SESSION X AN INTRODUCTION TO TWO-DIMENSIONAL FAST FOURIER TRANSFORMS AND THEIR APPLICATIONS

Introduction to two-dimensional Fourier analysis

M. S. RZESZOTARSKI

Case Western Reserve University, Cleveland, Ohio 44106

F. L. ROYER

Cleveland Veterans Administration Medical Center, Cleveland, Ohio 44106

and

G. C. GILMORE

Case Western Reserve University, Cleveland, Ohio 44106

The application of two-dimensional Fourier analysis provides new avenues for research in visual perception. This tutorial serves as an introduction to some of the methods used in two-dimensional Fourier analysis and an introduction to two-dimensional image processing in general.

Two-dimensional Fourier analysis is a powerful tool that has seen increasing popularity in recent years due to rapid advancements in digital image processing hardware. The purpose of this paper is to present an introduction to two-dimensional Fourier analysis using numerous examples with minimal mathematical development. The presentation serves as an introduction to two-dimensional image processing using the two-dimensional Fourier transform as a tool to achieve that end. Specific applications of Fourier analysis to psychology are covered by Royer, Rzeszotarski, and Gilmore (1983).

DIGITAL IMAGE REPRESENTATIONS

An image can be described mathematically as some function $f(x,y)$, where x and y are the spatial coordinates of the picture and $f(x,y)$ represents the brightness of the image at the point (x,y) . Computers work with digital images in which $f(x,y)$ is a function with numeric values for x , y , and brightness. The brightness is referred

to as the gray level in digital images, and each element in the image is called a picture element, or pixel for short. A typical digital image may contain a 256 by 256 matrix or larger, with a fixed number of gray levels. The gray level represents the amount of light that is transmitted through a film containing the image. Typical ranges of gray levels are from 0 to 255, representing a percent transmittance through the film of from .1% to 100%.

Figure 1a shows an image in digital form with 256 gray levels and a matrix size of 256 by 256 pixels. Figures 1b-1d illustrate what happens if one samples more coarsely so that the picture element size becomes larger. The ability to resolve fine detail is lost in Figure 1d and reduced in Figure 1c. One must sample an image so that the required detail is present in the digital image. In this example, if one is trying to detect eyeballs, then Figure 1c may have sufficient resolution, but if one wants to examine the details in the hair, one must use much finer sampling (Figure 1a in our example). The sampling rate is determined by how much detail one must retain in an image in order to see the objects of interest. Clearly, there are tradeoffs here. The maximum image size is determined by available digital memory, whereas the minimum image size is determined by the required resolution.

The author wishes to thank Lotte Jacobi for the use of the photograph of Albert Einstein. Requests for reprints should be addressed to Fred L. Royer, Research Laboratory 151B, VA Medical Center, Brecksville, Ohio 44141. This work was supported by NIA Grant R01 AG 03178-01 and by the VA Medical Research Service.

Figure 1. Digital images with 256 by 256 samples, 64 by 64 samples, 32 by 32 samples, and 16 by 16 samples, each with 256 gray levels.

Figure 2. Digital image with 256 by 256 samples and 32 gray levels, 8 gray levels, 4 gray levels, and 2 gray levels.

Another factor in storing a digital image is how many gray levels should be stored for each picture element. The images shown in Figures 2a-2d illustrate storage in 32, 8, 4, and 2 gray levels. Note that Figure 1a, which has 256 gray levels, looks essentially the same as Figure 2a. The human visual system can detect only about 50 gray levels, so the use of 256 is beyond the range of human visual system detection. As one reduces the

number of gray levels in an image, contouring becomes evident, starting with Figure 2b. The image is broken into regions of constant gray level, which may produce undesirable artifacts in the resulting image. As the number of gray levels is reduced, the number of bits of digital storage is also greatly reduced. In Figure 1a, there are 16.78 million bits of digital information in the 256 gray-level 256 by 256 image. If a 128 by 128 image size with only 32 gray levels is used, then only .5 million bits are required. In applications illustrated in this paper, the images all contain 256 gray levels and use a matrix size of 256 by 256.

