

IMAGE ESTIMATION BY EXAMPLE
Geophysical Soundings Image Construction
(Madagascar version)

Jon F. Claerbout

© November 18, 2008

Contents

1	Basic operators and adjoints	1
1.0.1	Programming linear operators	3
1.1	FAMILIAR OPERATORS	4
1.1.1	Adjoint derivative	5
1.1.2	Transient convolution	7
1.1.3	Internal convolution	9
1.1.4	Zero padding is the transpose of truncation	10
1.1.5	Adjoint of products are reverse-ordered products of adjoints	11
1.1.6	Nearest-neighbor coordinates	11
1.1.7	Data-push binning	12
1.1.8	Linear interpolation	13
1.1.9	Spray and sum : scatter and gather	15
1.1.10	Causal and leaky integration	16
1.1.11	Backsolving, polynomial division and deconvolution	18
1.1.12	The basic low-cut filter	19
1.1.13	Nearest-neighbor normal moveout (NMO)	22
1.1.14	Coding chains and arrays	23
1.2	ADJOINT DEFINED: DOT-PRODUCT TEST	26
1.2.1	Definition of a vector space	26
1.2.2	Dot-product test for validity of an adjoint	27
1.2.3	The word “adjoint”	29
1.2.4	Matrix versus operator	29
1.2.5	Inverse operator	30
1.2.6	Automatic adjoints	31

2	Model fitting by least squares	33
2.1	HOW TO DIVIDE NOISY SIGNALS	33
2.1.1	Dividing by zero smoothly	33
2.1.2	Damped solution	34
2.1.3	Smoothing the denominator spectrum	35
2.1.4	Imaging	38
2.1.5	Formal path to the low-cut filter	38
2.2	MULTIVARIATE LEAST SQUARES	39
2.2.1	Inside an abstract vector	39
2.2.2	Normal equations	41
2.2.3	Differentiation by a complex vector	42
2.2.4	From the frequency domain to the time domain	42
2.3	KRYLOV SUBSPACE ITERATIVE METHODS	44
2.3.1	Sign convention	44
2.3.2	Method of random directions and steepest descent	45
2.3.3	Null space and iterative methods	47
2.3.4	Why steepest descent is so slow	48
2.3.5	Conjugate direction	48
2.3.6	Routine for one step of conjugate-direction descent	50
2.3.7	A basic solver program	50
2.3.8	Test case: solving some simultaneous equations	52
2.4	INVERSE NMO STACK	54
2.5	THE WORLD OF CONJUGATE GRADIENTS	56
2.5.1	Physical nonlinearity	56
2.5.2	Statistical nonlinearity	56
2.5.3	Coding nonlinear fitting problems	57
2.5.4	Standard methods	57
2.5.5	Understanding CG magic and advanced methods	58
2.6	REFERENCES	61
3	Empty bins and inverse interpolation	63
3.1	MISSING DATA IN ONE DIMENSION	64

CONTENTS

3.1.1	Missing-data program	67
3.2	WELLS NOT MATCHING THE SEISMIC MAP	71
3.3	SEARCHING THE SEA OF GALILEE	77
3.4	INVERSE LINEAR INTERPOLATION	80
3.4.1	Abandoned theory for matching wells and seismograms	82
3.5	PREJUDICE, BULLHEADEDNESS, AND CROSS VALIDATION	83
4	The helical coordinate	85
4.1	FILTERING ON A HELIX	85
4.1.1	Review of 1-D recursive filters	86
4.1.2	Multidimensional deconvolution breakthrough	87
4.1.3	Examples of simple 2-D recursive filters	88
4.1.4	Coding multidimensional de/convolution	90
4.1.5	Causality in two-dimensions	92
4.2	FINITE DIFFERENCES ON A HELIX	94
4.2.1	Matrix view of the helix	98
4.3	CAUSALITY AND SPECTRAL FACTORIZATION	99
4.3.1	The spectral factorization concept	99
4.3.2	Cholesky decomposition	101
4.3.3	Toeplitz methods	101
4.3.4	Kolmogoroff spectral factorization	102
4.4	WILSON-BURG SPECTRAL FACTORIZATION	102
4.4.1	Wilson-Burg theory	102
4.5	HELIX LOW-CUT FILTER	105
4.6	THE MULTIDIMENSIONAL HELIX	108
4.7	SUBSCRIPTING A MULTIDIMENSIONAL HELIX	109
5	Preconditioning	117
5.1	PRECONDITIONED DATA FITTING	117
5.1.1	Preconditioner with a starting guess	118
5.2	PRECONDITIONING THE REGULARIZATION	118
5.2.1	The second miracle of conjugate gradients	119
5.2.2	Importance of scaling	120

5.2.3	Statistical interpretation	121
5.2.4	The preconditioned solver	121
5.3	OPPORTUNITIES FOR SMART DIRECTIONS	122
5.4	NULL SPACE AND INTERVAL VELOCITY	122
5.4.1	Balancing good data with bad	123
5.4.2	Lateral variations	125
5.4.3	Blocky models	125
5.5	INVERSE LINEAR INTERPOLATION	126
5.6	EMPTY BINS AND PRECONDITIONING	128
5.6.1	SEABEAM: Filling the empty bins with a laplacian	129
5.6.2	Three codes for inverse masking	130
5.7	THEORY OF UNDERDETERMINED LEAST-SQUARES	132
5.8	SCALING THE ADJOINT	134
5.9	A FORMAL DEFINITION FOR ADJOINTS	136
6	Multidimensional autoregression	137
6.0.1	Time domain versus frequency domain	137
6.1	SOURCE WAVEFORM, MULTIPLE REFLECTIONS	138
6.2	TIME-SERIES AUTOREGRESSION	139
6.3	PREDICTION-ERROR FILTER OUTPUT IS WHITE	141
6.3.1	PEF whiteness proof in 1-D	143
6.3.2	Simple dip filters	144
6.3.3	PEF whiteness proof in 2-D	145
6.3.4	Examples of modeling and deconvolving with a 2-D PEF	147
6.3.5	Seismic field data examples	148
6.4	PEF ESTIMATION WITH MISSING DATA	155
6.4.1	Internal boundaries to multidimensional convolution	156
6.4.2	Finding the prediction-error filter	160
6.5	TWO-STAGE LINEAR LEAST SQUARES	160
6.5.1	Adding noise (Geostat)	161
6.5.2	Inversions with geostat	163
6.5.3	Infill of 3-D seismic data from a quarry blast	166

CONTENTS

6.5.4	Imposing prior knowledge of symmetry	166
6.5.5	Hexagonal coordinates	167
6.6	BOTH MISSING DATA AND UNKNOWN FILTER	168
6.6.1	Objections to interpolation error	169
6.6.2	Packing both missing data and filter into a vector	170
6.7	LEVELED INVERSE INTERPOLATION	172
6.7.1	Test results for leveled inverse interpolation	172
6.7.2	Analysis for leveled inverse interpolation	173
6.7.3	Seabeam: theory to practice	174
6.7.4	Risky ways to do nonlinear optimization	175
6.7.5	The bane of PEF estimation	175
6.8	MULTIVARIATE SPECTRUM	175
6.8.1	What should we optimize?	177
6.8.2	Confusing terminology for data covariance	178
6.8.3	Hermeneutics	178
7	Spatial aliasing and scale invariance	181
7.1	INTERPOLATION BEYOND ALIASING	181
7.1.1	Interlacing a filter	182
7.2	MULTISCALE, SELF-SIMILAR FITTING	183
7.2.1	Examples of scale-invariant filtering	185
7.2.2	Scale-invariance introduces more fitting equations	186
7.2.3	Coding the multiscale filter operator	186
7.3	References	188
8	Nonstationarity: patching	193
8.1	PATCHING TECHNOLOGY	194
8.1.1	Weighting and reconstructing	196
8.1.2	2-D filtering in patches	198
8.1.3	Designing a separate filter for each patch	200
8.1.4	Triangular patches	200
8.2	STEEP-DIP DECON	202
8.2.1	Dip rejecting known-velocity waves	202

8.2.2	Tests of steep-dip decon on field data	203
8.2.3	Are field arrays really needed?	206
8.2.4	Which coefficients are really needed?	206
8.3	INVERSION AND NOISE REMOVAL	207
8.4	SIGNAL-NOISE DECOMPOSITION BY DIP	207
8.4.1	Signal/noise decomposition examples	209
8.4.2	Spitz for variable covariances	210
8.4.3	Noise removal on Shearer's data	211
8.4.4	The human eye as a dip filter	212
8.5	SPACE-VARIABLE DECONVOLUTION	214
9	Plane waves in three dimensions	221
9.1	THE LEVELER: A VOLUME OR TWO PLANES?	221
9.1.1	PEFs overcome spatial aliasing of difference operators	223
9.1.2	My view of the world	224
9.2	WAVE INTERFERENCE AND TRACE SCALING	224
9.2.1	Computing the proper scale factor for a seismogram	225
9.3	LOCAL MONOPLANE ANNIHILATOR	226
9.3.1	Mono-plane deconvolution	227
9.3.2	Monoplanes in local windows	227
9.3.3	Crossing dips	227
9.3.4	Tests of 2-D LOMOPLAN on field data	229
9.4	GRADIENT ALONG THE BEDDING PLANE	230
9.4.1	Definition of LOMOPLAN in 3-D	231
9.4.2	The quarterdome 3-D synthetic (qdome)	231
9.5	3-D SPECTRAL FACTORIZATION	232
10	Some research examples	235
10.1	GULF OF MEXICO CUBE	235

Preface

The difference between theory and practice is smaller in theory than it is in practice. –folklore

We make discoveries about reality by examining the discrepancy between theory and practice. There is a well-developed theory about the difference between theory and practice, and it is called “geophysical inverse theory”. In this book we investigate the practice of the difference between theory and practice. As the folklore tells us, there is a big difference. There are already many books on the theory, and often as not, they end in only one or a few applications in the author’s specialty. In this book on practice, we examine data and results from many diverse applications. I have adopted the discipline of suppressing theoretical curiosities until I find data that requires it (except for a few concepts at chapter ends).

Books on geophysical inverse theory tend to address theoretical topics that are little used in practice. Foremost is probability theory. In practice, probabilities are neither observed nor derived from observations. For more than a handful of variables, it would not be practical to display joint probabilities, even if we had them. If you are data poor, you might turn to probabilities. If you are data rich, you have far too many more rewarding things to do. When you estimate a few values, you ask about their standard deviations. When you have an image making machine, you turn the knobs and make new images (and invent new knobs). Another theory not needed here is singular-value decomposition.

In writing a book on the “practice of the difference between theory and practice” there is no worry to be bogged down in the details of diverse specializations because the geophysical world has many interesting data sets that are easily analyzed with elementary physics and simple geometry. (My specialization, reflection seismic imaging, has a great many less easily explained applications too.) We find here many applications that have a great deal in common with one another, and that commonality is not a part of common inverse theory. Many applications draw our attention to the importance of two weighting functions (one required for data space and the other for model space). Solutions depend strongly on these weighting functions (eigenvalues do too!). Where do these functions come from, from what rationale or estimation procedure? We’ll see many examples here, and find that these functions are not merely weights but filters. Even deeper, they are generally a combination of weights and filters. We do some tricky bookkeeping and bootstrapping when we filter the multidimensional neighborhood of missing and/or suspicious data.

Are you aged 23? If so, this book is designed for you. Life has its discontinuities: when you enter school at age 5, when you marry, when you leave university, when you retire.

The discontinuity at age 23, mid graduate school, is when the world loses interest in your potential to learn. Instead the world wants to know what you are accomplishing right now! This book is about how to make images. It is theory and programs that you can use right now.

This book is not devoid of theory and abstraction. Indeed it makes an important new contribution to the theory (and practice) of data analysis: multidimensional autoregression via the helical coordinate system.

The biggest chore in the study of “the practice of the difference between theory and practice” is that we must look at algorithms. Some of them are short and sweet, but other important algorithms are complicated and ugly in any language. This book can be printed without the computer programs and their surrounding paragraphs, or you can read it without them. I suggest, however, you take a few moments to try to read each program. If you can write in any computer language, you should be able to read these programs well enough to grasp the concept of each, to understand what goes in and what should come out. I have chosen the computer language (more on this later) that I believe is best suited for our journey through the “elementary” examples in geophysical image estimation.

Besides the tutorial value of the programs, if you can read them, you will know exactly how the many interesting illustrations in this book were computed so you will be well equipped to move forward in your own direction.

THANKS

2004 is my twelfth year of working on this book and much of it comes from earlier work and the experience of four previous books. In this book, as in my previous books, I owe a great deal to the many students at the Stanford Exploration Project. I would like to mention some with particularly notable contributions (in approximate historical order).

The concept of this book began along with the PhD thesis of Jeff Thorson. Before that, we imagers thought of our field as “an hoc collection of good ideas” instead of as “adjoints of forward problems”. Bill Harlan understood most of the preconditioning issues long before I did. All of us have a longstanding debt to Rick Ottolini who built a cube movie program long before anyone else in the industry had such a blessing.

My first book was built with a typewriter and ancient technologies. In early days each illustration would be prepared without reusing packaged code. In assembling my second book I found I needed to develop common threads and code them only once and make this code systematic and if not idiot proof, then “idiot resistant”. My early attempts to introduce “seplib” were not widely welcomed until Stew Levin rebuilt everything making it much more robust. My second book was typed in the troff text language. I am indebted to Kamal Al-Yahya who not only converted that book to \LaTeX , but who wrote a general-purpose conversion program that became used internationally.

Early days were a total chaos of plot languages. I and all the others at SEP are deeply indebted to Joe Dellinger who starting from work of Dave Hale, produced our internal plot language “vplot” which gave us reproducibiliy and continuity over decades. Now, for example, our plots seamlessly may be directed to postscript (and PDF), Xwindow, or the web. My second book required that illustrations be literally taped onto the sheet containing

the words. All of us benefitted immensely from the work of Steve Cole who converted Joe's vplot language to postscript which was automatically integrated with the text.

When I began my third book I was adapting liberally from earlier work. I began to realize the importance of being able to reproduce any earlier calculation and began building rules and file-naming conventions for "reproducible research". This would have been impossible were it not for Dave Nichols who introduced **cake**, a variant of the UNIX software building program **make**. Martin Karrenbach continued the construction of our invention of "reproducible research" and extended it to producing reproducible research reports on CD-ROM, an idea well ahead of its time. Some projects were fantastic for their time but had the misfortune of not being widely adopted, ultimately becoming unsupportable. In this category was Dave and Martin's implementation **xtex**, a magnificent way of embedding reproducible research in an electronic textbook. When **cake** suffered the same fate as **xtex**, Matthias Schwab saved us from mainstream isolation by bringing our build procedures into the popular GNU world.

Coming to the present textbook I mention Bob Clapp. He made numerous contributions. When Fortran77 was replaced by Fortran90, he rewrote Ratfor. For many years I (and many of us) depended on Ratfor as our interface to Fortran and as a way of presenting uncluttered code. Bob rewrote Ratfor from scratch merging it with other SEP-specific software tools (Sat) making Ratfor90. Bob prepared the interval-velocity examples in this book. Bob also developed most of the "geostat" ideas and examples in this book. Morgan Brown introduced the texture examples that we find so charming. Paul Sava totally revised the book's presentation of least-squares solvers making them more palatable to students and making more honest our claim that in each case the results you see were produced by the code you see.

One name needs to be singled out. Sergey Fomel converted all the examples in this book from my original Fortran 77 to a much needed modern style of Fortran 90. After I discovered the helix idea and its wide-ranging utility, he adapted all the relevant examples in this book to use it. If you read Fomel's programs, you can learn effective application of that 1990's revolution in coding style known as "object orientation."

This electronic book, "Geophysical Exploration by Example," is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version. This book is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details. You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Massachusetts Ave., Cambridge, MA 02139, USA.

Jon Claerbout
April 27, 2004

Chapter 1

Basic operators and adjoints

A great many of the calculations we do in science and engineering are really matrix multiplication in disguise. The first goal of this chapter is to unmask the disguise by showing many examples. Second, we see how the **adjoint** operator (matrix transpose) back projects information from data to the underlying model.

Geophysical modeling calculations generally use linear operators that predict data from models. Our usual task is to find the inverse of these calculations; i.e., to find models (or make images) from the data. Logically, the adjoint is the first step and a part of all subsequent steps in this **inversion** process. Surprisingly, in practice the adjoint sometimes does a better job than the inverse! This is because the adjoint operator tolerates imperfections in the data and does not demand that the data provide full information.

Using the methods of this chapter, you will find that once you grasp the relationship between operators in general and their adjoints, you can obtain the adjoint just as soon as you have learned how to code the modeling operator.

If you will permit me a poet's license with words, I will offer you the following table of **operators** and their **adjoints**:

matrix multiply	conjugate-transpose matrix multiply
convolve	crosscorrelate
truncate	zero pad
replicate, scatter, spray	sum or stack
spray into neighborhoods	sum within bins
derivative (slope)	negative derivative
causal integration	anticausal integration
add functions	do integrals
assignment statements	added terms
plane-wave superposition	slant stack / beam form
superpose curves	sum along a curve
stretch	squeeze
scalar field gradient	negative of vector field divergence
upward continue	downward continue
diffraction modeling	imaging by migration
hyperbola modeling	CDP stacking
ray tracing	tomography

The left column above is often called “**modeling**,” and the adjoint operators on the right are often used in “data **processing**.”

When the adjoint operator is *not* an adequate approximation to the inverse, then you apply the techniques of fitting and optimization explained in Chapter 2. These techniques require iterative use of the modeling operator and its adjoint.

The adjoint operator is sometimes called the “**back projection**” operator because information propagated in one direction (earth to data) is projected backward (data to earth model). Using complex-valued operators, the transpose and complex conjugate go together; and in **Fourier analysis**, taking the complex conjugate of $\exp(i\omega t)$ reverses the sense of time. With more poetic license, I say that adjoint operators *undo* the time and phase shifts of modeling operators. The inverse operator does this too, but it also divides out the color. For example, when linear interpolation is done, then high frequencies are smoothed out, so inverse interpolation must restore them. You can imagine the possibilities for noise amplification. That is why adjoints are safer than inverses. But nature determines in each application what is the best operator to use, and whether to stop after the adjoint, to go the whole way to the inverse, or to stop partway.

The operators and adjoints above transform vectors to other vectors. They also transform data planes to model planes, volumes, etc. A mathematical operator transforms an “abstract vector” which might be packed full of volumes of information like television signals (time series) can pack together a movie, a sequence of frames. We can always think of the operator as being a matrix but the matrix can be truly huge (and nearly empty). When the vectors transformed by the matrices are large like geophysical data set sizes then the matrix sizes are “large squared,” far too big for computers. Thus although we can always think of an operator as a matrix, in practice, we handle an operator differently. Each practical application requires the practitioner to prepare two computer programs. One performs the matrix multiply $\mathbf{y} = \mathbf{A}\mathbf{x}$ and another multiplies by the transpose $\tilde{\mathbf{x}} = \mathbf{A}'\mathbf{y}$ (without ever having the matrix itself in memory). It is always easy to transpose a matrix. It is less easy to take a computer program that does $\mathbf{y} = \mathbf{A}\mathbf{x}$ and convert it to another to do $\tilde{\mathbf{x}} = \mathbf{A}'\mathbf{y}$. In this chapter are many examples of increasing complexity. At the end of the

chapter we will see a test for any program pair to see whether the operators \mathbf{A} and \mathbf{A}' are mutually adjoint as they should be. Doing the job correctly (coding adjoints without making approximations) will reward us later when we tackle model and image estimation problems.

1.0.1 Programming linear operators

The operation $y_i = \sum_j b_{ij}x_j$ is the multiplication of a matrix \mathbf{B} by a vector \mathbf{x} . The adjoint operation is $\tilde{x}_j = \sum_i b_{ij}y_i$. The operation adjoint to multiplication by a matrix is multiplication by the transposed matrix (unless the matrix has complex elements, in which case we need the complex-conjugated transpose). The following **pseudocode** does matrix multiplication $\mathbf{y} = \mathbf{B}\mathbf{x}$ and multiplication by the transpose $\tilde{\mathbf{x}} = \mathbf{B}'\mathbf{y}$:

```

if adjoint
    then erase x
if operator itself
    then erase y
do iy = 1, ny {
do ix = 1, nx {
    if adjoint
        x(ix) = x(ix) + b(iy,ix) × y(iy)
    if operator itself
        y(iy) = y(iy) + b(iy,ix) × x(ix)
    }}

```

Notice that the “bottom line” in the program is that x and y are simply interchanged. The above example is a prototype of many to follow, so observe carefully the similarities and differences between the adjoint and the operator itself.

Next we restate the matrix-multiply pseudo code in real code.

The module `matmult` for matrix multiply and its adjoint exhibits the style that we will use repeatedly. At last count there were 53 such routines (operator with adjoint) in this book alone.

We have a module with two entries, one named `_init` for the physical scientist who defines the physical problem by defining the matrix, and another named `_lop` for the least-squares problem solver, the computer scientist who will not be interested in how we specify \mathbf{B} , but who will be iteratively computing $\mathbf{B}\mathbf{x}$ and $\mathbf{B}'\mathbf{y}$ to optimize the model fitting.

To use `matmult`, two calls must be made, the first one

```
matmult_init( bb);
```

is done by the physical scientist after he or she has prepared the matrix. Most later calls will be done by numerical analysts in solving code like in Chapter 2. These calls look like

```
matmult_lop( adj, add, nx, ny, x, y);
```

```

                                user/fomels/matmult.c
33 void matmult_lop (bool adj, bool add,
34                   int nx, int ny, float* x, float*y)
35 /*< linear operator >*/
36 {
37     int ix, iy;
38     sf_adjnull (adj, add, nx, ny, x, y);
39     for (ix = 0; ix < nx; ix++) {
40         for (iy = 0; iy < ny; iy++) {
41             if (adj) x[ix] += B[iy][ix] * y[iy];
42             else    y[iy] += B[iy][ix] * x[ix];
43         }
44     }
45 }

```

where **adj** is the logical variable saying whether we desire the adjoint or the operator itself, and where **add** is a logical variable saying whether we want to accumulate like $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{B}\mathbf{x}$ or whether we want to erase first and thus do $\mathbf{y} \leftarrow \mathbf{B}\mathbf{x}$.

We split operators into two independent processes, the first is used for geophysical set up while the second is invoked by mathematical library code (introduced in the next chapter) to find the model that best fits the data. Here is why we do so. It is important that the math code contain nothing about the geophysical particulars. This enables us to use the same math code on many different geophysical problems. This concept of “information hiding” arrived late in human understanding of what is desirable in a computer language. Subroutines and functions are the way that new programs use old ones. Object modules are the way that old programs (math solvers) are able to use new ones (geophysical operators).

1.1 FAMILIAR OPERATORS

The simplest and most fundamental linear operators arise when a matrix operator reduces to a simple row or a column.

A **row** is a summation operation.

A **column** is an impulse response.

If the inner loop of a matrix multiply ranges within a

row, the operator is called *sum* or *pull*.

column, the operator is called *spray* or *push*.

A basic aspect of adjointness is that the adjoint of a row matrix operator is a column

matrix operator. For example, the row operator $[a, b]$

$$y = [a \ b] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = ax_1 + bx_2 \quad (1.1)$$

has an adjoint that is two assignments:

$$\begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix} y \quad (1.2)$$

The adjoint of a sum of N terms is a collection of N assignments.

1.1.1 Adjoint derivative

In numerical analysis we represent the derivative a time function by a finite difference. We do this by subtracting each two neighboring time points and then dividing by the sample interval Δt . This amounts to convolution with the filter $(1, -1)/\Delta t$. Omitting the Δt we express this concept as:


$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} -1 & 1 & . & . & . & . \\ . & -1 & 1 & . & . & . \\ . & . & -1 & 1 & . & . \\ . & . & . & -1 & 1 & . \\ . & . & . & . & -1 & 1 \\ . & . & . & . & . & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} \quad (1.3)$$

The **filter impulse response** is seen in any column in the middle of the matrix, namely $(1, -1)$. In the transposed matrix, the filter-impulse response is time-reversed to $(-1, 1)$. So, mathematically, we can say that the adjoint of the time derivative operation is the negative time derivative. This corresponds also to the fact that the complex conjugate of $-i\omega$ is $i\omega$. We can also speak of the adjoint of the boundary conditions: we might say that the adjoint of “no boundary condition” is a “specified value” boundary condition. The last row in equation (1.3) is optional. It may seem unnatural to append a null row, but it can be a small convenience (when plotting) to have the input and output be the same size.


Equation (1.3) is implemented by the code in module `igrad1` which does the operator itself (the forward operator) and its adjoint.

The adjoint code may seem strange. It might seem more natural to code the adjoint to be the negative of the operator itself and then make the special adjustments for the boundaries. The code given, however, is correct and requires no adjustments at the ends. To see why, notice for each value of i , the operator itself handles one row of equation (1.3) while for each i the adjoint handles one column. That’s why coding the adjoint in this way does not require any special work on the ends. The present method of coding reminds us that the adjoint of a sum of N terms is a collection of N assignments.

Figure 1.1 illustrates the use of module `igrad1` for each north-south line of a topographic map. We observe that the gradient gives an impression of illumination from a low sun angle.



Topographic map, Stanford area



Southward slope

Figure 1.1: Topography near Stanford (top) southward slope (bottom).

ajt/igrad1 stangrad

```
filt/lib/igrad1.c
```

```

25 void sf_igrad1_lop (bool adj, bool add,
26                    int nx, int ny, float *xx, float *yy)
27 /*< linear operator >*/
28 {
29     int i;
30
31     sf_adjnull (adj, add, nx, ny, xx, yy);
32     for (i=0; i < nx-1; i++) {
33         if (adj) {
34             xx[i+1] += yy[i];
35             xx[i] -= yy[i];
36         } else {
37             yy[i] += xx[i+1] - xx[i];
38         }
39     }
40 }

```

To apply `igrad1` along the 1-axis for each point on the 2-axis of a two-dimensional map, we use the loop

```

for (iy=0; iy < ny; iy++)
    igrad1_lop(adj, add, nx, nx, map[iy], ruf[iy])

```

1.1.2 Transient convolution

The next operator we examine is convolution. It arises in many applications; and it could be derived in many ways. A basic derivation is from the multiplication of two polynomials, say $X(Z) = x_1 + x_2Z + x_3Z^2 + x_4Z^3 + x_5Z^4 + x_6Z^5$ times $B(Z) = b_1 + b_2Z + b_3Z^2 + b_4Z^3$.¹ Identifying the k -th power of Z in the product $Y(Z) = B(Z)X(Z)$ gives the k -th row of the convolution transformation (1.4).

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix} = \begin{bmatrix} b_1 & 0 & 0 & 0 & 0 & 0 \\ b_2 & b_1 & 0 & 0 & 0 & 0 \\ b_3 & b_2 & b_1 & 0 & 0 & 0 \\ 0 & b_3 & b_2 & b_1 & 0 & 0 \\ 0 & 0 & b_3 & b_2 & b_1 & 0 \\ 0 & 0 & 0 & b_3 & b_2 & b_1 \\ 0 & 0 & 0 & 0 & b_3 & b_2 \\ 0 & 0 & 0 & 0 & 0 & b_3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \mathbf{B}\mathbf{x} \quad (1.4)$$

Notice that columns of equation (1.4) all contain the same signal, but with different shifts. This signal is called the filter's impulse response.

¹ This book is more involved with matrices than with Fourier analysis. If it were more Fourier analysis we would choose notation to begin subscripts from zero like this: $B(Z) = b_0 + b_1Z + b_2Z^2 + b_3Z^3$.

