


SUMAT

SPRINGER UNDERGRADUATE TEXTS IN
MATHEMATICS AND TECHNOLOGY

Timothy G. Feeman

The Mathematics of Medical Imaging

A Beginner's Guide

 Springer

Springer Undergraduate Texts
in Mathematics and Technology

For other titles published in this series, go to
<http://www.springer.com/series/7438>

Series Editors

Jonathan M. Borwein

Helge Holden

Editorial Board

Lisa Goldberg

Armin Iske

Palle E.T. Jorgensen

Stephen M. Robinson

Timothy G. Feeman

The Mathematics of Medical Imaging

A Beginner's Guide

 Springer

Timothy G. Feeman
Department of Mathematical Sciences
Villanova University
Villanova, PA 19085-1966
USA
timothy.feeman@villanova.edu

Series Editors

Jonathan M. Borwein
Computer Assisted Research Mathematics
and its Applications, CARMA
School of Mathematical & Physical Sciences
University of Newcastle
Callaghan NSW 2308, Australia
jonathan.borwein@newcastle.edu.au

Helge Holden
Department of Mathematical Sciences
Norwegian University Science
and Technology
Alfred Getz vei 1
NO-7491 Trondheim
Norway
holden@math.ntnu.no

ISSN 1867-5506 e-ISSN 1867-5514
ISBN 978-0-387-92711-4 e-ISBN 978-0-387-92712-1
DOI 10.1007/978-0-387-92712-1
Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2009939881

Mathematics Subject Classification (2000): 42A38, 42B10, 44A12, 65T50, 68U10, 94A08, 94A20

© Springer Science+Business Media, LLC 2010

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

For Max and Simon

Contents

Preface	xi
1 X-rays	1
1.1 Introduction	1
1.2 X-ray behavior and Beer's law	3
1.3 Lines in the plane	7
1.4 Exercises	9
2 The Radon Transform	11
2.1 Definition	11
2.2 Examples	12
2.3 Linearity	14
2.4 Phantoms	15
2.5 The domain of \mathcal{R}	17
2.6 Exercises	18
3 Back Projection	19
3.1 Definition and properties	19
3.2 Examples	20
3.3 Exercises	23
4 Complex Numbers	25
4.1 The complex number system	25
4.2 The complex exponential function	27
4.3 Wave functions	29
4.4 Exercises	30

5	The Fourier Transform	33
	5.1 Definition and examples	33
	5.2 Properties and applications	34
	5.3 Heaviside and Dirac δ	38
	5.4 Inversion of the Fourier transform	40
	5.5 Multivariable forms	43
	5.6 Exercises	45
6	Two Big Theorems	47
	6.1 The central slice theorem	47
	6.2 Filtered back projection	48
	6.3 The Hilbert transform	50
	6.4 Exercises	51
7	Filters and Convolution	53
	7.1 Introduction	53
	7.2 Convolution	55
	7.3 Filter resolution	58
	7.4 Convolution and the Fourier transform	58
	7.5 The Rayleigh–Plancherel theorem	60
	7.6 Convolution in 2-dimensional space	62
	7.7 Convolution, \mathcal{B} , and \mathcal{R}	62
	7.8 Low-pass filters	63
	7.9 Exercises	66
8	Discrete Image Reconstruction	67
	8.1 Introduction	67
	8.2 Sampling	68
	8.3 Discrete low-pass filters	70
	8.4 Discrete Radon transform	73
	8.5 Discrete functions and convolution	74
	8.6 Discrete Fourier transform	77
	8.7 Discrete back projection	82
	8.8 Interpolation	82
	8.9 Discrete image reconstruction	86
	8.10 Matrix forms	91
	8.11 FFT—the fast Fourier transform	92
	8.12 Fan beam geometry	97
	8.13 Exercises	98

9 Algebraic Reconstruction Techniques 101

9.1 Introduction 101

9.2 Least squares approximation 103

9.3 Kaczmarz’s method 107

9.4 ART in medical imaging 110

9.5 Variations of Kaczmarz’s method 111

9.6 ART or the Fourier transform? 112

9.7 Exercises 113

10 MRI—An Overview 115

10.1 Introduction 115

10.2 Basics 116

10.3 The Bloch equation 118

10.4 The RF field 119

10.5 RF pulse sequences; T_1 and T_2 122

10.6 Gradients and slice selection 124

10.7 The imaging equation 127

10.8 Exercises 128

Appendix A Integrability 129

A.1 Improper integrals 129

A.2 Iterated improper integrals 131

A.3 L^1 and L^2 132

A.4 Summability 133

Appendix B Topics for Further Study 135

References 137

Index 139

Preface

In 1979, the Nobel Prize for Medicine and Physiology was awarded jointly to Allan McLeod Cormack and Godfrey Newbold Hounsfield, the two pioneering scientist-engineers primarily responsible for the development, in the 1960s and early 1970s, of computerized axial tomography, popularly known as the CAT or CT scan. In his papers [13], Cormack, then a Professor at Tufts University, in Massachusetts, developed certain mathematical algorithms that, he envisioned, could be used to create an image from X-ray data. Working completely independently of Cormack and at about the same time, Hounsfield, a research scientist at EMI Central Research Laboratories in the United Kingdom, designed the first operational CT scanner as well as the first commercially available model. (See [22] and [23].)

Since 1980, the number of CT scans performed each year in the United States has risen from about 3 million to over 67 million. What few people who have had CT scans probably realize is that the fundamental problem behind this procedure is essentially mathematical: **If we know the values of the integral of a two- or three-dimensional function along all possible cross-sections, then how can we reconstruct the function itself?** This particular example of what is known as an *inverse problem* was studied by Johann Radon, an Austrian mathematician, in the early part of the twentieth century. Radon's work incorporated a sophisticated use of the theory of transforms and integral operators, and, by expanding the scope of that theory, contributed to the development of the rich and vibrant mathematical field of functional analysis. Cormack essentially rediscovered Radon's ideas, but did so at a time when technological applications were actually conceivable. The practical obstacles to implementing Radon's theories are several. First, Radon's inversion methods assume knowledge of the behavior of the function along every cross-section, while, in practice, only a discrete set of cross-sections can feasibly be sampled. Thus, it is possible to construct only an approximation of the solution. Second, the computational power needed to process a multitude of discrete measurements and, from them, to obtain a useful approximate solution has been available for

just a few decades. The response to these obstacles has been a rich and dynamic development both of theoretical approaches to approximation methods, including the use of interpolation and filters, and of computer algorithms to effectively implement the approximation and inversion strategies. Alongside these mathematical and computational advances, the machines that perform the scans have gone through several generations of improvements in both the speed of data collection and the accuracy of the images, while the range of applications has expanded well beyond the original focus on imaging of the brain. Other related processes, such as positron emission tomography (PET), have developed alongside the advances in CT.

Clearly, this subject crosses many disciplinary boundaries. Indeed, literature on technical aspects of medical imaging appears in journals published in engineering, mathematics, computer science, biomedical research, and physics. This book, which grew out of a course I gave for undergraduate mathematics majors and minors at Villanova University in 2008, addresses the mathematical fundamentals of the topic in a concise way at a relatively elementary level. The emphasis is on the mathematics of CT, though there is also a chapter on magnetic resonance imaging (MRI), another medical imaging process whose originators have earned Nobel prizes. The discussion includes the necessary theoretical background, but also the role of approximation methods and some attention to the computer implementation of the inversion algorithms. A working knowledge of multivariable calculus and basic vector and matrix methods should serve as adequate prerequisite mathematics.

I hope you will join me, then, in this quest to comprehend one of the most significant and beneficial technological advances of our time and to experience mathematics as an inextricable part of human culture.

Villanova University, 2009

X-rays

1.1 Introduction

A computerized axial tomography (CAT or CT) scan is generated from a set of thousands of X-ray beams, consisting of 160 or more beams at each of 180 directions. To comprehend this large collection of X-rays, we must first understand just one beam.

When a single X-ray beam of known intensity passes through a medium, such as muscle or brain tissue or an ancient Egyptian sarcophagus, some of the energy present in the beam is absorbed by the medium and some passes through. The intensity of the beam as it emerges from the medium can be measured by a detector. The difference between the initial and final intensities tells us about the ability of the medium to absorb energy. For the sort of X-rays one might get at the dentist's office or for a suspected broken bone, the detector is a piece of film. A fan- or cone-shaped set of X-rays is emitted from a machine and those photons that are not blocked or absorbed by teeth or bone expose the film, thus creating a picture of the medium. The picture essentially lacks depth since anything positioned behind a point where the photons are blocked will not be seen. This shortcoming highlights a significant difficulty in imaging, namely, that the medium through which the X-rays pass is not homogeneous. For instance, muscles are fibrous and denser in some parts than others; brain tissue is composed of grey matter, water, blood, neurons, and more; inside the sarcophagus is a mummified, partly decomposed body, but also remains of objects that were buried along with the deceased.

The idea behind the CT scan is that, by measuring the changes in the intensity of X-ray beams passing through the medium in different directions and, then, by comparing the measurements, we might be able to determine which locations within the medium are more absorbent or less absorbent than others.

To get an idea of how this works, let's start with a simple model. Suppose we have a one-centimeter-thick slice of material (the medium) in the shape of a square. The

square is divided into a 3-by-3 rectangular grid of smaller squares each of which is either black or white. Each white square absorbs “1 unit” of X-ray energy while the black squares do not absorb X-ray energy. (So the white squares act like bone, say, and the black ones act like air.) Suppose now that an X-ray beam passing through the first row of the grid loses 2 energy units. It follows that there must be two white squares, and one black square, in the first row of the grid. If an X-ray beam passing through the first column of the grid loses 1 unit of energy, then the first column would contain only one white square. At this point, there are only four possibilities for the configuration of the first row and column (instead of the initial $2^5 = 32$ possibilities for these five squares), namely, Row 1 = WWB and Column 1 = WBB; Row 1 = WBW and Column 1 = WBB; Row 1 = BWW and Column 1 = BWB; or Row 1 = BWW and Column 1 = BBW. Continuing in this way, suppose that we measure the following losses in X-ray energy for the various rows and columns of the grid: Row 1 \rightarrow 2 units lost; Row 2 \rightarrow 2 units; Row 3 \rightarrow 1 unit; Column 1 \rightarrow 1 unit; Column 2 \rightarrow 2 units; and Column 3 \rightarrow 2 units of energy lost. Figure 1.1 shows one of several possible configurations of black and white squares that are consistent with these measurements. What are the other




Fig. 1.1. This grid of white and black squares has a prescribed X-ray energy absorption for each row and column.

possible configurations? Is there an easy way to determine the total number of white squares, and, consequently, the total number of black squares, in the grid? Is there more than one consistent pattern? If so, what additional information might help to pinpoint a unique shading pattern? For a rich and highly entertaining elaboration on this model for thinking about CT scans, see [18].

In some sense, creating an image from a CT scan consists of carrying out a scheme like this on a rectangular image grid subdivided into thousands of small squares. Instead of just black or white, each square is assigned a grayscale value—a number between 0 and 1, where, by common practice, black is 0 and white is 1—based on the energy-absorption ability of the material located in that square in the grid. Each X-ray passing through the material is measured, and the change in intensity gives us information about the amount of grey encountered along the path of the beam. What is not known is

precisely where along the way the reduction in energy occurred. Once we know the changes in intensity for enough X-rays, we try to create an image whose greyscale values are consistent with the measurements. This approach to image construction is akin to the basic idea implemented by Hounsfield in his early CT scanners and will be discussed in a more refined way in Chapter 9.

1.2 X-ray behavior and Beer's law

To simplify the analysis, we will make some assumptions that present an idealized view of what an X-ray is and how it behaves. Specifically, in thinking of an X-ray beam as being composed of photons, we will assume that the beam is *monochromatic*. That is, each photon has the same energy level E and the beam propagates at a constant frequency, with the same number of photons per second passing through every centimeter of the path of the beam. If $N(x)$ denotes the number of photons per second passing through a point x , then the intensity of the beam at the point x is

$$I(x) = E \cdot N(x).$$

We also assume that an X-ray beam has *zero width* and that it is not subject to refraction or diffraction. That is, X-rays beams are not bent by the medium nor do they spread out as they propagate.

Every substance through which an X-ray passes has the property that each millimeter of the substance absorbs a certain proportion of the photons that pass through it. This proportion, which is specific to the substance, is called the *attenuation coefficient* of that material. The units of the attenuation coefficient are something like “proportion of photons absorbed per millimeter of the medium.” In general the attenuation coefficient is nonnegative and its value depends on the substance involved. Bone has a very high attenuation coefficient, air has a low coefficient, and water is somewhere in between. Different soft tissues have slightly different attenuation coefficients associated with them.

Radiologists actually use a variant of the attenuation coefficient in their work. Developed by Godfrey Hounsfield, the *Hounsfield unit* associated with a medium is a number that represents a comparison of the attenuation coefficient of the medium with that of water. Specifically, the Hounsfield unit of a medium is

$$H_{\text{medium}} := \frac{A_{\text{medium}} - A_{\text{water}}}{A_{\text{water}}}, \quad (1.1)$$

where A denotes the true attenuation coefficient. Table 1.1 gives the Hounsfield units of some typical organic substances.

substance	Hounsfield units	substance	Hounsfield units
bone	1000	kidney	30
liver	40 to 60	cerebrospinal fluid	15
white matter	20 to 30	water	0
grey matter	37 to 45	fat	-100 to -50
blood	40	air	-1000
muscle	10-40		

Table 1.1. Approximate Hounsfield units for certain organic substances.

Now suppose an X-ray beam passes through some medium located between the position x and the position $x + \Delta x$, and suppose that $A(x)$ is the attenuation coefficient of the medium located there. Then the proportion of all photons that will be absorbed in the interval $[x, x + \Delta x]$ is $p(x) = A(x) \cdot \Delta x$. Thus the number of photons that will be absorbed per second by the medium located in the interval $[x, x + \Delta x]$ is $p(x) \cdot N(x) = A(x) \cdot N(x) \cdot \Delta x$. If we multiply both sides by the energy level E of each photon, we see that the corresponding *loss of intensity* of the X-ray beam over this interval is

$$\Delta I \approx -A(x) \cdot I(x) \cdot \Delta x.$$

Let $\Delta x \rightarrow 0$ to get the differential equation known as *Beer's law*:

$$\frac{dI}{dx} = -A(x) \cdot I(x) \tag{1.2}$$

This may also be stated as follows.

Beer's law. The rate of change of intensity per millimeter of a nonrefractive, monochromatic, zero-width X-ray beam passing through a medium is jointly proportional to the intensity of the beam and to the attenuation coefficient of the medium. This condition is expressed by the differential equation (1.2).

The differential equation (1.2) is separable and can be written as

$$\frac{dI}{I} = -A(x) dx.$$

If the beam starts at location x_0 with initial intensity $I_0 = I(x_0)$ and is detected, after passing through the medium, at the location x_1 with final intensity $I_1 = I(x_1)$, then we get

$$\int_{x_0}^{x_1} \frac{dI}{I} = - \int_{x_0}^{x_1} A(x) dx,$$

from which it follows that

$$\ln(I(x_1)) - \ln(I(x_0)) = - \int_{x_0}^{x_1} A(x) dx.$$

Thus

$$\ln\left(\frac{I_1}{I_0}\right) = - \int_{x_0}^{x_1} A(x) dx.$$

Multiplying both sides by -1 yields the result

$$\int_{x_0}^{x_1} A(x) dx = \ln\left(\frac{I_0}{I_1}\right). \quad (1.3)$$

This is a little bit “backwards” from what we often encounter in textbook problems in differential equations. There, we would typically know the coefficient function and use integration to find the function I . Here, however, we know the initial and final values of I , and it is the coefficient function A , which expresses an essential property of the medium being sampled by the X-ray, that is unknown. Thus, we see that from the measured intensity of the X-ray we are able to determine not the values of A itself, but rather the value of the integral of A along the line of the X-ray.

Example 1.1. For a simple example, suppose the attenuation-coefficient function A is constant throughout a sample. Then, the amount of absorption along any given X-ray beam depends only on the width of the sample along the line of the beam. So, if the beam is travelling along the x -axis and enters the sample at x_0 , say, and leaves the sample at x_1 , then the amount of absorption is $A(x_1 - x_0)$. It follows from (1.3) that

$$A = \frac{\ln\left(\frac{I_0}{I_1}\right)}{x_1 - x_0},$$

where I_0 and I_1 are the initial and final intensities of the X-ray.

Example 1.2. Suppose that

$$A(x) := \begin{cases} 1 - |x| & \text{if } |x| \leq 1, \\ 0 & \text{if } |x| > 1, \end{cases}$$

and suppose that an X-ray with initial intensity I_0 is emitted at the point x_0 with $x_0 < -1$, passes through the sample, and has final intensity I_1 , as measured by a detector at the point x_1 with $x_1 > 1$. Since

$$\int_{x_0}^{x_1} (1 - |x|) dx = 1,$$

we see that

$$\ln\left(\frac{I_0}{I_1}\right) = 1,$$

from which it follows that $I_1 = e^{-1} \cdot I_0$. For instance, if $I_0 = 1$, then $I_1 = e^{-1}$. Solving the differential equation $\frac{dI}{dx} = -A(x) \cdot I(x)$, with the initial condition $I_0 = I(-1) = 1$, yields

$$I(x) := \begin{cases} 1 & \text{if } x \leq -1, \\ e^{-x - \frac{1}{2} - \frac{1}{2}x^2} & \text{if } -1 \leq x \leq 0, \\ e^{-x - \frac{1}{2} + \frac{1}{2}x^2} & \text{if } 0 \leq x \leq 1, \\ e^{-1} & \text{if } x \geq 1. \end{cases}$$

Example 1.2: Alternate version. Suppose that the intersection of a sample with the x -axis lies entirely inside the interval $[-1, 1]$ and suppose that an X-ray with initial intensity $I_0 = 1$ is emitted at the point x_0 with $x_0 < -1$, passes through the sample, and has final intensity $I_1 = e^{-1}$, as measured by a detector at the point x_1 with $x_1 > 1$. Moreover, imagine that we somehow knew that the intensity function of the beam was given by

$$I(x) := \begin{cases} 1 & \text{if } x \leq -1, \\ e^{-x - \frac{1}{2} - \frac{1}{2}x^2} & \text{if } -1 \leq x \leq 0, \\ e^{-x - \frac{1}{2} + \frac{1}{2}x^2} & \text{if } 0 \leq x \leq 1, \\ e^{-1} & \text{if } x \geq 1. \end{cases}$$

Beer's law (1.2) now yields

$$A(x) := \begin{cases} 1 - |x| & \text{if } |x| \leq 1, \\ 0 & \text{if } |x| > 1. \end{cases}$$

With actual X-ray detection equipment, though, we would not know the function $I(x)$ at all values of x , only at the points of emission and detection. The condition $I_1 = e^{-1} \cdot I_0$ tells us only that $\int_{x_0}^{x_1} A(x) dx = 1$, not what formula $A(x)$ has.

For instance, we would not be able to distinguish $A_1(x) := \begin{cases} 1 - |x| & \text{if } |x| \leq 1, \\ 0 & \text{if } |x| > 1 \end{cases}$ from $A_2(x) := \begin{cases} 1 & \text{if } |x| \leq 1/2, \\ 0 & \text{if } |x| > 1/2 \end{cases}$ or, for that matter, from any other attenuation function having an integral equal to 1. A single X-ray can measure only the integral of the attenuation function, not its shape.

What we can measure: We can design an X-ray emission/detection machine that can measure the values of I_0 and I_1 . Hence, from (1.3), we can compute $\int_{x_0}^{x_1} A(x) dx$, the integral of the (unknown) attenuation-coefficient function along the path of the X-ray.

What we want to know: The value of $A(x)$ at each location depends on the nature of the matter located at the point x . It is precisely the function A itself that we wish to know.

Two- or three-dimensional interpretation. Suppose a sample of material occupies a finite region in space. At each point (x, y, z) within the sample, the material there has

an attenuation coefficient value $A(x, y, z)$. An X-ray beam passing through the sample follows a line ℓ from an initial point P (assumed to be outside the region) to a final point Q (also assumed to be outside the region). The emission/detection machine measures the initial and final intensities of the beam at P and Q , respectively, from which the value $\ln(I_{\text{initial}}/I_{\text{final}})$ is calculated. According to (1.3), this is equal to the value of the integral $\int_{\overline{PQ}} A(x, y, z) ds$, where ds represents arclength units along the segment \overline{PQ} of the line ℓ . Thus, the measurement of each X-ray beam gives us information about the average value of A along the path of the beam.

In our study of CT scans, we will consider a two-dimensional slice of the sample, obtained as the intersection of the sample and some plane, which we will generally assume coincides with the xy -plane. In this context, we interpret the attenuation-coefficient function as a function $A(x, y)$ of two variables within the specific slice. Indeed, the word *tomography* is built on the Greek language root form *tomos* meaning “slice.”

The fundamental question of image reconstruction is this: *Can we reconstruct the function $A(x, y, z)$ (within some finite region) if we know the average value of A along every line that passes through the region?*

1.3 Lines in the plane

For simplicity, let us assume that we are interested only in the cross-section of a sample that lies in the xy -plane. Each X-ray will follow a segment of a line in the plane, so we would like to have a way of cataloguing all such lines. For instance, every nonvertical line has an equation of the form $y = mx + b$. So we could catalogue these lines using all possible pairs (m, b) . However, vertical lines would be excluded from this list. Instead, we can classify all lines by adopting a “point–normal” approach, in which every line in the plane is characterized by a pair consisting of a point that the line passes through and a vector that is normal (i.e., perpendicular) to the line.

For a vector $\vec{\mathbf{n}}$ normal to a given line ℓ , there is some angle θ (with $0 \leq \theta < 2\pi$, say) such that $\vec{\mathbf{n}}$ is parallel to the line radiating out from the origin at an angle of θ measured counterclockwise from the positive x -axis. This line, at angle θ , is also perpendicular to ℓ and, hence, intersects ℓ at some point whose coordinates in the xy -plane have the form $(t \cos(\theta), t \sin(\theta))$ for some real number t . In this way, the line ℓ is characterized by the values of t and θ and is accordingly denoted by $\ell_{t, \theta}$. This may also be stated as follows.

Definition 1.3. For any real numbers t and θ , the line $\ell_{t, \theta}$ is the line that passes through the point $(t \cos(\theta), t \sin(\theta))$ and is perpendicular to the unit vector $\vec{\mathbf{n}} = \langle \cos(\theta), \sin(\theta) \rangle$.

Every line in the plane can be characterized as $\ell_{t, \theta}$ for some values of t and θ .

For example, the line with Cartesian equation $x + y = \sqrt{2}$ is the same as the line $\ell_{1,\pi/4}$. Figure 1.2 shows an array of lines corresponding to eleven different values of t at each of ten angles θ .

Two relationships are apparent, namely,

$$\ell_{t,\theta+2\pi} = \ell_{t,\theta} \text{ and } \ell_{t,\theta+\pi} = \ell_{-t,\theta} \text{ for all } t, \theta.$$

This means that each line in the plane has many different representations of the form $\ell_{t,\theta}$. To avoid this, we can use either of the sets

$$\{\ell_{t,\theta} : t \text{ real}, 0 \leq \theta < \pi\} \text{ or } \{\ell_{t,\theta} : t \geq 0, 0 \leq \theta < 2\pi\}.$$

For the most part, we will use the former of these sets.

To parameterize a line $\ell_{t,\theta}$, observe that the unit vector $\langle -\sin(\theta), \cos(\theta) \rangle$ is perpendicular to the vector $\langle \cos(\theta), \sin(\theta) \rangle$. Thus, every point on $\ell_{t,\theta}$ has the form

$$\langle t \cos(\theta), t \sin(\theta) \rangle + s \cdot \langle -\sin(\theta), \cos(\theta) \rangle$$

for some real number s . That is, the line $\ell_{t,\theta}$ can be parameterized as $(x(s), y(s))$, where $x(s) = t \cos(\theta) - s \sin(\theta)$ and $y(s) = t \sin(\theta) + s \cos(\theta)$ for $-\infty < s < \infty$. Consequently, $\ell_{t,\theta}$ can be described as

$$\ell_{t,\theta} = \{(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) : -\infty < s < \infty\}. \quad (1.4)$$

Note that at every point $(x(s), y(s))$ on $\ell_{t,\theta}$ we get $x^2 + y^2 = t^2 + s^2$.




Fig. 1.2. The figure shows the lines $\ell_{t,\theta}$ for eleven different values of t and values of θ in increments of $\pi/10$.

For an arbitrary point (a, b) in the plane and a given value for θ , there is a unique value of t for which the line $\ell_{t,\theta}$ passes through (a, b) . More precisely, there are unique values of both t and s for which $a = t \cos(\theta) - s \sin(\theta)$ and $b = t \sin(\theta) + s \cos(\theta)$. This is a system of two equations in the two unknowns, t and s . The solutions are

$t = a \cos(\theta) + b \sin(\theta)$ and $s = -a \sin(\theta) + b \cos(\theta)$. Again, note that $t^2 + s^2 = a^2 + b^2$, a fact that will be used later to implement a change of coordinates from the (x, y) system to the (t, s) framework.

With the parameterization $x(s) = t \cos(\theta) - s \sin(\theta)$ and $y(s) = t \sin(\theta) + s \cos(\theta)$, the arclength element along the line $\ell_{t,\theta}$ is given by

$$\sqrt{\left(\frac{dx}{ds}\right)^2 + \left(\frac{dy}{ds}\right)^2} ds = \sqrt{(-\sin(\theta))^2 + (\cos(\theta))^2} ds = ds.$$

Therefore, for a given function $A(x, y)$ defined in the plane, we get

$$\int_{\ell_{t,\theta}} A(x, y) = \int_{s=-\infty}^{\infty} A(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \quad (1.5)$$

The value of this integral is exactly what an X-ray emission/detection machine measures when an X-ray is emitted along the line $\ell_{t,\theta}$.

We can now rephrase the fundamental question of image reconstruction to ask, “*Can we reconstruct the function $A(x, y)$ (within some finite region of the plane) if we know the value of $\int_{\ell_{t,\theta}} A(x, y)$ for every line $\ell_{t,\theta}$?*”

1.4 Exercises

1.1. Consider a 3×3 grid in which each of the nine squares is shaded either black or white.

- Find all possible grids for which Rows 1 and 2 each have two white squares and Row 3 has one white square; Column 1 has one white square and Columns 2 and 3 each have two white squares. (Figure 1.1 shows one solution.)
- Find all possible grids for which Rows 1 and 2 each have two white squares and Row 3 has one white square; Columns 1 and 3 each have one white square and Column 2 has three white squares.

1.2. Find two different 4×4 grids having the same “row and column scans;” that is, the first row of the first pattern has the same number of white squares as the first row of the second pattern has, and so on.

1.3. What additional information would help to identify a 3×3 or 4×4 grid uniquely?

1.4. What does the attenuation coefficient measure?

1.5. Referring to Table 1.1, why must an image be accurate to within about 10 Hounsfield units in order to be clinically useful?

1.6. Explain in a few sentences, and with a minimum of mathematical detail, why Beer's law is a plausible model for X-ray attenuation.

1.7. Let

$$A(x) := \begin{cases} 1 - |x| & \text{if } |x| \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

and

$$I(x) := \begin{cases} 1 & \text{if } x \leq -1, \\ e^{-x - \frac{1}{2} - \frac{1}{2}x^2} & \text{if } -1 \leq x \leq 0, \\ e^{-x - \frac{1}{2} + \frac{1}{2}x^2} & \text{if } 0 \leq x \leq 1, \\ e^{-1} & \text{if } x \geq 1. \end{cases}$$

(a) Evaluate $\int_{-1}^1 A(x) dx$ and $\ln\left(\frac{I(-1)}{I(1)}\right)$.

(b) Verify that these functions satisfy the differential equation

$$\frac{dI}{dx} = -A(x) \cdot I(x).$$

1.8. (a) Explain the parameterization

$$\ell_{t,\theta} = \{(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) : -\infty < s < \infty\}.$$

(b) Prove that $x^2 + y^2 = t^2 + s^2$ at every point (x, y) on $\ell_{t,\theta}$.

1.9. Find values of t and θ for which the line $\ell_{t,\theta}$ is the same as the line with equation $\sqrt{3}x + y = 4$.

1.10. Explain why $\ell_{t,\theta} = \ell_{-t,\theta+\pi}$ for all t and all θ . (A sketch might help.)

1.11. (a) Given a , b , and θ , find t and s so that

$$(a, b) = (t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)).$$

(b) Show that $t^2 + s^2 = a^2 + b^2$ for the values you found in part (a).

1.12. For a fixed value of $R > 0$, let $f(x, y) := \begin{cases} 1 & \text{if } x^2 + y^2 \leq R^2, \\ 0 & \text{otherwise.} \end{cases}$

$$\text{Show that } \int_{\ell_{t,\theta}} f ds = \begin{cases} 2\sqrt{R^2 - t^2} & \text{if } |t| \leq R, \\ 0 & \text{if } |t| > R. \end{cases}$$

1.13. Discussion: Why does the fundamental question of image reconstruction require that we consider so many lines? Why would a single set of parallel lines not suffice? (Hint: Think about the game we played with the 3-by-3 grid of black and white squares.)

The Radon Transform

2.1 Definition

For a given function f defined in the plane, which may represent, for instance, the attenuation-coefficient function in a cross section of a sample, the fundamental question of image reconstruction calls on us to consider the value of the integral of f along a typical line $\ell_{t,\theta}$. For each pair of values of t and θ , we will integrate f along a different line. Thus, we really have a new function on our hands, where the inputs are the values of t and θ and the output is the value of the integral of f along the corresponding line $\ell_{t,\theta}$. But even more is going on than that because we also wish to apply this process to a whole variety of functions f . So really we start by selecting a function f . Then, once f has been selected, we get a corresponding function of t and θ . Schematically,

$$\text{input } f \mapsto \text{output } \left\{ (t, \theta) \mapsto \int_{\ell_{t,\theta}} f ds \right\}.$$

This multi-step process is called the *Radon transform*, named for the Austrian mathematician Johann Karl August Radon (1887–1956) who studied its properties. For the input f , we denote by $\mathcal{R}(f)$ the corresponding function of t and θ shown in the schematic. That is, we make the following definition.

Definition 2.1. For a given function f , whose domain is the plane, the *Radon transform* of f is defined, for each pair of real numbers (t, θ) , by

$$\mathcal{R}f(t, \theta) := \int_{\ell_{t,\theta}} f ds = \int_{s=-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds. \quad (2.1)$$

A few immediate observations are that (i) both f and $\mathcal{R}f$ are functions; (ii) f is a function of the Cartesian coordinates x and y while $\mathcal{R}f$ is a function of the polar coordinates t and θ ; (iii) $\mathcal{R}f(t, \theta)$ is a number (the value of an integral); (iv) in the integral on the right, the variable of integration is s , while the values of t and θ are preselected and so should be treated as “constants” when evaluating the integral.

2.2 Examples

Example 2.2. As an example, let f be the function defined by

$$f(x, y) := \begin{cases} 1 - \sqrt{x^2 + y^2} & \text{if } x^2 + y^2 \leq 1, \\ 0 & \text{if } x^2 + y^2 > 1. \end{cases} \quad (2.2)$$

The graph of f is a cone, shown in Figure 2.1. We have already observed that, on the line $\ell_{t, \theta}$, we have

$$x^2 + y^2 = (t \cos(\theta) - s \sin(\theta))^2 + (t \sin(\theta) + s \cos(\theta))^2 = t^2 + s^2.$$

It follows that, on the line $\ell_{t, \theta}$, the function f is given by

$$f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) := \begin{cases} 1 - \sqrt{t^2 + s^2} & \text{if } t^2 + s^2 \leq 1, \\ 0 & \text{if } t^2 + s^2 > 1. \end{cases} \quad (2.3)$$

From this, we see that the value of $\mathcal{R}f(t, \theta)$ depends only on t and not on θ and that $\mathcal{R}f(t, \theta) = 0$ whenever $|t| > 1$. For a fixed value of t such that $|t| \leq 1$, the condition $t^2 + s^2 \leq 1$ will be satisfied provided that $s^2 \leq 1 - t^2$. Thus, for any value of θ and for t such that $|t| \leq 1$, we have

$$\begin{aligned} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) \\ := \begin{cases} 1 - \sqrt{t^2 + s^2} & \text{if } -\sqrt{1 - t^2} \leq s \leq \sqrt{1 - t^2}, \\ 0 & \text{otherwise;} \end{cases} \end{aligned} \quad (2.4)$$

whence

$$\int_{\ell_{t, \theta}} f ds = \int_{s=-\sqrt{1-t^2}}^{\sqrt{1-t^2}} (1 - \sqrt{t^2 + s^2}) ds. \quad (2.5)$$

This integral requires a trigonometric substitution for its evaluation. Sparing the details for now, we have

$$\int_{s=-\sqrt{1-t^2}}^{\sqrt{1-t^2}} (1 - \sqrt{t^2 + s^2}) ds = \sqrt{1-t^2} - \frac{1}{2} t^2 \ln \left(\frac{1 + \sqrt{1-t^2}}{1 - \sqrt{1-t^2}} \right). \quad (2.6)$$




Fig. 2.1. The figure shows the cone defined in (2.2) and the graph of its Radon transform for any fixed value of θ .

In conclusion, we have shown that the Radon transform of this function f is given by

$$\mathcal{R}f(t, \theta) := \begin{cases} \sqrt{1-t^2} - \frac{1}{2}t^2 \ln\left(\frac{1+\sqrt{1-t^2}}{1-\sqrt{1-t^2}}\right) & \text{if } -1 \leq t \leq 1, \\ 0 & \text{if } |t| > 1. \end{cases} \quad (2.7)$$

In this case, where $\mathcal{R}f$ is independent of θ , the value of $\mathcal{R}f(t, \theta)$ corresponds to the area under the vertical cross-section of the cone defined by $z = f(t, y)$. Several of these cross-sections are visible in Figure 2.1.

Example 2.3. Consider a crescent-shaped region inside the disc $x^2 + y^2 = 1/4$ and outside the disc $(x - 1/8)^2 + y^2 = 9/64$. Assign density 1 to points in the crescent, density 1/2 to points inside the smaller disc, and density 0 to points outside the larger disc. In other words, the attenuation function is

$$A(t, \theta) := \begin{cases} 1 & \text{if } x^2 + y^2 \leq 1/4 \text{ and } (x - 1/8)^2 + y^2 > 9/64; \\ 0.5 & \text{if } (x - 1/8)^2 + y^2 \leq 9/64; \\ 0 & \text{if } x^2 + y^2 > 1/4. \end{cases} \quad (2.8)$$




Fig. 2.2. The figure shows the graph of the attenuation function A defined in (2.8), alongside a sinogram of its Radon transform $\mathcal{R}A(t, \theta)$.

Figure 2.2 shows the graph of this attenuation function alongside a graph of its Radon transform in the (t, θ) plane. Such a graph is called a *sinogram* and essentially depicts all of the data generated by the X-ray emission/detection machine for the given slice of the sample. The function values are shown in greyscale, with white corresponding to 1, neutral grey to 0.5, and black to 0. Figure 2.3 shows graphs of the Radon transform for the angles $\theta = 0$ and $\theta = \pi/3$.




Fig. 2.3. For the function A defined in (2.8), the figure shows graphs of its Radon transform $\mathcal{R}A(t, 0)$ (left) and $\mathcal{R}A(t, \pi/3)$ (right), for $-1/2 \leq t \leq 1/2$.

2.3 Linearity

Suppose that two functions f and g are both defined in the plane. Then so is the function $f + g$. Since the integral of a sum of two functions is equal to the sum of the integrals of the functions separately, it follows that we get, for every choice of t and θ ,

$$\begin{aligned}
 \mathcal{R}(f + g)(t, \theta) &= \int_{s=-\infty}^{\infty} (f + g)(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \\
 &= \int_{s=-\infty}^{\infty} \{f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) \\
 &\quad + g(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta))\} ds \\
 &= \int_{s=-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \\
 &\quad + \int_{s=-\infty}^{\infty} g(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \\
 &= \mathcal{R}f(t, \theta) + \mathcal{R}g(t, \theta).
 \end{aligned}$$

In other words, $\mathcal{R}(f + g) = \mathcal{R}f + \mathcal{R}g$ as functions.

Similarly, when a function is multiplied by a constant and then integrated, the result is the same as if the function were integrated first and then that value multiplied by the constant; i.e., $\int \alpha f = \alpha \int f$. In the context of the Radon transform, this means that $\mathcal{R}(\alpha f) = \alpha \mathcal{R}f$.

We now have proven the following proposition.

Proposition 2.4. *For two functions f and g and any constants α and β ,*

$$\mathcal{R}(\alpha f + \beta g) = \alpha \mathcal{R}f + \beta \mathcal{R}g. \quad (2.9)$$

In the language of linear algebra, we say that the Radon transform is a *linear transformation*; that is, the Radon transform \mathcal{R} maps a linear combination of functions to the same linear combination of the Radon transforms of the functions separately. We also express this property by saying that “ \mathcal{R} preserves linear combinations.”

Example 2.5. For a fixed value of R , define

$$F_R(x, y) := \begin{cases} 1 & \text{if } x^2 + y^2 \leq R^2, \\ 0 & \text{otherwise.} \end{cases}$$

From the exercises, we know that

$$\mathcal{R}(F_R)(t, \theta) = \begin{cases} 2\sqrt{R^2 - t^2} & \text{if } |t| \leq R, \\ 0 & \text{if } |t| > R. \end{cases}$$

Now consider the function

$$f(x, y) := \begin{cases} 0.5 & \text{if } x^2 + y^2 \leq 0.25, \\ 1.0 & \text{if } 0.25 < x^2 + y^2 \leq 1.0, \\ 0 & \text{otherwise.} \end{cases}$$

That is, $f = F_1 - (0.5)F_{0.5}$. By (2.9), it follows that

$$\begin{aligned} \mathcal{R}f(t, \theta) &= \mathcal{R}(F_1)(t, \theta) - (0.5)\mathcal{R}(F_{0.5})(t, \theta) \\ &= \begin{cases} 2\sqrt{1 - t^2} - \sqrt{(0.25) - t^2} & \text{if } |t| \leq 0.5, \\ 2\sqrt{1 - t^2} & \text{if } (0.5) < |t| \leq 1, \\ 0 & \text{if } |t| > 1. \end{cases} \end{aligned}$$

2.4 Phantoms

The fundamental question of image reconstruction asks whether a picture of an attenuation-coefficient function can be generated from the values of the Radon transform of

that function. We will see eventually that the answer is “Yes” if all values of the Radon transform are available, but only “*Approximately yes*” in practice, where only a finite set of values of the Radon transform are measured by a scanning machine. Consequently, the nice solution that works in the presence of full information will splinter into a variety of approximation methods that can be implemented when only partial information is at hand.

One method for testing the accuracy of a particular image reconstruction algorithm, or for comparing algorithms, is simply to apply each algorithm to data taken from an actual human subject. The drawback of this approach is that usually we don’t know exactly what we ought to see in the reconstructed image. That is what we are trying to find out by creating an image in the first place. But without knowing what the real data are, there is no way to determine the accuracy of any particular image. To get around this, we can apply algorithms to data taken from a physical object whose internal structure is known. That way, we know what the reconstructed image ought to look like and we can recognize inaccuracies in a given algorithm or identify disparities between different algorithms. Nonetheless, this approach can be misleading. Although the internal structure of the object is known, there may be errors in the data that were collected to represent the object. In turn, these errors may lead to errors in the reconstructed image. We will not be able to distinguish these flaws from errors caused by the algorithm itself. To resolve this dilemma, Shepp and Logan (see [40]) introduced the concept of a *mathematical phantom*. This is a simulated object whose structure is completely defined by mathematical formulas. Thus, no errors occur in collecting the data from the object. When an algorithm is applied to produce a reconstructed image of the phantom, all inaccuracies are due to the algorithm. This makes it possible to compare different algorithms in a meaningful way.

Since measurement of the Radon transform of an object forms the basis for creating a CT image of the object, it makes sense to use phantoms for which the Radon transform is known exactly. We can then test a proposed algorithm by seeing how well it handles the data from such a phantom. For example, we have computed the Radon transform of a circular disc of constant density centered at the origin. Using the linearity of \mathcal{R} , we can compute the Radon transform of any collection of nested discs, each having constant density, centered at the origin. Such a phantom is too simplistic, though, to serve as a useful model for any serious application. An actual slice of brain tissue is likely to include interesting features in a variety of shapes and sizes located in all regions within the slice, not only near the center. To address these concerns, Shepp and Logan developed the phantom shown in Figure 2.4.