IMAGE PROCESSING HARDWARE

The process of converting a photograph or drawing to digital form requires some type of digital image processing system. A typical system consists of an image digitizer, a digital computer, a mass storage device, and a display device for output. The digitization hardware can be a television camera system, a rotating drum film scanner, or an x,y digitizer, if simple line drawings are used as input images. Most minicomputers can perform the analysis of digital images, although some are tailored for this purpose and work much more efficiently than others. The output device can be a video monitor (television) or a rotating drum scanner. A drum scanner rotates under computer control, and the light intensity through the film at discrete coordinates is recorded in digital form on magnetic tape as the drum spins and steps across the film. For output, a negative is mounted in a light-tight box that is then exposed by a beam of modulated light using the digital data on magnetic tape to determine the modulation. A video monitor is commonly used for output in image processing applications. The monitor can be photographed, as illustrated by most of the photos contained in this paper, or it can be used directly as the viewing screen.

TWO-DIMENSION FOURIER TRANSFORMS: AN INTUITIVE INTRODUCTION

The principle of Fourier analysis is based on the premise that any image (or signal if one-dimensional) can be equivalently represented in two different domains, a spatial (or time) domain and a frequency domain. Fourier stated that one can perform a linear transformation between the two domains and still maintain the uniqueness of the image (or signal). If you have a one-dimensional time domain signal represented by some number of samples N , then you can equivalently represent that signal in the frequency domain using sines and cosines of varying frequencies with N samples representing the amplitudes of the individual sine and cosine components. In two dimensions, the representation is the same, except there are N by N samples. This transformation between the two domains has some additional advantages due to changes in some of the properties of

images (or signals) when working in one domain or the other.

The two-dimensional Fourier transform is difficult to comprehend at first glance, but it can be better understood in terms of some simple examples using images with information in only one direction (one-dimensional images). These are similar to one-dimensional signals, and the Fourier transforms will be familiar to one who has worked with one-dimensional signals. An example of a one-dimensional image is the sine-wave grating illustrated in Figure 3. The image varies in gray level sinusoidally across the image but is constant in the vertical direction. This represents an image with frequency content in only one direction, and with only one frequency component present. We can examine Fourier's principle by constructing a familiar image using the summation of various sinusoids. Let us construct a one-dimensional rectangular pulse image in which we have a region of brightness corresponding to the positive part of the pulse and two dark regions representing the negative portions of the rectangular pulse. This can be constructed by summing together a series of sinusoids with varying frequencies and amplitudes. Figures 4-7 illustrate the successive summation of sinusoids. The upper curve in the images shows the resulting gray-level profile across the image. The resulting image looks more and more like a rectangular pulse as higher frequency sinusoids are added in. Note that the high-frequency components have the most effect on the transition region between black and white. In an image, regions of rapidly changing gray scale are referred to as edges, and it is the edge information that requires high-

Figure 3. Sine-wave grating image, 512 by 512 samples.

Figure 4. Rectangular pulse image with one sinusoid (fundamental).

Figure 5. Rectangular pulse image with two sinusoids summed.

frequency information to represent it. This is an important concept that will have implications in later discussions.

The extension to two dimensions is complicated by the multiplicative nature of two-dimensional images. In this case, the sinusoids in the horizontal direction are summed as are those in the vertical direction, and the resulting image value at some point is the product of these two sums of sinusoids. This can be better under-

Figure 6. Rectangular pulse image with three sinusoids summed.

Figure 7. Rectangular pulse image with 25 sinusoids summed.

stood by examining Figure 8, in which an image is illustrated with only one frequency component in each direction. In this case, the image is dark wherever one of the cosines is negative, and the image is bright if the two cosines are positive. In a typical image, there are many sinusoids in each direction so the image is far more complex. An example of this summation is illustrated in Figures 9-12, in which an image consisting of a sum of sinusoids in both directions is illustrated. The resulting image is a bright square surrounded by a dark background. The figures illustrate the successive addi-

tion of frequency components in each direction. The lower left quadrants of Figures 9-12 show the sinusoid currently being added to the images, whereas the upper left and lower left and lower right quadrants illustrate the summation in one dimension, as shown previously. The upper right image is the resulting product of the upper left and lower right images (which are sums of sinusoids in one direction) and shows the development of this bright square as more and more sinusoids are added.

Figure 8. Image constructed from one sinusoid in each direction.

Figure 9. Square pulse image using one sinusoid in each direction.

Figure 10. Square pulse image using two sinusoids in each direction.

Figure 11. Square pulse image using three sinusoids in each direction.

