

Chapter 4

The Spatial and Frequency Domains

4.1 Introduction

The interconversion between spatial and frequency domains using Fourier and other transforms is of critical importance in image processing and, for some imaging methods, the construction of images from raw scan data. A significant feature of the transforms is that we can convert back and forth between spatial and frequency domains without loss of information or introduction of noise.

4.2 Images in the Spatial and Frequency Domains

4.2.1 *The Spatial Domain*

The concept of the spatial domain requires little introduction. We live in it! Most of the images we are familiar with are spatial domain images – they display a matrix of color or gray scale intensities in a 2D spatial plane. They represent a discrete sampling of the change in intensity of a signal in space and there is a direct correspondence between the coordinates in the image and space in the ‘real world’.

We can perform image processing operations directly on these spatial domain images and we often do. Most domestic image processing software, for example Adobe Photoshop, operates exclusively in the spatial domain. However, there are image adjustments that are faster and more precise if we perform them after first transforming the spatial domain image into its frequency domain equivalent. Also there are some image adjustments that can *only* be performed in the frequency domain. Sometimes we acquire raw image data in the spatial frequency domain, most notably in MRI, and it must be converted into the spatial domain in order to create an interpretable anatomical image.

4.2.2 Common All-Garden Temporal Frequency

When we encounter the term *frequency* we usually think about regular oscillations. Some examples might include radio waves, audio waves, an ultrasound signal, or perhaps the waves at the beach. A more specific description of frequency in these contexts would be temporal frequency – the rate of repetition in time.

Figure 4.1 illustrates how we can often simplify the description of a complex *time domain* signal (e.g. a recorded sound wave) by representing it as a *spectrum* showing the relative intensities of its individual frequency components. In this figure we see just one time dimension – the x axis. Although less familiar in everyday life, 2D examples are common in science. In magnetic resonance spectroscopy for example, it is normal to add third and even higher order time dimensions to investigate molecular structures.

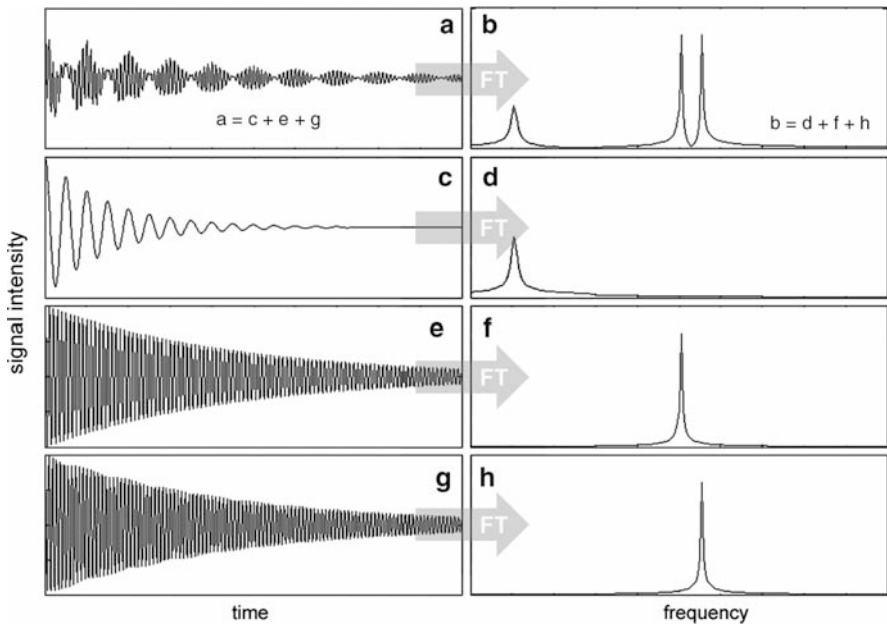


Fig. 4.1 Signals that have a regular periodic variation of intensity over time (*left*) are often represented as *spectra* (*right*) that show the intensities of specific frequency components. The most common method of conversion of a time domain signal to its frequency domain representation is the *Fourier transform (FT)*. In this example the rather complicated *time domain* signal (**a**) has a quite simple *frequency domain* representation (spectrum **b**) containing just three distinct frequency components. These three components, shown singly in spectra **d**, **f**, and **h**, correspond to the time domain signals **c**, **e**, and **g** respectively, and each of these is a simple decaying sinusoid. Signal **a** is the sum of the signals **c**, **e**, and **g**. Physical phenomena that might produce these signals could be as diverse as the vibrations of a musical instrument, or nuclear magnetic resonance in a solution of small molecules

Going beyond the familiarity of the idea of temporal frequency there is no reason we can't apply the same concepts (and the same maths) to signals that change with space. Our brains can do quite a good job of resolving sounds into their temporal frequency components, i.e. recognizing notes. Some people can do this with extreme precision. Human brains have, however, not evolved to perform spatial frequency analysis of what our eyes perceive. For that we need an external tool – preferably a computer.

4.2.3 The Concept of Spatial Frequency

In image processing we often use the term frequency to describe the rate of change of a signal in space, for example the rate at which the pixel intensity changes as we scan across or down an image. In this context we are talking about *spatial frequency* (in fact if we scanned the image and recorded the change of intensity then we would once again have a signal that changed with time and thus the spatial and temporal frequencies would be directly related).

The concept of spatial frequency is extremely useful in image processing. Many of the methods used in analog and digital signal processing (signals often described by their temporal frequency) have direct equivalents in image processing. This transfer of methods from the temporal frequency domain to the spatial frequency domain means much of the terminology has come along for the ride, with occasionally confusing consequences. We will encounter high-pass, low-pass, band-pass and ideal filters in image processing just as we would in (temporal) signal processing.

Let's start with a very simple spatial frequency example. Consider the image shown in Fig. 4.2. As we scan the image from left to right the intensity starts initially at mid gray, increases slowly to white, decreases slowly to black, and then increases again to mid gray. If we were to plot the intensity across one row of the image matrix (all the rows are identical in this image) we would see that the profile has a sinusoidal shape. It looks very similar to the trace we would see if we connected an oscilloscope to the domestic electricity supply and measured the changing voltage. We can say that the intensity changes with a particular spatial frequency – in this case the frequency is one cycle per image width. We can describe this particular intensity modulation with a very simple mathematical expression *of the form*:


$$I = \sin(\omega x) \tag{4.1}$$

where I is the intensity,
 x is the distance across the image,
and ω is the spatial frequency.

Since the value of $\sin(\omega x)$ ranges from -1 to $+1$, in Eq. 4.2 an intensity of $+1$ would represent white, and -1 would represent black.

If the bit depth of our image data was 8 we would have $2^8 = 256$ possible intensities and we would normally use the convention $0 = \text{black}$ and $255 = \text{white}$

Fig. 4.2 This is a very unusual image because the intensity profile in the x direction (a plot of the pixel intensities in any one row of the image matrix) is a perfect sinusoid with wavelength identical to the width of the image. The pixel intensities (I) can be described with a simple expression of the form: $I = \sin(\omega x)$ where ω is the *spatial frequency* of the intensity variation. In this particular image $\omega = \textit{exactly}$ one cycle per image width



when displaying the image. In this case the precise expression for the x -direction change in intensity in Fig. 4.2 would be:

$$I = 127.5 \times \sin\left(2\pi\omega \frac{x}{m}\right) + 127.5 \quad (4.2)$$

where x is the distance across the image in pixels,
 m is the width of the image in pixels,
 and the factor 2π is introduced to convert the spatial frequency from cycles per image width to radians per image width.

To keep this introduction as simple as possible we will just say that the maximum of the expression represents black and the minimum white. Also we won't worry about the need to convert cycles to radians so can leave out the $\frac{2\pi}{m}$. We can ignore these details for now as we are mainly interested in the general form of the intensity change as we move through the space of the image.