Equation (1.4) could be rewritten as

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix} = \begin{bmatrix} x_1 & 0 & 0 \\ x_2 & x_1 & 0 \\ x_3 & x_2 & x_1 \\ x_4 & x_3 & x_2 \\ x_5 & x_4 & x_3 \\ x_6 & x_5 & x_4 \\ 0 & x_6 & x_5 \\ 0 & 0 & x_6 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} = \mathbf{X}\mathbf{b} \quad (1.5)$$

In applications we can choose between $\mathbf{y} = \mathbf{X}\mathbf{b}$ and $\mathbf{y} = \mathbf{B}\mathbf{x}$. In one case the output \mathbf{y} is dual to the filter \mathbf{b} , and in the other case the output \mathbf{y} is dual to the input \mathbf{x} . Sometimes we must solve for \mathbf{b} and sometimes for \mathbf{x} ; so sometimes we use equation (1.5) and sometimes (1.4). Such solutions begin from the adjoints. The adjoint of (1.4) is

$$\begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \\ \hat{x}_3 \\ \hat{x}_4 \\ \hat{x}_5 \\ \hat{x}_6 \end{bmatrix} = \begin{bmatrix} b_1 & b_2 & b_3 & 0 & 0 & 0 & 0 & 0 \\ 0 & b_1 & b_2 & b_3 & 0 & 0 & 0 & 0 \\ 0 & 0 & b_1 & b_2 & b_3 & 0 & 0 & 0 \\ 0 & 0 & 0 & b_1 & b_2 & b_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & b_1 & b_2 & b_3 & 0 \\ 0 & 0 & 0 & 0 & 0 & b_1 & b_2 & b_3 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix} \quad (1.6)$$

The adjoint *crosscorrelates* with the filter instead of convolving with it (because the filter is backwards). Notice that each row in equation (1.6) contains all the filter coefficients and there are no rows where the filter somehow uses zero values off the ends of the data as we saw earlier. In some applications it is important not to assume zero values beyond the interval where inputs are given.

The adjoint of (1.5) crosscorrelates a fixed portion of filter input across a variable portion of filter output.

$$\begin{bmatrix} \hat{b}_1 \\ \hat{b}_2 \\ \hat{b}_3 \end{bmatrix} = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & x_5 & x_6 & 0 & 0 \\ 0 & x_1 & x_2 & x_3 & x_4 & x_5 & x_6 & 0 \\ 0 & 0 & x_1 & x_2 & x_3 & x_4 & x_5 & x_6 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix} \quad (1.7)$$

Module `tcail` is used for $\mathbf{y} = \mathbf{B}\mathbf{x}$ and module `tcaf1` is used for $\mathbf{y} = \mathbf{X}\mathbf{b}$.

The polynomials $X(Z)$, $B(Z)$, and $Y(Z)$ are called Z transforms. An important fact in real life (but not important here) is that the Z transforms are Fourier transforms in disguise. Each polynomial is a sum of terms and the sum amounts to a Fourier sum when we take $Z = e^{i\omega\Delta t}$. The very expression $Y(Z) = B(Z)X(Z)$ says that a product in the frequency domain (Z has a numerical value) is a convolution in the time domain (that's how we multiply polynomials, convolve their coefficients).

user/gee/tcail.c

```

45  for( b=0; b < nb; b++) {
46      for( x=0; x < nx; x++) { y = x + b;
47          if( adj) xx[x] += yy[y] * bb[b];
48          else    yy[y] += xx[x] * bb[b];
49      }
50  }

```

user/gee/tcaf1.c

```

45  for (b=0; b < nb; b++) {
46      for (x=0; x < nx; x++) { y = x + b;
47          if( adj) bb[b] += yy[y] * xx[x];
48          else    yy[y] += bb[b] * xx[x];
49      }
50  }

```

1.1.3 Internal convolution

Convolution is the computational equivalent of ordinary linear differential operators (with constant coefficients). Applications are vast, and end effects are important. Another choice of data handling at ends is that zero data not be assumed beyond the interval where the data is given. This is important in data where the crosscorrelation changes with time. Then it is sometimes handled as constant in short time windows. Care must be taken that zero signal values not be presumed off the ends of those short time windows; otherwise, the many ends of the many short segments can overwhelm the results.

In the sets (1.4) and (1.5), the top two equations explicitly assume that the input data vanishes before the interval on which it is given, and likewise at the bottom. Abandoning the top two and bottom two equations in (1.5) we get:

$$\begin{bmatrix} y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} x_3 & x_2 & x_1 \\ x_4 & x_3 & x_2 \\ x_5 & x_4 & x_3 \\ x_6 & x_5 & x_4 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \quad (1.8)$$

The adjoint is

$$\begin{bmatrix} \hat{b}_1 \\ \hat{b}_2 \\ \hat{b}_3 \end{bmatrix} = \begin{bmatrix} x_3 & x_4 & x_5 & x_6 \\ x_2 & x_3 & x_4 & x_5 \\ x_1 & x_2 & x_3 & x_4 \end{bmatrix} \begin{bmatrix} y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} \quad (1.9)$$

The difference between (1.9) and (1.7) is that here the adjoint crosscorrelates a fixed portion of *output* across a variable portion of *input*, whereas with (1.7) the adjoint crosscorrelates a fixed portion of *input* across a variable portion of *output*.

In practice we typically allocate equal space for input and output. Because the output is shorter than the input, it could slide around in its allocated space, so its location is specified

by an additional parameter called its **lag**. The value of **lag** always used in this book is

```
user/gee/icafl.c
```

```

45  for( b=0; b < nb; b++) {
46      for( y = nb; y <= ny; y++) { x = y - b - 1;
47          if( adj) bb[b] += yy[y-lag] * xx[x];
48          else    yy[y-lag] += bb[b] * xx[x];
49      }
50  }
```


lag=1. For **lag=1** the module **icafl** implements not equation (1.8) but (1.10):

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ x_3 & x_2 & x_1 \\ x_4 & x_3 & x_2 \\ x_5 & x_4 & x_3 \\ x_6 & x_5 & x_4 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \quad (1.10)$$

It may seem a little odd to put the required zeros at the beginning of the output, but filters are generally designed so that their strongest coefficient is the first, namely **bb(1)** so the alignment of input and output in equation (1.10) is the most common one.

The **end effects** of the convolution modules are summarized in Figure 5.6.

Figure 1.2: Example of convolution end-effects. From top to bottom: input; filter; output of **tcaif1()**; output of **icaif1()** also with (**lag=1**).
ajt/conv conv



1.1.4 Zero padding is the transpose of truncation

Surrounding a dataset by zeros (**zero padding**) is adjoint to throwing away the extended data (**truncation**). Let us see why this is so. Set a signal in a vector **x**, and then to make a longer vector **y**, add some zeros at the end of **x**. This zero padding can be regarded as the matrix multiplication

$$\mathbf{y} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \mathbf{x} \quad (1.11)$$

The matrix is simply an identity matrix **I** above a zero matrix **0**. To find the transpose to zero-padding, we now transpose the matrix and do another matrix multiply:

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \mathbf{y} \quad (1.12)$$

So the transpose operation to zero padding data is simply *truncating* the data back to its original length. Module `zpad1` below pads zeros on both ends of its input. Modules for two- and three-dimensional padding are in the library named `zpad2()` and `zpad3()`.

```
user/gee/zpad1.c
```

27	for (d=0; d < nd; d++) {
28	p = d + (np-nd)/2;
29	if (adj) data[d] += padd[p];
30	else padd[p] += data[d];
31	}

1.1.5 Adjoint of products are reverse-ordered products of adjoints

Here we examine an example of the general idea that adjoints of products are reverse-ordered products of adjoints. For this example we use the Fourier transformation. No details of **Fourier transformation** are given here and we merely use it as an example of a square matrix \mathbf{F} . We denote the complex-conjugate transpose (or **adjoint**) matrix with a prime, i.e., \mathbf{F}' . The adjoint arises naturally whenever we consider energy. The statement that Fourier transforms conserve energy is $\mathbf{y}'\mathbf{y} = \mathbf{x}'\mathbf{x}$ where $\mathbf{y} = \mathbf{F}\mathbf{x}$. Substituting gives $\mathbf{F}'\mathbf{F} = \mathbf{I}$, which shows that the inverse matrix to Fourier transform happens to be the complex conjugate of the transpose of \mathbf{F} .

With Fourier transforms, **zero padding** and **truncation** are especially prevalent. Most modules transform a dataset of length of 2^n , whereas dataset lengths are often of length $m \times 100$. The practical approach is therefore to pad given data with zeros. Padding followed by Fourier transformation \mathbf{F} can be expressed in matrix algebra as

$$\text{Program} = \mathbf{F} \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \quad (1.13)$$

According to matrix algebra, the transpose of a product, say $\mathbf{AB} = \mathbf{C}$, is the product $\mathbf{C}' = \mathbf{B}'\mathbf{A}'$ in reverse order. So the adjoint routine is given by

$$\text{Program}' = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \mathbf{F}' \quad (1.14)$$

Thus the adjoint routine *truncates* the data *after* the inverse Fourier transform. This concrete example illustrates that common sense often represents the mathematical abstraction that adjoints of products are reverse-ordered products of adjoints. It is also nice to see a formal mathematical notation for a practical necessity. Making an approximation need not lead to collapse of all precise analysis.

1.1.6 Nearest-neighbor coordinates

In describing physical processes, we often either specify models as values given on a uniform mesh or we record data on a uniform mesh. Typically we have a function f of time t or depth z and we represent it by $\mathbf{f}(\mathbf{iz})$ corresponding to $f(z_i)$ for $i = 1, 2, 3, \dots, n_z$ where

$z_i = z_0 + (i - 1)\Delta z$. We sometimes need to handle depth as an integer counting variable i and we sometimes need to handle it as a floating-point variable z . Conversion from the counting variable to the floating-point variable is exact and is often seen in a computer idiom such as either of

```
for (iz=0; iz < nz; iz++) {    z = z0 + iz * dz;
for (i3=0, i3 < n3; i3++) {    x3 = o3 + i3 * d3;
```

The reverse conversion from the floating-point variable to the counting variable is inexact. The easiest thing is to place it at the nearest neighbor. This is done by solving for iz , then adding one half, and then rounding down to the nearest integer. The familiar computer idioms are:

```
iz = 0.5 + ( z - z0) / dz;
i3 = 0.5 + (x3 - o3) / d3;
```

A small warning is in order: People generally use positive counting variables. If you also include negative ones, then to get the nearest integer, you should do your rounding with the C function `floor()`.

1.1.7 Data-push binning

Binning is putting data values in bins. Nearest-neighbor binning is an operator. There is both a forward operator and its adjoint. Normally the model consists of values given on a uniform mesh, and the data consists of pairs of numbers (ordinates at coordinates) sprinkled around in the continuum (although sometimes the data is uniformly spaced and the model is not).

In both the forward and the adjoint operation, each data coordinate is examined and the nearest mesh point (the bin) is found. For the forward operator, the value of the bin is added to that of the data. The adjoint is the reverse: we add the value of the data to that of the bin. Both are shown in two dimensions in module `bin2`. The most typical

user/gee/bin2.c

```

46     for (id=0; id < nd; id++) {
47         i1 = 0.5 + (xy[0][id]-o1)/d1;
48         i2 = 0.5 + (xy[1][id]-o2)/d2;
49         if (0<=i1 && i1<m1 &&
50             0<=i2 && i2<m2) {
51             im = i1+i2*m1;
52             if (adj) mm[im] += dd[id];
53             else    dd[id] += mm[im];
54         }
55     }
```

application requires an additional step, inversion. In the inversion applications each bin

contains a different number of data values. After the adjoint operation is performed, the inverse operator divides the bin value by the number of points in the bin. It is this inversion operator that is generally called binning. To find the number of data points in a bin, we can simply apply the adjoint of `bin2` to pseudo data of all ones. To capture this idea in an equation, let \mathbf{B} denote the linear operator in which the bin value is sprayed to the data values. The inverse operation, in which the data values in the bin are summed and divided by the number in the bin, is represented by

$$\mathbf{m} = \text{diag}(\mathbf{B}'\mathbf{1})^{-1}\mathbf{B}'\mathbf{d} \quad (1.15)$$

Empty bins, of course, leave us a problem. That we'll address in chapter 3. In Figure 1.3, the empty bins contain zero values.




Figure 1.3: Binned depths of the Sea of Galilee. ajt/galilee/galbin


1.1.8 Linear interpolation

The **linear interpolation** operator is much like the binning operator but a little fancier. When we perform the forward operation, we take each data coordinate and see which two model bin centers bracket it. Then we pick up the two bracketing model values and weight each of them in proportion to their nearness to the data coordinate, and add them to get the data value (ordinate). The adjoint operation is adding a data value back into the model vector; using the same two weights, the adjoint distributes the data ordinate value between the two nearest bins in the model vector. For example, suppose we have a data point near

each end of the model and a third data point exactly in the middle. Then for a model space 6 points long, as shown in Figure 1.4, we have the operator in (1.16).

Figure 1.4: Uniformly sampled model space and irregularly sampled data space corresponding to (1.16).

ajt/XFig helgerud



$$\begin{bmatrix} d_0 \\ d_1 \\ d_2 \end{bmatrix} \approx \begin{bmatrix} .7 & .3 & . & . & . & . \\ . & . & 1 & . & . & . \\ . & . & . & . & .5 & .5 \end{bmatrix} \begin{bmatrix} m_0 \\ m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \end{bmatrix} \quad (1.16)$$

The two weights in each row sum to unity. If a binning operator were used for the same data and model, the binning operator would contain a “1.” in each row. In one dimension (as here), data coordinates are often sorted into sequence, so that the matrix is crudely a diagonal matrix like equation (1.16). If the data coordinates covered the model space uniformly, the adjoint would roughly be the inverse. Otherwise, when data values pile up in some places and gaps remain elsewhere, the adjoint would be far from the inverse.

Module `lint1` does linear interpolation and its adjoint. In chapters 3 and 6 we build inverse operators.

user/gee/lint1.c

```

45  for (id=0; id < nd; id++) {
46      f = (coord[id]-o1)/d1;
47      im=floorf(f);
48      if (0 <= im && im < nm-1) {
49          fx=f-im;
50          gx=1.-fx;
51
52          if(adj) {
53              mm[im] += gx * dd[id];
54              mm[im+1] += fx * dd[id];
55          } else {
56              dd[id] += gx * mm[im] + fx * mm[im+1];
57          }
58      }
59  }
```

1.1.9 Spray and sum : scatter and gather

Perhaps the most common operation is the summing of many values to get one value. Its adjoint operation takes a single input value and throws it out to a space of many values. The **summation operator** is a row vector of ones. Its adjoint is a column vector of ones. In one dimension this operator is almost too easy for us to bother showing a routine. But it is more interesting in three dimensions, where we could be summing or spraying on any of three subscripts, or even summing on some and spraying on others. In module **spraysum**, both input and output are taken to be three-dimensional arrays. Externally, however, either could be a scalar, vector, plane, or cube. For example, the internal array **xx(n1,1,n3)** could be externally the matrix **map(n1,n3)**. When module **spraysum** is given the input dimensions and output dimensions stated below, the operations stated alongside are implied.

(n1,n2,n3)	(1,1,1)	Sum a cube into a value.
(1,1,1)	(n1,n2,n3)	Spray a value into a cube.
(n1,1,1)	(n1,n2,1)	Spray a column into a matrix.
(1,n2,1)	(n1,n2,1)	Spray a row into a matrix.
(n1,n2,1)	(n1,n2,n3)	Spray a plane into a cube.
(n1,n2,1)	(n1,1,1)	Sum rows of a matrix into a column.
(n1,n2,1)	(1,n2,1)	Sum columns of a matrix into a row.
(n1,n2,n3)	(n1,n2,n3)	Copy and add the whole cube.

If an axis is not of unit length on either input or output, then both lengths must be the same; otherwise, there is an error. Normally, after (possibly) erasing the output, we simply loop over all points on each axis, adding the input to the output. This implements either a copy or an add, depending on the **add** parameter. It is either a spray, a sum, or a copy, according to the specified axis lengths.

user/gee/spraysum.c

```

42  for (i3=0; i3 < SF_MAX(n3,m3); i3++) {
43      x = SF_MIN(i3 ,n3-1);
44      y = SF_MIN(i3 ,m3-1);
45      for (i2=0; i2 < SF_MAX(n2,m2); i2++) {
46          x = x*n2 + SF_MIN(i2 ,n2-1);
47          y = y*m2 + SF_MIN(i2 ,m2-1);
48          for (i1=0; i1 < SF_MAX(n1,m1); i1++) {
49              x = x*n1 + SF_MIN(i1 ,n1-1);
50              y = y*m1 + SF_MIN(i1 ,m1-1);
51
52              if( adj )  xx[x]  +=  yy[y];
53              else      yy[y]  +=  xx[x];
54          }
55      }
56  }
```

1.1.10 Causal and leaky integration

Causal integration is defined as

$$y(t) = \int_{-\infty}^t x(\tau) d\tau \quad (1.17)$$

Leaky integration is defined as

$$y(t) = \int_0^{\infty} x(t - \tau) e^{-\alpha\tau} d\tau \quad (1.18)$$

As $\alpha \rightarrow 0$, leaky integration becomes causal integration. The word “leaky” comes from electrical circuit theory where the voltage on a capacitor would be the integral of the current if the capacitor did not leak electrons.

Sampling the time axis gives a matrix equation that we should call causal summation, but we often call it causal integration. Equation (1.19) represents causal integration for $\rho = 1$ and leaky integration for $0 < \rho < 1$.

$$\mathbf{y} = \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \rho & 1 & 0 & 0 & 0 & 0 & 0 \\ \rho^2 & \rho & 1 & 0 & 0 & 0 & 0 \\ \rho^3 & \rho^2 & \rho & 1 & 0 & 0 & 0 \\ \rho^4 & \rho^3 & \rho^2 & \rho & 1 & 0 & 0 \\ \rho^5 & \rho^4 & \rho^3 & \rho^2 & \rho & 1 & 0 \\ \rho^6 & \rho^5 & \rho^4 & \rho^3 & \rho^2 & \rho & 1 \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \mathbf{C}\mathbf{x} \quad (1.19)$$

(The discrete world is related to the continuous by $\rho = e^{-\alpha\Delta\tau}$ and in some applications, the diagonal is 1/2 instead of 1.) Causal integration is the simplest prototype of a recursive operator. The coding is trickier than that for the operators we considered earlier. Notice when you compute y_5 that it is the sum of 6 terms, but that this sum is more quickly computed as $y_5 = \rho y_4 + x_5$. Thus equation (1.19) is more efficiently thought of as the recursion

$$y_t = \rho y_{t-1} + x_t \quad t \text{ increasing} \quad (1.20)$$

(which may also be regarded as a numerical representation of the **differential equation** $dy/dt + y(1 - \rho)/\Delta t = x(t)$.)

When it comes time to think about the adjoint, however, it is easier to think of equation (1.19) than of (1.20). Let the matrix of equation (1.19) be called \mathbf{C} . Transposing to get \mathbf{C}' and applying it to \mathbf{y} gives us something back in the space of \mathbf{x} , namely $\tilde{\mathbf{x}} = \mathbf{C}'\mathbf{y}$. From it we see that the adjoint calculation, if done recursively, needs to be done backwards, as in

$$\tilde{x}_{t-1} = \rho \tilde{x}_t + y_{t-1} \quad t \text{ decreasing} \quad (1.21)$$

Thus the adjoint of causal integration is **anticausal integration**.

A module to do these jobs is `leakint`. The code for anticausal integration is not obvious from the code for integration and the adjoint coding tricks we learned earlier. To understand


filt/lib/causint.c

```

35     t = 0.;
36     if (adj) {
37         for (i=nx-1; i >= 0; i--) {
38             t += yy[i];
39             xx[i] += t;
40         }
41     } else {
42         for (i=0; i <= nx-1; i++) {
43             t += xx[i];
44             yy[i] += t;
45         }
46     }

```

Figure 1.5: `in1` is an input pulse. `C in1` is its causal integral. `C' in1` is the anticausal integral of the pulse. `in2` is a separated doublet. Its causal integration is a box and its anticausal integration is a negative box. `CC in2` is the double causal integral of `in2`. How can an equilateral triangle be built? ajt/causint causint



the adjoint, you need to inspect the detailed form of the expression $\tilde{\mathbf{x}} = \mathbf{C}'\mathbf{y}$ and take care to get the ends correct. Figure 1.5 illustrates the program for $\rho = 1$.

Later we will consider equations to march wavefields up towards the earth's surface, a layer at a time, an operator for each layer. Then the adjoint will start from the earth's surface and march down, a layer at a time, into the earth.

EXERCISES:

- 1 Consider the matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \rho & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (1.22)$$

and others like it with ρ in other locations. Show what combination of these matrices will represent the leaky integration matrix in equation (1.19). What is the adjoint?

- 2 Modify the calculation in Figure 1.5 so that there is a triangle waveform on the bottom row.
- 3 Notice that the triangle waveform is not time aligned with the input `in2`. Force time alignment with the operator $\mathbf{C}'\mathbf{C}$ or $\mathbf{C}\mathbf{C}'$.
- 4 Modify `causint` on page 17 by changing the diagonal to contain 1/2 instead of 1. Notice how time alignment changes in Figure 1.5.

1.1.11 Backsolving, polynomial division and deconvolution

Ordinary differential equations often lead us to the backsolving operator. For example, the damped harmonic oscillator leads to a special case of equation (1.23) where $(a_3, a_4, \dots) = 0$. There is a huge literature on finite-difference solutions of ordinary differential equations that lead to equations of this type. Rather than derive such an equation on the basis of many possible physical arrangements, we can begin from the filter transformation \mathbf{B} in (1.4) but put the matrix on the other side of the equation so our transformation can be called one of inversion or backsubstitution. Let us also force the matrix \mathbf{B} to be a square matrix by truncating it with $\mathbf{T} = [\mathbf{I} \ 0]$, say $\mathbf{A} = [\mathbf{I} \ 0]\mathbf{B} = \mathbf{T}\mathbf{B}$. To link up with applications in later chapters, I specialize to 1's on the main diagonal and insert some bands of zeros.

$$\mathbf{A}\mathbf{y} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ a_1 & 1 & 0 & 0 & 0 & 0 & 0 \\ a_2 & a_1 & 1 & 0 & 0 & 0 & 0 \\ 0 & a_2 & a_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & a_2 & a_1 & 1 & 0 & 0 \\ a_5 & 0 & 0 & a_2 & a_1 & 1 & 0 \\ 0 & a_5 & 0 & 0 & a_2 & a_1 & 1 \end{bmatrix} \begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \mathbf{x} \quad (1.23)$$

Algebraically, this operator goes under the various names, “backsolving”, “polynomial division”, and “deconvolution”. The leaky integration transformation (1.19) is a simple example of backsolving when $a_1 = -\rho$ and $a_2 = a_5 = 0$. To confirm this, you need to verify that the matrices in (1.23) and (1.19) are mutually inverse.

A typical row in equation (1.23) says

$$x_t = y_t + \sum_{\tau>0} a_\tau y_{t-\tau} \quad (1.24)$$

Change the signs of all terms in equation (1.24) and move some terms to the opposite side

$$y_t = x_t - \sum_{\tau>0} a_\tau y_{t-\tau} \quad (1.25)$$

Equation (1.25) is a recursion to find y_t from the values of y at earlier times.

In the same way that equation (1.4) can be interpreted as $Y(Z) = B(Z)X(Z)$, equation (1.23) can be interpreted as $A(Z)Y(Z) = X(Z)$ which amounts to $Y(Z) = X(Z)/A(Z)$. Thus, convolution amounts to polynomial multiplication while the backsubstitution we are doing here is called deconvolution, and it amounts to polynomial division.

A causal operator is one that uses its present and past inputs to make its current output. Anticausal operators use the future but not the past. Causal operators are generally associated with lower triangular matrices and positive powers of Z , whereas anticausal operators are associated with upper triangular matrices and negative powers of Z . A transformation like equation (1.23) but with the transposed matrix would require us to run the recursive solution the opposite direction in time, as we did with leaky integration.

A module to backsolve (1.23) is `recfilt`.

The more complicated an operator, the more complicated is its adjoint. Given a transformation from \mathbf{x} to \mathbf{y} that is $\mathbf{TB}\mathbf{y} = \mathbf{x}$, we may wonder if the adjoint transform really is $(\mathbf{TB})'\hat{\mathbf{x}} = \mathbf{y}$. It amounts to asking if the adjoint of $\mathbf{y} = (\mathbf{TB})^{-1}\mathbf{x}$ is $\hat{\mathbf{x}} = ((\mathbf{TB})')^{-1}\mathbf{y}$. Mathematically we are asking if the inverse of a transpose is the transpose of the inverse. This is so because in $\mathbf{AA}^{-1} = \mathbf{I} = \mathbf{I}' = (\mathbf{A}^{-1})'\mathbf{A}'$ the parenthesized object must be the inverse of its neighbor \mathbf{A}' .

The adjoint has a meaning which is nonphysical. It is like the forward operator except that we must begin at the final time and revert towards the first. The adjoint pendulum damps as we compute it backward in time, but that, of course, means that the adjoint pendulum diverges as it is viewed moving forward in time.

1.1.12 The basic low-cut filter

Many geophysical measurements contain very low-frequency noise called “drift.” For example, it might take some months to survey the depth of a lake. Meanwhile, rainfall or evaporation could change the lake level so that new survey lines become inconsistent with old ones. Likewise, gravimeters are sensitive to atmospheric pressure, which changes with the weather. A magnetic survey of an archeological site would need to contend with the fact that the earth’s main magnetic field is changing randomly through time while the survey is being done. Such noises are sometimes called “secular noise.”

flt/lib/recflt.c

```

48
49  for (ix=0; ix < nx; ix++) {
50      tt[ix] = 0.;
51  }
52
53  if (adj) {
54      for (ix = nx-1; ix >= 0; ix--) {
55          tt[ix] = yy[ix];
56          for (ia = 0; ia < SF_MIN(na,ny-ix-1); ia++) {
57              iy = ix + ia + 1;
58              tt[ix] -= aa[ia] * tt[iy];
59          }
60      }
61      for (ix=0; ix < nx; ix++) {
62          xx[ix] += tt[ix];
63      }
64  } else {
65      for (iy = 0; iy < ny; iy++) {
66          tt[iy] = xx[iy];
67          for (ia = 0; ia < SF_MIN(na,iy); ia++) {
68              ix = iy - ia - 1;
69              tt[iy] -= aa[ia] * tt[ix];
70          }
71      }
72      for (iy=0; iy < ny; iy++) {
73          yy[iy] += tt[iy];
74      }

```

The simplest way to eliminate low frequency noise is to take a time derivative. A disadvantage is that the derivative changes the waveform from a pulse to a doublet (finite difference). Here we examine the most basic low-cut filter. It preserves the waveform at high frequencies; it has an adjustable parameter for choosing the bandwidth of the low cut; and it is causal (uses the past but not the future).

We make our causal lowcut filter (highpass filter) by two stages which can be done in either order.

1. Apply a time derivative, actually a finite difference, convolving the data with $(1, -1)$.
2. Integrate, actually to do a leaky integration, to deconvolve with $(1, -\rho)$ where numerically, ρ is slightly less than unity.

The convolution ensures that the zero frequency is removed. The leaky integration almost undoes the differentiation (but does not restore the zero frequency). Adjusting the numerical value of ρ adjusts the cutoff frequency of the filter. To learn the impulse response of the combined processes, we need to convolve the finite difference $(1, -1)$ with the leaky integration $(1, \rho, \rho^2, \rho^3, \rho^4, \dots)$. The result is $(1, \rho, \rho^2, \rho^3, \rho^4, \dots)$ minus $(0, 1, \rho, \rho^2, \rho^3, \dots)$. We can think of this as $(1, 0, 0, 0, 0, \dots)$ minus $(1 - \rho)(1, \rho, \rho^2, \rho^3, \dots)$. In other words the impulse response is an impulse followed by the negative of a weak $(1 - \rho)$ decaying exponential ρ^t . Roughly speaking, the cutoff frequency of the filter corresponds to matching one wavelength to the exponential decay time.

Some exercise with Fourier transforms or Z -transforms², shows the Fourier transform of this highpass filter filter to be

$$H(Z) = \frac{1 - Z}{1 - \rho Z} = 1 - (1 - \rho)[Z^1 + \rho Z^2 + \rho^2 Z^3 + \rho^3 Z^4 \dots] \quad (1.26)$$

where the unit-delay operator is $Z = e^{i\omega\Delta t}$ and where ω is the frequency. A symmetrical (noncausal) lowcut filter would filter once forward with $H(Z)$ and once backwards (adjoint) with $H(1/Z)$. This is not the place for a detailed Fourier analysis of this filter but it is the place to mention that a cutoff filter is typically specified by its cutoff frequency, a frequency that separates the pass and reject region. For this filter, the cutoff frequency ω_0 would correspond to matching a quarter wavelength of a sinusoid to the exponential decay length of ρ^k , namely, say the value of k for which $\rho^k \approx 1/2$

Seismological data is more complex. A single “measurement” consists of an explosion and echo signals recorded at many locations. As before, a complete survey is a track (or tracks) of explosion locations. Thus, in seismology, data space is higher dimensional. Its most troublesome noise is not simply low frequency; it is low velocity. We will do more with multidimensional data in later chapters.