THE FOURIER ENERGY SPECTRUM

The frequency domain representation of an image or signal is frequently displayed as an energy spectrum. Since one has a series of sines and cosines with varying amplitudes at each frequency, it is useful to represent

the distribution of amplitudes using the sum of the squared amplitudes for the sines and cosines at each frequency. This is referred to as the magnitude squared response and is directly related to the energy of the image at a given frequency. In the one-dimensional case, a plot of the magnitude squared vs. frequency is often provided so one can determine where the energy in a signal is located. When working with images, one has two-dimensional information that cannot easily be represented graphically. As a result, one commonly displays the energy spectrum as an image in which the brighter regions of the image correspond to higher energies. Further details of this display are provided later.

To better understand the Fourier energy spectrum, several one-dimensional examples will be examined first. Figure 13 shows two sine-wave gratings on the left as the spatial domain images, and two energy spectra on the right, representing the frequency domain image. Since each of the grating images contain only a single frequency component, one would expect to see only a single bright spot corresponding to that frequency. However, the Fourier representation of a signal yields both positive and negative frequency information, so the sine-wave frequency component is split equally between the positive and negative components. The bright spot in the center corresponds to the zero-frequency location in the image, or in our case, the mean luminance of the image, which is positive for any image (unless it is totally black). Horizontal sine-wave gratings yield a horizontal spectrum, and vertical gratings yield a vertical spectrum. Also, the sine-wave grating with higher frequency content (more sine-wave cycles in the image) has its components displaced further from the dc

Figure 12. Square pulse image using 25 sinusoids in each direction.

Figure 13. Sine-wave gratings and their energy spectra.

Figure 14. Image and spectrum using one sinusoid in each direction.

or zero-frequency point. The highest frequency one can achieve is the case in which the grating consists of alternating rows of black and white lines, and the resulting spectrum is located at the ends of the energy spectrum lines in the spectral images.

The extension of energy spectra to two dimensions is straightforward if one keeps in mind that it is now a multiplicative process. Figure 14 shows a spatial image in the lower left that consists of the product of two single frequencies, one in each direction. The corresponding energy spectrum is illustrated in the upper right. Note that there are nine discrete points at which the energy

is significant. The central point is again the dc or mean luminance of the image. Since the Fourier transform yields both positive and negative frequencies, we must explain the presence of four pairs of frequency components. These can be explained if we indulge in some mathematics to describe the process. In the X direction, the signal is of the form $A(1 - \cos(x))$, where A is the mean value for that particular direction, and in the Y direction, $B(1 - \cos(y))$. Since we have a multiplicative process, the result is some $A(1 - \cos(x))$ times $B(1 - \cos(y))$, which with some manipulation yields $AB(1 - \cos(x) - \cos(y) + \cos(x - y)/2 + \cos(x + y)/2)$. The math is not so important as is the finding that we now obtain four pairs of frequency components in the energy spectrum. Two pairs of components are in the respective directions of the original cosines in the spatial image, and two sum and difference pairs arise due to the multiplicative nature of the process.

A typical image contains many frequency components in each orientation, so the resulting energy spectrum is fairly complex. We can take a look at the square pulse as an example of a more complex image that is still simple enough to understand. Figure 15 illustrates the energy spectrum for both an ideal square pulse and a square pulse constructed with only a few frequency components. The energy spectrum for the constructed image contains only relatively low-frequency information. The ideal square pulse contains energy all the way out to the highest resolvable frequency.

TWO-DIMENSIONAL FOURIER TRANSFORMS: THEORY

The basic theory of the two-dimensional Fourier transform is presented with a few equations and many

Figure 15. Square pulse and spectrum, ideal and limited number of sinusoids cases.

examples of the images and energy spectra that are obtained when one examines some of the properties of the Fourier transform. For more detailed mathematics, refer to Andrews (1978), Brigham (1974), Dainty and Shaw (1974), Gonzalez and Wintz (1977), Huang (1975), Oppenheim and Schaffer (1975), Rabiner and Gold (1975), and Rosenfeld and Kak (1976).