Now let's look at the slightly more complex image shown in Fig. 4.3. This image is similar to Fig. 4.2 but we now see small modulations in the intensity superimposed on the single cycle that spans the image. Now the intensity is modulated with two different spatial frequencies – a low frequency of 1 cycle per image width, and a higher frequency of 10 cycles per image width. Note that the amplitude of the higher frequency modulation is about one fifth of that of the low frequency modulation. In this case the (simplified) expression for the intensity modulation would have two terms, one for each frequency:

$$I = \sin(\omega_1 x) + \frac{1}{5} \sin(\omega_2 x) \quad (4.3)$$

Fig. 4.3 A modified version of Fig. 4.2. The original intensity profile has a superimposed 'ripple' that can also be described by a sinusoid. The ripple pattern has *higher spatial frequency* and *lower amplitude* than the underlying profile. In this image the intensity profile can be described as the sum of the two sinusoids of different spatial frequency and amplitude:

$$I = \sin(\omega_1 x) + \frac{1}{5} \sin(\omega_2 x),$$

where $\omega_2 = 10\omega_1$



where $\omega_2 = 10\omega_1 = 10$ cycles per image width

Once again ω_1 has the value of one cycle per image width, or $\frac{1}{m}$ cycles/pixel. The astute will notice that this expression has a maximum value of about 1.2 – whiter than white in our $1 = \text{white}$, $-1 = \text{black}$ scheme. We can ignore this for now as we are mainly interested in the general form of the intensity change as we move through the space of the image.

Look closely at the intensity profile in Fig. 4.3. The addition of the higher spatial frequency (10 cycles per image width) means that the intensity changes much more quickly in the x direction than it does in Fig. 4.2 – the higher spatial frequency represents more rapid changes in image intensity as we scan a row of the image matrix. Put another way, the maximum *steepness* of the intensity gradient increases with the spatial frequency.

If we add a still higher spatial frequency, as in Fig. 4.4, we obtain a pattern that resembles a series of narrow black and white lines. As the spatial frequency increased so too did the maximum rate of change of image intensity.

Because the intensity modulations in the above images have sinusoidal profiles the mathematical expressions for the modulation are particularly simple. In fact the choice of sinusoidal intensity modulations in these introductory images is very deliberate. They illustrate the idea of describing changes of intensity in space with a sine wave. We could also make up simple images in which the intensity change was most simply described as the cosine of a spatial frequency and the position in space. A more complex image might be described by a mixture of sine and cosine terms.

This leads us to a VERY important principle of image processing. *No matter what the intensity profile of an image might be it is possible to describe it as the sum of a collection of sine and/or cosine waves of different frequencies and amplitudes.* Before we introduce this idea more formally, let's consider how a collection of sinusoids can describe an intensity profile that seems very different from a sinusoid – Fig. 4.5.

Fig. 4.4 The profile of this image can be described as the sum of two sinusoids of spatial frequencies 1 and 20 cycles per image width: $I = \sin(\omega x) + \sin(20\omega x)$. In this example both the spatial frequencies have equal amplitude. Notice that the higher spatial frequency component resembles the distinct edge detail seen when black lines are drawn on a white or gray background





Fig. 4.5 The intensity profile of this image does not resemble a sinusoid, yet, as the following images demonstrate, it can be accurately described as the sum of a long series of sinusoids of increasing spatial frequency and decreasing amplitude

Here the intensity profile has a step-like change from white to black exactly midway across the image. Although the profile looks nothing like a sinusoid we will soon see that it can be represented as the sum of a large number of sinusoids of progressively increasing frequency.

The first step in the process is to find a *single* low frequency sinusoid that gives a rough approximation to the step shaped profile. For this particular example the sinusoidal profile of Fig. 4.2 is just what we are looking for – it is bright or white in its left half and dark or black in its right half – not a very accurate representation of Fig. 4.5, but not too bad either. How can we improve it? Looking at the sinusoidal profile we could say that we need to flatten out the top of the white hump and the bottom of the black trough. We can do this by adding another sinusoid of *exactly* three times the frequency of the original as shown in Fig. 4.6. To achieve the optimum flattening of the white hump and the black trough the higher spatial frequency component has to have one third of the amplitude of the original. We still




Fig. 4.6 The square intensity profile of Fig. 4.5 can be approximated by adding two sinusoids of frequency ω and 3ω , where $\omega = 1$ cycle per image width. Addition of the higher frequency term flattens the top and bottom of the profile and steepens the sides. The best approximation to the square profile is obtained when the amplitude of the higher frequency is $\frac{1}{3}$ of the amplitude of the lower frequency




Fig. 4.7 Five different sinusoid-based profiles that approximate Fig. 4.5. As each successively higher spatial frequency term is added the accuracy of the approximation improves

a $I = \sin(\omega x)$

b $I = \sin(\omega x) + \frac{1}{3}\sin(3\omega x)$

c $I = \sin(\omega x) + \frac{1}{3}\sin(3\omega x) + \frac{1}{5}\sin(5\omega x)$

d $I = \sin(\omega x) + \frac{1}{3}\sin(3\omega x) + \frac{1}{5}\sin(5\omega x) + \frac{1}{7}\sin(7\omega x)$

e $I = \sin(\omega x) + \frac{1}{3}\sin(3\omega x) + \frac{1}{5}\sin(5\omega x) + \frac{1}{7}\sin(7\omega x) + \frac{1}{9}\sin(9\omega x)$

have humps and troughs but their depth, or magnitude, is smaller than in our first, single sinusoid approximation to Fig. 4.5.

Once again we can improve the fit by flattening humps and troughs with higher frequency terms. The first five steps of this process are illustrated in images and profiles in Fig. 4.7. As we progressively add more terms of higher frequency we achieve an ever closer approximation to the step-shaped profile of Fig. 4.5. In *this particular case* (two bands of identical width) each new higher frequency term is an odd multiple of the first single-cycle approximation, and each new term has a smaller amplitude than the previous one. As we add each new term the ‘roundness’ of the intensity profile decreases and so too does the amplitude of the ripples. Notice that as each odd frequency term is added the spatial position of the mid gray intensity level remains unchanged and coincides exactly with the black/white step and the edges of the original image.

In this process the choices of frequency and amplitude are not arbitrary – they are each chosen to give the best possible approximation to the original step shape. In this particular example the spatial frequency terms are *odd harmonics* - exactly analogous to the harmonics observed in a vibrating string.

Of course the rectangular ‘square wave’ profile seen in Fig. 4.5 would be highly unusual in a medical image. To describe intensity profiles seen in conventional images we would need more frequencies than just the odd harmonics. In fact, even Fig. 4.5 is a special case because the white to black transition occurs precisely midway across the image. If it did not then the odd harmonics alone would not be sufficient for an accurate representation – as we shall see later. (Note: Don’t confuse the frequencies and amplitudes we used to represent of Fig. 4.5 with those that describe Figs. 4.3 and 4.4 which have quite arbitrary spatial frequency components.)

We have just examined some simple images and seen how we could describe their intensity profiles as a single sinusoid or the sum of a series of sinusoids of particular amplitude and spatial frequency. We only looked at the profile in the x direction as these images had no variations in intensity in the y direction. Most images aren't so simple – the intensity is likely to be highly variable in both the x and y directions, so for a proper description of an image's spatial frequency components we need a 2D method. There are several to choose from. In this text we will concentrate on the Fourier Transform because of its versatility and wide usage in image construction and processing. However, the 2D Fourier Transform will be a little easier to relate to the discussion we've just had if we first glance over the Cosine and Hartley Transforms. Think of it as learning to wrestle a salt water crocodile by first playing with a couple of goannas.

