

1

Tomography

“Your Majesties, Your Royal Highnesses, Ladies and Gentlemen, Neither of this year’s laureates in physiology or medicine is a medical doctor. Nevertheless, they have achieved a revolution in the field of medicine. It is sometimes said that this new X-ray method that they have developed — computerised tomography — has ushered medicine into the space age. Few medical achievements have received such immediate acceptance and met with such unreserved enthusiasm as computerised tomography. It literally swept the world”.

Speech by Professor Torgny Greitz at the Nobel prize ceremony in 1979 at which A. M. Cormack and G. N. Hounsfield received the prize for Physiology or Medicine with the citation,

“for the development of computer assisted tomography”.

1.1 Introduction

Very often great ideas have the irritating habit of seeming obvious once they have been explained. Computerised Tomography is no exception. Behind the expensive hardware and the several thousands of lines of computer code that comprise the modern medical scanner there sits a very simple idea, tomography. The word ‘tomography’ was constructed using the

Greek root ‘*tomos*’ meaning to slice. Computerised Tomography, be it using X-rays, gamma photons or MRI, always proceeds by a series of slicing steps that reduce a 3D volume, first to a collection of 2D sections or slices, then each slice is cross-hatched with lines to produce a collection of 1D strips. Each stage reduces the amount of information that ends up in the detectors to a level that a reconstruction algorithm can handle without ambiguity.

Although crude methods of tomography existed in the early 1950’s, it was Hounsfield’s demonstration of a practical computerised tomographic X-ray method in 1972, which marked the start of what we call the ‘Second Revolution in Medical Imaging’. In principle the method could have been carried out using the old-fashioned hand mechanical calculator, in much the same way that early X-ray crystal diffraction data were manipulated to obtain a crystal structure. In practice however the method demands the use of a digital electronic computer to collect and process the relatively large amounts of data, which contribute to a high-resolution image. The second revolution could not really get going before the introduction of small, dedicated digital computers into industry, science and medicine.

In this chapter we describe the underlying principles of tomography, largely divorced from any one of its several applications within medicine. Each technology has its own combination of special bits of physics and engineering that together create a particular image contrast mechanism. However once the data has been collected, essentially the same mathematical method is used to create the final image in each case. We start the chapter by describing the simplest approximate tomographical reconstruction method that we call simple backprojection. This can be understood, and indeed performed, using just the mathematical operations of division and addition. It turns out that, although this scheme is actually remarkably successful, given its simplicity, it is not good enough for high-resolution reconstruction and is not ever used just on its own. In order to go further we have to make use of that branch of applied mathematics called Fourier analysis. In principle this provides an exact, as opposed to an approximate route to image reconstruction. A general introduction to Fourier Analysis is given in Appendix A, and Appendix B provides a derivation of the most important result for tomography, the Central Slice

Theorem. This, in one form or another, is at the heart of all tomographic reconstruction techniques. It is hoped that this chapter, taken in conjunction with the two appendices, will provide the reader with a basic understanding of much of the common jargon of tomographic methods and a single framework for understanding the details of the individual tomographic methods that we describe in later chapters.

1.2 Simple Backprojection Reconstruction

Projections

All tomographic image reconstruction involves the three basic operations illustrated in Figure 1.1. The first step consists of carving the object into a set of parallel slices. In the second step, each two-dimensional (2D) slice is analysed in turn by acquiring a set of projections of that slice. In the final step the set of projections are recombined, according to a particular mathematical algorithm, to reconstruct an image of the 2D slice. Thus our first task is to define what we mean by a projection of a slice.

A projection can be thought of as a shadow. Consider the child's game of creating animal head shadows on the wall, by clasping the hands in a particular way in front of a light bulb. The hands are three-dimensional objects that entirely stop any light that falls directly onto them. The image on the wall, however, is the shadow cast by the outline of the hands. It is two-dimensional and thus can be made to resemble an outline of some other three-dimensional object, for example, a rabbit head with long floppy ears! The contrast in the projection is all or nothing, formed by light absorption in the hands. If we imagine a series of straight lines or rays from the light bulb to the projecting wall, some lines stop at the hands, creating darkness on the wall and some miss altogether, creating light on the wall. The boundary between the two regions is the particular outline, formed from the lines that just miss or graze the hands and then continue to the wall. If the hands are rotated, the image on the wall changes, similarly if the hands are held fixed but the position of the light bulb is changed, so the shape of the shadow or projection on the wall changes.

