Continuing Education

Filtering in Frequency Space

James R. Galt, H. Lee Hise, Ernest V. Garcia, and David J. Nowak[†]

Emory University School of Medicine, Atlanta, Georgia; and [†]General Electric Company Medical Systems Group, Milwaukee, Wisconsin

This is the third in a series of four continuing education articles on computers in nuclear medicine. After studying this article, the reader should be able to: 1) understand filtering concepts and frequency space; and 2) be able to determine the appropriate filter selection and its application for a given nuclear medicine procedure.

Filtering is used extensively on nuclear medicine images to reduce statistical noise, enhance edges for edge detection, and in the reconstruction of tomographic images. As explained in a previous continuing education article (*1*), filtering can be performed in either the spatial domain or in frequency space. While that article dealt primarily with operations in the spatial domain, the purpose of this manuscript is to develop the reader's intuition as to how to use frequency space without understanding complicated mathematical formulas. By exercising this intuition, the reader will be able to select filters for given applications, determine proper cutoff frequencies, and even design filters to meet the requirements of new applications.

THE FREQUENCY DOMAIN

Frequency Space

Frequency domain methods of image enhancement deal with the spatial frequencies (cycles/centimeter or cycles/pixel) which make up the image. An image can be decomposed into a summation of different spatial frequencies. Spatial frequencies relate to the rate of change in intensity (counts) with distance. The high frequency components define edges, areas where there is a rapid change in intensity from bright to dark. Regions in an image where the changes are more gradual are primarily caused by the lower frequency components. An example of such a region is the central area of the lobe of the liver from a high-count liver scan. While the concept of frequency space may appear difficult to grasp at first, when intuition is developed the principles can be more straightforward to apply than those used in the spatial domain (*2*).

In order to convert an image into its frequency components it is necessary to use an image transform, such as a Fourier transform, to transfer the spatial image into frequency space. The Fourier transform yields the frequencies that are present

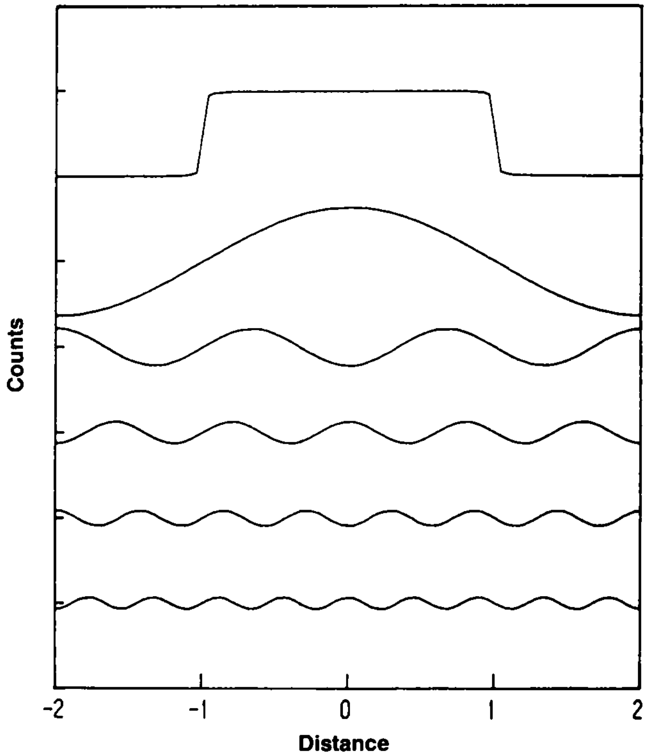

FIG. 1. Example of how the Fourier transform can decompose a square curve into sinusoidal frequency components. The top curve is a simulation of the count profile of an ideal flood-source image. The lower curves show the five lowest frequency components that make up the image. The summation of these five components, along with higher frequency (and lower amplitude) components gives the top curve.

For reprints contact: Ernest V. Garcia, Emory University School of Medicine, 1364 Clifton Road, N.E., Atlanta, GA 30322.

in the image and the amplitude of each frequency. An inverse Fourier transform transfers the frequency image back into the spatial domain, the conventional format for representing images. The method used by the computer (often with the aid of an array processor) to implement the Fourier transform is called a Fast Fourier Transform (FFT). The FFT is a computer algorithm which has been optimized for speed.

An example of how an image can be decomposed into its various frequency components can be seen in figure 1. The top curve is a simulation of a count profile through the center of the image of a radioactive flood source. The simulation assumes that the flood was imaged for a large number of counts (assuming no statistical fluctuations) using a high resolution collimator (the pixel count drops off very rapidly at the edge of the source). The lower sinusoidal curves represent the five lowest frequencies that make up the square-count profile. If these frequencies along with several higher frequencies are summed, the result is the top curve (note that the higher the frequency the smaller the amplitude).

Filtering

Filtering in frequency space is accomplished by removing, or altering the magnitude of selected frequency components. Each specific frequency is multiplied by a factor assigned by the filter. If a particular frequency needs to be totally suppressed, the factor assigned to that frequency is set to zero. Filters that place emphasis on the low frequency components while reducing the high frequency components are called low-pass filters. If the opposite effect is desired, a high-pass filter is applied which reduces the low frequencies. A band-pass filter reduces both low and high frequencies allowing only a band of frequencies in between to remain.

Figure 2 demonstrates how the higher and lower frequencies affect the image. The square curve again represents a profile of a flood source. The rounded curve is the lowest frequency component of the square profile. Thus, if a low-pass filter was applied to the flood image which only allowed this low frequency to go through, the result would appear as a very smooth image as evidenced by the disappearance of the sharp square edges. The remaining curve is the original count profile with the lowest frequency component subtracted out (a high-pass filter has been applied). It can be seen that removing the low frequency components emphasizes the change in intensity in the resulting profile at the location where the sharp square edges were defined in the original count profile.