4.2.4 The Cosine and Hartley Transforms

The *Cosine Transform* is a rigorous method for doing what we have just done by intuition with the main difference being that it describes a profile as the sum of a series of cosine terms. Just for completeness, and not for memorization, here is the mathematical description of the discrete 1D Cosine transform:

$$F(\omega) = \sqrt{\frac{2}{N}} \sum_{x=0}^{N-1} f(x) \cdot c_{\omega} \cos\left(\pi \frac{\omega(2x+1)}{2N}\right) \quad (4.4)$$

This equation says that if we have a digital image matrix N pixels wide then we can accurately describe the x profile of any row of the image matrix by adding together N different spatial frequencies. An analogous *Sine Transform* also exists, and although it more obviously resembles the intuitive process we followed in the last section, it is rarely used in image processing because it has mathematical disadvantages that make the Cosine Transform more useful.

Since we are dealing with digital images which are 2D arrays of discrete intensities the relevant implementation is the 2D *Discrete Cosine Transform* or DCT. The DCT is widely used in image processing – most commonly as a central part of the JPEG image compression method in which it is used to identify low amplitude high spatial frequency components that can be discarded because they have low visibility to humans.

In contrast to the Cosine Transform the *Hartley Transform* describes an intensity profile using both sine and cosine terms:

$$F(\omega) = \sum_{n=0}^{N-1} x_n \left[\cos\left(\frac{2\pi n\omega}{N}\right) + \sin\left(\frac{2\pi n\omega}{N}\right) \right] \quad (4.5)$$

Using both sine and cosine terms may seem like an unnecessary complication but we introduce it here because of the close relationship between the Hartley Transform and the Fourier Transform. In ImageJ (the image processing software used to illustrate examples in this text) the Fourier Transform is computed from a Hartley Transform.

4.3 Fourier Transforms and Fourier Spectra

With the preceding introductory background we can now discuss the use of the Fourier transform (FT) to resolve the spatial frequency components of 2D images.

4.3.1 1D Fourier Transforms

Put simply the Fourier transform states that any periodic function can be expressed as the sum of an infinite series of sines and cosines:

$$F(\omega) = \int_{-\infty}^{\infty} f(x) (\cos(2\pi\omega x) - i\sin(2\pi\omega x)) dx \quad (4.6)$$

This, just one of several possible expressions for a 1D Fourier Transform, is included here mainly for the sake of completeness. You don't have to remember it, nor understand it in detail, to understand its use in image processing. Since an image is two dimensional we need to apply Fourier transforms in both directions, a complication that will be described soon. For now we need only consider 1D transforms.

A major difference between the Fourier Transform and the Hartley Transform is that the FT is *complex* – it includes the *imaginary* number $i = \sqrt{-1}$. If you are not familiar with the idea of complex numbers don't panic. For the curious there is a basic introduction to complex numbers included in Appendix C, but you don't need to know this to get a feel for what the Fourier transform is doing.

Why have we now introduced yet another headache – complex numbers? What is the point of imaginary numbers? After all, images are composed of real numbers – the intensities of pixels.

If we were *only* interested in processing images to determine their spatial frequency components then the Cosine or Hartley Transforms would indeed be sufficient. However, image processing and image construction needs more versatility than this. Although you might argue that imaginary numbers don't exist (you can't count or measure objects with them), the complex number formalism is indispensable. We need complex numbers and Fourier Transforms to make MRI images and they are very handy for image analysis. Take a look at Appendix C for some more examples.

The following discussion deals mainly with the *magnitude* or *absolute value* of the Fourier Transform. The FT of a purely real function (one that does not involve multiples of i) such as a digital image, has complex terms, but we need not concern ourselves with this technicality just yet.

Working in the opposite direction to the Fourier transform above, we can synthesize an arbitrary periodic function by *Inverse Fourier Transformation* of its frequency components:

$$F(x) = \int_{-\infty}^{\infty} f(\omega) (\cos(2\pi\omega x) + i\sin(2\pi\omega x)) d\omega \quad (4.7)$$

This is very similar to the forward or direct Fourier Transform, the difference being only the sign of the second term.

Equations 4.6 and 4.7 are descriptions that apply to *continuous* functions and are unsuitable for digital data which is *discrete*, not continuous, and *finite*, not infinite. To deal with discrete finite data sets such as digital images we use a modified version of the Fourier transform – the *Discrete Fourier Transform* (DFT):

$$F(\omega) = \sum_{x=0}^{N-1} f(x) \left(\cos\left(\frac{2\pi\omega x}{N}\right) - i\sin\left(\frac{2\pi\omega x}{N}\right) \right) \quad (4.8)$$

In this expression x represents the pixel position where there are N pixels in one row or column of the image, and ω represents a specific spatial frequency. Notice that in the discrete Fourier Transform the integration from $-\infty$ to $+\infty$ has been replaced by a summation from 0 to $N - 1$ because we only have a finite set of data to deal with.

There is also a discrete form of the inverse Fourier Transform:

$$F(x) = \frac{1}{N} \sum_{\omega=0}^{N-1} f(\omega) \left(\cos\left(\frac{2\pi\omega x}{N}\right) + i\sin\left(\frac{2\pi\omega x}{N}\right) \right) \quad (4.9)$$

Using the DFT and inverse DFT we can convert an image back and forth between the spatial and frequency domains as many times as we like *without any loss of information*. The transforms are precise descriptions, not approximations. The ability to use such transforms without fear of data loss permits the processing of images in either the spatial or frequency domains. The only limitation to this claim is the internal precision of a computer's calculations. These are more than adequate for most image processing tasks.

The summations expressed in the DFT formulae can easily be performed by a computer. The DFT is still a complex function, and you might wonder how the computer represents $\sqrt{-1}$. It doesn't have to. We only need to store and manipulate the coefficients of the imaginary terms and those coefficients are real numbers. This, believe it or not, is how an anatomical image is constructed from raw MRI data in which the real and imaginary coefficients represent voltages measured on different axes of the MRI scanner.

The DFT has a special characteristic which is very important in image processing. *The DFT treats its finite set of discrete input data as if the data repeated itself infinitely.* For a digital image this means the image data is treated as if the image were tiled infinitely in space – a concept that will be illustrated shortly.

4.3.2 2D Fourier Transforms

For digital images (2D matrices of real numbers) we use a 2D Fourier transform to resolve the x and y spatial frequency components. The 2D FT produces a two layer matrix – one layer representing the coefficients of the real terms, and the other layer the coefficients of the imaginary terms.

Let's consider a digital image – an m column \times n row 2D matrix of pixel intensities. To perform the 2D Fourier transform of the image we can first compute the 1D DFT of each row of the image matrix and store these transformed rows in a new two layer complex matrix representing the real and imaginary coefficients. This new matrix now has both space and spatial frequency dimensions. Each row represents the series of spatial frequencies present in the corresponding row of the original image. Each column represents the coefficients of a particular x direction spatial frequency, and these may vary according to the y coordinate (row number) in the original image. Since the DFT is a summation of a finite number of terms (N in Eq. 4.8), not an infinite series, the DFT of a row of m pixel intensities has m terms representing m distinct spatial frequencies.

To complete the 2D Fourier transform we now compute the 1D DFT of each of the m columns of the complex matrix of row transforms. The output of this process is another complex matrix of spatial frequency coefficients. This matrix has spatial frequency dimensions in *both* directions – x and y . The elements of any row or column of this matrix represent coefficients of spatial frequency in the original image *but they no longer have any direct correspondence to a particular row or column in the original image.* We need the inverse Fourier Transform to decode this information.

We could also have first FT'ed the columns of the image matrix, and then the rows of the matrix of column FTs, with the same final result. The actual algorithm used in most software differs a little from this rows-then-columns method for the sake of computational efficiency. It has the same effect except that the complex FT matrix is square and has dimensions $m' \times m'$ where m' is a power of two – the $m' \times m'$ square is simply the smallest one that can enclose the original image.