The discrete two-dimensional Fourier transform of an image $f(x,y)$ is

$$F(u,v) = \frac{1}{N^2} \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} f(x,y) \text{EXP}[-j2\pi(ux/N + vy/N)], \quad (1)$$

where it is assumed the image is of size N by N such that $f(x,y)$ is defined for $x = 0, 1, 2, \dots, N-1$ and $y = 0, 1, 2, \dots, N-1$, and j is assumed to be the square root of -1 . The complex exponential has been represented here as a summation of cosines and sines at varying frequencies. The transform can be factored into a separable form that is computationally more appealing. It is equivalent to Equation 1 mathematically and is

$$F(u,v) = \frac{1}{N^2} \sum_{x=0}^{N-1} \text{EXP}[-j2\pi ux/N] \sum_{y=0}^{N-1} f(x,y) \text{EXP}[-j2\pi vy/N]. \quad (2)$$

The resulting Fourier transform $F(u,v)$ is defined for $u = 0, 1, 2, \dots, N-1$ and $v = 0, 1, 2, \dots, N-1$, where u and v are frequency indexes. The samples in the original image are spaced some dx in distance apart along the x -axis. The y -axis samples are also spaced some dx in distance apart (equal x and y sample spacing). In the frequency domain, the frequency components are spaced du apart, where du is defined as $(1/N)dx$. (The same expression applies to dv .) The sample spacing establishes the resolution in frequency one can achieve. For example, to attain better frequency resolution (smaller du), one must either take more samples (increase N) or sample more finely (reduce sample spacing dx). This has implications that will be discussed later.

While the equations appear complicated at first, they can be manipulated (Rabiner & Gold, 1975) to make them more manageable. If we make use of the separability demonstrated in Equation 2, then the two-dimensional discrete Fourier transform can be computed by doing two sets of one-dimensional transforms. This simplifies things to the point that almost any small minicomputer can perform the transformations if enough memory is available. One computes the transform (or inverse transform) by first doing one-dimensional transforms along each row in the image matrix. Then each column of the matrix that was previously row transformed is again Fourier transformed, yielding the two-dimensional Fourier transform. This factorization of

the Fourier transform enables rapid computation since only $2N$ one-dimensional Fourier transformations are required, and these can be computed using a fast Fourier transform algorithm. A matrix transpose is usually performed between transform phases, so only row transforms are done. The resulting Fourier transform matrix is complex, requiring two floating-point N by N matrices (unless clever storage schemes are employed). The individual Fourier transforms are very fast, and the major speed limitations are due to matrix storage and transposition problems.

The inverse Fourier transform also exists and is defined by

$$f(x,y) = \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} F(u,v) \text{EXP}[j2\pi(ux/N + vy/N)]. \quad (3)$$

It is used to transform an image from the frequency domain back into the spatial domain. The inverse transform can also be converted to a more efficient form, similar to Equation 2, enabling efficient computation. Inverse Fourier transforms can also be used to construct images by specifying only the frequency components that are to be included. Figures 3-15 illustrate examples that can be constructed using inverse transformations (i.e., starting with a Fourier spectrum).

TRANSFORM PROPERTIES

In the frequency domain, there are a number of properties that make the Fourier transform a useful tool. Here, the basic transform properties are examined to better understand what happens in each of the domains.

The Fourier transform of an image is often represented in an equivalent form known as magnitude and phase. The Fourier transform $F(u,v)$ is a complex function that can be represented as $F(u,v) = \text{Re}(F(u,v)) + j\text{Im}(F(u,v))$, or the real and imaginary parts of the complex function. The real part corresponds to the cosine terms, and the imaginary part represents the sine terms. The magnitude and phase representation is a polar coordinate transformation in which each sine/cosine pair is specified by a corresponding magnitude and phase. The magnitude function is the square root of the sum of the squares of the real and imaginary parts of $F(u,v)$. The phase is computed by taking the arctangent of $\text{Im}(F(u,v))/\text{Re}(F(u,v))$.

The energy spectrum of an image is defined as the magnitude spectrum squared and is commonly used to represent the Fourier transform of an image. Figure 16 shows two simple images and the energy spectra of those images. To display an image, we have only 256 gray levels, so to see more detail in the energy spectrum it is usually compressed using the function $D(u,v) = \log(1 + \text{magnitude squared})$, which yields a nonnegative logarithmic function with considerable dynamic range.

Figure 16. Square and rectangle and their energy spectra.

Figure 17. Small and large circles and their energy spectra.

The spectrum displayed has dc (zero frequency) positioned in the center of the image, with positive frequencies to the right and up and negative frequencies to the left and down. Both positive and negative frequencies are present in the spectrum, and the lighter regions represent higher energy than do the darker regions. The spectrum can be displayed using only two quadrants rather than four, but historically the entire spectrum has been displayed. Because of symmetry in the spectrum, the original image can be reconstructed if only two quadrants are used, much like the one-dimensional transform.