FILTERS

Figures 3 through 7 illustrate the results of a computer program developed for the purpose of demonstrating how filters may be used to modify images. Each of the figures, which have the same format, is divided into four sections. Two 64 $\times$ 64 pixel images take up the top half of each figure. The image on the left is the input (unfiltered) image and the image on the right is the filtered image. The lower left quadrant of the figure shows the filter used to modify the image. The frequency spectra of the input and filtered images are shown in the lower right quadrant of the figure. The frequencies that make up the input image are represented in graphical form by the input spectrum (lower right) where the X axis represents the spatial frequency in cycles/pixel and the Y axis represents the magnitude of that frequency (associated with image contrast or count difference between black and white). The frequencies of the filtered image are represented by the filtered spectrum after the input spectrum has been multiplied by the filter function on the lower left. Note that both the filter function and the frequency spectrum are two-dimensional functions. The one-dimensional plots shown here are simplifications which permit an easier description of the process.

The upper image is made up of four spatial frequencies that

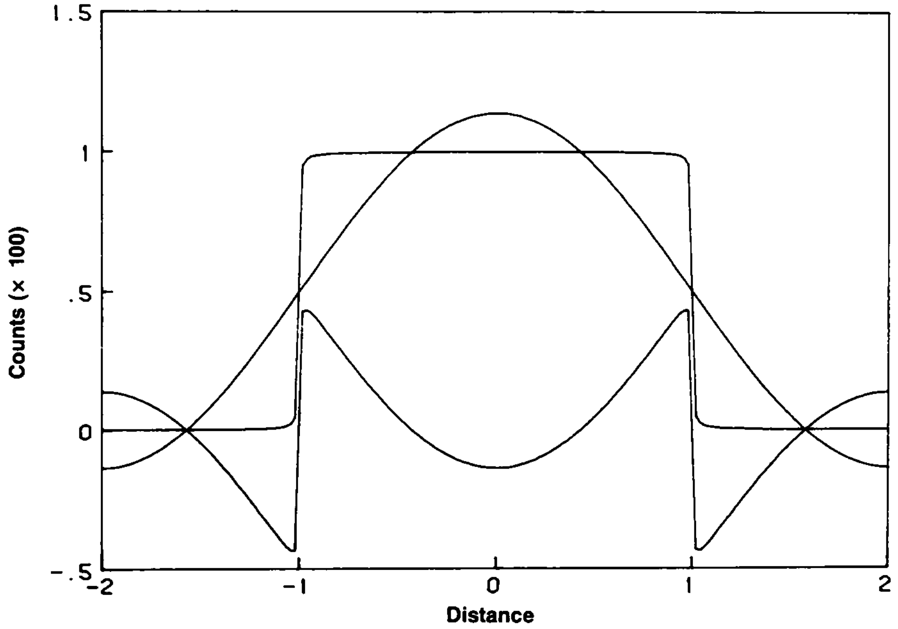

FIG. 2. High and low frequency contributions to the ideal flood source count profile shown in Figure 1. The smooth curve is the lowest frequency component in the profile. The lower curve demonstrates the application of a high-pass filter. The lowest frequency component has been removed from the profile leaving only the high frequency components. The rapid change in amplitude at the edges of the flood source in this curve show how high-pass filters may be used to emphasize edges in images.

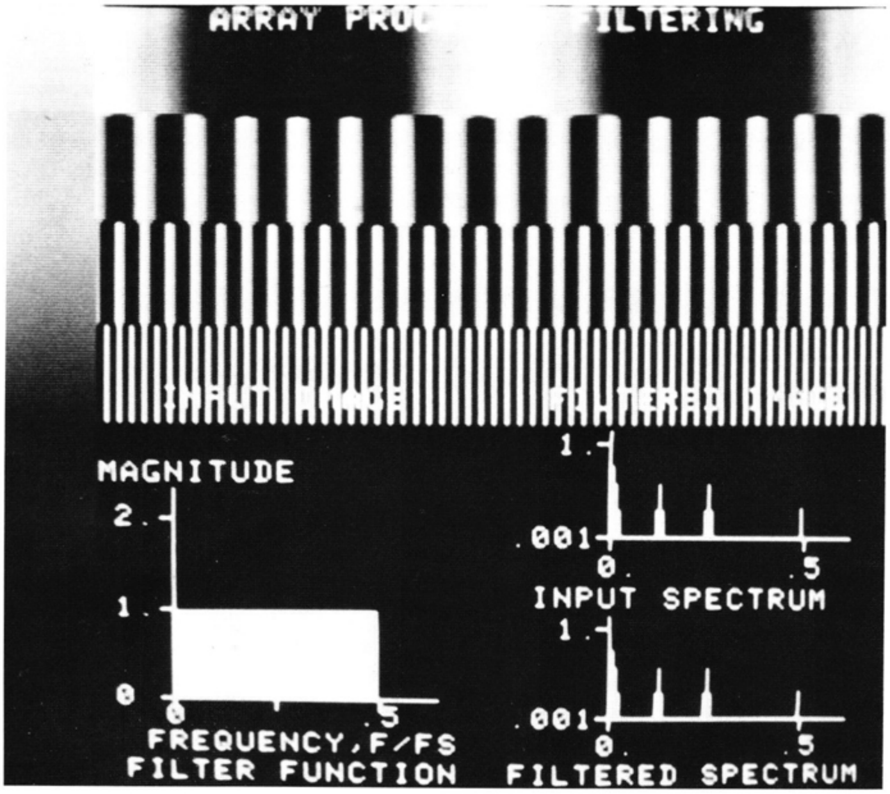

FIG. 3. Demonstration of Fourier filtering of an ideal image with no statistical noise and where the spatial frequencies are separated and discrete. The input image is made up of four components: the highest (64 × 64) frequency appears in the lower row corresponding to 32 bars per 64 pixels (0.5 cycles/pixel). The next two rows above it correspond to 16 bars per 64 pixels (0.25 cycles/pixel) and 8 bars per 64 pixels (0.125 cycles/pixel). The top row represents one complete cycle over the width of the image or one cycle for 64 pixels. The discrete frequencies plotted in the input spectrum and which appear in the input image are multiplied by the filter function to yield the filtered spectrum used to create the filtered image. The filter function is set to one from the DC frequency component (0 cycles per pixel) to the Nyquist frequency (0.5 cycles/pixel). Thus, the output spectrum and image are identical to the input spectrum and image.