In a more serious context this shadow method can be used to reconstruct the 3D surface shapes of objects. Again, we set up a light bulb with an object in front of it, but this time put the object on a rotating table and use a computer to record the shape of the shadow. If we collect lots of shadows, each with the object at a slightly different table angle, then we could use the computer to reconstruct a 3D rendering of the surface of the object.

In the medical applications of tomography we want to see beneath the surface of patients and so X-rays are used. Now we obtain more complicated shadows, formed by objects inside the body, as well as the body surface. A proportion of the X-rays pass straight through, and it is these X-rays that are detected to form the shadowgraph. If we collect lots of X-ray shadows, taken at different angles, then we can reconstruct the shape of the body including details of objects on the inside. In general X-ray shadows for tomography are obtained from one particular thin slice at a time because X-rays are scattered as well as stopped by tissue. X-rays scattered through a wide angle from one region would confound the

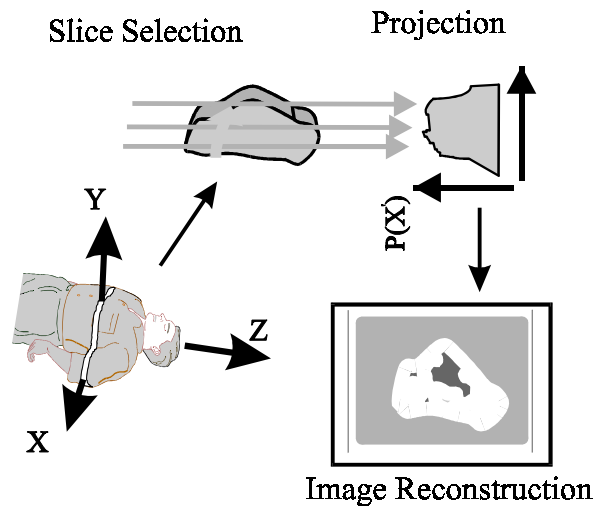


Figure 1.1 The tomographic process generally consists of three steps. First a single slice is selected. Next a complete set of projections of that slice is obtained. Finally the set of projections are recombined, using a mathematical recipe, to form a reconstructed image of the slice.

interpretation of the shadow from another. By concentrating on one thin slice at time, this confusion can be avoided. Thus the starting point for most tomographic image reconstruction is a set of shadows or projections of a particular thin 2D slice through the patient.

The Numbers in a Box Puzzle

The process of tomographic image reconstruction can be thought of as a way of solving the following puzzle. Consider a square array or matrix of positive numbers whose values we would like to know. Imagine that the array is sealed inside an opaque box so that we can't actually see it but that we have a way of finding the sums of the hidden numbers along the rows, columns and principal diagonals. In other words we have a means of obtaining a set of projections of the box of numbers.

What we want to do is reconstruct the array of numbers inside the box, using just the projections. Figure 1.2 illustrates this idea. All the medical technologies produce essentially the same hidden number puzzle. In each technology the mechanism of image contrast is of course different so that the numbers in the box (a, \dots, y) arise from very different physical mechanisms and represent very different aspects of the underlying tissue. In X-ray CT, each number corresponds to the linear X-ray attenuation

a	b	c	d	e	0
f	g	h	i	j	50
k	l	m	n	o	10
p	q	r	s	t	10
u	v	w	x	y	10
10 10 40 10 10					

Figure 1.2 The numbers in a box problem. Each set of numbers outside the box represents one projection. Each number within a projection is a sum through the box along a horizontal or vertical line

coefficient averaged over the particular little volume of tissue. In practice each slice must have a small but finite slice width and thus each of our numbers actually represents a small element of volume. With small arrays it is clearly possible to solve the puzzle, using trial and error, beginning at the corners. For the large arrays, 256×256 or bigger, used in medical tomography the guessing could go on for a very long time before we get an answer. Fortunately there is a very simple procedure that will rapidly produce an approximate answer, and an exact method that uses the properties of Fourier transforms.