EXERCISES:

- 1 Give an analytic expression for the waveform of equation (2.13).
- 2 Define a low-pass filter as $1 - H(Z)$. What is the low-pass impulse response?

² An introduction to Z -transforms is found in my earlier books, FGDP and ESA-PVI.

- 3 Put Galilee data on a coarse mesh. Consider north-south lines as one-dimensional signals. Find the value of ρ for which H is the most pleasing filter.
- 4 Find the value of ρ for which $\bar{H}H$ is the most pleasing filter.
- 5 Find the value of ρ for which H applied to Galilee has minimum energy. (Experiment with a range of about ten values around your favorite value.)
- 6 Find the value of ρ for which $\bar{H}H$ applied to Galilee has minimum energy.
- 7 Repeat above for east-west lines.




Figure 1.6: The depth of the Sea of Galilee after roughening. ajt/galilee galocut


1.1.13 Nearest-neighbor normal moveout (NMO)

Normal-moveout correction (**NMO**) is a geometrical correction of reflection seismic data that stretches the time axis so that data recorded at nonzero separation x_0 of shot and receiver, after stretching, appears to be at $x_0 = 0$. NMO correction is roughly like time-to-depth conversion with the equation $v^2 t^2 = z^2 + x_0^2$. After the data at x_0 is stretched from t to z , it should look like stretched data from any other x (assuming these are plane horizontal reflectors, etc.). In practice, z is not used; rather, **traveltime depth** τ is used, where $\tau = z/v$; so $t^2 = \tau^2 + x_0^2/v^2$. (Because of the limited alphabet of programming languages, I often use the keystroke **z** to denote τ .)

Typically, many receivers record each shot. Each seismogram can be transformed by NMO and the results all added. This is called “**stacking**” or “**NMO stacking**.” The adjoint to this operation is to begin from a model which ideally is the zero-offset trace and spray this model to all offsets. From a matrix viewpoint, stacking is like a *row* vector of normal moveout operators and modeling is like a *column*. An example is shown in Figure 1.7.

Figure 1.7: Hypothetical model, synthetic data, and model image.

ajt/cunha cunha



A module that does reverse moveout is `hypotenusei`. Given a zero-offset trace, it makes another at non-zero offset. The adjoint does the usual normal moveout correction. (My 1992 textbook (PVI) illustrates many additional features of normal moveout.) A companion routine `imospray` loops over offsets and makes a trace for each. The adjoint of `imospray` is the industrial process of moveout and stack.

This is the first time we have built an operator (moveout and stack) from a collection of other operators (moveout at various offsets) and there are new possibilities for confusion and program bugs. As earlier with the matrix multiplication operator, to use the `imospray` operator, there are two steps, one that we use to set things up

```
imospray_init( slow, x0,dx, t0,dt, nt,nx);
```

and another step that we use a lot in the next chapter for analysis and data fitting.

```
imospray_lop( adj, add, n1, n2, stack, gather);
```

Later we’ll see programs that are not operators. Consider the the adjoint of spraying which is stacking. Here the occurrence of the `add=true` in `imospray` assures we do not erase the stack each time we add in another trace.

1.1.14 Coding chains and arrays

With a collection of operators, we can build more elaborate operators. One method is chaining. For example, the operator product $\mathbf{A} = \mathbf{BC}$ is represented in the subroutine `chain2(op1, op2, ...)`. Likewise the operator product $\mathbf{A} = \mathbf{BCD}$ is represented in the in the subroutine `chain3(op1, op2, op3, ...)`. Another way to make more elaborate operators is to put operators in a matrix such as subroutine `array` also in module `chain`.

filt/proc/hypotenusei.c

```

42 void hypotenusei_set(float t0 /* time origin */,
43                    float dt /* time sampling */,
44                    float xs /* offset times slowness */)
45 /*< set up >*/
46 {
47     int it;
48     float t, z2;
49
50     for (it=0; it < nt; it++) {
51         t = t0 + dt*it;
52         z2 = t * t - xs * xs;
53         iz[it] = ( z2 >= 0.)? 0.5 + (sqrtf(z2) - t0) /dt: -1;
54     }
55 }
56
57 void hypotenusei_lop(bool adj, bool add,
58                    int n1, int n2, float *zz, float *tt)
59 /*< linear operator >*/
60 {
61     int it;
62
63     sf_adjnull(adj, add, n1, n2, zz, tt);
64
65     for (it=0; it < nt; it++) {
66         if (iz[it] < 0) continue;
67
68         if (adj)
69             zz[iz[it]] += tt[it];
70         else
71             tt[it] += zz[iz[it]];
72     }
73 }

```

filt/proc/imospray.c

```

47 void imospray_lop(bool adj, bool add, int n1, int n2,
48                  float *stack, float *gather)
49 /*< linear operator >*/
50 {
51     int ix;
52     float x;
53
54     sf_adjnull(adj, add, n1, n2, stack, gather);
55
56     for (ix=0; ix < nx; ix++) {
57         x = x0 + dx*ix;
58
59         hypotenusei_set (t0, dt, x);
60         hypotenusei_lop (adj, true, nt, nt, stack, gather+ix*nt);
61     }
62 }

```

filt/lib/chain.c

```

26 void sf_chain( sf_operator oper1      /* outer operator */,
27               sf_operator oper2      /* inner operator */,
28               bool adj                /* adjoint flag */,
29               bool add                /* addition flag */,
30               int nm                  /* model size */,
31               int nd                  /* data size */,
32               int nt                  /* intermediate size */,
33               /*@out@*/ float* mod   /* [nm] model */,
34               /*@out@*/ float* dat   /* [nd] data */,
35               float* tmp              /* [nt] intermediate */)
36 /*< Chains two operators, computing oper1{oper2{mod}}
37    or its adjoint. The tmp array is used for temporary storage. >*/
38 {
39     if (adj) {
40         oper1 (true, false, nt, nd, tmp, dat);
41         oper2 (true, add, nm, nt, mod, tmp);
42     } else {
43         oper2 (false, false, nm, nt, mod, tmp);
44         oper1 (false, add, nt, nd, tmp, dat);
45     }
46 }

```

```

                                filt/lib/chain.c
70 void sf_array( sf_operator oper1      /* top operator */,
71               sf_operator oper2      /* bottom operator */,
72               bool adj                /* adjoint flag */,
73               bool add                /* addition flag */,
74               int nm                  /* model size */,
75               int nd1                 /* top data size */,
76               int nd2                 /* bottom data size */,
77               /*@out@*/ float* mod   /* [nm] model */,
78               /*@out@*/ float* dat1  /* [nd1] top data */,
79               /*@out@*/ float* dat2  /* [nd2] bottom data */)
80 /*< Constructs an array of two operators,
81    computing {oper1{mod},oper2{mod}} or its adjoint. >*/
82 {
83     if (adj) {
84         oper1 (true, add, nm, nd1, mod, dat1);
85         oper2 (true, true, nm, nd2, mod, dat2);
86     } else {
87         oper1 (false, add, nm, nd1, mod, dat1);
88         oper2 (false, add, nm, nd2, mod, dat2);
89     }
90 }

```

1.2 ADJOINT DEFINED: DOT-PRODUCT TEST

Having seen many examples of **spaces**, operators, and adjoints, we should now see more formal definitions because abstraction helps us push these concepts to their limits.

1.2.1 Definition of a vector space

An operator transforms a **space** to another space. Examples of spaces are model space **m** and data space **d**. We think of these spaces as vectors whose components are packed with numbers, either real or complex numbers. The important practical concept is that not only does this packing include one-dimensional spaces like signals, two-dimensional spaces like images, 3-D movie cubes, and zero-dimensional spaces like a data mean, etc, but spaces can be sets of all the above. One space that is a set of three cubes is the earth’s magnetic field, which has three components; and each component is a function of a three-dimensional space. (The 3-D *physical space* we live in is not the abstract **vector space** of models and data so abundant in this book. Here the word “space” without an adjective means the vector space.)

A more heterogeneous example of a vector space is **data tracks**. A depth-sounding survey of a lake can make a vector space that is a collection of tracks, a vector of vectors (each vector having a different number of components, because lakes are not square). This

vector space of depths along tracks in a lake contains the depth values only. The (x, y) -coordinate information locating each measured depth value is (normally) something outside the vector space. A data space could also be a collection of echo soundings, waveforms recorded along tracks.

We briefly recall information about vector spaces found in elementary books: Let α be any scalar. Then if \mathbf{d}_1 is a vector and \mathbf{d}_2 is conformable with it, then other vectors are $\alpha\mathbf{d}_1$ and $\mathbf{d}_1 + \mathbf{d}_2$. The size measure of a vector is a positive value called a norm. The norm is usually defined to be the **dot product** (also called the L_2 **norm**), say $\mathbf{d} \cdot \mathbf{d}$. For complex data it is $\bar{\mathbf{d}} \cdot \mathbf{d}$ where $\bar{\mathbf{d}}$ is the complex conjugate of \mathbf{d} . In theoretical work the “length of a vector” means the vector’s norm. In computational work the “length of a vector” means the number of components in the vector.

Norms generally include a **weighting function**. In physics, the norm generally measures a conserved quantity like energy or momentum, so, for example, a weighting function for magnetic flux is permittivity. In data analysis, the proper choice of the weighting function is a practical statistical issue, discussed repeatedly throughout this book. The algebraic view of a weighting function is that it is a diagonal matrix with positive values $w(i) \geq 0$ spread along the diagonal, and it is denoted $\mathbf{W} = \text{diag}[w(i)]$. With this weighting function the L_2 norm of a data space is denoted $\mathbf{d}'\mathbf{W}\mathbf{d}$. Standard notation for norms uses a double absolute value, where $\|\mathbf{d}\| = \sqrt{\mathbf{d}'\mathbf{W}\mathbf{d}}$. A central concept with norms is the triangle inequality, $\|\mathbf{d}_1 + \mathbf{d}_2\| \leq \|\mathbf{d}_1\| + \|\mathbf{d}_2\|$ whose proof you might recall (or reproduce with the use of dot products).

1.2.2 Dot-product test for validity of an adjoint

There is a huge gap between the conception of an idea and putting it into practice. During development, things fail far more often than not. Often, when something fails, many tests are needed to track down the cause of failure. Maybe the cause cannot even be found. More insidiously, failure may be below the threshold of detection and poor performance suffered for years. The **dot-product test** enables us to ascertain whether the program for the adjoint of an operator is precisely consistent with the operator itself. It can be, and it should be.

Conceptually, the idea of matrix transposition is simply $a'_{ij} = a_{ji}$. In practice, however, we often encounter matrices far too large to fit in the memory of any computer. Sometimes it is also not obvious how to formulate the process at hand as a matrix multiplication. (Examples are differential equations and fast Fourier transforms.) What we find in practice is that an application and its adjoint amounts to two routines. The first routine amounts to the matrix multiplication $\mathbf{F}\mathbf{x}$. The adjoint routine computes $\mathbf{F}'\mathbf{y}$, where \mathbf{F}' is the **conjugate-transpose** matrix. In later chapters we will be solving huge sets of simultaneous equations, in which both routines are required. If the pair of routines are inconsistent, we are doomed from the start. The dot-product test is a simple test for verifying that the two routines are adjoint to each other.

The associative property of linear algebra says that we do not need parentheses in a vector-matrix-vector product like $\mathbf{y}'\mathbf{F}\mathbf{x}$ because we get the same result no matter where we

put the parentheses. They serve only to determine the sequence of computation. Thus,

$$\mathbf{y}'(\mathbf{F}\mathbf{x}) = (\mathbf{y}'\mathbf{F})\mathbf{x} \quad (1.27)$$

$$\mathbf{y}'(\mathbf{F}\mathbf{x}) = (\mathbf{F}'\mathbf{y})'\mathbf{x} \quad (1.28)$$

(In general, the matrix is not square.) To perform the dot-product test, load the vectors \mathbf{x} and \mathbf{y} with random numbers. Using your program for \mathbf{F} , compute the vector $\tilde{\mathbf{y}} = \mathbf{F}\mathbf{x}$, and using your program for \mathbf{F}' , compute $\tilde{\mathbf{x}} = \mathbf{F}'\mathbf{y}$. Inserting these into equation (1.28) gives you two scalars that should be equal.

$$\mathbf{y}'(\mathbf{F}\mathbf{x}) = \mathbf{y}'\tilde{\mathbf{y}} = \tilde{\mathbf{x}}'\mathbf{x} = (\mathbf{F}'\mathbf{y})'\mathbf{x} \quad (1.29)$$

The left and right sides of this equation will be computationally equal only if the program doing \mathbf{F}' is indeed adjoint to the program doing \mathbf{F} (unless the random numbers do something miraculous). A final word: If an operator uses complex arithmetic then both its input and output are “in the field of complex numbers.” The conversion between real and complex numbers is not a linear operator despite its seeming similarity to truncation and zero padding.

The program for applying the dot product test is `dottest` on page 28. The C way of passing a linear operator as an argument is to specify the function interface. Fortunately, we have already defined the interface for a generic linear operator. To use the `dottest` program, you need to initialize an operator with specific arguments (the `_init` subroutine) and then pass the operator itself (the `_lop` function) to the test program. You also need to specify the sizes of the model and data vectors so that temporary arrays can be constructed. The program runs the dot product test twice, second time with `add = true` to test if the operator can be used properly for accumulating the result like $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{B}\mathbf{x}$.

filr/lib/dottest.c

```

50     oper(false, false, nm, nd, mod1, dat1);
51     dot1[0] = cblas_sdot( nd, dat1, 1, dat2, 1);
52
53     oper(true, false, nm, nd, mod2, dat2);
54     dot1[1] = cblas_sdot( nm, mod1, 1, mod2, 1);
55
56     oper(false, true, nm, nd, mod1, dat1);
57     dot2[0] = cblas_sdot( nd, dat1, 1, dat2, 1);
58
59     oper(true, true, nm, nd, mod2, dat2);
60     dot2[1] = cblas_sdot( nm, mod1, 1, mod2, 1);

```

I tested (1.29) on many operators and was surprised and delighted to find that it is often satisfied to an accuracy near the computing precision. I do not doubt that larger rounding errors could occur, but so far, every time I encountered a relative discrepancy of 10^{-5} or more, I was later able to uncover a conceptual or programming error. Naturally, when I do dot-product tests, I scale the implied matrix to a small dimension in order to speed things along, and to be sure that boundaries are not overwhelmed by the much larger interior.

Do not be alarmed if the operator you have defined has **truncation** errors. Such errors in the definition of the original operator should be identically matched by truncation errors in the adjoint operator. If your code passes the **dot-product test**, then you really have coded the adjoint operator. In that case, to obtain inverse operators, you can take advantage of the standard methods of mathematics.

We can speak of a **continuous function** $f(t)$ or a **discrete function** f_t . For continuous functions we use integration, and for discrete ones we use summation. In formal mathematics, the dot-product test *defines* the adjoint operator, except that the summation in the dot product may need to be changed to an integral. The input or the output or both can be given either on a continuum or in a discrete domain. So the dot-product test $\mathbf{y}'\tilde{\mathbf{y}} = \tilde{\mathbf{x}}'\mathbf{x}$ could have an integration on one side of the equal sign and a summation on the other. Linear-operator theory is rich with concepts not developed here.

1.2.3 The word “adjoint”

In mathematics the word “**adjoint**” has three meanings. One of them, the so-called **Hilbert adjoint**, is the one generally found in physics and engineering and it is the one used in this book. In linear algebra is a different matrix, called the **adjugate** matrix. It is a matrix whose elements are signed cofactors (minor determinants). For invertible matrices, this matrix is the **determinant** times the **inverse matrix**. It can be computed without ever using division, so potentially the adjugate can be useful in applications where an inverse matrix does not exist. Unfortunately, the adjugate matrix is sometimes called the adjoint matrix, particularly in the older literature. Because of the confusion of multiple meanings of the word adjoint, in the first printing of PVI, I avoided the use of the word and substituted the definition, “**conjugate transpose**”. Unfortunately this was often abbreviated to “conjugate,” which caused even more confusion. Thus I decided to use the word adjoint and have it always mean the Hilbert adjoint found in physics and engineering.

1.2.4 Matrix versus operator

Here is a short summary of where we have been and where we are going: Start from the class of linear operators, add subscripts and you get matrices. Examples of operators without subscripts are routines that solve differential equations and routines that do fast Fourier transform. What people call “sparse matrices” are often not really matrices but operators, because they are not defined by data structures but by routines that apply them to a vector. With sparse matrices you easily can do $\mathbf{A}(\mathbf{B}(\mathbf{C}\mathbf{x}))$ but not $(\mathbf{ABC})\mathbf{x}$.

Although a linear operator does not have defined subscripts, you can determine what would be the operator value at any subscript: by applying the operator to an impulse function, you would get a matrix column. The adjoint operator is one from which we can extract the transpose matrix. For large spaces this extraction is unwieldy, so to test the validity of adjoints, we probe them with random vectors, say \mathbf{x} and \mathbf{y} , to see whether $\mathbf{y}'(\mathbf{Ax}) = (\mathbf{A}'\mathbf{y})'\mathbf{x}$. Mathematicians define adjoints by this test, except that instead of using random vectors, they say “for all functions,” which includes the continuum.

This defining test makes adjoints look mysterious. Careful inspection of operator adjoints, however, generally reveals that they are built up from simple matrices. Given adjoints

\mathbf{A}' , \mathbf{B}' , and \mathbf{C}' , the adjoint of \mathbf{ABC} is $\mathbf{C}'\mathbf{B}'\mathbf{A}'$. Fourier transforms and linear-differential-equation solvers are chains of matrices, so their adjoints can be assembled by the application of adjoint components in reverse order. The other way we often see complicated operators being built from simple ones is when operators are put into components of matrices, typically a 1×2 or 2×1 matrix containing two operators. An example of the adjoint of a two-component column operator is

$$\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}' = \begin{bmatrix} \mathbf{A}' & \mathbf{B}' \end{bmatrix} \quad (1.30)$$

Although in practice an operator might be built from matrices, fundamentally, a matrix is a data structure whereas an operator is a procedure. A matrix is an operator if its subscripts are hidden but it can be applied to a space, producing another space.

As matrices have inverses, so do linear operators. You don't need subscripts to find an inverse. The conjugate-gradient method and conjugate-direction method explained in the next chapter are attractive methods of finding them. They merely apply \mathbf{A} and \mathbf{A}' and use inner products to find coefficients of a polynomial in \mathbf{AA}' that represents the inverse operator.

Whenever we encounter a **positive-definite** matrix we should recognize its likely origin in a nonsymmetric matrix \mathbf{F} times its adjoint. Those in natural sciences often work on solving simultaneous equations but fail to realize that they should return to the origin of the equations which is often a fitting goal; i.e., applying an operator to a model should yield data, i.e., $\mathbf{d} \approx \mathbf{d}_0 + \mathbf{F}(\mathbf{m} - \mathbf{m}_0)$ where the operator \mathbf{F} is a partial derivative matrix (and there are potential underlying nonlinearities). This begins another story with new ingredients, weighting functions and statistics.

1.2.5 Inverse operator

A common practical task is to fit a vector of observed data \mathbf{d}_{obs} to some theoretical data $\mathbf{d}_{\text{theor}}$ by the adjustment of components in a vector of model parameters \mathbf{m} .

$$\mathbf{d}_{\text{obs}} \approx \mathbf{d}_{\text{theor}} = \mathbf{Fm} \quad (1.31)$$

A huge volume of literature establishes theory for two estimates of the model, $\hat{\mathbf{m}}_1$ and $\hat{\mathbf{m}}_2$, where

$$\hat{\mathbf{m}}_1 = (\mathbf{F}'\mathbf{F})^{-1}\mathbf{F}'\mathbf{d} \quad (1.32)$$

$$\hat{\mathbf{m}}_2 = \mathbf{F}'(\mathbf{FF}')^{-1}\mathbf{d} \quad (1.33)$$

Some reasons for the literature being huge are the many questions about the existence, quality, and cost of the inverse operators. Before summarizing that, let us quickly see why these two solutions are reasonable. Inserting equation (1.31) into equation (1.32), and inserting equation (1.33) into equation (1.31), we get the reasonable statements:

$$\hat{\mathbf{m}}_1 = (\mathbf{F}'\mathbf{F})^{-1}(\mathbf{F}'\mathbf{F})\mathbf{m} = \mathbf{m} \quad (1.34)$$

$$\hat{\mathbf{d}}_{\text{theor}} = (\mathbf{FF}')(\mathbf{FF}')^{-1}\mathbf{d} = \mathbf{d} \quad (1.35)$$

Equation (1.34) says that the estimate $\hat{\mathbf{m}}_1$ gives the correct model \mathbf{m} if you start from the theoretical data. Equation (1.35) says that the model estimate $\hat{\mathbf{m}}_2$ gives the theoretical data if we derive $\hat{\mathbf{m}}_2$ from the theoretical data. Both of these statements are delightful. Now let us return to the problem of the inverse matrices.

Strictly speaking, a rectangular matrix does not have an inverse. Surprising things often happen, but commonly, when \mathbf{F} is a tall matrix (more data values than model values) then the matrix for finding $\hat{\mathbf{m}}_1$ is invertible while that for finding $\hat{\mathbf{m}}_2$ is not, and when the matrix is wide instead of tall (the number of data values is less than the number of model values) it is the other way around. In many applications neither $\mathbf{F}'\mathbf{F}$ nor $\mathbf{F}\mathbf{F}'$ is invertible. This difficulty is solved by “**damping**” as we will see in later chapters. The point to notice in this chapter on adjoints is that in any application where $\mathbf{F}\mathbf{F}'$ or $\mathbf{F}'\mathbf{F}$ equals \mathbf{I} (unitary operator), that the adjoint operator \mathbf{F}' is the inverse \mathbf{F}^{-1} by either equation (1.32) or (1.33).

Theoreticians like to study inverse problems where \mathbf{m} is drawn from the field of continuous functions. This is like the vector \mathbf{m} having infinitely many components. Such problems are hopelessly intractable unless we find, or assume, that the operator $\mathbf{F}'\mathbf{F}$ is an identity or diagonal matrix.

In practice, theoretical considerations may have little bearing on how we proceed. Current computational power limits matrix inversion jobs to about 10^4 variables. This book specializes in big problems, those with more than about 10^4 variables, but the methods we learn are also excellent for smaller problems.

1.2.6 Automatic adjoints

Computers are not only able to perform computations; they can do mathematics. Well known software is Mathematica and Maple. Adjoints can also be done by symbol manipulation. For example Ralf Giering offers a program for converting linear operator programs into their adjoints.

EXERCISES:

- 1 Suppose a linear operator \mathbf{F} has its input in the discrete domain and its output in the continuum. How does the operator resemble a matrix? Describe the operator \mathbf{F}' that has its input in the discrete domain and its output in the continuum. To which do you apply the words “scales and adds some functions,” and to which do you apply the words “does a bunch of integrals”? What are the integrands?

Chapter 2

Model fitting by least squares

The first level of computer use in science and engineering is **modeling**. Beginning from physical principles and design ideas, the computer mimics nature. After this, the worker looks at the result and thinks a while, then alters the modeling program and tries again. The next, deeper level of computer use is that the computer itself examines the results of modeling and reruns the modeling job. This deeper level is variously called “**fitting**” or “**estimation**” or “**inversion**.” We inspect the **conjugate-direction method** of fitting and write a subroutine for it that will be used in most of the examples in this monograph.

2.1 HOW TO DIVIDE NOISY SIGNALS

If “inversion” is dividing by a matrix, then the place to begin is dividing one number by another, say one function of frequency by another function of frequency. A single parameter fitting problem arises in Fourier analysis, where we seek a “best answer” at each frequency, then combine all the frequencies to get a best signal. Thus emerges a wide family of interesting and useful applications. However, Fourier analysis first requires us to introduce complex numbers into statistical estimation.

Multiplication in the Fourier domain is **convolution** in the time domain. Fourier-domain division is time-domain **deconvolution**. This division is challenging when the divisor has observational error. Failure erupts if zero division occurs. More insidious are the poor results we obtain when zero division is avoided by a near miss.

2.1.1 Dividing by zero smoothly

Think of any real numbers x , y , and f where $y = xf$. Given y and f we see a computer program containing $x = y/f$. How can we change the program so that it never divides by zero? A popular answer is to change $x = y/f$ to $x = yf/(f^2 + \epsilon^2)$, where ϵ is any tiny value. When $|f| \gg |\epsilon|$, then x is approximately y/f as expected. But when the divisor f vanishes, the result is safely zero instead of infinity. The transition is smooth, but some criterion is needed to choose the value of ϵ . This method may not be the only way or the best way to cope with **zero division**, but it is a good way, and it permeates the subject of signal analysis.

To apply this method in the Fourier domain, suppose that X , Y , and F are complex numbers. What do we do then with $X = Y/F$? We multiply the top and bottom by the complex conjugate \bar{F} , and again add ϵ^2 to the denominator. Thus,

$$X(\omega) = \frac{\overline{F(\omega)} Y(\omega)}{\overline{F(\omega)} F(\omega) + \epsilon^2} \quad (2.1)$$

Now the denominator must always be a positive number greater than zero, so division is always safe. Equation (2.1) ranges continuously from **inverse filtering**, with $X = Y/F$, to filtering with $X = \bar{F}Y$, which is called “**matched filtering**.” Notice that for any complex number F , the phase of $1/F$ equals the phase of \bar{F} , so the filters $1/F$ and \bar{F} have inverse amplitudes but identical phase.

2.1.2 Damped solution

Equation (2.1) is the solution to an optimization problem that arises in many applications. Now that we know the solution, let us formally define the problem. First, we will solve a simpler problem with real values: we will choose to minimize the **quadratic function** of x :

$$Q(x) = (fx - y)^2 + \epsilon^2 x^2 \quad (2.2)$$

The second term is called a “**damping factor**” because it prevents x from going to $\pm\infty$ when $f \rightarrow 0$. Set $dQ/dx = 0$, which gives

$$0 = f(fx - y) + \epsilon^2 x \quad (2.3)$$

This yields the earlier answer $x = fy/(f^2 + \epsilon^2)$.

With Fourier transforms, the signal X is a complex number at each frequency ω . So we generalize equation (2.2) to

$$Q(\bar{X}, X) = (\overline{FX - Y})(FX - Y) + \epsilon^2 \bar{X}X = (\bar{X}\bar{F} - \bar{Y})(FX - Y) + \epsilon^2 \bar{X}X \quad (2.4)$$

To minimize Q we could use a real-values approach, where we express $X = u + iv$ in terms of two real values u and v and then set $\partial Q/\partial u = 0$ and $\partial Q/\partial v = 0$. The approach we will take, however, is to use complex values, where we set $\partial Q/\partial X = 0$ and $\partial Q/\partial \bar{X} = 0$. Let us examine $\partial Q/\partial \bar{X}$:

$$\frac{\partial Q(\bar{X}, X)}{\partial \bar{X}} = \bar{F}(FX - Y) + \epsilon^2 X = 0 \quad (2.5)$$

The derivative $\partial Q/\partial X$ is the complex conjugate of $\partial Q/\partial \bar{X}$. So if either is zero, the other is too. Thus we do not need to specify both $\partial Q/\partial X = 0$ and $\partial Q/\partial \bar{X} = 0$. I usually set $\partial Q/\partial \bar{X}$ equal to zero. Solving equation (2.5) for X gives equation (2.1).

Equation (2.1) solves $Y = XF$ for X , giving the solution for what is called “the **deconvolution** problem with a known wavelet F .” Analogously we can use $Y = XF$ when the filter F is unknown, but the input X and output Y are given. Simply interchange X and F in the derivation and result.

2.1.3 Smoothing the denominator spectrum

Equation (2.1) gives us one way to divide by zero. Another way is stated by the equation

$$X(\omega) = \frac{\overline{F(\omega)} Y(\omega)}{\langle \overline{F(\omega)} F(\omega) \rangle} \quad (2.6)$$

where the strange notation in the denominator means that the spectrum there should be smoothed a little. Such smoothing fills in the holes in the spectrum where zero-division is a danger, filling not with an arbitrary numerical value ϵ but with an average of nearby spectral values. Additionally, if the denominator spectrum $\overline{F(\omega)} F(\omega)$ is rough, the smoothing creates a shorter autocorrelation function.

Both divisions, equation (2.1) and equation (2.6), irritate us by requiring us to specify a parameter, but for the latter, the parameter has a clear meaning. In the latter case we smooth a spectrum with a smoothing window of width, say $\Delta\omega$ which this corresponds inversely to a time interval over which we smooth. Choosing a numerical value for ϵ has not such a simple interpretation.

