wavelet is. It resembles a ripple in a pond, which is where the name comes from. Wavelet analysis uses functions like this to split an image into different wavelet scales or levels, each corresponding to a specific band of spatial frequencies. Wavelets operate locally on an image (this is the spatial component). Wavelets also let you process images in frequency space similar to the Fourier Transform.

For this exercise, make sure that Auto Min/Max is turned on in the Image Display Control.

Step 1: Load the image “m27stack.fts”. This is an image of the Dumbbell Nebula taken with the Alta CCD. It is available in the General Observations section of the web page.

Step 2: Click on Enhance→Wavelet Spatial Filter. The Wavelet Spatial Filter tool will open. There are four panels with the following controls:

- 1) Maximum Wavelet Scale. This sets the limit of the largest pixel scale to be processed.
- 2) Wavelet Profile sliders. Set the amount of enhancement for each pixel scale.
- 3) Add Constant. Allows you to add a constant to the entire image to help keep the values from becoming negative. This is a similar idea to the bias level of the CCD itself.
- 4) No Dark Rings around Stars. Self explanatory.
- 5) Select/Unselect All. Controls whether the entire range of pixel scales is processed or not
- 6) Settings. Allows you to save and recall your favorite filter settings.
- 7) Apply Filter button. This is the one you click when you have everything set up.

Step 3: There is a predefined filter for this exercise. Click the Recall button to load the filter. The filter is in the Filters directory and is called “Sharpen.wlt”. Select it.

Step 4: Click the Apply Filter button. A status bar will appear to apprise you of the progress of the calculations. This is a computationally intense technique, so if your computer is a little older, it can take several minutes.

- 1) Describe the resulting image.
- 2) Include a screen shot of your image in your report.

Step 5: Return to the original image of M27. Click Enhance→ Unsharp Mask. Set the radius to 2.1 and leave the other values as default. Click Apply. Include a screen shot of your image in your report

Step 6: Click Multi-Image→Image Math. Subtract the Unsharp Mask image from the Wavelet Image.

- 1) What is left?
- 2) Describe the differences between the Unsharp Mask and the Wavelet images
- 3) Include a screen shot of this image in your report.

Step 7: Return to the original M27 image. Click Enhance→Wavelet Spatial Filter. This time, load the filter “DeepSkyHump.wlt”. This filter has the maximum pixel size set to 6. Apply the filter. Include a screen shot of your image in your report.

Step 8: Click Multi-Image→Image Math. Subtract the wavelet image using Sharpen.wlt from the wavelet image using custom1.wlt.

- 1) Describe the resulting image.
- 2) What effect did changing the maximum pixel size have?
- 3) Include a screen shot of your image in your report.

Close all your images when you are done with this section.

Step 9: Experiment using your own Wavelet Filter settings. Record the settings and results you find for three different examples.

Morphological Processing

Morphological processes are not the usual techniques of astronomical image processing, but some do have their uses, especially in the area of feature detection. Here we will work with the Contour Map.

Step 1: Load an averaged stacked image of an object we have observed. Nebulae and galaxies usually work better than images of star clusters. Click Enhance→Morphological Operations→Contour mapping. The Contour Map Tool will appear. Choose “Outline Regions of Equal Area”, and set the number of lines to 15.

Step 2: Click Apply and a window will appear showing the progress of the calculations. A new image will appear containing the contour map of your object.

Step 3: Use the Image Display Control to remove the sky background. You do this by increasing the Black value until the background darkens.

- 1) What Black value did you decide on?
- 2) Include a screen shot of your image in your report.

Step 4: Return to the original image. Click Enhance→Blur . Use the Blur Tool to smooth the image. Select a Radius of 3 pixels and click Apply. A fuzzy version of M51 will appear.

Step 5: Rerun the contour map using this blurred image.

- 1) Describe the contour image.
- 2) What are the differences between this contour image and your previous one?
- 3) Include a screen shot of this image in your report.

Step 6: Return to the original image. Click on Transform→Resample. Resample the image to 200%. Rerun the contour map using this image. Compare your results with the previous two contour maps. Include screen shots of your images.

Step 7: Open a stacked image of a globular cluster such as M13 or M5”. Repeat steps 1-3 with this image.

- 1) Describe the resulting contour map.
- 2) Resample the image to 200% and rerun the contour map. Change the Black value to eliminate the gray background. Describe the resulting contour map.
- 3) Do you think contour maps are helpful in detecting faint objects? Explain.
- 4) Include screen shots of your images in your report.