4.3.3 Fourier Spectra

In Fig. 4.7 we partially resolved an image into its spatial frequency components (in one dimension). We stopped our approximation when we got to five terms, but

adding more higher frequency terms would have increased the accuracy of the approximation. We could have used a 2D Fourier transform to perform the same task, including calculation of all the higher order terms, with a result as represented in Fig. 4.8. The FT is a complex function and often produces negative coefficients for the real and imaginary terms. To illustrate these negative coefficients in Fig. 4.7 the real and imaginary parts of the complex data are represented by two special images (c, d) in which zero is represented by mid-gray rather than black. Positive coefficients are represented by lighter pixels and negative coefficients by darker pixels. While we have illustrated the real and imaginary data separately in this figure, they are not normally viewed in this way. The convention is to show the *magnitude* or *Fourier spectrum* which represents the square root of the sum of the squares of the real and imaginary coefficients of the complex matrix (for more explanation of this see Appendix C). The magnitude is thus always a positive number (we squared the coefficients, not the imaginary numbers) and is thus easy to display. This is what is shown in Fig. 4.8b. Take note of the *spectrum* part of the description ‘Fourier spectrum’. A spectrum is a display of intensity as a function of frequency – in our case,

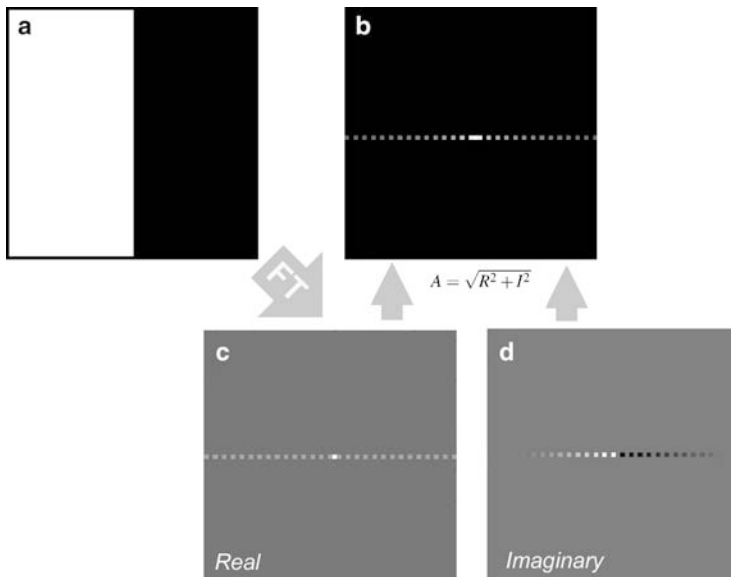


Fig. 4.8 2D Fourier transform of Fig. 4.5. The Fourier spectrum (b) represents the *magnitude* of the Real and Imaginary parts, represented here by images c and d. The Fourier spectrum amplitude $A = \sqrt{R^2 + I^2}$ where R and I are the coefficients of the real and imaginary terms. Note that in c and d mid-gray represents zero. Positive coefficients are represented by lighter pixels and negative coefficients by darker pixels. For *this particular* image all the real coefficients are positive. The Fourier spectrum of a purely real matrix, such as a digital image, is always symmetrical – even if the image is not symmetrical. Detail of the center of this Fourier spectrum is shown in Fig. 4.9




Fig. 4.9 An alternative representation of Fig. 4.8b. On the *left* the intensity of the pixels is plotted as a 3D bar chart. In a Fourier spectrum it is conventional to plot the log of the complex amplitude so that the normally large zero and low frequency amplitudes does not obscure the display of the terms with smaller amplitudes. On the *right* we see an enlargement of the center of the Fourier spectrum image. Note that the intensities of the pixels representing amplitudes of the spatial frequencies 1, 3, 5, 7, etc. *decrease* as frequency increases – just as we saw in Fig. 4.7 (the height of the bars, and corresponding pixel intensities, decrease slowly because it is the log of the amplitude that is plotted). The bright zero frequency point at the center of the Fourier spectrum represents (but is not equal to) the average of all pixel intensities in the original image. Because it has the largest amplitude it is always shown as a white pixel in the Fourier spectrum

spatial frequency. This should be a reminder that the x and y axes of the Fourier spectrum image represent x spatial frequency and y spatial frequency, NOT spatial distances.

Figure 4.9 shows two alternative views of the Fourier spectrum. The 3D bar chart is intended to emphasize that the pixel intensities in the Fourier spectrum represent the amplitudes of the individual spatial frequency terms. Notice that the first five on either side of the center correspond exactly to the frequencies and amplitudes of our intuitive approximation (Fig. 4.7). The height of the bars, and corresponding pixel intensities, do not drop off in the sequence $1, \frac{1}{3}, \frac{1}{5}, \frac{1}{7}, \dots$ because it is the log of the amplitude that is plotted.

There are some important points to note about the Fourier spectrum display:

- The origin (‘zero frequency point’, see below) lies by convention at the center of the FT image matrix. The image is symmetrical because in the Fourier spectrum of an image there is no difference between positive and negative frequencies.
- *For this example* all the non-zero data lies along a central row at right angles to the black/white edge in the original spatial domain image.
- *For this example* the non-zero data appears in the first and then every second pixel as we count outwards from the central pixel. These are the odd harmonics mentioned previously.
- The pixel intensity decreases with increasing distance from the center. This represents the progressively decreasing amplitude of the higher order harmonics, also seen in Fig. 4.7.

- *For this example* there are no non-zero elements in the vertical direction apart from the central row. This is because there is no modulation of the intensity in the vertical direction, i.e. for any column of the image matrix for Fig. 4.5 all the elements are identical.

In general, if the dimensions of a digital image are $n \times n$ pixels, then there will be n terms (i.e. n discrete spatial frequencies) in each dimension of its Fourier transform. Many of these spatial frequencies may have zero amplitude. The equations in the caption of Fig. 4.7 only show the non-zero terms.

4.3.4 The Zero Frequency or ‘DC’ Term

What about the central white pixel we see in all the Fourier spectra? As mentioned above, the center of the FT matrix represents the ‘zero frequency’ and is often called a DC term because of its *direct current* (DC) equivalent in electrical signal processing. It represents a signal that does not vary with time or space. In the Fourier transform of an image the DC term can be thought of as representing the average intensity value of the whole image. The only image that will have a zero amplitude for the zero frequency is a completely black image.

The real explanation for the zero frequency term is hidden in Eq. 4.2. Remember that the value of sines and cosines range from -1 to $+1$ but in real images the pixel intensities are all positive numbers or zero. To keep our mathematical description of image profiles simple we temporarily adopted the convention of $-1 = \text{black}$ and $+1 = \text{white}$. Going back to the $0 = \text{black}$ and $255 = \text{white}$ convention for an 8 bit image we need to add a constant to the sum of all the sine terms to make sure the total is non-negative. In Eq. 4.2 we did this by adding 127.5 to the sine term that describes the profile of Fig. 4.2. Notice that 127.5 is the average pixel intensity. Thus the DC term represents a signal that does not vary with space and its effect is to ‘offset’ the output of the non-zero frequency terms.

Because the amplitude of the zero frequency term is usually very much higher than the amplitude of any of the other frequency terms it is common to plot the log of the amplitudes to display an image of the Fourier spectrum. This makes it easier to see variations in the amplitude data.