The Shepp–Logan phantom is composed of eleven ellipses of various sizes, eccentricities, and locations. (The MATLAB^R version shown here does not include an ellipse that models a blood clot in the lower right near the boundary.) The densities are assigned so that they fall into the ranges typically encountered in a clinical setting. Since




Fig. 2.4. The Shepp–Logan phantom is used as a mathematical facsimile of a brain for testing image reconstruction algorithms. This version of the phantom is stored in MATLAB[®]. On the right is a sinogram of the phantom’s Radon transform, with θ in increments of $\pi/18$ (10°).

we can compute exactly the Radon transform of any ellipse, the Shepp–Logan phantom has proven to be a reliable model on which to test reconstruction algorithms.

2.5 The domain of \mathcal{R}

As we can see from the definition (2.1), the Radon transform $\mathcal{R}f$ of a function f is defined provided that the integral of f along $\ell_{t,\theta}$ exists for every pair of values of t and θ . Each of these integrals is ostensibly an improper integral evaluated on an infinite interval. Thus, in general, the function f must be integrable along every such line, as discussed in greater detail in Appendix A.

In the context of medical imaging, the function f represents the density or attenuation-coefficient function of a slice of whatever material is being imaged. Thus, the function has *compact support*, meaning that there is some finite disc outside of which the function has the value 0. In this case, the improper integrals $\int_{\ell_{t,\theta}} f ds$ become regular integrals over finite intervals. The only requirement, then, for the existence of $\mathcal{R}f$ is that f be integrable over the finite disc on which it is supported. This will be the case, for instance, if f is piecewise continuous on the disc.

For a wealth of information about the Radon transform and its generalizations, as well as an extensive list of references on this topic, see the monograph [20].

2.6 Exercises

2.1. Evaluate the integral $\int_{s=-\sqrt{1-t^2}}^{\sqrt{1-t^2}} (1 - \sqrt{t^2 + s^2}) ds$ from (2.5).

2.2. Use a computer algebra system to generate pictures such as those in Figures 2.1 and 2.2.

2.3. Consider the function $f(x, y) := \begin{cases} 1 & \text{if } |x| \leq 1 \text{ and } |y| \leq 1, \\ 0 & \text{otherwise.} \end{cases}$ (That is, f has the value 1 inside the square where $-1 \leq x \leq 1$ and $-1 \leq y \leq 1$, and the value 0 outside this square.)

- (a) Sketch the graph of the function $\mathcal{R}f(t, 0)$, the Radon transform of f corresponding to the angle $\theta = 0$.
- (b) Sketch the graph of the function $\mathcal{R}f(t, \pi/4)$, the Radon transform of f corresponding to the angle $\theta = \pi/4$.

Back Projection

3.1 Definition and properties

Let us begin the process of trying to recover the values of an attenuation-coefficient function $f(x, y)$ from the values of its Radon transform $\mathcal{R}f$.

Suppose we select some point in the plane, call it (x_0, y_0) . This point lies on many different lines in the plane. In fact, for each value of θ , there is *exactly one* real number t for which the line $\ell_{t, \theta}$ passes through (x_0, y_0) . Specifically, the value $t = x_0 \cos(\theta) + y_0 \sin(\theta)$ is the one that works, which is to say that, for any given values of x_0 , y_0 , and θ , the line $\ell_{(x_0 \cos(\theta) + y_0 \sin(\theta)), \theta}$ passes through the point (x_0, y_0) . The proof of this fact is left as an exercise.

Figure 3.1 shows a network of back-projection lines through a selection of points in the first quadrant.




Fig. 3.1. For an array of points in the first quadrant, the figure shows the network of the back-projection lines corresponding to values of θ in increments of $\pi/9$.

In practice, whatever sort of matter is located at the point (x_0, y_0) in some sample affects the intensity of any X-ray beam that passes through that point. We now see that each such beam follows a line of the form $\ell_{(x_0 \cos(\theta) + y_0 \sin(\theta)), \theta}$ for some angle θ . In other words, the attenuation coefficient $f(x_0, y_0)$ of whatever is located at the point (x_0, y_0) is accounted for in the value of the Radon transform $\mathcal{R}f(x_0 \cos(\theta) + y_0 \sin(\theta), \theta)$, for each angle θ .

The first step in recovering $f(x_0, y_0)$ is to compute the *average value* of these line integrals averaged over all lines that pass through (x_0, y_0) . That is, we compute

$$\frac{1}{\pi} \int_{\theta=0}^{\pi} \mathcal{R}f(x_0 \cos(\theta) + y_0 \sin(\theta), \theta) d\theta. \quad (3.1)$$

Formally, this integral provides the motivation for a transform called the back projection, or the back-projection transform.

Definition 3.1. Let $h = h(t, \theta)$ be a function whose inputs are polar coordinates. The *back projection* of h at the point (x, y) is defined by

$$\mathcal{B}h(x, y) := \frac{1}{\pi} \int_{\theta=0}^{\pi} h(x \cos(\theta) + y \sin(\theta), \theta) d\theta. \quad (3.2)$$

Note that the inputs for $\mathcal{B}h$ are Cartesian coordinates while those of h are polar coordinates.

The proof of the following proposition is left as an exercise.

Proposition 3.2. *The back projection is a linear transformation. That is, for any two functions h_1 and h_2 and arbitrary constants c_1 and c_2 ,*

$$\mathcal{B}(c_1 h_1 + c_2 h_2)(x, y) = c_1 \mathcal{B}h_1(x, y) + c_2 \mathcal{B}h_2(x, y) \quad (3.3)$$

for all values of x and y .

Example 3.3. Back projection of \mathcal{R} . In the context of medical imaging, the integral in (3.1) represents the back projection of the Radon transform of the attenuation-coefficient function f :

$$\mathcal{B}\mathcal{R}f(x, y) = \frac{1}{\pi} \int_{\theta=0}^{\pi} \mathcal{R}f(x \cos(\theta) + y \sin(\theta), \theta) d\theta. \quad (3.4)$$

3.2 Examples

Before we rush to the assumption that (3.4) gives us $f(x_0, y_0)$ back again, let us analyze more closely. Each of the numbers $\mathcal{R}f(x_0 \cos(\theta) + y_0 \sin(\theta), \theta)$, themselves the values

of integrals, really measures the total accumulation of the attenuation-coefficient function f along a particular line. Hence, the value of the Radon transform along a given line would be the same if all of the matter there were replaced with a homogeneous sample whose attenuation coefficient was the average of the actual sample. The integral in (3.4) is computing the average value of those averages. Thus, we see that this does not give us the value of $f(x_0, y_0)$. Instead, it gives us an “averaged out” or “smoothed out” version of the attenuation coefficient. For this reason, (3.4) is only a first step in recovering f .

Example 3.4. Suppose f_1 is the attenuation-coefficient function corresponding to a disc of radius $1/2$ centered at the origin and with constant density 1. Then, for every line $\ell_{0,\theta}$ through the origin, we have $\mathcal{R}f_1(0, \theta) = 1$. Consequently, $\mathcal{B}\mathcal{R}f_1(0, 0) = 1$.

Now suppose that f_2 is the attenuation-coefficient function for a ring of width $1/2$ consisting of all points at distances between $1/4$ and $3/4$ from the origin and with constant density 1 in this ring. Then, again, for every line $\ell_{0,\theta}$ through the origin, we get $\mathcal{R}f_2(0, \theta) = 1$. So, again, $\mathcal{B}\mathcal{R}f_2(0, 0) = 1$.

Thus, $\mathcal{B}\mathcal{R}f_1(0, 0) = \mathcal{B}\mathcal{R}f_2(0, 0) = 1$, even though $f_1(0, 0) = 1$ and $f_2(0, 0) = 0$. This illustrates the fact that the back projection of the Radon transform of a function does not necessarily reproduce the original function.

Example 3.5. We previously considered the function

$$f(x, y) := \begin{cases} 1 & \text{if } x^2 + y^2 \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

and computed that

$$\mathcal{R}f(t, \theta) = \int_{\ell_{t,\theta}} f ds = \begin{cases} 2\sqrt{1-t^2} & \text{if } |t| \leq 1, \\ 0 & \text{if } |t| > 1. \end{cases}$$

From this it follows that, for each point (x, y) , we have

$$\begin{aligned} & \mathcal{R}f(x\cos(\theta) + y\sin(\theta), \theta) \\ &= \begin{cases} 2\sqrt{1 - (x\cos(\theta) + y\sin(\theta))^2} & \text{if } |x\cos(\theta) + y\sin(\theta)| \leq 1, \\ 0 & \text{if } |x\cos(\theta) + y\sin(\theta)| > 1. \end{cases} \end{aligned}$$

It can be pretty difficult to figure out which values of θ satisfy the inequality $|x\cos(\theta) + y\sin(\theta)| \leq 1$ for an arbitrary point (x, y) . However, the maximum possible value of the expression $|x\cos(\theta) + y\sin(\theta)|$ is $\sqrt{x^2 + y^2}$, and we already know that we only care about points for which $\sqrt{x^2 + y^2} \leq 1$. Hence, $|x\cos(\theta) + y\sin(\theta)| \leq 1$ will hold for all the points (x, y) that we care about.

Now apply the back projection. Assuming, as we are, that $x^2 + y^2 \leq 1$, then

$$\begin{aligned}\mathcal{B}\mathcal{R}f(x, y) &= \frac{1}{\pi} \int_{\theta=0}^{\pi} \mathcal{R}f(x \cos(\theta) + y \sin(\theta), \theta) d\theta \\ &= \frac{1}{\pi} \int_{\theta=0}^{\pi} 2\sqrt{1 - (x \cos(\theta) + y \sin(\theta))^2} d\theta.\end{aligned}\quad (3.5)$$

As Figure 3.2 illustrates, where the graph of f is a circular column of height 1 with a flat top, the graph of $\mathcal{B}\mathcal{R}f$ is a circular column with the top rounded off. This is due to the “smoothing” effect of the back projection.




Fig. 3.2. For a test function whose graph is a cylinder or a cone, the back projection of the Radon transform of the function yields a rounded-off cylinder or a rounded-off cone.

Example 3.6. This time, let

$$f(x, y) := \begin{cases} 1 - \sqrt{x^2 + y^2} & \text{if } x^2 + y^2 \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

We previously computed that, for $-1 \leq t \leq 1$,

$$\begin{aligned}\mathcal{R}f(t, \theta) &= \int_{s=-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \\ &= \int_{s=-\sqrt{1-t^2}}^{\sqrt{1-t^2}} (1 - \sqrt{t^2 + s^2}) ds \\ &= \sqrt{1-t^2} - \frac{1}{2}t^2 \ln \left(\frac{1 + \sqrt{1-t^2}}{1 - \sqrt{1-t^2}} \right).\end{aligned}$$

As in the previous example, the inequality $|x \cos(\theta) + y \sin(\theta)| \leq 1$ is satisfied for all values of θ as long as we only look at points (x, y) in the unit disc. With that assumption, applying the back projection yields

$$\begin{aligned}\mathcal{B}\mathcal{R}f(x, y) &= \int_{\theta=0}^{\pi} \mathcal{R}f(x\cos(\theta) + y\sin(\theta), \theta) d\theta \\ &= \int_{\theta=0}^{\pi} \left\{ \sqrt{1 - (t_{\theta}(x, y))^2} - \frac{1}{2}(t_{\theta}(x, y))^2 \ln \left(\frac{1 + \sqrt{1 - (t_{\theta}(x, y))^2}}{1 - \sqrt{1 - (t_{\theta}(x, y))^2}} \right) \right\} d\theta,\end{aligned}$$

where, for brevity, $t_{\theta}(x, y) = x\cos(\theta) + y\sin(\theta)$. If we settle for the use of a Riemann sum, then we can approximate the values of $\mathcal{B}\mathcal{R}f(x, y)$ and generate an approximate graph of $\mathcal{B}\mathcal{R}f(x, y)$ over the unit disc.

In this example, shown in Figure 3.2, the graph of f is a cone while the graph of $\mathcal{B}\mathcal{R}f$ is a cone that has been rounded off, again illustrating the smoothing effect of the back projection.

In the last two examples above, we restricted the back projection to points (x, y) for which $x^2 + y^2 \leq 1$. For those points, the inequality $|x\cos(\theta) + y\sin(\theta)| \leq 1$ holds for every value of θ . In reality, even for points outside the unit circle, that is, even if $x^2 + y^2 > 1$, there are always *some values* of θ for which $|x\cos(\theta) + y\sin(\theta)| \leq 1$. (For instance, this follows from Exercise 3.2 below.) For those values of θ , the corresponding line $\ell_{(x\cos(\theta)+y\sin(\theta)),\theta}$ passes through the unit disc, and so provides a nonzero value of $\mathcal{R}f(x\cos(\theta) + y\sin(\theta), \theta)$. In turn, this nonzero value of the Radon transform contributes to a nonzero value of the back projection $\mathcal{B}\mathcal{R}f(x, y)$. Thus, in the examples above, we have effectively truncated the back projection by “filtering out” points outside the unit circle and excluding them from the smoothing process. The resulting images of $\mathcal{B}\mathcal{R}f$ are actually closer to the original f than if we had not done this filtering. This analysis raises the questions of how we can describe this filtering process mathematically and of whether there are other forms of filtering that will enhance the effort to recover the original attenuation-coefficient function.

Example 3.7. A phantom. As a final example for now, Figure 3.3 shows the back projection of the Radon transform of the Shepp–Logan phantom, introduced in Figure 2.4. This should reinforce the fact that the back projection is not, by itself, the inverse of the Radon transform. Nonetheless, as we shall see in Chapter 6, the back projection is a crucial ingredient in the image reconstruction process.

3.3 Exercises

3.1. Verify that, for given values of a , b , and θ , the line $\ell_{(a\cos(\theta)+b\sin(\theta)),\theta}$ passes through the point (a, b) .

3.2. Verify that, for every pair of real numbers a and b , the set of points in the plane that satisfy the polar-coordinate equation



Fig. 3.3. The back projection of the Radon transform of the Shepp–Logan phantom.

$$r = (a \cos(\theta) + b \sin(\theta)) \text{ for } 0 \leq \theta \leq \pi$$

forms a circle that passes through the origin as well as through the point with Cartesian coordinates (a, b) . Find the radius of the circle and the location of its center.

3.3. According to Exercise 3.2, for a function $F = F(r, \theta)$ whose inputs are polar coordinates, the value of $\mathcal{B}F(x, y)$ is the average value of the function F on the circle determined by the polar-coordinate equation $r = (x \cos(\theta) + y \sin(\theta))$ for $0 \leq \theta \leq \pi$. Use this to compute the following back projections.

- (i) Compute $\mathcal{B}g(x, y)$, where $g(r, \theta) := r \cos(\theta)$.
- (ii) Compute $\mathcal{B}h(x, y)$, where $h(r, \theta) := r \sin(\theta)$.

3.4. Prove Proposition 3.2, which asserts that the back projection is a linear transformation. That is, show that, for any two functions h_1 and h_2 and arbitrary constants c_1 and c_2 ,

$$\mathcal{B}(c_1 h_1 + c_2 h_2)(x, y) = c_1 \mathcal{B}h_1(x, y) + c_2 \mathcal{B}h_2(x, y)$$

for all values of x and y .

Complex Numbers

There is no real number a for which $a^2 + 1 = 0$. In order to develop an expanded number system that includes solutions to this simple quadratic equation, we *define* the “imaginary number” $i = \sqrt{-1}$. That is, this new number i is defined by the condition that $i^2 + 1 = 0$.

Since $i^2 = -1$, it follows that $i^3 = i^2 \cdot i = -1 \cdot i = -i$. Similarly, $i^4 = i^3 \cdot i = -i \cdot i = -i^2 = -(-1) = 1$, and $i^5 = i^4 \cdot i = 1 \cdot i = i$.

As a quick observation, notice that the equation $a^4 = 1$ now has not only the familiar solutions $a = \pm 1$ but also two “imaginary” solutions $a = \pm i$. Thus, there are four *fourth roots of unity*, namely, ± 1 and $\pm i$. The inclusion of the number i in our number system provides us with new solutions to many simple equations.

4.1 The complex number system

The complex number system, denoted by \mathbf{C} , is defined to be the set

$$\mathbf{C} = \{a + b \cdot i : a \text{ and } b \text{ are real numbers}\}.$$

To carry out arithmetic operations in \mathbf{C} , use the usual rules of commutativity, associativity, and distributivity, along with the definition $i^2 = -1$. Thus, $(a + bi) + (c + di) = (a + c) + (b + d)i$ and $(a + bi) \cdot (c + di) = (ac - bd) + (ad + bc)i$. Also, $-(a + bi) = -a + (-b)i$. Momentarily, we will look at division of one complex number by another.

A geometric view of \mathbf{C} . Each complex number $z = a + bi$ is determined by two real numbers. The number a is called the *real part* of z and is denoted by $\Re z = a$, while the number b is called the *imaginary part* and is denoted by $\Im z = b$. (It is important to keep in mind that the imaginary part of a complex number is actually a real number, which is the coefficient of i in the complex number.) In this sense, the complex number system

is “two-dimensional,” and so can be represented geometrically in the xy -plane, where we plot the real part of a complex number as the x -coordinate and the imaginary part of the complex number as the y -coordinate.

A real number a can also be written as $a = a + 0 \cdot i$, and so corresponds to the point $(a, 0)$ on the x -axis, which is known therefore as the *real axis*. Similarly, a purely imaginary number $b \cdot i = 0 + b \cdot i$ corresponds to the point $(0, b)$ on the y -axis, which is called the *imaginary axis*.

The distance to the origin from the point (a, b) , corresponding to the complex number $a + bi$, is equal to $\sqrt{a^2 + b^2}$. Accordingly, we make the following definition.

Definition 4.1. The *modulus* of the complex number $a + bi$, denoted by $|a + bi|$, is defined by

$$|a + bi| = \sqrt{a^2 + b^2}. \quad (4.1)$$

The modulus of a complex number is analogous to the absolute value of a real number. Indeed, for a real number $a = a + 0 \cdot i$, we get that $|a + 0 \cdot i| = \sqrt{a^2 + 0^2} = |a|$ in the usual sense. Note also that $|a + bi| = 0$ if, and only if, $a = b = 0$. A central observation is that

$$(a + bi) \cdot (a - bi) = a^2 + b^2 = |a + bi|^2.$$

With this in mind, we make another definition.

Definition 4.2. The *conjugate* of a complex number $a + bi$, denoted by $\overline{a + bi}$, is defined by

$$\overline{a + bi} = a - bi. \quad (4.2)$$

The central property, to repeat, is that

$$(a + bi) \cdot \overline{(a + bi)} = (a + bi) \cdot (a - bi) = a^2 + b^2 = |a + bi|^2. \quad (4.3)$$

The conjugate of a complex number is the key ingredient when it comes to the arithmetic operation of division. For example, notice that

$$(5 + 12i)(5 - 12i) = 5^2 + 12^2 = 13^2 = 169.$$

It follows from this that

$$\frac{1}{5 + 12i} = \frac{5 - 12i}{169} = \frac{5}{169} - \frac{12}{169}i,$$

whence

$$\frac{3 + 4i}{5 + 12i} = (3 + 4i) \cdot \frac{1}{5 + 12i} = (3 + 4i) \left(\frac{5}{169} - \frac{12}{169}i \right) = \frac{63}{169} - \frac{16}{169}i.$$

In general, we get that

$$\frac{1}{a+bi} = \frac{1}{a^2+b^2}(a-bi).$$

This enables us to divide any complex number by $(a+bi)$, provided that a and b are not both 0. The act of dividing by the nonzero complex number $(a+bi)$ is re-expressed as multiplication by $(a-bi)$ and division by the nonzero real number (a^2+b^2) .

4.2 The complex exponential function

Consider these well-known Taylor series:

$$\begin{aligned}\cos(x) &= \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n}}{(2n)!} = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \dots, \\ \sin(x) &= \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)!} = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \dots,\end{aligned}$$

and

$$\exp(x) = e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!} = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \frac{x^4}{4!} + \dots.$$

Substitute $x = i\theta$ (where θ is assumed to be a real number) into the series for $\exp(x)$ to get (formally)

$$\begin{aligned}\exp(i\theta) &= e^{i\theta} = \sum_{n=0}^{\infty} \frac{(i\theta)^n}{n!} \\ &= 1 + (i\theta) + \frac{(i\theta)^2}{2!} + \frac{(i\theta)^3}{3!} + \frac{(i\theta)^4}{4!} + \frac{(i\theta)^5}{5!} + \dots \\ &= 1 + i\theta - \frac{\theta^2}{2!} - i\frac{\theta^3}{3!} + \frac{\theta^4}{4!} + i\frac{\theta^5}{5!} - \dots \\ &= \left[1 - \frac{\theta^2}{2!} + \frac{\theta^4}{4!} - \frac{\theta^6}{6!} + \dots \right] + i \cdot \left[\theta - \frac{\theta^3}{3!} + \frac{\theta^5}{5!} - \dots \right] \\ &= \cos(\theta) + i \cdot \sin(\theta).\end{aligned}$$

Euler's formula. The remarkable relationship

$$e^{i\theta} = \cos(\theta) + i \cdot \sin(\theta) \tag{4.4}$$

between an imaginary power of e and the sine and cosine functions is known as *Euler's formula* after its discoverer, Leonhard Euler (1707–1783).

Some examples of Euler's formula that are of special interest are

$$\begin{aligned} e^{i\pi} &= \cos(\pi) + i \sin(\pi) = -1 + i(0) = -1, \\ e^{i\pi/2} &= \cos(\pi/2) + i \sin(\pi/2) = 0 + i(1) = i, \text{ and} \\ e^{2\pi i} &= \cos(2\pi) + i \sin(2\pi) = 1. \end{aligned}$$

In general, for any real number θ , $\Re(e^{i\theta}) = \cos(\theta)$ and $\Im(e^{i\theta}) = \sin(\theta)$. Hence,

$$|e^{i\theta}| = \sqrt{\cos^2(\theta) + \sin^2(\theta)} = 1$$

regardless of the value of θ .

Geometrically, for real numbers r and θ , we get

$$re^{i\theta} = r(\cos(\theta) + i \sin(\theta)) = r \cos(\theta) + ir \sin(\theta)$$

so that the complex number $re^{i\theta}$ corresponds to the point in the xy -plane with Cartesian coordinates $(r \cos(\theta), r \sin(\theta))$. This same point has *polar coordinates* r and θ . For this reason, the form $re^{i\theta}$ is called the *polar form* of the complex number $(r \cos(\theta) + i \sin(\theta))$.

Observe that $|re^{i\theta}| = |r|$. So, $|r|$ is the modulus of $re^{i\theta}$. The number θ , viewed as an angle now, is called the *argument* of the complex number $re^{i\theta}$.

A simple computation shows that $(re^{i\theta}) \cdot (Re^{i\phi}) = r \cdot R \cdot e^{i(\theta+\phi)}$. Thus, when we multiply two complex numbers, expressed here in their polar forms, the modulus of the product is equal to the product of the individual moduli and the argument of the product is the *sum* of the individual arguments.

DeMoivre's law. When the equation $(e^{i\theta})^n = e^{in\theta}$ is translated into standard complex number form, we get

$$[\cos(\theta) + i \sin(\theta)]^n = \cos(n\theta) + i \sin(n\theta). \quad (4.5)$$

This is called *DeMoivre's law*.

It is just a short step now to define the exponential function for every complex number. Namely, for any complex number $z = a + bi$,

$$e^z = e^{a+bi} = e^a \cdot e^{bi} = e^a \cdot (\cos(b) + i \sin(b)). \quad (4.6)$$

The complex exponential function has many interesting and important properties, not least of which is that it is a conformal mapping. (See [15] and [37], for example.) For our purposes, the periodicity property of the exponential function is central.

4.3 Wave functions

Waves, periodicity, and frequency. For a fixed real number ω , the functions $t \mapsto \cos(\omega t)$ and $t \mapsto \sin(\omega t)$ have period $2\pi/\omega$ and frequency $\omega/(2\pi)$. If the value of ω is large, then these functions have high frequency and short wavelength; lower values of ω yield functions with longer waves and lower frequencies.

For a given real number ω , consider the function $E_\omega(t) = e^{i\omega t}$. Then

$$E_\omega(t + 2\pi/\omega) = e^{i\omega(t+2\pi/\omega)} = e^{i\omega t + 2\pi i} = e^{i\omega t} \cdot e^{2\pi i} = e^{i\omega t} = E_\omega(t),$$

where we have used the fact that $e^{2\pi i} = 1$. Thus, the function E_ω is *periodic* with period $2\pi/\omega$ and frequency $\omega/(2\pi)$.

Another point of view is to see that

$$E_\omega(t) = e^{i\omega t} = \cos(\omega t) + i \sin(\omega t)$$

which is a sum of two periodic functions each having period $2\pi/\omega$. So the sum also is periodic with that same period.

A signal, such as a radio, light, or sound wave, can, in principle, be decomposed into its components of specific frequencies. For example, we might try to decompose the sound wave from a musical instrument into the high notes, the mid-tones, bass, and so on. That is, a signal, viewed as a function propagated over time, might also be viewed as a composite of functions of the form $\cos(\omega t)$ and $\sin(\omega t)$, or, more compactly, $E_\omega(t) = e^{i\omega t}$, for various values of ω . The pertinent issue becomes how to determine which values of ω correspond to the different frequency components of a given signal $f(t)$, and, for each such ω , to find the amplitude of the associated component.

Fourier analysis. As a first step toward finding the different frequencies that make up a given signal, consider the “harmonic frequencies” $n/(2\pi)$ where n is an integer. The basic periodic functions having these frequencies are of the form $E_n(t) = e^{int}$. Two basic computations involving these functions are the following.

$$\text{For } m \neq n, \int_0^{2\pi} e^{imt} \cdot e^{-int} dt = \int_0^{2\pi} e^{i(m-n)t} dt = \left. \frac{e^{i(m-n)t}}{i(m-n)} \right|_0^{2\pi} = 0,$$

since $e^{i(m-n)2\pi} = e^0 = 1$.

$$\text{For } m = n, \int_0^{2\pi} e^{imt} \cdot e^{-int} dt = \int_0^{2\pi} e^0 dt = 2\pi.$$

These computations mean that, if a signal has the form

$$f(t) = \sum_{n \in \Lambda} c_n e^{int}$$

for integers n in some set Λ , then the amplitudes $\{c_n\}$ are given by

$$c_n = \frac{1}{2\pi} \int_0^{2\pi} f(t) \cdot e^{-int} dt \text{ for each } n \in \Lambda.$$

Applying this analysis to an arbitrary function f that is periodic on the interval $0 \leq t \leq 2\pi$, we *define* the n th Fourier coefficient of f by

$$\hat{f}(n) = \frac{1}{2\pi} \int_0^{2\pi} f(t) \cdot e^{-int} dt. \quad (4.7)$$

We have in mind that f somehow can be represented, at least on the interval $0 \leq t \leq 2\pi$, by the sum $\sum \hat{f}(n)e^{int}$. Determining precisely when such a representation is valid (and even what is meant by “representation” or “valid”) is part of the subject of Fourier series, initiated in 1811 by Jean-Baptiste Joseph Fourier (1768–1830).

In our analysis of X-ray attenuation-coefficient functions, we will use a different tool, also pioneered by Fourier, called the *Fourier transform*. The Fourier transform is analogous to the Fourier series but it allows for all possible frequencies and does not assume that the signal is periodic.

4.4 Exercises

4.1. (a) Let z and w be complex numbers. Show that $\overline{(z+w)} = \bar{z} + \bar{w}$ and that $\overline{z \cdot w} = \bar{z} \cdot \bar{w}$.

(b) Show that $\overline{re^{i\theta}} = re^{-i\theta}$.

(c) Use (4.3) to express $\frac{1}{3+4i}$ in the form $a+bi$.

4.2. Evaluate the integral $\int_0^{2\pi} e^{i\theta} d\theta$.

4.3. (a) Show that $\int_0^{\infty} e^{-\lambda x} \cdot e^{-i\omega x} dx = \frac{\lambda - i\omega}{\lambda^2 + \omega^2}$, where λ and ω are real numbers with $\lambda \neq 0$.

(b) Use part (a) to show that $\int_{-\infty}^{\infty} e^{-\lambda|x|} \cdot e^{-i\omega x} dx = \frac{2\lambda}{\lambda^2 + \omega^2}$.

4.4. (a) Use Euler’s formula to show that

$$e^{-i\omega T} - e^{i\omega T} = -2i \sin(\omega T)$$

for all real numbers ω and T .

(b) Now show that $\int_{-T}^T e^{-i\omega x} dx = 2 \frac{\sin(\omega T)}{\omega}$ for $\omega \neq 0$.

(c) Plot the graph of $S(\omega) = 2 \frac{\sin(\omega T)}{\omega}$, as a function of the real variable ω , for $T = 1$, for $T = 0.5$, and for $T = 0.2$. (Note that setting $S(0) = 2T$ yields a continuous function on the real line.)

4.5. Use a change of variables to show that

$$\int_{-\infty}^{\infty} f(x - \alpha) e^{-i\omega x} dx = e^{-i\omega\alpha} \int_{-\infty}^{\infty} f(x) e^{-i\omega x} dx$$

for all real numbers ω and α .

The Fourier Transform

5.1 Definition and examples

For a given function f such that $\int_{-\infty}^{\infty} |f(x)| dx < \infty$, the *Fourier transform* of f is defined, for each real number ω , by

$$\mathcal{F}f(\omega) := \int_{-\infty}^{\infty} f(x) e^{-i\omega x} dx. \quad (5.1)$$

The idea behind this definition is that, for each value of ω , the value of $\mathcal{F}f(\omega)$ captures the component of f that has the frequency $\omega/(2\pi)$ (and period $2\pi/\omega$).

Example 5.1. The Fourier transform of a Gaussian. Let $f(x) = e^{-Ax^2}$, for some positive constant $A > 0$. Then we have

$$\mathcal{F}f(\omega) = \sqrt{\frac{\pi}{A}} e^{-\frac{\omega^2}{4A}}. \quad (5.2)$$

To prove this, we first need the following fact.

Lemma 5.2. For $A \neq 0$, we have $\int_{-\infty}^{\infty} e^{-Ax^2} dx = \sqrt{\frac{\pi}{A}}$.

Proof. Squaring the integral, we get

$$\begin{aligned} \left(\int_{-\infty}^{\infty} e^{-Ax^2} dx \right)^2 &= \left(\int_{-\infty}^{\infty} e^{-Ax^2} dx \right) \left(\int_{-\infty}^{\infty} e^{-Ax^2} dx \right) \\ &= \left(\int_{-\infty}^{\infty} e^{-Ax^2} dx \right) \left(\int_{-\infty}^{\infty} e^{-Ay^2} dy \right) \end{aligned}$$

$$\begin{aligned}
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-A(x^2+y^2)} dx dy \\
\text{(polar coordinates)} &= \int_{\theta=0}^{2\pi} \int_{r=0}^{\infty} e^{-Ar^2} r dr d\theta \\
&= \int_{\theta=0}^{2\pi} \left(\lim_{b \rightarrow \infty} \frac{1 - e^{-Ab^2}}{2A} \right) d\theta \\
&= \int_{\theta=0}^{2\pi} \frac{1}{2A} d\theta \\
&= \frac{\pi}{A}.
\end{aligned}$$

Taking square roots proves the lemma. \square

Now to compute the Fourier transform for this example. For each ω ,

$$\begin{aligned}
\mathcal{F}f(\omega) &= \int_{-\infty}^{\infty} e^{-Ax^2} e^{-i\omega x} dx \\
&= \int_{-\infty}^{\infty} e^{-A(x^2+i\omega x/A)} dx \\
\text{(complete the square)} &= \int_{-\infty}^{\infty} e^{-A(x^2+i\omega x/A+(i\omega/2A)^2)} e^{A(i\omega/2A)^2} dx \\
&= e^{-\omega^2/4A} \int_{-\infty}^{\infty} e^{-A(x+i\omega/2A)^2} dx \\
&= e^{-\omega^2/4A} \int_{-\infty}^{\infty} e^{-Au^2} du \quad \text{with } u = x + i\omega/2A \\
&= \sqrt{\frac{\pi}{A}} e^{-\omega^2/4A} \quad \text{by the lemma.}
\end{aligned}$$

This establishes the result we were after. \square

Observe that if we take $A = 1/2$, then $f(x) = e^{-x^2/2}$ and $\mathcal{F}f(\omega) = \sqrt{2\pi}e^{-\omega^2/2}$, a constant multiple of f itself.

Additional examples are considered in the exercises. Let us look at some basic properties of the Fourier transform.

5.2 Properties and applications

Additivity. Because the integral of a sum of functions is equal to the sum of the integrals of the functions separately, it follows that

$$\mathcal{F}(f+g)(\omega) = \mathcal{F}f(\omega) + \mathcal{F}g(\omega), \quad (5.3)$$

for all integrable functions f and g and every real number ω .

Constant multiples. Because the integral of $k \cdot f$ is equal to k times the integral of f , we see that

$$\mathcal{F}(kf)(\omega) = k \cdot \mathcal{F}f(\omega), \quad (5.4)$$

for all integrable functions f , all (complex) numbers k , and every real number ω .

The two properties (5.3) and (5.4) taken together prove the following proposition.

Proposition 5.3. *The Fourier transform acts as a linear transformation on the space of all absolutely integrable functions. That is, for two such functions f and g and any constants α and β , we get that*

$$\mathcal{F}(\alpha f + \beta g) = \alpha \mathcal{F}(f) + \beta \mathcal{F}(g). \quad (5.5)$$

Shifting/translation. For an integrable function f and fixed real number α , let $g(x) = f(x - \alpha)$. (So the graph of g is the graph of f shifted or translated to the right by α units.) Then

$$\mathcal{F}g(\omega) = e^{-i\omega\alpha} \mathcal{F}f(\omega). \quad (5.6)$$

Proof. For each ω , we get

$$\begin{aligned} \mathcal{F}g(\omega) &= \int_{-\infty}^{\infty} f(x - \alpha) e^{-i\omega x} dx \\ (\text{let } u = x - \alpha) &= \int_{-\infty}^{\infty} f(u) e^{-i\omega(u+\alpha)} du \\ &= e^{-i\omega\alpha} \int_{-\infty}^{\infty} f(u) e^{-i\omega u} du \\ &= e^{-i\omega\alpha} \mathcal{F}f(\omega) \text{ as claimed.} \quad \square \end{aligned}$$

Since the graph of g is a simple translation of the graph of f , the magnitude of the component of g at any given frequency is the same as that of f . However, the components occur at different places in the two signals, so there is a phase shift or delay in the Fourier transform. Also, the fixed translation α encompasses more cycles of a wave at a higher frequency than at a lower one. Therefore, the larger the value of α is relative to the wavelength $2\pi/\omega$ the larger will be the phase delay in the transform. That is, the phase delay in the transform is proportional to ω , which explains the factor of $e^{-i\omega\alpha}$ in the transform of the shifted function g .

Shifting/modulation. For a given function f and a fixed real number ω_0 , let $h(x) = e^{i\omega_0 x} f(x)$. (So h “modulates” f by multiplying f by a periodic function of a fixed frequency $\omega_0/2\pi$.) Then

$$\mathcal{F}h(\omega) = \mathcal{F}f(\omega - \omega_0). \quad (5.7)$$

Proof. For each ω , we get

$$\begin{aligned}\mathcal{F}h(\omega) &= \int_{-\infty}^{\infty} e^{i\omega_0 x} f(x) e^{-i\omega x} dx \\ &= \int_{-\infty}^{\infty} f(x) e^{-i(\omega - \omega_0)x} dx \\ &= \mathcal{F}f(\omega - \omega_0) \quad (\text{by definition!}).\end{aligned}\quad \square$$

These two shifting properties show that a translation of f results in a modulation of $\mathcal{F}f$ while a modulation of f produces a translation of $\mathcal{F}f$.

Scaling. For a given function f and a fixed real number $a \neq 0$, let $\phi(x) = f(ax)$. Then

$$\mathcal{F}\phi(\omega) = \frac{1}{|a|} \mathcal{F}f(\omega/a). \quad (5.8)$$

Proof. Assume that $a > 0$ for now. For each ω , we get

$$\begin{aligned}\mathcal{F}\phi(\omega) &= \int_{-\infty}^{\infty} f(ax) e^{-i\omega x} dx \\ (\text{let } u = ax) &= \frac{1}{a} \int_{-\infty}^{\infty} f(u) e^{-i\omega u/a} du \\ &= \frac{1}{a} \int_{-\infty}^{\infty} f(u) e^{-i(\omega/a)u} du \\ &= \frac{1}{a} \mathcal{F}f(\omega/a) \quad (\text{by definition!}).\end{aligned}$$

A similar argument applies when $a < 0$. (Why do we get a factor of $1/|a|$ in this case?) \square

Even and odd functions. A function f , defined on the real line, is *even* if $f(-x) = f(x)$ for every x . Similarly, a function g is *odd* if $g(-x) = -g(x)$ for every x . For example, the *cosine* function is even while the *sine* function is odd.

Using Euler's formula (4.4), $e^{i\theta} = \cos(\theta) + i\sin(\theta)$, we may write the Fourier transform of a suitable real-valued function f as

$$\mathcal{F}f(\omega) = \int_{-\infty}^{\infty} f(x) \cos(\omega x) dx - i \cdot \int_{-\infty}^{\infty} f(x) \sin(\omega x) dx.$$

Now, if f is even, then, for fixed ω , the function $x \mapsto f(x) \sin(\omega x)$ is odd, whence $\int_{-\infty}^{\infty} f(x) \sin(\omega x) dx = 0$. Thus, *an even function has a real-valued Fourier transform.*

Similarly, if f is odd, then $x \mapsto f(x) \cos(\omega x)$ is also odd for each fixed ω . Thus, $\int_{-\infty}^{\infty} f(x) \cos(\omega x) dx = 0$. It follows that *an odd function has a purely imaginary Fourier transform.*

Transform of the complex conjugate. For a complex-number-valued function f defined on the real line \mathbf{R} , the complex conjugate of f is the function \overline{f} defined by

$$\overline{f}(x) = \overline{f(x)} \text{ for every real number } x. \quad (5.9)$$

To uncover the relationship between the Fourier transform of \overline{f} and that of f , let ω be an arbitrary real number. Then we have

$$\begin{aligned} \mathcal{F}\overline{f}(\omega) &= \int_{-\infty}^{\infty} \overline{f(x)} e^{-i\omega x} dx \\ &= \int_{-\infty}^{\infty} \overline{f(x)} e^{i(-\omega)x} dx \\ &= \int_{-\infty}^{\infty} \overline{f(x)} \overline{e^{-i(-\omega)x}} dx \\ &= \overline{\int_{-\infty}^{\infty} f(x) e^{-i(-\omega)x} dx} \\ &= \overline{\mathcal{F}f(-\omega)} \\ &= \overline{\mathcal{F}\overline{f}(-\omega)}. \end{aligned}$$

This proves the following proposition.

Proposition 5.4. For an integrable function f defined on the real line, and for every real number ω ,

$$\mathcal{F}\overline{f}(\omega) = \overline{\mathcal{F}\overline{f}(-\omega)}. \quad (5.10)$$

Example 5.5. Let $f(x) = \begin{cases} 1 & \text{if } -1 \leq x \leq 1, \\ 0 & \text{if } |x| > 1. \end{cases}$

Then the Fourier transform of f is $\mathcal{F}f(\omega) = 2 \sin(\omega)/\omega$. Now let $\phi(x) = f(ax)$, where $a > 0$. That is, $\phi(x) = \begin{cases} 1 & \text{if } -1/a \leq x \leq 1/a, \\ 0 & \text{if } |x| > 1/a. \end{cases}$

So we already know from earlier work that $\mathcal{F}\phi(\omega) = 2 \sin(\omega/a)/\omega$, which is the same as $(1/a)\mathcal{F}f(\omega/a)$. This agrees with the scaling result (5.8).