We jump from simple mathematical theorizing towards a genuine practical application when I grab some real data, a function of time and space from another textbook. Let us call this data $f(t, x)$ and its 2-D Fourier transform $F(\omega, k_x)$. The data and its autocorrelation are in Figure 2.1.

The autocorrelation $a(t, x)$ of $f(t, x)$ is the inverse 2-D Fourier Transform of $\overline{F(\omega, k_x)} F(\omega, k_x)$. Autocorrelations $a(x, y)$ satisfy the symmetry relation $a(x, y) = a(-x, -y)$. Figure 2.2 shows only the interesting quadrant of the two independent quadrants. We see the autocorrelation of a 2-D function has some resemblance to the function itself but differs in important ways.

Instead of messing with two different functions X and Y to divide, let us divide F by itself. This sounds like $1 = F/F$ but we will watch what happens when we do the division carefully avoiding zero division in the ways we usually do.

Figure 2.2 shows what happens with

$$1 = F/F \approx \frac{\overline{F} F}{\overline{F} F + \epsilon^2} \quad (2.7)$$

and with

$$1 = F/F \approx \frac{\overline{F} F}{\langle \overline{F} F \rangle} \quad (2.8)$$

From Figure 2.2 we notice that both methods of avoiding zero division give similar results. By playing with the ϵ and the smoothing width the pictures could be made even more similar. My preference, however, is the smoothing. It is difficult to make physical sense of choosing a numerical value for ϵ . It is much easier to make physical sense of choosing a smoothing window. The smoothing window is in (ω, k_x) space, but Fourier transformation tells us its effect in (t, x) space.




Figure 2.1: 2-D data (right) and a quadrant of its autocorrelation (left). Notice the longest nonzero time lag on the data is about 5.5 sec which is the latest nonzero signal on the autocorrelation. lsq/antoine antoine10




Figure 2.2: Equation (2.7) (left) and equation (2.8) (right). Both ways of dividing by zero give similar results. lsq/antoine antoine11

2.1.4 Imaging

The example of dividing a function by itself ($1 = F/F$) might not seem to make much sense, but it is very closely related to estimation often encountered in imaging applications. It's not my purpose here to give a lecture on imaging theory, but here is an overbrief explanation.

Imagine a downgoing wavefield $D(\omega, x, z)$ and scatterer that from the downgoing wavefield creates an upgoing wavefield $U(\omega, x, z)$. Given U and D , if there is a strong temporal correlation between them at any (x, z) it likely means there is a reflector nearby that is manufacturing U from D . This reflectivity could be quantified by U/D . At the earth's surface the surface boundary condition says something like $U = D$ or $U = -D$. Thus at the surface we have something like F/F . As we go down in the earth, the main difference is that U and D get time shifted in opposite directions, so U and D are similar but for that time difference. Thus, a study of how we handle F/F is worthwhile.

2.1.5 Formal path to the low-cut filter

This book defines many geophysical estimation problems. Many of them amount to statement of two goals. The first goal is a data fitting goal, the goal that the model should imply some observed data. The second goal is that the model be not too big or too wiggly. We will state these goals as two residuals, each of which is ideally zero. A very simple data fitting goal would be that the model m equals the data d , thus the difference should vanish, say $0 \approx m - d$. A more interesting goal is that the model should match the data especially at high frequencies but not necessarily at low frequencies.

$$0 \approx -i\omega(m - d) \quad (2.9)$$

A danger of this goal is that the model could have a zero-frequency component of infinite magnitude as well as large amplitudes for low frequencies. To suppress this, we need the second goal, a model residual which is to be minimized. We need a small number ϵ . The model goal is

$$0 \approx \epsilon m \quad (2.10)$$

To see the consequence of these two goals, we add the squares of the residuals

$$Q(m) = \omega^2(m - d)^2 + \epsilon^2 m^2 \quad (2.11)$$

and then we minimize $Q(m)$ by setting its derivative to zero

$$0 = \frac{dQ}{dm} = 2\omega^2(m - d) + 2\epsilon^2 m \quad (2.12)$$

or

$$m = \frac{\omega^2}{\omega^2 + \epsilon^2} d \quad (2.13)$$

which is a low-cut filter with a cutoff frequency of $\omega_0 = \epsilon$.

Of some curiosity and significance is the numerical choice of ϵ . The general theory says we need an epsilon, but does not say how much. For now let us simply rename $\epsilon = \omega_0$ and think of it as a “cut off frequency”.

2.2 MULTIVARIATE LEAST SQUARES

2.2.1 Inside an abstract vector

In engineering uses, a vector has three scalar components that correspond to the three dimensions of the space in which we live. In least-squares data analysis, a vector is a one-dimensional array that can contain many different things. Such an array is an “**abstract vector**.” For example, in earthquake studies, the vector might contain the time an earthquake began, as well as its latitude, longitude, and depth. Alternatively, the abstract vector might contain as many components as there are seismometers, and each component might be the arrival time of an earthquake wave. Used in signal analysis, the vector might contain the values of a signal at successive instants in time or, alternatively, a collection of signals. These signals might be “**multiplexed**” (interlaced) or “**demultiplexed**” (all of each signal preceding the next). When used in image analysis, the one-dimensional array might contain an image, which could itself be thought of as an array of signals. Vectors, including abstract vectors, are usually denoted by **boldface letters** such as \mathbf{p} and \mathbf{s} . Like physical vectors, abstract vectors are **orthogonal** when their dot product vanishes: $\mathbf{p} \cdot \mathbf{s} = 0$. Orthogonal vectors are well known in physical space; we will also encounter them in abstract vector space.

We consider first a hypothetical application with one data vector \mathbf{d} and two fitting vectors \mathbf{f}_1 and \mathbf{f}_2 . Each fitting vector is also known as a “**regressor**.” Our first task is to approximate the data vector \mathbf{d} by a scaled combination of the two regressor vectors. The scale factors x_1 and x_2 should be chosen so that the model matches the data; i.e.,

$$\mathbf{d} \approx \mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 \quad (2.14)$$

Notice that we could take the partial derivative of the data in (2.14) with respect to an unknown, say x_1 , and the result is the regressor \mathbf{f}_1 .

The **partial derivative** of all theoretical data with respect to any model parameter gives a **regressor**. A **regressor** is a column in the matrix of partial-derivatives, $\partial d_i / \partial m_j$.

The fitting goal (2.14) is often expressed in the more compact mathematical matrix notation $\mathbf{d} \approx \mathbf{F}\mathbf{x}$, but in our derivation here we will keep track of each component explicitly and use mathematical matrix notation to summarize the final result. Fitting the observed data $\mathbf{d} = \mathbf{d}^{\text{obs}}$ to its two theoretical parts $\mathbf{f}_1 x_1$ and $\mathbf{f}_2 x_2$ can be expressed as minimizing the length of the residual vector \mathbf{r} , where

$$\mathbf{0} \approx \mathbf{r} = \mathbf{d}^{\text{theor}} - \mathbf{d}^{\text{obs}} \quad (2.15)$$

$$\mathbf{0} \approx \mathbf{r} = \mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d} \quad (2.16)$$

We use a dot product to construct a sum of squares (also called a “**quadratic form**”) of the components of the residual vector:

$$Q(x_1, x_2) = \mathbf{r} \cdot \mathbf{r} \quad (2.17)$$

$$= (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) \cdot (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) \quad (2.18)$$

To find the gradient of the quadratic form $Q(x_1, x_2)$, you might be tempted to expand out the dot product into all nine terms and then differentiate. It is less cluttered, however, to remember the product rule, that

$$\frac{d}{dx} \mathbf{r} \cdot \mathbf{r} = \frac{d\mathbf{r}}{dx} \cdot \mathbf{r} + \mathbf{r} \cdot \frac{d\mathbf{r}}{dx} \quad (2.19)$$

Thus, the gradient of $Q(x_1, x_2)$ is defined by its two components:

$$\frac{\partial Q}{\partial x_1} = \mathbf{f}_1 \cdot (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) + (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) \cdot \mathbf{f}_1 \quad (2.20)$$

$$\frac{\partial Q}{\partial x_2} = \mathbf{f}_2 \cdot (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) + (\mathbf{f}_1 x_1 + \mathbf{f}_2 x_2 - \mathbf{d}) \cdot \mathbf{f}_2 \quad (2.21)$$

Setting these derivatives to zero and using $(\mathbf{f}_1 \cdot \mathbf{f}_2) = (\mathbf{f}_2 \cdot \mathbf{f}_1)$ etc., we get

$$(\mathbf{f}_1 \cdot \mathbf{d}) = (\mathbf{f}_1 \cdot \mathbf{f}_1)x_1 + (\mathbf{f}_1 \cdot \mathbf{f}_2)x_2 \quad (2.22)$$

$$(\mathbf{f}_2 \cdot \mathbf{d}) = (\mathbf{f}_2 \cdot \mathbf{f}_1)x_1 + (\mathbf{f}_2 \cdot \mathbf{f}_2)x_2 \quad (2.23)$$

We can use these two equations to solve for the two unknowns x_1 and x_2 . Writing this expression in matrix notation, we have

$$\begin{bmatrix} (\mathbf{f}_1 \cdot \mathbf{d}) \\ (\mathbf{f}_2 \cdot \mathbf{d}) \end{bmatrix} = \begin{bmatrix} (\mathbf{f}_1 \cdot \mathbf{f}_1) & (\mathbf{f}_1 \cdot \mathbf{f}_2) \\ (\mathbf{f}_2 \cdot \mathbf{f}_1) & (\mathbf{f}_2 \cdot \mathbf{f}_2) \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (2.24)$$

It is customary to use matrix notation without dot products. To do this, we need some additional definitions. To clarify these definitions, we inspect vectors \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{d} of three components. Thus

$$\mathbf{F} = [\mathbf{f}_1 \quad \mathbf{f}_2] = \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \\ f_{31} & f_{32} \end{bmatrix} \quad (2.25)$$

Likewise, the *transposed* matrix \mathbf{F}' is defined by

$$\mathbf{F}' = \begin{bmatrix} f_{11} & f_{21} & f_{31} \\ f_{12} & f_{22} & f_{32} \end{bmatrix} \quad (2.26)$$

The matrix in equation (2.24) contains dot products. Matrix multiplication is an abstract way of representing the dot products:

$$\begin{bmatrix} (\mathbf{f}_1 \cdot \mathbf{f}_1) & (\mathbf{f}_1 \cdot \mathbf{f}_2) \\ (\mathbf{f}_2 \cdot \mathbf{f}_1) & (\mathbf{f}_2 \cdot \mathbf{f}_2) \end{bmatrix} = \begin{bmatrix} f_{11} & f_{21} & f_{31} \\ f_{12} & f_{22} & f_{32} \end{bmatrix} \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \\ f_{31} & f_{32} \end{bmatrix} \quad (2.27)$$

Thus, equation (2.24) without dot products is

$$\begin{bmatrix} f_{11} & f_{21} & f_{31} \\ f_{12} & f_{22} & f_{32} \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \end{bmatrix} = \begin{bmatrix} f_{11} & f_{21} & f_{31} \\ f_{12} & f_{22} & f_{32} \end{bmatrix} \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \\ f_{31} & f_{32} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (2.28)$$

which has the matrix abbreviation

$$\mathbf{F}'\mathbf{d} = (\mathbf{F}'\mathbf{F})\mathbf{x} \quad (2.29)$$

Equation (2.29) is the classic result of least-squares fitting of data to a collection of regressors. Obviously, the same matrix form applies when there are more than two regressors and each vector has more than three components. Equation (2.29) leads to an **analytic solution** for \mathbf{x} using an inverse matrix. To solve formally for the unknown \mathbf{x} , we premultiply by the inverse matrix $(\mathbf{F}' \mathbf{F})^{-1}$:

$$\mathbf{x} = (\mathbf{F}' \mathbf{F})^{-1} \mathbf{F}' \mathbf{d} \quad (2.30)$$

Equation (2.30) is the central result of **least-squares** theory. We see it everywhere.

In our first manipulation of matrix algebra, we move around some parentheses in (2.29):

$$\mathbf{F}' \mathbf{d} = \mathbf{F}' (\mathbf{F} \mathbf{x}) \quad (2.31)$$

Moving the parentheses implies a regrouping of terms or a reordering of a computation. You can verify the validity of moving the parentheses if you write (2.31) in full as the set of two equations it represents. Equation (2.29) led to the “analytic” solution (2.30). In a later section on conjugate directions, we will see that equation (2.31) expresses better than (2.30) the philosophy of iterative methods.

Notice how equation (2.31) invites us to cancel the matrix \mathbf{F}' from each side. We cannot do that of course, because \mathbf{F}' is not a number, nor is it a square matrix with an inverse. If you really want to cancel the matrix \mathbf{F}' , you may, but the equation is then only an approximation that restates our original goal (2.14):

$$\mathbf{d} \approx \mathbf{F} \mathbf{x} \quad (2.32)$$

A speedy problem solver might ignore the mathematics covering the previous page, study his or her application until he or she is able to write the **statement of goals** (2.32) = (2.14), premultiply by \mathbf{F}' , replace \approx by $=$, getting (2.29), and take (2.29) to a simultaneous equation-solving program to get \mathbf{x} .

What I call “**fitting goals**” are called “**regressions**” by statisticians. In common language the word regression means to “trend toward a more primitive perfect state” which vaguely resembles reducing the size of (energy in) the residual $\mathbf{r} = \mathbf{F} \mathbf{x} - \mathbf{d}$. Formally this is often written as:

$$\min_{\mathbf{x}} ||\mathbf{F} \mathbf{x} - \mathbf{d}|| \quad (2.33)$$

The notation above with two pairs of vertical lines looks like double absolute value, but we can understand it as a reminder to square and sum all the components. This formal notation is more explicit about what is constant and what is variable during the fitting.

2.2.2 Normal equations

An important concept is that when energy is minimum, the residual is orthogonal to the fitting functions. The fitting functions are the column vectors \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_3 . Let us verify only that the dot product $\mathbf{r} \cdot \mathbf{f}_2$ vanishes; to do this, we’ll show that those two vectors are orthogonal. Energy minimum is found by

$$0 = \frac{\partial}{\partial x_2} \mathbf{r} \cdot \mathbf{r} = 2 \mathbf{r} \cdot \frac{\partial \mathbf{r}}{\partial x_2} = 2 \mathbf{r} \cdot \mathbf{f}_2 \quad (2.34)$$

(To compute the derivative refer to equation (2.16).) Equation (2.34) shows that the residual is orthogonal to a fitting function. The fitting functions are the column vectors in the fitting matrix.

The basic least-squares equations are often called the “**normal**” equations. The word “normal” means perpendicular. We can rewrite equation (2.31) to emphasize the perpendicularity. Bring both terms to the left, and recall the definition of the residual \mathbf{r} from equation (2.16):

$$\mathbf{F}'(\mathbf{F}\mathbf{x} - \mathbf{d}) = \mathbf{0} \quad (2.35)$$

$$\mathbf{F}'\mathbf{r} = \mathbf{0} \quad (2.36)$$

Equation (2.36) says that the **residual** vector \mathbf{r} is perpendicular to each row in the \mathbf{F}' matrix. These rows are the **fitting functions**. Therefore, the residual, after it has been minimized, is perpendicular to *all* the fitting functions.

2.2.3 Differentiation by a complex vector

Complex numbers frequently arise in physical problems, particularly those with Fourier series. Let us extend the multivariable least-squares theory to the use of complex-valued unknowns \mathbf{x} . First recall how complex numbers were handled with single-variable least squares; i.e., as in the discussion leading up to equation (2.5). Use a prime, such as \mathbf{x}' , to denote the complex conjugate of the transposed vector \mathbf{x} . Now write the positive **quadratic form** as

$$Q(\mathbf{x}', \mathbf{x}) = (\mathbf{F}\mathbf{x} - \mathbf{d})'(\mathbf{F}\mathbf{x} - \mathbf{d}) = (\mathbf{x}'\mathbf{F}' - \mathbf{d}')(\mathbf{F}\mathbf{x} - \mathbf{d}) \quad (2.37)$$

After equation (2.4), we minimized a quadratic form $Q(\bar{X}, X)$ by setting to zero both $\partial Q/\partial \bar{X}$ and $\partial Q/\partial X$. We noted that only one of $\partial Q/\partial \bar{X}$ and $\partial Q/\partial X$ is necessarily zero because they are conjugates of each other. Now take the derivative of Q with respect to the (possibly complex, row) vector \mathbf{x}' . Notice that $\partial Q/\partial \mathbf{x}'$ is the complex conjugate transpose of $\partial Q/\partial \mathbf{x}$. Thus, setting one to zero sets the other also to zero. Setting $\partial Q/\partial \mathbf{x}' = \mathbf{0}$ gives the normal equations:

$$\mathbf{0} = \frac{\partial Q}{\partial \mathbf{x}'} = \mathbf{F}'(\mathbf{F}\mathbf{x} - \mathbf{d}) \quad (2.38)$$

The result is merely the complex form of our earlier result (2.35). Therefore, differentiating by a complex vector is an abstract concept, but it gives the same set of equations as differentiating by each scalar component, and it saves much clutter.

2.2.4 From the frequency domain to the time domain

Equation (2.4) is a frequency-domain quadratic form that we minimized by varying a single parameter, a Fourier coefficient. Now we will look at the same problem in the time domain. We will see that the time domain offers flexibility with boundary conditions, constraints, and weighting functions. The notation will be that a filter f_t has input x_t and output y_t . In Fourier space this is $Y = XF$. There are two problems to look at, unknown filter F and unknown input X .

Unknown filter

When inputs and outputs are given, the problem of finding an unknown filter appears to be overdetermined, so we write $\mathbf{y} \approx \mathbf{X}\mathbf{f}$ where the matrix \mathbf{X} is a matrix of downshifted columns like (1.5). Thus the quadratic form to be minimized is a restatement of equation (2.37) with filter definitions:

$$Q(\mathbf{f}', \mathbf{f}) = (\mathbf{X}\mathbf{f} - \mathbf{y})'(\mathbf{X}\mathbf{f} - \mathbf{y}) \quad (2.39)$$

The solution \mathbf{f} is found just as we found (2.38), and it is the set of simultaneous equations $\mathbf{0} = \mathbf{X}'(\mathbf{X}\mathbf{f} - \mathbf{y})$.

Unknown input: deconvolution with a known filter

For solving the unknown-input problem, we put the known filter f_t in a matrix of downshifted columns \mathbf{F} . Our statement of wishes is now to find x_t so that $\mathbf{y} \approx \mathbf{F}\mathbf{x}$. We can expect to have trouble finding unknown inputs x_t when we are dealing with certain kinds of filters, such as **bandpass filters**. If the output is zero in a frequency band, we will never be able to find the input in that band and will need to prevent x_t from diverging there. We do this by the statement that we wish $\mathbf{0} \approx \epsilon \mathbf{x}$, where ϵ is a parameter that is small and whose exact size will be chosen by experimentation. Putting both wishes into a single, partitioned matrix equation gives

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix} \approx \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \epsilon \mathbf{I} \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix} \quad (2.40)$$

To minimize the residuals \mathbf{r}_1 and \mathbf{r}_2 , we can minimize the scalar $\mathbf{r}'\mathbf{r} = \mathbf{r}'_1\mathbf{r}_1 + \mathbf{r}'_2\mathbf{r}_2$. This is

$$\begin{aligned} Q(\mathbf{x}', \mathbf{x}) &= (\mathbf{F}\mathbf{x} - \mathbf{y})'(\mathbf{F}\mathbf{x} - \mathbf{y}) + \epsilon^2 \mathbf{x}'\mathbf{x} \\ &= (\mathbf{x}'\mathbf{F}' - \mathbf{y}')(\mathbf{F}\mathbf{x} - \mathbf{y}) + \epsilon^2 \mathbf{x}'\mathbf{x} \end{aligned} \quad (2.41)$$

We solved this minimization in the frequency domain (beginning from equation (2.4)).

Formally the solution is found just as with equation (2.38), but this solution looks unappealing in practice because there are so many unknowns and because the problem can be solved much more quickly in the Fourier domain. To motivate ourselves to solve this problem in the time domain, we need either to find an approximate solution method that is much faster, or to discover that constraints or time-variable weighting functions are required in some applications. This is an issue we must be continuously alert to, whether the cost of a method is justified by its need.

EXERCISES:

- 1 In 1695, 150 years before Lord Kelvin's absolute temperature scale, 120 years before Sadi Carnot's PhD thesis, 40 years before Anders Celsius, and 20 years before Gabriel Fahrenheit, the French physicist Guillaume **Amontons**, deaf since birth, took a mercury manometer (pressure gauge) and sealed it inside a glass pipe (a constant volume of air). He heated it to the boiling point of water at 100°C. As he lowered the temperature to freezing at 0° C, he observed the pressure dropped by 25% . He could not drop

the temperature any further but he supposed that if he could drop it further by a factor of three, the pressure would drop to zero (the lowest possible pressure) and the temperature would have been the lowest possible temperature. Had he lived after Anders Celsius he might have calculated this temperature to be -300° C (Celsius). Absolute zero is now known to be -273° C.

It is your job to be Amontons' lab assistant. Your i th measurement of temperature T_i you make with Issac Newton's thermometer and you measure pressure P_i and volume V_i in the metric system. Amontons needs you to fit his data with the regression $0 \approx \alpha(T_i - T_0) - P_i V_i$ and calculate the temperature shift T_0 that Newton should have made when he defined his temperature scale. Do not solve this problem! Instead, cast it in the form of equation (2.14), identifying the data \mathbf{d} and the two column vectors \mathbf{f}_1 and \mathbf{f}_2 that are the fitting functions. Relate the model parameters x_1 and x_2 to the physical parameters α and T_0 . Suppose you make ALL your measurements at room temperature, can you find T_0 ? Why or why not?

2.3 KRYLOV SUBSPACE ITERATIVE METHODS

The **solution time** for simultaneous **linear equations** grows cubically with the number of unknowns. There are three regimes for solution; which one is applicable depends on the number of unknowns m . For m three or less, we use analytical methods. We also sometimes use analytical methods on matrices of size 4×4 when the matrix contains enough zeros. Today in year 2001, a desktop workstation, working an hour solves about a 4000×4000 set of simultaneous equations. A square image packed into a 4096 point vector is a 64×64 array. The computer power for linear algebra to give us solutions that fit in a $k \times k$ image is thus proportional to k^6 , which means that even though computer power grows rapidly, imaging resolution using "exact numerical methods" hardly grows at all from our 64×64 current practical limit.

The retina in our eyes captures an image of size about 1000×1000 which is a lot bigger than 64×64 . Life offers us many occasions where final images exceed the 4000 points of a 64×64 array. To make linear algebra (and inverse theory) relevant to such problems, we investigate special techniques. A numerical technique known as the "**conjugate-direction method**" works well for all values of m and is our subject here. As with most simultaneous equation solvers, an exact answer (assuming exact arithmetic) is attained in a finite number of steps. And if n and m are too large to allow enough iterations, the iterative methods can be interrupted at any stage, the partial result often proving useful. Whether or not a partial result actually is useful is the subject of much research; naturally, the results vary from one application to the next.

2.3.1 Sign convention

On the last day of the survey, a storm blew up, the sea got rough, and the receivers drifted further downwind. The data recorded that day had a larger than usual difference from that predicted by the final model. We could call $(\mathbf{d} - \mathbf{F}\mathbf{m})$ the *experimental error*. (Here \mathbf{d} is data, \mathbf{m} is model parameters, and \mathbf{F} is their linear relation).

The alternate view is that our theory was too simple. It lacked model parameters for the

waves and the drifting cables. Because of this model oversimplification we had a **modeling error** of the opposite polarity ($\mathbf{F}\mathbf{m} - \mathbf{d}$).

A strong experimentalist prefers to think of the error as experimental error, something for him or her to work out. Likewise a strong analyst likes to think of the error as a theoretical problem. (Weaker investigators might be inclined to take the opposite view.)

Regardless of the above, and opposite to common practice, I define the **sign convention** for the error (or residual) as $(\mathbf{F}\mathbf{m} - \mathbf{d})$. When we choose this sign convention, our hazard for analysis errors will be reduced because \mathbf{F} is often complicated and formed by combining many parts.

Beginners often feel disappointment when the data does not fit the model very well. They see it as a defect in the data instead of an opportunity to design a stronger theory.

2.3.2 Method of random directions and steepest descent

Let us minimize the sum of the squares of the components of the **residual** vector given by

$$\text{residual} = \text{transform} \quad \text{model space} - \text{data space} \quad (2.42)$$

$$\begin{bmatrix} \mathbf{r} \end{bmatrix} = \begin{bmatrix} \mathbf{F} \end{bmatrix} \begin{bmatrix} \mathbf{x} \end{bmatrix} - \begin{bmatrix} \mathbf{d} \end{bmatrix} \quad (2.43)$$

A **contour plot** is based on an altitude function of space. The altitude is the **dot product** $\mathbf{r} \cdot \mathbf{r}$. By finding the lowest altitude, we are driving the residual vector \mathbf{r} as close as we can to zero. If the residual vector \mathbf{r} reaches zero, then we have solved the simultaneous equations $\mathbf{d} = \mathbf{F}\mathbf{x}$. In a two-dimensional world the vector \mathbf{x} has two components, (x_1, x_2) . A contour is a curve of constant $\mathbf{r} \cdot \mathbf{r}$ in (x_1, x_2) -space. These contours have a statistical interpretation as contours of uncertainty in (x_1, x_2) , with measurement errors in \mathbf{d} .

Let us see how a random search-direction can be used to reduce the residual $0 \approx \mathbf{r} = \mathbf{F}\mathbf{x} - \mathbf{d}$. Let $\Delta\mathbf{x}$ be an abstract vector with the same number of components as the solution \mathbf{x} , and let $\Delta\mathbf{x}$ contain arbitrary or random numbers. We add an unknown quantity α of vector $\Delta\mathbf{x}$ to the vector \mathbf{x} , and thereby create \mathbf{x}_{new} :

$$\mathbf{x}_{\text{new}} = \mathbf{x} + \alpha\Delta\mathbf{x} \quad (2.44)$$

This gives a new residual:

$$\mathbf{r}_{\text{new}} = \mathbf{F} \mathbf{x}_{\text{new}} - \mathbf{d} \quad (2.45)$$

$$\mathbf{r}_{\text{new}} = \mathbf{F}(\mathbf{x} + \alpha\Delta\mathbf{x}) - \mathbf{d} \quad (2.46)$$

$$\mathbf{r}_{\text{new}} = \mathbf{r} + \alpha\Delta\mathbf{r} = (\mathbf{F}\mathbf{x} - \mathbf{d}) + \alpha\mathbf{F}\Delta\mathbf{x} \quad (2.47)$$

which defines $\Delta\mathbf{r} = \mathbf{F}\Delta\mathbf{x}$.

Next we adjust α to minimize the dot product: $\mathbf{r}_{\text{new}} \cdot \mathbf{r}_{\text{new}}$

$$(\mathbf{r} + \alpha \Delta \mathbf{r}) \cdot (\mathbf{r} + \alpha \Delta \mathbf{r}) = \mathbf{r} \cdot \mathbf{r} + 2\alpha(\mathbf{r} \cdot \Delta \mathbf{r}) + \alpha^2 \Delta \mathbf{r} \cdot \Delta \mathbf{r} \quad (2.48)$$

Set to zero its derivative with respect to α using the chain rule

$$0 = (\mathbf{r} + \alpha \Delta \mathbf{r}) \cdot \Delta \mathbf{r} + \Delta \mathbf{r} \cdot (\mathbf{r} + \alpha \Delta \mathbf{r}) = 2(\mathbf{r} + \alpha \Delta \mathbf{r}) \cdot \Delta \mathbf{r} \quad (2.49)$$

which says that the new residual $\mathbf{r}_{\text{new}} = \mathbf{r} + \alpha \Delta \mathbf{r}$ is perpendicular to the “fitting function” $\Delta \mathbf{r}$. Solving gives the required value of α .

$$\alpha = -\frac{(\mathbf{r} \cdot \Delta \mathbf{r})}{(\Delta \mathbf{r} \cdot \Delta \mathbf{r})} \quad (2.50)$$

A “computation **template**” for the method of random directions is

```

 $\mathbf{r} \leftarrow \mathbf{F}\mathbf{x} - \mathbf{d}$ 
iterate {
     $\Delta \mathbf{x} \leftarrow$  random numbers
     $\Delta \mathbf{r} \leftarrow \mathbf{F} \Delta \mathbf{x}$ 
     $\alpha \leftarrow -(\mathbf{r} \cdot \Delta \mathbf{r})/(\Delta \mathbf{r} \cdot \Delta \mathbf{r})$ 
     $\mathbf{x} \leftarrow \mathbf{x} + \alpha \Delta \mathbf{x}$ 
     $\mathbf{r} \leftarrow \mathbf{r} + \alpha \Delta \mathbf{r}$ 
}
```

A nice thing about the method of random directions is that you do not need to know the adjoint operator \mathbf{F}' .