4.3.5 Fourier Spectra of More Complex Images

In Fig. 4.10 we see a slightly more complicated version of Fig. 4.5 and its Fourier spectrum. At first sight this Fourier spectrum looks identical to the one shown in Fig. 4.10b, but if we look closely we see that in Fig. 4.10b the pixels immediately on either side of the central (zero frequency) pixel are black. There is a good reason for this. The first non-zero term in the sinusoidal approximation to the intensity profile of Fig. 4.10a has twice the frequency of that for Fig. 4.5 because in Fig. 4.10a




Fig. 4.10 Another simple image and its Fourier spectrum. For this image the spatial frequency ± 1 cycle per image width has amplitude zero – it is not useful in describing the x -direction profile. This profile is best described by summation of spatial frequencies $\pm 2, \pm 6, \pm 10, \pm 14$, etc




Fig. 4.11 A more complex image and its Fourier spectrum. For this image the spatial frequencies from ± 1 to ± 7 cycles per image width have amplitude zero – they are not useful in describing the x -direction profile. This profile is best described by summation of spatial frequencies $\pm 8, \pm 24, \pm 40$, etc. For clarity image **b** shows an enlargement of the central part of the Fourier spectrum

we have two black and two white bands. The amplitude of the spatial frequency 1 cycle/image width is zero. If we increase the number of black and white bands further (Fig. 4.11) the first non-zero terms in the FT occur, as we now expect, further from the center. We still have a ‘square wave’ intensity profile so the next non-zero frequency is the next odd harmonic of the ‘primary’ frequency.

Let’s look again at why there is only one row of non-zero data in these Fourier spectra. Remember that to perform the 2D FT we first transformed all the *rows* of the original spatial domain image and then transformed all the *columns* of the intermediate row-FT matrix. Since all the rows of the original image are identical to each other, all the rows of the intermediate row-FT matrix will be identical to each other. Thus all the elements in any single *column* of the row-FT matrix will be identical to each other. When we transform these columns of the intermediate matrix we produce only DC values because all the elements in any single column are identical. The FT of a series of constant non-zero terms has just one non-zero term, the DC term. In the Fourier spectrum the y direction DC term lies midway




Fig. 4.12 Most images have intensity changes in both the x and y directions. In this image we have a y -direction spatial frequency of exactly 8 cycles per image height, and an x direction frequency of exactly 2 cycles per image width. For clarity image **b** shows an enlargement of the center of the Fourier spectrum. The size of the original spatial domain image matrix (**a**) was 128×128 . If the full 128×128 Fourier spectrum were shown in **b** it would be difficult to see the details that are very close to the center

between the top and bottom. Of course this lack of spatial frequency modulation in the y direction is also evident in the original image.

Now let's look at an image (Fig. 4.12) in which we have intensity changes in both the horizontal and vertical directions. Once again we have an intensity profile that is sinusoidal - two cycles per image width in the horizontal direction and 8 cycles per image width in the vertical direction. We see that the non-zero terms in the FT are displaced 2 pixels from the center horizontally and eight pixels vertically.

It should now seem reasonably obvious why we have the non-zero elements in the central row and central column of Fig. 4.12b, but what about the extra four white spots that have appeared off the center lines? We can understand the origin of these if we look at the 'halfway point' of the 2D Fourier transform process - the row-FT matrix. Figure 4.13 illustrates the stepwise creation of the Fourier spectrum of Fig. 4.12a. Since each of the Fig. 4.12a rows (except the ones that are completely black) has a sinusoidally modulated intensity with spatial frequency two cycles per image width the 1D FT has non-zero elements 2 pixels either side of the center. The completely black rows have 1D FTs that are all zeros. If we complete the 2D FT process by performing a 1D FT on every column of Fig. 4.13b (strictly speaking we mean on the complex matrix whose magnitude is Fig. 4.13b) we would find that, because only the central column and the two columns two pixels either side of the center have non-zero data, there will be only these same three columns with non-zero data in the final Fourier spectrum. Furthermore, this non-zero data in the intermediate matrix, as can be seen in Fig. 4.13b, has a sinusoidal profile with frequency eight cycles per image height. Thus in the final Fourier spectrum (d) we get non-zero elements eight pixels above and below of the center - exactly as we saw in Fig. 4.12b.

On the basis of the discussion so far we might assume that the Fourier spectrum of an image becomes more complicated as the complexity of intensity changes in




Fig. 4.13 The 2D Fourier transform can be formed by stepwise 1D Fourier transformation. Either the rows or the columns can be transformed first to create intermediate ‘partial Fourier spectra’ (**b** and **c**). The coordinates of images **b** and **c** thus represent *spatial frequency* in one direction, and *space* in the other direction

the image increases, however, this is not always the case. Consider the image and Fourier spectrum shown in Fig. 4.14. Although this image appears quite simple in comparison with Fig. 4.12a its Fourier spectrum has so many non-zero terms that it looks like a white blur.

The explanation of this lies in our earlier statement that when we perform a 2D Fourier transform on an image the image is treated as if it were tiled infinitely in all directions. For the images we have discussed so far this tiling does not result in any change to the modulation of the intensity that we see within the image itself. Figure 4.14a is different from the previous images in this respect. Here we have a spatial frequency of 0.5 cycles per image width and 2.5 cycles per image height. Since the intensity modulation is not an integer number of full cycles in either direction the intensity profile of the infinite tiling does not have a regular sinusoidal profile. Figure 4.15 shows the tiling effect with just four copies of Fig. 4.14a. There are sharp discontinuities in the sinusoidal profile at the joins between tiles.




Fig. 4.14 An apparently simple image may have a complicated Fourier spectrum containing non-zero amplitudes for a large number of spatial frequencies. This is because the Fourier transform treats an image as if it were tiled infinitely in all directions. Tiling of this particular image, as shown in Fig. 4.15 below, results in sharp discontinuities in the sinusoidal profile at the joins between tiles




Fig. 4.15 Partial tiling of Fig. 4.14a. The Fourier transform treats an image as if it were tiled infinitely in all directions. When Fig. 4.14a is tiled we find that the profiles in both the x and y directions feature sharp discontinuities where edges of the original image join. These discontinuities can only be described by the sum of a very large number of spatial frequencies, as evident in the Fourier spectrum (Fig. 4.14b)

The FTs of these ‘corrupted’ sinusoids contain a large number of non-zero amplitudes of many frequencies in order to account for the discontinuity in the intensity profile. The Fourier spectrum (Fig. 4.14b) looks like a white blur because many spatial frequencies have non-zero amplitudes.

Now consider Fig. 4.16a which looks exactly the same as Fig. 4.5 but has a very different Fourier spectrum – it appears to have no non-zero terms in the central row. This tells us something about Fig. 4.16a that we probably can’t see with the naked eye. The black and white stripes are not exactly the same width (the black stripe in the original image was 33 pixels wide and the white stripe 31 pixels). Now we can no longer approximate the intensity profile with a series of odd harmonics because the black/white step does not occur precisely at a null (zero) point of the odd




Fig. 4.16 A slightly asymmetrical version of Fig. 4.5 and its Fourier spectrum. In this image the *white band* is slightly narrower than the *black band*. The Fourier spectrum shows non-zero amplitudes for *all* x -direction spatial frequencies. Unlike Fig. 4.5, this image cannot be described by the odd spatial frequencies alone




Fig. 4.17 An image of the letter ‘A’ and its Fourier spectrum. Each straight edge in the spatial domain image gives rise to a linear feature (a long line of non-zero spatial frequency amplitudes) in the Fourier spectrum. The Fourier spectrum feature is always oriented perpendicular to the straight edge it describes in the original image. This is demonstrated more clearly for the ‘letter A’ image in Fig. 4.18

harmonics. In fact we need both sine *and* cosine profiles. This example demonstrates how examination of the Fourier spectrum can sometimes tell us things about the spatial domain image that are difficult or impossible to perceive directly.

Applying what we have learnt about interpretation of Fourier spectra we can now explain the features of Fig. 4.17. Because the edges of the letter are very sharp there are a large number of non-zero terms in the Fourier spectrum. Each of the straight lines that form the letter ‘A’ give rise to a blurred white line in the Fourier spectrum, and this line lies perpendicular to the edge from which it originated in the spatial domain image.