Simple Backprojection

In Figure 1.2, each of the numbers outside the box is the sum of box numbers, along a particular direction, through the array. Since we are going to consider a number of projections, each one comprising a number of elements, we need a systematic labelling convention for a particular element in a particular projection. A projection is labelled P_i where $i = 1, 2, 3, 4$ etc. denotes the angle at which the projection was taken. Thus the first element, counting from the left, in Projection 1 is written $P_1(1)$, the second element in projection 1 is $P_1(2)$ and so on. In the bottom outside row of Figure 1.2 we see that:

$$\begin{aligned} P_1(1) &= 10 = a + f + k + p + u \\ P_1(2) &= 10 = b + g + l + q + v \\ &\dots\dots\dots \\ P_1(5) &= 10 = e + j + o + t + y \end{aligned}$$

and in the right-hand outside column, we have that

$$\begin{aligned} P_3(1) &= 0 = a + b + c + d + e \\ P_3(2) &= 50 = f + g + h + i + j \\ &\dots\dots\dots \\ P_3(5) &= 10 = u + v + w + x + y \end{aligned} \tag{1.1}$$

Each set of external numbers represents a single projection through the box and of course it contains information about the values and the arrangement of the hidden numbers. In the X-ray case the analogy with our puzzle is quite clear. The first number in Projection 1 is obtained by shining

a fine pencil beam of X-rays through the box from the top to the bottom along the leftmost column. The relative number of X-ray photons that survive the trip through the box is here represented by $P_1(1) = 10$. Each little volume will attenuate the beam by an amount determined by the local attenuation, so that each number in the projection represents the total attenuation through the box, along a particular line, the sum of box numbers along that line. $P_1(2)$, the second number in the first projection, is obtained by shifting the beam by one column to the right. The different projections arise from different choices of beam direction.

If we pick any box number, m , for example, it is clear that it contributes to all projections. At every angle there will be one line of sight that passes through the m cell. Since this is true of all the box numbers, we can think of the reconstruction algorithm simply as a way of solving a set of simultaneous equations. In this case we have 25 numbers and this requires 25 equations to get a solution. If we were only dealing with very small arrays then the reconstruction could proceed in this manner. In a real application, providing a spatial resolution of a few mm in a cross-section of the human torso, the size of the array will be 256×256 or larger. This would require the simultaneous solution of 2^{16} equations. Although this is by no means impossible it would be extremely time-consuming and fortunately is unnecessary.

Simple backprojection begins with an empty matrix that represents the image of the object slice and the idea is to fill in, or reconstruct, the unknown number in each cell. Consider the first number $P_1(1)$, in the first projection. During the data collection this was built up from contributions from all cells, along the line of sight that ends at this point in this projection. As a first approximation we can think of each cell along this line as contributing the same average amount to the sum. Thus we put back into each of the corresponding empty cells the number

$$\frac{P_1(1)}{5} = \frac{a + f + k + p + u}{5} \quad (1.2)$$

and repeat this for all the other numbers in the first projection. The backprojection process is repeated for the second and subsequent projections, but at each stage we add the new average numbers from the

current projection to those already in the matrix. After two projections have been backprojected, the centre cell of the reconstruction array will contain:

$$\frac{P_1(3)}{5} + \frac{P_3(3)}{5} = \frac{c + b + m + r + w}{5} + \frac{k + l + m + n + o}{5} \quad (1.3)$$

We see that the real contribution to the centre cell, m , enters twice into the temporary answer, whereas all the other contributions only enter once. After N projections we will have $Nm/5$ coming from the target cell together with a sum of other contributions. Figure 1.3 shows the reconstruction process interrupted after two stages of backprojecting the sums. As we proceed, the image gradually condenses into a reasonably good approximation to the original array.

Simple backprojection has a severe limitation as far as real imaging is concerned. Regions of the slice that contribute no signal to the projection actually end up with a definite positive intensity after backprojection. We




Figure 1.3 Reconstructing the numbers in a box. The original 'image' together with two projections is at the top left. The array, after backprojection of the vertical projection, is at bottom left. After adding the second, horizontal backprojection the array is at the top right. Even just two projections begin to reveal the 'T' in the original image.

noted above that the simple process puts back into each box a multiple of the real original number together with other contributions. These unwanted contributions are all positive and thus keep on adding to our answer as we step through the projections. This leads to an incorrect, gradual fogging of the reconstructed image that is unacceptable. A partial remedy is to calculate the average value of all the entries currently in the reconstruction matrix at each stage and subtract that from each entry. In general simple backprojection on its own is seldom if ever used, since the better Fourier method or its derivative Filtered Backprojection produces a very much more reliable result. We have dwelt on the simple process in order to show that tomographic reconstruction is basically very easy to understand and requires no mathematics beyond averaging and addition.