In practice, random directions are rarely used. It is more common to use the **gradient** direction than a random direction. Notice that a vector of the size of $\Delta \mathbf{x}$ is

$$\mathbf{g} = \mathbf{F}'\mathbf{r} \quad (2.51)$$

Notice also that this vector can be found by taking the gradient of the size of the residuals:

$$\frac{\partial}{\partial \mathbf{x}'} \mathbf{r} \cdot \mathbf{r} = \frac{\partial}{\partial \mathbf{x}'} (\mathbf{x}' \mathbf{F}' - \mathbf{d}') (\mathbf{F} \mathbf{x} - \mathbf{d}) = \mathbf{F}' \mathbf{r} \quad (2.52)$$

Choosing $\Delta \mathbf{x}$ to be the gradient vector $\Delta \mathbf{x} = \mathbf{g} = \mathbf{F}'\mathbf{r}$ is called “the method of **steepest descent**.”

Starting from a model $\mathbf{x} = \mathbf{m}$ (which may be zero), below is a **template** of pseudocode for minimizing the residual $\mathbf{0} \approx \mathbf{r} = \mathbf{F}\mathbf{x} - \mathbf{d}$ by the steepest-descent method:

```

 $\mathbf{r} \leftarrow \mathbf{F}\mathbf{x} - \mathbf{d}$ 
iterate {
     $\Delta \mathbf{x} \leftarrow \mathbf{F}' \mathbf{r}$ 
     $\Delta \mathbf{r} \leftarrow \mathbf{F} \Delta \mathbf{x}$ 
     $\alpha \leftarrow -(\mathbf{r} \cdot \Delta \mathbf{r})/(\Delta \mathbf{r} \cdot \Delta \mathbf{r})$ 
     $\mathbf{x} \leftarrow \mathbf{x} + \alpha \Delta \mathbf{x}$ 
     $\mathbf{r} \leftarrow \mathbf{r} + \alpha \Delta \mathbf{r}$ 
}
```

2.3.3 Null space and iterative methods

In applications where we fit $\mathbf{d} \approx \mathbf{F}\mathbf{x}$, there might exist a vector (or a family of vectors) defined by the condition $\mathbf{0} = \mathbf{F}\mathbf{x}_{\text{null}}$. This family is called a **null space**. For example, if the operator \mathbf{F} is a time derivative, then the null space is the constant function; if the operator is a second derivative, then the null space has two components, a constant function and a linear function, or combinations of them. The null space is a family of model components that have no effect on the data.

When we use the steepest-descent method, we iteratively find solutions by this updating:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha \Delta \mathbf{x} \quad (2.53)$$

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha \mathbf{F}' \mathbf{r} \quad (2.54)$$

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha \mathbf{F}'(\mathbf{F}\mathbf{x} - \mathbf{d}) \quad (2.55)$$

After we have iterated to convergence, the gradient $\Delta \mathbf{x}$ vanishes as does $\mathbf{F}'(\mathbf{F}\mathbf{x} - \mathbf{d})$. Thus, an iterative solver gets the same solution as the long-winded theory leading to equation (2.30).

Suppose that by adding a huge amount of \mathbf{x}_{null} , we now change \mathbf{x} and continue iterating. Notice that $\Delta \mathbf{x}$ remains zero because $\mathbf{F}\mathbf{x}_{\text{null}}$ vanishes. Thus we conclude that any null space in the initial guess \mathbf{x}_0 will remain there unaffected by the gradient-descent process.

Linear algebra theory enables us to dig up the entire null space should we so desire. On the other hand, the computer demands might be vast. Even the memory for holding the many \mathbf{x} vectors could be prohibitive. A much simpler and more practical goal is to find out if the null space has any members, and if so, to view some of them. To try to see a member of the null space, we take two starting guesses and we run our iterative solver for each of them. If the two solutions, \mathbf{x}_1 and \mathbf{x}_2 , are the same, there is no null space. If the solutions differ, the difference is a member of the null space. Let us see why: Suppose after iterating to minimum residual we find

$$\mathbf{r}_1 = \mathbf{F}\mathbf{x}_1 - \mathbf{d} \quad (2.56)$$

$$\mathbf{r}_2 = \mathbf{F}\mathbf{x}_2 - \mathbf{d} \quad (2.57)$$

We know that the residual squared is a convex quadratic function of the unknown \mathbf{x} . Mathematically that means the minimum value is unique, so $\mathbf{r}_1 = \mathbf{r}_2$. Subtracting we find $\mathbf{0} = \mathbf{r}_1 - \mathbf{r}_2 = \mathbf{F}(\mathbf{x}_1 - \mathbf{x}_2)$ proving that $\mathbf{x}_1 - \mathbf{x}_2$ is a model in the null space. Adding $\mathbf{x}_1 - \mathbf{x}_2$ to any to any model \mathbf{x} will not change the theoretical data. Are you having trouble visualizing \mathbf{r} being unique, but \mathbf{x} not being unique? Imagine that \mathbf{r} happens to be independent of one of the components of \mathbf{x} . That component is nonunique. More generally, it is some linear combination of components of \mathbf{x} that \mathbf{r} is independent of.

A practical way to learn about the existence of null spaces and their general appearance is simply to try gradient-descent methods beginning from various different starting guesses.

“Did I fail to run my iterative solver long enough?” is a question you might have. If two residuals from two starting solutions are not equal, $\mathbf{r}_1 \neq \mathbf{r}_2$, then you should be running your solver through more iterations.

If two different starting solutions produce two different residuals, then you didn't run your solver through enough iterations.

2.3.4 Why steepest descent is so slow

Before we can understand why the **conjugate-direction method** is so fast, we need to see why the **steepest-descent method** is so slow. Imagine yourself sitting on the edge of a circular bowl. If you jump off the rim, you'll slide straight to the bottom at the center. Now imagine an ellipsoidal bowl of very large ellipticity. As you jump off the rim, you'll first move in the direction of the gradient. This is not towards the bottom at the center of the ellipse (unless you were sitting on the major or minor axis).

We can formalize the situation. A parametric equation for a line is $\mathbf{x} = \mathbf{x}_{\text{old}} + \alpha \Delta \mathbf{x}$ where α is the parameter for moving on the line. The process of selecting α is called "**line search**." Think of a two-dimensional example where the vector of unknowns \mathbf{x} has just two components, x_1 and x_2 . Then the size of the residual vector $\mathbf{r} \cdot \mathbf{r}$ can be displayed with a contour plot in the plane of (x_1, x_2) . Our ellipsoidal bowl has ellipsoidal contours of constant altitude. As we move in a line across this space by adjusting α , equation(2.48) gives our altitude. This equation has a unique minimum because it is a parabola in α . As we approach the minimum, our trajectory becomes tangential to a contour line in (x_1, x_2) -space. This is where we stop. Now we compute our new residual \mathbf{r} and we compute the new gradient $\Delta \mathbf{x} = \mathbf{g} = \mathbf{F}'\mathbf{r}$. OK, we are ready for the next slide down. When we turn ourselves from "parallel to a contour line" to the direction of $\Delta \mathbf{x}$ which is "perpendicular to that contour", we are turning 90° . Our path to the bottom of the bowl will be made of many segments, each turning 90° from the previous. We will need an infinite number of such steps to reach the bottom. It happens that the amazing conjugate-direction method would reach the bottom in just two jumps (because (x_1, x_2) is a two dimensional space.)

2.3.5 Conjugate direction

In the **conjugate-direction method**, not a line, but rather a plane, is searched. A plane is made from an arbitrary linear combination of two vectors. One vector will be chosen to be the gradient vector, say \mathbf{g} . The other vector will be chosen to be the previous descent step vector, say $\mathbf{s} = \mathbf{x}_j - \mathbf{x}_{j-1}$. Instead of $\alpha \mathbf{g}$ we need a linear combination, say $\alpha \mathbf{g} + \beta \mathbf{s}$. For minimizing quadratic functions the plane search requires only the solution of a two-by-two set of linear equations for α and β . The equations will be specified here along with the program. (For *nonquadratic* functions a plane search is considered intractable, whereas a line search proceeds by bisection.)

For use in linear problems, the conjugate-direction method described in this book follows an identical path with the more well-known conjugate-gradient method. We use the conjugate-direction method for convenience in exposition and programming.

The simple form of the conjugate-direction algorithm covered here is a sequence of steps. In each step the minimum is found in the plane given by two vectors: the gradient vector and the vector of the previous step.

Given the linear operator \mathbf{F} and a generator of solution steps (random in the case of random directions or gradient in the case of steepest descent), we can construct an optimally convergent iteration process, which finds the solution in no more than n steps, where n is the size of the problem. This result should not be surprising. If \mathbf{F} is represented by a full matrix, then the cost of direct inversion is proportional to n^3 , and the cost of matrix multiplication is n^2 . Each step of an iterative method boils down to a matrix multiplication. Therefore, we need at least n steps to arrive at the exact solution. Two circumstances make large-scale optimization practical. First, for sparse convolution matrices the cost of matrix multiplication is n instead of n^2 . Second, we can often find a reasonably good solution after a limited number of iterations. If both these conditions are met, the cost of optimization grows linearly with n , which is a practical rate even for very large problems.

Fourier-transformed variables are often capitalized. This convention will be helpful here, so in this subsection only, we capitalize vectors transformed by the \mathbf{F} matrix. As everywhere, a matrix such as \mathbf{F} is printed in **boldface** type but in this subsection, vectors are *not* printed in boldface print. Thus we define the solution, the solution step (from one iteration to the next), and the gradient by

$$X = \mathbf{F} x \quad \text{solution} \quad (2.58)$$

$$S_j = \mathbf{F} s_j \quad \text{solution step} \quad (2.59)$$

$$G_j = \mathbf{F} g_j \quad \text{solution gradient} \quad (2.60)$$

A linear combination in solution space, say $s + g$, corresponds to $S + G$ in the conjugate space, because $S + G = \mathbf{F}s + \mathbf{F}g = \mathbf{F}(s + g)$. According to equation (2.43), the residual is the theoretical data minus the observed data.

$$R = \mathbf{F}x - D = X - D \quad (2.61)$$

The solution x is obtained by a succession of steps s_j , say

$$x = s_1 + s_2 + s_3 + \cdots \quad (2.62)$$

The last stage of each iteration is to update the solution and the residual:

$$\text{solution update :} \quad x \leftarrow x + s \quad (2.63)$$

$$\text{residual update :} \quad R \leftarrow R + S \quad (2.64)$$

The *gradient* vector g is a vector with the same number of components as the solution vector x . A vector with this number of components is

$$g = \mathbf{F}' R = \text{gradient} \quad (2.65)$$

$$G = \mathbf{F} g = \text{conjugate gradient} \quad (2.66)$$

The gradient g in the transformed space is G , also known as the **conjugate gradient**.

The minimization (2.48) is now generalized to scan not only the line with α , but simultaneously another line with β . The combination of the two lines is a plane:

$$Q(\alpha, \beta) = (R + \alpha G + \beta S) \cdot (R + \alpha G + \beta S) \quad (2.67)$$

The minimum is found at $\partial Q/\partial\alpha = 0$ and $\partial Q/\partial\beta = 0$, namely,

$$0 = G \cdot (R + \alpha G + \beta S) \quad (2.68)$$

$$0 = S \cdot (R + \alpha G + \beta S) \quad (2.69)$$

The solution is

$$\begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \frac{-1}{(G \cdot G)(S \cdot S) - (G \cdot S)^2} \begin{bmatrix} (S \cdot S) & -(S \cdot G) \\ -(G \cdot S) & (G \cdot G) \end{bmatrix} \begin{bmatrix} (G \cdot R) \\ (S \cdot R) \end{bmatrix} \quad (2.70)$$

This may look complicated. The steepest descent method requires us to compute only the two dot products $\mathbf{r} \cdot \Delta\mathbf{r}$ and $\Delta\mathbf{r} \cdot \Delta\mathbf{r}$ while equation (2.67) contains five dot products, but the extra trouble is well worth while because the “conjugate direction” is such a much better direction than the gradient direction.

The many applications in this book all need to find α and β with (2.70) and then update the solution with (2.63) and update the residual with (2.64). Thus we package these activities in a subroutine named **cgstep**. To use that subroutine we will have a computation **template** like we had for steepest descents, except that we will have the repetitive work done by subroutine **cgstep**. This template (or pseudocode) for minimizing the residual $\mathbf{0} \approx \mathbf{r} = \mathbf{F}\mathbf{x} - \mathbf{d}$ by the conjugate-direction method is

```

 $\mathbf{r} \leftarrow \mathbf{F}\mathbf{x} - \mathbf{d}$ 
iterate {
     $\Delta\mathbf{x} \leftarrow \mathbf{F}' \mathbf{r}$ 
     $\Delta\mathbf{r} \leftarrow \mathbf{F} \Delta\mathbf{x}$ 
     $(\mathbf{x}, \mathbf{r}) \leftarrow \text{cgstep}(\mathbf{x}, \Delta\mathbf{x}, \mathbf{r}, \Delta\mathbf{r})$ 
}
```

where the subroutine **cgstep()** remembers the previous iteration and works out the step size and adds in the proper proportion of the $\Delta\mathbf{x}$ of the previous step.

2.3.6 Routine for one step of conjugate-direction descent

The conjugate vectors G and S in the program are denoted by **gg** and **ss**. The inner part of the conjugate-direction task is in function **cgstep()**.

Observe the **cgstep()** function has a logical parameter called **forget**. This parameter does not need to be input. In the normal course of things, **forget** will be true on the first iteration and false on subsequent iterations. This refers to the fact that on the first iteration, there is no previous step, so the conjugate direction method is reduced to the steepest descent method. At any iteration, however, you have the option to set **forget=true** which amounts to restarting the calculation from the current location, something we rarely find reason to do.

2.3.7 A basic solver program

There are many different methods for iterative least-square estimation some of which will be discussed later in this book. The conjugate-gradient (CG) family (including the first

filt/lib/cgstep.c

```

51  if (forget) {
52      for (i = 0; i < nx; i++) S[i] = 0.;
53      for (i = 0; i < ny; i++) Ss[i] = 0.;
54      beta = 0.0;
55      alfa = cblas_dsdot( ny, gg, 1, gg, 1);
56      if (alfa <= 0.) return;
57      alfa = - cblas_dsdot( ny, gg, 1, rr, 1) / alfa;
58  } else {
59      /* search plane by solving 2-by-2
60       G . (R - G*alfa - S*beta) = 0
61       S . (R - G*alfa - S*beta) = 0 */
62      gdg = cblas_dsdot( ny, gg, 1, gg, 1);
63      sds = cblas_dsdot( ny, Ss, 1, Ss, 1);
64      gds = cblas_dsdot( ny, gg, 1, Ss, 1);
65      if (gdg == 0. || sds == 0.) return;
66      determ = 1.0 - (gds/gdg)*(gds/sds);
67      if (determ > EPSILON) determ *= gdg * sds;
68      else determ = gdg * sds * EPSILON;
69      gdr = - cblas_dsdot( ny, gg, 1, rr, 1);
70      sdr = - cblas_dsdot( ny, Ss, 1, rr, 1);
71      alfa = ( sds * gdr - gds * sdr ) / determ;
72      beta = (-gds * gdr + gdg * sdr ) / determ;
73  }
74  cblas_sscal(nx,beta,S,1);
75  cblas_saxpy(nx,alfa,g,1,S,1);
76
77  cblas_sscal(ny,beta,Ss,1);
78  cblas_saxpy(ny,alfa,gg,1,Ss,1);
79
80  for (i = 0; i < nx; i++) {
81      x[i] += S[i];
82  }
83  for (i = 0; i < ny; i++) {
84      rr[i] += Ss[i];
85  }

```

order conjugate-direction method described above) share the property that theoretically they achieve the solution in n iterations, where n is the number of unknowns. The various CG methods differ in their numerical errors, memory required, adaptability to non-linear optimization, and their requirements on accuracy of the adjoint. What we do in this section is to show you the generic interface.

None of us is an expert in both geophysics and in optimization theory (OT), yet we need to handle both. We would like to have each group write its own code with a relatively easy interface. The problem is that the OT codes must invoke the physical operators yet the OT codes should not need to deal with all the data and parameters needed by the physical operators.

In other words, if a practitioner decides to swap one solver for another, the only thing needed is the name of the new solver.

The operator entrance is for the geophysicist, who formulates the estimation problem. The solver entrance is for the specialist in numerical algebra, who designs a new optimization method. The C programming language allows us to achieve this design goal by means of generic function interfaces.

A basic solver is `tinysolver`.

The two most important arguments in `tinysolver()` are the operator function `Fop`, which is defined by the interface from Chapter 1, and the stepper function `stepper`, which implements one step of an iterative estimation. For example, a practitioner who chooses to use our new `cgstep()` on page 51 for iterative solving the operator `matmult` on the current page would write the call

```
tinysolver ( matmult_lop, cgstep, ...
```

so while you are reading the `tinysolver` module, you should visualize the `Fop()` function as being `matmult_lop`, and you should visualize the `stepper()` function as being `cgstep`.

The other required parameters to `tinysolver()` are `d` (the data we want to fit), `m` (the model we want to estimate), and `niter` (the maximum number of iterations). There are also a couple of optional arguments. For example, `m0` is the starting guess for the model. If this parameter is omitted, the model is initialized to zero. To output the final residual vector, we include a parameter called `resd`, which is optional as well. We will watch how the list of optional parameters to the generic solver routine grows as we attack more and more complex problems in later chapters.

2.3.8 Test case: solving some simultaneous equations

Now we assemble a module `cgtest` for solving simultaneous equations. Starting with the conjugate-direction module `cgstep` on page 51 we insert the module `matmult` on page 4 as the linear operator.

Below shows the solution to 5×4 set of simultaneous equations. Observe that the “exact” solution is obtained in the last step. Because the data and answers are integers, it is quick to check the result manually.

filt/lib/tinysolver.c

```

23 void sf_tinysolver (sf_operator Fop      /* linear operator */,
24                    sf_solverstep stepper /* stepping function */,
25                    int nm                /* size of model */,
26                    int nd                /* size of data */,
27                    float* m              /* estimated model */,
28                    const float* m0       /* starting model */,
29                    const float* d        /* data */,
30                    int niter              /* iterations */)
31 /*< Generic linear solver. Solves oper{x} =~ dat >*/
32 {
33     int i, iter;
34     float *g, *rr, *gg;
35
36     g = sf_floatalloc (nm);
37     rr = sf_floatalloc (nd);
38     gg = sf_floatalloc (nd);
39
40     for (i=0; i < nd; i++) rr[i] = - d[i];
41     if (NULL==m0) {
42         for (i=0; i < nm; i++) m[i] = 0.0;
43     } else {
44         for (i=0; i < nm; i++) m[i] = m0[i];
45         Fop (false, true, nm, nd, m, rr);
46     }
47
48     for (iter=0; iter < niter; iter++) {
49         Fop (true, false, nm, nd, g, rr);
50         Fop (false, false, nm, nd, g, gg);
51
52         stepper (false, nm, nd, m, g, rr, gg);
53     }
54
55     free (g);
56     free (rr);
57     free (gg);

```

```

                                user/fomels/cgtest.c
23 void cgtest(int nx, int ny, float *x,
24             const float *yy, float **fff, int niter)
25 /*< testing conjugate gradients with matrix multiplication >*/
26 {
27     matmult_init( fff);
28     sf_tinysolver( matmult_lop, sf_cgstep,
29                  nx, ny, x, NULL, yy, niter);
30     sf_cgstep_close();
31 }

```

```

d transpose
    3.00    3.00    5.00    7.00    9.00

```

```

F transpose
    1.00    1.00    1.00    1.00    1.00
    1.00    2.00    3.00    4.00    5.00
    1.00    0.00    1.00    0.00    1.00
    0.00    0.00    0.00    1.00    1.00

```

```

for iter = 0, 4
x   0.43457383  1.56124675  0.27362058  0.25752524
res -0.73055887  0.55706739  0.39193487 -0.06291389 -0.22804642
x   0.51313990  1.38677299  0.87905121  0.56870615
res -0.22103602  0.28668585  0.55251014 -0.37106210 -0.10523783
x   0.39144871  1.24044561  1.08974111  1.46199656
res -0.27836466 -0.12766013  0.20252672 -0.18477242  0.14541438
x   1.00001287  1.00004792  1.00000811  2.00000739
res  0.00006878  0.00010860  0.00016473  0.00021179  0.00026788
x   1.00000024  0.99999994  0.99999994  2.00000024
res -0.00000001 -0.00000001  0.00000001  0.00000002 -0.00000001

```

EXERCISES:


- 1 One way to remove a mean value m from signal $s(t) = \mathbf{s}$ is with the fitting goal $\mathbf{0} \approx \mathbf{s} - m$. What operator matrix is involved?
- 2 What linear operator subroutine from Chapter 1 can be used for finding the mean?
- 3 How many CD iterations should be required to get the exact mean value?
- 4 Write a mathematical expression for finding the mean by the CG method.

2.4 INVERSE NMO STACK

To illustrate an example of solving a huge set of simultaneous equations without ever writing down the matrix of coefficients we consider how *back projection* can be upgraded towards *inversion* in the application called **moveout and stack**.

Figure 2.3: Top is a model trace \mathbf{m} . Next are the synthetic data traces, $\mathbf{d} = \mathbf{M}\mathbf{m}$. Then, labeled `niter=0` is the **stack**, a result of processing by adjoint modeling. Increasing values of `niter` show \mathbf{x} as a function of iteration count in the fitting goal $\mathbf{d} \approx \mathbf{M}\mathbf{x}$. (Carlos Cunha-Filho)

lsq/invstack invstack



The seismograms at the bottom of Figure 2.3 show the first four iterations of conjugate-direction inversion. You see the original rectangle-shaped waveform returning as the iterations proceed. Notice also on the **stack** that the early and late events have unequal amplitudes, but after enough iterations they are equal, as they began. Mathematically, we can denote the top trace as the model \mathbf{m} , the synthetic data signals as $\mathbf{d} = \mathbf{M}\mathbf{m}$, and the stack as $\mathbf{M}'\mathbf{d}$. The conjugate-gradient algorithm optimizes the fitting goal $\mathbf{d} \approx \mathbf{M}\mathbf{x}$ by variation of \mathbf{x} , and the figure shows \mathbf{x} converging to \mathbf{m} . Because there are 256 unknowns in \mathbf{m} , it is gratifying to see good convergence occurring after the first four iterations. The fitting is done by module `invstack`, which is just like `cgtest` on page 54 except that the matrix-multiplication operator `matmult` on page 4 has been replaced by `imospray`. Studying the program, you can deduce that, except for a scale factor, the output at `niter=0` is identical to the stack $\mathbf{M}'\mathbf{d}$. All the signals in Figure 2.3 are intrinsically the same scale.

filt/proc/invstack.c

```

24 void invstack(int nt, float *model, int nx, const float *gather,
25               float t0, float x0,
26               float dt, float dx, float slow, int niter)
27 /*< NMO stack by inverse of forward modeling */
28 {
29     imospray_init( slow, x0, dx, t0, dt, nt, nx);
30     sf_tinysolver( imospray_lop, sf_cgstep,
31                   nt, nt*nx, model, NULL, gather, niter);
32     sf_cgstep_close ();
33     imospray_close (); /* garbage collection */
34 }

```

This simple inversion is inexpensive. Has anything been gained over conventional stack? First, though we used **nearest-neighbor** interpolation, we managed to preserve the spectrum of the input, apparently all the way to the Nyquist frequency. Second, we preserved the true amplitude scale without ever bothering to think about (1) dividing by the number of contributing traces, (2) the amplitude effect of NMO stretch, or (3) event truncation.

With depth-dependent velocity, wave fields become much more complex at wide offset. NMO soon fails, but wave-equation forward modeling offers interesting opportunities for inversion.

2.5 THE WORLD OF CONJUGATE GRADIENTS

Nonlinearity arises in two ways: First, theoretical data might be a nonlinear function of the model parameters. Second, observed data could contain imperfections that force us to use **nonlinear methods** of statistical estimation.

2.5.1 Physical nonlinearity

When standard methods of physics relate theoretical data $\mathbf{d}_{\text{theor}}$ to model parameters \mathbf{m} , they often use a nonlinear relation, say $\mathbf{d}_{\text{theor}} = \mathbf{f}(\mathbf{m})$. The power-series approach then leads to representing theoretical data as

$$\mathbf{d}_{\text{theor}} = \mathbf{f}(\mathbf{m}_0 + \Delta\mathbf{m}) \approx \mathbf{f}(\mathbf{m}_0) + \mathbf{F}\Delta\mathbf{m} \quad (2.71)$$

where \mathbf{F} is the matrix of partial derivatives of data values by model parameters, say $\partial d_i / \partial m_j$, evaluated at \mathbf{m}_0 . The theoretical data $\mathbf{d}_{\text{theor}}$ minus the observed data \mathbf{d}_{obs} is the residual we minimize.

$$\mathbf{0} \approx \mathbf{d}_{\text{theor}} - \mathbf{d}_{\text{obs}} = \mathbf{F}\Delta\mathbf{m} + [\mathbf{f}(\mathbf{m}_0) - \mathbf{d}_{\text{obs}}] \quad (2.72)$$

$$\mathbf{r}_{\text{new}} = \mathbf{F}\Delta\mathbf{m} + \mathbf{r}_{\text{old}} \quad (2.73)$$

It is worth noticing that the residual updating (2.73) in a nonlinear problem is the same as that in a linear problem (2.47). If you make a large step $\Delta\mathbf{m}$, however, the new residual will be different from that expected by (2.73). Thus you should always re-evaluate the residual vector at the new location, and if you are reasonably cautious, you should be sure the residual norm has actually decreased before you accept a large step.

The pathway of inversion with physical nonlinearity is well developed in the academic literature and Bill **Symes** at Rice University has a particularly active group.

2.5.2 Statistical nonlinearity

The data itself often has **noise bursts** or **gaps**, and we will see later in Chapter ?? that this leads us to readjusting the **weighting function**. In principle, we should fix the weighting function and solve the problem. Then we should revise the weighting function and solve the problem again. In practice we find it convenient to change the weighting function during the optimization descent. Failure is possible when the weighting function is changed too rapidly or drastically. (The proper way to solve this problem is with robust estimators. Unfortunately, I do not yet have an all-purpose robust solver. Thus we are (temporarily, I hope) reduced to using crude reweighted least-squares methods. Sometimes they work and sometimes they don't.)

2.5.3 Coding nonlinear fitting problems

We can solve nonlinear least-squares problems in about the same way as we do iteratively reweighted ones. A simple adaptation of a linear method gives us a **nonlinear solver** if the residual is recomputed at each iteration. Omitting the weighting function (for simplicity) the **template** is:

```

iterate {
     $\mathbf{r} \leftarrow \mathbf{f}(\mathbf{m}) - \mathbf{d}$ 
    Define  $\mathbf{F} = \partial \mathbf{d} / \partial \mathbf{m}$ .
     $\Delta \mathbf{m} \leftarrow \mathbf{F}' \mathbf{r}$ 
     $\Delta \mathbf{r} \leftarrow \mathbf{F} \Delta \mathbf{m}$ 
     $(\mathbf{m}, \mathbf{r}) \leftarrow \text{step}(\mathbf{m}, \mathbf{r}, \Delta \mathbf{m}, \Delta \mathbf{r})$ 
}

```

A formal theory for the optimization exists, but we are not using it here. The assumption we make is that the step size will be small, so that familiar line-search and plane-search approximations should succeed in reducing the residual. Unfortunately this assumption is not reliable. What we should do is test that the residual really does decrease, and if it does not we should revert to steepest descent with a smaller step size. Perhaps we should test an incremental variation on the status quo: where inside **solver** on page 53, we check to see if the residual diminished in the *previous* step, and if it did not, restart the iteration (choose the *current* step to be steepest descent instead of CD). I am planning to work with some mathematicians to gain experience with other solvers.