Just as the spatial domain image is treated by the 2D Fourier transform as if it were tiled infinitely in space the frequency domain data is also, effectively, infinitely tiled in frequency space. This means that the long diagonal features in the Fourier spectrum are ‘wrapped’ at the edges – a feature that extends to the edge of the Fourier spectrum continues, with decreasing intensity, from the opposite edge. This characteristic appears more obvious when we look at the Fourier spectra of the individual parts of the letter ‘A’ image (Fig. 4.18). Most importantly, note that the




Fig. 4.18 Fourier spectra of the individual parts of the letter 'A' image. Each long diagonal feature in the Fourier spectrum 'wraps' to the opposite side of the spectrum. Note that the mathematical processes represented in this diagram are essentially identical to those in Fig. 4.1. This diagram illustrates *2D spatial frequency* spectra and their corresponding *spatial* domain signals. Figure 4.1 illustrates *1D temporal frequency* spectra and their corresponding *time* domain signals

extended Fourier spectrum feature does not ‘bounce’ or ‘reflect’ at the edge of the Fourier spectrum. It is merely the symmetry of the letter ‘A’ image that leads to this illusion. As we will see later, this ‘data wrapping’ phenomenon is important in the formation and interpretation of some MRI artifacts.

If we remove the sharp edges from the spatial domain image by application of a blurring filter we might get something similar to Fig. 4.19a. Now we find that the Fourier spectrum loses most of its high frequency terms – their amplitudes are zero or very close to zero. The only significant amplitudes lie in a cluster around the central zero frequency. The blurred image has no rapid changes in its intensity profile – there are no high spatial frequencies in this image.

Conversely, we could zero the low frequency data from the FT of Fig. 4.17b as shown in Fig. 4.20a, where the black circle represents the area in which the complex FT data has been set to zero. Now when we perform an inverse FT on the edited frequency domain data we form the spatial domain image shown in Fig. 4.20b. The result is that we have only edge detail. The tonal detail, the white infill of the letter A, is lost.




Fig. 4.19 A blurred version of the ‘letter A’ image and its Fourier spectrum. There are no rapid changes in the intensity profile so only low spatial frequencies have non-zero amplitudes




Fig. 4.20 Edited FT data from Fig. 4.17b (a). All of the low frequency amplitudes under the black circle have been zeroed. Inverse Fourier transformation of this edited data produces a spatial domain image (b) that contains only edge information

4.3.6 How Many Spatial Frequencies are Needed?

2D Fourier spectra are images in which the pixel intensities represent the amplitudes of complex spatial frequency data. By convention the zero frequency point is plotted in the center of the Fourier spectrum. The pixels immediately adjacent to the central pixel represent the amplitudes, in the x and y directions, of the spatial frequency 1 cycle per image width. The next pixels represent the amplitudes of spatial frequency 2 cycles per image width, and so on. The amplitudes of the highest spatial frequencies are represented by the pixel intensities at the edges of the Fourier spectrum. What is the highest spatial frequency available? Is it high enough to describe the most sudden intensity change in the image?

Imagine that we have a square digital image which has dimensions $m \times m$ pixels. The Fourier spectrum of this image will have the same dimensions so the *maximum* spatial frequency will be $f_{max} = \frac{m-1}{2}$ cycles per image width (we have to subtract 1 because one pixel is assigned to the zero frequency ‘DC’ term). For a 65×65 pixel image this means the highest spatial frequency in the Fourier spectrum is 32 cycles per image width. Is this high enough?

Now imagine the highest possible rate of intensity change in a 65×65 pixel image. This would be when adjacent pixels in the image had the maximum possible intensity difference – such an image might look like Fig. 4.21, which is made up of alternating rows of black and white pixels. The total number of *line pairs* is 32.5. A *theoretical* way to describe this pattern with sinusoids would be to start with the spatial frequency 32.5 cycles per image width and add its odd harmonics, starting with $3 \times 32.5 = 97.5$. But we don’t have *any* of these spatial frequencies available! Our highest is 32.

The solution is to use a sinusoid of frequency 32 and its *lower* frequency neighbors. This time the required amplitudes, as shown in Fig. 4.21, decrease as the spatial frequency *decreases*. In an $m \times m$ pixel image there can *never* be an intensity pattern denser than $\frac{m}{2}$ line pairs per image width, and this pattern can be described *perfectly* by addition of sinusoids of spatial frequency zero to $\frac{m-1}{2}$ cycles per image width.




Fig. 4.21 The maximum possible rate of intensity change in a 65×65 pixel image is 32.5 line pairs per image width (or, in this case, height). Only the spatial frequencies 0 to 32 cycles are available, but these are sufficient to perfectly describe the line pairs pattern as the sum of a series of sinusoids

Up to the limit of precision of the calculations, the 2D discrete Fourier transform is a 100% accurate description of a digital image.

4.3.7 *Fourier Spectra of Lines*

Based on the discussion so far we could reasonably come to the conclusion that the edge detail in an image is encoded in the outer high frequency part of the Fourier spectrum and the tonal detail, or shading, is encoded by the central low frequency part of the spectrum. What then would the Fourier spectrum of a narrow line look like? When we zeroed the low frequency amplitudes in Fig. 4.20 we were left with just the outline of the original bold ‘A’ shape – more or less a line drawing of the solid shape. It looks like we just need the medium and high frequencies to define a line drawing but the reality is a little more complicated.

Instead of going straight to the Fourier spectrum of a narrow line let’s look at what happens to the Fourier spectrum of our original black and white stripe image (Fig. 4.8a) as we make the white stripe progressively narrower. The effect is illustrated, in three steps, in Fig. 4.22. As the white stripe gets narrower its Fourier spectrum continues to have significant amplitudes across the range of frequencies. The original very simple harmonic pattern seen in Fig. 4.8b becomes more complex and there are actually *more* non-zero low frequency amplitudes. When the white stripe is very narrow (effectively a line) *all* the spatial frequencies have large amplitudes. In the extreme case, when the line is infinitely narrow (in a digital image this means one pixel wide), the amplitudes of all spatial frequencies are identical – the profile of the Fourier spectrum is a flat horizontal line of constant intensity.

Our original idea that we only need high frequency terms to describe a line has turned out to be inaccurate. In fact we need *the full range* of available frequencies to describe a narrow line. The reason that we didn’t need the low frequencies to show the outline of the letter A is that this is, in fact, not a very accurate outline. Close inspection of the original of Fig. 4.20b would show a series of faint gray lines running parallel to all of the white lines. These ‘ringing’ artifacts are due to the fact that we removed the low frequency information. We will discuss them more in Chapter 7.

4.4 The Complex Data Behind Fourier Spectra

Because the Fourier transform treats image data as if the image were tiled infinitely in space it is quite possible for two different images to have identical Fourier spectra. For example, Fig. 4.23a shows a modified version of Fig. 4.8a in which the black region has been moved from the right side to the center of the image. The Fourier spectra of the original and the modified image are identical. In this example we moved the black region but we didn’t change its width – it remains the same as the




Fig. 4.22 Progressive narrowing of the white stripe of Fig. 4.8a changes the simple harmonic pattern in the Fourier spectrum. When the white stripe is just one pixel wide *all* x direction spatial frequencies have identical amplitude. In the case, not shown here, of an image containing only a single non-zero pixel intensity (in other words a point) all x and y direction spatial frequencies have identical non-zero amplitude, and thus all the pixels in the Fourier spectrum have identical intensity

total width of the two white regions. So long as the black and white areas have the same width, no matter where we put the black stripe the infinitely tiled patterns will appear identical. The spatial frequency amplitudes needed to describe the patterns will be identical so the Fourier spectra will also be identical.

The information that describes the spatial *position* of the image data is encoded in the difference between the coefficients of the real and imaginary data. Comparing the representations of these coefficients in Fig. 4.8c & d with Fig. 4.23c & d we see that these parts are not identical.