1.3 The Fourier Method

Fourier methods are very widely used in science and engineering, particularly in signal and image processing. One reason for their use is that difficult problems in real space can often be solved, using simple algebra, in Fourier space or K -space. The idea is very similar to that of using log tables to do long multiplication and division. It will be recalled that the logarithm of a product of two numbers is just the sum of the individual logarithms. Thus to multiply two numbers we take their logs, add the logs together then take the antilogarithm of the sum to obtain the product. Fourier methods have similar steps: the original problem is transformed into a new space, K -space, where some algebraic manipulation takes place. Finally the answer is obtained by an inverse Fourier transformation back to real space.

K -Space

A more detailed description of waves, Fourier analysis and K -space is given in Appendix A. Those readers with no background in undergraduate science are encouraged to read this before continuing with this section. The Fourier method can be summarised in words as a decomposition of an image into a set of hypothetical simple waves with different frequencies. For an image with edges $L \times L$, the smallest

frequency is given by $K_0 = 2\pi/L$. All the other waves in the set have frequencies $2K_0, 3K_0, 4K_0, \dots$. The term 'K-space' refers to a map of the amplitudes of all these hypothetical waves that can be considered to make up the image. They are laid out on a Cartesian grid with K_x , parallel to x and K_y parallel to y . The amplitude at each point can be represented by a line along Z (perpendicular to the page) or more conventionally by a colour scale. The general features of Fourier analysis are illustrated in Figure 1.4. In general, although a K -space map is entirely mathematically equivalent

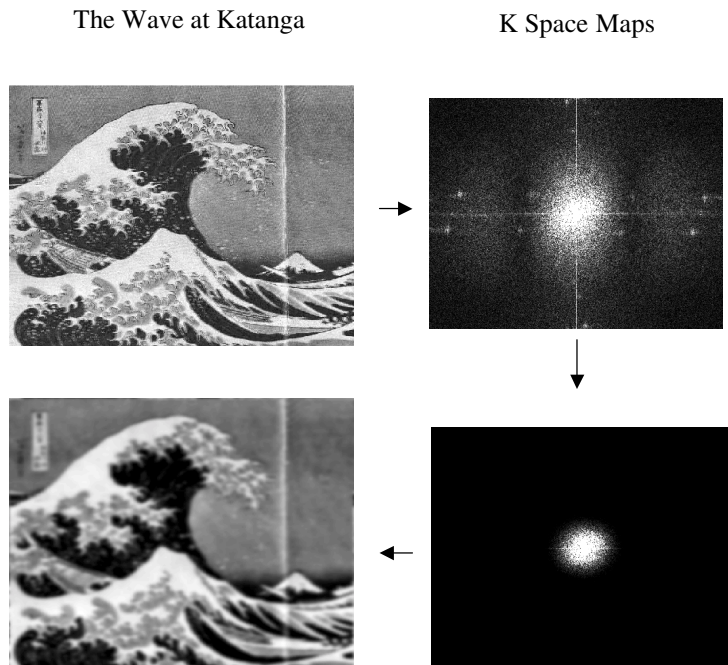


Figure 1.4 The monochrome Hokusai, in the top left panel, is Fourier transformed to produce its K -space map in the top right-hand panel. The K -space map is filtered, to reduce the amplitudes of the larger K values (shorter waves) in the bottom right-hand panel. The filtered K -space waves are then recombined to produce the filtered (here blurred) image in the bottom left-hand panel. Here we have obliterated all but a tiny central portion of the K -space map (all short waves) but still the reconstructed image is still easily recognisable as the Hokusai.

to its real parent image, its visual appearance is generally less than illuminating, as can be seen in Figure 1.4. There are however a few general rules of thumb which are useful guides to its appreciation. The region close to $K = 0$ is by far the most important part. In general, for medical images, the map has its biggest amplitude there. So that overall the K -space map generally looks like a single, roughly symmetrical peaked mesa, tailing off in height as K increases in all directions away from $K = 0$. It is here that the overall form of the slice is represented in waves. Small K corresponds to long wavelengths, the smallest being the overall extent, L of the image. Waves further from $K = 0$ represent finer and finer details within the slice.

If there were to be any roughly repeating pattern within the slice, for example like a chessboard, then the Fourier map would have little extra peaks at positions in K -space corresponding to the size of the pattern motif. Hokusai's print certainly does not contain a chessboard but the K -space map does have bright spots well away from $K = 0$. These arise from a roughly repeating texture in the paper that has been picked out by the reproduction process.