Experience shows that nonlinear problems have many pitfalls. Start with a linear problem, add a minor physical improvement or unnormal noise, and the problem becomes nonlinear and probably has another solution far from anything reasonable. When solving such a nonlinear problem, we cannot arbitrarily begin from zero as we do with linear problems. We must choose a reasonable starting guess, and then move in a stable and controlled manner. A simple solution is to begin with several steps of steepest descent and then switch over to do some more steps of CD. Avoiding CD in earlier iterations can avoid instability. Strong linear “regularization” discussed later can also reduce the effect of nonlinearity.

2.5.4 Standard methods

The conjugate-direction method is really a family of methods. Mathematically, where there are n unknowns, these algorithms all converge to the answer in n (or fewer) steps. The various methods differ in numerical accuracy, treatment of underdetermined systems, accuracy in treating ill-conditioned systems, space requirements, and numbers of dot products. Technically, the method of CD used in the **cgstep** module on page 51 is not the conjugate-gradient method itself, but is equivalent to it. This method is more properly called the **conjugate-direction method** with a memory of one step. I chose this method for its clarity and flexibility. If you would like a free introduction and summary of conjugate-gradient methods, I particularly recommend *An Introduction to Conjugate Gradient Method Without Agonizing Pain* by Jonathon Shewchuk, which you can download¹.

¹<http://www.cs.cmu.edu/afs/cs/project/quake/public/papers/painless-conjugate-gradient.ps>

I suggest you skip over the remainder of this section and return after you have seen many examples and have developed some expertise, and have some technical problems.

The **conjugate-gradient method** was introduced by **Hestenes** and **Stiefel** in 1952. To read the standard literature and relate it to this book, you should first realize that when I write fitting goals like

$$0 \approx \mathbf{W}(\mathbf{F}\mathbf{m} - \mathbf{d}) \quad (2.74)$$

$$0 \approx \mathbf{A}\mathbf{m}, \quad (2.75)$$

they are equivalent to minimizing the quadratic form:

$$\mathbf{m} : \quad \min_{\mathbf{m}} Q(\mathbf{m}) = (\mathbf{m}'\mathbf{F}' - \mathbf{d}')\mathbf{W}'\mathbf{W}(\mathbf{F}\mathbf{m} - \mathbf{d}) + \mathbf{m}'\mathbf{A}'\mathbf{A}\mathbf{m} \quad (2.76)$$

The optimization theory (OT) literature starts from a minimization of

$$\mathbf{x} : \quad \min_{\mathbf{x}} Q(\mathbf{x}) = \mathbf{x}'\mathbf{H}\mathbf{x} - \mathbf{b}'\mathbf{x} \quad (2.77)$$

To relate equation (2.76) to (2.77) we expand the parentheses in (2.76) and abandon the constant term $\mathbf{d}'\mathbf{d}$. Then gather the quadratic term in \mathbf{m} and the linear term in \mathbf{m} . There are two terms linear in \mathbf{m} that are transposes of each other. They are scalars so they are equal. Thus, to invoke “standard methods,” you take your problem-formulation operators \mathbf{F} , \mathbf{W} , \mathbf{A} and create two modules that apply the operators

$$\mathbf{H} = \mathbf{F}'\mathbf{W}'\mathbf{W}\mathbf{F} + \mathbf{A}'\mathbf{A} \quad (2.78)$$

$$\mathbf{b}' = 2(\mathbf{F}'\mathbf{W}'\mathbf{W}\mathbf{d})' \quad (2.79)$$

The operators \mathbf{H} and \mathbf{b}' operate on model space. Standard procedures do not require their adjoints because \mathbf{H} is its own adjoint and \mathbf{b}' reduces model space to a scalar. You can see that computing \mathbf{H} and \mathbf{b}' requires one temporary space the size of data space (whereas **cgstep** requires two).

When people have trouble with conjugate gradients or conjugate directions, I always refer them to the **Paige and Saunders algorithm LSQR**. Methods that form \mathbf{H} explicitly or implicitly (including both the standard literature and the book3 method) square the condition number, that is, they are twice as susceptible to rounding error as is **LSQR**.

2.5.5 Understanding CG magic and advanced methods

This section includes Sergey Fomel’s explanation on the “magic” convergence properties of the conjugate-direction methods. It also presents a classic version of conjugate gradients, which can be found in numerous books on least-square optimization.

The key idea for constructing an optimal iteration is to update the solution at each step in the direction, composed by a linear combination of the current direction and all previous solution steps. To see why this is a helpful idea, let us consider first the method of random directions. Substituting expression (2.50) into formula (2.48), we see that the residual power decreases at each step by

$$\mathbf{r} \cdot \mathbf{r} - \mathbf{r}_{\text{new}} \cdot \mathbf{r}_{\text{new}} = \frac{(\mathbf{r} \cdot \Delta \mathbf{r})^2}{(\Delta \mathbf{r} \cdot \Delta \mathbf{r})}. \quad (2.80)$$

To achieve a better convergence, we need to maximize the right hand side of (2.80). Let us define a new solution step \mathbf{s}_{new} as a combination of the current direction $\Delta\mathbf{x}$ and the previous step \mathbf{s} , as follows:

$$\mathbf{s}_{\text{new}} = \Delta\mathbf{x} + \beta\mathbf{s} . \quad (2.81)$$

The solution update is then defined as

$$\mathbf{x}_{\text{new}} = \mathbf{x} + \alpha\mathbf{s}_{\text{new}} . \quad (2.82)$$

The formula for α (2.50) still holds, because we have preserved in (2.82) the form of equation (2.44) and just replaced $\Delta\mathbf{x}$ with \mathbf{s}_{new} . In fact, formula (2.50) can be simplified a little bit. From (2.49), we know that \mathbf{r}_{new} is orthogonal to $\Delta\mathbf{r} = \mathbf{F}\mathbf{s}_{\text{new}}$. Likewise, \mathbf{r} should be orthogonal to $\mathbf{F}\mathbf{s}$ (recall that \mathbf{r} was \mathbf{r}_{new} and \mathbf{s} was \mathbf{s}_{new} at the previous iteration). We can conclude that

$$(\mathbf{r} \cdot \Delta\mathbf{r}) = (\mathbf{r} \cdot \mathbf{F}\mathbf{s}_{\text{new}}) = (\mathbf{r} \cdot \mathbf{F}\Delta\mathbf{x}) + \beta(\mathbf{r} \cdot \mathbf{F}\mathbf{s}) = (\mathbf{r} \cdot \mathbf{F}\Delta\mathbf{x}) . \quad (2.83)$$

Comparing (2.83) with (2.80), we can see that adding a portion of the previous step to the current direction does not change the value of the numerator in expression (2.80). However, the value of the denominator can be changed. Minimizing the denominator maximizes the residual increase at each step and leads to a faster convergence. This is the denominator minimization that constrains the value of the adjustable coefficient β in (2.81).

The procedure for finding β is completely analogous to the derivation of formula (2.50). We start with expanding the dot product $(\Delta\mathbf{r} \cdot \Delta\mathbf{r})$:

$$(\mathbf{F}\mathbf{s}_{\text{new}} \cdot \mathbf{F}\mathbf{s}_{\text{new}}) = \mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\Delta\mathbf{x} + 2\beta(\mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\mathbf{s}) + \beta^2 \mathbf{F}\mathbf{s} \cdot \mathbf{F}\mathbf{s} . \quad (2.84)$$

Differentiating with respect to β and setting the derivative to zero, we find that

$$0 = 2(\mathbf{F}\Delta\mathbf{x} + \beta\mathbf{F}\mathbf{s}) \cdot \mathbf{F}\mathbf{s} . \quad (2.85)$$

Equation (2.85) states that the *conjugate direction* $\mathbf{F}\mathbf{s}_{\text{new}}$ is orthogonal (perpendicular) to the previous conjugate direction $\mathbf{F}\mathbf{s}$. It also defines the value of β as

$$\beta = -\frac{(\mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\mathbf{s})}{(\mathbf{F}\mathbf{s} \cdot \mathbf{F}\mathbf{s})} . \quad (2.86)$$

Can we do even better? The positive quantity that we minimized in (2.84) decreased by

$$\mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\Delta\mathbf{x} - \mathbf{F}\mathbf{s}_{\text{new}} \cdot \mathbf{F}\mathbf{s}_{\text{new}} = \frac{(\mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\mathbf{s})^2}{(\mathbf{F}\mathbf{s} \cdot \mathbf{F}\mathbf{s})} \quad (2.87)$$

Can we decrease it further by adding another previous step? In general, the answer is positive, and it defines the method of conjugate directions. I will state this result without a formal proof (which uses the method of mathematical induction).

- If the new step is composed of the current direction and a combination of all the previous steps:

$$\mathbf{s}_n = \Delta\mathbf{x}_n + \sum_{i < n} \beta_i \mathbf{s}_i , \quad (2.88)$$

then the optimal convergence is achieved when

$$\beta_i = -\frac{(\mathbf{F}\Delta\mathbf{x}_n \cdot \mathbf{F}\mathbf{s}_i)}{(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} . \quad (2.89)$$

- The new conjugate direction is orthogonal to the previous ones:

$$(\mathbf{F}\mathbf{s}_n \cdot \mathbf{F}\mathbf{s}_i) = 0 \quad \text{for all } i < n \quad (2.90)$$

To see why this is an optimally convergent method, it is sufficient to notice that vectors $\mathbf{F}\mathbf{s}_i$ form an orthogonal basis in the data space. The vector from the current residual to the smallest residual also belongs to that space. If the data size is n , then n basis components (at most) are required to represent this vector, hence no more than n conjugate-direction steps are required to find the solution.

The computation template for the method of conjugate directions is

```

r  ← Fx - d
iterate {
     $\Delta\mathbf{x}$  ← random numbers
    s  ←  $\Delta\mathbf{x} + \sum_{i < n} \beta_i \mathbf{s}_i$  where  $\beta_i = -\frac{(\mathbf{F}\Delta\mathbf{x} \cdot \mathbf{F}\mathbf{s}_i)}{(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)}$ 
     $\Delta\mathbf{r}$  ← Fs
     $\alpha$   ←  $-(\mathbf{r} \cdot \Delta\mathbf{r}) / (\Delta\mathbf{r} \cdot \Delta\mathbf{r})$ 
    x  ← x +  $\alpha\mathbf{s}$ 
    r  ← r +  $\alpha\Delta\mathbf{r}$ 
}

```

What happens if we “feed” the method with gradient directions instead of just random directions? It turns out that in this case we need to remember from all the previous steps \mathbf{s}_i only the one that immediately precedes the current iteration. Let us derive a formal proof of that fact as well as some other useful formulas related to the method of *conjugate gradients*.

According to formula (2.49), the new residual \mathbf{r}_{new} is orthogonal to the conjugate direction $\Delta\mathbf{r} = \mathbf{F}\mathbf{s}_{\text{new}}$. According to the orthogonality condition (2.90), it is also orthogonal to all the previous conjugate directions. Defining $\Delta\mathbf{x}$ equal to the gradient $\mathbf{F}'\mathbf{r}$ and applying the definition of the adjoint operator, it is convenient to rewrite the orthogonality condition in the form

$$0 = (\mathbf{r}_n \cdot \mathbf{F}\mathbf{s}_i) = (\mathbf{F}'\mathbf{r}_n \cdot \mathbf{s}_i) = (\Delta\mathbf{x}_{n+1} \cdot \mathbf{s}_i) \quad \text{for all } i \leq n \quad (2.91)$$

According to formula (2.88), each solution step \mathbf{s}_i is just a linear combination of the gradient $\Delta\mathbf{x}_i$ and the previous solution steps. We deduce from formula (2.91) that

$$0 = (\Delta\mathbf{x}_n \cdot \mathbf{s}_i) = (\Delta\mathbf{x}_n \cdot \Delta\mathbf{x}_i) \quad \text{for all } i < n \quad (2.92)$$

In other words, in the method of conjugate gradients, the current gradient direction is always orthogonal to all the previous directions. The iteration process constructs not only an orthogonal basis in the data space but also an orthogonal basis in the model space, composed of the gradient directions.

Now let us take a closer look at formula (2.89). Note that $\mathbf{F}\mathbf{s}_i$ is simply related to the residual step at i -th iteration:

$$\mathbf{F}\mathbf{s}_i = \frac{\mathbf{r}_i - \mathbf{r}_{i-1}}{\alpha_i} . \quad (2.93)$$

Substituting relationship (2.93) into formula (2.89) and applying again the definition of the adjoint operator, we obtain

$$\beta_i = -\frac{\mathbf{F}\Delta\mathbf{x}_n \cdot (\mathbf{r}_i - \mathbf{r}_{i-1})}{\alpha_i(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} = -\frac{\Delta\mathbf{x}_n \cdot \mathbf{F}'(\mathbf{r}_i - \mathbf{r}_{i-1})}{\alpha_i(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} = -\frac{\Delta\mathbf{x}_n \cdot (\Delta\mathbf{x}_{i+1} - \Delta\mathbf{x}_i)}{\alpha_i(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} \quad (2.94)$$

Since the gradients $\Delta\mathbf{x}_i$ are orthogonal to each other, the dot product in the numerator is equal to zero unless $i = n - 1$. It means that only the immediately preceding step \mathbf{s}_{n-1} contributes to the definition of the new solution direction \mathbf{s}_n in (2.88). This is precisely the property of the conjugate gradient method we wanted to prove.

To simplify formula (2.94), rewrite formula (2.50) as

$$\alpha_i = -\frac{(\mathbf{r}_{i-1} \cdot \mathbf{F}\Delta\mathbf{x}_i)}{(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} = -\frac{(\mathbf{F}'\mathbf{r}_{i-1} \cdot \Delta\mathbf{x}_i)}{(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} = -\frac{(\Delta\mathbf{x}_i \cdot \Delta\mathbf{x}_i)}{(\mathbf{F}\mathbf{s}_i \cdot \mathbf{F}\mathbf{s}_i)} \quad (2.95)$$

Substituting (2.95) into (2.94), we obtain

$$\beta = -\frac{(\Delta\mathbf{x}_n \cdot \Delta\mathbf{x}_n)}{\alpha_{n-1}(\mathbf{F}\mathbf{s}_{n-1} \cdot \mathbf{F}\mathbf{s}_{n-1})} = \frac{(\Delta\mathbf{x}_n \cdot \Delta\mathbf{x}_n)}{(\Delta\mathbf{x}_{n-1} \cdot \Delta\mathbf{x}_{n-1})}. \quad (2.96)$$

The computation template for the method of conjugate gradients is then

```

r  ← Fx - d
β  ← 0
iterate {
    Δx ← F'r
    if not the first iteration β ←  $\frac{(\Delta\mathbf{x} \cdot \Delta\mathbf{x})}{\gamma}$ 
    γ  ←  $(\Delta\mathbf{x} \cdot \Delta\mathbf{x})$ 
    s  ← Δx + βs
    Δr ← Fs
    α  ←  $-\gamma/(\Delta\mathbf{r} \cdot \Delta\mathbf{r})$ 
    x  ← x + αs
    r  ← r + αΔr
}

```

2.6 REFERENCES

- Hestenes, M.R., and Stiefel, E., 1952, Methods of conjugate gradients for solving linear systems: J. Res. Natl. Bur. Stand., **49**, 409-436.
- Paige, C.C., and Saunders, M.A., 1982a, LSQR: an algorithm for sparse linear equations and sparse least squares: Assn. Comp. Mach. Trans. Mathematical Software, **8**, 43-71.
- Paige, C.C., and Saunders, M.A., 1982b, Algorithm 583, LSQR: sparse linear equations and least squares problems: Assn. Comp. Mach. Trans. Mathematical Software, **8**, 195-209.

Chapter 3

Empty bins and inverse interpolation

Let us review the big picture. In Chapter 1 we developed adjoints and in Chapter 2 we developed inverse operators. Logically, correct solutions come only through inversion. Real life, however, seems nearly the opposite. This is puzzling but intriguing.

Every time you fill your car with gasoline, it derives much more from the adjoint than from inversion. I refer to the fact that “practical seismic data processing” relates much more to the use of adjoints than of inverses. It has been widely known for about the last 15 years that medical imaging and all basic image creation methods are like this. It might seem that an easy path to fame and profit would be to introduce the notion of inversion, but it is not that easy. Both cost and result quality enter the picture.

First consider cost. For simplicity, consider a data space with N values and a model (or image) space of the same size. The computational cost of applying a dense adjoint operator increases in direct proportion to the number of elements in the matrix, in this case N^2 . To achieve the minimum discrepancy between theoretical data and observed data (inversion) theoretically requires N iterations raising the cost to N^3 .

Consider an image of size $m \times m = N$. Continuing, for simplicity, to assume a dense matrix of relations between model and data, the cost for the adjoint is m^4 whereas the cost for inversion is m^6 . We’ll consider computational costs for the year 2000, but noticing that costs go as the sixth power of the mesh size, the overall situation will not change much in the foreseeable future. Suppose you give a stiff workout to a powerful machine; you take an hour to invert a 4096×4096 matrix. The solution, a vector of 4096 components could be laid into an image of size $64 \times 64 = 2^6 \times 2^6 = 4096$. Here is what we are looking at for costs:

adjoint cost	$(m \times m)^2$	$(512 \times 512)^2$	$(2^9 2^9)^2$	2^{36}
inverse cost	$(m \times m)^3$	$(64 \times 64)^3$	$(2^6 2^6)^3$	2^{36}

These numbers tell us that for applications with dense operators, the biggest images that we are likely to see coming from inversion methods are 64×64 whereas those from adjoint methods are 512×512 . For comparison, the retina of your eye is comparable to your

computer screen at 1000×1000 . We might summarize by saying that while adjoint methods are less than perfect, inverse methods are “legally blind” :-)

<http://sepwww.stanford.edu/sep/jon/family/jos/gifmovie.html> holds a movie blinking between Figures 3.1 and 3.2.

Figure 3.1: Jos greets Andrew, “Welcome back Andrew” from the Peace Corps. At a resolution of 512×512 , this picture is about the same as the resolution as the paper it is printed on, or the same as your viewing screen, if you have scaled it to 50% of screen size. iin/. 512x512



This cost analysis is oversimplified in that most applications do not require dense operators. With sparse operators, the cost advantage of adjoints is even more pronounced since for adjoints, the cost savings of operator sparseness translate directly to real cost savings. The situation is less favorable and much more muddied for inversion. The reason that Chapter 2 covers iterative methods and neglects exact methods is that in practice iterative methods are not run to their theoretical completion but they run until we run out of patience.

Cost is a big part of the story, but the story has many other parts. Inversion, while being the only logical path to the best answer, is a path littered with pitfalls. The first pitfall is that the data is rarely able to determine a complete solution reliably. Generally there are aspects of the image that are not learnable from the data.

In this chapter we study the simplest, most transparent example of data insufficiency. Data exists at irregularly spaced positions in a plane. We set up a cartesian mesh and we discover that some of the bins contain no data points. What then?

3.1 MISSING DATA IN ONE DIMENSION

A method for restoring **missing data** is to ensure that the restored data, after specified filtering, has minimum energy. Specifying the filter chooses the interpolation philosophy. Generally the filter is a **roughening** filter. When a roughening filter goes off the end of smooth data, it typically produces a big end transient. Minimizing energy implies a choice for unknown data values at the end, to minimize the transient. We will examine five cases

Figure 3.2: Jos greets Andrew, “Welcome back Andrew” again. At a resolution of 64×64 the pixels are clearly visible. From far the pictures are the same. From near, examine their glasses. iin/. 64x64




and then make some generalizations.

A method for restoring missing data is to ensure that the restored data, after specified filtering, has **minimum energy**.

Let u denote an unknown (missing) value. The dataset on which the examples are based is $(\dots, u, u, 1, u, 2, 1, 2, u, u, \dots)$. Theoretically we could adjust the missing u values (each different) to minimize the energy in the unfiltered data. Those adjusted values would obviously turn out to be all zeros. The unfiltered data is data that has been filtered by an impulse function. To find the missing values that minimize energy out of other filters, we can use subroutine `mis1()` on page 70. Figure 3.3 shows interpolation of the dataset with $(1, -1)$ as a roughening filter. The interpolated data matches the given data where they overlap.

Figure 3.3: Top is given data. Middle is given data with interpolated values. Missing values seem to be interpolated by straight lines. Bottom shows the filter $(1, -1)$, whose output has minimum energy.

iin/miss1figs mlines



Figures 3.3–3.6 illustrate that the rougher the filter, the smoother the interpolated data, and vice versa. Let us switch our attention from the residual spectrum to the residual itself. The residual for Figure 3.3 is the *slope* of the signal (because the filter $(1, -1)$ is a

Figure 3.4: Top is the same input data as in Figure 3.3. Middle is interpolated. Bottom shows the filter $(-1, 2, -1)$. The missing data seems to be interpolated by parabolas.




Figure 3.5: Top is the same input. Middle is interpolated. Bottom shows the filter $(1, -3, 3, -1)$. The missing data is very smooth. It shoots upward high off the right end of the observations, apparently to match the data slope there.





Figure 3.6: Bottom shows the filter $(1, 1)$. The interpolation is rough. Like the given data itself, the interpolation has much energy at the Nyquist frequency. But unlike the given data, it has little zero-frequency energy.




first derivative), and the slope is constant (uniformly distributed) along the straight lines where the least-squares procedure is choosing signal values. So these examples confirm the idea that the **least-squares method** abhors large values (because they are squared). Thus, least squares tends to distribute residuals uniformly in both time and frequency to the extent allowed by the **constraints**.

This idea helps us answer the question, what is the best filter to use? It suggests choosing the filter to have an amplitude spectrum that is inverse to the spectrum we want for the interpolated data. A systematic approach is given in chapter 6, but I offer a simple subjective analysis here: Looking at the data, we see that all points are positive. It seems, therefore, that the data is rich in low frequencies; thus the filter should contain something like $(1, -1)$, which vanishes at zero frequency. Likewise, the data seems to contain Nyquist frequency, so the filter should contain $(1, 1)$. The result of using the filter $(1, -1) * (1, 1) = (1, 0, -1)$ is shown in Figure 3.7. This is my best subjective interpolation based on the idea that the missing data should look like the given data. The **interpolation** and **extrapolations** are so good that you can hardly guess which data values are given and which are interpolated.

Figure 3.7: Top is the same as in Figures 3.3 to 3.6. Middle is interpolated. Bottom shows the filter $(1, 0, -1)$, which comes from the coefficients of $(1, -1) * (1, 1)$. Both the given data and the interpolated data have significant energy at both zero and Nyquist frequencies.

`iin/miss1figs mbest`



3.1.1 Missing-data program

Now let us see how Figures 3.3-3.7 could have been calculated and how they were calculated. They could have been calculated with matrices, in which matrices were pulled apart according to subscripts of known or missing data; instead I computed them with operators, and applied only operators and their adjoints. First we inspect the matrix approach because it is more conventional.

Matrix approach to missing data

Customarily, we have referred to data by the symbol \mathbf{d} . Now that we are dividing the data space into two parts, known and unknown (or missing), we will refer to this complete space as the model (or map) space \mathbf{m} .

There are 15 data points in Figures 3.3-3.7. Of the 15, 4 are known and 11 are missing. Denote the known by k and the missing by u . Then the sequence of missing and known is $(u, u, u, u, k, u, k, k, k, u, u, u, u, u, u)$. Because I cannot print 15×15 matrices, please allow me to describe instead a data space of 6 values $(m_1, m_2, m_3, m_4, m_5, m_6)$ with known values only m_2 and m_3 , that is arranged like (u, k, k, u, u, u) .

Our approach is to minimize the energy in the residual, which is the filtered map (model) space. We state the fitting goals $\mathbf{0} \approx \mathbf{F}\mathbf{m}$ as

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \approx \mathbf{r} = \begin{bmatrix} a_1 & 0 & 0 & 0 & 0 & 0 \\ a_2 & a_1 & 0 & 0 & 0 & 0 \\ a_3 & a_2 & a_1 & 0 & 0 & 0 \\ 0 & a_3 & a_2 & a_1 & 0 & 0 \\ 0 & 0 & a_3 & a_2 & a_1 & 0 \\ 0 & 0 & 0 & a_3 & a_2 & a_1 \\ 0 & 0 & 0 & 0 & a_3 & a_2 \\ 0 & 0 & 0 & 0 & 0 & a_3 \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \\ m_6 \end{bmatrix} \quad (3.1)$$

We rearrange the above fitting goals, bringing the columns multiplying known data values (m_2 and m_3) to the left, getting $\mathbf{y} = -\mathbf{F}_k\mathbf{m}_k \approx \mathbf{F}_u\mathbf{m}_u$.

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix} = - \begin{bmatrix} 0 & 0 \\ a_1 & 0 \\ a_2 & a_1 \\ a_3 & a_2 \\ 0 & a_3 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} m_2 \\ m_3 \end{bmatrix} \approx \begin{bmatrix} a_1 & 0 & 0 & 0 \\ a_2 & 0 & 0 & 0 \\ a_3 & 0 & 0 & 0 \\ 0 & a_1 & 0 & 0 \\ 0 & a_2 & a_1 & 0 \\ 0 & a_3 & a_2 & a_1 \\ 0 & 0 & a_3 & a_2 \\ 0 & 0 & 0 & a_3 \end{bmatrix} \begin{bmatrix} m_1 \\ m_4 \\ m_5 \\ m_6 \end{bmatrix} \quad (3.2)$$

This is the familiar form of an overdetermined system of equations $\mathbf{y} \approx \mathbf{F}_u\mathbf{m}_u$ which we could solve for \mathbf{m}_u as illustrated earlier by conjugate directions, or by a wide variety of well-known methods.

The trouble with this matrix approach is that it is awkward to program the partitioning of the operator into the known and missing parts, particularly if the application of the operator uses arcane techniques, such as those used by the fast-Fourier-transform operator or various numerical approximations to differential or partial differential operators that depend on regular data sampling. Even for the modest convolution operator, we already have a library of convolution programs that handle a variety of end effects, and it would be much nicer to use the library as it is rather than recode it for all possible geometrical arrangements of missing data values.

Note: Here I take the main goal to be the clarity of the code, not the efficiency or accuracy of the solution. So, if your problem consumes too many resources, and if you have many more known points than missing ones, maybe you should fit $\mathbf{y} \approx \mathbf{F}_u\mathbf{m}_u$ and ignore the suggestions below.