have been separated into four distinct horizontal rows. This pattern is similar to that found when imaging bar phantoms for scintillation camera resolution quality control. The frequencies have been separated, instead of being presented on top of each other, so that the effects of filtering may be more easily demonstrated. The lower row of the image represents the highest frequency that can be accurately reproduced, alternating bright and dark pixels yielding 32 bars per 64 pixels or one cycle for each two pixels (0.5 cycles/pixel). The next two rows represent one cycle for each four pixels (0.25 cycles/pixel) and one cycle for each eight pixels (0.125 cycles/pixel), respectively. The top row represents one complete cycle over the width of the image or one cycle for 64 pixels.

The highest possible frequency which may be faithfully displayed (0.5 cycles/pixel) is called the Nyquist frequency. If the source image has more variations than the Nyquist frequency it will not be faithfully reproduced, and some high frequency information will be lost. This loss of information is called aliasing. A common example of aliasing is the way the spokes of wagon wheels in old westerns on television appear to rotate backwards while the wagon is going forwards. This occurs because the sampling rate of video (around 30 frames per sec) is not fast enough to show a true representation of the wheel going forward. The positions of the spokes of the wheels change more rapidly than the sampling rate. In the same way, if the image source has changes in intensity over a distance shorter than two pixels (a higher frequency than the Nyquist frequency) the result will be an image that does not accurately reflect the radioactive source distribution.

As was mentioned before, filters may be applied in frequency space by multiplying the magnitude of a given frequency component by the filter magnitude at that frequency. Notice that in figure 3, the filter has a value of one over the entire applicable frequency range (all frequencies below the Nyquist frequency). Each frequency component is multiplied by one and the output image and spectrum are identical to the input image and spectrum.

von Hann Filter

A common filter used in nuclear medicine image processing is the von Hann filter shown in figure 4. This low-pass filter is usually called a Hanning filter in the nuclear medicine field. The filter has a magnitude of one at the lowest frequencies and decreases to zero at a frequency known as the cutoff frequency. The Hann filter in figure 4 has a cutoff frequency of 0.5 cycles/pixel and removes this frequency component from the image. When this filter multiplies the input image spectrum, it can be seen that the 0.5 cycles/pixel frequency bars of the input image have been removed in both the filtered spectrum and image. The amplitude of the next two middle frequency bars have been lessened by the product of the filter value at those frequencies, and these bars do not have as much contrast in the output image as in the input image.

Another Hann filter, with a cutoff frequency of 0.25 cycles/pixel, is shown in figure 5. It can be seen that now two rows of bars (0.25 and 0.5 cycles/pixel) of the output image are removed. The next lower frequency bar shows diminished contrast caused by the lessening of the amplitude of the frequency component. As before, the top bars with the lowest frequency show little change.

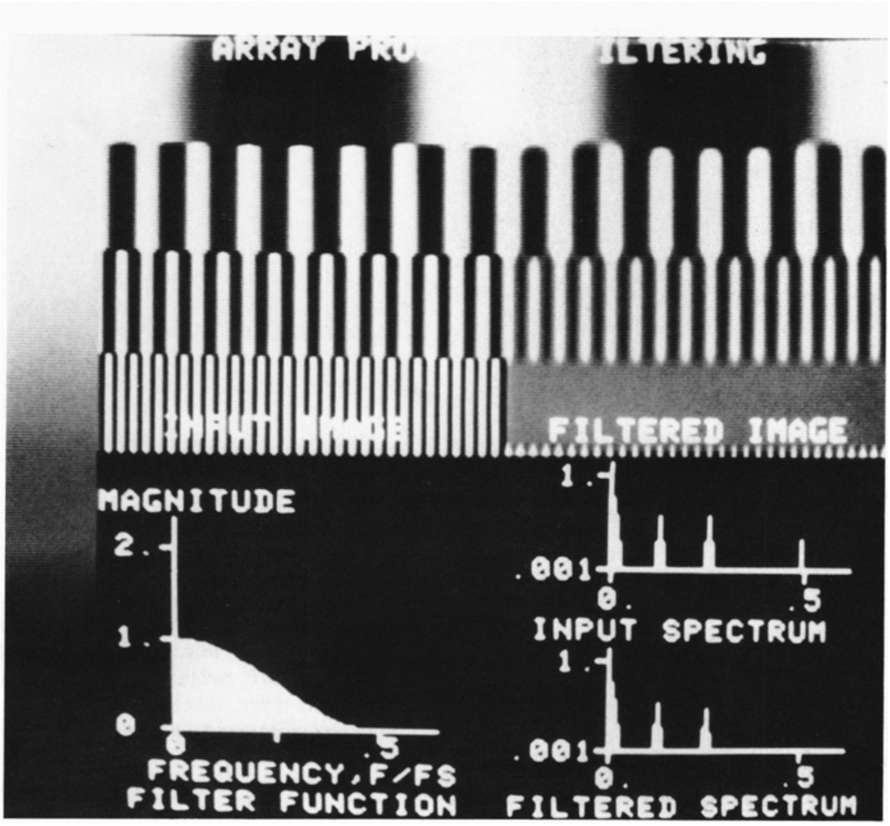

FIG. 4. Demonstration of Fourier filtering using a Hann filter with a cutoff frequency of 0.5 cycles/pixel. The discrete frequencies plotted in the input spectrum are multiplied by the filter function to yield the filtered spectrum used to create the filtered image. Note that the 0.5 cycles/pixel frequency, which is the highest frequency in the input image, after it is multiplied by the zero value of the filter at that frequency, totally disappears in both the output spectrum and filtered image. Note that the amplitude (contrast) of the next higher frequency is also reduced by the filter.