How well the Fourier method reconstructs the slice depends critically on how far out into K -space the data are accurate. Often inevitable random noise, collected along with the image data, varies rapidly with position, that is to say with small wavelengths, so that this is represented at larger K . Since the real image data is decreasing in amplitude with increasing K so finer details eventually become masked by noise. The filtering used in Figure 1.4 illustrates this. After removing Fourier amplitudes at larger K values, the resulting reconstructed image becomes blurred; it has a reduced spatial resolution. It has also lost the 'image artefacts' we associated above with paper texture.

The K -space map forms an essential intermediate step in all tomography. All of the technologies aim to get a good estimate of the K -space map of the slice under investigation, extending in K as far as noise allows and sampled as finely as is practical. MRI is deliberately set up to acquire the data in the first place as a K -space map. Once a good K map has been obtained then it is a relatively simple matter to recombine the waves and recover a picture of the original slice. First the connections between

projections, their K -space representation and the K -space map of the slice have to be established through the central slice theorem.

The Central Slice Theorem

The Fourier solution to the tomography problem relies on the central slice theorem. Appendix B gives a simple mathematical derivation of the theorem. Here we will show qualitatively how this is used to provide an (in principle) exact method of image reconstruction.

Once sufficient projection data from a real image slice has been collected, we have our numbers in a box problem but on a very much larger scale. We superimpose an imaginary grid on the slice and the local contrast value in each grid square is denoted by μ . This varies continuously with position X, Y , and so we write $\mu = \mu(X, Y)$. One element of a projection, that is to say one of the numbers outside the box in Figure 1.2 consists of a sum of all the values of $\mu(X, Y)$ along a particular line of sight, chosen for that projection, across the slice. If we denote the position of this number within the projection as x' , we can write $P(x') = \int \mu ds$, where the element ds is an element of length along the particular line of sight across the slice. As we move along the projection, we vary x' , and build up the profile, $P(x')$ of this particular projection. The next step requires Fourier analysis. We make the Fourier transform of our projection — $F(P(x'))$. The central slice theorem of Fourier theory allows to say that this is exactly the same thing as one profile of the 2D Fourier transform of $\mu(X, Y)$, or put another way, one line of our required K -space map of the slice. Figure 1.5 illustrates this process. The direction of this line in K -space is not at all arbitrary but exactly parallel to the direction of the chosen projection. We now know how the slice can be decomposed into waves along a particular direction. Slices are however two-dimensional and thus we need an array of different Fourier profiles in K -space before we can obtain a picture of μ . If we change the angle by a small amount and obtain another profile then we can add another line to our K -space map (see Figure 1.6). If we obtain projections covering 360° then we will have sampled K -space on a closely-spaced set of lines or spokes. We now know how the slice can be decomposed into waves along any direction.

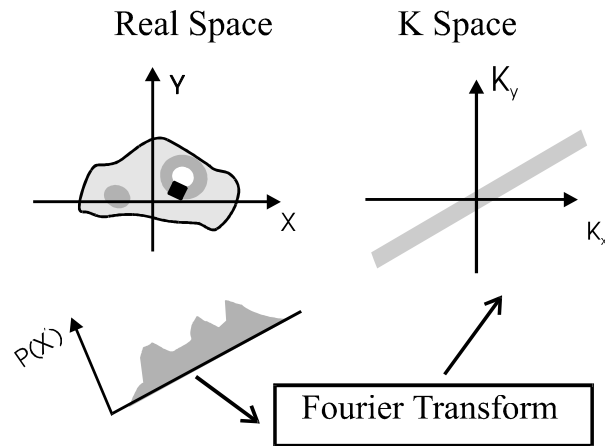


Figure 1.5 The Central Slice Theorem. The projection, $P(x')$, of the slice through the object is obtained by moving a fine X-ray beam across the slice at the angle shown in the left panel. A 1D Fourier transform of $P(x')$ yields one profile through the 2D Fourier transform of $\mu(X, Y)$ as shown in the right panel. The line in Fourier space makes the same angle with the K axes as the original projection direction does to a set of axes established in the real space slice.

Fourier analysis provides the means not only to decompose an image into waves but, starting with the set of waves, the K -space map, to recombine them into a picture of the original slice. This is the inverse Fourier transform. If we make the 2D inverse Fourier transform of our K -space map then we obtain an estimate or picture of $\mu(X, Y)$.

Digital Reconstruction

In practice all medical tomography is carried out with discrete data using a digital computer. This leads to a number of slight changes from the description that we have just given and some constraints. Each projection is measured using detectors that have a finite resolution. In X-ray CT the X-ray photons are recorded using an array of detector elements, each of which is about 5–10 mm across. Thus each profile is actually sampled at a finite number, 256 (say) points. Similarly in MRI the projection profile is obtained from a time signal that is sampled at definite, discrete intervals in

time. In addition there will be a finite number, e.g. 256, of different projections that together span K -space.