Operator approach to missing data

For the operator approach to the fitting goal $-\mathbf{F}_k\mathbf{m}_k \approx \mathbf{F}_u\mathbf{m}_u$ we rewrite it as $-\mathbf{F}_k\mathbf{m}_k \approx \mathbf{FJm}$ where

$$-\mathbf{F}_k \mathbf{m}_k \approx \begin{bmatrix} a_1 & 0 & 0 & 0 & 0 & 0 \\ a_2 & a_1 & 0 & 0 & 0 & 0 \\ a_3 & a_2 & a_1 & 0 & 0 & 0 \\ 0 & a_3 & a_2 & a_1 & 0 & 0 \\ 0 & 0 & a_3 & a_2 & a_1 & 0 \\ 0 & 0 & 0 & a_3 & a_2 & a_1 \\ 0 & 0 & 0 & 0 & a_3 & a_2 \\ 0 & 0 & 0 & 0 & 0 & a_3 \end{bmatrix} \begin{bmatrix} 1 & . & . & . & . & . \\ . & 0 & . & . & . & . \\ . & . & 0 & . & . & . \\ . & . & . & 1 & . & . \\ . & . & . & . & 1 & . \\ . & . & . & . & . & 1 \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \\ m_6 \end{bmatrix} = \mathbf{FJm} \quad (3.3)$$

Notice the introduction of the new diagonal matrix \mathbf{J} , called a **masking** matrix or a **constraint-mask** matrix because it multiplies constrained variables by zero leaving freely adjustable variables untouched. Experience shows that a better name than “mask matrix” is “**selector** matrix” because what comes out of it, that which is selected, is a less-confusing name for it than which is rejected. With a selector matrix the whole data space seems freely adjustable, both the missing data values and known values. We see that the CD method does not change the known (constrained) values. In general, we derive the fitting goal (3.3) by

$$\mathbf{0} \approx \mathbf{Fm} \quad (3.4)$$

$$\mathbf{0} \approx \mathbf{F}(\mathbf{J} + (\mathbf{I} - \mathbf{J}))\mathbf{m} \quad (3.5)$$

$$\mathbf{0} \approx \mathbf{FJm} + \mathbf{F}(\mathbf{I} - \mathbf{J})\mathbf{m} \quad (3.6)$$

$$\mathbf{0} \approx \mathbf{FJm} + \mathbf{Fm}_{\text{known}} \quad (3.7)$$

$$\mathbf{0} \approx \mathbf{r} = \mathbf{FJm} + \mathbf{r}_0 \quad (3.8)$$

As usual, we find a direction to go $\Delta\mathbf{m}$ by the gradient of the residual energy.

$$\Delta\mathbf{m} = \frac{\partial}{\partial \mathbf{m}'} \mathbf{r}'\mathbf{r} = \left(\frac{\partial}{\partial \mathbf{m}'} \mathbf{r}' \right) \mathbf{r} = \left(\frac{\partial}{\partial \mathbf{m}'} (\mathbf{m}'\mathbf{J}'\mathbf{F}' + \mathbf{r}'_0) \right) \mathbf{r} = \mathbf{J}'\mathbf{F}'\mathbf{r} \quad (3.9)$$

We begin the calculation with the known data values where missing data values are replaced by zeros, namely $(\mathbf{I} - \mathbf{J})\mathbf{m}$. Filter this data, getting $\mathbf{F}(\mathbf{I} - \mathbf{J})\mathbf{m}$, and load it into the residual \mathbf{r}_0 . With this initialization completed, we begin an iteration loop. First we compute Δm from equation (3.9).

$$\Delta\mathbf{m} \leftarrow \mathbf{J}'\mathbf{F}'\mathbf{r} \quad (3.10)$$

\mathbf{F}' applies a *crosscorrelation* of the filter to the residual and then \mathbf{J}' sets to zero any changes proposed to known data values. Next, compute the change in residual $\Delta\mathbf{r}$ from the proposed change in the data $\Delta\mathbf{m}$.

$$\Delta\mathbf{r} \leftarrow \mathbf{FJ}\Delta\mathbf{m} \quad (3.11)$$

This applies the filtering again. Then use the method of steepest descent (or conjugate direction) to choose the appropriate scaling (or inclusion of previous step) of $\Delta\mathbf{m}$ and $\Delta\mathbf{r}$, and update \mathbf{m} and \mathbf{r} accordingly and iterate.

The subroutine to find missing data is `mis1()`. It assumes that zero values in the input data correspond to missing data locations. It uses our convolution operator `tcail()` on page 9. You can also check the Index for other **operators** and **modules**.

```

user/gee/misl.c
54 void misl(int niter      /* number of iterations */,
55          float *xx      /* data/model */,
56          const bool *known /* mask for known data */,
57          const char *step /* solver */)
58 /*< interpolate >*/
59 {
60     switch (step[1]) {
61         case 'g': /* conjugate gradients */
62             sf_solver (tcail_lop, sf_cgstep, nx, ny, xx, zero,
63                      niter, "x0", xx, "known", known, "end");
64             sf_cgstep_close();
65             break;
66         case 'd': /* conjugate directions */
67             sf_cdstep_init();
68             sf_solver (tcail_lop, sf_cdstep, nx, ny, xx, zero,
69                      niter, "x0", xx, "known", known, "end");
70             sf_cdstep_close();
71             break;
72         default:
73             sf_error("%s: unknown step %s", __FILE__, step);
74             break;
75     }
76 }

```

I sought reference material on conjugate gradients with constraints and didn't find anything, leaving me to fear that this chapter was in error and that I had lost the magic property of convergence in a finite number of iterations. I tested the code and it did converge in a finite number of iterations. The explanation is that these constraints are almost trivial. We pretended we had extra variables, and computed a $\Delta \mathbf{m} = \mathbf{g}$ for each of them. Then we set the $\Delta \mathbf{m} = \mathbf{g}$ to zero, hence making no changes to anything, like as if we had never calculated the extra $\Delta \mathbf{m}$'s.

EXERCISES:

- 1 Figures 3.3–3.6 seem to extrapolate to vanishing signals at the side boundaries. Why is that so, and what could be done to leave the sides unconstrained in that way?
- 2 Show that the interpolation curve in Figure 3.4 is not parabolic as it appears, but cubic. (HINT: First show that $(\nabla^2)' \nabla^2 u = \mathbf{0}$.)
- 3 Verify by a program example that the number of iterations required with simple constraints is the number of free parameters.
- 4 A signal on a uniform mesh has missing values. How should we estimate the mean?

3.2 WELLS NOT MATCHING THE SEISMIC MAP

Accurate knowledge comes from a **well**, but wells are expensive and far apart. Less accurate knowledge comes from surface seismology, but this knowledge is available densely in space and can indicate significant **trends** between the wells. For example, a prospective area may contain 15 wells but 600 or more seismic stations. To choose future well locations, it is helpful to match the known well data with the seismic data. Although the seismic data is delightfully dense in space, it often mismatches the wells because there are systematic differences in the nature of the measurements. These discrepancies are sometimes attributed to velocity **anisotropy**. To work with such measurements, we do not need to track down the physical model, we need only to merge the information somehow so we can appropriately **map** the trends between wells and make a proposal for the next drill site. Here we consider only a scalar value at each location. Take \mathbf{w} to be a vector of 15 components, each component being the seismic travel time to some fixed depth in a well. Likewise let \mathbf{s} be a 600-component vector each with the seismic travel time to that fixed depth as estimated wholly from surface seismology. Such empirical corrections are often called “**fudge factors**”. An example is the Chevron oil field in Figure 3.8. The binning of the seismic data in Figure 3.8 is not really satisfactory when we have available the techniques of missing data estimation to fill the empty bins. Using the ideas of subroutine `mis1()` on page 70, we can extend the seismic data into the empty part of the plane. We use the same principle that we minimize the energy in the filtered map where the map must match the data where it is known. I chose the filter $\mathbf{A} = \nabla' \nabla = -\nabla^2$ to be the Laplacian operator (actually, its negative) to obtain the result in Figure 3.9.

Figure 3.9 also involves a **boundary condition calculation**. Many differential equations have a solution that becomes infinite at infinite distance, and in practice this means that the largest solutions may often be found on the boundaries of the plot, exactly where




Figure 3.8: Binning by data push. Left is seismic data. Right is well locations. Values in bins are divided by numbers in bins. (Toldi) `iin/chevron wellseis`




Figure 3.9: Seismic binned (left) and extended (right) by minimizing energy in $\nabla^2 \mathbf{s}$. `iin/chevron misseis`

there is the least information. To obtain a more pleasing result, I placed artificial “average” data along the outer boundary. Each boundary point was given the value of an average of the interior data values. The average was weighted, each weight being an inverse power of the separation distance of the boundary point from the interior point.

Parenthetically, we notice that all the unknown interior points could be guessed by the same method we used on the outer boundary. After some experience guessing what inverse power would be best for the weighting functions, I do not recommend this method. Like gravity, the forces of interpolation from the weighted sums are not blocked by intervening objects. But the temperature in a house is not a function of temperature in its neighbor’s house. To further isolate the more remote points, I chose weights to be the inverse fourth power of distance.

The first job is to fill the gaps in the seismic data. We just finished doing a job like this in one dimension. I’ll give you more computational details later. Let us call the extended seismic data \mathbf{s} .

Think of a map of a model space \mathbf{m} of infinitely many hypothetical wells that must match the real wells, where we have real wells. We must find a map that matches the wells exactly and somehow matches the seismic information elsewhere. Let us define the vector \mathbf{w} as shown in Figure 3.8 so \mathbf{w} is observed values at wells and zeros elsewhere.

Where the seismic data contains sharp bumps or streaks, we want our final earth model to have those features. The wells cannot provide the rough features because the wells are too far apart to provide high spatial frequencies. The well information generally conflicts with the seismic data at low spatial frequencies because of systematic discrepancies between the two types of measurements. Thus we must accept that \mathbf{m} and \mathbf{s} may differ at low spatial frequencies (where gradient and Laplacian are small).

Our final map \mathbf{m} would be very unconvincing if it simply jumped from a well value at one point to a seismic value at a neighboring point. The map would contain discontinuities around each well. Our philosophy of finding an earth model \mathbf{m} is that our earth map should contain no obvious “footprint” of the data acquisition (well locations). We adopt the philosophy that the difference between the final map (extended wells) and the seismic information $\mathbf{x} = \mathbf{m} - \mathbf{s}$ should be smooth. Thus, we seek the minimum residual \mathbf{r} which is the roughened difference between the seismic data \mathbf{s} and the map \mathbf{m} of hypothetical omnipresent wells. With roughening operator \mathbf{A} we fit

$$\mathbf{0} \approx \mathbf{r} = \mathbf{A}(\mathbf{m} - \mathbf{s}) = \mathbf{A}\mathbf{x} \quad (3.12)$$

along with the constraint that the map should match the wells at the wells. We could write this as $\mathbf{0} = (\mathbf{I} - \mathbf{J})(\mathbf{m} - \mathbf{w})$. We honor this constraint by initializing the map $\mathbf{m} = \mathbf{w}$ to the wells (where we have wells, and zero elsewhere). After we find the gradient direction to suggest some changes to \mathbf{m} , we simply will not allow those changes at well locations. We do this with a mask. We apply a “missing data selector” to the gradient. It zeros out possible changes at well locations. Like with the goal (3.7), we have

$$\mathbf{0} \approx \mathbf{r} = \mathbf{A}\mathbf{J}\mathbf{x} + \mathbf{A}\mathbf{x}_{\text{known}} \quad (3.13)$$

After minimizing \mathbf{r} by adjusting \mathbf{x} , we have our solution $\mathbf{m} = \mathbf{x} + \mathbf{s}$.

Now we prepare some roughening operators \mathbf{A} . We have already coded a 2-D gradient operator `igrad2` on page 79. Let us combine it with its adjoint to get the 2-D laplacian

operator. (You might notice that the laplacian operator is “self-adjoint” meaning that the operator does the same calculation that its adjoint does. Any operator of the form $\mathbf{A}'\mathbf{A}$ is self-adjoint because $(\mathbf{A}'\mathbf{A})' = \mathbf{A}'\mathbf{A}'' = \mathbf{A}'\mathbf{A}$.)

Subroutine `lapfill()` on page 76 is the same idea as `mis1()` on page 70 except that the filter \mathbf{A} has been specialized to the laplacian implemented by module `laplac2` on page 75.

Subroutine `lapfill()` can be used for each of our two problems, (1) extending the seismic data to fill space, and (2) fitting the map exactly to the wells and approximately to the seismic data. When extending the seismic data, the initially non-zero components $\mathbf{s} \neq \mathbf{0}$ are fixed and cannot be changed.

The final map is shown in Figure 3.10.




Figure 3.10: Final map based on Laplacian roughening. iin/chevron finalmap

Results can be computed with various filters. I tried both ∇^2 and ∇ . There are disadvantages of each, ∇ being too cautious and ∇^2 perhaps being too aggressive. Figure 3.11 shows the difference \mathbf{x} between the extended seismic data and the extended wells. Notice that for ∇ the difference shows a localized “tent pole” disturbance about each well. For ∇^2 there could be large overshoot between wells, especially if two nearby wells have significantly different values. I don’t see that problem here.

My overall opinion is that the Laplacian does the better job in this case. I have that opinion because in viewing the extended gradient I can clearly see where the wells are. The wells are where we have acquired data. We’d like our map of the world to not show where we acquired data. Perhaps our estimated map of the world cannot help but show where we have and have not acquired data, but we’d like to minimize that aspect.

A good image of the earth hides our data **acquisition footprint**.

To understand the behavior theoretically, recall that in one dimension the filter ∇ in-


```

                                user/fomels/laplac2.c
34 void laplac2_lop (bool adj, bool add,
35                  int np, int nr, float *p, float *r)
36 /*< linear operator >*/
37 {
38     int i1, i2, j;
39
40     sf_adjnull (adj, add, np, nr, p, r);
41
42     for (i2=0; i2 < n2; i2++) {
43         for (i1=0; i1 < n1; i1++) {
44             j = i1+i2*n1;
45             if (i1 > 0) {
46                 if (adj) {
47                     p[j-1] -= r[j];
48                     p[j]   += r[j];
49                 } else {
50                     r[j] += p[j] - p[j-1];
51                 }
52             }
53             if (i1 < n1-1) {
54                 if (adj) {
55                     p[j+1] -= r[j];
56                     p[j]   += r[j];
57                 } else {
58                     r[j] += p[j] - p[j+1];
59                 }
60             }
61             if (i2 > 0) {
62                 if (adj) {
63                     p[j-n1] -= r[j];
64                     p[j]   += r[j];
65                 } else {
66                     r[j] += p[j] - p[j-n1];
67                 }
68             }
69             if (i2 < n2-1) {
70                 if (adj) {
71                     p[j+n1] -= r[j];
72                     p[j]   += r[j];
73                 } else {
74                     r[j] += p[j] - p[j+n1];
75                 }
76             }
77         }
78     }

```

```

user/fomels/lapfill.c
63 void lapfill(int niter /* number of iterations */,
64              float* mm /* model [m1*m2] */,
65              bool *known /* mask for known data [m1*m2] */)
66 /*< interpolate >*/
67 {
68     if (grad) {
69         sf_solver (igrad2_lop, sf_cgstep, n12, 2*n12, mm, zero,
70                  niter, "x0", mm, "known", known, "end");
71     } else {
72         sf_solver (laplac2_lop, sf_cgstep, n12, n12, mm, zero,
73                  niter, "x0", mm, "known", known, "end");
74     }
75     sf_cgstep_close ();
76 }

```




Figure 3.11: Difference between wells (the final map) and the extended seismic data. Left is plotted at the wells (with gray background for zero). Center is based on gradient roughening and shows tent-pole-like residuals at wells. Right is based on Laplacian roughening.

terpolates with straight lines and ∇^2 interpolates with cubics. This is because the fitting goal $\mathbf{0} \approx \nabla \mathbf{m}$, leads to $\frac{\partial}{\partial \mathbf{m}} \mathbf{m}' \nabla' \nabla \mathbf{m} = \mathbf{0}$ or $\nabla' \nabla \mathbf{m} = \mathbf{0}$, whereas the fitting goal $\mathbf{0} \approx \nabla^2 \mathbf{m}$ leads to $\nabla^4 \mathbf{m} = \mathbf{0}$ which is satisfied by cubics. In two dimensions, minimizing the output of ∇ gives us solutions of Laplace's equation with sources at the known data. It is as if ∇ stretches a rubber sheet over poles at each well, whereas ∇^2 bends a stiff plate.

Just because ∇^2 gives smoother maps than ∇ does not mean those maps are closer to reality. This is a deeper topic, addressed in Chapter 6. It is the same issue we noticed when comparing figures 3.3-3.7.

3.3 SEARCHING THE SEA OF GALILEE

Figure 3.12 shows a bottom-sounding survey of the Sea of Galilee¹ at various stages of processing. The ultimate goal is not only a good map of the depth to bottom, but images useful for the purpose of identifying **archaeological**, geological, or geophysical details of the sea bottom. The Sea of Galilee is unique because it is a *fresh*-water lake *below* sea-level. It seems to be connected to the great rift (pull-apart) valley crossing east Africa. We might delineate the Jordan River delta. We might find springs on the water bottom. We might find archaeological objects.

The raw data is 132,044 triples, (x_i, y_i, z_i) , where x_i ranges over about 12 km and where y_i ranges over about 20 km. The lines you see in Figure 3.12 are sequences of data points, i.e., the track of the survey vessel. The depths z_i are recorded to an accuracy of about 10 cm.

The first frame in Figure 3.12 shows simple binning. A coarser mesh would avoid the empty bins but lose resolution. As we refine the mesh for more detail, the number of empty bins grows as does the care needed in devising a technique for filling them. This first frame uses the simple idea from Chapter 1 of spraying all the data values to the nearest bin with `bin2()` on page 12 and dividing by the number in the bin. Bins with no data obviously need to be filled in some other way. I used a missing data program like that in the recent section on “wells not matching the seismic map.” Instead of roughening with a Laplacian, however, I used the gradient operator `igrad2` on page 79. The solver is `grad2fill()`.

The output of the roughening operator is an image, a filtered version of the depth, a filtered version of something real. Such filtering can enhance the appearance of interesting features. For example, scanning the shoreline of the roughened image (after missing data was filled), we see several ancient shorelines, now submerged.

The adjoint is the easiest image to build. The roughened map is often more informative than the map itself.

The views expose several defects of the data acquisition and of our data processing. The impulsive glitches (St. Peter's fish?) need to be removed but we must be careful not to throw out the sunken ships along with the bad data points. Even our best image shows clear evidence of the recording vessel's tracks. Strangely, some tracks are deeper than others.

¹ Data collected by Zvi **ben Avraham**, TelAviv University. Please communicate with him zvi@jupiter1.tau.ac.il for more details or if you make something publishable with his data.




Figure 3.12: Views of the bottom of the Sea of Galilee. iin/galilee locfil

user/fomels/igrad2.c

```

45     for (i2=0; i2 < n2-1; i2++) {
46         for (i1=0; i1 < n1-1; i1++) {
47             i = i1+i2*n1;
48             if (adj == true) {
49                 p[i+1] += r[i];
50                 p[i+n1] += r[i+n12];
51                 p[i] -= (r[i] + r[i+n12]);
52             } else {
53                 r[i] += (p[i+1] - p[i]);
54                 r[i+n12] += (p[i+n1] - p[i]);
55             }
56         }
57     }

```

user/fomels/grad2fill.c

```

52 void grad2fill(int niter /* number of iterations */,
53               float* mm /* estimated model */,
54               bool *known /* mask */)
55 /*< Run optimization >*/
56 {
57     sf_solver (igrad2_lob, sf_cgstep, n12, 2*n12, mm, zero, niter,
58               "x0", mm, "known", known, "end");
59     sf_cgstep_close ();
60 }

```

Perhaps the survey is assembled from work done in different seasons and the water level varied by season. Perhaps some days the vessel was more heavily loaded and the depth sounder was on a deeper keel. As for the navigation equipment, we can see that some data values are reported outside the lake!

We want the sharpest possible view of this classical site. A treasure hunt is never easy and no one guarantees we will find anything of great value but at least the exercise is a good warm-up for submarine petroleum exploration.

3.4 INVERSE LINEAR INTERPOLATION

In Chapter 1 we defined **linear interpolation** as the extraction of values from between mesh points. In a typical setup (occasionally the role of data and model are swapped), a model is given on a uniform mesh and we solve the easy problem of extracting values between the mesh points with subroutine `lint1()` on page 14. The genuine problem is the inverse problem, which we attack here. Data values are sprinkled all around, and we wish to find a function on a uniform mesh from which we can extract that data by **linear interpolation**. The adjoint operator for subroutine `lint1()` simply piles data back into its proper location in model space without regard to how many data values land in each region. Thus some model values may have many data points added to them while other model values get none. We could interpolate by minimizing the energy in the model gradient, or that in the second derivative of the model, or that in the output of any other roughening filter applied to the model.

Formalizing now our wish that data \mathbf{d} be extractable by **linear interpolation** \mathbf{F} , from a model \mathbf{m} , and our wish that application of a roughening filter with an operator \mathbf{A} have minimum energy, we write the fitting goals:

$$\begin{aligned} \mathbf{0} &\approx \mathbf{Fm} - \mathbf{d} \\ \mathbf{0} &\approx \mathbf{Am} \end{aligned} \tag{3.14}$$

Suppose we take the roughening filter to be the second difference operator $(1, -2, 1)$ scaled by a constant ϵ , and suppose we have a data point near each end of the model and a third data point exactly in the middle. Then, for a model space 6 points long, the fitting goal could look like

$$\left[\begin{array}{cccccc} .8 & .2 & . & . & . & . \\ . & . & 1 & . & . & . \\ . & . & . & . & .5 & .5 \\ \hline \epsilon & . & . & . & . & . \\ -2\epsilon & \epsilon & . & . & . & . \\ \epsilon & -2\epsilon & \epsilon & . & . & . \\ . & \epsilon & -2\epsilon & \epsilon & . & . \\ . & . & \epsilon & -2\epsilon & \epsilon & . \\ . & . & . & \epsilon & -2\epsilon & \epsilon \\ . & . & . & . & \epsilon & -2\epsilon \\ . & . & . & . & . & \epsilon \end{array} \right] \begin{bmatrix} m_0 \\ m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \end{bmatrix} - \begin{bmatrix} d_0 \\ d_1 \\ d_2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \mathbf{r}_d \\ \mathbf{r}_m \end{bmatrix} \approx \mathbf{0} \tag{3.15}$$

The residual vector has two parts, a data part \mathbf{r}_d on top and a model part \mathbf{r}_m on the

bottom. The data residual should vanish except where contradictory data values happen to lie in the same place. The model residual is the roughened model.

After all the definitions, we load the negative of the data into the residual. If a starting model \mathbf{m}_0 is present, then we update the data part of the residual $\mathbf{r}_d = \mathbf{F}\mathbf{m}_0 - \mathbf{d}$ and we load the model part of the residual $\mathbf{r}_m = \mathbf{A}\mathbf{m}_0$. Otherwise we begin from a zero model $\mathbf{m}_0 = \mathbf{0}$ and thus the model part of the residual \mathbf{r}_m is also zero. After this initialization, subroutine `solver_reg()` begins an iteration loop by first computing the proposed model perturbation $\Delta\mathbf{m}$ (called `g` in the program) with the adjoint operator:

$$\Delta\mathbf{m} \leftarrow \begin{bmatrix} \mathbf{F}' & \mathbf{A}' \end{bmatrix} \begin{bmatrix} \mathbf{r}_d \\ \mathbf{r}_m \end{bmatrix} \quad (3.16)$$

Using this value of $\Delta\mathbf{m}$, we can find the implied change in residual $\Delta\mathbf{r}$ as

$$\Delta \begin{bmatrix} \mathbf{r}_d \\ \mathbf{r}_m \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{F} \\ \mathbf{A} \end{bmatrix} \Delta\mathbf{m} \quad (3.17)$$

and the last thing in the loop is to use the optimization step function `stepper()` to choose the length of the step size and to choose how much of the previous step to include.

An example of using the new solver is subroutine `invint1`. I chose to implement the model roughening operator \mathbf{A} with the convolution subroutine `tcai1()` on page 9, which has transient end effects (and an output length equal to the input length plus the filter length). The adjoint of subroutine `tcai1()` suggests perturbations in the convolution input (not the filter).

```
user/gee/invint1.c
```


```

24 void invint1(int niter                /* number of iterations */,
25              int nd                  /* data size */,
26              float *coord            /* data coordinates */,
27              const float *dd         /* data values */,
28              int n1, float ol, float dl /* model grid */,
29              int na, const float *aa  /* filter */,
30              float *mm               /* estimated model */,
31              float eps               /* regularization */)
32 /*< inverse interpolation >*/
33 {
34     lint1_init( ol, dl, coord); /* interpolation */
35     tcail_init( na, aa);        /* filtering */
36
37     sf_solver_reg(lint1_lop, sf_cgstep, tcail_lop,
38                  n1+na, n1, nd, mm, dd, niter, eps);
39     sf_cgstep_close( );
40 }

```

Figure 3.13 shows an example for a $(1, -2, 1)$ filter with $\epsilon = 1$. The continuous curve representing the model \mathbf{m} passes through the data points. Because the models are computed


Figure 3.13: Sample points and estimation of a continuous function through them. `iin/invint im1-2+1`



with transient convolution end-effects, the models tend to damp linearly to zero outside the region where signal samples are given.

To show an example where the result is clearly a theoretical answer, I prepared another figure with the simpler filter $(1, -1)$. When we minimize energy in the first derivative of the waveform, the residual distributes itself uniformly between data points so the solution there is a straight line. Theoretically it should be a straight line because a straight line has a vanishing second derivative, and that condition arises by differentiating by \mathbf{x}' , the minimized quadratic form $\mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x}$, and getting $\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{0}$. (By this logic, the curves between data points in Figure 3.13 must be cubics.) The $(1, -1)$ result is shown in Figure 3.14.

Figure 3.14: The same data samples and a function through them that minimizes the energy in the first derivative. `iin/invint im1-1a`



The example of Figure 3.14 has been a useful test case for me. You'll see it again in later chapters. What I would like to show you here is a movie showing the convergence to Figure 3.14. Convergence occurs rapidly where data points are close together. The large gaps, however, fill at a rate of one point per iteration.

3.4.1 Abandoned theory for matching wells and seismograms

Let us consider theory to construct a map \mathbf{m} that fits dense seismic data \mathbf{s} and the well data \mathbf{w} . The first goal $\mathbf{0} \approx \mathbf{L}\mathbf{m} - \mathbf{w}$ says that when we linearly interpolate from the map, we should get the well data. The second goal $\mathbf{0} \approx \mathbf{A}(\mathbf{m} - \mathbf{s})$ (where \mathbf{A} is a roughening operator like ∇ or ∇^2) says that the map \mathbf{m} should match the seismic data \mathbf{s} at high frequencies but need not do so at low frequencies.

$$\begin{aligned} \mathbf{0} &\approx \mathbf{L}\mathbf{m} - \mathbf{w} \\ \mathbf{0} &\approx \mathbf{A}(\mathbf{m} - \mathbf{s}) \end{aligned} \tag{3.18}$$

Although (3.18) is the way I originally formulated the well-fitting problem, I abandoned

it for several reasons: First, the map had ample pixel resolution compared to other sources of error, so I switched from linear interpolation to binning. Once I was using binning, I had available the simpler empty-bin approaches. These have the further advantage that it is not necessary to experiment with the relative weighting between the two goals in (3.18). A formulation like (3.18) is more likely to be helpful where we need to handle rapidly changing functions where binning is inferior to linear interpolation, perhaps in reflection seismology where high resolution is meaningful.

EXERCISES:

- 1 It is desired to find a compromise between the Laplacian roughener and the gradient roughener. What is the size of the residual space?
- 2 Like the seismic prospecting industry, you have solved a huge problem using binning. You have computer power left over to do a few iterations with linear interpolation. How much does the cost per iteration increase? Should you refine your model mesh, or can you use the same model mesh that you used when binning?

3.5 PREJUDICE, BULLHEADEDNESS, AND CROSS VALIDATION

First we first look at data \mathbf{d} . Then we think about a model \mathbf{m} , and an operator \mathbf{L} to link the model and the data. Sometimes the operator is merely the first term in a series expansion about $(\mathbf{m}_0, \mathbf{d}_0)$. Then we fit $\mathbf{d} - \mathbf{d}_0 \approx \mathbf{L}(\mathbf{m} - \mathbf{m}_0)$. To fit the model, we must reduce the fitting residuals. Realizing that the importance of a data residual is not always simply the size of the residual but is generally a function of it, we conjure up (topic for later chapters) a weighting function (which could be a filter) operator \mathbf{W} . This defines our data residual:

$$\mathbf{r}_d = \mathbf{W}[\mathbf{L}(\mathbf{m} - \mathbf{m}_0) - (\mathbf{d} - \mathbf{d}_0)] \quad (3.19)$$

Next we realize that the data might not be adequate to determine the model, perhaps because our comfortable dense sampling of the model ill fits our economical sparse sampling of data. Thus we adopt a fitting goal that mathematicians call “regularization” and we might call a “model style” goal or more simply, a quantification of our prejudice about models. We express this by choosing an operator \mathbf{A} , often simply a roughener like a gradient (the choice again a topic in this and later chapters). It defines our model residual by $\mathbf{A}\mathbf{m}$ or $\mathbf{A}(\mathbf{m} - \mathbf{m}_0)$, say we choose

$$\mathbf{r}_m = \mathbf{A}\mathbf{m} \quad (3.20)$$

In an ideal world, our model prejudice would not conflict with measured data, however, life is not so simple. Since conflicts between data and preconceived notions invariably arise (and they are why we go to the expense of acquiring data) we need an adjustable parameter that measures our “bullheadedness”, how much we intend to stick to our preconceived notions in spite of contradicting data. This parameter is generally called epsilon ϵ because we like to imagine that our bullheadedness is small. (In mathematics, ϵ is often taken to be an infinitesimally small quantity.) Although any bullheadedness seems like a bad thing,

it must be admitted that measurements are imperfect too. Thus as a practical matter we often find ourselves minimizing

$$\min := \mathbf{r}_d \cdot \mathbf{r}_d + \epsilon^2 \mathbf{r}_m \cdot \mathbf{r}_m \quad (3.21)$$

and wondering what to choose for ϵ . I have two suggestions: My simplest suggestion is to choose ϵ so that the residual of data fitting matches that of model styling. Thus

$$\epsilon = \sqrt{\frac{\mathbf{r}_d \cdot \mathbf{r}_d}{\mathbf{r}_m \cdot \mathbf{r}_m}} \quad (3.22)$$

My second suggestion is to think of the force on our final solution. In physics, force is associated with a gradient. We have a gradient for the data fitting and another for the model styling:

$$\mathbf{g}_d = \mathbf{L}'\mathbf{W}'\mathbf{r}_d \quad (3.23)$$

$$\mathbf{g}_m = \mathbf{A}'\mathbf{r}_m \quad (3.24)$$

We could balance these forces by the choice

$$\epsilon = \sqrt{\frac{\mathbf{g}_d \cdot \mathbf{g}_d}{\mathbf{g}_m \cdot \mathbf{g}_m}} \quad (3.25)$$

Although we often ignore ϵ in discussing the formulation of a problem, when time comes to solve the problem, reality intercedes. Generally, \mathbf{r}_d has different physical units than \mathbf{r}_m (likewise \mathbf{g}_d and \mathbf{g}_m) and we cannot allow our solution to depend on the accidental choice of units in which we express the problem. I have had much experience choosing ϵ , but it is only recently that I boiled it down to the above two suggestions. Normally I also try other values, like double or half those of the above choices, and I examine the solutions for subjective appearance. If you find any insightful examples, please tell me about them.