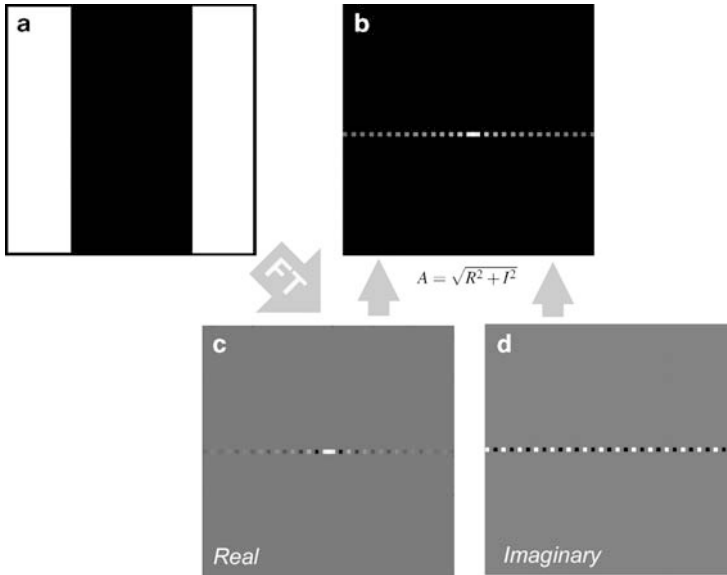


Fig. 4.23 Shifting the position of the black stripe in Fig. 4.8a, so that it lies in the middle of the image rather than at the right side, has no effect on the Fourier spectrum so long as the width of the black and white areas remain identical. This is because there is no change to the pattern that is formed when the image is tiled infinitely. However, the *position* of the stripe is encoded in the complex data (images **c** and **d**). Note the difference between the real and imaginary data in this figure when compared with Fig. 4.8c and d

The difference between the coefficients of the real and imaginary data that we are interested in is called the *phase* angle (ϕ):

$$\phi = \tan^{-1} \left(\frac{I}{R} \right) \quad (4.10)$$

When a complex number is plotted as a vector in a complex plane diagram the phase angle is the angle between the positive real axis and the vector. The magnitude, what we show in the Fourier spectrum, is the length of the vector. See Appendix C for a graphical illustration.

The forward and inverse Fourier transforms can operate on complex data expressed in either real/imaginary format or magnitude/phase format. The two formats contain equivalent information and can be interconverted (Fig. 4.24).

When displaying representations of complex frequency domain data most of the time we can safely ignore the phase information and just look at Fourier spectra. This is not to say that phase data is unimportant and can be discarded. It cannot, but we *mostly* let the computer deal with it and don't bother to create displays of it. Figure 4.25 illustrates the effect of applying the inverse Fourier transform to only the magnitude data, or only the phase data when attempting to reconstruct Fig. 4.8a from its complex frequency domain equivalent.

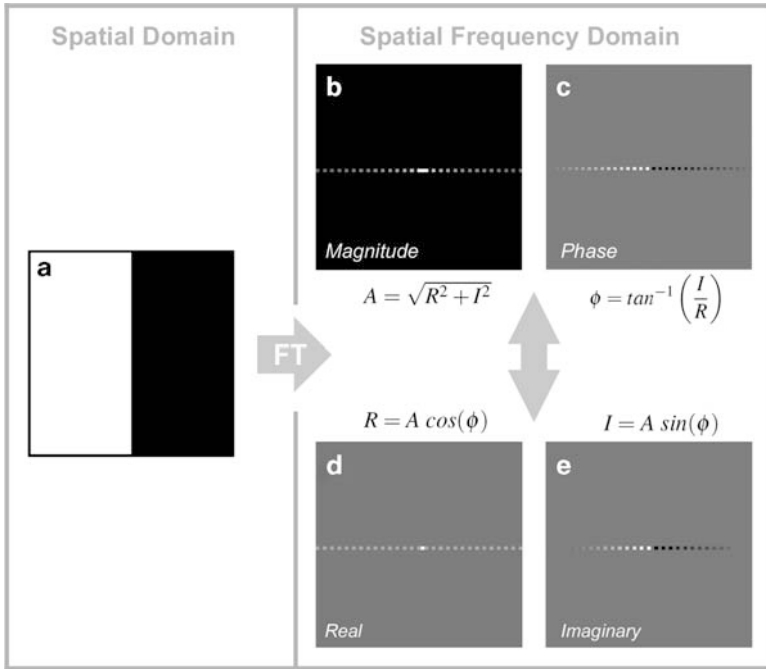


Fig. 4.24 The forward and inverse Fourier transforms can operate on complex data expressed in either *Real + Imaginary* format, or *Magnitude + Phase* format. The two formats contain equivalent information and can be interconverted




Fig. 4.25 When an image is converted to its frequency domain equivalent by Fourier transformation both real and imaginary components are created. To recreate the spatial domain image by inverse Fourier transformation both the real and imaginary data are required. Image **a** here illustrates the result of inverse transformation of just the magnitude data from the Fourier transform of Fig. 4.8a. Image **b** illustrates the result of inverse transformation of just the phase data

Phase information is critical in the creation of magnetic resonance images – *phase encoding* is routinely used to describe spatial position in at least one dimension of an MR image. Raw MRI data always contains matrices of real and imaginary coefficients. Inverse Fourier transformation of the raw data creates *another* complex data set and it is the *magnitude* of this data that forms the spatial domain anatomical image. Remember, however, that although we display magnitude data this is *not* a Fourier spectrum – the coordinates now are space, not frequency.

Notice that the image Fig. 4.25b created from phase data mainly comprises a central white line – at the position of the boundary between the black and white regions in the original image. So the phase image is something like an image of the edges in the original data. This information can be used to identify tissue boundaries and measure the velocity of fluid (e.g. blood) flow in MRI.

The relationships of image data converted between the spatial and frequency domains by the Fourier transform are summarized in Fig. 4.26. In this diagram complex data is described in terms of magnitude and phase. It could alternatively be described in terms of real and imaginary coefficients. The descriptions are equivalent provided that the appropriate versions of the forward and inverse Fourier transforms are applied.

The Fourier transform treats *all* data as if it were complex. When a Fourier transform is used to convert an image to its spatial frequency domain equivalent the image data is treated as if either: (a) it is in Real/Imaginary format and all the

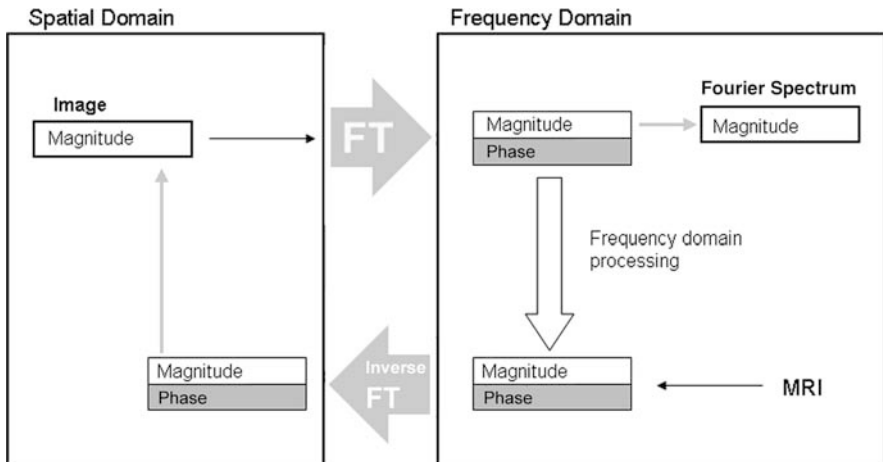


Fig. 4.26 The relationships of image data converted between the spatial and frequency domains by the Fourier transform. A spatial domain image is comprised of magnitude (intensity) data only. Fourier transformation of this magnitude data produces a complex frequency domain data set that can be expressed in terms of magnitude and phase, as shown here, or real and imaginary coefficients (not shown). The Fourier spectrum is a display of the magnitude part of the frequency domain data. Inverse Fourier transformation of complex frequency domain data produces a complex spatial domain data set. We make an image from the magnitude of this complex data. In MRI raw data is acquired in the spatial frequency domain. An MR image is a magnitude display of the inverse Fourier transform of the raw data

imaginary coefficients are zero; or (b) it is in Magnitude/Phase format and all the phase angles are zero. In the former case the pixel intensities represent the real coefficients, and in the latter they represent the magnitude. These two representations are in fact identical.