In X-ray CT the projections are lines through the origin of K -space and so the complete sampling is achieved by assembling a set of ‘spokes’ with a set of ‘beads’ on each spoke. Each bead corresponds to one value of K_x, K_y and at each of these points there is one value of $F(K_x, K_y)$. The complete set constitutes a sampled version of the 2D Fourier transform of $\mu(X, Y)$. In MRI there is considerably more flexibility in the way in which K -space is actually sampled. Conventionally in MRI, the values of $F(K_x, K_y)$ profiles are collected along lines parallel to K_x at constant values of K_y , each projection then corresponds to a different value of K_y . Rapid imaging MRI schemes are now used that have a variety of more exotic paths through K -space, including spirals. In X-ray CT, the order in which discrete profiles are gathered is immaterial but in MRI, particularly in rapid imaging schemes, the order and trajectory in K -space can have important consequences on the resulting image. This is discussed further in Chapter 6.

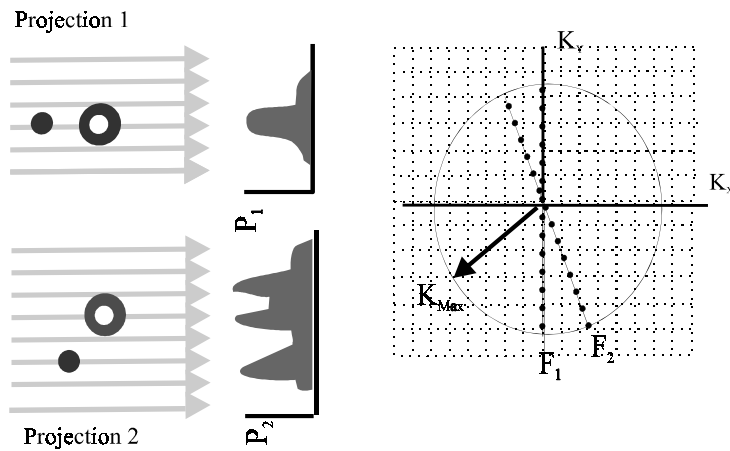


Figure 1.6 Two projections P_1, P_2 of the slice in the left panel provide two lines in K -space as shown in the right panel. Each of the filled ‘beads’ in Fourier space represents one Fourier amplitude obtained from the 1D FT of the corresponding projection data sampled at discrete intervals. Notice that at large K , Fourier space is less well-sampled than at low K . This leads to a loss in spatial resolution. The circle in Fourier space represents the maximum possible value of K that can be obtained when real detectors of finite width or finite sampling intervals are used.

The final step of inverse 2D Fourier transform of the sampled version of $F(K_x, K_y)$ is in principle straightforward. In practice there are some important points, arising from the way the data is collected, that need to be considered if we want to understand the final image and be able to distinguish real from spurious features or artefacts. Each technology using tomography has its own particular set of limitations brought about by engineering compromises. These will be discussed in the appropriate subsequent chapters. Here we will give a brief description of some general limitations created by sequential data gathering, digitisation and the intervening excursion to and from K -space.

Reconstruction Variations

Although all the medical imaging technologies use a reconstruction method that is based on the Fourier method, there are differences in its implementation in PET, SPECT and CT on the one hand and MRI on the other. We will show in Chapter 6 that in MRI the measured signal is in fact already a 1D Fourier transform corresponding to one line in K -space. Thus in this case there is no need for an explicit forward transformation of profile data. Rather the K -space data is collected line by line and then a single final inverse 2D transform provides an estimate of the reconstructed image. Thus reconstruction in MRI corresponds most closely to the Fourier method.

CT, SPECT and PET use a slightly different method that has particular advantages for these technologies. Here the method is filtered backprojection. This is described in Appendix B and its results are illustrated in Figure 1.7. Filtered backprojection is rather similar to our simple scheme since it doesn't involve the explicit use of a Fourier transform-inverse manipulation. Instead it makes use of a mathematical operation called convolution, but derived directly from Fourier. In filtered backprojection each projection is first multiplied by (convoluted with) a mathematical function before backprojection. The function, in its simplest form, is the Radon filter shown in Figure B.3. More complex filter functions are used in gamma imaging in order to remove the effects of noise on the final image. Some common composite filters are illustrated in Figure B.4.