Example 5.6. Let $f(x) = \begin{cases} 1 & \text{if } -1 \leq x \leq 1, \\ 0 & \text{if } |x| > 1 \end{cases}$ as in the previous example. So, again, the Fourier transform of f is $\mathcal{F}f(\omega) = 2 \sin(\omega)/\omega$. Now let

$$g(x) = f(x-2) = \begin{cases} 1 & \text{if } 1 \leq x \leq 3, \\ 0 & \text{if } x < 1 \text{ or } x > 3. \end{cases}$$

By the shifting/translation result (5.6), we get

$$\mathcal{F}g(\omega) = e^{-2i\omega} \mathcal{F}f(\omega) = 2e^{-2i\omega} \frac{\sin(\omega)}{\omega}.$$

Example 5.7. As an application of the shifting/modulation property (5.7), observe that

$$\cos(\omega_0 x) = (1/2)(e^{i\omega_0 x} + e^{-i\omega_0 x}).$$

Thus, for any suitable f , take $h(x) = f(x) \cos(\omega_0 x)$. That is,

$$h(x) = (1/2)e^{i\omega_0 x} f(x) + (1/2)e^{i(-\omega_0)x} f(x).$$

It follows that

$$\mathcal{F}h(\omega) = (1/2)\mathcal{F}f(\omega - \omega_0) + (1/2)\mathcal{F}f(\omega + \omega_0).$$

For a specific example, let $f(x) = \begin{cases} 1 & \text{if } -1 \leq x \leq 1, \\ 0 & \text{if } |x| > 1 \end{cases}$ as in the previous examples. So, again, the Fourier transform of f is $\mathcal{F}f(\omega) = 2 \sin(\omega)/\omega$. With $h(x) = f(x) \cos(\omega_0 x)$, we get

$$\begin{aligned} \mathcal{F}h(\omega) &= (1/2)\mathcal{F}f(\omega - \omega_0) + (1/2)\mathcal{F}f(\omega + \omega_0) \\ &= (1/2) \left\{ 2 \frac{\sin(\omega - \omega_0)}{(\omega - \omega_0)} + 2 \frac{\sin(\omega + \omega_0)}{(\omega + \omega_0)} \right\} \\ &= \frac{\sin(\omega - \omega_0)}{(\omega - \omega_0)} + \frac{\sin(\omega + \omega_0)}{(\omega + \omega_0)} \end{aligned}$$

The graph of $\mathcal{F}f$ in this example has a main peak of height 2 centered at $\omega = 0$ with smaller ripples at the edges. The graph of $\mathcal{F}h$ has two main peaks, each of height 1, centered at $\omega = \pm\omega_0$ with smaller ripples at the edges.

5.3 Heaviside and Dirac δ

Definition 5.8. The *Heaviside function* H is defined by

$$H(x) = \begin{cases} 0 & \text{if } x < 0, \\ 1 & \text{if } x > 0. \end{cases} \quad (5.11)$$

Technically, $H(0)$ is not defined, which, nonetheless, does not stop us from writing formulas like

$$H(x-a) - H(x-b) = \begin{cases} 1 & \text{if } a < x < b, \\ 0 & \text{otherwise.} \end{cases}$$

Formally, the value of the Heaviside function jumps by 1 at 0; so we think of the differential dH as having the value 1 at 0 and the value 0 everywhere else. Formally,

this suggests that $\int_{-\infty}^{\infty} f(x) dH = f(0)$ for any function f . (The intuitive idea is that dH is 0 except at 0 where $dH = 1$, so the value of the integral is just $f(0) \cdot 1 = f(0)$; this is an example of a general theory known as *Riemann–Stieltjes integration*.)

Definition 5.9. The *Dirac δ function*, denoted simply by δ , is defined to be the formal derivative $\delta = dH/dx$ of the Heaviside function with respect to x .

Of course, since H is constant except at 0, we get that $\delta(x) = 0$ except at $x = 0$ where H does not have a derivative (or the derivative is essentially ∞). So apparently $\delta(0)$ does not make sense. Alternatively, we can think of $\delta(0)$ as being equal to ∞ . So the “graph” of δ basically consists of an infinitely tall spike at 0. (*Warning:* We should not expect a computer to graph this for us!) The Dirac δ function is also known as the *impulse function*, perhaps because its graph evokes the image of a sudden impulse of energy that immediately dies out again.

In the context of integration, δ has a property of central importance. Since $\delta = dH/dx$, then it follows that $dH = \delta(x) dx$. Hence, for any f , we get

$$\int_{-\infty}^{\infty} f(x) \delta(x) dx = \int_{-\infty}^{\infty} f(x) dH = f(0). \quad (5.12)$$

In particular, substituting $f(x) = 1$ into (5.12) yields the formula

$$\int_{-\infty}^{\infty} \delta(x) dx = 1. \quad (5.13)$$

Thus, the graph of δ is an infinitely tall spike with no width but which has underneath it an area of 1!

We use (5.12) to compute the Fourier transform of the δ function. Namely,

$$\mathcal{F} \delta(\omega) = \int_{-\infty}^{\infty} \delta(x) e^{-i\omega x} dx = e^{-i\omega \cdot 0} = e^0 = 1. \quad (5.14)$$

So the Fourier transform of δ is the constant function equal to 1.

Neither H nor δ can exist in reality as a physical signal. The δ impulse would have to appear out of nowhere (like a Big Bang?) and then die instantly. The Heaviside signal would also have to appear out of nowhere and then propagate without loss of energy for eternity. Moreover, neither of these functions is even properly defined as a *function*. Instead, they fall into a class known as *generalized functions* or *distributions*. Nonetheless, they are useful idealized mathematical and physical constructs and we shall use them without further troubling ourselves over their precise natures.

There are other approaches to the definition of δ , for instance as the integral in (5.19), or as the limit

$$\delta(x) = \lim_{a \rightarrow 0^+} \left(\frac{1}{2a} \right) \square_a(x), \quad (5.15)$$

where $\square_a(x) = \begin{cases} 1 & \text{if } -a \leq x \leq a, \\ 0 & \text{if } |x| > a. \end{cases}$ We will not concern ourselves here with these alternatives. (The function \square_a is also called the *characteristic function* of the interval $[-a, a]$ and is often denoted by $\chi_{[-a, a]}$. We will use the symbol \square to remind us of the shape of its graph.)

5.4 Inversion of the Fourier transform

In the definition of the Fourier transform (5.1), there may seem to be no reason *a priori* to have used the exponent $-i\omega x$ rather than its opposite $+i\omega x$. This prompts the following definition.

Definition 5.10. For a function g for which $\int_{-\infty}^{\infty} |g(\omega)| dx < \infty$, the *inverse Fourier transform* of g is defined, for each real number x , by

$$\mathcal{F}^{-1}g(x) := \frac{1}{2\pi} \int_{\omega=-\infty}^{\infty} g(\omega) e^{i\omega x} d\omega. \quad (5.16)$$

The reasons for the factor $1/2\pi$ as well as for the name of this transform are made clear by the following essential theorem.

Theorem 5.11. Fourier inversion theorem. *If f is continuous on the real line and $\int_{-\infty}^{\infty} |f(x)| dx < \infty$, then*

$$\mathcal{F}^{-1}(\mathcal{F}f)(x) = f(x) \text{ for all } x. \quad (5.17)$$

To prove this, we first need a lemma.

Lemma 5.12. *For real numbers t and x ,*

$$\int_{\omega=-\infty}^{\infty} e^{i\omega(t-x)} d\omega = \lim_{\varepsilon \rightarrow 0} \sqrt{\frac{\pi}{\varepsilon}} \cdot e^{-(x-t)^2/(4\varepsilon)}. \quad (5.18)$$

Proof of Lemma 5.12. For each real number ω we have $\lim_{\varepsilon \rightarrow 0} e^{-\varepsilon\omega^2} = 1$. Thus, for any real numbers t and x ,

$$\begin{aligned} \int_{\omega=-\infty}^{\infty} e^{i\omega(t-x)} d\omega &= \lim_{\varepsilon \rightarrow 0} \int_{\omega=-\infty}^{\infty} e^{-\varepsilon\omega^2} \cdot e^{i\omega(t-x)} d\omega \\ &= \lim_{\varepsilon \rightarrow 0} \int_{\omega=-\infty}^{\infty} e^{-\varepsilon\omega^2} \cdot e^{-i\omega(x-t)} d\omega \end{aligned}$$

Previously, we computed the Fourier transform of a Gaussian function e^{-Ax^2} to be $\sqrt{\pi/A} e^{-t^2/(4A)}$. Applying this formula to the Gaussian function $e^{-\varepsilon\omega^2}$ (so $A = \varepsilon$ and ω is in place of x) and evaluating the resulting Fourier transform at $(x-t)$ yields

$$\int_{\omega=-\infty}^{\infty} e^{-\varepsilon\omega^2} \cdot e^{-i\omega(x-t)} d\omega = \sqrt{\frac{\pi}{\varepsilon}} e^{-(x-t)^2/(4\varepsilon)}.$$

Substituting this into the previous calculation, we get

$$\begin{aligned} \int_{\omega=-\infty}^{\infty} e^{i\omega(t-x)} d\omega &= \lim_{\varepsilon \rightarrow 0} \int_{\omega=-\infty}^{\infty} e^{-\varepsilon\omega^2} \cdot e^{-i\omega(x-t)} d\omega \\ &= \lim_{\varepsilon \rightarrow 0} \sqrt{\frac{\pi}{\varepsilon}} e^{-(x-t)^2/(4\varepsilon)}, \end{aligned}$$

which establishes the lemma. \square

Proof of Theorem 5.11. For a continuous function f that is absolutely integrable on the real line, let F denote the Fourier transform $\mathcal{F}f$. For any real number t , we get

$$\begin{aligned} \mathcal{F}^{-1}F(t) &= \frac{1}{2\pi} \int_{\omega=-\infty}^{\infty} F(\omega) e^{i\omega t} d\omega \\ &= \frac{1}{2\pi} \int_{\omega=-\infty}^{\infty} \int_{x=-\infty}^{\infty} f(x) e^{-i\omega x} e^{i\omega t} dx d\omega \quad (\text{by definition of } F) \\ &= \frac{1}{2\pi} \int_{x=-\infty}^{\infty} \int_{\omega=-\infty}^{\infty} f(x) e^{i\omega(t-x)} d\omega dx \\ &= \frac{1}{2\pi} \int_{x=-\infty}^{\infty} f(x) \left(\int_{\omega=-\infty}^{\infty} e^{i\omega(t-x)} d\omega \right) dx \\ &= \frac{1}{2\pi} \cdot \lim_{\varepsilon \rightarrow 0} \int_{x=-\infty}^{\infty} f(x) \sqrt{\frac{\pi}{\varepsilon}} e^{-(x-t)^2/(4\varepsilon)} dx \quad (\text{by the lemma}). \end{aligned}$$

Set $y = (x-t)/(2\sqrt{\varepsilon})$, so $x = t + 2\sqrt{\varepsilon}y$ and $dx = 2\sqrt{\varepsilon}dy$:

$$\begin{aligned} \mathcal{F}^{-1}F(t) &= \frac{1}{2\pi} \cdot \lim_{\varepsilon \rightarrow 0} \left(\sqrt{\frac{\pi}{\varepsilon}} \cdot 2\sqrt{\varepsilon} \right) \int_{y=-\infty}^{\infty} f(t + 2\sqrt{\varepsilon}y) e^{-y^2} dy \\ &= \frac{1}{\sqrt{\pi}} \cdot f(t) \cdot \int_{y=-\infty}^{\infty} e^{-y^2} dy \quad (\text{since } \lim_{\varepsilon \rightarrow 0} f(t + 2\sqrt{\varepsilon}y) = f(t) \forall y) \\ &= \frac{1}{\sqrt{\pi}} \cdot f(t) \cdot \sqrt{\pi} \\ &= f(t). \end{aligned}$$

This is the desired result. \square

A bit more generally, if f has a point of discontinuity at α and if the one-sided limits $\lim_{x \rightarrow \alpha^-} f(x)$ and $\lim_{x \rightarrow \alpha^+} f(x)$ both exist, then

$$\mathcal{F}^{-1}(\mathcal{F}f)(\alpha) = (1/2) \left(\lim_{x \rightarrow \alpha^-} f(x) + \lim_{x \rightarrow \alpha^+} f(x) \right).$$

The proof of this claim is left as an exercise.

Example 5.13. Right away let's apply the inverse Fourier transform to a function that is not even a function! We saw before that the Fourier transform of the δ function is the constant function to 1. That means that the δ function is the inverse Fourier transform of the constant function 1. That is,

$$\delta(x) = \mathcal{F}^{-1}1(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{i\omega x} d\omega \quad (5.19)$$

for all x . Interestingly, had we elected to *define* the δ function by this integral, then we could have used that to provide a simpler proof of the Fourier inversion theorem, Theorem 5.11. (See (2.12) in [9].)

Example 5.14. Inverse Fourier transform of a Gaussian. We have already seen that the Fourier transform of a Gaussian is another Gaussian. Thus, the same will be true for the inverse Fourier transform of a Gaussian. Specifically,

$$\begin{aligned} \mathcal{F}^{-1}(e^{-B\omega^2})(x) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-B\omega^2} e^{i\omega x} d\omega \\ &= \sqrt{\frac{1}{4\pi B}} e^{-x^2/(4B)}. \end{aligned} \quad (5.20)$$

In particular, if we take $B = 1/(4A)$ and multiply both sides by $\sqrt{\pi/A}$ we see that we get e^{-Ax^2} , illustrating again the inverse relationship between the transforms (5.1) and (5.16).

Example 5.15. Inverse Fourier transform of a square wave. As in (5.15), let

$$\square_a(x) = \begin{cases} 1 & \text{if } -a \leq x \leq a, \\ 0 & \text{if } |x| > a. \end{cases} \quad (5.21)$$

Then the inverse Fourier transform of \square_a is given by

$$\begin{aligned}
\mathcal{F}^{-1}\square_a(x) &= \frac{1}{2\pi} \int_{\omega=-\infty}^{\infty} \square_a(\omega) e^{i\omega x} d\omega \\
&= \frac{1}{2\pi} \int_{\omega=-a}^a e^{i\omega x} d\omega \\
&= \frac{1}{2\pi} \frac{e^{iax} - e^{-iax}}{ix} \\
&= \frac{1}{\pi} \frac{\sin(ax)}{x},
\end{aligned} \tag{5.22}$$

where Euler's formula (4.4) was used in the final step.

5.5 Multivariable forms

We will need two generalizations of the Fourier transform in order to apply it within the context of medical imaging, where the functions involved are defined in the plane using either Cartesian or polar coordinates.

First, for a function $h(r, \theta)$ defined using polar coordinates in the plane, we simply apply the one-variable Fourier transform in the radial variable (r) only. This gives the following definitions.

Definition 5.16. For $h(r, \theta)$, we define the Fourier transform of h at a point (t, θ) (also in polar coordinates and with the same angle θ) by

$$\mathcal{F}h(t, \theta) = \int_{r=-\infty}^{\infty} h(r, \theta) e^{-irt} dr. \tag{5.23}$$

Definition 5.17. Similarly, the inverse Fourier transform of h is given by

$$\mathcal{F}^{-1}h(t, \theta) = \frac{1}{2\pi} \int_{r=-\infty}^{\infty} h(r, \theta) e^{irt} dr. \tag{5.24}$$

The inverse relationship between these transforms is clear because the variable θ is treated as a constant in the computations. Significantly, this generalization of the Fourier transform and its inverse applies to the Radon transform of a function f since $\mathcal{R}f$ is defined at the points (t, θ) corresponding to the lines $\ell_{t, \theta}$.

The second generalization of the Fourier transform is applied to functions g defined using Cartesian coordinates in the plane. For instance, $g(x, y)$ might represent the X-ray attenuation coefficient of a tissue sample located at the point (x, y) .

Definition 5.18. For such a function g , we define the Fourier transform of g evaluated at the point (X, Y) (also in Cartesian coordinates in the plane) by

$$\mathcal{F}g(X, Y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) e^{-i(xX+yY)} dx dy. \quad (5.25)$$

Definition 5.19. Similarly, the inverse Fourier transform of g at (x, y) is given by

$$\mathcal{F}^{-1}g(x, y) = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(X, Y) e^{i(xX+yY)} dx dy. \quad (5.26)$$

As in the one-variable setting, we have the inverse relationship

$$\mathcal{F}^{-1}(\mathcal{F}g)(x, y) = \mathcal{F}(\mathcal{F}^{-1}g)(x, y) = g(x, y).$$

Of course, some assumptions and restrictions on g are necessary for these integrals to make sense, but we will gloss over these concerns here. (See Appendix A for details.)

A geometric viewpoint. Consider a vector $\langle X, Y \rangle = \langle r \cos(\theta), r \sin(\theta) \rangle$ in \mathbf{R}^2 . For arbitrary real numbers x and y , let

$$\begin{aligned} t &= x \cos(\theta) + y \sin(\theta), \\ s &= -x \sin(\theta) + y \cos(\theta), \text{ and} \\ \langle x_1, y_1 \rangle &= \langle -s \sin(\theta), s \cos(\theta) \rangle. \end{aligned}$$

Then

$$\langle x, y \rangle = \langle t \cos(\theta), t \sin(\theta) \rangle + \langle x_1, y_1 \rangle,$$

and, because $\langle x_1, y_1 \rangle \cdot \langle \cos(\theta), \sin(\theta) \rangle = 0$, it follows that

$$xX + yY = rt \cos^2(\theta) + rt \sin^2(\theta) + r(x_1 \cos(\theta) + y_1 \sin(\theta)) = rt.$$

Thus, the expression $e^{-i(xX+yY)}$ in the definition of the 2-dimensional Fourier transform is the same as the expression e^{-irt} that appears in the one-variable Fourier transform. This function is periodic with period $2\pi/r$ but, in this setting, adding a multiple of $2\pi/r$ to the value of t amounts to moving the point (x, y) by a distance of $2\pi/r$ in the direction of the vector $\langle \cos(\theta), \sin(\theta) \rangle$. In other words, the function $e^{-i(xX+yY)}$ oscillates with period $2\pi/\sqrt{X^2+Y^2}$ ($= 2\pi/r$) in the direction of the vector $\langle X, Y \rangle$.

In higher dimensions, suppose the function g is defined in n -dimensional space and represent points in n -dimensional space as vectors, such as \mathbf{u} or \mathbf{v} . The n -dimensional Fourier transform is defined by

$$\mathcal{F}g(\mathbf{u}) = \int_{\mathbf{R}^n} g(\mathbf{v}) e^{-i(\mathbf{u} \cdot \mathbf{v})} d\mathbf{v}. \quad (5.27)$$

We will not have occasion to use the generalized Fourier transform for $n > 2$.

5.6 Exercises

5.1. Use the “shifting rules,” (5.6) and (5.7), to examine the following examples. (In each case, the Fourier transform of the given function has been computed previously.)

(a) Let $f(x) = e^{-|x|}$. Then the Fourier transform of f is

$$\mathcal{F}f(\omega) = \frac{2}{1 + \omega^2}.$$

- (i) Compute the Fourier transform of $g(x) = f(x - \alpha)$, where α is a constant;
- (ii) Compute the Fourier transform of $h(x) = e^{i\omega_0 x} f(x)$, where ω_0 is a constant;
- (iii) Plot the Fourier transform of f , the Fourier transform of g with $\alpha = 1$, and the Fourier transform of h with $\omega_0 = \pi/2$.

(b) For $\Pi_{1/2}(x) = \begin{cases} 1 & \text{if } -\frac{1}{2} \leq x \leq \frac{1}{2}, \\ 0 & \text{if } |x| > \frac{1}{2}, \end{cases}$ the Fourier transform is

$$\mathcal{F}\Pi_{1/2}(\omega) = \frac{\sin(\omega/2)}{(\omega/2)}.$$

- (i) Compute the Fourier transform of $g(x) = \Pi_{1/2}(x - \alpha)$, where α is a constant;
- (ii) Compute the Fourier transform of $h(x) = e^{i\omega_0 x} \Pi_{1/2}(x)$, where ω_0 is a constant;
- (iii) Plot the Fourier transform of $\Pi_{1/2}$, the Fourier transform of g with $\alpha = 1$, and the Fourier transform of h with $\omega_0 = \pi/2$.

(c) For $f(x) := e^{-x^2/2}$, the Fourier transform of f is

$$\mathcal{F}f(\omega) = \sqrt{2\pi} e^{-\omega^2/2}.$$

- (i) Compute the Fourier transform of $g(x) = f(x - \alpha)$, where α is a constant;
- (ii) Compute the Fourier transform of $h(x) = e^{i\omega_0 x} f(x)$, where ω_0 is a constant;
- (iii) Plot the Fourier transform of f , the Fourier transform of g with $\alpha = 1$, and the Fourier transform of h with $\omega_0 = \pi/2$.

5.2. Recall Euler’s Formula: $e^{it} = \cos(t) + i \sin(t)$ for every real number t .

- (a) Prove that $\cos(ax) = (e^{iax} + e^{-iax})/2$ for all real numbers a and x .
- (b) Prove that the function $G(x) = e^{iax}$ has period $2\pi/a$.
- (c) Explain, based on parts (a) and (b), why the Fourier transform of $f(x) = \cos(ax)$ consists of two impulses, located at the frequencies $\omega_1 = \frac{a}{2\pi}$ and $\omega_2 = -\frac{a}{2\pi}$.

5.3. With the assumption that $\lim_{x \rightarrow \pm\infty} f(x) e^{-i\omega x} = 0$ for all ω , show that

$$\mathcal{F}(f')(\omega) = i\omega \mathcal{F}f(\omega) \text{ for all } \omega.$$

5.4. Again using the fact that $F(\omega) = \frac{2}{1 + \omega^2}$ is the Fourier transform of $f(x) = e^{-|x|}$, compute the inverse Fourier transform of

$$H(\omega) = \frac{1}{1 + B^2(\omega - \omega_0)^2},$$

where B and ω_0 are (real) constants. (Such a function H is called a *Lorentzian*.)

5.5. A common type of signal is a *decaying wave*. Compute the Fourier transforms of these two decaying waves:

(a) (two-way decaying wave) $f(x) = e^{-\lambda|x|} \cos(\omega_0 x)$;

(b) (one-way decaying wave) $g(x) = \begin{cases} e^{-\lambda x} \cos(\omega_0 x) & \text{if } x \geq 0, \\ 0 & \text{if } x < 0. \end{cases}$

5.6. Show that the inverse Fourier transform of an even function is a real-valued function while the inverse Fourier transform of an odd function is purely imaginary (has real part equal to 0).

5.7. Let f be absolutely integrable on the real line and piecewise continuous with a point of discontinuity at α . In particular, the one-sided limits $\lim_{x \rightarrow \alpha^-} f(x)$ and $\lim_{x \rightarrow \alpha^+} f(x)$ both exist. Prove that

$$\mathcal{F}^{-1}(\mathcal{F}f)(\alpha) = (1/2) \left(\lim_{x \rightarrow \alpha^-} f(x) + \lim_{x \rightarrow \alpha^+} f(x) \right).$$

Two Big Theorems

The ideas discussed in this chapter involve interactions between three transforms—the Radon transform, the Fourier transform, and the back-projection transform. Each of these transforms is defined in terms of improper integrals on infinite intervals. This raises the somewhat technical matter of determining which functions may appropriately be considered, an issue that is addressed in Appendix A. For the time being, we assume that any function being considered here meets the requirements. For those functions that arise in the practical world of medical imaging this is certainly the case.

6.1 The central slice theorem

The interaction between the Fourier transform and the Radon transform is expressed in an equation known as the *central slice theorem* (also called the *central projection theorem*).

In this presentation, the symbols \mathcal{F} and \mathcal{F}_2 are used to denote the 1- and 2-dimensional Fourier transforms, respectively. The Radon transform is denoted by \mathcal{R} . The function f , representing, say, an X-ray attenuation coefficient, is a function of 2-dimensional Cartesian coordinates.

Theorem 6.1. Central slice theorem. *For any suitable function f defined in the plane and all real numbers S and θ ,*

$$\mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) = \mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.1)$$

Proof. Given f defined in the plane and real numbers S and θ , the definition of the 2-dimensional Fourier transform gives

$$\mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-iS(x \cos(\theta) + y \sin(\theta))} dx dy \quad (6.2)$$

Now, instead of integrating separately over $-\infty < x < \infty$ and $-\infty < y < \infty$, we can reorganize the points in the xy -plane according to the value of $x \cos(\theta) + y \sin(\theta)$. Specifically, for each real number t , gather together all of the points (x, y) in the plane for which $x \cos(\theta) + y \sin(\theta) = t$. *This is exactly the line $\ell_{t, \theta}$!* From our earlier analysis, we know that, for each point (x, y) on $\ell_{t, \theta}$, the real number $s = -x \sin(\theta) + y \cos(\theta)$ satisfies $x = t \cos(\theta) - s \sin(\theta)$ and $y = t \sin(\theta) + s \cos(\theta)$. Moreover, the matrix $\begin{bmatrix} \partial x / \partial t & \partial x / \partial s \\ \partial y / \partial t & \partial y / \partial s \end{bmatrix}$ has determinant 1, so that $ds dt = dx dy$ when the change of variables is put in place. With these changes, the right-hand side of (6.2) becomes

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) e^{-iSt} ds dt. \quad (6.3)$$

The factor e^{-iSt} in the integrand of (6.3) does not depend on s , so it may be factored out of the inner integral. Thus, (6.3) becomes

$$= \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \right) e^{-iSt} dt. \quad (6.4)$$

The inner integral in (6.4) is exactly the definition of $\mathcal{R}f(t, \theta)$, the Radon transform of the function f evaluated at the point (t, θ) . That is, (6.4) is the same as

$$\int_{-\infty}^{\infty} (\mathcal{R}f(t, \theta)) e^{-iSt} dt. \quad (6.5)$$

Finally, the integral (6.5) is the definition of the Fourier transform of $\mathcal{R}f$ evaluated at (S, θ) . That is, (6.5) is equal to

$$\mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.6)$$

We have established, as we set out to do, the equality

$$\mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) = \mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.7)$$

□

6.2 Filtered back projection

The back projection served as a first attempt at inverting the Radon transform and recovering the X-ray attenuation-coefficient function. The result was not the original function

but a smoothed-out version of it. The next theorem, called the *filtered back-projection formula*, shows how to correct for the smoothing effect and recover the original function.

Theorem 6.2. The filtered back-projection formula. *For a suitable function f defined in the plane and real numbers x and y ,*

$$f(x, y) = \frac{1}{2} \mathcal{B} \left\{ \mathcal{F}^{-1} [|S| \mathcal{F} (\mathcal{R}f) (S, \theta)] \right\} (x, y). \quad (6.8)$$

Proof. The (2-dimensional) Fourier transform and its inverse transform are just that—inverses. Hence, for any suitable function f and any point (x, y) in the plane, we get

$$f(x, y) = \mathcal{F}_2^{-1} \mathcal{F}_2 f(x, y). \quad (6.9)$$

Applying the definition of the inverse 2-dimensional Fourier transform, the right-hand side of (6.9) becomes

$$\frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathcal{F}_2 f(X, Y) e^{+i(xX+yY)} dX dY. \quad (6.10)$$

Now change variables from Cartesian coordinates (X, Y) to polar coordinates (S, θ) , where $X = S \cos(\theta)$ and $Y = S \sin(\theta)$. Rather than use the usual intervals $0 \leq S < \infty$ and $0 \leq \theta \leq 2\pi$, however, allow S to be any real number and restrict θ to $0 \leq \theta \leq \pi$. With this variation, we get $dX dY = |S| dS d\theta$ (rather than $S dS d\theta$). With these changes, (6.10) becomes

$$\frac{1}{4\pi^2} \int_0^{\pi} \int_{-\infty}^{\infty} \mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) e^{+iS(x \cos(\theta) + y \sin(\theta))} |S| dS d\theta. \quad (6.11)$$

The factor $\mathcal{F}_2 f(S \cos(\theta), S \sin(\theta))$ in the integrand of (6.11) is, according to the central slice theorem (Theorem 6.1), the same as $\mathcal{F} (\mathcal{R}f) (S, \theta)$. Thus, (6.11) is the same as

$$\frac{1}{4\pi^2} \int_0^{\pi} \int_{-\infty}^{\infty} \mathcal{F} (\mathcal{R}f) (S, \theta) e^{+iS(x \cos(\theta) + y \sin(\theta))} |S| dS d\theta. \quad (6.12)$$

The inner integral (with respect to S) in (6.12) is, by definition, 2π times the inverse Fourier transform of the function $|S| \mathcal{F} (\mathcal{R}f) (S, \theta)$ evaluated at the point $(x \cos(\theta) + y \sin(\theta), \theta)$. That is, (6.12) is the same as

$$\frac{1}{2\pi} \int_0^{\pi} \mathcal{F}^{-1} [|S| \mathcal{F} (\mathcal{R}f) (S, \theta)] (x \cos(\theta) + y \sin(\theta), \theta) d\theta. \quad (6.13)$$

Finally, the integral in (6.13) is the one used in the back projection of the (admittedly somewhat elaborate) function $\mathcal{F}^{-1} [|S| \mathcal{F} (\mathcal{R}f) (S, \theta)]$. Hence, (6.13) is equal to

$$\frac{1}{2} \mathcal{B} \left\{ \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)] \right\} (x, y). \quad (6.14)$$

We have successfully established the desired formula

$$f(x, y) = \frac{1}{2} \mathcal{B} \left\{ \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)] \right\} (x, y). \quad (6.15)$$

□

Without the factor of $|S|$ in the formula, the Fourier transform and its inverse would cancel out and the result would be simply the back projection of the Radon transform of f , which we know does not lead to recovery of f . Thus, the essential element in the formula is to multiply the Fourier transform of $\mathcal{R}f(S, \theta)$ by the absolute-value function $|S|$ *before* the inverse Fourier transform is applied. In the language of signal processing, we say that the Fourier transform of $\mathcal{R}f$ is *filtered* by multiplication by $|S|$. That is why the formula is called the filtered back-projection formula. We will discuss filters in greater detail in the next chapter.

The filtered back-projection formula is the fundamental basis for image reconstruction. However, it assumes that the values of $\mathcal{R}f(S, \theta)$ are known for all possible lines $\ell_{S, \theta}$. In practice, of course, this is not the case. Only a finite number of X-ray samples are taken and we must approximate an image from the resulting data. Thus, we will turn our attention to the practical implementation of the filtered back-projection formula. But first, let us look at a different formula for recovering f that was presented by Radon in 1917.

6.3 The Hilbert transform

A property of the Fourier transform that was not considered above concerns the interaction of the Fourier transform with the derivative of a function. To wit,

$$\mathcal{F} \left(\frac{df}{dx} \right) (\omega) = i\omega \mathcal{F}(f)(\omega). \quad (6.16)$$

Applied to the Radon transform, (6.16) yields

$$\mathcal{F} \left(\frac{\partial(\mathcal{R}f)(t, \theta)}{\partial t} \right) (S, \theta) = iS \mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.17)$$

Now, $|S| = S \cdot \text{sgn}(S)$, where

$$\operatorname{sgn}(S) = \begin{cases} 1 & \text{if } S > 0, \\ 0 & \text{if } S = 0, \\ -1 & \text{if } S < 0. \end{cases}$$

Thus, from (6.17),

$$i \cdot \operatorname{sgn}(S) \cdot \mathcal{F} \left(\frac{\partial(\mathcal{R}f)(t, \theta)}{\partial t} \right) (S, \theta) = -|S| \mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.18)$$

It now follows from Theorem 6.2, the filtered back-projection formula, that

$$f(x, y) = \frac{-1}{2} \mathcal{B} \left\{ \mathcal{F}^{-1} \left[i \cdot \operatorname{sgn}(S) \cdot \mathcal{F} \left(\frac{\partial(\mathcal{R}f)(t, \theta)}{\partial t} \right) (S, \theta) \right] \right\} (x, y). \quad (6.19)$$

That is quite a pile of symbols, so, historically, it has been simplified by defining a new transform, called the *Hilbert transform*, named for David Hilbert (1862–1943).

Definition 6.3. For a suitable function g defined on the real line, the *Hilbert transform* of g , denoted by $\mathcal{H}g$, is the function whose Fourier transform is equal to $i \cdot \operatorname{sgn} \cdot \mathcal{F}g$. That is, for each real number t , we define

$$\mathcal{H}g(t) = \mathcal{F}^{-1} [i \cdot \operatorname{sgn}(\omega) \cdot \mathcal{F}g(\omega)](t). \quad (6.20)$$

With this definition, (6.19) simplifies to the formula

$$f(x, y) = \frac{-1}{2} \mathcal{B} \left[\mathcal{H} \left(\frac{\partial(\mathcal{R}f)(t, \theta)}{\partial t} \right) (S, \theta) \right] (x, y). \quad (6.21)$$

This is Radon's original inversion formula, though expressed here in contemporary notation.

6.4 Exercises

6.1. Provide the logical explanation for each step in the proof of Theorem 6.1, the central slice theorem. (Note: The symbols \mathcal{F} and \mathcal{F}_2 are used to denote the 1- and 2-dimensional Fourier transforms, respectively. The Radon transform is denoted by \mathcal{R} . In this problem, f is a function of 2-dimensional Cartesian coordinates, and $(S \cos(\theta), S \sin(\theta))$ is a typical point in 2-dimensional Cartesian space.)

$$\mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-iS(x \cos(\theta) + y \sin(\theta))} dx dy \quad (6.22)$$

$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) e^{-iSt} ds dt \quad (6.23)$$

$$= \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds \right) e^{-iSt} dt \quad (6.24)$$

$$= \int_{-\infty}^{\infty} (\mathcal{R}f(t, \theta)) e^{-iSt} dt \quad (6.25)$$

$$= \mathcal{F}(\mathcal{R}f)(S, \theta). \quad (6.26)$$

6.2. Provide the logical explanation for each step in the derivation of Theorem 6.2, the filtered back-projection formula. (Note: The symbols \mathcal{F} and \mathcal{F}_2 are used to denote the 1- and 2-dimensional Fourier transforms, respectively. The Radon transform is denoted by \mathcal{R} and the back projection by \mathcal{B} . In this problem, f is a function of 2-dimensional Cartesian coordinates and (x, y) is a typical point in the Cartesian plane.)

$$f(x, y) = \mathcal{F}_2^{-1} \mathcal{F}_2 f(x, y) \quad (6.27)$$

$$= \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathcal{F}_2 f(X, Y) e^{+i(xX+yY)} dX dY \quad (6.28)$$

$$= \frac{1}{4\pi^2} \int_0^\pi \int_{-\infty}^{\infty} \mathcal{F}_2 f(S \cos(\theta), S \sin(\theta)) e^{+iS(x \cos(\theta) + y \sin(\theta))} |S| dS d\theta \quad (6.29)$$

$$= \frac{1}{4\pi^2} \int_0^\pi \int_{-\infty}^{\infty} \mathcal{F}(\mathcal{R}f)(S, \theta) e^{+iS(x \cos(\theta) + y \sin(\theta))} |S| dS d\theta \quad (6.30)$$

$$= \frac{1}{2\pi} \int_0^\pi \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)](x \cos(\theta) + y \sin(\theta), \theta) d\theta \quad (6.31)$$

$$= \frac{1}{2} \mathcal{B} \{ \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)] \}(x, y). \quad (6.32)$$

Filters and Convolution

7.1 Introduction

Of constant concern in the analysis of signals is the presence of *noise*, a term which here means more or less any effect that corrupts a signal. This corruption may arise from background radiation, stray signals that interfere with the main signal, errors in the measurement of the actual signal, or what have you. In order to remove the effects of noise and form a clearer picture of the actual signal, a *filter* is applied.

For a first example of a filter, consider that the noise present in a signal is often *random*. That means that the *average* amount of noise over time should be 0. Consider also that noise often has a high frequency, so the graph of the noise signal is fuzzy and jaggedy. That means that the amount of noise should average out to 0 over a fairly short time interval. So, let $T > 0$ be a positive real number and let f represent a noisy signal. For each fixed value of x , the average value of f over the interval $x - T \leq t \leq x + T$ is given by

$$f_{\text{ave}}(x) = \frac{1}{2T} \int_{t=x-T}^{x+T} f(t) dt. \quad (7.1)$$

The function f_{ave} that has just been defined represents a *filtered* version of the original signal f . For an appropriate value of T , the noise should average out to 0 over the interval, so f_{ave} would be close to the noise-free signal that we are trying to recover. If the value of T is too large, then some interesting features of the true signal may get smoothed out too much. If the choice of T is too small, then the time interval may be too short for the randomized noise to average out to 0.

A deeper analysis of (7.1) suggests that we consider the function $\phi = \square_T / (2T)$, where

$$\square_T(t) = \begin{cases} 1 & \text{if } -T \leq t \leq T, \\ 0 & \text{if } |t| > T. \end{cases} \quad (7.2)$$

Notice that

$$\int_{-\infty}^{\infty} \phi(t) dt = \frac{1}{2T} \cdot \int_{-T}^T \square_T(t) dt = 1.$$

Also, for each fixed value of x , we get

$$\phi(x-t) = \begin{cases} \frac{1}{2T} & \text{if } x-T \leq t \leq x+T, \\ 0 & \text{otherwise.} \end{cases} \quad (7.3)$$

Hence, for any given function f and any fixed value of x , we get

$$f(t)\phi(x-t) = \begin{cases} \frac{1}{2T}f(t) & \text{if } x-T \leq t \leq x+T, \\ 0 & \text{otherwise,} \end{cases} \quad (7.4)$$

from which it follows that the integral in (7.1) is the same as the integral

$$f_{\text{ave}}(x) = \int_{t=-\infty}^{\infty} f(t)\phi(x-t) dt. \quad (7.5)$$

Computationally, the function f_{ave} represents a moving average of the value of f over intervals of width $2T$. This technique is used for the analysis of all sorts of signals—radio, electrical, microwave, audio—and also for things we might not think of as being signals, like long-term behavior of stock market prices.

Graphically, the graph of $\phi(x-t)$ as a function of t is obtained by flipping the graph of ϕ over from right to left and then sliding this flipped graph along the t -axis until it is centered at x instead of at 0. This reflected-and-translated version of ϕ is then superimposed on the graph of f , and the area under the graph of the resulting product is computed. To generate the graph of f_{ave} , we reflect the graph of ϕ and then slide the reflected graph across the graph of f , stopping at each x value to compute the area underneath the product where the graphs overlap.

Example 7.1. Consider a simple “square wave:”

$$f(t) = \square_a(t) = \begin{cases} 1 & \text{if } |t| \leq a, \\ 0 & \text{if } |t| > a. \end{cases}$$

Take $\phi = \square_T/(2T)$ as above. Then the product $f(t) \cdot \phi(x-t)$ is equal to 0 whenever $|x| > T+a$, since then it is impossible to have *both* $|t| \leq a$ and $|x-t| \leq T$. For x such that $-T-a \leq x \leq 0$, we get $f(t) \cdot \phi(x-t) = 1/(2T)$ provided that $-a \leq t \leq x+T$. Similarly, for x such that $0 \leq x \leq T+a$, we get $f(t) \cdot \phi(x-t) = 1/(2T)$ provided that $x-T \leq t \leq a$. It follows now that, for x such that $-T-a \leq x \leq 0$, we get

$$\int_{t=-\infty}^{\infty} f(t) \cdot \phi(x-t) dt = \int_{t=-a}^{x+T} \left(\frac{1}{2T} \right) dt = \frac{1}{2T} \cdot (T+a+x).$$

Similarly, for x such that $0 \leq x \leq T+a$, we get

$$\int_{t=-\infty}^{\infty} f(t) \cdot \phi(x-t) dt = \int_{t=x-T}^a \left(\frac{1}{2T} \right) dt = \frac{1}{2T} \cdot (T+a-x).$$

Thus, the filtered function f_{ave} , as in (7.5), is given by

$$f_{\text{ave}}(x) = \begin{cases} \frac{1}{2T} \cdot (T+a-|x|) & \text{if } |x| \leq T+a, \\ 0 & \text{if } |x| > T+a. \end{cases} \quad (7.6)$$

This shows that, where the graph of f is a “box” (square wave) on the interval $[-a, a]$, the graph of f_{ave} is a “tent” on the interval $[-T-a, T+a]$. So the “tent” represents a smoothed-out “box”—instead of the signal starting and ending abruptly as with the box, the smoothed-out signal, the tent, gradually fades in and then gradually fades out. Perhaps we can visualize filling a rectangular box with dry sand. When the box is turned upside down and lifted away, the pile of sand will lose its box shape and collapse into a cone shape. The box becomes a tent.