Poisson Noise

One of the primary reasons for applying low-pass (smoothing) filters is to reduce the high-frequency random noise associated with Poisson noise (statistical count fluctuations). The effect of filtering on Poisson noise is shown in figure 6 where the same input image used in figure 3 is degraded by the addition of random noise. The same Hann filter used in figure 4 is applied. The Hann filter removes the high frequency random noise resulting in a better definition of the 0.25 and 0.125 cycles/pixel bars, but also removes the highest (0.5 cycles/pixel) frequency component of the image and markedly reduces contrast in the 0.125 cycles/pixel bars. The Hann filter appears to be better suited for studies where higher statistical accuracy is needed at the expense of a loss in spatial resolution.

Butterworth Filter

Another low-pass filter, the Butterworth filter, is demonstrated in figure 7. Two parameters are needed to define the Butterworth filter, the cutoff frequency and the order of the filter. The order of the filter is related to how fast the filter is cut off, the higher the order the sharper the cutoff. The Butterworth filter may be applied using a sharper cutoff than the Hann filter, retaining contrast at higher frequencies while still eliminating the Poisson noise. The filter in figure 7 has a cutoff of 0.4 cycles/pixel and an order of 25. The Butterworth filter is better suited for studies where higher resolution needs to be preserved at the expense of higher statistical count fluctuations.

Definition Ambiguities

Before going on to clinical applications of filtering, a word of caution should be introduced about the terminology used by the nuclear medicine industry in describing filters. While some manufacturers define the cutoff frequency at the point where a filter value drops to zero, others define the cutoff frequency as the point where the filter magnitude drops below a given value. The meaning of the term cutoff frequency may also differ for different kinds of filters (even by the same manufacturer). It is also important to define the units being used when speaking of frequencies. Frequencies may be defined as cycles/pixel, cycles/centimeter, or in any number of other units. Unless it is clear which units are being used, and how the cutoff frequency is defined, inappropriate values may be used to define filters, resulting in improper filtering.

CLINICAL APPLICATIONS

Planar Imaging

As mentioned earlier, the inclusion of random noise in nuclear medicine images requires the application of digital image processing techniques for removal. The characteristics which are modified by the filter operator are the spatial and contrast resolution and the signal-to-noise ratio (*3*). Application of a filter to improve one of these characteristics, whether in the spatial or frequency domain, will affect the others. The choice of the type and design of filter to use is based on the tradeoffs of these characteristics, which depend on the type of imaging and information to be extracted. Therefore, care must be taken in selecting the filter of choice so that erroneous results are not produced.

Most nuclear medicine computer systems offer a selection of image processing filters that fall into the spatial or frequency

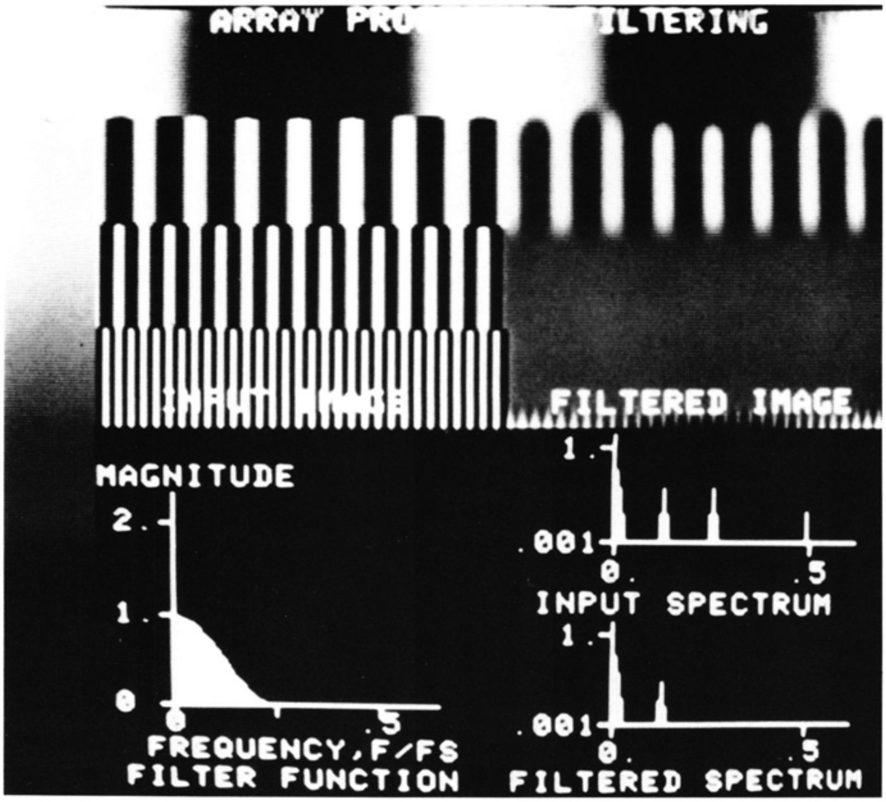

FIG. 5. Demonstration of Fourier filtering using a Hann filter with a cutoff of 0.25 cycles/pixel. Note that the two highest frequencies (0.5 cycles/pixel and 0.25 cycles/pixel) totally disappear in both the output spectrum and filtered image after applying the filter. Note that the ampliture (contrast) of the next higher frequency is also reduced by the filter.