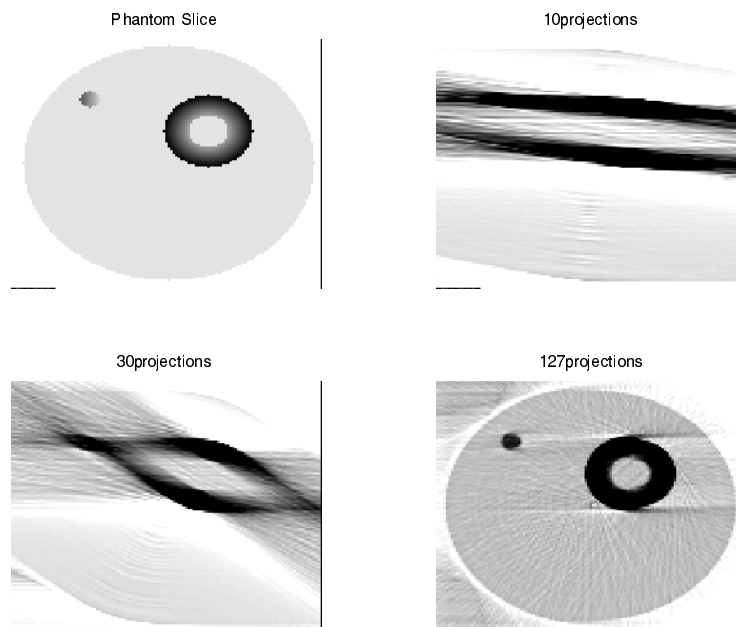


Figure 1.7 Tomographic reconstruction of the model slice used in Figure 1.6, using the filtered backprojection method. The model consists of a small homogeneous disc, and a toroid. Different values of ‘attenuation coefficient’ correspond to different greys. The reconstruction has been interrupted after 10, 30, 127 projections, obtained at intervals of $1 \dots 2^\circ$, have been added. A crude reconstruction of the form of the phantom is clearly visible even after 30 projections. Notice that clinically unacceptable ‘streak’ artefacts remain in the final image, arising here from the deliberate relatively coarse sampling in angle.

Filtered backprojection is computationally more efficient, especially for CT, since reconstruction can begin as soon as the first projection has been obtained, allowing the data acquisition and reconstruction steps in the imaging process to be interleaved in time. The method also lends itself rather easily to modification, along the lines described in Appendix B, in order to deal with the effects of statistical noise that are common in gamma imaging.

1.4 Practical Issues in Tomography

Sequential Data Collection and Patient Movement

The projection reconstruction scheme that we have just described necessarily requires that the data be acquired over a finite period of time, dictated by mechanical, physical or even physiological constraints that enter a particular method. The original Hounsfield EMI CT scanner took about an hour to collect enough data for one slice. Each element in each projection required the mechanical positioning of an X-ray source and an X-ray detector and then a counting time at each position. Modern X-ray CT machines are very much faster and collect slice data in about 0.2 s. Anatomical MRI on the other hand can still take about five minutes per slice for some examinations. During the time that data is being acquired in a sequential fashion, great care has to be taken to limit movement of the patient. Any movement will alter the relative positions of a subset of the projections. At best this will ultimately limit the spatial resolution of the image, leading to some smearing. At worst completely fictitious features can appear in the image because some of the projection data is misplaced with respect to the rest. The Fourier reconstruction process can be unforgiving of this type of data error.

Digitisation of the Projections

As we have already said the projection profiles, $P(x')$ are not obtained as continuous functions of position but rather as a finite number of samples. In the case of X-ray CT the distance between successive samples in any one projection is determined by the size of the X-ray detector element used. Each little detector sums the collected X-ray intensity over its width and averages over any spatial variations whose scale is smaller than the detector size. It thus acts like a low pass spatial filter. Finite detectors or indeed any finite sampling of the projections sets a definite lower limit on the final spatial resolution in the image. The finite sampling interval in real space corresponds to the maximum values of K_x , K_y at which data are collected and thus the fine spatial detail that can be resolved.

In MRI the data are acquired as a time signal that is sampled at a discrete number of time points. Here care has to be taken to avoid the aliasing problem (see Appendix A) by ensuring that the signal, before sampling, does not contain any temporal frequency components that are more than twice the sampling rate. Failure to observe this will result in gross image artefacts as a direct consequence of the Fourier scheme.