7.2 Convolution

When some other function g is used in place of $\square_T/(2T)$ in the integral in (7.5), then the resulting function is not a simple moving average of the value of f over successive intervals. But we do get a modified version of f that has been “filtered” in a way that is determined by the function g . We make the following formal definition.

Definition 7.2. Given two functions f and g (defined and integrable on the real line), the *convolution* of f and g is denoted by $f * g$ and defined by

$$(f * g)(x) := \int_{t=-\infty}^{\infty} f(t)g(x-t) dt \quad \text{for } x \in \mathbf{R}. \quad (7.7)$$

For instance, the function f_{ave} in (7.5) is the same as the convolution $f * \phi$, where $\phi = \square_T/(2T)$. Graphically, the graph of $f * g$ can be obtained by reflecting the graph of g across the y -axis, then sliding the reflected graph across the graph of f , stopping at each x to compute the integral of the product where the two graphs overlap.

Example 7.3. The computation in (7.6) gives the convolution $\square_a * (\square_T/(2T))$.

Example 7.4. For the “tent” function $\wedge(t) = \begin{cases} 1 - |t| & \text{if } -1 \leq t \leq 1, \\ 0 & \text{if } |t| > 1, \end{cases}$ the convolution $\wedge * \wedge$ is piecewise cubic on the interval $-2 \leq x \leq 2$ and vanishes outside that interval.

In general, it is not so easy to compute the convolution of two functions “by hand.” The most manageable situation occurs if one of the functions is a “box” function $k \cdot \square_T$. Another helpful observation is that, if f vanishes outside the interval $[a, b]$ and g vanishes outside the interval $[c, d]$, then the convolution $f * g$ vanishes outside the interval $[a + c, b + d]$. The proof of this is left as an exercise.

7.2.1 Some properties of convolution

Commutativity. For suitable functions f and g , we get

$$f * g = g * f.$$

Proof. For each real number x , by definition

$$(f * g)(x) = \int_{t=-\infty}^{\infty} f(t) g(x-t) dt.$$

Make a change of variables with $u = x - t$. Then $du = -dt$ and $t = (x - u)$. Also, when $t = -\infty$, then $u = \infty$ and, when $t = \infty$, then $u = -\infty$. (Remember that x is fixed throughout this process). Thus, the previous integral becomes

$$\int_{u=\infty}^{-\infty} f(x-u) g(u) (-du)$$

which is the same as the integral

$$\int_{u=-\infty}^{\infty} g(u) f(x-u) du.$$

This last integral is exactly the definition of $(g * f)(x)$. Thus, $f * g = g * f$ as claimed. \square

Linearity. For suitable functions f , g_1 , and g_2 , and for scalars α and β , we get

$$f * (\alpha g_1 + \beta g_2) = \alpha(f * g_1) + \beta(f * g_2).$$

This property follows immediately from the fact that integration is linear. Combining this with the commutativity result, we also get that

$$(\alpha g_1 + \beta g_2) * f = \alpha(g_1 * f) + \beta(g_2 * f).$$

Shifting. Given a function f and a real number a , let f_a denote the shifted (translated) function

$$f_a(x) = f(x - a).$$

Then, for suitable g , we get

$$\begin{aligned} (g * f_a)(x) &= \int_{t=-\infty}^{\infty} g(t) f_a(x-t) dt \\ &= \int_{t=-\infty}^{\infty} g(t) f(x-t-a) dt \\ &= \int_{t=-\infty}^{\infty} g(t) f((x-a)-t) dt \\ &= (g * f)(x-a) \\ &= (g * f)_a(x). \end{aligned}$$

Similarly,

$$\begin{aligned} (g_a * f)(x) &= \int_{t=-\infty}^{\infty} g_a(t) f(x-t) dt \\ &= \int_{t=-\infty}^{\infty} g(t-a) f(x-t) dt \\ &= \int_{t=-\infty}^{\infty} g(t-a) f((x-a)-(t-a)) dt \\ &= \int_{s=-\infty}^{\infty} g(s) f((x-a)-s) ds \text{ where } s = t-a \\ &= (g * f)(x-a) \\ &= (g * f)_a(x). \end{aligned}$$

Convolution with δ . The convolution of an arbitrary function with the Dirac delta function yields an interesting result—it isolates the value of the function at a specific point. Specifically, for each real number x , compute

$$(f * \delta)(x) = \int_{t=-\infty}^{\infty} f(t) \delta(x-t) dt = f(x),$$

where we have used the facts that $\delta(x-t) = 0$ unless $t = x$ and that $\int_{-\infty}^{\infty} \delta(s) ds = 1$. In other words, convolution with δ acts like the identity map:

$$(f * \delta)(x) = f(x) \text{ for all } x; \quad f * \delta = f. \quad (7.8)$$

7.3 Filter resolution

The convolution of a function f with the δ function reproduces f exactly; so this “filter” has perfect resolution. More generally, let ϕ be a nonnegative function with a single maximum value M attained at $x = 0$. Suppose also that ϕ is increasing for $x < 0$ and decreasing for $x > 0$. (For example, ϕ could be a Gaussian or a “tent.”) Let the numbers x_1 and x_2 satisfy $x_1 < 0 < x_2$ and $\phi(x_1) = \phi(x_2) = M/2$, half the maximum value of ϕ . The distance $(x_2 - x_1)$ is called the *full width half maximum* of the function ϕ , denoted $\text{FWHM}(\phi)$. For the filter of “convolution with ϕ ,” the resolution of the filter is *defined* to be equal to $\text{FWHM}(\phi)$.

The idea is that a function ϕ having a smaller FWHM is pointier or spikier than a function with a larger FWHM and, hence, looks more like the δ function. So the resolution is better if the filter function ϕ has a smaller FWHM.

Here is another way to see how the resolution of a filter is related to the FWHM of the filter function. Suppose that a function f consists of two unit impulses, two instantaneous “blips,” separated by a distance of d . It can be shown, for instance by using the graphical approach of sliding the reflected graph of ϕ across the graph of f , that, if $d > \text{FWHM}(\phi)$, then the graph of $(f * \phi)$ has two peaks, like f . But if $d \leq \text{FWHM}(\phi)$, then the graph of $(f * \phi)$ has only one peak, so the detail in the graph of f is lost, or blurred, by convolution with ϕ . In other words, the smallest distance between distinct features of f that can still be seen in the filtered signal $(f * \phi)$ is $\text{FWHM}(\phi)$. The choice of the filter function has a direct effect on the resolution of the filtered signal.

Example 7.5. FWHM of a Gaussian. For ω real, let $F(\omega) = e^{-B\omega^2}$, where B is a positive constant. The maximum value of F is $F(0) = 1$. Thus, half maximum is achieved when $e^{-B\omega^2} = 1/2$, or when $\omega = \pm\sqrt{\ln(2)/B}$. Therefore,

$$\text{FWHM} = 2\sqrt{\ln(2)/B}$$

for this function.

7.4 Convolution and the Fourier transform

For suitable functions f and g , the product of their respective Fourier transforms, evaluated at ω , is

$$\mathcal{F}f(\omega) \cdot \mathcal{F}g(\omega) = \int_{x=-\infty}^{\infty} f(x) e^{-i\omega x} dx \cdot \int_{y=-\infty}^{\infty} g(y) e^{-i\omega y} dy. \quad (7.9)$$

Keep in mind that x and y are just “dummy” variables here. Now, introduce a new

variable s such that, for each fixed value of x , we have $y = s - x$. Thus $dy = ds$, and hence the right-hand side of (7.9) becomes

$$\int_{x=-\infty}^{\infty} f(x) e^{-i\omega x} dx \cdot \int_{s=-\infty}^{\infty} g(s-x) e^{-i\omega(s-x)} ds. \quad (7.10)$$

The first integral in (7.10) is independent of s , so we may move it inside the other integral. Since $e^{-i\omega x} \cdot e^{-i\omega(s-x)} = e^{-i\omega s}$, this yields

$$\int_{s=-\infty}^{\infty} \left(\int_{x=-\infty}^{\infty} f(x) g(s-x) dx \right) e^{-i\omega s} ds. \quad (7.11)$$

Notice now that the inner integral in (7.11) is exactly $(f * g)(s)$, while the outer integral is the Fourier transform of the inner integral, evaluated at ω . That is, (7.11) is the same as

$$\int_{s=-\infty}^{\infty} (f * g)(s) e^{-i\omega s} ds = \mathcal{F}(f * g)(\omega). \quad (7.12)$$

Thus, we have established the following property.

Theorem 7.6. For suitable functions f and g ,

$$\mathcal{F} f \cdot \mathcal{F} g = \mathcal{F}(f * g). \quad (7.13)$$

So, we see that the Fourier transform of a convolution is just the product of the individual transforms. This relationship will play a significant role in what is to come. We might wonder as well what happens if we apply the Fourier transform to a product. The result is almost as clean as in the previous theorem, except for a factor of $1/2\pi$ that creeps in.

Theorem 7.7. For suitable functions f and g ,

$$\mathcal{F}(f \cdot g) = \frac{1}{2\pi} (\mathcal{F} f) * (\mathcal{F} g). \quad (7.14)$$

Proof. Given f and g , for simplicity denote $\mathcal{F} f$ and $\mathcal{F} g$ by F and G , respectively. The Fourier transform of the product $f \cdot g$, evaluated at an arbitrary real number ω , is

$$\begin{aligned} & \mathcal{F}(f \cdot g)(\omega) \\ &= \int_{x=-\infty}^{\infty} f(x) g(x) e^{-i\omega x} dx \\ &= \int_{x=-\infty}^{\infty} \left(\frac{1}{2\pi} \int_{v=-\infty}^{\infty} F(v) e^{+ivx} dv \right) \left(\frac{1}{2\pi} \int_{\zeta=-\infty}^{\infty} G(\zeta) e^{+i\zeta x} d\zeta \right) e^{-i\omega x} dx \\ &= \left(\frac{1}{2\pi} \right)^2 \int_{v=-\infty}^{\infty} F(v) \left(\int_{\zeta=-\infty}^{\infty} G(\zeta) \left(\int_{x=-\infty}^{\infty} e^{+i(v+\zeta-\omega)x} dx \right) d\zeta \right) dv. \end{aligned}$$

Observe that the inner integral, $\int_{x=-\infty}^{\infty} e^{+i(v+\zeta-\omega)x} dx$, represents (2π) times the inverse Fourier transform of the constant function 1 evaluated at $(v + \zeta - \omega)$. But $\mathcal{F}^{-1}(1) = \delta$, so

$$\int_{x=-\infty}^{\infty} e^{+i(v+\zeta-\omega)x} dx = (2\pi) \cdot \delta(v + \zeta - \omega).$$

Moreover, $\delta(v + \zeta - \omega) = 0$ except when $\zeta = \omega - v$. Hence, from (5.12),

$$\begin{aligned} \int_{\zeta=-\infty}^{\infty} G(\zeta) \left(\int_{x=-\infty}^{\infty} e^{+i(v+\zeta-\omega)x} dx \right) d\zeta &= 2\pi \int_{\zeta=-\infty}^{\infty} G(\zeta) \delta(v + \zeta - \omega) d\zeta \\ &= 2\pi G(\omega - v). \end{aligned}$$

Continuing from where we left off, we now see that

$$\begin{aligned} \mathcal{F}(f \cdot g)(\omega) &= \left(\frac{1}{2\pi} \right)^2 \int_{v=-\infty}^{\infty} F(v) \left(\int_{\zeta=-\infty}^{\infty} G(\zeta) \left(\int_{x=-\infty}^{\infty} e^{+i(v+\zeta-\omega)x} dx \right) d\zeta \right) dv \\ &= \left(\frac{1}{2\pi} \right) \int_{v=-\infty}^{\infty} F(v) G(\omega - v) dv \\ &= \left(\frac{1}{2\pi} \right) F * G(\omega), \end{aligned}$$

which establishes the claim. \square

7.5 The Rayleigh–Plancherel theorem

Theorem 7.8. Rayleigh–Plancherel. *Let f be an integrable function. If either the function f or its Fourier transform $\mathcal{F}f$ is square-integrable on the real line, then so is the other and*

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\mathcal{F}f(\omega)|^2 d\omega. \quad (7.15)$$

Before looking at a proof of this statement, note that both integrands are nonnegative functions even if f or $\mathcal{F}f$ has complex number values. The function $|\mathcal{F}f(\omega)|^2$ is called the *power spectrum* of f , and in some physical applications actually does represent the total power (measured in watts, for example) of a signal at a given frequency. In the same sort of setting, the integral on the right in (7.15) represents a measure of the total amount of power present in the system.

Computationally, the value of the theorem is that one of the integrals might be comparatively easy to evaluate while the other, on its own, may be difficult.

Example 7.9. Take

$$f(x) = \begin{cases} 1 & \text{if } -1 \leq x \leq 1, \\ 0 & \text{if } |x| > 1. \end{cases}$$

Then

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \int_{-1}^1 1 dx = 2.$$

We have seen before that $\mathcal{F}f(\omega) = 2 \frac{\sin(\omega)}{\omega}$. Hence, by (7.15),

$$\int_{-\infty}^{\infty} \frac{\sin^2(\omega)}{\omega^2} d\omega = \pi.$$

Rayleigh established this result in 1889 and used it in his analysis of black body radiation. Where Rayleigh tacitly assumed that both integrals would be finite, Plancherel proved, in 1910, that the existence of one or the other integral is indeed the only hypothesis required. In other words, the relation (7.15) holds whenever either integral exists.

Now for the proof.

Proof of Theorem 7.8. Let f be an integrable function and suppose that either f or $\mathcal{F}f$ is square-integrable. We have

$$\begin{aligned} \int_{-\infty}^{\infty} |f(x)|^2 dx &= \int_{-\infty}^{\infty} f(x) \overline{f(x)} dx \\ &= \int_{-\infty}^{\infty} f(x) \overline{f(x)} e^{-i(0)x} dx \\ &= \mathcal{F}(f \cdot \bar{f})(0) \\ &= \frac{1}{2\pi} [(\mathcal{F}f) * (\mathcal{F}\bar{f})](0) \text{ by (7.14)} \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathcal{F}f(\omega) \mathcal{F}\bar{f}(0 - \omega) d\omega \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathcal{F}f(\omega) \overline{\mathcal{F}f(\omega)} d\omega \text{ by (5.10)} \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} |\mathcal{F}f(\omega)|^2 d\omega. \end{aligned}$$

Hence, provided one of these integrals is finite, then so is the other and the desired relation holds. \square

7.6 Convolution in 2-dimensional space

For two functions whose inputs are polar coordinates in the plane, the convolution is defined in terms of the radial variable only. That is, for $f(t, \theta)$ and $g(t, \theta)$, we define

$$(f * g)(t, \theta) = \int_{s=-\infty}^{\infty} f(s, \theta) \cdot g(t - s, \theta) ds. \quad (7.16)$$

For two functions whose inputs are Cartesian coordinates in the plane, the convolution incorporates both variables. That is, for $F(x, y)$ and $G(x, y)$, we define

$$(F * G)(x, y) = \int_{t=-\infty}^{\infty} \int_{s=-\infty}^{\infty} F(s, t) \cdot G(x - s, y - t) ds dt. \quad (7.17)$$

As in one dimension, convolution is *commutative*: $f * g = g * f$ and $F * G = G * F$.

7.7 Convolution, \mathcal{B} , and \mathcal{R}

Given functions $g(t, \theta)$ (in polar coordinates) and $f(x, y)$ (in Cartesian coordinates), recall that the back projection and Radon transform are defined by

$$\mathcal{B}g(x, y) = \frac{1}{\pi} \int_{\theta=0}^{\pi} g(x \cos(\theta) + y \sin(\theta), \theta) d\theta \quad (7.18)$$

and

$$\mathcal{R}f(t, \theta) = \int_{s=-\infty}^{\infty} f(t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta)) ds. \quad (7.19)$$

Proposition 7.10. (See [30], Theorem 1.3) For suitable functions $g(t, \theta)$ and $f(x, y)$, and arbitrary real numbers X and Y , we have

$$(\mathcal{B}g * f)(X, Y) = \mathcal{B}(g * \mathcal{R}f)(X, Y). \quad (7.20)$$

Proof. From (7.17) and (7.18), we compute

$$\begin{aligned} & (\mathcal{B}g * f)(X, Y) \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathcal{B}g(X - x, Y - y) \cdot f(x, y) dx dy \\ &= \frac{1}{\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left[\int_0^{\pi} g((X - x) \cos(\theta) + (Y - y) \sin(\theta), \theta) d\theta \right] f(x, y) dx dy. \end{aligned}$$

Now substitute $x = t \cos(\theta) - s \sin(\theta)$ and $y = t \sin(\theta) + s \cos(\theta)$. Keeping in mind that $ds dt = dx dy$ and using (7.19), the preceding integral becomes

$$\begin{aligned} &= \frac{1}{\pi} \int_0^\pi \int_{-\infty}^{\infty} g(X \cos(\theta) + Y \sin(\theta) - t, \theta) \cdot \mathcal{R}f(t, \theta) dt d\theta \\ &= \frac{1}{\pi} \int_0^\pi (g * \mathcal{R}f)(X \cos(\theta) + Y \sin(\theta), \theta) d\theta \text{ by (7.16)} \\ &= \mathcal{B}(g * \mathcal{R}f)(X, Y) \text{ by (7.18)}. \end{aligned}$$

This proves the claim. □

7.8 Low-pass filters

Let us return our attention to the filtered back-projection formula,

$$f(x, y) = \frac{1}{2} \mathcal{B} \{ \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)] \} (x, y), \quad (7.21)$$

where f is some suitable function defined in the xy -plane.

Suppose there were a function $\phi(t)$ whose Fourier transform satisfied $\mathcal{F}\phi(S) = |S|$. That is, suppose $\mathcal{F}\phi$ were equal to the absolute-value function. Then we would get

$$|S| \mathcal{F}(\mathcal{R}f)(S, \theta) = [\mathcal{F}\phi \cdot \mathcal{F}(\mathcal{R}f)](S, \theta).$$

This is the product of two Fourier transforms, which we now know is equal to the Fourier transform of the convolution of the two functions. So we would have

$$|S| \mathcal{F}(\mathcal{R}f)(S, \theta) = \mathcal{F}(\phi * \mathcal{R}f)(S, \theta).$$

Hence, we would have

$$\begin{aligned} \mathcal{F}^{-1} [|S| \mathcal{F}(\mathcal{R}f)(S, \theta)] &= \mathcal{F}^{-1} [\mathcal{F}(\phi * \mathcal{R}f)(S, \theta)] \\ &= (\phi * \mathcal{R}f)(t, \theta). \end{aligned}$$

Substituting this into (7.21), we would get

$$f(x, y) = \frac{1}{2} \mathcal{B}(\phi * \mathcal{R}f)(x, y). \quad (7.22)$$

Thus, reconstruction of f would require the convolution, or filtering, of $\mathcal{R}f$, the data from the X-ray machine, with ϕ followed by an application of the back projection to that. That doesn't sound so terrible, *except that there is no such function ϕ . The absolute-value function is not the Fourier transform of any function.*

So we might as well ignore all of the preceding discussion, right? Well, not so fast. The filtered back-projection formula (7.21) gives a recipe for reconstructing f , but in practice the recipe has a problem. The function $|S|\mathcal{F}(\mathcal{R}f)(S, \theta)$ in the formula is highly sensitive to noise. The value of S represents a frequency that is present in a signal. So, if the Radon transform $\mathcal{R}f$, representing the X-ray data, has a component at a high frequency, then that component is magnified by the factor $|S|$. That means that noise present in the data gets exaggerated, an effect which corrupts the reconstructed image. Thus, in practice, we don't really want to use the factor $|S|$ anyway.

In place of $|S|$, we use a function that is close to the absolute-value function for S near 0 but that vanishes when the value of $|S|$ is "large." Such a function is called a *low-pass filter* because the lower frequencies are not affected by its presence while higher frequencies, including noise, get cut off. Also, in order to use the modification (7.22) of the filtered back-projection formula, we want the function that replaces $|S|$ to be the Fourier transform of something. That is, we want to replace $|S|$ with a function of the form $A = \mathcal{F}\phi$, where A is nonzero on some finite interval and zero outside that interval.

Definition 7.11. A function ϕ whose Fourier transform is nonzero on some finite interval and zero outside that interval is called a *band-limited function*.

So, using this terminology, we want to replace $|S|$ in the filtered back-projection formula by a low-pass filter that is the Fourier transform of a band-limited function. The price of doing this is that *the formula (7.22) is no longer exact*, but gives only an approximation for f :

$$f(x, y) \approx \frac{1}{2} \mathcal{B}(\mathcal{F}^{-1}A * \mathcal{R}f)(x, y). \quad (7.23)$$

To design a low-pass filter to replace the absolute-value function, we typically use a function of the form

$$A(\omega) = |\omega| \cdot F(\omega) \cdot \square_L(\omega), \quad (7.24)$$

for some number $L > 0$. Thus, $A(\omega)$ vanishes for $|\omega| > L$ and has the value $|\omega| \cdot F(\omega)$ when $|\omega| \leq L$. Near the origin, the value of A should be close to the absolute value, so we want F to be an even function for which $F(0) = 1$. Also, choosing A to be an even function guarantees that $\phi = \mathcal{F}^{-1}A$ is real-valued. In general, the function $F \cdot \square_L$, by which $|\cdot|$ is multiplied, is called the *window function*.

Example 7.12. Low-pass filters. Here are some of the low-pass filters most commonly used in medical imaging. We will analyze them more closely in Section 8.2.

- The *Ram–Lak filter*:

$$A_1(\omega) = |\omega| \cdot \square_L(\omega) = \begin{cases} |\omega| & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L, \end{cases} \quad (7.25)$$

where $L > 0$. This is simply a truncation of the absolute-value function to a finite interval. It is the Fourier transform of a band-limited function. This filter was used by Bracewell and Riddle [7, 8] as well as by Ramachandran and Lakshminarayanan [35].

- The *Shepp–Logan filter*:

$$\begin{aligned} A_3(\omega) &= |\omega| \cdot \left(\frac{\sin(\pi\omega/(2L))}{\pi\omega/(2L)} \right) \cdot \square_L(\omega) \\ &= \begin{cases} \frac{2L}{\pi} \cdot |\sin(\pi\omega/(2L))| & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L. \end{cases} \end{aligned} \quad (7.26)$$

This filter was introduced by Shepp and Logan [40].

- The *low-pass cosine filter*:

$$\begin{aligned} A_2(\omega) &= |\omega| \cdot \cos(\pi\omega/(2L)) \cdot \square_L(\omega) \\ &= \begin{cases} |\omega| \cos(\pi\omega/(2L)) & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L. \end{cases} \end{aligned} \quad (7.27)$$

This filter is commonly used in signal analysis.

We conclude this chapter by looking at the general form of a band-limited function. So, let f be band-limited and select a number $L > 0$ such that $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > L$. In particular, this means that there is some function G such that

$$\mathcal{F}f(\omega) = \square_L(\omega) \cdot G(\omega),$$

where, again, $\square_L(\omega) = 1$ when $|\omega| \leq L$ and $\square_L(\omega) = 0$ when $|\omega| > L$. From Example 5.15 and Theorem 5.11, we know that \square_L is the Fourier transform of the function $\phi(x) = \sin(Lx)/(\pi x)$. If we let $g = \mathcal{F}^{-1}G$, then we see that $\mathcal{F}f = \mathcal{F}\phi \cdot \mathcal{F}g$, the product of two Fourier transforms. Hence, from Theorem 7.6, it follows that $\mathcal{F}f = \mathcal{F}(\phi * g)$ and, hence that $f = \phi * g$. This is therefore the general form that a band-limited function must have: the convolution of the function ϕ , for some value of L , and some integrable function g .

7.9 Exercises

7.1. Use a computer algebra system to compute and plot the convolutions of various pairs of functions. (Note: It is perfectly reasonable to form the convolution of a function with itself!)

7.2. In this exercise, let the filter function be $\Pi_{1/2}$. (Recall that $\Pi_{1/2}(x) = 1$ for $-1/2 \leq x \leq 1/2$ and $\Pi_{1/2}(x) = 0$ when $|x| > 1/2$.)

(a) For $g(x) = \cos(x)$, compute $(\Pi_{1/2} * g)(x)$ for x real.

(b) For $h(x) = \sin(x)$, compute $(\Pi_{1/2} * h)(x)$ for x real.

(c) For $F(x) = e^{-|x|}$, compute $(\Pi_{1/2} * F)(x)$ for x real.

7.3. Suppose the function f vanishes outside the interval $[a, b]$ and the function g vanishes outside the interval $[c, d]$. Show that the convolution $f * g$ vanishes outside the interval $[a + c, b + d]$.

7.4. Apply the inverse Fourier transform \mathcal{F}^{-1} to both sides of (7.13) and (7.14) to provide companion statements about the inverse Fourier transform of a convolution and of a product.

7.5. Show that, if f and g are suitable functions of two real variables, then

$$\mathcal{R}(f * g)(t, \theta) = (\mathcal{R}f * \mathcal{R}g)(t, \theta) \quad (7.28)$$

for all values of t and θ . (*Hint:* This can be proven either as a consequence of the central slice theorem (Theorem 6.1) or directly from the definitions, using (7.16) and (7.17).)

7.6. (a) Compute the FWHM (full width half maximum) of the Lorentzian signal

$$g(\omega) = \frac{T_2}{1 + 4\pi^2 T_2^2 (\omega - \omega_0)^2},$$

where T_2 is a constant (one of the *relaxation constants* related to magnetic resonance imaging) and the signal is centered around the angular frequency ω_0 .

(b) Find a value for B so that the Gaussian signal

$$f(\omega) = T_2 e^{-B(\omega - \omega_0)^2}$$

has the same FWHM as the Lorentzian signal in part (a).

(c) For several values of the relaxation constant T_2 , plot and compare the graphs of the signals in parts (a) and (b).

Discrete Image Reconstruction

8.1 Introduction

We have seen that, when complete continuous X-ray data are available, then an attenuation-coefficient function $f(x, y)$ can be reconstructed exactly using the filtered back-projection formula, Theorem 6.2. To repeat,

$$f(x, y) = \frac{1}{2} \mathcal{B} \left\{ \mathcal{F}^{-1} [|S| \mathcal{F} (\mathcal{R}f) (S, \theta)] \right\} (x, y). \quad (8.1)$$

In the previous chapter, in Section 7.8, we discussed the practice of replacing the absolute-value function $|\cdot|$ in (8.1) with a low-pass filter A , obtained by multiplying the absolute value by a window function that vanishes outside some finite interval. Then, in place of (8.1), we use the approximation

$$f(x, y) \approx \frac{1}{2} \mathcal{B} \left(\mathcal{F}^{-1} A * \mathcal{R}f \right) (x, y). \quad (8.2)$$

The starting point in the implementation of (8.2) in the practical reconstruction of images from X-ray data is that only a finite number of values of $\mathcal{R}f(S, \theta)$ are available in a real study. This raises both questions of accuracy—*How many values are needed in order to generate a clinically useful image?*—and problems of computation—*What do the various components of the formula (8.2) mean in a discrete setting?* Also, as we shall see, one step in the algorithm requires essentially that we “fill in” some missing values. This is done using a process called interpolation that we will discuss later. Each different method of interpolation has its advantages and disadvantages. And, of course, the choice of the low-pass filter affects image quality.

8.2 Sampling

The term *sampling* refers to the situation where the values of a function that presumably is defined on the whole real line are known or are computed only at a discrete set of points. For instance, we might know the values of the function at all points of the form $k \cdot d$, where d is called the *sample spacing*. The basic problem is to determine how many sampled values are enough, or what value d should have, to form an accurate description of the function as a whole.

For example, consider a string of sampled values $\{f(k) = 1\}$, sampled at every integer (so $d = 1$). These could be the values of the constant function $f(x) \equiv 1$, or of the periodic function $\cos(2\pi x)$. If it is the latter, then we obviously would need more samples to be certain.

Heuristically, we might think of a signal as consisting of a compilation of sine or cosine waves of various frequencies and amplitudes. The narrowest “bump” in this compilation is due to the wave that has the shortest wavelength and also constitutes the smallest “feature” that is present in the signal. The reciprocal of the shortest wavelength present in the signal is the maximum frequency that is present in the Fourier transform of the signal. So this signal is a *band-limited* function—its Fourier transform is zero outside a finite interval.

Now, suppose that f is a band-limited function for which $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > L$. By definition (see (4.7)), the *Fourier series coefficients* of $\mathcal{F}f$ (where we act as if $\mathcal{F}f$ has been extended beyond the interval $[-L, L]$ to be periodic on the real line) are given by

$$C_n = \left(\frac{1}{2L} \right) \int_{-L}^L \mathcal{F}f(\omega) e^{-i\pi n\omega/L} d\omega, \quad (8.3)$$

where n is any integer. For each integer n , then,

$$\begin{aligned} (2\pi) \cdot f(\pi n/L) &= 2\pi \cdot \mathcal{F}^{-1}(\mathcal{F}f)(\pi n/L) \\ &= \int_{-\infty}^{\infty} \mathcal{F}f(\omega) e^{i\omega\pi n/L} d\omega \\ &= \int_{-L}^L \mathcal{F}f(\omega) e^{i\omega\pi n/L} d\omega \quad \text{since } f \text{ is band-limited} \\ &= (2L) \cdot C_{-n}. \end{aligned}$$

That is,

$$C_{-n} = (\pi/L) \cdot f(\pi n/L) \text{ for every integer } n.$$

Assuming that $\mathcal{F}f$ is continuous, it follows from standard results of Fourier series that

$$\mathcal{F}f(\omega) = \sum_{n=-\infty}^{\infty} C_{-n} e^{-in\pi\omega/L} = \left(\frac{\pi}{L} \right) \cdot \sum_{n=-\infty}^{\infty} f(\pi n/L) e^{-in\pi\omega/L}. \quad (8.4)$$

Finally, all of this means that, for every x , we have

$$\begin{aligned}
 f(x) &= \mathcal{F}^{-1}(\mathcal{F}f)(x) \\
 &= \left(\frac{1}{2\pi}\right) \int_{-L}^L \mathcal{F}f(\omega) e^{i\omega x} dx \\
 &= \left(\frac{1}{2\pi}\right) \left(\frac{\pi}{L}\right) \int_{-L}^L \left[\sum_{n=-\infty}^{\infty} f(\pi n/L) e^{-in\pi\omega/L} \right] e^{i\omega x} dx \text{ from (8.4)} \\
 &= \left(\frac{1}{2L}\right) \sum_{n=-\infty}^{\infty} \left[f(\pi n/L) \cdot \int_{-L}^L e^{i\omega(Lx-n\pi)/L} d\omega \right] \\
 &= \left(\frac{1}{2L}\right) \sum_{n=-\infty}^{\infty} f(\pi n/L) \cdot (2L) \cdot \left(\frac{\sin(Lx-n\pi)}{Lx-n\pi}\right) \\
 &= \sum_{n=-\infty}^{\infty} f(\pi n/L) \cdot \left(\frac{\sin(Lx-n\pi)}{Lx-n\pi}\right).
 \end{aligned}$$

Thus we see that, for a band-limited function f whose Fourier transform vanishes outside the interval $[-L, L]$, the function f can be reconstructed exactly from the values $\{f(n\pi/L) : -\infty < n < \infty\}$. In other words, the appropriate sample spacing for the function f is $d = \pi/L$. Since L represents the maximum value of $|\omega|$ present in the Fourier transform $\mathcal{F}f$, the value $2\pi/L$ represents the smallest wavelength present in the signal f . Therefore, *the optimal sample spacing is equal to half of the size of the smallest detail present in the signal*. This result is known as *Nyquist's theorem* and the value $d = \pi/L$ is called the *Nyquist distance*.

To sum up, we have established the following.

Theorem 8.1. Nyquist's theorem. *If f is a square-integrable band-limited function such that the Fourier transform $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > L$, then, for every real number x ,*

$$f(x) = \sum_{n=-\infty}^{\infty} f(\pi n/L) \cdot \frac{\sin(Lx - n\pi)}{Lx - n\pi}. \quad (8.5)$$

Nyquist's theorem is also sometimes referred to as *Shannon–Whittaker interpolation* since it asserts that *any* value of the function f can be interpolated from the values $\{f(n\pi/L)\}$.

A heuristic approach to Nyquist's theorem. As above, assume that the function f is band-limited with $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > L$. We want to extend $\mathcal{F}f$ to be periodic on the whole line, so we can take the period to be $2L$, the length of the interval $-L \leq \omega \leq L$. Now, each of the functions $\omega \mapsto e^{-in\pi\omega/L}$, where n is an integer, has period $2L$. From the definition of the Fourier transform,

$$\mathcal{F}f(\omega) = \int_{x=-\infty}^{\infty} f(x) e^{-i\omega x} dx. \quad (8.6)$$

Approximate this integral using a Riemann sum with $dx = \pi/L$. That is, the Riemann sum will use a partition of the line that includes all points of the form $n\pi/L$, where n is an integer. This results in the approximation

$$\mathcal{F}f(\omega) \approx \left(\frac{\pi}{L}\right) \cdot \sum_{n=-\infty}^{\infty} f(n\pi/L) e^{-in\pi\omega/L}.$$

This approximates $\mathcal{F}f$ in a way that is periodic on the whole line with period $2L$. Notice that this is the same as (8.4) except for having “ \approx ” instead of “ $=$.” In the proof above, we appealed to results in the theory of Fourier series that assert that this approximation is actually an equality. Without that knowledge, we instead substitute the approximation for $\mathcal{F}f$ into the rest of the proof of Nyquist’s theorem and end up with the approximation

$$f(x) \approx \sum_{n=-\infty}^{\infty} f(\pi n/L) \cdot \frac{\sin(Lx - n\pi)}{Lx - n\pi}. \quad (8.7)$$

Nyquist’s theorem, which builds on the results concerning Fourier series, asserts that this is in fact an equality.

Oversampling. The interpolation formula (8.5) is an infinite series where, in practice, we would only use a partial sum. However, the series (8.5) may converge fairly slowly because the expression $\sin(Lx - n\pi)/(Lx - n\pi)$ is on the order of $(1/n)$ for large values of n and the harmonic series $\sum 1/n$ diverges. That means that a partial sum might require a large number of terms in order to achieve a good approximation to $f(x)$.

To address this difficulty, notice that, if $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > L$, and if $R > L$, then $\mathcal{F}f(\omega) = 0$ whenever $|\omega| > R$ as well. Thus, we can use Shannon–Whittaker interpolation on the interval $[-R, R]$ instead of $[-L, L]$. This requires that the function f be sampled at the Nyquist distance π/R instead of π/L . Since $\pi/R < \pi/L$, this results in what is called *oversampling* of the function f . So there is a computational price to pay for oversampling, but the improvement in the results, when using a partial sum to approximate $f(x)$, may be worth that price.

8.3 Discrete low-pass filters

The image reconstruction formula (8.2) involves the inverse Fourier transform $\mathcal{F}^{-1}A$ for the low-pass filter A . In practice, this too will be sampled, just like the Radon transform. Nyquist’s theorem, Theorem 8.1, tells us how many sampled values are needed to get an accurate representation of $\mathcal{F}^{-1}A$. Here, we investigate how this works for two particular low-pass filters.




Fig. 8.1. The Shepp–Logan filter, with $L = 10$, and a piecewise linear rendering of its inverse Fourier transform from the sampled values.

Example 8.2. The *Shepp–Logan filter*, introduced in (7.26), is defined by

$$\begin{aligned}
 A(\omega) &= |\omega| \cdot \left(\frac{\sin(\pi\omega/(2L))}{\pi\omega/(2L)} \right) \cdot \square_L(\omega) \\
 &= \begin{cases} \frac{2L}{\pi} \cdot |\sin(\pi\omega/(2L))| & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L, \end{cases} \quad (8.8)
 \end{aligned}$$

for some choice of $L > 0$.

The inverse Fourier transform of A is a band-limited function. From the fact that A is an even function that vanishes outside the interval $[-L, L]$, we get, for each real number x ,

$$\begin{aligned}
 (\mathcal{F}^{-1}A)(x) &= \frac{1}{\pi} \int_0^L \frac{2L}{\pi} \cdot \sin(\pi\omega/(2L)) \cdot \cos(x\omega) d\omega \\
 &= \left(\frac{L}{\pi^2} \right) \cdot \left\{ \frac{\cos((x - \pi/(2L))\omega)}{x - \pi/(2L)} - \frac{\cos((x + \pi/(2L))\omega)}{x + \pi/(2L)} \right\} \Big|_0^L \\
 &= \left(\frac{L}{\pi^2} \right) \cdot \left\{ \left[\frac{\cos(Lx - \pi/2)}{x - \pi/(2L)} - \frac{\cos(Lx + \pi/2)}{x + \pi/(2L)} \right] \right. \\
 &\quad \left. - \left[\frac{1}{x - \pi/(2L)} - \frac{1}{x + \pi/(2L)} \right] \right\}. \quad (8.9)
 \end{aligned}$$

According to Nyquist’s theorem, the function $(\mathcal{F}^{-1}A)$ can be reconstructed exactly from its values taken in increments of the Nyquist distance π/L . Setting $x = \pi n/L$ in (8.9) yields

$$\begin{aligned}
(\mathcal{F}^{-1}A)(\pi n/L) &= \left(\frac{L}{\pi^2}\right) \cdot \left\{ \left[\frac{\cos(\pi n - \pi/2)}{\pi n/L - \pi/(2L)} - \frac{\cos(\pi n + \pi/2)}{\pi n/L + \pi/(2L)} \right] \right. \\
&\quad \left. - \left[\frac{1}{\pi n/L - \pi/(2L)} - \frac{1}{\pi n/L + \pi/(2L)} \right] \right\} \\
&= \left(\frac{L}{\pi^2}\right) \cdot \left\{ \frac{1}{\pi n/L + \pi/(2L)} - \frac{1}{\pi n/L - \pi/(2L)} \right\} \\
&= \frac{4L^2}{\pi^3(1 - 4n^2)}. \tag{8.10}
\end{aligned}$$

Example 8.3. The Ram–Lak filter, defined in (7.25), has the formula

$$A(\omega) = |\omega| \cdot \Gamma_L(\omega) = \begin{cases} |\omega| & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L. \end{cases} \tag{8.11}$$

Proceeding as in the previous example, we find that the inverse Fourier transform of the Ram–Lak filter satisfies

$$\begin{aligned}
(\mathcal{F}^{-1}A)(x) &= \frac{1}{\pi} \int_0^L \omega \cdot \cos(x\omega) d\omega \\
&= \left(\frac{1}{\pi}\right) \cdot \left\{ \frac{\cos(x\omega) + (x\omega) \cdot \sin(x\omega)}{x^2} \right\} \Big|_0^L \\
&= \left(\frac{1}{\pi}\right) \cdot \left\{ \frac{\cos(Lx) + (Lx) \cdot \sin(Lx) - 1}{x^2} \right\} \\
&= \left(\frac{1}{\pi}\right) \cdot \left\{ \frac{(Lx) \cdot \sin(Lx)}{x^2} - \frac{2 \cdot \sin^2(Lx/2)}{x^2} \right\}, \tag{8.12}
\end{aligned}$$

where the trigonometric identity $\cos(\theta) = 1 - 2 \cdot \sin^2(\theta/2)$ was used in the last step.

As before, set $x = \pi n/L$ to evaluate $(\mathcal{F}^{-1}A)$ at multiples of the Nyquist distance π/L . Thus,

$$\begin{aligned}
(\mathcal{F}^{-1}A)(\pi n/L) &= \left(\frac{1}{\pi}\right) \cdot \left\{ \frac{(\pi n) \cdot \sin(\pi n)}{(\pi n/L)^2} - \frac{2 \cdot \sin^2(\pi n/2)}{(\pi n/L)^2} \right\} \\
&= \frac{L^2}{2\pi} \cdot \left\{ \frac{2 \cdot \sin(\pi n)}{\pi n} - \left[\frac{\sin(\pi n/2)}{(\pi n/2)} \right]^2 \right\}. \tag{8.13}
\end{aligned}$$

For the Ram–Lak filter, we have $(\mathcal{F}^{-1}A)(0) = L^2/(2\pi)$. For nonzero even values of n , we get $(\mathcal{F}^{-1}A)(\pi n/L) = 0$, while for n odd the value is given by $(\mathcal{F}^{-1}A)(\pi n/L) = -2L^2/(\pi^3 \cdot n^2)$.