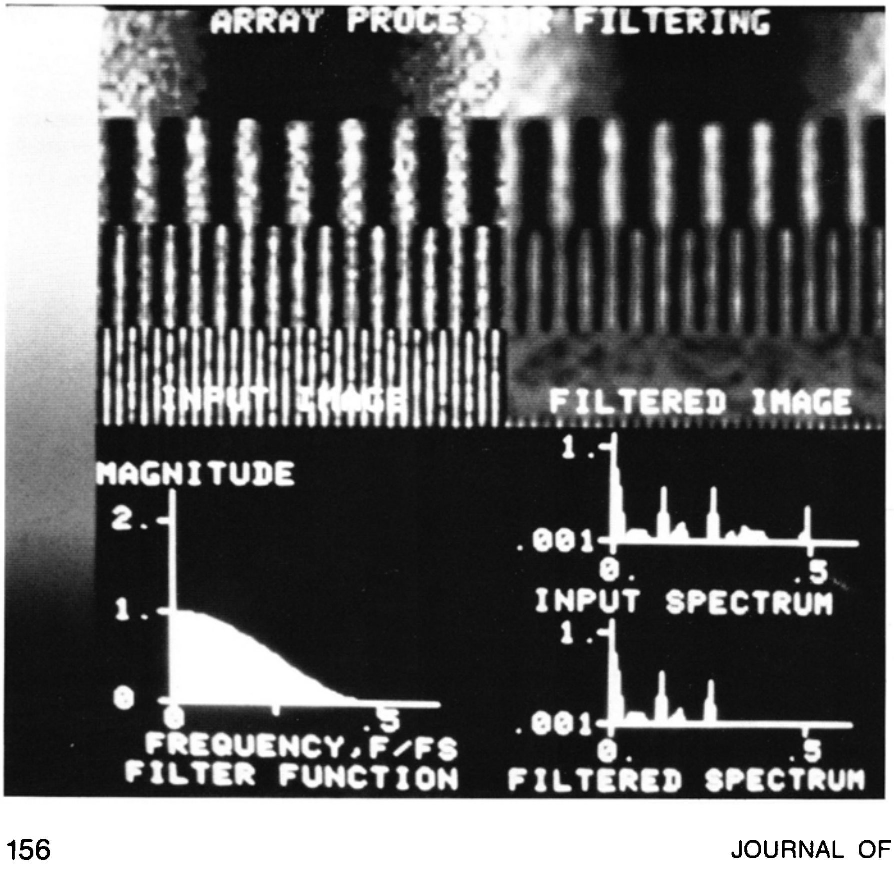

FIG. 6. Demonstration of Fourier filtering of the image in Figure 3 where statistical noise has been added. The filter function is a Hann filter with a cutoff frequency of 0.5 cycles/pixel. Much of the high frequency noise, along with the highest frequency (0.5 cycles/per pixel) disappears in both the output spectrum and filtered image. Note that the amplitude (contrast) of the next higher frequency is also reduced by the filter.

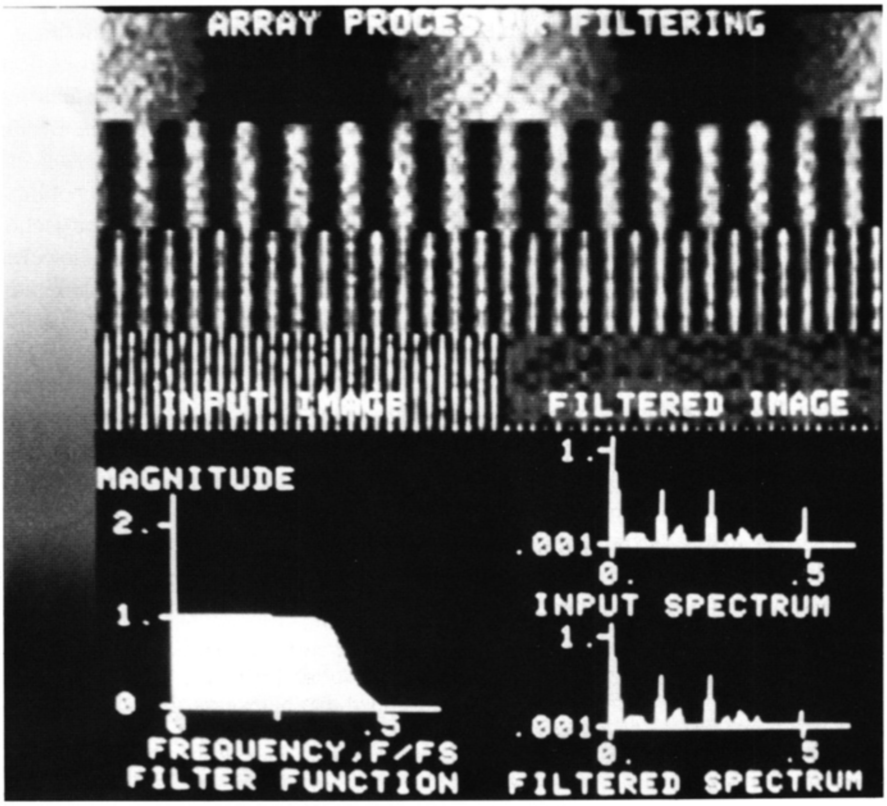

FIG. 7. Demonstration of Fourier filtering of the input image from Figure 6 using a Butterworth filter. As was the case for the Hanning filter in Figure 6, much of the noise, along with the highest frequency (0.5 cycles/pixel) disappears in both the output spectrum and filtered image. Note that the amplitude (contrast) of the next higher frequency is much greater than with the use of the Hann filter. Thus, contrast has been retained, while the noise has been effectively removed.