The Number of Projections

In every technology a particular expanse of a slice, called the field of view (FOV), is selected for image reconstruction. Whatever the scheme used to sample Fourier space, each projection will span the field of view in one direction with, say N , discrete points. If the FOV has dimensions $L \times L$, then the interval between points, along one direction will be L/N . The smallest K value will be $K_{\min} = 2\pi/L$ and the maximum value of K will be $K_{\max} = \pi N/L$. Generally K -space will be sampled in the other orthogonal direction in increments of K_{\min} . The aim is to sample K -space as uniformly as possible so that there are approximately equal intervals between points in both the K_x and the K_y directions. Once the digitisation step along a projection is established, by for example the choice of detector size, then the interval between successive projections in K -space is arranged to achieve uniform sampling. Any under-sampling in a particular direction will introduce artefacts again as a direct result of the Fourier reconstruction process.

Signal to Noise Ratio

All medical imaging schemes suffer to a greater or lesser extent from the effects of random noise. Both X-ray and gamma ray detection is a random process which has an inevitable associated statistical uncertainty. If N counts are collected in a detection process then there is an intrinsic uncertainty of \sqrt{N} in the number of counts. Thus the signal to noise ratio is $N/\sqrt{N} = \sqrt{N}$. The relative accuracy of any atomic counting measurement increases with the total number of counts collected. X-ray CT is the least affected because extremely large numbers of X-ray photons are

collected to give each element in a projection. Both SPECT and PET, on the other hand, are limited by random fluctuations in the photon counts making up each projection element. In MRI the radio frequency signal is intrinsically weak in comparison with thermal noise as shown in Chapters 2 and 6 and thus it too is eventually limited by random fluctuations in the signal. In addition to degrading the imaging signal, random noise can also upset the intervening Fourier steps in tomography, introducing extra problems. For example large amounts of noise can mimic high frequency signals, leading to aliasing and eventually to image artefacts as well as a general reduction in the sharpness of the final image.

Tomographic Image Artefacts

Tomography, can introduce false detail and distortion, that is to say image artefacts, into the final image as a result of faulty equipment, errors in the stored data and patient movement. Because of its digital nature, tomography can also produce generic artefacts. Patient movement is always a problem when data acquisition times exceed a few seconds. The movement can simply blur the image or introduce more troublesome discrete artefacts, when the patient motion modulates the projection data with a particular temporal frequency. The interpretation of any tomographic image can be confused by partial volume effects. These arise directly from the use of finite image voxel dimensions. Even a very high resolution image, made up from cubical cells about 1 mm on a side may well comprise some cells covering a range of tissue types, bone, blood vessel, muscle and fat. The signal from these cells will reflect the average rather than just one tissue type. At the edges of very high contrast, such as bony structures in CT, there will be very abrupt changes in the numerical values making up the projections. Image reconstruction using a finite number of projections cannot faithfully reproduce the sharp edge; rather it creates a halo close to the edge in the image.

Key Concepts in Chapter 1

Concept	Section	Cross Ref
The three steps in all Tomography	1.2	4.5, 6.4, B1
<i>K</i> -space	1.3	6.4, 6.5, A.1, A.3, B.1
Central Slice theorem	1.3	4.5, 6.4, B1
Tomographic image artefacts	1.4	4.5, 6.7

Questions and Problems

1.
 - a) In the context of X-ray CT explain the term projection.
 - b) Describe the method of simple backprojection.
 - c) State the central slice theorem and explain how it may be used to reconstruct tomographic images. (see Appendix B)
2. A phantom object used to check X-ray CT consists of a solid cylinder of bone of radius R cm, embedded in a very weakly attenuating block of resin. Show that the path length L , of an X-ray beam passing across the cross-section of the cylinder, a distance x from the cylinder axis is given by:

$$L = 2\sqrt{R^2 - x^2}$$

Sketch the profile of a projection obtained by scanning a pencil beam across the diameter of the cylinder and use a superposition principle to deduce, without further calculation, the profile of a hollow cylinder with inner and outer radii 5 and 8 cm.

3. A square chessboard, 0.25 m on a side, has 64 equal adjacent alternating black and white squares. What are the lowest and the most probable spatial frequencies associated with the board?
4. The empty reconstruction matrix for a cross section of the human thorax consists of a square array of 512×512 cells. If the field of view just contains the entire thorax estimate the dimensions of each cell. Can blood vessels be imaged with this field of view?
5. The detectors used in a certain CT machine have a width w . How many samples are required to obtain the highest possible resolution in a projection of an object with a maximum dimension D ?