Example 8.4. The low-pass cosine filter, defined in (7.27), is left for the exercises.




Fig. 8.2. The Ram–Lak filter, with $L = 10$, and a piecewise linear rendering of its inverse Fourier transform from the sampled values.

8.4 Discrete Radon transform

In the context of a CT scan, the X-ray machine does not assess the attenuation along every line $\ell_{t,\theta}$. Instead, in the model we have been using, the Radon transform is sampled for a finite number of angles θ between 0 and π and, at each of these angles, for some finite number of values of t . Both the angles and the t values are evenly spaced. So, in this model, the X-ray sources rotate by a fixed angle from one set of readings to the next and, within each setting, the individual X-ray beams are evenly spaced. This is called the *parallel beam geometry*.

If the machine takes scans at N different angles, then $\Delta\theta = \pi/N$ and the specific values of θ that occur are $\{k\pi/N : 0 \leq k \leq N-1\}$. We are assuming that the beams at each angle form a set of parallel lines. The spacing between these beams, say d , is called the *sample spacing*. For instance, suppose there are $2 \cdot M + 1$ parallel X-ray beams at each angle. With the object to be scanned centered at the origin, the corresponding values of t are $\{j \cdot d : -M \leq j \leq M\}$. (The specific values of M and d essentially depend on the design of the machine itself and on the sizes of the objects the machine is designed to scan.) Thus, the continuous Radon transform $\mathcal{R}f$ is replaced by the discrete function $\mathcal{R}_D f$ defined by

$$\mathcal{R}_D f_{j,k} = \mathcal{R}f(jd, k\pi/N), \quad (8.14)$$

for $-M \leq j \leq M$ and $0 \leq k \leq (N-1)$.

To implement formula (8.2), we must now decide what we mean by the convolution of two functions for which we have only sampled values. At the same time, we will adopt some rules and conventions for dealing with discretely defined functions in general.

8.5 Discrete functions and convolution

A discrete function of one variable is a mapping from the integers into the set of real or complex numbers. In other words, a discrete function g may be thought of as a two-way infinite list, or sequence, of real or complex numbers $\{\dots, g(-2), g(-1), g(0), g(1), g(2), \dots\}$. For a discrete function g , we customarily identify $g_n = g(n)$ for all integers n .

Definition 8.5. In analogy with the integral that defines continuous convolution, the *discrete convolution* of two discrete functions f and g , denoted by $f \bar{*} g$, is defined by

$$(f \bar{*} g)_m = \sum_{j=-\infty}^{\infty} f_j \cdot g_{(m-j)} \quad \text{for each integer } m. \quad (8.15)$$

As with infinite integrals, there is the issue, which we evade for the time being, of whether this sum converges.

Proposition 8.6. *A few properties of discrete convolution are*

- (i) $f \bar{*} g = g \bar{*} f$,
- (ii) $f \bar{*} (g + h) = f \bar{*} g + f \bar{*} h$, and
- (iii) $f \bar{*} (\alpha g) = \alpha (f \bar{*} g)$,

for suitable functions f , g , and h , and for any constant α .

In practice, we typically know the values of a function f only at a finite set of points, say $\{k \cdot d : k = 0, 1, \dots, (N-1)\}$, where N is the total number of points at which f has been computed and d is the sample spacing. For simplicity, let $f_k = f(kd)$ for each k between 0 and $(N-1)$ and use the set of values $\{f_1, f_2, \dots, f_{N-1}\}$ to represent f as a discrete function. There are two useful ways to extend this sequence to one that is defined for all integers. The simplest approach is just to pad the sequence with zeros by setting $f_k = 0$ whenever k is not between 0 and $N-1$. A more intriguing, and ultimately more useful, idea is to extend the sequence to be *periodic* with period N . Specifically, for any given integer m , there is a unique integer n such that $0 \leq m + n \cdot N \leq (N-1)$. We define $f_m = f_{m+n \cdot N}$. Thus, $f_N = f_0$, $f_{-1} = f_{N-1}$, $f_{N+1} = f_1$, and so on. This defines the discrete function $f = \{f_k\}$ as a periodic function on the set of all integers. We will refer to such a function as an *N -periodic discrete function*.

Definition 8.7. For two N -periodic discrete functions $f = \{f_k : 0 \leq k \leq N-1\}$ and $g = \{g_k : 0 \leq k \leq N-1\}$, the discrete convolution, denoted as before by $f \bar{*} g$, is defined by

$$(f \bar{*} g)_m := \sum_{j=0}^{N-1} f_j \cdot g_{(m-j)} \quad \text{for each integer } m. \quad (8.16)$$

This is obviously similar to Definition 8.5 except that the sum extends only over one full period rather than the full set of integers. Notice that when m and j satisfy $0 \leq m, j \leq (N-1)$, then the difference $(m-j)$ satisfies $|m-j| \leq (N-1)$. The periodicity of the discrete functions enables us to assign values to the discrete convolution function outside the range 0 to $(N-1)$, so that $f \bar{*} g$ also has period N .

The two versions of discrete convolution in (8.15) and (8.16) can be reconciled if one of the two functions has only finitely many nonzero values.

Proposition 8.8. *Let f and g be two-way infinite discrete functions and suppose there is some natural number K such that $g_k = 0$ whenever $k < 0$ or $k \geq K$. Let M be an integer satisfying $M \geq K-1$ and let \tilde{f} and \tilde{g} be the $(2M+1)$ -periodic discrete functions defined by $\tilde{f}(m) = f(m)$ and $\tilde{g}(m) = g(m)$ for $-M \leq m \leq M$. Then, for all m satisfying $0 \leq m \leq K-1$,*

$$(f \bar{*} g)_m = \left(\tilde{f} \bar{*} \tilde{g} \right) (m). \quad (8.17)$$

Proof. Let f, g, K, M, \tilde{f} , and \tilde{g} be as stated. For each pair of integers m and j satisfying $0 \leq m, j \leq K-1$, it follows that $|m-j| \leq K-1$. Thus, since $M \geq K-1$, we get that $f(m-j) = \tilde{f}(m-j)$. Now fix a value of m with $0 \leq m \leq K-1$. From the definition (8.15) and the properties of $\bar{*}$, we have

$$\begin{aligned} (f \bar{*} g)_m &= \sum_{j=-\infty}^{\infty} f(m-j) \cdot g(j) \\ &= \sum_{j=0}^{K-1} f(m-j) \cdot g(j) \\ &= \sum_{j=0}^{K-1} \tilde{f}(m-j) \cdot g(j) \\ &= \sum_{j=-M}^M \tilde{f}(m-j) \cdot \tilde{g}(j) \\ &= \left(\tilde{f} \bar{*} \tilde{g} \right) (m). \quad \square \end{aligned}$$

What is really going on in this proposition? Well, there are a few worries when it comes to using periodic discrete functions. One concern is that it may not be clear what the appropriate period is. So, we might sample a finite set of values of a continuous periodic function, but the values we compute might not correspond to a full period. Then, when we extend the data to form a discrete periodic function, we have the wrong one. Or, the function whose values we have sampled might not be periodic at all. Then we ought not to use a periodic discrete function to model it. The proposition offers a remedy that is known as *zero padding*: we can take a finite set of values of a function (g) and, by padding the sequence of values with a lot of zeros, form a periodic discrete

function (\tilde{g}) in such a way that the periodic discrete convolution gives the same value as the “true” discrete convolution, at least at the points where the value has been sampled.

For an illustration of the benefits of zero padding, consider the function $\square_{1/2}$ defined, as before, by

$$\square_{1/2}(x) = \begin{cases} 1 & \text{if } -1/2 \leq x < 1/2, \\ 0 & \text{otherwise.} \end{cases}$$

Suppose we sample $\square_{1/2}$ at two points $x = -1/2$ and $x = 0$ and, so generate the 2-periodic discrete function f with $f_0 = f_1 = 1$. For the periodic discrete convolution $f \bar{*} f$, we get

$$\begin{aligned} (f \bar{*} f)_0 &= 1 \cdot 1 + 1 \cdot 1 = 2 \text{ and} \\ (f \bar{*} f)_1 &= 1 \cdot 1 + 1 \cdot 1 = 2. \end{aligned}$$

Now pad the function f with a couple of zeros to get the 4-periodic function $\tilde{f} = \{1, 1, 0, 0\}$. (We could also think of this as sampling the function $\square_{1/2}$ at the x values $-1/2, 0, 1/2$, and 1.) Then we get

$$\begin{aligned} (\tilde{f} \bar{*} \tilde{f})_0 &= 1 \cdot 1 + 1 \cdot 0 + 0 \cdot 0 + 0 \cdot 1 = 1, \\ (\tilde{f} \bar{*} \tilde{f})_1 &= 1 \cdot 1 + 1 \cdot 1 + 0 \cdot 0 + 0 \cdot 0 = 2, \\ (\tilde{f} \bar{*} \tilde{f})_2 &= 1 \cdot 0 + 1 \cdot 1 + 0 \cdot 1 + 0 \cdot 1 = 1, \text{ and} \\ (\tilde{f} \bar{*} \tilde{f})_3 &= 1 \cdot 0 + 1 \cdot 0 + 0 \cdot 1 + 0 \cdot 1 = 0. \end{aligned}$$

Instead of accurately representing $\square_{1/2}$, the function f acts like a constant, and so $f \bar{*} f$ is constant, too. But \tilde{f} does a better job of representing $\square_{1/2}$, and the convolution $\tilde{f} \bar{*} \tilde{f}$ is a discrete version of the function $2(\square_{1/2} * \square_{1/2}) = 2(1 - |x|)$, for $-1 \leq x < 1$, sampled at the x values $-1/2, 0, 1/2$, and 1.

Importantly, Proposition 8.8 applies to the discrete convolution of the sampled band-limited function $\mathcal{F}^{-1}A$, where A is a low-pass filter, and the sampled Radon transform $\mathcal{R}_D f$, where f is the attenuation function we wish to reconstruct. Since the scanned object is finite in size, we can set $\mathcal{R}_D f(j, \theta) = 0$ whenever $|j|$ is sufficiently large. Thus, with enough zero padding, the discrete Radon transform (8.14) can be extended to be periodic in the radial variable (jd), and (8.17) shows how to compute the desired discrete convolution.

For discrete functions defined using polar coordinates, the discrete convolution is carried out in the radial variable only. In particular, for a given filter A , we compute the discrete convolution of the sampled inverse Fourier transform of A with the discrete Radon transform of f as

$$(\mathcal{F}^{-1}A \bar{*} \mathcal{R}_D f)_{m, \theta} = \sum_{j=0}^{N-1} \mathcal{F}^{-1}A_j \cdot \mathcal{R}_D f_{m-j, \theta}. \quad (8.18)$$

8.6 Discrete Fourier transform

For a continuous function f , the Fourier transform is defined by

$$\mathcal{F}f(\omega) = \int_{-\infty}^{\infty} f(x) e^{-i\omega x} dx$$

for every real number ω . For a discrete analogue to this, we consider an N -periodic discrete function f . Replacing the integral in the Fourier transform by a sum and summing over one full period of f yields the expression

$$\sum_{k=0}^{N-1} f_k e^{-i\omega k}.$$

Now take ω to have the form $\omega_j = 2\pi j/N$. Then, for each choice of k , we get $e^{-i\omega_j k} = e^{-i2\pi k j/N}$. This is a periodic function with period N in the j variable, just as f_k has period N in the k variable. In this way, the Fourier transform of the discrete function f is also a discrete function, defined for ω_j , with $j = 0, 1, \dots, (N-1)$, and extended by periodicity to every integer.

Definition 8.9. The *discrete Fourier transform*, denoted by \mathcal{F}_D , transforms an N -periodic discrete function f into another N -periodic discrete function $\mathcal{F}_D f$ defined by

$$(\mathcal{F}_D f)_j = \sum_{k=0}^{N-1} f_k e^{-i2\pi k j/N} \quad \text{for } j = 0, 1, \dots, (N-1). \quad (8.19)$$

For other integer values of j , the value of $(\mathcal{F}_D f)_j$ is defined by the periodicity requirement.

Remark 8.10. In the summation (8.19), we can replace the range $0 \leq k \leq (N-1)$ by any string of the form $M \leq k \leq (M+N-1)$, where M is an integer. This is due to the periodicity of the discrete function f .

Example 8.11. Fix a natural number M and set

$$f_k = \begin{cases} 1 & \text{if } -M \leq k \leq M, \\ 0 & \text{otherwise.} \end{cases}$$

Now let $N > 2M$ and think of f as being an N -periodic discrete function. (Thus, f consists of a string of 1s possibly padded with some 0s.) We might think of f as a discrete sampling of the characteristic function of a finite interval. The continuous Fourier transform of a characteristic function is a function of the form $a \sin(bx)/x$.




Fig. 8.3. Continuous and discrete Fourier transforms of a square wave. Zero padding has been used in the sampling of the square wave. The “wrap around” in the second picture illustrates the N -periodic behavior of the discrete transform.

In the discrete setting, use the periodicity of f to compute, for each $j = 0, 1, \dots, N-1$,

$$\begin{aligned}
 (\mathcal{F}_D f)_j &= \sum_{k=0}^M e^{-2\pi i k j / N} + \sum_{k=N-M}^{N-1} e^{-2\pi i k j / N} \\
 &= e^0 + \sum_{k=1}^M \left(e^{-2\pi i k j / N} + e^{-2\pi i (N-k) j / N} \right) \\
 &= 1 + 2 \cdot \sum_{k=1}^M \cos(2\pi k j / N) \\
 &= 2 \cdot (1/2 + \cos(2\pi j / N) + \cos(2\pi 2 j / N) + \dots + \cos(2\pi M j / N)) \\
 &= 2 \cdot \frac{\sin((M+1/2) \cdot 2\pi j / N)}{2 \cdot \sin(\pi j / N)} \\
 &= \frac{\sin((2M+1)\pi j / N)}{\sin(\pi j / N)}.
 \end{aligned}$$

We have used the identity

$$1/2 + \cos(\theta) + \cos(2\theta) + \dots + \cos(M\theta) = \frac{\sin((M+1/2)\theta)}{2 \cdot \sin(\theta/2)}. \quad (8.20)$$

Figure 8.3 shows a comparison, with $M = 8$ and $N = 20$, of the graph of

$$y = \frac{\sin((2 \cdot M + 1)x/2)}{x/2}$$

with the set of points

$$\left(2\pi j / N, \frac{\sin((2 \cdot M + 1)\pi j / N)}{\sin(\pi j / N)} \right),$$

where the point $(0, (2 \cdot M + 1))$ corresponds to $j = 0$. The diagram on the left shows the situation for $0 \leq x \leq 2\pi$ and $j = 0, 1, \dots, N - 1$. While the curve continues its decay, the points representing the discrete Fourier transform reveal the N -periodic behavior. In the diagram on the right, the interval is $-\pi \leq x \leq \pi$, corresponding to $j = -N/2, \dots, N/2$. In comparing the two diagrams, we can see how the N -periodic behavior of the discrete Fourier transform works.

Where there is a discrete Fourier transform, there must also be an inverse transform. Indeed, the formula (8.2) demands it. To this end, we have the following.

Definition 8.12. For an N -periodic discrete function g , the *discrete inverse Fourier transform* of g is the N -periodic function defined by

$$(\mathcal{F}_D^{-1}g)_n = \frac{1}{N} \sum_{k=0}^{N-1} g_k e^{i2\pi kn/N}, \text{ for } n = 0, \dots, N - 1. \quad (8.21)$$

For example, we have already computed the discrete inverse Fourier transforms for the Shepp–Logan and Ram–Lak low-pass filters, respectively, in (8.10) and (8.13).

The discrete analogue of the Fourier inversion theorem (Theorem 5.11) holds.

Theorem 8.13. For a discrete function f with period N ,

$$\mathcal{F}_D^{-1}(\mathcal{F}_D f)_n = f_n \text{ for all integers } n. \quad (8.22)$$

Before we prove this theorem, we need the following fact about roots of unity.

Lemma 8.14. For any nonzero integers M and N ,

$$\sum_{k=0}^{N-1} e^{i2\pi Mk/N} = \begin{cases} N & \text{if } M/N \text{ is an integer,} \\ 0 & \text{if } M/N \text{ is not an integer.} \end{cases} \quad (8.23)$$

Proof. If M/N is an integer, then $e^{i2\pi Mk/N} = 1$ for every integer k , from which it follows that $\sum_{k=0}^{N-1} e^{i2\pi Mk/N} = N$.

If M/N is not an integer, then $(e^{i2\pi M/N})^N - 1 = 0$, but $e^{i2\pi M/N} \neq 1$. Since the expression $x^N - 1$ factors as $x^N - 1 = (x - 1)(1 + x + x^2 + \dots + x^{N-1})$, it follows that

$$1 + e^{i2\pi M/N} + \left(e^{i2\pi M/N}\right)^2 + \dots + \left(e^{i2\pi M/N}\right)^{N-1} = 0.$$

That is,

$$\sum_{k=0}^{N-1} e^{i2\pi Mk/N} = 0$$

as claimed. \square

Proof of Theorem 8.13. For a given integer n with $0 \leq n \leq (N-1)$, compute

$$\begin{aligned} \mathcal{F}_D^{-1}(\mathcal{F}_D f)_n &= \frac{1}{N} \sum_{k=0}^{N-1} (\mathcal{F}_D f)_k e^{i2\pi kn/N} \\ &= \frac{1}{N} \sum_{k=0}^{N-1} \left(\sum_{m=0}^{N-1} f_m e^{-i2\pi mk/N} \right) e^{i2\pi kn/N} \\ &= \frac{1}{N} \sum_{m=0}^{N-1} \left[f_m \left(\sum_{k=0}^{N-1} e^{i2\pi(n-m)k/N} \right) \right]. \end{aligned}$$

When $m = n$, we have $e^{i2\pi(n-m)k/N} = 1$ for every k . So

$$\sum_{k=0}^{N-1} e^{i2\pi(n-m)k/N} = N \text{ when } m = n.$$

However, when $m \neq n$, then Lemma 8.14 asserts that

$$\sum_{k=0}^{N-1} e^{i2\pi(n-m)k/N} = 0.$$

It follows that

$$\frac{1}{N} \sum_{m=0}^{N-1} \left[f_m \left(\sum_{k=0}^{N-1} e^{i2\pi(n-m)k/N} \right) \right] = \frac{1}{N} [f_n \cdot N] = f_n.$$

That is, $\mathcal{F}_D^{-1}(\mathcal{F}_D f)_n = f_n$ for $0 \leq n \leq (N-1)$ and, by periodicity, for all integers n . \square

Were we so inclined, we could establish properties about linearity, shifting, and so on for the discrete Fourier transform and its inverse. However, our focus is on the implementation of the formula (8.2), so our priority must be to study the interaction between the discrete Fourier transform and discrete convolution. In particular, the following theorem holds.

Theorem 8.15. For two discrete functions $f = \{f_k : 0 \leq k \leq N-1\}$ and $g = \{g_k : 0 \leq k \leq N-1\}$ with the same period, we have

$$\mathcal{F}_D(f \bar{*} g) = (\mathcal{F}_D f) \cdot (\mathcal{F}_D g). \quad (8.24)$$

In words, the discrete Fourier transform of a convolution is the product of the discrete transforms individually.

Proof. For each integer n such that $0 \leq n \leq (N-1)$,

$$\begin{aligned}
(\mathcal{F}_D f)_n \cdot (\mathcal{F}_D g)_n &= \left(\sum_{k=0}^{N-1} f_k e^{-2\pi i n k / N} \right) \cdot \left(\sum_{\ell=0}^{N-1} g_\ell e^{-2\pi i n \ell / N} \right) \\
&= \sum_{k=0}^{N-1} f_k \left(\sum_{\ell=-k}^{-k+N-1} g_\ell e^{-2\pi i n \ell / N} \right) e^{-2\pi i n k / N} \quad (\text{by periodicity of } g) \\
&= \sum_{k=0}^{N-1} f_k \left(\sum_{j=0}^{N-1} g_{j-k} e^{-2\pi i n (j-k) / N} \right) e^{-2\pi i n k / N} \quad (\text{where } j = k + \ell) \\
&= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \left(f_k \cdot g_{j-k} e^{-2\pi i n j / N} \right) \\
&= \sum_{j=0}^{N-1} \left(\sum_{k=0}^{N-1} f_k \cdot g_{j-k} \right) e^{-2\pi i n j / N} \\
&= \sum_{j=0}^{N-1} (f \bar{*} g)_j e^{-2\pi i n j / N} \\
&= [\mathcal{F}_D (f \bar{*} g)]_n,
\end{aligned}$$

which is the desired result. \square

We conclude this section by mentioning a few results that are discrete versions of some of the results from earlier chapters. The proofs are left as exercises.

Proposition 8.16. *In analogy to (7.14), for two discrete N -periodic functions f and g ,*

$$\mathcal{F}_D (f \cdot g) = \frac{1}{N} (\mathcal{F}_D f) \bar{*} (\mathcal{F}_D g). \quad (8.25)$$

For a discrete function f , define the conjugate function \bar{f} by setting $\bar{f}_k = \overline{f_k}$ for every k . In analogy to (5.10), we have the following fact.

Proposition 8.17. *For an N -periodic discrete function f ,*

$$(\mathcal{F}_D \bar{f})_j = \overline{(\mathcal{F}_D f)_{-j}} \quad (8.26)$$

for all j .

There is also a discrete version of the Rayleigh–Plancherel theorem, Theorem 7.8:

Proposition 8.18. *For an N -periodic discrete function f ,*

$$\sum_{n=0}^{N-1} |f_n|^2 = \frac{1}{N} \sum_{n=0}^{N-1} |(\mathcal{F}_D f)_n|^2. \quad (8.27)$$

8.7 Discrete back projection

In the continuous setting, the back projection is defined by

$$\mathcal{B}h(x, y) := \frac{1}{\pi} \int_{\theta=0}^{\pi} h(x \cos(\theta) + y \sin(\theta), \theta) d\theta. \quad (8.28)$$

Definition 8.19. In the discrete setting, the continuously variable angle θ is replaced by the discrete set of angles $\{k\pi/N : 0 \leq k \leq N-1\}$. So the value of $d\theta$ becomes π/N and the back-projection integral is replaced by the sum

$$\mathcal{B}_D h(x, y) = \left(\frac{1}{N}\right) \sum_{k=0}^{N-1} h(x \cos(k\pi/N) + y \sin(k\pi/N), k\pi/N). \quad (8.29)$$

This is to be applied to $h = (\mathcal{F}_D^{-1}A) \bar{*}(\mathcal{R}_D f)$. The obstacle here is that the reconstruction grid, within which the final image is to be presented, is a rectangular array of pixels, located at points $\{(x_m, y_n)\}$, each of which is to be assigned a color or greyscale value. Thus, the discrete back projection (8.29) requires the values of $(\mathcal{F}_D^{-1}A) \bar{*}(\mathcal{R}_D f)$ at the corresponding points $\{(x_m \cos(k\pi/N) + y_n \sin(k\pi/N), k\pi/N)\}$. However, the Radon transform samples and the values of $(\mathcal{F}_D^{-1}A) \bar{*}(\mathcal{R}_D f)$ are known only at the points $\{(jd, k\pi/N)\}$. These points are arranged in a polar grid and generally do not match up with the points needed.

To overcome this obstacle, observe that, for a given (x, y) and a given k , the number $x \cos(k\pi/N) + y \sin(k\pi/N)$ will lie in between two multiples of d . That is, there is some value n such that

$$nd \leq x \cos(k\pi/N) + y \sin(k\pi/N) < (n+1)d.$$

Hence, we wish to *interpolate* a value for $(\mathcal{F}_D^{-1}A) \bar{*}(\mathcal{R}_D f)$ at the point $(x \cos(k\pi/N) + y \sin(k\pi/N), k\pi/N)$, using the known values at nearby points.

8.8 Interpolation

When we are given (or have computed) a discrete set of values $f_k = f(x_k)$ at a finite set of inputs $\{x_k\}$, the process of somehow assigning values $f(x)$ for inputs x between the distinct $\{x_k\}$ in order to create a continuous or, at least, piecewise continuous function defined on an interval is called *interpolation*. Exactly how the interpolated values are assigned depends on factors such as what additional properties (continuous, differentiable, infinitely differentiable, *et cetera*) the function is to have, the degree of

computational difficulty incurred in assigning the new values, and more. Let us look at some interpolation schemes that are commonly used.

Nearest neighbor. Given the values $\{f(x_k)\}$, the value $f(x)$ is assigned to be the same as the value of f at the input x_k that is nearest to x . This creates a step function that steps up or down halfway between successive values of x_k . This is computationally simple, but the interpolated function is not continuous, and what to do at the halfway points between the x_k s is not clear.

Linear. Literally just “connect the dots,” thus creating a continuous function composed of line segments. Computationally, for x lying between x_k and x_{k+1} , we set

$$f(x) = \left(\frac{f(x_{k+1}) - f(x_k)}{x_{k+1} - x_k} \right) \cdot (x - x_k) + f(x_k). \quad (8.30)$$

This method is often used in medical imaging. It is computationally simple and produces a continuous function f .

Cubic spline. Successive points $(x_k, f(x_k))$ and $(x_{k+1}, f(x_{k+1}))$ are connected by part of a cubic curve, $y = ax^3 + bx^2 + cx + d$. The pieces are joined together so that the resulting overall curve has a continuous second-order derivative. In addition to the values $f(x_k)$, values for the derivative $f'(x_k)$ are prescribed. Cubic splines are computationally more complicated than linear interpolation, but are still manageable. Computing the coefficients of each cubic piece involves solving a system of linear equations. Splines are commonly used in a variety of engineering applications, including medical imaging.

Lagrange, or polynomial, interpolation. When $N + 1$ data points are known, then it is possible to fit a polynomial of degree N to the data. For instance, three noncollinear points determine a parabola. Generally, a polynomial of degree N has $N + 1$ coefficients. The $N + 1$ data points provide $N + 1$ equations that these coefficients must satisfy. A general formula, given values of $f(x_k)$ for $k = 1, \dots, N + 1$, is

$$f(x) = \sum_{j=1}^{N+1} f(x_j) \cdot \frac{\prod_{k \neq j} (x - x_k)}{\prod_{k \neq j} (x_j - x_k)}. \quad (8.31)$$

Notice that, if we take $x = x_n$ for some n , then the formula (8.31) yields $f(x) = f(x_n)$ as it should. So the formula agrees with the data.

To see how the interpolation process can be generalized, take a closer look at nearest-neighbor interpolation. We start with a discrete function g , where $g(m)$ denotes the value of g at the sample point $m \cdot d$, and we want to construct an interpolated function $\mathcal{I}(g)$ using the nearest-neighbor method. For a given value x , the sample point $m \cdot d$ is the nearest neighbor to x exactly when $|x - m \cdot d| < d/2$, or, equivalently, when $|(x/d) - m| < 1/2$. It follows that

$$\mathcal{I}(g)(x) = \sum_m g(m) \cdot \square_{1/2}\left(\frac{x}{d} - m\right), \text{ for all } x, \quad (8.32)$$

where, as usual,

$$\square_{1/2}(x) = \begin{cases} 1 & \text{if } |x| < 1/2, \\ 0 & \text{if } |x| > 1/2. \end{cases}$$

(For $x = \pm 1/2$, make a slight modification by assigning the value $+1$ to one and 0 to the other, depending whether we want to consider $-1/2$ or $+1/2$ as “closer” to 0 .) This approach offers a nice compact formula for nearest-neighbor interpolation.

Adopting a similar approach to linear interpolation, suppose that a given x satisfies

$$m_* \cdot d \leq x < (m_* + 1) \cdot d$$

for some integer m_* . Then

$$\begin{aligned} \left| \frac{x}{d} - m_* \right| &\leq 1, \\ \left| \frac{x}{d} - (m_* + 1) \right| &\leq 1, \end{aligned}$$

and, for all integers m other than m_* and $(m_* + 1)$,

$$\left| \frac{x}{d} - m \right| > 1.$$

Consider again the “tent” function \wedge , defined by

$$\wedge(x) = \begin{cases} 1 - |x| & \text{if } |x| \leq 1, \\ 0 & \text{if } |x| > 1. \end{cases}$$

For a discrete function g , evaluated at the sample points $\{m \cdot d\}$, and the corresponding function $\mathcal{I}(g)$ obtained from g by linear interpolation, we get, for x satisfying $m_* \cdot d \leq x < (m_* + 1) \cdot d$,

$$\begin{aligned} \mathcal{I}(g)(x) &= \frac{g(m_* + 1) - g(m_*)}{d} \cdot (x - m_* \cdot d) + g(m_*) \\ &= (g(m_* + 1) - g(m_*)) \cdot \left(\frac{x}{d} - m_*\right) + g(m_*) \\ &= g(m_*) \left(1 - \left(\frac{x}{d} - m_*\right)\right) + g(m_* + 1) \left(\frac{x}{d} - m_*\right) \\ &= g(m_*) \left(1 - \left|\frac{x}{d} - m_*\right|\right) + g(m_* + 1) \left(1 - \left|\frac{x}{d} - (m_* + 1)\right|\right) \\ &= \sum_m g(m) \cdot \wedge\left(\frac{x}{d} - m\right). \end{aligned} \quad (8.33)$$

Generalizing this approach, we can design an interpolation method like so.

Definition 8.20. For a selected weighting function W , satisfying the conditions below, the W -interpolation $\mathcal{I}_W(g)$ of a discrete function g is defined by

$$\mathcal{I}_W(g)(x) = \sum_m g(m) \cdot W\left(\frac{x}{d} - m\right) \text{ for } -\infty < x < \infty. \quad (8.34)$$

Of course, for this to be reasonable, the weighting function W should satisfy a few conditions. For instance, if g is real-valued, bounded, or even, then $\mathcal{I}_W(g)$ should be too, which means that W should be real-valued, bounded, and even. Also, at the sample points, we expect the interpolation to be exact. That is, for any integer k , we expect $\mathcal{I}_W(g)(k \cdot d) = g(k)$, which implies that $W(0) = 1$ and that $W(m) = 0$ for all integers $m \neq 0$. Lastly, for purposes of integration, it would be nice for W -interpolation to “preserve areas” in the sense that the integral of the W -interpolated function $\mathcal{I}_W(g)$ is actually *equal to* the approximation we get when we apply the trapezoidal rule to the sampled points $\{(m \cdot d, g(m))\}$. That is, we would like to have

$$\int_{-\infty}^{\infty} \mathcal{I}_W(g)(x) dx = d \cdot \sum_m g(m). \quad (8.35)$$

To see what this implies about W , observe that, for each integer m ,

$$\begin{aligned} \int_{-\infty}^{\infty} W\left(\frac{x}{d} - m\right) dx &= d \cdot \int_{-\infty}^{\infty} W(u - m) du \text{ where } u = x/d \\ &= d \cdot \int_{-\infty}^{\infty} W(u) du. \end{aligned}$$

In particular, the value of the integral is independent of m . From (8.34), we now compute

$$\begin{aligned} \int_{-\infty}^{\infty} \mathcal{I}_W(g)(x) dx &= \int_{-\infty}^{\infty} \left\{ \sum_m g(m) \cdot W\left(\frac{x}{d} - m\right) \right\} dx \\ &= \sum_m g(m) \cdot \int_{-\infty}^{\infty} W\left(\frac{x}{d} - m\right) dx \\ &= d \cdot \int_{-\infty}^{\infty} W(u) du \cdot \sum_m g(m). \end{aligned}$$

(The interchange of the summation and the integral is valid because we are summing over finitely many m .) Thus, for (8.35) to hold, we want W to satisfy

$$\int_{-\infty}^{\infty} W(u) du = 1.$$

In addition to the functions $\square_{1/2}$ and \wedge considered above, the function $x \mapsto \frac{\sin(\pi x)}{(\pi x)}$ satisfies these conditions on the weighting function.

Interpolation and convolution. For a discrete function g and a weighting function W , let $\mathcal{I}_W(g)$ be as in (8.34). Suppose also that $g = \phi \bar{*} f$ for some discrete functions ϕ and f . That is, suppose that, for each m ,

$$g(m) = \sum_k \phi(m-k) \cdot f(k).$$

Then, for each x ,

$$\begin{aligned} \mathcal{I}_W(g)(x) &= \sum_m \left(\sum_k \phi(m-k) \cdot f(k) \right) \cdot W\left(\frac{x}{d} - m\right) \\ &= \sum_k \left(\sum_m \phi(m-k) \cdot W\left(\frac{x}{d} - m\right) \right) \cdot f(k) \\ &= \sum_k \left(\sum_m \phi(m-k) \cdot W\left(\frac{x-kd}{d} - (m-k)\right) \right) f(k). \end{aligned} \quad (8.36)$$

It is tempting to mistake the sum $\sum_m \phi(m-k) \cdot W\left(\frac{x-kd}{d} - (m-k)\right)$ in (8.36) for the interpolated function $\mathcal{I}_W(\phi)(x-kd)$. The discrete function ϕ is periodic after all, so shifting the summation from summing over m to summing over $(m-k)$ shouldn't matter. However, W is not periodic, so there is a glitch when $k = N-1$ for N -periodic functions. That's where ϕ "wraps around" but W does not. Nonetheless, it is *almost correct* to say that these two expressions are the same. So, from (8.36), we can write

$$\mathcal{I}_W(\phi \bar{*} f)(x) \approx \sum_k \mathcal{I}_W(\phi)(x-kd) \cdot f(k). \quad (8.37)$$

The expression on the left in (8.37) represents an interpolation of the filtered form of f . On the right, the interpolation is brought inside and the expression has the form of a weighted sum of the values $\{f(k)\}$, where the weight depends on the distance between x and the sampled point kd .

8.9 Discrete image reconstruction

We have now examined the discrete versions of all of the elements of the formula (8.2), used to approximate the image f starting from a discrete set of samples of its Radon transform.

Image Reconstruction Algorithm I. In the implementation of the discrete filtered back-projection algorithm (8.2), the function to be interpolated is $(\mathcal{F}_D^{-1}A) \bar{*} (\mathcal{R}_D f)$, where A is the low-pass filter used in place of the absolute-value function. For simplicity of notation, let us denote the interpolated function by \mathcal{I} . That is, the value $\mathcal{I}(t, k\pi/N)$ is interpolated, according to the selected interpolation method, from the computed values $(\mathcal{F}_D^{-1}A) \bar{*} (\mathcal{R}_D f)(jd, k\pi/N)$. For each lattice point (x_m, y_n) in the

reconstruction grid, we then approximate

$$\begin{aligned} f(x_m, y_n) &\approx \left(\frac{1}{2}\right) \mathcal{B}_D \mathcal{I}(x_m, y_n) \\ &= \left(\frac{1}{2N}\right) \sum_{k=0}^{N-1} \mathcal{I}\left(x_m \cos\left(\frac{k\pi}{N}\right) + y_n \sin\left(\frac{k\pi}{N}\right), \frac{k\pi}{N}\right). \end{aligned} \quad (8.38)$$

Image Reconstruction Algorithm II. Alternatively, we can use the approach indicated by (8.37). Instead of interpolating the filtered Radon transform $(\mathcal{F}_D^{-1}A) \ast (\mathcal{R}_D f)$, we interpolate the filter itself and then form a weighted sum of the sampled values of the Radon transform. That is, we replace each value of \mathcal{I} in (8.38) by the weighted sum

$$\mathcal{W}(k) = \sum_j \mathcal{I}_{\mathcal{F}_D^{-1}A}\left(x_m \cos\left(\frac{k\pi}{N}\right) + y_n \sin\left(\frac{k\pi}{N}\right) - jd\right) \cdot \mathcal{R}_D f\left(jd, \frac{k\pi}{N}\right).$$

This leads to a different approximation to the discrete filtered back projection given by

$$f(x_m, y_n) \approx \left(\frac{1}{2N}\right) \sum_{k=0}^{N-1} \mathcal{W}(k). \quad (8.39)$$

Now let's try to put this all together.

Example 8.21. As a first example, define a radially symmetric attenuation function f by

$$f(x, y) = \begin{cases} 1/2 & \text{if } x^2 + y^2 \leq 1/16, \\ 1/4 & \text{if } 1/16 < x^2 + y^2 \leq 1/4, \\ 1 & \text{if } 1/4 < x^2 + y^2 \leq 9/16, \\ 0 & \text{if } 9/16 < x^2 + y^2. \end{cases}$$

Recalling that we computed the Radon transform of a disc in Chapter 2, and using the linearity of the Radon transform, we get that

$$\mathcal{R}f(t) = \begin{cases} (1/2)\sqrt{9-16t^2} - (3/4)\sqrt{1-4t^2} + (1/8)\sqrt{1-16t^2} & \text{if } |t| \leq 1/4, \\ (1/2)\sqrt{9-16t^2} - (3/4)\sqrt{1-4t^2} & \text{if } 1/4 < |t| \leq 1/2, \\ (1/2)\sqrt{9-16t^2} & \text{if } 1/2 < |t| \leq 3/4, \\ 0 & \text{if } 3/4 < |t|. \end{cases}$$

This is independent of θ because of the radial symmetry of f . The graphs of f and of $\mathcal{R}f(t)$ are shown in Figure 8.4.

In the rest of this example, the sample spacing is taken to be $d = 0.05$. The corresponding values of the Radon transform are $\mathcal{R}f(0.05 \cdot k)$ for $-20 \leq k \leq 20$.

For the filter, we will use the Shepp–Logan filter defined in (8.8). Its discrete inverse Fourier transform is given in (8.10). The sample spacing used in that computation




Fig. 8.4. A radially symmetric attenuation function and its Radon transform.

was π/L . However, those values were computed from the continuous inverse Fourier transform, as in (5.16). To compensate for the presence of the additional factor of 2π in the complex exponential in the definition (8.21), we should modify the sample spacing from π/L to $\pi/(2\pi L)$, or simply $1/(2L)$. To match up the sample spacing of this filter with that of the Radon transform, we want $1/(2L) = d = 0.05$; so $L = 10$. That is, for this example we use the Shepp–Logan filter given by

$$A(\omega) = \begin{cases} (20/\pi) \cdot |\sin(.05\pi\omega)| & \text{if } |\omega| \leq 10, \\ 0 & \text{if } |\omega| > 10. \end{cases}$$

Its discrete inverse Fourier transform is given by

$$(\mathcal{F}_D^{-1}A)_n = \frac{400}{\pi^3(1-4n^2)}.$$

The graphs of A and $\mathcal{F}_D^{-1}A$ are shown in Figure 8.1.

Next, we compute the discrete convolution $(\mathcal{F}_D^{-1}A) \star \mathcal{R}f$, as defined in (8.16). Denote this discrete function by γ , for short. So, for $-19 \leq m \leq 20$,

$$\begin{aligned} \gamma(m) &= \sum_{j=-19}^{20} (\mathcal{F}_D^{-1}A)_{(m-j)} \cdot \mathcal{R}f(.05 \cdot j) \\ &= \sum_{j=-19}^{20} \frac{400}{\pi^3(1-4(m-j)^2)} \cdot \mathcal{R}f(.05 \cdot j). \end{aligned}$$

Applying linear interpolation to γ as in (8.33), the interpolated function h is given by

$$h(t) = \sum_{m=-19}^{20} \gamma(m) \cdot \wedge(20t - m) \text{ for } -1 \leq t \leq 1.$$




Fig. 8.5. The linear interpolation h of $(\mathcal{F}_D^{-1}A) \star \mathcal{R}f$.

Of course, both γ and h are independent of θ because $\mathcal{R}f$ is. The graph of h is shown in Figure 8.5.

Finally, we implement the formula (8.38), from Image Reconstruction Algorithm I, with the angle θ sampled in increments of $\pi/18$. We get

$$f(x, y) \approx \left(\frac{1}{36} \right) \sum_{k=0}^{17} h(x \cos(k\pi/18) + y \sin(k\pi/18)).$$

On a 100×100 grid in the square $\{(x, y) : -1 \leq x \leq 1, -1 \leq y \leq 1\}$ this yields the picture shown on the left in Figure 8.6. Doubling the number of angle samples to 36 yields the picture shown on the right.