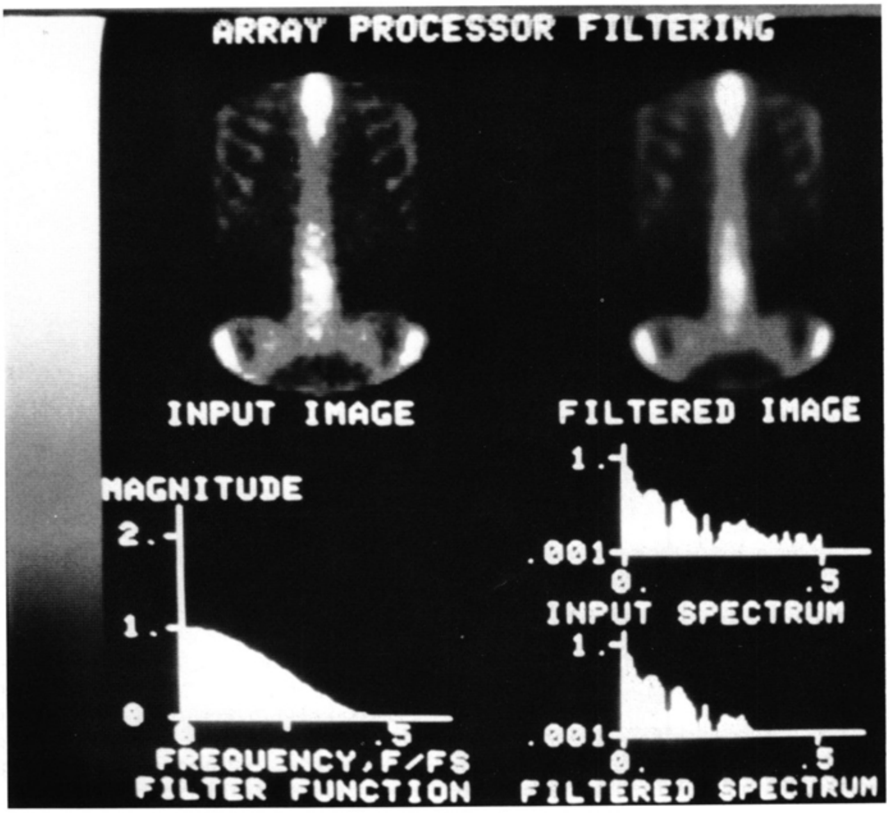

FIG. 8. Representative planar bone scan (upper left), that has been filtered in frequency space (upper right) using a Hanning filter with cutoff frequency of 0.5 cycles/pixel (lower left). Spectra of amplitude versus frequency for the unfiltered and filtered image are shown (lower right).

domain. The selection of which to use depends on two factors, the degree of accuracy and the amount of time available for filtering. For example a 3 × 3 or 5 × 5 spatial kernel may yield acceptable results for viewing a 64 × 64 multigated cardiac study. But when more exacting solutions are required, as in tomography, the spatial kernel becomes excessively large and takes significant time to apply to an image. There exists a break-even point at which the use of a frequency domain filter becomes not only more exact but faster (*4,5*). Because we are discussing frequency filtering, we will assume that the break point has been exceeded.

Figure 8 is formatted in the same way as figure 3. Figure 8 shows an example of a planar bone image before (upper left) and after (upper right) application of a Hanning filter with a cutoff frequency of 0.5 cycles/pixel. In the spectrum display (lower right) one can see the effect of the filter on the higher frequencies. While the filtered image shows a retention of rib detail, the high frequency variations (graininess) in the image have been removed. In this case the signal-to-noise ratio is improved, with a subsequent loss in some spatial resolution.

Figure 9 is an example of the same planar bone image before and after the application of a Butterworth filter with a cutoff frequency of 0.1 cycles/pixel. While not clinically useful for this image, it demonstrates effectively the point of creating erroneous results. In the filtered image, one can see the vastly improved signal-to-noise ratio, and the significant loss of spatial resolution. In this case, not only has the higher frequency noise been removed, but clinical information has been filtered out, causing loss of detail in the thorax and spine and the removal of the frequency defining the ribs.

Tomography

The basic process for reconstructing tomographic images has been presented elsewhere (*6*). Tomographic projections are nothing more than a series of planar images taken at different angles around the patient. These images are then backprojected into transaxial images. The transaxial images can then be re-oriented to produce sagittal, coronal, or oblique angle images. As with planar images, tomographic projection data includes a certain amount of noise. This noise, however, is amplified in the backprojection technique and causes the resultant transaxial information to appear different from the radioactive distribution being studied (*1*).

More specifically, the goal of tomographic reconstruction is to provide a blur-corrected transaxial image from processed planar projections. The solution to this problem in frequency space involves the application of a ramp filter (Fig. 10) to the frequency components of each projection. Once filtered, each projection is transformed back into the conventional spatial domain and then backprojected to form the blur-corrected transaxial tomograms. Unfortunately, this technique may be time consuming.

The ramp filter, as shown in Figure 10, is a high-pass filter. Previously it was explained that high-pass filters enhance the edges of the radioactive distribution in the image. Intuitively, the reconstruction process known as filtered backprojection may be thought of as first extracting the edges of a three-dimensional radioactive source from different angles and second backprojecting the edges from the different angles to generate the count distribution in a transaxial image. The ramp filter, being a high-pass filter that linearly enhances higher