Figure 17 illustrates the change in the energy spectrum as the size of an image is increased. This illustrates

several important points. The height of the circle viewed at some distance D defines the visual angle V for the object. If the circle is K samples tall, then we have K samples per V degrees visual angle. This defines our sampling rate in terms of samples per degree visual angle. One cycle in an image requires two samples (one black and one white), so the highest frequency that can be resolved is $K/2V$. Hence our power spectrum extends from dc in the center of the image to the Nyquist frequency ($K/2V$) at the edges of the image. In psychological research, the frequency axis is usually labeled in cycles per degree visual angle. The number of samples used to describe the image over some visual

Figure 18. Example of image translation.

Figure 19. Example of image rotation.

angle sets the upper limit on the resolution in the frequency domain.

Several properties of the Fourier transform are of interest in two-dimensional Fourier analysis of images. Specifically, translation, rotation, distributivity, scaling, correlation, and convolution properties will be discussed.

Figure 18 illustrates the case in which an object within an image is moved about or translated. The energy spectrum remains the same, but the phase information is shifted by a linear phase shift. Since we normally observe the energy spectrum for qualitative comparisons, the translation is not reflected in the images, so we can look at images with variable translations and not be concerned with it. It is, however, important if the phase information is to be utilized.

Rotation of an image in the spatial domain results in rotation of the energy spectrum in the frequency domain. Figure 19 is an example of an image rotated by 45 deg and the corresponding spectra. One observes that the spectrum has been rotated by 45 deg, just as the image has.

The distributivity property of two-dimensional Fourier analysis is useful in interpreting energy spectra. It states that the Fourier transform of two images summed together spatially is the same as the sum of the Fourier transforms of the individual images. This is illustrated in Figure 20, in which one can see the individual components of a plus sign. (Recall the vertical bar spectrum is the same as the horizontal bar spectrum rotated 90 deg.) The corresponding spectrum is the complex addition of the two components, so the spectrum looks somewhat different from the individual sections, but the total energy is summed.

The scaling properties relate scaling in both spatial and frequency domains. The scaling relationship states

Figure 20. Example of image superposition.

Figure 21. Example of image scaling.

that as one enlarges an image in the spatial domain, one compresses the Fourier domain image and vice versa. Figures 17 and 21 reflect what happens when one enlarges an image. The scaling property suggests that if one has a very compact spatial image, it will require high-frequency information to represent it. On the other hand, a spatial image that varies slowly over a large region can be represented by a compact spectrum centered over the lower frequencies. It is analogous to the standing waves produced in a pipe organ. If one has a large organ pipe, then the standing waves are very long, producing only low frequencies. If one has a very compact organ pipe, then the sounds it produces are high in frequency.

Another useful property of two-dimensional Fourier transforms is correlation. Correlation in the spatial domain is often used to do template matching, in which one compares the alignment of two images or portions of images. If the image being correlated is larger than 64 by 64 in size, then the Fourier approach may be more efficient than calculating in the spatial domain because of the number of computations required (Rabiner & Gold, 1975). To do correlation, one performs Fourier transformations on the two images, then multiplies one image times the complex conjugate of the other image, and then performs an inverse Fourier transform on the resulting product. (Recall that the complex conjugate of a matrix is just the matrix with the sign of the imaginary elements of the matrix reversed.) Although there is considerable computation involved, it is still more efficient to do it this way rather than shift and multiply in the spatial domain.

Correlation images are useful when performing template matching comparisons. One can also compute correlation images in which the brightness of the image

is proportional to the correlation between two images. This is useful when one has some set of figures (e.g., letters of the alphabet) and is trying to determine which letters would be the most confusing (most correlated with each other). The auto- and cross-correlation images can be computed to see which parts of the letters cause the confusion, and different figures can be selected to minimize this correlation between letters, if desired.

Another function commonly used in Fourier analysis is the coherence function. It is similar to the correlation coefficient commonly used in one-dimensional signal comparisons. In this case, a coherence coefficient is computed in the Fourier domain over some band of frequencies. Hence, one can examine the spectral similarity between two images over some specified range of frequencies. The interested reader should consult the references for further details.