Example 8.22. Now consider the crescent-shaped phantom illustrated in Figure 2.2 and defined by the attenuation function in (2.8). This time, we use the Ram-Lak filter




Fig. 8.6. The discrete filtered back-projection reconstruction of the “bull’s-eye.” Algorithm (8.38) was used with sample spacing $d = 0.05$. The angle was sampled in increments of $\pi/18$ on the left and $\pi/36$ on the right.




Fig. 8.7. The discrete filtered back-projection algorithm (8.39) applied to a crescent-shaped phantom with sample spacing $d = 0.05$. The angle is sampled in increments of $\pi/18$ on the left and $\pi/36$ on the right.

defined in (8.11). As in the previous example, in order to match the sample spacing of the discrete inverse transform of the low-pass filter with that of the Radon transform, we set $L = 1/(2d)$. Taking $d = 0.05$ as before, we therefore take $L = 10$ in (8.13). That is, the filter is given by

$$A(\omega) = \begin{cases} |\omega| & \text{if } |\omega| \leq 10, \\ 0 & \text{if } |\omega| > 10, \end{cases}$$

and its discrete inverse Fourier transform is given by

$$(\mathcal{F}_D^{-1}A)_n = \frac{50}{\pi} \cdot \left\{ \frac{2 \cdot \sin(\pi n)}{\pi n} - \left[\frac{\sin(\pi n/2)}{(\pi n/2)} \right]^2 \right\}.$$

This is pictured in Figure 8.2.

When we apply Image Reconstruction Algorithm II, given by (8.39), with linear interpolation, the results are shown in Figure 8.7.

Example 8.23. Image comparison. By altering the choices of algorithm, low-pass filter, interpolation method, and sample spacings, we can create a variety of reconstructions of the same phantom. Comparing the results helps us to assess how our choices affect the results. For example, the illustrations on the left sides of Figures 8.7 and 8.8 differ only in the choice of low-pass filter—Ram-Lak compared to Shepp-Logan. The picture on the right in Figure 8.8 depicts the difference between the two images. Locations where the two agree are shown as neutral grey; lighter or darker gradations of grey indicate locations where the two reconstructions differ. For instance, the differences are most pronounced at the outer ring of the crescent, where we see both a bright band and a dark band in the picture.




Fig. 8.8. On the left, algorithm (8.39) is applied to the crescent-shaped phantom using the Shepp–Logan filter with sample spacing $d = 0.05$ and angle increments of $\pi/18$. The figure on the right depicts the difference between this reconstruction and the one on the left in Figure 8.7, which used the Ram–Lak filter.

8.10 Matrix forms

For a discrete function $\mathbf{f} = \langle f_0, \dots, f_{N-1} \rangle$ having period N , we have defined the discrete Fourier transform of \mathbf{f} to be the discrete function $\mathbf{F} = \langle F_0, \dots, F_{N-1} \rangle$, also having period N , satisfying

$$F_j = \sum_{k=0}^{N-1} f_k e^{-i2\pi kj/N} \text{ for } j = 0, 1, \dots, (N-1).$$

Setting $w_N := e^{-i2\pi/N}$, we have $e^{-i2\pi kj/N} = (w_N)^{kj}$. Thus, in matrix form,

$$F_j = \begin{bmatrix} 1 & (w_N)^j & (w_N)^{2j} & \dots & (w_N)^{(N-1)j} \end{bmatrix} \begin{bmatrix} f_0 \\ f_1 \\ f_2 \\ \vdots \\ f_{N-1} \end{bmatrix}. \quad (8.40)$$

Denote by \mathbf{W}_N the $N \times N$ matrix whose entry in row j and column k is the number $(w_N)^{kj}$. It follows from (8.40) that the discrete Fourier transform of \mathbf{f} is given by

$$\mathcal{F}_D \mathbf{f} = \mathbf{W}_N \mathbf{f}, \quad (8.41)$$

where \mathbf{f} is viewed as a column vector in this context.

If we let $\overline{\mathbf{W}_N}$ denote the matrix obtained by taking the complex conjugates of the entries of \mathbf{W}_N , then it follows from Lemma 8.14 that $\overline{\mathbf{W}_N} \mathbf{W}_N = N \cdot I_N$. That is,

$$(\mathbf{W}_N)^{-1} = \frac{1}{N} \overline{\mathbf{W}_N},$$

from which we see that the inverse discrete Fourier transform can be expressed in matrix form as

$$\mathcal{F}_D^{-1} \mathbf{g} = \frac{1}{N} \overline{\mathbf{W}_N} \mathbf{g}, \quad (8.42)$$

where \mathbf{g} is an N -periodic discrete function viewed as a column vector.

Example 8.24. This example illustrates the relation (8.24) in the matrix setting. With $N = 3$, we get $w_3 = e^{-i2\pi/3}$. For two discrete 3-periodic functions $\mathbf{f} = \langle f_0, f_1, f_2 \rangle$ and $\mathbf{g} = \langle g_0, g_1, g_2 \rangle$, the convolution is

$$\mathbf{f} \bar{*} \mathbf{g} = \langle f_0 g_0 + f_1 g_2 + f_2 g_1, f_0 g_1 + f_1 g_0 + f_2 g_2, f_0 g_2 + f_1 g_1 + f_2 g_0 \rangle.$$

Hence,

$$\begin{aligned} \mathbf{W}_3(\mathbf{f} \bar{*} \mathbf{g}) &= \begin{bmatrix} 1 & 1 & 1 \\ 1 & w_3 & (w_3)^2 \\ 1 & (w_3)^2 & w_3 \end{bmatrix} \begin{bmatrix} f_0 g_0 + f_1 g_2 + f_2 g_1 \\ f_0 g_1 + f_1 g_0 + f_2 g_2 \\ f_0 g_2 + f_1 g_1 + f_2 g_0 \end{bmatrix} \\ &= \begin{bmatrix} (f_0 + f_1 + f_2) \cdot (g_0 + g_1 + g_2) \\ (f_0 + w_3 f_1 + (w_3)^2 f_2) \cdot (g_0 + w_3 g_1 + (w_3)^2 g_2) \\ (f_0 + (w_3)^2 f_1 + w_3 f_2) \cdot (g_0 + (w_3)^2 g_1 + w_3 g_2) \end{bmatrix} \\ &= (\mathbf{W}_3 \mathbf{f}) \cdot (\mathbf{W}_3 \mathbf{g}), \end{aligned}$$

where the product in the last step is the entry-by-entry product of two vectors, each viewed as a discrete 3-periodic function.

8.11 FFT—the fast Fourier transform

For an N -periodic discrete function \mathbf{f} , the computation of the discrete Fourier transform of \mathbf{f} requires N^2 multiplications, each of the form $f_k \cdot e^{-2\pi i k j / N}$ with $0 \leq k, j \leq N - 1$. For the large values of N encountered in medical imaging, this implies a significant amount of computational time. With the introduction, in [12], of the *fast Fourier transform*, Cooley and Tukey showed that the computation time can be reduced substantially—by a factor of $\log(N)/N$ —if N is a power of 2.

For starters, suppose that N is an even number, say $N = 2 \cdot M$. Then the N th roots of unity can be divided into two sets. Those of the form $e^{-2\pi i (2k) j / N} = e^{-2\pi i k j / M}$, for $0 \leq k \leq (M - 1)$, are also M th roots of unity. The rest have the form $e^{-2\pi i (2k+1) j / N} = e^{-2\pi i j / N} \cdot e^{-2\pi i k j / M}$, for $0 \leq k \leq (M - 1)$. Thus, for each j between 0 and $(N - 1)$, the corresponding component of the discrete Fourier transform of \mathbf{f} can be expressed as

$$(\mathcal{F}_D \mathbf{f})_j = \sum_{k=0}^{M-1} f_{2k} \cdot e^{-2\pi i k j / M} + \left(e^{-2\pi i j / N} \right)^j \cdot \sum_{k=0}^{M-1} f_{2k+1} \cdot e^{-2\pi i k j / M}. \quad (8.43)$$

Moreover, for $0 \leq k \leq (M-1)$ and $M \leq j \leq (N-1)$, we get

$$e^{-2\pi i k j / M} = e^{-2\pi i k (j-M) / M} \cdot e^{-2\pi i k M / M} = e^{-2\pi i k (j-M) / M}. \quad (8.44)$$

In other words, for $M \leq j \leq (N-1)$, we may replace j with $(j-M)$ in the sums on the right-hand side of (8.43). Also, $(e^{-2\pi i / N})^j = -(e^{-2\pi i / N})^{j-M}$. So now we can express the discrete Fourier transform of \mathbf{f} in an even simpler way. Namely, for $0 \leq j \leq (M-1)$, we have

$$(\mathcal{F}_D \mathbf{f})_j = \sum_{k=0}^{M-1} f_{2k} \cdot e^{-2\pi i k j / M} + (e^{-2\pi i / N})^j \cdot \sum_{k=0}^{M-1} f_{2k+1} \cdot e^{-2\pi i k j / M} \quad (8.45)$$

and

$$(\mathcal{F}_D \mathbf{f})_{j+M} = \sum_{k=0}^{M-1} f_{2k} \cdot e^{-2\pi i k j / M} - (e^{-2\pi i / N})^j \cdot \sum_{k=0}^{M-1} f_{2k+1} \cdot e^{-2\pi i k j / M}. \quad (8.46)$$

The beauty of this is that, for $N = 2M$, we have now represented the discrete Fourier transform of the N -periodic discrete function \mathbf{f} in terms of the discrete Fourier transforms of *two* M -periodic discrete functions $\mathbf{f}^0 = \{f_{2k} : 0 \leq k \leq M-1\}$ and $\mathbf{f}^1 = \{f_{2k+1} : 0 \leq k \leq M-1\}$. In symbols,

$$(\mathcal{F}_D \mathbf{f})_j = (\mathcal{F}_D \mathbf{f}^0)_j + (e^{-2\pi i / N})^j \cdot (\mathcal{F}_D \mathbf{f}^1)_j \quad (8.47)$$

and

$$(\mathcal{F}_D \mathbf{f})_{j+M} = (\mathcal{F}_D \mathbf{f}^0)_j - (e^{-2\pi i / N})^j \cdot (\mathcal{F}_D \mathbf{f}^1)_j, \quad (8.48)$$

for $0 \leq j \leq (M-1)$. Assuming that $(\mathcal{F}_D \mathbf{f}^0)$ and $(\mathcal{F}_D \mathbf{f}^1)$ have already been computed, then only two multiplications are required to compute each component of $(\mathcal{F}_D \mathbf{f})$, for a total of $2N$ multiplications. (For each component, there is one multiplication by the factor 1 and another by the factor $\pm e^{-2\pi i j / N}$.)

Now, if $M = N/2$ also happens to be an even number, then we can apply the same reasoning to split each of the M -periodic discrete functions \mathbf{f}^0 and \mathbf{f}^1 into two $M/2$ -periodic functions. The discrete Fourier transforms of \mathbf{f}^0 and \mathbf{f}^1 can be computed from those of the corresponding pair using $2M$ multiplications each, for a total of $2 \cdot 2M = 2N$ multiplications. We now see that, if $N = 2^p$ is a power of 2, then we can continue this process of splitting discrete functions into pairs of discrete functions, halving the period at each stage, until we are left with 2^p separate 1-periodic discrete functions—the individual components of the original function \mathbf{f} . Each of these components is equal to its own discrete Fourier transform (apply the definition (8.19)!), so we can start from there, forming pairs, computing discrete Fourier transforms, and building back up to \mathbf{f} . Each stage in the process will require $2N$ multiplications, and there are p stages until we reach the discrete Fourier transform of \mathbf{f} . The total number of multiplications

required is, therefore, $2Np = 2N \log_2(N)$, which, for large values of N , compares favorably to the N^2 products required if (8.19) is used. For example, if $N = 2^{16} = 65536$, then $2N \log_2(N)$ is less than 0.05% of N^2 . Moreover, many of the multiplications involved are actually just multiplication by the factor 1; so the savings in computation time may be even greater than it initially appears. This method is called the *fast Fourier transform*.

Some care must be taken in how the individual components of \mathbf{f} are paired up if we are eventually to arrive at the functions \mathbf{f}^0 and \mathbf{f}^1 and then, finally, at \mathbf{f} . In the first stage, form separate lists of the components with even subscripts (f_0, f_2 , and so on) and those with odd subscripts (f_1, f_3 , etc.). Then pair the components in the first half of each list with those in the second half of the list. Thus, f_0 is paired with $f_{N/2}$, f_2 with $f_{N/2+2}$, and so on up to the pair $f_{N/2-2}$ and f_{N-2} . Similarly, f_1 is paired with $f_{N/2+1}$, f_3 with $f_{N/2+3}$, up to the pair $f_{N/2-1}$ and f_{N-1} . The discrete Fourier transform of each pair is computed from the transforms of its two components. For the next stage, the pairs in the even list are paired up according the same scheme: those in the first half of the list are paired with those in the second half. In the same way, the pairs in the first half of the odd list are matched up with the pairs in the second half of the odd list. This pattern then replicates from stage to stage until, in the penultimate stage, all of the components with even subscripts are combined in the same $N/2$ -periodic discrete function, and those with odd subscripts make up another discrete function. In the final stage, the evens and odds are reunited at last.

Example 8.25. Fast Fourier transform for an 8-periodic discrete function. Let \mathbf{f} be an 8-periodic discrete function. To compute the discrete Fourier transform of \mathbf{f} via the fast Fourier transform algorithm, begin with the separate components of \mathbf{f} and pair them up as (f_0, f_4) , (f_2, f_6) , (f_1, f_5) , and (f_3, f_7) . The discrete Fourier transforms of these pairs are, respectively,

$$\{f_0 + f_4, f_0 - f_4\}, \{f_2 + f_6, f_2 - f_6\}, \{f_1 + f_5, f_1 - f_5\}, \text{ and } \{f_3 + f_7, f_3 - f_7\}.$$

For the next stage, merge pairs of 2-periodic functions to form two 4-periodic functions: $\{(f_0, f_4), (f_2, f_6)\}$ and $\{(f_1, f_5), (f_3, f_7)\}$. By (8.47) and (8.48), the discrete transforms of these pairs are, respectively,

$$\begin{aligned} &\{(f_0 + f_4) + (f_2 + f_6), (f_0 - f_4) - i \cdot (f_2 - f_6), \\ &\quad (f_0 + f_4) - (f_2 + f_6), (f_0 - f_4) + i \cdot (f_2 - f_6)\} \end{aligned}$$

and

$$\begin{aligned} &\{(f_1 + f_5) + (f_3 + f_7), (f_1 - f_5) - i \cdot (f_3 - f_7), \\ &\quad (f_1 + f_5) - (f_3 + f_7), (f_1 - f_5) + i \cdot (f_3 - f_7)\}. \end{aligned}$$

Finally, combine the two 4-periodic functions into the single 8-periodic function $\mathbf{f} = \{[(f_0, f_4), (f_2, f_6)], [(f_1, f_5), (f_3, f_7)]\}$. Again following from (8.47) and (8.48), the discrete transform is

$$\begin{aligned} \mathcal{F}_D \mathbf{f} = \{ & [(f_0 + f_4) + (f_2 + f_6)] + [(f_1 + f_5) + (f_3 + f_7)], \\ & [(f_0 - f_4) - i \cdot (f_2 - f_6)] + e^{-\pi i/4} [(f_1 - f_5) - i \cdot (f_3 - f_7)], \\ & [(f_0 + f_4) - (f_2 + f_6)] - i[(f_1 + f_5) - (f_3 + f_7)], \\ & [(f_0 - f_4) + i \cdot (f_2 - f_6)] + e^{-3\pi i/4} [(f_1 - f_5) + i \cdot (f_3 - f_7)], \\ & [(f_0 + f_4) + (f_2 + f_6)] - [(f_1 + f_5) + (f_3 + f_7)], \\ & [(f_0 - f_4) - i \cdot (f_2 - f_6)] - e^{-\pi i/4} [(f_1 - f_5) - i \cdot (f_3 - f_7)], \\ & [(f_0 + f_4) - (f_2 + f_6)] + i[(f_1 + f_5) - (f_3 + f_7)], \\ & [(f_0 - f_4) + i \cdot (f_2 - f_6)] - e^{-3\pi i/4} [(f_1 - f_5) + i \cdot (f_3 - f_7)] \}. \end{aligned}$$

This agrees with (8.19), and if one counts every instance of a coefficient other than 1, the total number of multiplications involved is only about half of $8^2 = 64$ (though, admittedly, $N = 8$ is too small for there to be much of a savings here).

If this is applied to the 8-periodic discrete function with $f_k = \cos(\pi k/2)$ for $0 \leq k \leq 7$, the resulting discrete Fourier transform is

$$\mathcal{F}_D \mathbf{f} = \{0, 0, 4, 0, 0, 0, 4, 0\}.$$

Thus, the amplitude of the *cosine* wave is evenly divided between two opposite frequencies. This corresponds to the continuous setting in which the continuous Fourier transform of the *cosine* function consists of two impulses at opposite frequencies.

The fast Fourier transform has a matrix implementation as well. This amounts to factoring the matrix \mathbf{W}_N in (8.41) in the case where $N = 2^p$. As in (8.41), an N -periodic discrete function \mathbf{f} is viewed as a column vector,

$$\mathbf{f} = \begin{bmatrix} f_0 \\ f_1 \\ f_2 \\ \vdots \\ f_{N-1} \end{bmatrix}.$$

The initial step of rearranging the components of \mathbf{f} for the first round of pairing is implemented using a matrix obtained by rearranging the rows of the $N \times N$ identity matrix. Specifically, if \mathbf{e}_k denotes the k th row of the identity matrix, then let E_N denote

the matrix

$$E_N := \begin{bmatrix} \mathbf{e}_0 \\ \mathbf{e}_{N/2} \\ \mathbf{e}_2 \\ \mathbf{e}_{N/2+2} \\ \vdots \\ \mathbf{e}_{N/2-2} \\ \mathbf{e}_{N-2} \\ \mathbf{e}_1 \\ \mathbf{e}_{N/2+1} \\ \vdots \\ \mathbf{e}_{N/2-1} \\ \mathbf{e}_{N-1} \end{bmatrix}.$$

Then $E_N \mathbf{f}$ achieves the desired reordering of the components of \mathbf{f} .

As in (8.41), let $w_N = e^{-2\pi i/N}$ and, for each integer k with $1 \leq k \leq p$, let R_k be the $2^{k-1} \times 2^{k-1}$ diagonal matrix with entries $\left((w_N)^{2^{p-k}} \right)^j$ in the j th diagonal position, for $0 \leq j \leq 2^{k-1} - 1$. Let I denote the identity matrix (in this case, of dimensions $2^{k-1} \times 2^{k-1}$) and define a $2^k \times 2^k$ matrix B_k by

$$B_k := \left[\begin{array}{c|c} I & R_k \\ \hline I & -R_k \end{array} \right]. \quad (8.49)$$

Multiplication by B_k computes the discrete Fourier transform of a 2^k -periodic discrete function from the transforms of a pair of 2^{k-1} -periodic functions. Thus, defining a $2^p \times 2^p$ block-diagonal matrix T_k by arranging 2^{p-k} copies of B_k along the diagonal, it follows that multiplication by T_k computes the discrete transforms of 2^{p-k} different 2^k -periodic discrete functions. This corresponds to the k th stage of the fast Fourier transform algorithm. The end result of this algorithm is therefore described by the product

$$\mathcal{F}_D \mathbf{f} = T_p \cdot T_{p-1} \cdots T_1 \cdot E_N \mathbf{f}. \quad (8.50)$$

For example, when $N = 4$, we get

$$E_4 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad T_1 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{bmatrix},$$

and

$$T_2 = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -i \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & i \end{bmatrix}.$$

With $\mathbf{f} = \{f_0, f_1, f_2, f_3\}$, we have

$$\begin{aligned} \mathcal{F}_D \mathbf{f} &= T_2 \cdot T_1 \cdot E_4 \mathbf{f} \\ &= \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -i \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & i \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f_0 \\ f_1 \\ f_2 \\ f_3 \end{bmatrix}. \end{aligned}$$

Each of the matrices T_k has the property that there are only two nonzero entries in each row and each column. Thus, the matrix product in (8.50) requires only $2N$ multiplications for each factor. There are p factors in all, not counting E_N (which requires no multiplications), for a total of $2Np = 2N \log_2(N)$ multiplications, as we found before. Again, many of these multiplications involve the factor 1.

8.12 Fan beam geometry

In order to implement the discrete filtered back-projection formulas (8.2) and (8.38), it is necessary to know the values of $\mathcal{R}f(t, \theta)$ for a variety of choices of t for each of the selected angles θ . Conceptually, we have envisioned a scanning machine that sends out a set of parallel X-ray beams at each selected angle and records the corresponding values of the Radon transform. Such a machine would have to have a strip of distinct transmitters spaced at an appropriate sample spacing and able to rotate as a single unit during the scanning process. Each setting would correspond to a particular value of $\theta = k\pi/N$, and, once the readings had been taken, the corresponding summand in (8.38) could be calculated. This corresponds to the parallel beam geometry that has been the basis of our analysis all along.

In practice, however, it is easier to design a machine that has a single X-ray beam transmitter that emits a fan of beams. An arc of detectors on the other side measures the values of the Radon transform along the lines corresponding to the beams in the fan. The trouble is that each transmission includes beams at a variety of values of θ . Observe, though, that if two beams in the fan make an angle of ϕ with each other when they are emitted, then, when the transmitter itself has been rotated by the same angle ϕ and the two beams are emitted again, one of the new beams will be parallel to one of the beams from the previous transmission. (See Figure 8.9 for an illustration of this.) In other words, once the scanning process has been completed, it is possible to

reorganize the fan beam data into an equivalent collection of parallel beam data. The image reconstruction algorithm can then be applied to this reorganized data to produce an image. This approach is called the *fan beam geometry*.




Fig. 8.9. Different fans yield parallel beams.

Rather than go into more detail here, we leave an examination of the fan beam geometry as a topic for further study. Other scanning geometries that have been developed for use in later generation scanning machines include the spiral beam and cone beam geometries, both of which facilitate the collection of data for more than one slice at the same time. The overarching goal in the design of new geometries is to lessen the radiation exposure of the patient by increasing the efficiency of the collection of the X-ray data.

8.13 Exercises

8.1. Consider the low-pass cosine filter (7.27):

$$\begin{aligned} A(\omega) &= |\omega| \cdot \cos(\pi\omega/(2L)) \cdot \Gamma_L(\omega) \\ &= \begin{cases} |\omega| \cos(\pi\omega/(2L)) & \text{if } |\omega| \leq L, \\ 0 & \text{if } |\omega| > L. \end{cases} \end{aligned}$$

Following the format used for the Shepp–Logan and Ram–Lak filters in Section 8.3,

- (a) compute the inverse Fourier transform $\mathcal{F}^{-1}A$, and
- (b) show that

$$(\mathcal{F}^{-1}A)(\pi n/L) = \left(\frac{2L}{\pi^2}\right) \cdot \left(\frac{\pi \cdot \cos(\pi n)}{(1-4n^2)} - \frac{2 \cdot (1+4n^2)}{(1-4n^2)^2}\right)$$

for all integers n .

(c) Use a graphing device to plot the filter A and the sampled points of its inverse Fourier transform.

8.2. For discrete functions f , g , and h , and any constant α , prove that

(a) $f \bar{*} g = g \bar{*} f$,

(b) $f \bar{*} (g + h) = f \bar{*} g + f \bar{*} h$, and

(c) $f \bar{*} (\alpha g) = \alpha (f \bar{*} g)$.

(Assume that all of the sums converge.)

8.3. Prove Proposition 8.16: For two discrete N -periodic functions f and g ,

$$\mathcal{F}_D(f \cdot g) = \frac{1}{N} (\mathcal{F}_D f) \bar{*} (\mathcal{F}_D g). \quad (8.51)$$

Note that this is analogous to (7.14).

8.4. For a discrete function f , define the conjugate function \bar{f} by setting $\bar{f}_k = \overline{f_k}$ for every k . In analogy to (5.10), prove Proposition 8.17:

$$(\mathcal{F}_D \bar{f})_j = \overline{(\mathcal{F}_D f)_{-j}} \quad (8.52)$$

for all j . (Note that $(\mathcal{F}_D f)_{-j} = (\mathcal{F}_D f)_{N-j}$ by periodicity.)

8.5. Prove Proposition 8.18, the discrete Rayleigh–Plancherel theorem: For an N -periodic discrete function f ,

$$\sum_{n=0}^{N-1} |f_n|^2 = \frac{1}{N} \sum_{n=0}^{N-1} |(\mathcal{F}_D f)_n|^2. \quad (8.53)$$

8.6. Prove the identity (8.20):

$$1/2 + \cos(\theta) + \cos(2\theta) + \cdots + \cos(M\theta) = \frac{\sin((M+1/2)\theta)}{2 \cdot \sin(\theta/2)},$$

for all natural numbers M and all real θ .

8.7. On the interval $0 \leq x \leq 1$, write down a formula for the piecewise linear function F determined by the values $F(0) = 0.5$, $F(1/3) = 0.3$, $F(2/3) = 0.6$, and $F(1) = 0.5$. (The formula should consist of three line segments.) Then use this formula to compute the values $F(0.2)$ and $F(0.7)$.

8.8. Recall the formula (8.31) for Lagrange interpolation of a function f using $N+1$ data points $(x_1, f(x_1)), \dots, (x_{N+1}, f(x_{N+1}))$:

$$f(x) = \sum_{j=1}^{N+1} f(x_j) \cdot \frac{\prod_{k \neq j} (x - x_k)}{\prod_{k \neq j} (x_j - x_k)}.$$

Verify that, if $x = x_n$ for some n between 1 and $N + 1$, then the formula yields $f(x) = f(x_n)$. (Thus, the formula agrees with the data.)

Algebraic Reconstruction Techniques

9.1 Introduction

To this point, we have studied how Fourier transform methods are used in image reconstruction. This is the approach taken in the seminal work of Cormack [13] and used in the algorithms of today's CT scan machines. However, the first CT scanner, designed in the late 1960s by Godfrey Hounsfield, used an approach grounded in linear algebra and matrix theory to generate an image from the machine readings. Algorithms that adopt this point of view are known as *algebraic reconstruction techniques*, or ART, for short. In this chapter, we look at a few basic mathematical elements of ART.

Where the Fourier transform methods begin with a continuous theory—the filtered back-projection formula of Theorem 6.2—which is then modeled using discrete methods, ART treats the problem of image reconstruction as a discrete problem from the start. Any image that we produce will be constructed inside a rectangular grid of picture elements, or pixels. The number of pixels in a given image may be large, but it is nonetheless finite, typically on the order of 10^5 . For example, there are 65536 pixels in a 256-by-256 grid. To form an image, a specific color value is assigned to each pixel. For instance, the color value assigned to a given pixel might be a *greyscale* value, a number between 0 (= black) and 1 (= white), that represents the density or attenuation coefficient of the matter in the sample at the location of the given pixel. ART techniques use a system of constraints derived from the machine readings to compute these color values.

So, suppose an image is to be constructed in a K -by- K grid of pixels. Each pixel is really a small square in the plane. For convenience, number the pixels like so: 1 through K from left to right across the top row; $K + 1$ through $2K$ across the second row; and so on until, in the bottom row, we find pixels numbered $(K - 1)K + 1$ through K^2 . Next, define the *pixel basis functions* b_1, \dots, b_{K^2} by

$$b_k(x, y) = \begin{cases} 1 & \text{if } (x, y) \text{ lies inside pixel number } k, \\ 0 & \text{if } (x, y) \text{ does not lie inside pixel number } k \end{cases} \quad (9.1)$$

for $k = 1, 2, \dots, K^2$ and points (x, y) in the plane. If we assign the color value x_k to the k th pixel, then the resulting image will be represented by the function

$$\tilde{f}(x, y) = \sum_{k=1}^{K^2} x_k \cdot b_k(x, y) \quad (9.2)$$

for each point (x, y) lying inside the overall region covered by the grid. Applying the Radon transform \mathcal{R} to both sides of this equation, and using the linearity of \mathcal{R} , we get, for each choice of t and θ ,

$$\mathcal{R}\tilde{f}(t, \theta) = \sum_{k=1}^{K^2} x_k \cdot \mathcal{R}b_k(t, \theta). \quad (9.3)$$

In practice, the X-ray machine gives us the values of $\mathcal{R}\tilde{f}(t, \theta)$ for some finite set of lines $\ell_{t, \theta}$. For convenience, let's say these known values correspond to (t_1, θ_1) , (t_2, θ_2) , \dots , (t_J, θ_J) for some positive integer J . Then,

$$\text{let } p_j = \mathcal{R}\tilde{f}(t_j, \theta_j), \text{ for } j = 1, 2, \dots, J.$$

The system of equations (9.3) can now be written as

$$p_j = \sum_{k=1}^{K^2} x_k \cdot \mathcal{R}b_k(t_j, \theta_j) \text{ for } j = 1, \dots, J. \quad (9.4)$$

Our next observation is that, since the pixel basis function b_k has the value 1 on its pixel and 0 elsewhere, the value of the integral $\mathcal{R}b_k(t_j, \theta_j)$ is equal to the *length of the intersection* of the line ℓ_{t_j, θ_j} with pixel number k . In principle, these values are easy to compute. (*Caveat:* If we allow finite-width X-ray beams, rather than zero-width, then this computation becomes more complicated.) So, let's denote by r_{jk} the length of the intersection of the line ℓ_{t_j, θ_j} with pixel number k ; that is,

$$\text{let } r_{jk} = \mathcal{R}b_k(t_j, \theta_j) \text{ for } j = 1, \dots, J \text{ and } k = 1, \dots, K^2. \quad (9.5)$$

With this notation, the system (9.3) can be written as

$$p_j = \sum_{k=1}^{K^2} x_k \cdot r_{jk} \text{ for } j = 1, \dots, J. \quad (9.6)$$

This is a system of J linear equations in K^2 unknowns (x_1, \dots, x_{K^2}) . Typically, both J and K^2 are on the order of 10^5 , so the system is large. However, any particular line ℓ_{t_j, θ_j} passes through relatively few of the pixels in the grid, on the order of K out of K^2 total pixels. Thus, most of the values r_{jk} are equal to 0, meaning that the system (9.6) is large but *sparse*. With typical values of J and K , only about one percent of the entries in the coefficient matrix of the system are nonzero.

Some computational concerns arise from this approach to the image reconstruction problem. For one thing, the system of equations we have to solve is *large*—typically on the order of 10^5 equations. Each sampling of the Radon transform produces an equation in the system while each pixel corresponds to an unknown, the color value for that pixel. If the system of equations is overdetermined, with more equations than unknowns, then the system likely does not have an exact solution. If the system is underdetermined, with more unknowns than equations, then there may be infinitely many solutions, only one of which could possibly be the correct solution. A typical scan might include 200 X-ray measurements at each of 180 different directions, for a total of 36000 equations in the system. A grid of 160×160 pixels gives 25600 unknowns and an overdetermined system. To get an image with higher resolution, though, we may want to use a grid of 256×256 pixels, or 65536 unknowns. This results in a system that is heavily underdetermined, so the iterative algorithms discussed below are ineffective. For this reason, among others, iterative algorithms are no longer used in commercial CT machines. In any case, due to errors in measurement in sampling the Radon transform, the equations are only estimates to begin with. So, again, the system is not likely to have an exact solution. The fact that the coefficient matrix is *sparse*, with only a small proportion of nonzero entries, also has a direct effect on the computational complexity.

We now look at two different methods for arriving at an approximate solution to a system of linear equations. These methods are least squares approximation and Kaczmarz's method.

9.2 Least squares approximation

For a given $M \times N$ matrix A and a given vector \mathbf{p} in \mathbf{R}^M , the system $A\mathbf{x} = \mathbf{p}$ may not have a solution. One approach to finding an approximate solution is to find the vector $\hat{\mathbf{y}}$ of the form $\hat{\mathbf{y}} = A\hat{\mathbf{x}}$ such that $\hat{\mathbf{y}}$ is *closest* to \mathbf{p} amongst all vectors of this form. Thus, the goal of least squares approximation is to find a vector $\hat{\mathbf{x}}$ in \mathbf{R}^N such that

$$\|A\hat{\mathbf{x}} - \mathbf{p}\| = \min_{\mathbf{x} \in \mathbf{R}^N} \|A\mathbf{x} - \mathbf{p}\|. \quad (9.7)$$

In other words, we wish to find the vector $A\hat{\mathbf{x}}$ in the range of the matrix A that is closest to \mathbf{p} amongst all elements of the range of A . Because the magnitude of a vector can

be expressed as a sum of squares, this form of approximation is called *least squares approximation*.

Now for the technicalities.

Definition 9.1. For a given $M \times N$ matrix A , the *range* of A is the set $\{A\mathbf{x} \mid \mathbf{x} \in \mathbf{R}^N\}$. Notice that this is a subset of \mathbf{R}^M .

Theorem 9.2. For a given $M \times N$ matrix A , the range of A is a subspace of \mathbf{R}^M .

Proof. If \mathbf{x}_1 and \mathbf{x}_2 are any elements of \mathbf{R}^N and λ is any constant (scalar), then

$$\lambda A\mathbf{x}_1 + A\mathbf{x}_2 = A(\lambda\mathbf{x}_1 + \mathbf{x}_2)$$

by the properties of matrix arithmetic. Thus, the range of A is closed under the operations of vector addition and scalar multiplication. \square

Definition 9.3. For a given $M \times N$ matrix A , the *transpose* of A , denoted by A^T , is the $N \times M$ matrix whose entry in row k and column j is the same as the entry of A in row j and column k . In other words, column j of A^T is the same as row j of A , and row k of A^T is the same as column k of A . Notice that, in general, $(A^T)^T = A$.

Example 9.4. If $A = \begin{bmatrix} 4 & 0 \\ -3 & 2 \\ 1 & 5 \end{bmatrix}$, then $A^T = \begin{bmatrix} 4 & -3 & 1 \\ 0 & 2 & 5 \end{bmatrix}$.

The interaction between A , A^T , and the dot product, stated in the next theorem, is of fundamental importance in many applications of linear algebra.

Theorem 9.5. Let A be an $M \times N$ matrix, let \mathbf{x} be a vector in \mathbf{R}^N , and let \mathbf{y} be a vector in \mathbf{R}^M . Then

$$A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A^T \mathbf{y}. \quad (9.8)$$

Notice that the dot product on the left-hand side involves two vectors in \mathbf{R}^M while the dot product on the right-hand side involves two vectors in \mathbf{R}^N .

Proof. Let us consider an example to illustrate. Take $A = \begin{bmatrix} 4 & 0 \\ -3 & 2 \\ 1 & 5 \end{bmatrix}$, $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ in

\mathbf{R}^2 , and $\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$ in \mathbf{R}^3 . Then we compute:

$$\begin{aligned}
\mathbf{Ax} \cdot \mathbf{y} &= \begin{bmatrix} 4 & 0 \\ -3 & 2 \\ 1 & 5 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \cdot \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} \\
&= \begin{bmatrix} 4x_1 \\ -3x_1 + 2x_2 \\ x_1 + 5x_2 \end{bmatrix} \cdot \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} \\
&= 4x_1y_1 + (-3x_1 + 2x_2)y_2 + (x_1 + 5x_2)y_3 \\
&= x_1(4y_1 - 3y_2 + y_3) + x_2(2y_2 + 5y_3) \\
&= \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \cdot \begin{bmatrix} 4y_1 - 3y_2 + y_3 \\ 2y_2 + 5y_3 \end{bmatrix} \\
&= \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \cdot \begin{bmatrix} 4 & -3 & 1 \\ 0 & 2 & 5 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} \\
&= \mathbf{x} \cdot \mathbf{A}^T \mathbf{y} \text{ as desired.}
\end{aligned}$$

The general proof follows this same pattern and is left as an exercise. \square

Corollary 9.6. *If \mathbf{y} is orthogonal to the range of A , then \mathbf{y} is in the nullspace of A^T .*

Proof. If \mathbf{y} is orthogonal to the range of A , then $\mathbf{Ax} \cdot \mathbf{y} = 0$ for every vector \mathbf{x} in \mathbf{R}^N . Hence, it follows from Theorem 9.5 that $\mathbf{x} \cdot \mathbf{A}^T \mathbf{y} = 0$ for every \mathbf{x} in \mathbf{R}^N . In other words, the vector $\mathbf{A}^T \mathbf{y}$ is orthogonal to *every* vector in \mathbf{R}^N . In particular, the vector $\mathbf{A}^T \mathbf{y}$ is *orthogonal to itself*, so that $\mathbf{A}^T \mathbf{y} \cdot \mathbf{A}^T \mathbf{y} = 0$. That is, $\|\mathbf{A}^T \mathbf{y}\|^2 = 0$, from which it follows that $\mathbf{A}^T \mathbf{y} = \mathbf{0}$. In other words, \mathbf{y} is in the nullspace of \mathbf{A}^T as claimed. \square

Now we return to the least squares approximation problem itself. Recall that our goal is to find the vector $\widehat{\mathbf{Ax}}$ in the range of the matrix A that is closest to \mathbf{p} of all such vectors. Geometrically, if we take an arbitrary element \mathbf{Ax} in the range of A and consider the projection of the vector $(\mathbf{Ax} - \mathbf{p})$ onto the range of A , then the foot of this projection will be a vector in the range of A that is at least as close to \mathbf{p} as the original vector \mathbf{Ax} was. Thus, the closest element to \mathbf{p} in the range of A is the element $\widehat{\mathbf{Ax}}$ for which the foot of the projection of $(\widehat{\mathbf{Ax}} - \mathbf{p})$ onto the range of A is the vector $\widehat{\mathbf{Ax}}$ itself. This means that the vector $(\widehat{\mathbf{Ax}} - \mathbf{p})$ must be orthogonal to the range of A . By Corollary 9.6 above, this means that the vector $(\widehat{\mathbf{Ax}} - \mathbf{p})$ must lie in the nullspace of the matrix \mathbf{A}^T . In other words, the vector $\widehat{\mathbf{x}}$ must satisfy the equation

$$\mathbf{A}^T (\widehat{\mathbf{Ax}} - \mathbf{p}) = \mathbf{0}, \quad (9.9)$$

or what is the same thing,

$$\mathbf{A}^T \widehat{\mathbf{Ax}} = \mathbf{A}^T \mathbf{p}. \quad (9.10)$$

This last equation is called the *normal equation*.

Definition 9.7. A *least squares solution* to the equation $A\mathbf{x} = \mathbf{p}$ is a vector $\hat{\mathbf{x}}$ that satisfies the normal equation (9.10). In this case, the vector $A\hat{\mathbf{x}}$ is the closest element in the range of A to the vector \mathbf{p} .

Notice that, in the case where the matrix $A^T A$ is invertible, then the normal equation (9.10) yields the least squares solution

$$\hat{\mathbf{x}} = (A^T A)^{-1} A^T \mathbf{p}. \quad (9.11)$$

In general, however, $A^T A$ need not be invertible, in which case the least squares solution $\hat{\mathbf{x}}$ is not unique. Nonetheless, the range element $\hat{\mathbf{y}} = A\hat{\mathbf{x}}$ that is closest to \mathbf{p} is unique.

Example 9.8. Take $A = \begin{bmatrix} -1 & 2 \\ 2 & -3 \\ -1 & 3 \end{bmatrix}$ and $\mathbf{p} = \begin{bmatrix} 4 \\ 1 \\ 2 \end{bmatrix}$. Show that $A^T A$ is invertible and

that the (unique) least squares solution, computed from (9.11), is $\hat{\mathbf{x}} = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$. Compute $\|A\hat{\mathbf{x}} - \mathbf{p}\|$ in this case.

Example 9.9. Take $A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 0 & 1 \end{bmatrix}$ and $\mathbf{p} = \begin{bmatrix} 1 \\ 3 \\ 8 \\ 2 \end{bmatrix}$. Show that $A^T A$ is not invertible.

Find all solutions to the normal equation (9.10) and show that the (unique) element

closest to \mathbf{p} in the range of A is $\hat{\mathbf{y}} = \begin{bmatrix} 2 \\ 2 \\ 5 \\ 5 \end{bmatrix}$. Compute $\|\hat{\mathbf{y}} - \mathbf{p}\|$ in this case.