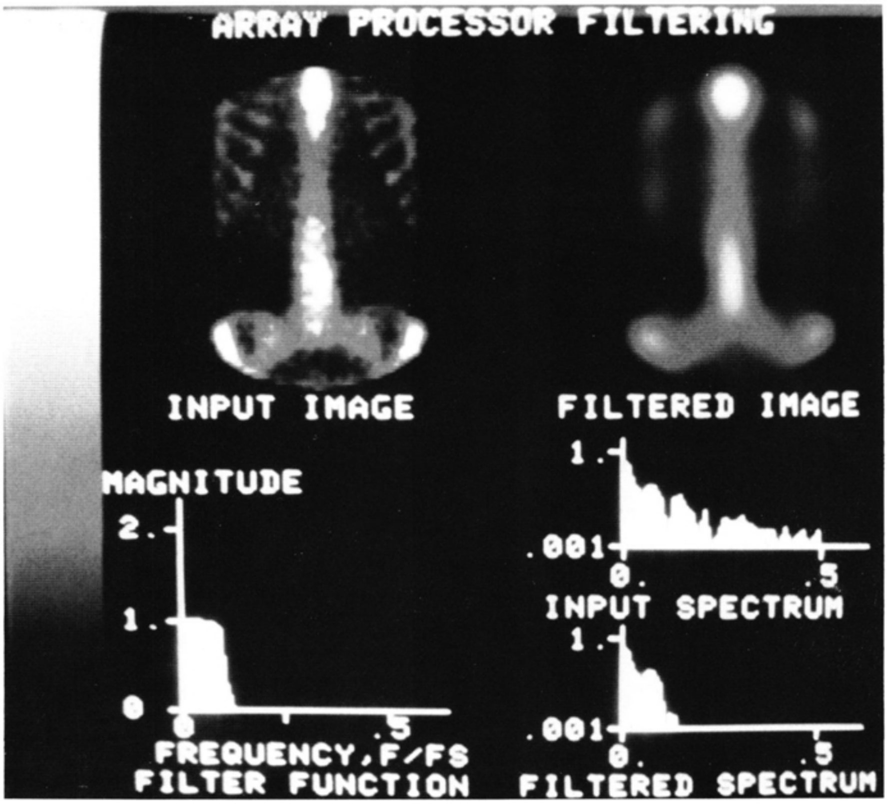

FIG. 9. Representative planar bone scan (upper left), that has been over filtered in frequency space (upper right) using a Butterworth filter with cutoff frequency of 0.125 cycles/pixel (lower left). Spectra of amplitude versus frequency for the unfiltered and filtered image is shown (lower right). Note that the frequency which defines the ribs (the middle peak) has been removed and is the same for the ribs in the resultant planar image.

frequencies, yields the highest resolution possible in a reconstruction but also propagates the high frequency noise associated with low-count statistics. This propagation of noise often results in clinically uninterpretable images (Fig. 11, unfiltered-high).

Modification to the ramp filter must be made to compensate for the undesirable noise. This is done by combining the ramp characteristics with those of a low-pass filter or window. Hanning and Butterworth filters are two commonly used windows that can be multiplied by a ramp filter to yield different degrees of trade-off between reduction of statistical noise versus degradation of spatial and contrast resolution.

Figure 10A, illustrates examples of a rectangular filter, a Hanning filter with the cutoff frequency of 0.5 cycles/pixel, and a Hanning filter with the cutoff frequency set to 0.375, which may be used to window the ramp filter. When these windows are multiplied by the ramp, the filter function then becomes those in figure 10B. When this is done, the filter is referred to as a ramp-Hanning filter for clarification and is somewhat characteristic of a band-pass filter.

Noise can be removed from the final images either by applying the filters to the planar projections prior to backprojection (filtered backprojection) or afterwards by filtering the transaxial slices (post-processing filtering). Whether to filter before or after backprojection, remains debatable. It may be argued that filtering prior to backprojection is more desirable for two reasons. First, it reduces the propagation of noise at an earlier stage in the image formation process. Second, it promotes the implementation of a filter symmetric in three dimensions (same resolution in the X, Y, and Z directions). Proponents of post-processing would argue that the same results could be obtained by the careful selection of filters applied to the transaxial images.

One of the more interesting properties of frequency filtering is that once a frequency has been removed from an image by a filter, further processing will not bring that frequency back to the image. Figure 11, which contains tomographic images from a ^{201}Tl study, helps to demonstrate this point.

The three filters in figure 10 were separately applied to the unprocessed planar ^{201}Tl projections and backprojected (Fig. 11, unfiltered). These same filters were applied to planar projections smoothed with a nine-point smooth operator (1) which duplicates the effects of a Hann filter (with a cut-off of 0.5 cycles/pixel) and backprojected (Fig. 11, filtered). One can notice the projection of noise in the image labeled unfiltered-high and acceptable levels of resolution and noise suppression in the unfiltered medium and low (resolution) images. Of interest are the images in the filtered section (right) which show little or no change when refiltered. Because the projections were prefiltered, the application of another filter with a higher cutoff frequency has minimal effect. This principle is important in predicting the effect of applying multiple filters at different stages of processing.

SUMMARY

Nuclear medicine images (planar and tomographic) can be decomposed into their corresponding frequency patterns by