Last and perhaps most useful is the convolution property. Convolution of two digital images is commonly called digital filtering; the spectrum of the original image is modified by some weighting function (a digital filter). Convolution is computed by a series of transforms and multiplications. Two spatial images are forward transformed, the spectra are multiplied together, and then the product image is inverse transformed to obtain the digitally filtered spatial image.

The multiplication is usually performed by weighting an image with a particular type of weighting function. There are four basic types of weighting functions, and each is illustrated in Figure 22. In Figure 22a (lower left) a low-pass filter function is shown in which only those frequency components that fall within the white area are passed. Figure 22b (upper left) illustrates a bandpass filter in which a band of frequencies is passed and all other frequencies are attenuated. Figure 22c

Figure 22. Four typical weighting functions used for convolution. The white areas have a weighting of unity, and the black areas have a weighting of zero.

Figure 23. Original and low-pass filtered images and energy spectra.

Figure 24. Original and high-pass filtered images and energy spectra.

(lower right) is a high-pass filter, and Figure 22d (upper right) is a notch filter in which frequencies shown in white are passed and those in black are attenuated. The calculation of a proper weighting function depends on a number of factors, many of which are beyond the scope of this paper. The interested reader should consult the references previously cited for in depth discussions of the details.

Convolution is a powerful tool for modifying the spectral composition of an image to tailor it for specific applications. Figures 23-26 illustrate the effects of applying the four filter functions defined previously on

Figure 25. Original and bandpass-filtered images and energy spectra.

Figure 26. Original and notch-filtered images and energy spectra.

the Einstein image. The original image and its spectrum are shown below, and the filtered image and corresponding spectrum are illustrated in the upper part of the images. The low-pass filtered image (Figure 23) has a loss of sharpness in the image due to the lack of high frequencies needed to delineate the edges. The high-pass filter (Figure 24) contains only the edge information, since only regions of rapidly changing contrast are

visible. The bandpass (Figure 25) and notch-filtered (Figure 26) images illustrate the effect of examining or removing a selected band of frequencies in an image.

The process of digital filtering an image can be employed to compensate for some degradation in an imaging system or to match two groups of subjects who might "see" things differently. Royer et al. (1983) elaborates this point, but the frequency response that is used in doing the multiplication in the frequency domain is an arbitrary function (with some constraints on symmetry, realizability, etc.). Hence, if one group of subjects responds to a set of images differently from another group, and the differences can be quantitated in the frequency domain (using contrast sensitivity curves), then an inverse filter can be generated that will compensate or "correct" the image so that the two groups will now "see" the image with the same frequency response. Also, there is some evidence that the human visual system responds differently for horizontal, vertical, and diagonal information, so one could construct digital filters to study these properties as well.

CONCLUSIONS

The recent advances in computer and image processing technologies have opened up new avenues for psychological research. Two-dimensional Fourier analysis is a powerful tool for examining the mechanisms for visual processing in the brain, and for better image understanding. This brief tutorial merely skims the surface of two-dimensional analysis, and the interested reader is advised to obtain several of the many excellent books that have been written in this area, including Andrews (1978), Brigham (1974), Dainty and Shaw (1974), Gonzalez and Wintz (1977), Huang (1975), and Rosenfeld and Kak (1976).

REFERENCES

- ANDREWS, H. C. *Tutorial and selected papers in digital image processing*. New York: IEEE Computer Society, 1978.
- BRIGHAM, E. O. *The fast Fourier transform*. Englewood Cliffs, N.J.: Prentice-Hall, 1974.
- DAINTY, J. C., & SHAW, R. *Image science*. New York: Academic Press, 1974.
- GONZALEZ, R. C., & WINTZ, P. *Digital image processing*. Reading, Mass: Addison-Wesley, 1977.
- HUANG, T. S. *Picture processing and digital filtering*. New York: Springer-Verlag, 1975.
- OPPENHEIM, A. V., & SCHAFER, R. W. *Digital signal processing*. Englewood Cliffs, N.J.: Prentice-Hall, 1975.
- RABINER, L. R., & GOLD, B. *Theory and application of digital signal processing*. Englewood Cliffs, N.J.: Prentice-Hall, 1975.
- ROSENFELD, A., & KAK, A. C. *Digital picture processing*. New York: Academic Press, 1976.
- ROYER, F. L., RZESZOTARSKI, M. S., & GILMORE, G. C. Application of two-dimensional Fourier transforms to problems of visual perception. *Behavior Research Methods & Instrumentation*, 1983, 15, 319-326.