Here is an alternate approach to solving the least squares approximation problem that leads us again to the normal equation (9.10). For a given $M \times N$ matrix A and a given vector \mathbf{p} in \mathbf{R}^M , define a real-number-valued function $F : \mathbf{R}^N \rightarrow \mathbf{R}$ by

$$F(\mathbf{x}) := \|A\mathbf{x} - \mathbf{p}\|^2 = (A\mathbf{x} - \mathbf{p}) \cdot (A\mathbf{x} - \mathbf{p}), \text{ for all } \mathbf{x} \text{ in } \mathbf{R}^N.$$

One can show that the *gradient vector* of F satisfies

$$\nabla F(\mathbf{x}) = 2(A^T A\mathbf{x} - A^T \mathbf{p}), \text{ for all } \mathbf{x} \text{ in } \mathbf{R}^N. \quad (9.12)$$

The least squares solution $\hat{\mathbf{x}}$ produces a minimum value for F and, hence, the gradient of F at $\hat{\mathbf{x}}$ must vanish. That is, $\nabla F(\hat{\mathbf{x}}) = \mathbf{0}$ for the least squares solution $\hat{\mathbf{x}}$. According to (9.12), then, the least squares solution $\hat{\mathbf{x}}$ satisfies the normal equation $A^T A\hat{\mathbf{x}} = A^T \mathbf{p}$.

One computational concern associated with the least squares method in the context of imaging is that the matrix $A^T A$ in the normal equation (9.10) is huge, on the order

of K^2 -by- K^2 . Also, the matrix $A^T A$ might not be sparse even though A is. Another concern is that the matrix $A^T A$ might not be invertible or that its inverse might be difficult to compute. Moreover, round-off errors can be fatal. For instance, a matrix entry that ought to be 0 but shows up as a tiny nonzero number can wreck the process of inverting a matrix. Similarly, an entry that ought to be a tiny nonzero number but shows up as 0 can also have a deleterious effect.

9.3 Kaczmarz's method

Kaczmarz's method is an iterative procedure, or algorithm, for approximating a solution to a linear system $A\mathbf{x} = \mathbf{p}$. If we denote by \mathbf{r}_i the i th row of the matrix A and by p_i the i th coordinate of the vector \mathbf{p} , then the system $A\mathbf{x} = \mathbf{p}$ is the same as having $\mathbf{r}_i \cdot \mathbf{x} = p_i$ for every value of i . Kaczmarz's method works by producing a sequence of vectors each of which satisfies one of the individual equations $\mathbf{r}_i \cdot \mathbf{x} = p_i$. The first research article to explore the application of algebraic reconstruction techniques to medical imaging was [19], which appeared in 1970. The principal technique employed therein for the reconstruction of images turned out to be the same as Kaczmarz's method.

9.3.1 Affine spaces

Before looking at Kaczmarz's method itself, we make the following definition.

Definition 9.10. For a fixed n -dimensional vector \mathbf{r} and a number p , the *affine space* $\mathcal{S}_{\mathbf{r},p}$ is defined by

$$\mathcal{S}_{\mathbf{r},p} = \{\mathbf{x} \in \mathbf{R}^n : \mathbf{r} \cdot \mathbf{x} = p\}.$$

Note that the affine space $\mathcal{S}_{\mathbf{r},p}$ is a subspace of \mathbf{R}^n if, and only if, $p = 0$.

We have already seen an important example of affine spaces in the lines $\ell_{t,\theta}$. Each point on $\ell_{t,\theta}$ is the terminal point of a vector $\mathbf{x} = \langle t \cos(\theta) - s \sin(\theta), t \sin(\theta) + s \cos(\theta) \rangle$ for some value of the parameter s . Letting $\mathbf{r} = \langle \cos(\theta), \sin(\theta) \rangle$, we compute

$$\begin{aligned} \mathbf{r} \cdot \mathbf{x} &= t \cos^2(\theta) - s \sin(\theta) \cos(\theta) + t \sin^2(\theta) + s \cos(\theta) \sin(\theta) \\ &= t(\cos^2(\theta) + \sin^2(\theta)) \\ &= t \end{aligned}$$

regardless of the value of s . Thus, each line $\ell_{t,\theta}$ is an affine space. If $t = 0$, then $\ell_{0,\theta}$ is a line through the origin and, so, is a subspace of \mathbf{R}^2 .

Every affine space can be viewed as a copy of a subspace that has been shifted by a fixed vector. For example, for each fixed value of θ , the line

$$\ell_{0,\theta} = \{ \langle -s \sin(\theta), s \cos(\theta) \rangle : -\infty < s < \infty \}$$

forms a subspace of \mathbf{R}^2 . For each real number t , we then have that

$$\ell_{t,\theta} = \langle t \cos(\theta), t \sin(\theta) \rangle + \ell_{0,\theta}.$$

To formulate the same notion slightly differently, let $\mathbf{r} = \langle \cos(\theta), \sin(\theta) \rangle$ and $\mathbf{x}_0 = \langle t \cos(\theta), t \sin(\theta) \rangle$. Then $\mathbf{r} \cdot \mathbf{x}_0 = t$. The subspace $\ell_{0,\theta}$ consists of the set of all vectors \mathbf{x} for which $\mathbf{r} \cdot \mathbf{x} = 0$. Thus, the line $\ell_{t,\theta}$ consists of the terminal points of all vectors of the form $\mathbf{x}_0 + \mathbf{x}$, where $\mathbf{r} \cdot \mathbf{x} = 0$.

Similarly, for any vector \mathbf{r} and any real number p , consider the affine space $\mathcal{S}_{\mathbf{r},p}$ and the subspace $\mathcal{S}_{\mathbf{r},0}$. Observe that, for \mathbf{x}_0 and \mathbf{x}_1 in $\mathcal{S}_{\mathbf{r},p}$, the vector $\mathbf{x}_h = \mathbf{x}_1 - \mathbf{x}_0$ satisfies the homogeneous equation $\mathbf{r} \cdot \mathbf{x}_h = 0$. That is, \mathbf{x}_h is in the subspace $\mathcal{S}_{\mathbf{r},0}$. Since $\mathbf{x}_1 = \mathbf{x}_h + \mathbf{x}_0$, it follows that every element of the affine space $\mathcal{S}_{\mathbf{r},p}$ can be obtained by adding the vector \mathbf{x}_0 to some element of the subspace $\mathcal{S}_{\mathbf{r},0}$. This is reminiscent of the use of a “particular solution” together with the general homogeneous solution to obtain the general solution to a nonhomogeneous linear differential equation or a nonhomogeneous system of linear equations.

A final crucial observation is that, since the vector \mathbf{r} is orthogonal to the subspace $\mathcal{S}_{\mathbf{r},0}$, and, since the affine space $\mathcal{S}_{\mathbf{r},p}$ is a parallel translation of this subspace, then it follows that the vector \mathbf{r} itself is orthogonal to the affine space $\mathcal{S}_{\mathbf{r},p}$.

Definition 9.11. Affine projection. Given a vector \mathbf{u} and an affine space $\mathcal{S}_{\mathbf{r},p}$ for some vector \mathbf{r} and some number p , the *affine projection* of \mathbf{u} in $\mathcal{S}_{\mathbf{r},p}$ is the vector \mathbf{u}_* in $\mathcal{S}_{\mathbf{r},p}$ that is *closest* to \mathbf{u} amongst all vectors in $\mathcal{S}_{\mathbf{r},p}$.

Now, in order to move from \mathbf{u} to the closest point in the affine space, it is evident that we should move *orthogonally* to the affine space. According to our previous observations, this means that we should move in the direction of the vector \mathbf{r} itself. Thus, the vector \mathbf{u}_* that we seek should have the form $\mathbf{u}_* = \mathbf{u} - \lambda \mathbf{r}$ for some number λ .

Substituting $\mathbf{u}_* = \mathbf{u} - \lambda \mathbf{r}$ into the equation $\mathbf{r} \cdot \mathbf{u}_* = p$ and solving for λ yields

$$\lambda = \frac{(\mathbf{r} \cdot \mathbf{u}) - p}{\mathbf{r} \cdot \mathbf{r}}.$$

Thus, we have proven the following proposition.

Proposition 9.12. *The affine projection \mathbf{u}_* of the vector \mathbf{u} in the affine space $\mathcal{S}_{\mathbf{r},p}$ is given by*

$$\mathbf{u}_* = \mathbf{u} - \left(\frac{(\mathbf{r} \cdot \mathbf{u}) - p}{\mathbf{r} \cdot \mathbf{r}} \right) \mathbf{r}. \quad (9.13)$$

9.3.2 Kaczmarz's method

Now we put our knowledge of affine spaces to work. Recall that our goal is to find an approximate solution to a linear system $A\mathbf{x} = \mathbf{p}$. Again denote the i th row of A by \mathbf{r}_i and the i th coordinate of \mathbf{p} by p_i . Then each of the equations $\mathbf{r}_i \cdot \mathbf{x} = p_i$ describes an affine space. Kaczmarz's method proceeds by starting with an initial guess at a solution, a vector of prospective color values, and then computing the affine projection of this initial guess onto the first affine space in our list. This projection is then projected onto the next affine space in the list, and so on until we have gone through the entire list of affine spaces. This constitutes one "iteration" and the result of this iteration becomes the starting point for the next iteration.

In detail, the method proceeds as follows.

- (i) Select a starting "guess" for \mathbf{x} ; call it \mathbf{x}^0 .
- (ii) Next set $\mathbf{x}^{0,0} = \mathbf{x}^0$.
- (iii) The inductive step is this: Once the vector $\mathbf{x}^{0,i-1}$ has been determined, define

$$\mathbf{x}^{0,i} = \mathbf{x}^{0,i-1} - \left(\frac{\mathbf{x}^{0,i-1} \cdot \mathbf{r}_i - p_i}{\mathbf{r}_i \cdot \mathbf{r}_i} \right) \mathbf{r}_i. \quad (9.14)$$

We have used the affine projection formula (9.13) from Proposition 9.12.

- (iv) Note that if the matrix A has J rows, then the vectors $\mathbf{x}^{0,1}, \mathbf{x}^{0,2}, \dots, \mathbf{x}^{0,J}$ will be computed.
- (v) Once $\mathbf{x}^{0,J}$ has been computed, define $\mathbf{x}^1 = \mathbf{x}^{0,J}$ and begin the process again starting with \mathbf{x}^1 . That is, now set $\mathbf{x}^{1,0} = \mathbf{x}^1$ and compute the vectors $\mathbf{x}^{1,1}, \mathbf{x}^{1,2}, \dots, \mathbf{x}^{1,J}$, as in (9.14).
- (vi) Then let $\mathbf{x}^2 = \mathbf{x}^{1,J}$ and repeat the process starting with $\mathbf{x}^{2,0} = \mathbf{x}^2$.
- (vii) Stop when we've had enough!

There are, as we might expect, some computational concerns with this method. In principle, the successive vectors $\mathbf{x}^0, \mathbf{x}^1, \mathbf{x}^2, \dots$ should get closer to a vector that satisfies the original system $A\mathbf{x} = \mathbf{p}$. However, the convergence may be quite slow, meaning that many steps of the iteration would have to be applied to get a good approximant. Also, if the system has no solution, then the vectors computed from the algorithm might settle into a specific pattern, called an "attractor" in dynamical systems theory, or might even exhibit "chaotic" behavior.

Example 9.13. Apply Kaczmarz's method to the system consisting of just the two lines $x + 2y = 5$ and $x - y = 1$. So $\mathbf{r}_1 = \langle 1, 2 \rangle$, $\mathbf{r}_2 = \langle 1, -1 \rangle$, $p_1 = 5$, and $p_2 = 1$. With the initial guess $\mathbf{x}^0 = \langle 0.5, 0.5 \rangle$, the diagram on the left in Figure 9.1 shows that the solution $\langle 7/3, 4/3 \rangle$ is quickly found.




Fig. 9.1. On the left, Kaczmarz's method converges quickly to the point of intersection of two lines. On the right, the successive iterations settle into a triangle pattern for a system of three lines.

However, when we include the third line $4x + y = 6$ (so $\mathbf{r}_3 = \langle 4, 1 \rangle$ and $p_3 = 6$), then the successive iterations settle into a triangle pattern, shown in the diagram on the right in the figure.

Before looking at an image created using Kaczmarz's method, we state, without proof, the main convergence theorem for this algorithm. (For a proof see [14] or [30].)

Theorem 9.14. *If the linear system $A\mathbf{x} = \mathbf{p}$ has at least one solution, then Kaczmarz's method converges to a solution of this system. Moreover, if \mathbf{x}^0 is in the range of A^T , then Kaczmarz's method converges to the solution of minimum norm.*

9.4 ART in medical imaging

The convergence result in Theorem 9.14 generally is not relevant in medical imaging, where the linear systems encountered tend to be indeterminate and where, in any case, it is computationally feasible to run only a few iterations of Kaczmarz's method.

The image in Figure 9.2 is a reconstruction of the crescent-shaped phantom illustrated in Figure 2.2 and defined by the attenuation function in (2.8). The image here is divided into a 25×25 grid of "pixels," so there are 625 pixel basis functions. The Radon transform of the phantom was sampled with a sample spacing of 0.05 on the interval $-0.6 \leq t \leq 0.6$ and with angle increments of $\pi/18$, for a total of 432 different values. The resulting system, as in (9.6), has 432 equations in 625 unknowns. Five full iterations of Kaczmarz's method were applied with an initial vector \mathbf{x}^0 having every coordinate equal to 0.5. (In other words, the starting point was taken to be a neutral grey image.) This image shows the crescent, but it does not compare particularly favorably to the images in Figures 8.7 and 8.8, even though the computer processing time required was greater.




Fig. 9.2. Kaczmarz's method applied to a crescent-shaped phantom, using a 25×25 image grid and 432 values of the Radon transform.

The image in Figure 9.2 raises some concerns about the use of Kaczmarz's method in medical imaging. As we have mentioned before, the sheer size of the system of equations involved can be an impediment. In the context of tomography, each of the values r_{ij} represents the length of the intersection of one of the X-ray beams with one of the pixel squares in the image grid. Thus, as observed before, most of these values are zeros. The formula for computing $\mathbf{x}^{k,i}$ from $\mathbf{x}^{k,i-1}$ only alters $\mathbf{x}^{k,i-1}$ in those components that correspond to the pixels through which the i th beam passes. To streamline the computation of $\mathbf{x}^{k,i-1} \cdot \mathbf{r}_i$, we could store the locations of the nonzero entries of \mathbf{r}_i .

Another concern is that of the rate of convergence. Adjacent X-ray beams, transmitted along the lines $\ell_{t,\theta}$ for nearby values of t and θ , will intersect many of the same pixels. Thus, the corresponding affine spaces $\{\mathbf{x} \cdot \mathbf{r} = p\}$ will be almost parallel. This might tend to slow down the convergence of the algorithm and increase the number of iterations required to get an acceptable image. One way to address this is to use more and smaller pixels, though this adds to the processing time.

9.5 Variations of Kaczmarz's method

Perhaps the most commonly employed variation of Kaczmarz's method involves the introduction of so-called *relaxation parameters* in the crucial step (9.14) of the algorithm. Specifically, for each i and k , let λ_{ik} satisfy $0 < \lambda_{ik} < 2$ and replace the formula in (9.14) with the formula

$$\mathbf{x}^{k,i} = \mathbf{x}^{k,i-1} - \lambda_{ik} \cdot \left(\frac{\mathbf{x}^{k,i-1} \cdot \mathbf{r}_i - p_i}{\mathbf{r}_i \cdot \mathbf{r}_i} \right) \mathbf{r}_i. \quad (9.15)$$

For $\lambda_{ik} = 1$ this is the same as before. For $0 < \lambda_{ik} < 1$, the vector $\mathbf{x}^{k,i-1}$ is projected only part of the way to the affine space $\mathcal{S}_{\mathbf{r}_i, p_i}$. When $1 < \lambda_{ik} < 2$, the vector $\mathbf{x}^{k,i-1}$ is projected to the other side of $\mathcal{S}_{\mathbf{r}_i, p_i}$. Note that, if $\lambda_{ik} = 2$, then the vector $\mathbf{x}^{k,i}$ is just the reflection of $\mathbf{x}^{k,i-1}$ across $\mathcal{S}_{\mathbf{r}_i, p_i}$ and there is no improvement in the proximity to a solution. This is the reason for the restriction $\lambda_{ik} < 2$. In fact, the usual requirement is that the value of λ_{ik} be bounded away from 0 and 2; that is, there should be numbers α and β such that $0 < \alpha \leq \lambda_{ik} \leq \beta < 2$ for all i and k .

The additional control over the projection offered by the relaxation parameters can be used to facilitate finding an acceptable approximate solution to an indeterminate system. Alternatively, we can use this additional control to modify the original system of linear equations by replacing it with a system of inequalities instead. To do this, we select, for each i , a (small) positive number ε_i and consider the inequalities

$$p_i - \varepsilon_i \leq \mathbf{r}_i \cdot \mathbf{x} \leq p_i + \varepsilon_i.$$

A solution to this system of inequalities is a vector \mathbf{x}^* that, instead of lying in the intersection of some collection of affine spaces, lies in close proximity to those spaces. Geometrically, if we think of an affine space as a higher-dimensional plane sitting inside a many-dimensional space, then vectors in proximity to an affine space form a higher-dimensional “slab,” with thickness $2 \cdot \varepsilon_i$. The solution vector \mathbf{x}^* would lie inside these slabs. The use of relaxation parameters in the variation of Kaczmarz’s method enables us to control the projection of each successive vector in the iteration into the next slab. When inequalities are used instead of equations, the problem is called a *feasibility problem* rather than an optimization problem.

The article [11] provides a nice introduction to these ideas, while [10] and [19] offer more details of these methods. A much more general approach to feasibility problems that includes Kaczmarz’s method and its variants can be found in the comprehensive article [3].

9.6 ART or the Fourier transform?

As mentioned at the start of this chapter, the first CT scanner, invented at EMI by Hounsfield, essentially used an ART approach for its images. However, Fourier transform methods, such as the filtered back-projection formula, are generally faster to implement on a computer. Consequently, today’s commercial scanners are programmed to use transform methods. The iterative algorithms of ART simply converge too slowly, while the filtered back projection, which is based on a continuous model, can be adapted fairly easily to any desired level of accuracy.

It is nonetheless worth studying ART, and not only for its intrinsic mathematical interest. For instance, in some nonmedical applications of CT, such as nondestructive

material testing, abrupt changes in the density of the material being scanned require image reconstruction methods that can provide high contrast. ART turns out to be useful in this regard. Also, ART can have a role to play in single-photon emission computerized tomography (SPECT) and positron emission tomography (PET), where difficulty in the measurement of the attenuation can sometimes render transform methods less reliable than usual. Finally, we mention the problem of incomplete data collection, which can occur, for instance, if the range of angles used in a CT scan is restricted in order to limit the patient's exposure to X-rays. Transform methods that rely on convolution require the completion of the data, whereas the iterative ART methods simply get applied to a smaller set of equations.

9.7 Exercises

9.1. In each case, find all least squares solutions to the system $A\mathbf{x} = \mathbf{b}$.

$$(a) \quad A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}; \quad \mathbf{b} = \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}.$$

$$(b) \quad A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix}; \quad \mathbf{b} = \begin{bmatrix} -3 \\ -1 \\ 0 \\ 2 \\ 5 \\ 1 \end{bmatrix}.$$

Warning: $A^T A$ is not invertible!

9.2. For the system of two lines $x_1 - x_2 = 0$ and $x_1 + x_2 = 5$ and the starting point $\mathbf{x}^0 = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$, apply Kaczmarz's method to compute $\mathbf{x}^{0,1}$ and $\mathbf{x}^{0,2}$. Show that the vector $\mathbf{x}^{0,2}$ lies on both lines.

9.3. For the system of three lines $x_1 - x_2 = 1$, $x_2 = 1$, and $x_1 = 0$ and the starting point $\mathbf{x}^0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, apply Kaczmarz's method to compute \mathbf{x}^1 and \mathbf{x}^2 . (That is, apply two full cycles of the iteration.) What happens? What happens if we start instead at $\mathbf{x}^0 = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}$?

9.4. Prove Theorem 9.5: Let A be an $M \times N$ matrix, let \mathbf{x} be a vector in \mathbf{R}^N , and let \mathbf{y} be a vector in \mathbf{R}^M . Then

$$A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A^T \mathbf{y}.$$

9.5. Let A be an $M \times N$ matrix. Prove that, if \mathbf{y} is in the nullspace of A^T , then \mathbf{y} is orthogonal to the range of A . In conjunction with the Corollary 9.6, this shows that the orthogonal complement of the range of A coincides with the nullspace of A^T .

9.6. For a given $M \times N$ matrix A and a given vector \mathbf{p} in \mathbf{R}^M , let the function $F : \mathbf{R}^N \rightarrow \mathbf{R}$ be defined by

$$F(\mathbf{x}) := \|\mathbf{Ax} - \mathbf{p}\|^2 = (\mathbf{Ax} - \mathbf{p}) \cdot (\mathbf{Ax} - \mathbf{p}), \text{ for all } \mathbf{x} \text{ in } \mathbf{R}^N.$$

Show that the *gradient vector* of F satisfies

$$\nabla F(\mathbf{x}) = 2(A^T \mathbf{Ax} - A^T \mathbf{p}), \text{ for all } \mathbf{x} \text{ in } \mathbf{R}^N,$$

as claimed in (9.12).

9.7. Let $A = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$ and $\mathbf{p} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$. Define F by

$$F(x, y) = \left\| A \begin{bmatrix} x \\ y \end{bmatrix} - \mathbf{p} \right\|^2.$$

By solving the equation $\nabla F(x, y) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, show that the vector \mathbf{q} in the range of A that

is closest to \mathbf{p} is $\mathbf{q} = \begin{bmatrix} 1.5 \\ 1.5 \end{bmatrix}$.

MRI—An Overview

10.1 Introduction

Magnetic resonance imaging, or MRI, is an imaging technique that has grown alongside CT and that, like CT, has produced Nobel laureates of its own. Where the physics of CT is fairly straightforward—X-rays are emitted and their changes in intensity measured—MRI is based on the generation of a complex of overlapping, fluctuating electromagnetic fields that must be precisely controlled. Mathematically, the effects of the electromagnetic fields on the atomic nuclei in the sample being studied are modeled with differential equations. The Fourier transform is the primary tool for analyzing the electrical signals generated by the motions of atomic nuclei under the influence of these fields.

Clinically, MRI is safer than CT for most patients since it involves no radiation. The magnetic fields involved operate at frequencies in the radio band range. (In fact, to the patient undergoing an MRI exam, it sounds like listening to a very loud, very weird radio station.) In order to emphasize the safety and to discourage confusion, the original appellation of *nuclear magnetic resonance imaging* (nMRI) was shortened. On the downside, an MRI machine is expensive to purchase, operate, and maintain. Also, the intensity of the magnetic fields can rule out the procedure for some patients, including those with certain metallic implants.

Magnetic resonance imaging is a wide-ranging and continually developing field of study and practice, and is the subject of an extensive body of literature. Consequently, in this chapter we present only a brief overview of some of the basic principles involved in MRI, emphasizing aspects of the underlying mathematics. For a reader wishing to undertake a more intensive investigation of MRI, some possible starting points are the article [21] and the books [17], [5], [16], and [25].

Two basic descriptions of the phenomenon known as nuclear magnetic resonance (NMR) were published in 1946, one by a team of researchers led by Felix Bloch

(1905–1983) and the other by a team headed by Edward Purcell (1912–1997). Bloch's point of view is based on principles from classical physics and adopts an aggregate approach, looking at the net magnetization of the nuclei in a sample. Purcell's description, grounded in quantum physics, examines the magnetic effects at the level of an individual nucleus. It is perhaps ironic that Bloch was trained as a quantum physicist (his doctoral advisor was Werner Heisenberg), while Purcell was a classical physicist (his doctoral advisor was John van Vleck). In 1952, Bloch and Purcell were joint recipients of the Nobel Prize for Physics.

For several decades after Bloch and Purcell established the physical basis for studying NMR, the primary application was to chemical spectroscopy, and it was only around 1970 that the possibility of using NMR for imaging was realized. Paul Lauterbur (1929–2007) is credited with introducing the idea of using gradient magnetic fields to achieve spatial resolution of the radio signal emitted by a magnetized sample. Applying his technology to a setup consisting of test tubes of heavy water sitting inside a beaker of regular water, he produced the first images that could distinguish between two different kinds of water. Peter Mansfield (1933–) advanced Lauterbur's work by developing techniques for mathematically analyzing the radio signals, including a technique known as echo-planar imaging that speeds up the imaging process. Lauterbur and Mansfield were jointly awarded the 2003 Nobel Prize for Medicine and Physiology.

10.2 Basics

The nucleus of a hydrogen atom, a chemical element found in abundance in the human body, possesses a property known as spin. Conceptually, one can think of the single proton that comprises this nucleus as a tiny spinning top, rotating about an axis. This property brings with it a magnetic effect, whereby the spinning proton behaves like a bar magnet with north and south poles. As the little magnet spins, it generates an electrical signal. In the absence of other factors, there is no preferred choice for the axis around which the proton spins nor for the orientation of this axis within three-dimensional space. Within a sample of hydrogen-rich tissue, then, the distribution of spins will be random and the resulting signals will cancel each other out.

If, however, the sample is immersed in a strong external magnetic field having a fixed direction, then the spin axes of the hydrogen protons will tend to align either in the same direction as the external field or in the opposite direction. There will be some of each, but measurably more in the same direction as the field, as this state involves a lower energy level. To be more precise, the axes of the spinning hydrogen protons will not align exactly with the external field but will precess, or wobble, about it, much as a spinning top whose axis is not vertical precesses about the vertical as it spins. Due to the precession, the magnetic moment of any one particular nucleus will have a vector

component that is perpendicular to the direction of the external field. However, because there is no preferred phase for the precession, the phases of the various nuclei in the sample will be randomly distributed and, as a result, the sum of these components will cancel out. Over the whole sample, then, the effect is of an aggregate nuclear magnetic moment that is aligned in the same direction as the external magnetic field.

In this equilibrium state, with all of the nuclei behaving in basically the same way, exhibiting a sort of “herd mentality,” no useful electrical signal will be generated. To create a signal, a second magnetic field is introduced, one that oscillates in a plane perpendicular to the static external field. This new field causes the alignment of the nuclei to flip out of the equilibrium state. As the nuclei begin to precess about a new axis, the specifics of which depend on the strength and duration of the new field, the aggregate nuclear magnetic moment develops a nonzero net component transverse to the static magnetic field. After a certain amount of time, this second field is cut off and the nuclei relax back to their previous state.

There are two aspects to this relaxation.

As the net transverse component of the magnetic moment precesses, it induces an electromotive force (emf) in a coil that surrounds the sample. At the same time, with only the static field in effect, the nuclei gradually become de-phased in the transverse direction as they move toward equilibrium. Essentially, the net transverse magnetization describes a decaying wave, the rate of decay of which can be measured by analyzing the resulting induced emf. This process is known as *spin-spin*, or T_2 , relaxation.

While the transverse component of the nuclear magnetic moment decays to zero, the component in the direction of the static magnetic field returns, asymptotically, to its equilibrium level. This process is called *spin-lattice*, or T_1 , relaxation. As we shall see, the rate at which this process evolves can be measured through careful manipulation of the second magnetic field and the analysis of an electrical signal induced by the motion of the aggregate nuclear magnetic moment.

Different tissue types, or, more precisely, magnetized nuclei contained inside different chemical environments, have different T_1 and T_2 relaxation rates, the measurements of which reveal the types of material present within the sample. To resolve this information spatially, so that an image can be created showing the location within the sample of each type of tissue, additional magnetic fields, known as gradients, are introduced to the experiment. This is discussed in what follows.

At the molecular and atomic level, quantum effects certainly exist, but the classical approach outlined here is more feasible when it comes to designing practical machines to implement the system. So, with this conceptual framework in mind, we turn our attention to the mathematical model introduced by Bloch in 1946.

10.3 The Bloch equation

In an MRI machine, a strong, uniform, and steady magnetic field, \mathbf{B}_0 , is generated by running an electrical current through a large coil. In clinical applications, the direction of the field \mathbf{B}_0 is aligned along the length of the patient's body, which direction is taken as the z -axis of a Cartesian coordinate system and is also referred to as the *longitudinal* direction. The magnitude or strength of \mathbf{B}_0 , denoted by B_0 (without the boldface type), is usually about 0.5 tesla in practice.

Denote the aggregate magnetic moment of the nuclei in a sample by $\mathbf{M}(t, \mathbf{p})$, or simply by \mathbf{M} if the context is understood. This is a vector function that depends both on time t and on the location \mathbf{p} within the sample. The coordinate functions of \mathbf{M} are denoted by M_x , M_y , and M_z . That is, for each time t and each point \mathbf{p} in the sample,

$$\mathbf{M}(t, \mathbf{p}) = \langle M_x(t, \mathbf{p}), M_y(t, \mathbf{p}), M_z(t, \mathbf{p}) \rangle.$$

As discussed above, in the presence of the static magnetic field \mathbf{B}_0 alone, the equilibrium nuclear magnetization of the sample is directed along the z -axis. That is, at equilibrium, $\mathbf{M} = \langle 0, 0, M_{eq} \rangle$, where M_{eq} is the magnitude of this vector.

In addition to the steady magnetic field \mathbf{B}_0 , a variety of other magnetic fields are introduced. These additional fields alter the magnitude and alignment of the nuclear magnetization of the sample and vary both temporally and spatially. Here, we use $\mathbf{B} = \mathbf{B}(t, \mathbf{p})$ to denote the total external magnetic field experienced by the sample at time t and location \mathbf{p} .

The *Bloch equation* models the rate of change over time of the magnetic moment of the nuclei at each point in the sample. With the notation just introduced, the equation is

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} - \frac{\langle M_x, M_y, 0 \rangle}{T_2} - \frac{\langle 0, 0, M_z - M_{eq} \rangle}{T_1}, \quad (10.1)$$

where γ , T_1 , and T_2 are constants. For reasons that will be clear soon, the value of γ is related to the resonant frequency of the system, while T_1 and T_2 are called the *relaxation times*.

In the presence only of the static magnetic field $\mathbf{B}_0 = \langle 0, 0, B_0 \rangle$, directed along the z -axis, the Bloch equation simplifies to the system of equations

$$\begin{aligned} \frac{dM_x}{dt} &= \gamma B_0 M_y(t) - \frac{M_x(t)}{T_2}, \\ \frac{dM_y}{dt} &= -\gamma B_0 M_x(t) - \frac{M_y(t)}{T_2}. \\ \frac{dM_z}{dt} &= -\frac{M_z(t) - M_{eq}}{T_1}. \end{aligned} \quad (10.2)$$

The first two of these equations define a first-order linear system with a constant coefficient matrix that can be solved using standard eigenvalue–eigenvector methods. The third equation can be treated either as a separable equation or as a first order linear equation. The upshot of this analysis is that

$$\begin{aligned} M_x(t) &= e^{-t/T_2} (M_x(0) \cos(\omega_0 t) - M_y(0) \sin(\omega_0 t)), \\ M_y(t) &= e^{-t/T_2} (M_x(0) \sin(\omega_0 t) + M_y(0) \cos(\omega_0 t)), \text{ and} \\ M_z(t) &= M_z(0) e^{-t/T_1} + M_{eq} (1 - e^{-t/T_1}), \end{aligned} \quad (10.3)$$

where $\omega_0 = -\gamma B_0$. Thus, we see that, for times t that are large compared to the value of T_1 , the longitudinal component $M_z(t)$ tends toward the equilibrium magnetization M_{eq} . Meanwhile, in the transverse plane, as the xy -plane is called in this context, the magnetic moment of the nuclei rotates, or precesses, about the z -axis with angular frequency $\omega_0 = -\gamma B_0$ radians per second, known as the *Larmor frequency*. The constant γ , called the *gyromagnetic ratio*, typically has a value of about 2.68×10^8 radians per second per tesla, or, equivalently, about 42.6 megahertz (MHz) per tesla. (One hertz is one cycle per second, or 2π radians per second.) Thus, in a typical MRI experiment, the Larmor frequency lies in the radio frequency band. In comparison, X-rays have a much higher frequency of about 3 gigahertz.

As $t \rightarrow \infty$, the transverse component of the nuclear magnetization tends to zero, which does not mean, however, that the precession of the individual nuclei about the z -axis ceases. Rather, these rotations go out of phase with each other so that the distribution of the individual moments becomes random and the aggregate transverse component tends to zero. In the equilibrium state, each nucleus tends to precess with angular frequency ω_0 about the z -axis with longitudinal component M_{eq} . A useful image, again, is that of a spinning top whose axis is not vertical but precesses about the vertical with a fixed frequency.

The constant T_1 is called the *spin–lattice* relaxation time and reflects the dissipation in energy away from the spinning nuclei (the spin system) as the atomic and molecular structure (the lattice) of the sample settles into the equilibrium state. The *spin–spin* relaxation time, as T_2 is known, reflects the randomization of the phases of the spinning nuclei as the aggregate transverse component goes to zero. Thus, T_2 reflects the dissipation of energy within the spin system.

10.4 The RF field

No signal is emitted by the atomic nuclei so long as they are subjected only to the static magnetic field \mathbf{B}_0 . To knock them out of this equilibrium, a radio frequency (RF) transmitter is used to apply a linearly polarized RF magnetic field

$$\mathbf{B}_1 = \langle 2B_1 \cos(\omega t), 0, 0 \rangle.$$

This field is generated by sending an oscillating electrical current through a transmitting coil that surrounds the sample. The field \mathbf{B}_1 oscillates along the x -axis with frequency ω , called the *irradiation frequency*, and is effectively the sum of two circularly polarized fields that oscillate in the xy -plane with the same frequency but in opposite directions. Namely,

$$\mathbf{B}_1 = \langle B_1 \cos(\omega t), B_1 \sin(\omega t), 0 \rangle + \langle B_1 \cos(\omega t), -B_1 \sin(\omega t), 0 \rangle.$$

Physically, the nuclei, and, hence, the aggregate magnetic moment, are significantly affected only by the circular field that oscillates in the same direction as the precession. This means that we may take

$$\mathbf{B}_1 = \langle B_1 \cos(\omega t), B_1 \sin(\omega t), 0 \rangle \quad (10.4)$$

to be the effective RF magnetic field.

The magnetic field in the longitudinal direction consists of the static field \mathbf{B}_0 and the contribution from the gradient field, \mathbf{B}_G , discussed below. The overall magnetic field applied to the sample is, then,

$$\mathbf{B} = \mathbf{B}_1 + \mathbf{B}_0 + \mathbf{B}_G = \langle B_1 \cos(\omega t), B_1 \sin(\omega t), B_0 + B_G \rangle. \quad (10.5)$$

The duration of the RF pulse is short relative to the values of T_1 and T_2 , sufficiently so that we may ignore the T_1 and T_2 terms when we analyze the Bloch equation (10.1) during this time interval. So, in this context, the Bloch equation leads to a system of linear differential equations with nonconstant coefficients; specifically,

$$\begin{aligned} \frac{dM_x}{dt} &= \gamma(B_0 + B_G)M_y - \gamma B_1 M_z \sin(\omega t), \\ \frac{dM_y}{dt} &= -\gamma(B_0 + B_G)M_x + \gamma M_z B_1 \cos(\omega t), \\ \frac{dM_z}{dt} &= \gamma B_1 M_x \sin(\omega t) - \gamma B_1 M_y \cos(\omega t). \end{aligned} \quad (10.6)$$

To render this system more amenable to solution, it is convenient to introduce a rotating coordinate frame for the transverse plane, rather than the usual x - and y -coordinates. To this end, let $\mathbf{e}_1 = \langle \cos(\omega t), \sin(\omega t) \rangle$ and $\mathbf{e}_2 = \langle -\sin(\omega t), \cos(\omega t) \rangle$, and set

$$\begin{aligned} u(t) &= M_x(t) \cos(\omega t) + M_y(t) \sin(\omega t) \text{ and} \\ v(t) &= M_y(t) \cos(\omega t) - M_x(t) \sin(\omega t). \end{aligned} \quad (10.7)$$

Thus,

$$\begin{aligned}M_x &= u \cos(\omega t) - v \sin(\omega t) \text{ and} \\M_y &= u \sin(\omega t) + v \cos(\omega t).\end{aligned}\tag{10.8}$$

The vector $\langle M_x(t), M_y(t) \rangle$ in the standard coordinate frame is the same as the vector $u(t)\mathbf{e}_1 + v(t)\mathbf{e}_2$ in the rotating frame.

Translated into the rotating frame, (10.6) yields the system

$$\begin{aligned}\frac{du}{dt} &= [\gamma(B_0 + B_G) + \omega]v, \\ \frac{dv}{dt} &= -[\gamma(B_0 + B_G) + \omega]u + \gamma B_1 M_z, \\ \frac{dM_z}{dt} &= -\gamma B_1 v.\end{aligned}\tag{10.9}$$

Suppose now that the RF transmitter is set so that the irradiation frequency matches the Larmor frequency. That is, suppose that we set $\omega = \omega_0$. Then the RF magnetic field oscillates in resonance with the natural frequency of the nuclei in the presence of the static field. With the rotating reference frame rotating about the z -axis at the same frequency as that at which the nuclei precess about the z -axis, it appears, from the point of view of the nuclei themselves, that they are not precessing at all. In other words, the apparent effect is as if there were no static field. This is what is behind the use of the word *resonance* in the terms *nuclear magnetic resonance* and *magnetic resonance imaging*. To see how this affects the mathematical model, take $\omega = \omega_0 = -\gamma B_0$ in (10.9) to get the system

$$\begin{aligned}\frac{du}{dt} &= \gamma B_G v, \\ \frac{dv}{dt} &= -\gamma B_G u + \gamma B_1 M_z, \\ \frac{dM_z}{dt} &= -\gamma B_1 v.\end{aligned}\tag{10.10}$$

If, in addition, we set $B_G = 0$ and assume that B_1 is constant, then (10.10) has constant coefficients and its solution is

$$\begin{aligned}u(t) &= u(0), \\ v(t) &= v(0) \cos(-\gamma B_1 t) - M_z(0) \sin(-\gamma B_1 t), \text{ and} \\ M_z(t) &= M_z(0) \cos(-\gamma B_1 t) + v(0) \sin(-\gamma B_1 t).\end{aligned}\tag{10.11}$$

Thus, when viewed in the rotating coordinate frame, the aggregate magnetization vector of the nuclei in the sample precesses around the \mathbf{e}_1 -axis with frequency $\omega_1 = -\gamma B_1$. Meanwhile, the \mathbf{e}_1 -axis itself rotates around the z -axis at the Larmor frequency.

10.5 RF pulse sequences; T_1 and T_2

To generate a signal, suppose that the RF transmitter, whose irradiation frequency matches the Larmor frequency, is cut off after a time of τ seconds. The nuclear magnetic moment will then begin to relax from its state of precession about the \mathbf{e}_1 -axis back towards its equilibrium state of precession about the z -axis. As mentioned before, the aggregate nuclear magnetic moment now has a nonzero component transverse to the static magnetic field. As this net transverse component of the magnetic moment relaxes, an emf is induced in a coil surrounding the sample.

For instance, suppose that the oscillating RF magnetic field is cut off after time τ_1 , where $-\gamma B_1 \tau_1 = \pi/2$. This is called a $\pi/2$ pulse or a 90° pulse. From (10.11), we see that the net magnetization of the sample at time τ_1 , when viewed in the rotating frame, is given by

$$u(\tau_1) = u(0), \quad v(\tau_1) = -M_z(0), \quad \text{and} \quad M_z(\tau_1) = v(0). \quad (10.12)$$

That is, the effect of this RF field is that the orientation of the aggregate magnetization vector has been flipped by an angle of $\pi/2$, or 90° , from alignment with the z -axis into alignment with the \mathbf{e}_2 -axis.

A second important RF pulse is the π , or 180° , pulse, in which an RF magnetic field, oscillating at the Larmor frequency, is cut off after time τ_2 , with $-\gamma B_1 \tau_2 = \pi$. (Obviously, $\tau_2 = 2\tau_1$.) At the instant when the field is cut off, the net magnetization in the sample is

$$u(\tau_2) = u(0), \quad v(\tau_2) = -v(0), \quad \text{and} \quad M_z(\tau_2) = -M_z(0). \quad (10.13)$$

Thus, the orientation of the aggregate magnetization vector has been flipped by an angle of π , or 180° , from alignment with the positive z -axis into alignment with the negative z -axis.