4.5 Two Practical Applications of Fourier Transforms

So far our discussions of the concept of spatial frequency, and transformations between the spatial and spatial frequency domain, have been illustrated with deliberately simple artificial images rather than medical images. We looked at simple images so that the discussion would not be sidetracked or confused by too many effects being present simultaneously. Now that we have a basic understanding of how the Fourier transform crocodile behaves, and hopefully it is a little less scary, we can look at what happens when we toss a medical image into the enclosure.

4.5.1 *How Does the Focal Spot of an X-Ray Tube Affect Image Resolution?*

One of the determinants of X-ray and CT image quality is the physical size of the source of the X-rays – the focal spot on the anode of the X-ray tube. A large focal spot will produce more blurry images than a small spot. A lot of engineering energy has gone into development of X-ray tubes that will produce very intense X-ray beams from the smallest possible focal spot. Nevertheless, since about 98% of the electron beam energy gets turned into anode heat rather than X-ray photons, there is always a compromise between small focal spot size and beam intensity. Many X-ray machines permit the adjustment of focal spot size according to the resolution and speed requirements of the imaging study. What effect does the size and shape of the focal spot have on the spatial resolution of an X-ray image? We can quantify the answer to this question by careful measurement and use of a Fourier transform.

We know the size and shape of the focal spot affect the sharpness of images but how do we measure this? We start with an image of the focal spot – easily produced with a pinhole camera. In this case the pinhole is a very small hole in a sheet of lead and we form the image on a high resolution sensor which could be a digital flat panel or a low speed film (without a fluorescent intensifying screen that would introduce extraneous blur). Suitably magnified we would expect to see the image looking something like Fig. 4.27a. The spot image is not a neat circular blob but an irregular blurred rectangle with two bright edges – images of the X-ray tube filament (a long coil of wire) that result from imperfect focusing of the electron beam. Since the spot image is a rectangle it should be intuitively obvious that the amount of image blur due to the finite size of the spot is going to be greater along one axis of the imaging plane than it is along the perpendicular axis.




Fig. 4.27 A pinhole image of the focal spot of an X-ray tube (a) and its Fourier spectrum (b). The profile of the Fourier spectrum (c) shows the way contrast is lost with increasing spatial frequency in the x and y directions. Low and medium spatial frequencies, represented by the bright center of the Fourier spectrum, have high amplitude and consequently good contrast. High spatial frequencies, represented by the dark outer areas of the Fourier spectrum, have low amplitude and consequently very poor contrast. Notice that the slightly asymmetrical shape of the Fourier spectrum is consistent with the 2D profile of the focal spot image. The bright horizontal bars in image a are, effectively, two images of the X-ray tube filament that result from imperfect focusing of the electron beam onto the anode of the X-ray tube

The blurring caused by the finite sized focal spot will blur every point in all imaged objects. Put another way, the image of every point in the imaged object will be a more or less bright spot having the same size and shape as the focal spot image. Small well-defined objects, which as we now know are described by a large range of spatial frequencies, will appear in the image with low contrast. The way to quantify the loss of contrast *for each spatial frequency* is to Fourier transform the image of the focal spot and examine the Fourier spectrum – Fig. 4.27b. The Fourier spectrum has a bright central region corresponding to high amplitudes for the low and medium spatial frequencies. So low and medium spatial frequencies will, in the absence of other blurring factors such as an intensifying screen, have high amplitudes in an image produced with this focal spot. All the high frequencies in the Fourier spectrum have very low amplitude. So high spatial frequencies will have very low contrast in an image produced with this focal spot – all small well-defined objects and all sharp edges will be blurred.

The profile of a line drawn from the center of the Fourier spectrum (Fig. 4.27c) tells us the degree of attenuation of spatial frequencies in the direction of the line. The profile would be symmetrical about the zero frequency point because the Fourier spectra of images (2D arrays of real numbers) are always symmetrical. The shape of the profile of the line is called the *Modulation Transfer Function*. The MTF is used extensively in quantitative description of imaging system performance and we will discuss it more fully in Chapter 5.

4.5.2 Making Diagnostic Images from Raw MRI Data

Finally, let's look at MRI data, which is directly acquired in the spatial frequency domain. Figure 4.28 shows the Fourier spectrum of raw complex data from a




Fig. 4.28 Fourier spectrum of raw frequency domain MRI data (a), and the spatial domain diagnostic image (b) formed by inverse Fourier transformation

scan of a human brain. The magnitude of the inverse Fourier transform gives us the diagnostic spatial domain image.

Although the Fourier spectrum looks rather complicated we can interpret some of its features in terms of the resultant spatial domain image. We see plenty of signal at all spatial frequencies, consistent with the presence of both edge and tonal detail in the spatial domain image. There are a number of weak radial lines in the Fourier spectrum encoding several distinct edges in the anatomical image – the back of the skull particularly. The narrow vertical line is consistent with the intensity difference between the top and bottom edges of the image.

There is another particularly interesting feature of the spatial domain image. The subject's nose, which is just outside the field of view on the left side of the image, has appeared behind his head! This artifact, called phase wrap, occurs because the signal from the nose, although outside the field of view, is nevertheless detected and recorded in the complex frequency domain raw data. It appears on the opposite side of the image because of the effective tiling of the frequency domain data.

4.6 Summary

- A digital image is a *spatial domain* data set. The coordinates of the image data represent distance or position in space.
- The term spatial frequency describes the rate of change of intensity in an image in terms of a sinusoidal intensity profile. Any intensity profile can be described as the sum of a series of sinusoids of appropriate spatial frequency and amplitude. Any image can be described as the sum of a series of x and y direction sinusoids.

- A 2D Fourier transform converts an image into its spatial frequency equivalent in the *spatial frequency domain*. The Fourier transform converts the *real* intensity values of a spatial domain image into a symmetrical *complex* frequency domain matrix. The Fourier transform of an image is an *exact* representation of the image. No information is lost in the process of Fourier transformation.
- Complex data can be expressed in two equivalent interconvertible forms: *Real* and *Imaginary* coefficients; or *Magnitude* and *Phase* coefficients.
- The *Fourier spectrum* of a digital image is a display of the *magnitude* of the spatial frequency domain version of the image data. The coordinates of the spectrum data represent differences in spatial frequency.
- By convention, 2D Fourier spectra are displayed with the lower frequency terms in the center of the spectrum. In the middle of the spectrum is the ‘zero frequency’ or ‘DC’ term. The magnitude or amplitude of the DC term represents the average pixel intensity in the spatial domain image. Because the amplitude of the DC term is normally much greater than the amplitude of the non-zero frequencies, it is normal to display the log of the amplitudes in the Fourier spectrum.
- Lines and sharp edges in images are characterized by non-zero amplitudes of many spatial frequencies. Straight lines and edges in images give rise to linear features in the Fourier spectrum, and these features lie perpendicular to their originating lines in the spatial domain image.
- The Fourier transform effectively treats 2D data (for example an image) as if it were tiled infinitely in space. Thus discontinuities in intensity at the joins between edges of the original data will give rise to linear features in the Fourier spectrum that are oriented perpendicular to the originating edges.
- An image that contains no sharp edge detail (and no edge discontinuities when tiled) can be described with low spatial frequencies only. The Fourier spectrum of such an image will have non-zero values only near its center.
- Inverse Fourier transformation of complex spatial frequency domain data produces complex spatial domain data. We normally view only the magnitude of this complex data. Raw MRI data is an example of a complex spatial frequency domain data set.