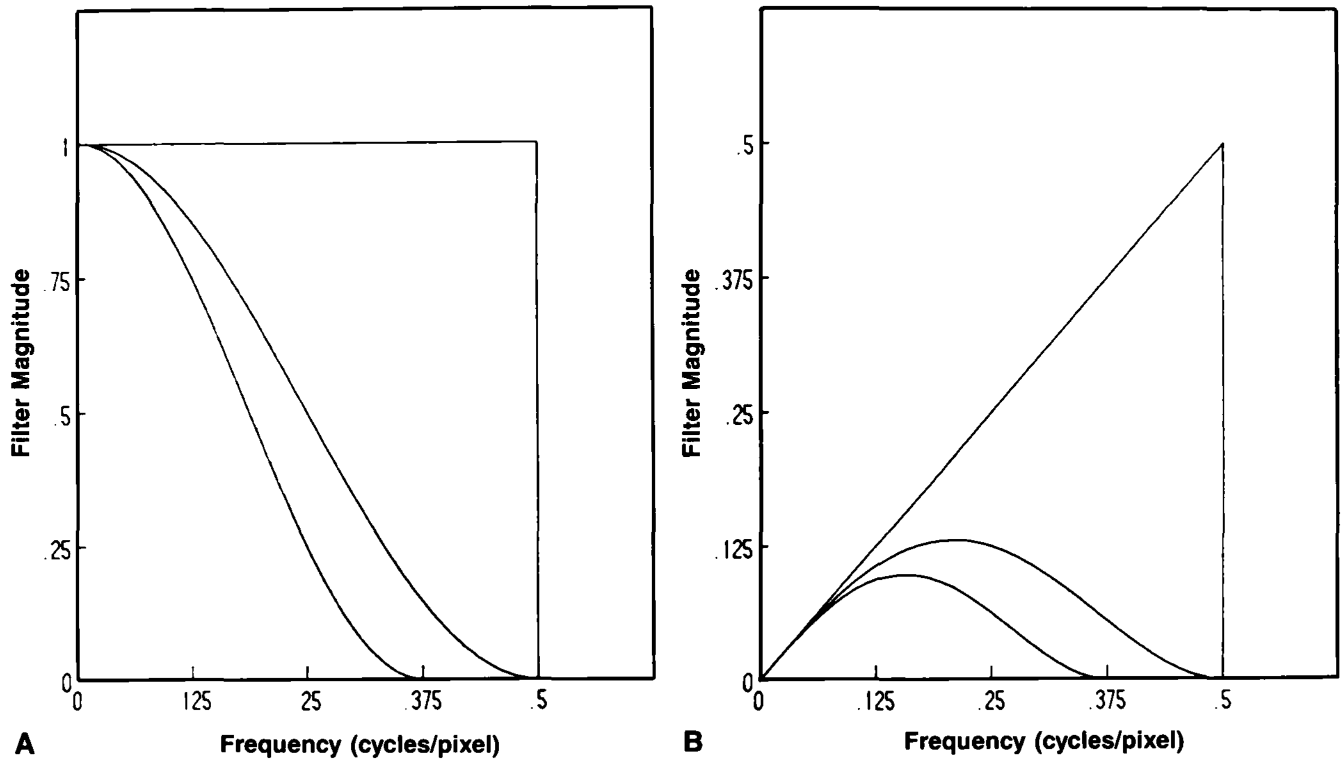

FIG. 10. (A) Plot of filter amplitude versus frequency for a rectangular Hanning with cutoff of 0.5 cycles/pixel and Hanning with cutoff of 0.375 cycles/pixel filter windows. The same filters (B) were then multiplied by a ramp filter and plotted as magnitude versus frequency.

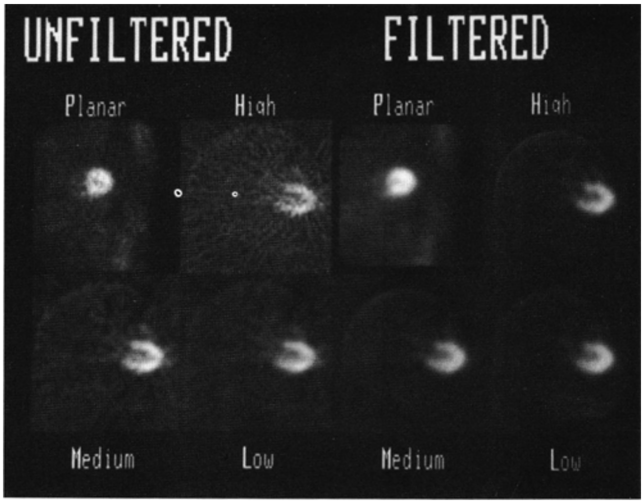

FIG. 11. The same set of ^{201}Tl raw projection images backprojected without prefiltering (unfiltered) and backprojected after prefiltering using a nine-point smoothing operator (filtered). The high, medium, and low-pass frequency filters defined in Figure 10B were then applied to representative transaxial slices. In the unfiltered data set, noise is progressively removed with the application of the filters. In the filtered data set, however, little or no change can be seen in transaxial data because noise along with higher frequencies had been removed previously by the prefilter.

use of FFTs. This transformed image can then be returned to the spatial domain using an inverse Fourier transform. Once an image is in the frequency domain, selected frequencies can be left alone, modified, or removed by application of filters.

Filters can be described by their characteristic shape (Hanning, Butterworth, ramp, etc.) and by their cutoff frequency. The highest frequency which can be faithfully reproduced in the filtered or unfiltered image is the Nyquist frequency or 0.5 cycles/pixel. A general way of describing filters is by the frequencies that are allowed to pass through them (high-pass, low-pass, band-pass). Different manufacturers describe the same filter in different ways.

Trade-offs in statistical noise as opposed to spatial and contrast resolution are the two main criteria for selecting filters.

The use of filters in frequency space is preferred over processing in the spatial domain because of their superior flexibility in defining the shape of the filter function, and because they are intuitively more straightforward.

The solution to the tomographic reconstruction problem is the application of the ramp filter to each projection prior to backprojection. This filter maintains the highest degree of spatial resolution but enhances noise in the resultant transaxial slices. Therefore, modification to the ramp filter is made using filter characteristics from the Hanning, Butterworth, or other filters. The resultant filter being referred to as Ramp-Hanning, for example.

Filtering may be done prior to or after backprojection. Once a frequency component has been removed from an image, however, there is no filter that can return it.

ACKNOWLEDGMENTS

The authors greatly acknowledge the suggestions provided by Jan Berg and the secretarial assistance of Sybil Sheppard.

REFERENCES

1. Garcia EV. Digital processing in nuclear medicine imaging. *J Nucl Med Technol* 1986;14:21–30.

2. Baxes GA. *Digital Imaging Processing.* Englewood Cliffs, NJ: Prentice Hall, 1984.

3. Garcia EV, Ezekiel A. Digital processing in cardiac imaging. *Int J Cardiac Imaging* 1985;1:5–27.

4. Huesman RH, Gullberg GT, Greenberg WL, et al. *User Manual: Donner Algorithms for Reconstruction Tomography.* Berkeley: Lawrence Berkeley Laboratory, 1977:47–59.

5. Gonzalez RC, Wintz PA. *Digital Imaging Processing.* Reading: Addison-Wesley, 1977:58–67.

6. Eisner RL. Principles of instrumentation in SPECT. *J Nucl Med Technol* 1985;13:23–31.