In the *inversion recovery* method for measuring T_1 , a 180° RF pulse is applied in order to align the aggregate nuclear magnetization with the negative z -axis. After the pulse has been cut off, a time τ is allowed to pass, during which the magnetization partially recovers back toward the equilibrium state. Then a 90° RF pulse is applied to flip the partially recovered magnetization into the xy -plane. Following this pulse, the resulting signal is acquired. The size of the signal depends on the value of $M_z(\tau)$. Once the signal has been acquired, the magnetization is allowed to relax all the way back to equilibrium.

From (10.3), we see that

$$(M_{eq} - M_z(\tau)) = (M_{eq} - M_z(0)) e^{-\tau/T_1}.$$

Hence,

$$\ln(M_{eq} - M_z(\tau)) = \ln(M_{eq} - M_z(0)) - \tau/T_1.$$

This means that, if we graph the value of $\ln(M_{eq} - M_z(\tau))$ against the value of τ , the result is a straight line of slope $-1/T_1$. The signal acquisition step tells us the value of $M_z(\tau)$, so, by applying the inversion recovery pulse sequence just described for a variety of values of τ , we can determine the value of T_1 .

One technical detail is that, because the signals involved are fairly weak, the pulse sequence should be repeated several times for each selected value of τ in order to increase the signal-to-noise ratio, and thus our confidence, in the measurement of $M_z(\tau)$.

The inversion recovery method is typically abbreviated as

$$(180^\circ - \tau - 90^\circ - AT - t_\infty)_n,$$

where AT refers to the signal acquisition time and the value of t_∞ is large compared to T_1 so that the magnetization has time to relax back to equilibrium. (Generally, take $t_\infty > 4 \cdot T_1$.) The subscript n indicates the number of times the sequence is to be repeated for each selected value of τ .

Another approach to measuring T_1 is *saturation recovery*, abbreviated as

$$(90^\circ - HS - \tau - 90^\circ - AT - HS)_n.$$

In this scheme, HS refers to the application of a magnetic field pulse that destroys the homogeneity of the static magnetic field and results in a system that is “saturated”—the nuclear magnetic moments are scattered. After time τ , the magnetization has partially recovered back toward equilibrium. This magnetization is flipped into the xy -plane by a 90° pulse, and the resulting signal, which depends on $M_z(\tau)$, is measured. Then the field homogeneity is destroyed again and the sequence is repeated. As with inversion recovery, a semi-log plot of $\ln(M_{eq} - M_z(\tau))$ as a function of τ produces a straight line of slope $-1/T_1$.

The principal method for measuring the value of T_2 is called *spin-echo*. First proposed, in 1950, by Hahn, this method is abbreviated

$$(90^\circ - \tau - 180^\circ)_n.$$

The initial 90° RF pulse flips the magnetization into the xy -plane. Once this pulse has ended, we would expect the magnetization to precess at the fixed frequency ω_0 as it relaxes back to equilibrium. This would mean a constant component of magnetization in the direction of the \mathbf{e}_2 -axis in the rotating frame. However, there is really a small

spread of frequencies, at least partly because of slight inhomogeneity in the static field. Thus, after some time τ has passed, the nuclear moments have fanned out a bit, with the ones that are rotating faster getting ahead of the slower ones. Application of a 180° RF pulse has the effect of flipping this fan over in the xy -plane, so that the faster rotators are now *behind* the slower ones. Thus, after another time lapse of duration τ , the moments will come back together, forming an *echo* of the magnetization that existed immediately following the original 90° pulse. The echo is weaker than the original, though, because of the random nature of spin–spin relaxation, which cannot be refocused, or undone, by the echo. During the time interval of length 2τ , the amplitude of the transverse magnetization will have diminished, or relaxed, by a factor of $e^{-2\tau/T_2}$. By repeating the process, we increase the signal-to-noise ratio in the measurements.

A variation on the spin-echo method is to use a train of 180° RF pulses spaced at intervals of 2τ , with signal acquisition midway between successive pulses.

10.6 Gradients and slice selection

The work we have done so far provides the foundation for the methods of NMR spectroscopy. There, the investigator's aim is basically to create a graph of the spectrum of the signal generated by the nuclear magnetic moments in a sample as they respond to some specific RF pulse sequence. As we shall discuss below, the different frequency components present in the signal, and their corresponding amplitudes, provide information about the variety of relaxation times and, hence, the variety of chemical environments present in the sample. In this way, the investigator can analyze the chemical composition of the sample.

The insight that earned a Nobel Prize for Lauterbur was that, by introducing yet another carefully controlled component to the magnetic field, it is possible to restrict the fluctuation in the aggregate magnetic moment of the sample to a specified slice within the sample. In other words, it is possible to localize the behavior of the nuclear magnetic moments and thereby generate an image of a specified slice that reveals the locations of the various chemical environments to be found there. This is the core idea behind magnetic resonance imaging, or MRI. A typical MRI study consists of creating a whole series of such images that, collectively, help the investigator to form a sense of the three-dimensional composition and structure of a sample.

The key to slice selection is to introduce a magnetic field, \mathbf{B}_G , that is oriented in the same direction as the static external field but that varies in magnitude according to the location within the sample. To achieve this, \mathbf{B}_G is actually composed of three separately controlled magnetic fields. The first of these three varies in magnitude in proportion to the x -coordinate of the point in the sample; the magnitude of the second field is proportional to the y -coordinate of the location; and the magnitude of the third

field is proportional to the z -coordinate. That is, there are constants G_1 , G_2 , and G_3 such that, at the point $\mathbf{p} = \langle x, y, z \rangle$,

$$\begin{aligned}\mathbf{B}_G(\mathbf{p}) &= \langle 0, 0, G_1x + G_2y + G_3z \rangle \\ &= \langle 0, 0, \mathbf{G} \cdot \mathbf{p} \rangle,\end{aligned}\tag{10.14}$$

where $\mathbf{G} = \langle G_1, G_2, G_3 \rangle$. This magnetic field is called a *gradient field* because the constants G_1 , G_2 , and G_3 measure the gradient of the field as the location varies within the sample. The magnitude of \mathbf{B}_G at the point \mathbf{p} is $B_G(\mathbf{p}) = |\mathbf{G} \cdot \mathbf{p}|$.

Geometrically, the set $\{\mathbf{p} : \mathbf{G} \cdot \mathbf{p} = 0\}$ defines a plane through the origin with normal vector \mathbf{G} . Now, let

$$G = |\mathbf{G}| = \sqrt{G_1^2 + G_2^2 + G_3^2}$$

and observe that, for any given number $\alpha > 0$, the vector $\mathbf{r}_\alpha = (\alpha/G)\mathbf{G}$ satisfies $\mathbf{r}_\alpha \cdot \mathbf{G} = \alpha G$ and $|\mathbf{r}_\alpha| = \alpha$. Thus, the set

$$\{\mathbf{p} : |\mathbf{G} \cdot \mathbf{p}| \leq \alpha G\}\tag{10.15}$$

defines a slice of thickness 2α centered at the origin and normal to \mathbf{G} . In this way, we can select a slice of the sample that we wish to image, identify an appropriate normal vector \mathbf{G} , and tailor the gradient magnetic field \mathbf{B}_G accordingly. (Experimentally, it is possible to locate the origin at any desired point within the sample, so it suffices to consider slices centered on the origin.)

The total magnetic field is now given by (10.5) and the Bloch equation has the form given in (10.6). Taking $\omega = \omega_0 = -\gamma B_0$, the Larmor frequency, and viewing the system in the rotating frame (10.7), the Bloch equation translates into the system in (10.10). That is, we have

$$\begin{aligned}\frac{du}{dt} &= \gamma B_G v, \\ \frac{dv}{dt} &= -\gamma B_G u + \gamma B_1 M_z, \\ \frac{dM_z}{dt} &= -\gamma B_1 v.\end{aligned}\tag{10.16}$$

The T_1 and T_2 terms are not present in this system, so we are assuming that the duration of the gradient magnetic field is short compared to the relaxation times. If we also assume that the RF field \mathbf{B}_1 is weak, then (10.16) yields $M_z(t) \approx M_z(0)$ and we can write the other two equations as

$$\begin{aligned}\frac{du}{dt} &= \gamma B_G v = -\omega_G v, \\ \frac{dv}{dt} &= -\gamma B_G u + \gamma B_1 M_z(0) = \omega_G u + \gamma B_1 M_z(0),\end{aligned}\quad (10.17)$$

where $\omega_G = -\gamma B_G$.

Now introduce the complex-number-valued function

$$\phi(t) = u(t) + i \cdot v(t). \quad (10.18)$$

Differentiating, and using (10.17), we get

$$\begin{aligned}\frac{d\phi}{dt} &= \frac{du}{dt} + i \cdot \frac{dv}{dt} \\ &= -\omega_G v + i \cdot \omega_G u + i \cdot \gamma B_1 M_z(0) \\ &= i \cdot \omega_G \cdot \phi(t) + i \cdot B_1 M_z(0).\end{aligned}\quad (10.19)$$

This is a first-order linear differential equation with integrating factor $e^{-i\omega_G t}$. With the initial condition $\phi(0) = 0$, the solution is

$$\phi(t) = i\gamma M_z(0) e^{i\omega_G t} \cdot \left\{ \int_0^t B_1(s) \cdot e^{-i\omega_G s} ds \right\}. \quad (10.20)$$

Substituting $\omega_G = -\gamma B_G = -\gamma(\mathbf{G} \cdot \mathbf{p})$ into (10.20) gives us

$$\phi(t, \mathbf{p}) = i\gamma M_z(0) e^{-i\gamma(\mathbf{G} \cdot \mathbf{p})t} \cdot \left\{ \int_0^t B_1(s) \cdot e^{i\gamma(\mathbf{G} \cdot \mathbf{p})s} ds \right\}. \quad (10.21)$$

Next, note that the RF pulse is exactly that—a pulse. So it shuts off after a certain amount of time τ . Thus, $B_1(s) = 0$ for $s \leq 0$ and for $s > \tau$. It follows that $\phi(t, \mathbf{p}) = \phi(\tau, \mathbf{p})$ for all $t > \tau$. Using (10.21) and a change of variables in the integration, we see that, for $t > \tau$,

$$\begin{aligned}\phi(t, \mathbf{p}) &= i\gamma M_z(0) e^{-i\gamma(\mathbf{G} \cdot \mathbf{p})\tau} \cdot \left\{ \int_0^\tau B_1(s) \cdot e^{i\gamma(\mathbf{G} \cdot \mathbf{p})s} ds \right\} \\ &= i\gamma M_z(0) e^{-i\gamma(\mathbf{G} \cdot \mathbf{p})\tau} \cdot \left\{ \int_{s=-\tau/2}^{\tau/2} B_1(s + \tau/2) \cdot e^{i\gamma(\mathbf{G} \cdot \mathbf{p})(s+\tau/2)} ds \right\} \\ &= i\gamma M_z(0) e^{-i\gamma(\mathbf{G} \cdot \mathbf{p})\tau/2} \cdot \left\{ \int_{s=-\tau/2}^{\tau/2} B_1(s + \tau/2) \cdot e^{i\gamma(\mathbf{G} \cdot \mathbf{p})s} ds \right\}.\end{aligned}\quad (10.22)$$

Moreover, it follows from (10.22) that

$$|\phi(t, \mathbf{p})| = \gamma |M_z(0)| \cdot \left| \int_{s=-\tau/2}^{\tau/2} B_1(s + \tau/2) \cdot e^{i\gamma(\mathbf{G} \cdot \mathbf{p})s} ds \right|. \quad (10.23)$$

The integral in (10.22) and (10.23) is approximately equal to the inverse Fourier transform of the function $B_1(s + \tau/2)$ evaluated at $\gamma(\mathbf{G} \cdot \mathbf{p})$, provided that the value of $B_1(s + \tau/2)$ is small when $|s| > \tau/2$. In particular, if B_1 is a Gaussian, then the inverse Fourier transform of B_1 is also a Gaussian (see (5.20)). Hence, from (10.23), $|\phi(t, \mathbf{p})|$ will be “large” when $|\mathbf{G} \cdot \mathbf{p}|$ is small and, just as importantly, $|\phi(t, \mathbf{p})|$ will be “small” when $|\mathbf{G} \cdot \mathbf{p}|$ is large.

Specifically, consider the slice of thickness 2α defined, as in (10.15), by

$$\{\mathbf{p} : |\mathbf{G} \cdot \mathbf{p}| \leq \alpha G\}.$$

Shaping the RF pulse so that

$$B_1(s + \tau/2) = e^{-(\alpha\gamma G s)^2/8} \quad (10.24)$$

results in

$$|\phi(t, \mathbf{p})| \text{ proportional to } e^{-2\left(\frac{\mathbf{G} \cdot \mathbf{p}}{\alpha G}\right)^2}. \quad (10.25)$$

Since $\int_{-1}^1 e^{-2x^2} dx \approx 0.9545 \cdot \int_{-\infty}^{\infty} e^{-2x^2} dx$, the Gaussian in (10.25) has about 95% of its area in the selected slice. Hence, the transverse nuclear magnetic moment, as measured by $|\phi(t, \mathbf{p})|$, is predominantly concentrated in that slice.

10.7 The imaging equation

Having explored how an array of magnetic fields can be carefully orchestrated to produce a spatially-encoded fluctuation in the transverse component of the nuclear magnetic moment of a sample, we will examine now, in only the coarsest fashion, how this leads to an image that is clinically useful.

The fluctuating magnetization of the nuclei in the sample induces an electromotive force (emf) in a coil surrounding the sample. Faraday’s law of induction shows how to express this induced emf in terms of the derivative of the magnetization. In practice, it is most convenient to represent the magnetization in the sample by the complex-number-valued function

$$M^*(t, \mathbf{p}) = M_x(t, \mathbf{p}) + i \cdot M_y(t, \mathbf{p}). \quad (10.26)$$

It follows from (10.7) and (10.18), that

$$\begin{aligned} |M^*(t, \mathbf{p})|^2 &= (M_x(t, \mathbf{p}))^2 + (M_y(t, \mathbf{p}))^2 \\ &= (u(t, \mathbf{p}))^2 + (v(t, \mathbf{p}))^2 \\ &= |\phi(t, \mathbf{p})|^2. \end{aligned}$$

Thus, in the presence of the gradient magnetic field \mathbf{B}_G , the value of M^* is “large” inside the selected slice and “small” outside the slice.

The signal that is induced in the coil is sent through a preamplifier and then is subjected to a variety of treatments, including a phase detection step that shifts the signal down in frequency by ω_0 so that the frequencies present in the modified signal are centered around 0. This modification simplifies the signal analysis in later steps. Application of a low-pass filter increases the signal-to-noise ratio. At this stage, the signal $S(t)$, induced by the nuclear magnetization and modified by the receiver system, can be represented as

$$S(t) = K \int M^*(t, \mathbf{p}) \exp(-i\omega t) d\mathbf{p}, \quad (10.27)$$

where K is some complex number constant. This formula is called the *imaging equation*. It expresses the signal $S(t)$ as the (2- or 3-dimensional) Fourier transform of the complex transverse magnetization M^* of the sample. Thus, the function M^* can be recovered, and an image of it created, by applying the inverse Fourier transform.

In practice, the signal is sampled at a discrete set of times $\{k \cdot \Delta t\}$ and the inversion is done on a digital computer. Thus, some of the techniques discussed in Chapter 8 come into play, including sampling, the discrete Fourier transform and its inverse, and the fast Fourier transform.

10.8 Exercises

10.1. Verify that (10.3) is the solution to the system (10.2).

10.2. Verify that the system (10.6) is equivalent to the system (10.9) when translated into the rotating coordinate frame for the transverse plane.

10.3. Verify that (10.11) gives the solution to the system (10.10).

10.4. Verify that (10.20) is the solution to the system (10.19).

Appendix A

Integrability

A.1 Improper integrals

The Radon transform, back projection, and Fourier transform all involve *improper integrals*, evaluated over infinite intervals. We have applied these concepts, computing examples and proving theorems, without considering the more technical questions of how these improper integrals are defined and for what functions they make sense. We ought not evade these questions completely, so let us now attend to them.

Consider a function f , defined on the real line and having either real or complex values, with the property that the integral $\int_a^b f(x) dx$ exists for every finite interval $[a, b]$. If the limit $\lim_{b \rightarrow \infty} \int_a^b f(x) dx$ exists, then we denote this limit by $\int_a^\infty f(x) dx$ and we say that this improper integral converges. That is,

$$\int_a^\infty f(x) dx = \lim_{b \rightarrow \infty} \int_a^b f(x) dx$$

provided the limit exists.

Similarly,

$$\int_{-\infty}^b f(x) dx = \lim_{a \rightarrow -\infty} \int_a^b f(x) dx$$

provided the limit exists, in which case we say that the improper integral $\int_{-\infty}^b f(x) dx$ converges.

If both of the improper integrals $\int_a^\infty f(x) dx$ and $\int_{-\infty}^a f(x) dx$ converge for some real number a , then we define

$$\int_{-\infty}^\infty f(x) dx = \int_{-\infty}^a f(x) dx + \int_a^\infty f(x) dx.$$

Again we say that the improper integral $\int_{-\infty}^{\infty} f(x) dx$ converges and the function f is said to be *integrable* on the real line. If $|f|$ is integrable on the real line, then we say that f is *absolutely integrable*.

Some facts about improper integrals are similar to facts about infinite series. Here are a few, with the proofs left as exercises.

- If f and g are both integrable on $[a, \infty)$ and if c is any number, then $f + cg$ is also integrable on $[a, \infty)$ and

$$\int_a^{\infty} (f + cg)(x) dx = \int_a^{\infty} f(x) dx + c \int_a^{\infty} g(x) dx.$$

- If $f \geq 0$ and the set $\mathcal{S} = \{\int_a^b f(x) dx : b \geq a\}$ is bounded above, then $\int_a^{\infty} f(x) dx$ converges to the least upper bound of the set \mathcal{S} .
- If $|f(x)| \leq g(x)$ for all real x and the integral $\int_a^{\infty} g(x) dx$ converges, then $\int_a^{\infty} f(x) dx$ also converges and $|\int_a^{\infty} f(x) dx| \leq \int_a^{\infty} g(x) dx$. (This is a form of the comparison test. To prove it for f real-valued, apply the previous facts to $f = f_+ - f_-$, where $f_+(x) = \max\{f(x), 0\}$ and $f_-(x) = -\min\{f(x), 0\}$.)
- If $\int_a^{\infty} |f(x)| dx$ converges, then so does $\int_a^{\infty} f(x) dx$. In words, if f is absolutely integrable on $[a, \infty)$, then f is integrable there. This is a corollary of the previous fact. The converse statement is not true.

Similar statements prevail regarding improper integrals on $(-\infty, b)$ or on $(-\infty, \infty)$.

One large class of functions to which we can look for examples are the piecewise continuous functions. A real- or complex-valued function f , defined on the real line, is *piecewise continuous* if, in every finite interval $[a, b]$, there are only a finite number of points at which f is discontinuous and if the one-sided limits $\lim_{x \rightarrow \alpha^-} f(x)$ and $\lim_{x \rightarrow \alpha^+} f(x)$ both exist at each point of discontinuity α . Hence, we see that a piecewise continuous function is integrable on every finite interval of the real line. This is the starting point for asking whether any of the improper integrals above converge. The class of piecewise continuous functions will be denoted by \mathcal{PC} . This class includes all functions that are continuous on the real line.

Lebesgue's theory of integration makes it possible to extend the notion of integrability on the real line to more functions than the piecewise continuous ones. We will not go into this far-reaching theory here, but we will borrow the notation L^1 to denote the class of all (Lebesgue integrable) functions that are absolutely integrable on the real line. Thus, for instance, the set of piecewise continuous functions that are absolutely integrable is denoted by $L^1 \cap \mathcal{PC}$.

Since $|f(x)e^{-i\omega x}| = |f(x)|$ for all real numbers x and ω , it follows from the above facts that every function in L^1 has a Fourier transform. The Fourier inversion theorem (Theorem 5.11), however, applies only to continuous functions in L^1 , though a modified version of it applies to functions in $L^1 \cap \mathcal{PC}$. This is ample motivation for restricting our attention mainly to functions in $L^1 \cap \mathcal{PC}$.

A.2 Iterated improper integrals

Another technical matter is the manipulation of iterated improper integrals over the plane. This shows up in the proof of the Fourier inversion theorem (Theorem 5.11), for instance, where the order of integration with respect to the two variables is switched in the middle of the proof. This begs the question of whether such a step is justified.

Suppose that g is a continuous function of two real variables such that, for some real number a , the improper integral $\int_a^\infty g(x, y) dx$ converges for every value of y in some interval $J = [\alpha, \beta]$. Then we say that this improper integral converges *uniformly* on J provided that, for every $\varepsilon > 0$, there exists a number B such that

$$\left| \int_a^b g(x, y) dx - \int_a^\infty g(x, y) dx \right| < \varepsilon \quad (\text{A.1})$$

for all $b > B$ and all $y \in J$.

The relevant facts, found in standard texts in elementary real analysis (such as [2], for example), are these.

- (i) With g as in the preceding paragraph, if $\int_a^\infty g(x, y) dx$ converges uniformly on the interval J , then the integral is a continuous function of y on J .

Proof. For each natural number n , let $G_n(y) = \int_a^{a+n} g(x, y) dx$. Each function G_n is continuous on J and the sequence $\{G_n\}$ converges uniformly on J to the function $G(y) = \int_a^\infty g(x, y) dx$. Hence, the function G is also continuous on J . \square

- (ii) Again supposing g to be a continuous function of two real variables, if the integral $\int_a^\infty g(x, y) dx$ converges uniformly on the interval $J = [\alpha, \beta]$, then the improper integral $\int_a^\infty \int_\alpha^\beta g(x, y) dy dx$ converges and

$$\int_a^\infty \int_\alpha^\beta g(x, y) dy dx = \int_\alpha^\beta \int_a^\infty g(x, y) dx dy.$$

Proof. With G_n and G as defined in the preceding proof, the uniform convergence implies that

$$\int_\alpha^\beta G(y) dy = \lim_{n \rightarrow \infty} \int_\alpha^\beta G_n(y) dy.$$

That is,

$$\int_\alpha^\beta \int_a^\infty g(x, y) dx dy = \lim_{n \rightarrow \infty} \int_\alpha^\beta \int_a^{a+n} g(x, y) dx dy.$$

The continuity of g implies that iterated integrals over finite rectangles can be evaluated in either order, so that

$$\int_\alpha^\beta \int_a^{a+n} g(x, y) dx dy = \int_a^{a+n} \int_\alpha^\beta g(x, y) dy dx.$$

Hence,

$$\lim_{n \rightarrow \infty} \int_{\alpha}^{\beta} \int_a^{a+n} g(x, y) dx dy = \int_a^{\infty} \int_{\alpha}^{\beta} g(x, y) dy dx$$

and the desired result follows. \square

- (iii) If g is continuous for $x \geq a$ and $y \geq \alpha$, and if the improper integrals $\int_a^{\infty} |g(x, y)| dx$ and $\int_{\alpha}^{\infty} |g(x, y)| dy$ converge uniformly on every finite interval, then if either of the integrals $\int_a^{\infty} \int_{\alpha}^{\infty} |g(x, y)| dy dx$ or $\int_{\alpha}^{\infty} \int_a^{\infty} |g(x, y)| dx dy$ converges,

$$\int_a^{\infty} \int_{\alpha}^{\infty} g(x, y) dy dx = \int_{\alpha}^{\infty} \int_a^{\infty} g(x, y) dx dy.$$

Proof. Suppose that $g \geq 0$ and that the integral $\int_{\alpha}^{\infty} \int_a^{\infty} g(x, y) dx dy$ converges. It follows from the previous result and the nonnegativity of g that, for each $b > a$,

$$\begin{aligned} \int_a^b \int_{\alpha}^{\infty} g(x, y) dy dx &= \int_{\alpha}^{\infty} \int_a^b g(x, y) dx dy \\ &\leq \int_{\alpha}^{\infty} \int_a^{\infty} g(x, y) dx dy. \end{aligned}$$

Hence, $\int_a^{\infty} \int_{\alpha}^{\infty} g(x, y) dy dx$ converges by the comparison test, and

$$\int_a^{\infty} \int_{\alpha}^{\infty} g(x, y) dy dx \leq \int_{\alpha}^{\infty} \int_a^{\infty} g(x, y) dx dy.$$

Reversing the above argument shows that

$$\int_a^{\infty} \int_{\alpha}^{\infty} g(x, y) dy dx \geq \int_{\alpha}^{\infty} \int_a^{\infty} g(x, y) dx dy,$$

from which it follows that the integrals are equal, as claimed.

In general, for a real- or complex-valued function g that satisfies the hypotheses, we may write $g = g_1 - g_2 + ig_3 - ig_4$, where each of the functions g_j satisfies $0 \leq g_j \leq |g|$. (For instance, take $g_1(x, y) = \max\{\Re g(x, y), 0\}$.) The result then applies to each g_j and, by the linearity of the integral, to g itself. \square

A.3 L^1 and L^2

A function f defined on the real line is in the class L^2 , and is said to be *square-integrable*, if the (improper) integral $\int_{-\infty}^{\infty} |f(x)|^2 dx$ is finite. Neither L^1 nor L^2 is a subset of the other. For instance, the function f given by

$$f(x) = \begin{cases} 1/x & \text{if } x \geq 1, \\ 0 & \text{if } x < 1 \end{cases}$$

is in L^2 but not in L^1 , while the function g defined by

$$g(x) = \begin{cases} 1/\sqrt{x} & \text{if } 0 < x \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

is in L^1 but not in L^2 .

The Rayleigh–Plancherel theorem (Theorem 7.8) applies to any absolutely integrable function f for which either f or its Fourier transform $\mathcal{F}f$ is square-integrable. In particular, for f in $L^1 \cap L^2$, it follows from (7.15) that $\mathcal{F}f$ is in L^2 . Moreover, the mapping $f \mapsto \mathcal{F}f$ is an *isometry* from $L^1 \cap L^2$ into L^2 .

It is a fact that $L^1 \cap L^2$ is dense in L^2 in the sense that, for every function f in L^2 , there exists a sequence of functions $\{f_k\}$ in $L^1 \cap L^2$ such that

$$\lim_{k \rightarrow \infty} \int_{-\infty}^{\infty} |f(x) - f_k(x)|^2 dx = 0.$$

(For instance, one may take f_k to be the restriction of f to the interval $[-k, k]$.) Each of the Fourier transforms $\mathcal{F}f_k$ is a well-defined function in L^2 , and we may therefore define the Fourier transform of f to be the limit of the transforms $\mathcal{F}f_k$. (This limit exists because of the isometry implied by (7.15) and the fact that the space L^2 is complete.)

A modified form of the Fourier inversion theorem (Theorem 5.11) also holds in the L^2 setting so that, in the end, the mapping $f \mapsto \mathcal{F}f$ defines an isometric mapping from L^2 onto L^2 . See Chapter 9 of [37] for a more in-depth discussion of these ideas.

A.4 Summability

Just as the Fourier transform and many of the theorems and computations accompanying it required the manipulation of improper integrals, so working with the discrete version of the Fourier transform can, in principle, involve working with infinite series and iterated infinite series. In practice, however, the discrete functions we use are always finite lists and the important theorems, like (8.22) and (8.24) to name two, use only finite sums.

The study of infinite series, Fourier series, and sequence spaces such as ℓ^1 and ℓ^2 is well worth the investment, but lies beyond the scope of our work here. The books [2] and [38] are good places to start.

Appendix B

Topics for Further Study

- For a wealth of information about the Radon transform and its generalizations, as well as an extensive list of references on this topic, see the monograph [20]. Investigate the interaction between the Radon transform and the derivative in order to better understand Radon's original inversion formula.
- The Fourier transform, like Fourier series, was developed originally in the study of differential equations related to the propagation and diffusion of heat. The interaction between the Fourier transform and derivatives was mentioned in the exercises but did not figure into the discussion of CT scans. Moreover, the Fourier transform is a primary tool of physicists, astronomers, and engineers that is used to tackle a broad range of problems. See Bracewell's definitive treatise [6] for much more on this topic.
- We have focused on the filtered back-projection algorithm as well as some basic ART techniques. We have not discussed so-called direct Fourier inversion, which takes the central slice theorem (Theorem 6.1) as its starting point. The filtered back projection (Theorem 6.2) is primarily what is used in current practice. Nonetheless, direct Fourier inversion is a worthwhile subject.
- Study the effect of incorporating finite (nonzero) X-ray beam width into the algorithms for CT.
- Investigate the *fan beam*, *spiral beam*, and *cone beam* geometries and their use in CT scan technology. The latter two methods can reduce the radiation exposure time of the patient by collecting data in all three dimensions at once.
- Investigate the use of wavelets, rather than the Fourier transform, in signal analysis. Wavelets are particularly useful for analyzing signals that are of short duration or that come in bursts. Wavelets are also at the core of the signal compression methods used, for example, in the creation of mp3 music and sound files.

- Study the evolution of the scanning machines themselves, from the earliest EMI scanner designed by Hounsfield to the 5th generation machines that employ electron beam CT.
- In positron emission tomography (PET) and single-photon emission computed tomography (SPECT), a radioactive isotope is introduced internally into the patient. This isotope tends to bind to areas where certain pathologies are present or certain physiological functions in effect. Positrons are emitted from the sample in pairs moving in opposite directions. When the intensities of a matching pair are measured by external detectors, the sum of the measurements yields a value of the Radon transform along the line defined by the paths of the two particles. Then a modified version of the CT analysis discussed here can be applied to form an image. So PET and SPECT are like “inside-out” versions of CT. See [24] for an introduction to these types of tomography.
- The study of Fourier series predates the development of the Fourier transform historically. Though we have alluded to the theory of Fourier series in only a few places in the present work, it nonetheless inspires and informs the transform theory and is a cornerstone of the branch of mathematics known as functional analysis.
- Functional MRI (fMRI) exploits the difference in the magnetic response of nuclei contained in oxygenated blood compared to those in deoxygenated blood. Increased neuronal activity requires a rapid influx of oxygen to the area of the activity, where the additional oxygen is consumed by the active neurons. MRI images that portray these magnetic variations are created during many repetitions of some activity or experience, such as performing mental arithmetic. Then statistical methods are used to determine which areas of the brain can reliably be said to be most active during the activity.
- For the strongest versions of many of the theorems in Chapters 5 and 7, as well as for a careful error analysis of the discrete approximations discussed in Chapter 8, we would require a deeper understanding of the spaces L^1 and L^2 of integrable functions and of the theory of integral operators and kernel functions.
- In this book, only a little attention has been given to the computer implementation of the image reconstruction algorithms discussed in Chapter 8. For instance, the imaging toolbox in MATLAB^R has subroutines for both the discrete Radon transform of a phantom and the discrete filtered back projection of data. Options in the latter subroutine include the choice of the interpolation method and of the filter function. Maple^R has both continuous and discrete versions of the Fourier transform and the inverse transform built in and was used to create many of the figures in this book. It would be interesting to look more closely at the computational process that is encoded into the scanning machines used in clinical practice.

References

1. Baillie, R., D. Borwein, and J. M. Borwein, Surprising sinc sums and integrals, *Amer. Math. Monthly* **115** (2008), 888–901.
2. Bartle, R. G., *The Elements of Real Analysis*, 2nd ed., John Wiley, New York, 1976.
3. Bauschke, H. H., and J. M. Borwein, On projection algorithms for solving convex feasibility problems, *SIAM Review* **38** (1996), 367–426.
4. Bloch, F., W. W. Hansen, and M. Packard, Nuclear induction, *Phys. Rev.* **69** (1946), 127.
5. Blümich, B., *NMR Imaging of Materials*, Oxford University Press, Oxford, 2000.
6. Bracewell, R. N., *The Fourier Transform and Its Applications*, 3rd ed., McGraw-Hill, Boston, 2000.
7. Bracewell, R. N., Image reconstruction in radio astronomy, in *Image Reconstruction from Projections: Implementation and Applications*, G. T. Herman, Editor, Topics in Applied Physics **32**, Springer, Berlin, 1979.
8. Bracewell, R. N., and A. C. Riddle, Inversion of fan-beam scans in radio astronomy, *Astrophysical J.* **150** (1967), 427–434.
9. Butz, T., *Fourier Transformation for Pedestrians*, Springer, Berlin, 2006.
10. Censor, Y., Row-action methods for huge and sparse systems and their applications, *SIAM Review* **23** (1981), 444–464.
11. Censor, Y., Finite series-expansion reconstruction methods, *Proc. IEEE* **71** (1983), 409–419.
12. Cooley, T. W., and J. W. Tukey, An algorithm for the machine calculation of complex Fourier series, *Math. Computation* **19** (1965), 297–301.
13. Cormack, A. M., Representation of a function by its line integrals, with some radiological applications I, II, *J. Appl. Phys.* **34** (1963), 2722–2727; **35** (1964), 2908–2912.
14. Epstein, C. L., *Introduction to the Mathematics of Medical Imaging*, 2nd ed., SIAM, Philadelphia, 2008.
15. Feeman, T. G., Conformality, the exponential function, and world map projections, *College Math Journal* **32** (2001), 334–342.
16. Freeman, R., *Magnetic Resonance in Chemistry and Medicine*, Oxford University Press, Oxford, 2003.
17. Gadian, D. G., *Nuclear Magnetic Resonance and Its Applications to Living Systems*, Oxford University Press, Oxford, 1982.
18. Gardner, M., Mathematical Recreations, *Scientific American* (September 1990).

19. Gordon, R., R. Bender, and G. T. Herman, Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography, *J. Theoretical Biology* **29** (1970), 471–481.
20. Helgason, S., *The Radon Transform*, 2nd ed., Birkhäuser, Boston, 1999.
21. Hinshaw, W. S., and A. H. Lent, An introduction to NMR imaging: From the Bloch equation to the imaging equation, *Proc. IEEE* **71** (1983), 338–350.
22. Hounsfield, G. N., Computerized transverse axial scanning tomography, *British J. Radiology* **46** (1973), 1016–1022.
23. Hounsfield, G. N., A method of and apparatus for examination of a body by radiation such as X or gamma radiation, The Patent Office, London, (1972) Patent Specification 1283915.
24. Knoll, G. F., Single-photon emission computed tomography, *Proc. IEEE* **71** (1983), 320–329.
25. Kuperman, V., *Magnetic Resonance Imaging: Physical Principles and Applications*, Academic Press, San Diego, 2000.
26. Lauterbur, P. C., Image formation by induced local interactions: examples employing nuclear magnetic resonance, *Nature* **242** (1973), 190.
27. Lewitt, R. M., Reconstruction algorithms: Transform methods, *Proc. IEEE* **71** (1983), 390–408.
28. Mansfield, P., Multi-planar image formation using NMR spin echoes, *J. Phys. C.* **10** L55 (1977).
29. Louis, A. K., and F. Natterer, Mathematical problems of computerized tomography, *Proc. IEEE* **71** (1983), 379–389.
30. Natterer, F., *The Mathematics of Computerized Tomography*, Classics in Applied Mathematics **32**, SIAM, Philadelphia, 2001.
31. Noble, B., and J. W. Daniel, *Applied Linear Algebra*, 3rd ed., Prentice-Hall, Englewood Cliffs, 1988.
32. *The Princeton Companion to Mathematics*, T. Gowers, Editor, Princeton University Press, Princeton, 2008.
33. Purcell, E. M., H. C. Torrey, and R. V. Pound, Resonance absorption by nuclear magnetic moments in a solid, *Phys. Rev.* **69** (1946), 37.
34. Radon, J., Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten, *Berichte Sächsische Akademie der Wissenschaften* **69** (1917), 262–277.
35. Ramachandran, G. N., and A. V. Lakshminarayanan, Three-dimensional reconstruction from radiographs and electron micrographs II: Application of convolutions instead of Fourier transforms, *Proc. Natl. Acad. Sci. US* **68** (1971), 2236–2240.
36. Rowland, S. W., Computer implementation of image reconstruction formulas, in *Image Reconstruction from Projections: Implementation and Applications*, G. T. Herman, Editor, Topics in Applied Physics **32**, Springer, Berlin, 1979.
37. Rudin, W., *Real and Complex Analysis*, 2nd ed., McGraw-Hill, New York, 1974.
38. Seeley, R., *An Introduction to Fourier Series and Integrals*, W. A. Benjamin, New York, 1966.
39. Shepp, L. A., and J. B. Kruskal, Computerized tomography: The new medical X-ray technology, *Amer. Math. Monthly* **34** (1978), 35–44.
40. Shepp, L. A., and B. F. Logan, The Fourier reconstruction of a head section, *IEEE Trans. Nucl. Sci.* **NS-21** (1974), 21–43.
41. Shingareva, I., and C. Lizárraga-Celaya, *Maple and Mathematica: A Problem Solving Approach for Mathematics*, Springer, Vienna, 2007.

Index

Symbols

\mathcal{B} 20
 \mathcal{F} 33
 $\ell_{t,\theta}$ 7
 \mathcal{R} 11
 \wedge (tent) 56
 δ (Dirac) 39
 \square_a 40

A

affine projection 108
affine space 107–108
ART 101–113
attenuation coefficient 3

B

back projection 19–23
 and convolution 62
 discrete 82
 filtered 49
band-limited function 64, 65
Beer's law 4–5
Bloch equation 118
Bloch, Felix 115

C

central slice theorem 47
complex exponential 27
complex numbers 25–30
convolution 55

2-dimensional 62
 and Fourier transform 58–60, 80
 discrete 74
Cormack, Allan McLeod xi
cubic spline 83

D

DeMoivre's law 28
Dirac δ 39
 Fourier transform of 39, 42
discrete function 74
 N -periodic 74
dynamical systems 109

E

Euler's formula 27
Euler, Leonhard 27
exponential function
 complex 27

F

fast Fourier transform *see* FFT
FFT 92–97
filter 53
 and medical imaging 64–65
 low-pass 64, 70–72
 low-pass cosine 65, 98
 Ram–Lak 65, 72
 resolution 58
 Shepp–Logan 65, 71
filtered back projection 49

discrete 67
 fMRI 136
 Fourier series 30, 68, 136
 Fourier transform 33–44, 135
 and L^2 133
 and convolution 58–60
 and MRI 128
 and Radon transform 47
 discrete 77, 91
 and convolution 80
 discrete inverse 79
 inverse of 40
 inversion theorem 40, 79
 multivariable 43–44
 Fourier, Jean-Baptiste Joseph 30
 FWHM 58
 of a Gaussian 58

G

Gaussian
 and MRI 127
 Fourier transform of 33, 42
 FWHM of 58

H

Heaviside function 38
 Heisenberg, Werner 116
 Hilbert transform 51
 Hilbert, David 51
 Hounsfield unit 3
 Hounsfield, Godfrey Newbold xi

I

image reconstruction 86
 fundamental question of 7, 9
 integration 129–132
 L^1, L^2 132–133, 136
 interpolation 82–86
 and convolution 85–86
 Shannon–Whittaker 69

K

Kaczmarz's method 103, 107–113

L

Larmor frequency 119, 121, 122, 125
 Lauterbur, Paul 116, 124

least squares 103–107
 normal equation 105, 106
 Lorentzian 46, 66

M

magnetic resonance 115–128
 T_1, T_2 66, 117–119, 122–124
 fMRI 136
 relaxation constants 66
 Mansfield, Peter 116

N

Nyquist distance 69
 Nyquist's theorem 69

P

PET xii, 136
 phantom 16
 Shepp–Logan 16, 23
 pixel basis 101
 power spectrum 60
 Purcell, Edward 116

R

Radon transform 11–17, 19
 and ART 102
 and back projection 20
 and convolution 62
 and Fourier transform 47
 discrete 73
 Radon, Johann xi, 11
 Ram–Lak filter 65, 72
 Rayleigh–Plancherel
 theorem 60, 81, 133

S

sampling 68–70
 oversampling 70
 sarcophagus 1
 Shepp–Logan
 filter 65, 71
 phantom 16, 23
 sinogram 14
 SPECT 136
 spin–lattice relaxation *see*
 magnetic resonance, T_1, T_2
 spin–spin relaxation *see*
 magnetic resonance, T_1, T_2

T

tomography 7

V

van Vleck, John 116

W

wavelets 135

window function 64

X

X-rays 1–5

beam width 3, 135

cone beam 135

fan beam 97–98, 135

parallel beam 73

spiral beam 135

Z

zero padding 75