Computationally, we could choose a new ϵ with each iteration, but it is more expeditious to freeze ϵ , solve the problem, recompute ϵ , and solve the problem again. I have never seen a case where more than one iteration was necessary.

People who work with small problems (less than about 10^3 vector components) have access to an attractive theoretical approach called cross-validation. Simply speaking, we could solve the problem many times, each time omitting a different data value. Each solution would provide a model that could be used to predict the omitted data value. The quality of these predictions is a function of ϵ and this provides a guide to finding it. My objections to cross validation are two-fold: First, I don't know how to apply it in the large problems like we solve in this book (I should think more about it); and second, people who worry much about ϵ , perhaps first should think more carefully about their choice of the filters \mathbf{W} and \mathbf{A} , which is the focus of this book. Notice that both \mathbf{W} and \mathbf{A} can be defined with a scaling factor which is like scaling ϵ . Often more important in practice, with \mathbf{W} and \mathbf{A} we have a scaling factor that need not be constant but can be a function of space or spatial frequency within the data space and/or model space.

Chapter 4

The helical coordinate

For many years it has been true that our most powerful signal-analysis techniques are in *one*-dimensional space, while our most important applications are in *multi*-dimensional space. The helical coordinate system makes a giant step towards overcoming this difficulty.

Many geophysical map estimation problems appear to be multidimensional, but actually they are not. To see the tip of the iceberg, consider this example: On a two-dimensional

cartesian mesh, the function

0	0	0	0
0	1	1	0
0	1	1	0
0	0	0	0

has the autocorrelation

1	2	1
2	4	2
1	2	1

Likewise, on a one-dimensional cartesian mesh,

the function $\mathcal{B} =$

1	1	0	0	...	0	1	1
---	---	---	---	-----	---	---	---

has the autocorrelation $\mathcal{R} =$

1	2	1	0	...	0	2	4	2	0	...	1	2	1
---	---	---	---	-----	---	---	---	---	---	-----	---	---	---

.

Observe the numbers in the one-dimensional world are identical with the numbers in the two-dimensional world. This correspondence is no accident.

4.1 FILTERING ON A HELIX

Figure 1 shows some two-dimensional shapes that are convolved together. The left panel shows an impulse response function, the center shows some impulses, and the right shows the superposition of responses.

A surprising, indeed amazing, fact is that Figure 1 was not computed with a two-dimensional convolution program. It was computed with a one-dimensional computer program. It could have been done with anybody's one-dimensional convolution program, either in the time domain or in the fourier domain. This magical trick is done with the helical coordinate system.




Figure 4.1: Two-dimensional convolution as performed in one dimension by module `helicon` `hlx/helicon diamond`

A basic idea of filtering, be it in one dimension, two dimensions, or more, is that you have some filter coefficients and some sampled data; you pass the filter over the data; at each location you find an output by crossmultiplying the filter coefficients times the underlying data and summing the terms.

The helical coordinate system is much simpler than you might imagine. Ordinarily, a plane of data is thought of as a collection of columns, side by side. Instead, imagine the columns stored end-to-end, and then coiled around a cylinder. This is the helix. Fortran programmers will recognize that fortran’s way of storing 2-D arrays in one-dimensional memory is exactly what we need for this helical mapping. Seismologists sometimes use the word “supertrace” to describe a collection of seismograms stored “end-to-end”.

Figure 2 shows a helical mesh for 2-D data on a cylinder. Darkened squares depict a 2-D filter shaped like the Laplacian operator $\partial_{xx} + \partial_{yy}$. The input data, the filter, and the output data are all on helical meshes all of which could be unrolled into linear strips. A compact 2-D filter like a Laplacian, on a helix is a sparse 1-D filter with long empty gaps.

Since the values output from filtering can be computed in any order, we can slide the filter coil over the data coil in any direction. The order that you produce the outputs is irrelevant. You could compute the results in parallel. We could, however, slide the filter over the data in the screwing order that a nut passes over a bolt. The screw order is the same order that would be used if we were to unwind the coils into one-dimensional strips and convolve them across one another. The same filter coefficients overlay the same data values if the 2-D coils are unwound into 1-D strips. The helix idea allows us to obtain the same convolution output in either of two ways, a one-dimensional way, or a two-dimensional way. I used the one-dimensional way to compute the obviously two-dimensional result in Figure 1.

4.1.1 Review of 1-D recursive filters

Convolution is the operation we do on polynomial coefficients when we multiply polynomials. Deconvolution is likewise for polynomial division. Often these ideas are described as polynomials in the variable Z . Take $X(Z)$ to denote the polynomial whose coefficients

Figure 4.2: Filtering on a helix. The same filter coefficients overlay the same data values if the 2-D coils are unwound into 1-D strips. (*Mathematica* drawing by Sergey Fomel)

hlx/Math sergey-helix

are samples of input data, and let $A(Z)$ likewise denote the filter. The convention I adopt here is that the first coefficient of the filter has the value $+1$, so the filter's polynomial is $A(Z) = 1 + a_1Z + a_2Z^2 + \dots$. To see how to convolve, we now identify the coefficient of Z^k in the product $Y(Z) = A(Z)X(Z)$. The usual case (k larger than the number N_a of filter coefficients) is

$$y_k = x_k + \sum_{i=1}^{N_a} a_i x_{k-i} \quad (4.1)$$

Convolution computes y_k from x_k whereas deconvolution (also called back substitution) does the reverse. Rearranging (1) we get

$$x_k = y_k - \sum_{i=1}^{N_a} a_i x_{k-i} \quad (4.2)$$

where now we are finding the output x_k from its past outputs x_{k-i} and from the present input y_k . We see that the deconvolution process is essentially the same as the convolution process, except that the filter coefficients are used with opposite polarity; and they are applied to the past *outputs* instead of the past *inputs*. That is why deconvolution must be done sequentially while convolution can be done in parallel.

4.1.2 Multidimensional deconvolution breakthrough

Deconvolution (polynomial division) can undo convolution (polynomial multiplication). A magical property of the helix is that we can consider 1-D convolution to be the same as 2-D

convolution. Hence is a second magical property: We can use 1-D *deconvolution* to undo convolution, whether that convolution was 1-D or 2-D. Thus, we have discovered how to undo 2-D convolution. We have discovered that 2-D deconvolution on a helix is equivalent to 1-D deconvolution. The helix enables us to do multidimensional deconvolution.

Deconvolution is recursive filtering. Recursive filter outputs cannot be computed in parallel, but must be computed sequentially as in one dimension, namely, in the order that the nut screws on the bolt.

Recursive filtering sometimes solves big problems with astonishing speed. It can propagate energy rapidly for long distances. Unfortunately, recursive filtering can also be unstable. The most interesting case, near resonance, is also near instability. There is a large literature and extensive technology about recursive filtering in one dimension. The helix allows us to apply that technology to two (and more) dimensions. It is a huge technological breakthrough.

In 3-D we simply append one plane after another (like a 3-D fortran array). It is easier to code than to explain or visualize a spool or torus wrapped with string, etc.

4.1.3 Examples of simple 2-D recursive filters

Let us associate x - and y -derivatives with a finite-difference stencil or template. (For simplicity take $\Delta x = \Delta y = 1$.)

$$\frac{\partial}{\partial x} = \begin{array}{|c|c|} \hline 1 & -1 \\ \hline \end{array} \quad (4.3)$$

$$\frac{\partial}{\partial y} = \begin{array}{|c|} \hline 1 \\ \hline -1 \\ \hline \end{array} \quad (4.4)$$

Convolving a data plane with the stencil (3) forms the x -derivative of the plane. Convolving a data plane with the stencil (4) forms the y -derivative of the plane. On the other hand, *deconvolving* with (3) integrates data along the x -axis for each y . Likewise, deconvolving with (4) integrates data along the y -axis for each x . Next we look at a fully two-dimensional operator (like the cross derivative ∂_{xy}).

A nontrivial two-dimensional convolution stencil is

$$\begin{array}{|c|c|} \hline 0 & -1/4 \\ \hline 1 & -1/4 \\ \hline -1/4 & -1/4 \\ \hline \end{array} \quad (4.5)$$

We will convolve and deconvolve a data plane with this operator. Although everything is shown on a plane, the actual computations are done in one dimension with equations (1) and (2). Let us manufacture the simple data plane shown on the left in Figure 3. Beginning with a zero-valued plane, we add in a copy of the filter (5) near the top of the frame. Nearby add another copy with opposite polarity. Finally add some impulses near the bottom boundary. The second frame in Figure 3 is the result of deconvolution by the filter (5) using the one-dimensional equation (2). Notice that deconvolution turns the filter itself into an impulse, while it turns the impulses into comet-like images. The use of a helix is evident by the comet images wrapping around the vertical axis.




Figure 4.3: Illustration of 2-D deconvolution. Left is the input. Right is after deconvolution with the filter (5) as preformed by by module `polydiv` `hlx/helicon wrap-four`

The filtering in Figure 3 ran along a helix from left to right. Figure 4 shows a second filtering running from right to left. Filtering in the reverse direction is the adjoint. After deconvolving both ways, we have accomplished a symmetrical smoothing. The final frame undoes the smoothing to bring us exactly back to where we started. The smoothing was done with two passes of *deconvolution* and it is undone by two passes of *convolution*. No errors, no evidence remains of any of the boundaries where we have wrapped and truncated.




Figure 4.4: Recursive filtering backwards (leftward on the space axis) is done by the *adjoint* of 2-D deconvolution. Here we see that 2-D *deconvolution* compounded with its adjoint is exactly inverted by 2-D *convolution* and its adjoint. `hlx/helicon back-four`

Chapter 5 explains the important practical role to be played by a multidimensional operator for which we know the exact inverse. Other than multidimensional Fourier transformation, transforms based on polynomial multiplication and division on a helix are the only known easily invertible linear operators.

In seismology we often have occasion to steer summation along beams. Such an impulse response is shown in Figure 6.

Of special interest are filters that destroy plane waves. The inverse of such a filter creates plane waves. Such filters are like wave equations. A filter that creates two plane waves is




Figure 4.5: A simple low-order 2-D filter whose inverse contains plane waves of two different dips. One of them is spatially aliased. `hlx/helicon wrap-waves`

illustrated in figure 5.




Figure 4.6: A simple low-order 2-D filter whose inverse times its inverse adjoint, is approximately a dipping seismic arrival. `hlx/helicon dip`

4.1.4 Coding multidimensional de/convolution

Let us unroll the filter helix seen in Figure 2 and see what we have. Start from the idea that a 2-D filter is generally made from a cluster of values near one another in two dimensions similar to the Laplacian operator in the figure. We see that in the helical approach, a 2-D filter is a 1-D filter containing some long intervals of zeros. The intervals are about the length of a 1-D seismogram.

Our program for 2-D convolution with a 1-D convolution program, could convolve with the somewhat long 1-D strip, but it is much more cost effective to ignore the many zeros, which is what we do. We do not multiply by the backside zeros, nor do we even store them in memory. Whereas an ordinary convolution program would do time shifting by a code line like `iy=ix+lag`, Module `helicon` ignores the many zero filter values on backside of the tube by using the code `iy=ix+lag[ia]` where a counter `ia` ranges over the nonzero filter coefficients. Before operator `helicon` is invoked, we need to prepare two lists, one list containing nonzero filter coefficients `flt[ia]`, and the other list containing the corresponding lags `lag[ia]` measured to include multiple wraps around the helix. For example, the 2-D

Laplace operator can be thought of as the 1-D filter

$$\begin{array}{|c|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 0 & \cdots & 0 & 1 & -4 & 1 & 0 & \cdots & 0 & 1 \\ \hline \end{array} \xrightarrow{\text{helical boundaries}} \begin{array}{|c|c|c|} \hline & 1 & \\ \hline 1 & -4 & 1 \\ \hline & 1 & \\ \hline \end{array} \quad (4.6)$$

The first filter coefficient in equation (6) is +1 as implicit to module `helicon`. To apply the Laplacian on a 1000×1000 mesh requires the filter inputs:

i	lag[i]	flt[i]
---	-----	-----
0	999	1
1	1000	-4
2	1001	1
3	2000	1

Here we choose to use “declaration of a type”, a modern computer language feature that is absent from Fortran 77. Fortran 77 has the built in complex arithmetic type. In module `helix` we define a type `filter`, actually, a helix filter. After making this definition, it will be used by many programs. The helix filter consists of three vectors, a real valued vector of filter coefficients, an integer valued vector of filter lags, and an optional vector that has logical values “`true`” for output locations that will not be computed (either because of boundary conditions or because of missing inputs). The filter vectors are the size of the nonzero filter coefficients (excluding the leading 1.) while the logical vector is long and relates to the data size. The `helix` module allocates and frees memory for a helix filter. By default, the logical vector is not allocated but is set to `null` with the `nullify` operator and ignored.

filt/lib/helix.c

```

29
30 typedef struct sf_helixfilter {
31     int      nh;
32     float*   flt;
33     int*     lag;
34     bool*    mis;

```

For those of you with no C experience, the “->” appearing in the helix module denotes a pointer. Fortran 77 has no pointers (or everything is a pointer). The behavior of pointers is somewhat different in each language. Never-the-less, the idea is simple. In module `helicon` you see the expression `aa->flt[ia]`. It refers to the filter named `aa`. Any filter defined by the `helix` module contains three vectors, one of which is named `flt`. The second component of the `flt` vector in the `aa` filter is referred to as `aa->flt[1]` which in the example above refers to the value 4.0 in the center of the laplacian operator. For data sets like above with 1000 points on the 1-axis, this value 4.0 occurs after 1000 lags, thus `aa->lag[1]=1000`.

Our first convolution operator `tcai1` on page 9 was limited to one dimension and a particular choice of end conditions. With the helix and C pointers, the operator `helicon` is a *multidimensional* filter with considerable flexibility (because of the `mis` vector) to work

```

                                filt/lib/helicon.c
35 void sf_helicon_lop( bool adj, bool add,
36                     int nx, int ny, float* xx, float*yy)
37 /*< linear operator >*/
38 {
39     int ia, iy, ix;
40
41     sf_copy_lop(adj, add, nx, nx, xx, yy);
42
43     for (ia = 0; ia < aa->nh; ia++) {
44         for (iy = aa->lag[ia]; iy < nx; iy++) {
45             if( aa->mis != NULL && aa->mis[iy]) continue;
46             ix = iy - aa->lag[ia];
47             if(adj) {
48                 xx[ix] += yy[iy] * aa->flt[ia];
49             } else {
50                 yy[iy] += xx[ix] * aa->flt[ia];
51             }
52         }
53     }
54 }

```

around boundaries and missing data. The code fragment `aa->lag[ia]` corresponds to `b-1` in `tcai1` on page 9.

Operator `helicon` did the convolution job for Figure 1. As with `tcai1` on page 9 the adjoint of filtering is filtering backwards which means unscrewing the helix.

The companion to convolution is deconvolution. The module `polydiv` on page 9 is essentially the same as `polydiv1` on the current page, but here it was coded using our new `filter` type in module `helix` on this page which will simplify our many future uses of convolution and deconvolution. Although convolution allows us to work around missing input values, deconvolution does not (any input affects all subsequent outputs), so `polydiv` never references `aa->mis[ia]`.

EXERCISES:

- 1 Observe the matrix (1.4) which corresponds to subroutine `tcai1` on page 9. What is the matrix corresponding to `helicon`?

4.1.5 Causality in two-dimensions

In one dimension, most filters of interest have a short memory. Significant filter coefficients are concentrated shortly after $t = 0$. The favorite example in Physics is the damped harmonic oscillator, all of which is packed into a two-lag filter (second order differential

user/gee/polydiv.c

```

36 void polydiv_lop( bool adj, bool add,
37                  int nx, int ny, float* xx, float*yy)
38 /*< linear operator >*/
39 {
40     int ia, iy, ix;
41
42     sf_adjnull( adj, add, nx, ny, xx, yy);
43
44     for (ix=0; ix < nx; ix++) tt[ix] = 0.;
45
46     if (adj) {
47         for (ix = nx-1; ix >= 0; ix--) {
48             tt[ix] = yy[ix];
49             for (ia = 0; ia < aa->nh; ia++) {
50                 iy = ix + aa->lag[ia];
51                 if( iy >= ny) continue;
52                 tt[ix] -= aa->flt[ia] * tt[iy];
53             }
54         }
55         for (ix=0; ix < nx; ix++) xx[ix] += tt[ix];
56     } else {
57         for (iy = 0; iy < ny; iy++) {
58             tt[iy] = xx[iy];
59             for (ia = 0; ia < aa->nh; ia++) {
60                 ix = iy - aa->lag[ia];
61                 if( ix < 0) continue;
62                 tt[iy] -= aa->flt[ia] * tt[ix];
63             }
64         }
65         for (iy=0; iy < ny; iy++) yy[iy] += tt[iy];
66     }
67 }

```

equation). The complete story is rich in mathematics and in concepts, but to sum up, filters fall into two categories according to the numerical values of their coefficients. There are filters for which equations (1) and (2) work as desired and expected. These filters are called “minimum phase”. There are also filters for which (2) is a disaster numerically, the feedback process diverging to infinity.

Divergent cases correspond to physical processes that require boundary conditions. Equation (2) only allows for initial conditions. I oversimplify by trying to collapse an entire book (FGDP) into a few sentences by saying here that for any fixed spectrum there exist many filters. Of these, only one has stable polynomial division. That filter has its energy compacted as soon as possible after the “1.0” at zero lag.

Now let us turn to two dimensions. Filters of interest will correspond to energy concentrated near the end of a helix. Let us examine the end of a helix. At the very end, as in the 1-D case, is a coefficient with the numerical value 1.0. Keeping only coefficients within two mesh points in any direction from the 1.0, we copy the coefficients from near the end of the helix to a cartesian mesh like this:

$$\begin{array}{ccccc}
 h & c & 0 & & h & c & \cdot & & \cdot & \cdot & 0 \\
 p & d & 0 & & p & d & \cdot & & \cdot & \cdot & 0 \\
 q & e & \mathbf{1} & = & q & e & \cdot & + & \cdot & \cdot & \mathbf{1} \\
 s & f & a & & s & f & a & & \cdot & \cdot & \cdot \\
 u & g & b & & u & g & b & & \cdot & \cdot & \cdot
 \end{array} \tag{4.7}$$

$$2 - \text{D filter} = \text{variable} + \text{constrained}$$

where a, b, c, \dots, u are adjustable coefficients.

Which side of the little rectangular patch of coefficients we choose to place the 1.0 is rather arbitrary. The important matter is that as a matter of principle, the 1.0 is expected to lie along one side of the little patch. It is rarely (if ever) found at a corner of the patch. It is important that beyond the 1.0 (in whatever direction that may be) the filter coefficients must be zero because in one dimension, these coefficients lie before zero lag. Our foundations, the basic convolution-deconvolution pair (1) and (2) are applicable only to filters with all coefficients *after* zero lag.

Time-series analysis is rich with concepts that the helix now allows us to apply to many dimensions. First is the notion of an impulse function. Observe that an impulse function on the 2-D surface of the helical cylinder maps to an impulse function on the 1-D line of the unwound coil. An autocorrelation function that is an impulse corresponds both to a white (constant) spectrum in 1-D and to a white (constant) spectrum in 2-D. Next we look at a particularly important autocorrelation function and see how 2-D is the same as 1-D.

4.2 FINITE DIFFERENCES ON A HELIX

The function

$$\mathcal{R} = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|} \hline -1 & 0 & \cdots & 0 & -1 & 4 & -1 & 0 & \cdots & 0 & -1 \\ \hline \end{array} \tag{4.8}$$

is an autocorrelation function. It is symmetrical about the “4” and its Fourier transform is positive for all frequencies. Digging out our old textbooks¹ we discover how to compute a causal wavelet with this autocorrelation. I used the “Kolmogoroff spectral-factorization method” to find this wavelet \mathcal{H} :

$$\mathcal{H} = \begin{bmatrix} 1.791 & -.651 & -.044 & -.024 & \cdots & \cdots & -.044 & -.087 & -.200 & -.558 \end{bmatrix} \quad (4.9)$$

According to the Kolmogoroff theory, if we form the autocorrelation of \mathcal{H} , we will get \mathcal{R} . This is not obvious from the numbers themselves because the computation requires a little work.

Let the time reversed version of \mathcal{H} be denoted \mathcal{H}' . This notation is consistent with an idea from Chapter 1 that the adjoint of a filter matrix is another filter matrix with a reversed filter. In engineering it is conventional to use the asterisk symbol “*” to denote convolution. Thus, the idea that the autocorrelation of a signal \mathcal{H} is a convolution of the signal \mathcal{H} with its time reverse (adjoint) can be written as $\mathcal{H}' * \mathcal{H} = \mathcal{H} * \mathcal{H}' = \mathcal{R}$.

Wind the signal \mathcal{R} around a vertical-axis helix to see its two-dimensional shape \mathcal{R}_2 :

$$\mathcal{R} \xrightarrow{\text{helical boundaries}} \begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix} = \mathcal{R}_2 \quad (4.10)$$

This 2-D filter is the negative of the finite-difference representation of the Laplacian operator, generally denoted $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$. Now for the magic: Wind the signal \mathcal{H} around the same helix to see its two-dimensional shape \mathcal{H}_2

$$\mathcal{H}_2 = \begin{bmatrix} & & & & 1.791 & -.651 & -.044 & -.024 & \cdots \\ \cdots & -.044 & -.087 & -.200 & -.558 & & & & \end{bmatrix} \quad (4.11)$$

In the representation (11) we see the coefficients diminishing rapidly away from maximum value 1.791. My claim is that the two-dimensional autocorrelation of (11) is (10). You verified this idea earlier when the numbers were all ones. You can check it again in a few moments if you drop the small values, say 0.2 and smaller.

Since the autocorrelation of \mathcal{H} is $\mathcal{H}' * \mathcal{H} = \mathcal{R} = -\nabla^2$ is a second derivative, the operator \mathcal{H} must be something like a first derivative. As a geophysicist, I found it natural to compare the operator $\frac{\partial}{\partial y}$ with \mathcal{H} by applying them to a local topographic map. The result shown in Figure 7 is that \mathcal{H} enhances drainage patterns whereas $\frac{\partial}{\partial y}$ enhances mountain ridges.

The operator \mathcal{H} has curious similarities and differences with the familiar gradient and divergence operators. In two-dimensional physical space, the gradient maps one field to *two* fields (north slope and east slope). The factorization of $-\nabla^2$ with the helix gives us the operator \mathcal{H} that maps one field to *one* field. Being a one-to-one transformation (unlike gradient and divergence) the operator \mathcal{H} is potentially invertible by deconvolution (recursive filtering).

I have chosen the name² “helix derivative” or “helical derivative” for the operator \mathcal{H} . A telephone pole has a narrow shadow behind it. The helix integral (middle frame of Figure 8)

¹ PVI or FGDP, for example, explain spectral factorization. More concisely in PVI, more richly in FGDP.

² Any fact this basic should be named in some earlier field of mathematics or theoretical physics. Admittedly, the concept exists on an infinite cartesian plane without a helix, but all my codes in a finite space involve the helix, and the helix concept led me to it.




Figure 4.7: Topography, helical derivative, slope south. $hlx/helocut$ $helocut$

and the helix derivative (left frame) show shadows with an angular bandwidth approaching 180° .

Our construction makes \mathcal{H} have the energy spectrum $k_x^2 + k_y^2$, so the magnitude of the Fourier transform is $\sqrt{k_x^2 + k_y^2}$. It is a cone centered and with value zero at the origin. By contrast, the components of the ordinary gradient have amplitude responses $|k_x|$ and $|k_y|$ that are lines of zero across the (k_x, k_y) -plane.

The rotationally invariant cone in the Fourier domain contrasts sharply with the nonrotationally invariant function shape in (x, y) -space. The difference must arise from the phase spectrum. The factorization (11) is nonunique in that causality associated with the helix mapping can be defined along either x - or y -axes; thus the operator (11) can be rotated or reflected.

This is where the story all comes together. One-dimensional theory, either the old Kolmogoroff spectral factorization, or the new Wilson-Burg spectral-factorization method produces not merely a causal wavelet with the required autocorrelation. It produces one that is stable in deconvolution. Using \mathcal{H} in one-dimensional polynomial division, we can solve many formerly difficult problems very rapidly. Consider the Laplace equation with sources (Poisson's equation). Polynomial division and its reverse (adjoint) gives us $\mathcal{P} = (\mathcal{Q}/\mathcal{H})/\mathcal{H}'$ which means that we have solved $\nabla^2 \mathcal{P} = -\mathcal{Q}$ by using polynomial division on a helix. Using the seven coefficients shown, the cost is fourteen multiplications (because we need to run both ways) per mesh point. An example is shown in Figure 8.




Figure 4.8: Deconvolution by a filter whose autocorrelation is the two-dimensional Laplacian operator. Amounts to solving the Poisson equation. Left is \mathcal{Q} ; Middle is \mathcal{Q}/\mathcal{H} ; Right is $(\mathcal{Q}/\mathcal{H})/\mathcal{H}'$. hlx/helicon lapfac

Figure 8 contains both the helix derivative and its inverse. Contrast them to the x - or y -derivatives (doublets) and their inverses (axis-parallel lines in the (x, y) -plane). Simple derivatives are highly directional whereas the helix derivative is only slightly directional achieving its meagre directionality entirely from its phase spectrum.

In practice we often require an isotropic filter. Such a filter is a function of $k_r = \sqrt{k_x^2 + k_y^2}$. It could be represented as a sum of helix derivatives to integer powers.

4.2.1 Matrix view of the helix

Physics on a helix can be viewed thru the eyes of matrices and numerical analysis. This is not easy because the matrices are so huge. Discretize the (x, y) -plane to an $N \times M$ array and pack the array into a vector of $N \times M$ components. Likewise pack the Laplacian operator $\partial_{xx} + \partial_{yy}$ into a matrix. For a 4×3 plane, that matrix is shown in equation (12).

$$-\nabla^2 = \left[\begin{array}{cccc|cccc|cccc} 4 & -1 & . & . & -1 & . & . & . & . & . & . & . \\ -1 & 4 & -1 & . & . & -1 & . & . & . & . & . & . \\ . & -1 & 4 & -1 & . & . & -1 & . & . & . & . & . \\ . & . & -1 & 4 & h & . & . & -1 & . & . & . & . \\ \hline -1 & . & . & h & 4 & -1 & . & . & -1 & . & . & . \\ . & -1 & . & . & -1 & 4 & -1 & . & . & -1 & . & . \\ . & . & -1 & . & . & -1 & 4 & -1 & . & . & -1 & . \\ . & . & . & -1 & . & . & -1 & 4 & h & . & . & -1 \\ \hline . & . & . & . & -1 & . & . & h & 4 & -1 & . & . \\ . & . & . & . & . & -1 & . & . & -1 & 4 & -1 & . \\ . & . & . & . & . & . & -1 & . & . & -1 & 4 & -1 \\ . & . & . & . & . & . & . & -1 & . & . & -1 & 4 \end{array} \right] \quad (4.12)$$

The two-dimensional matrix of coefficients for the Laplacian operator is shown in (12), where, on a cartesian space, $h = 0$, and in the helix geometry, $h = -1$. (A similar partitioned matrix arises from packing a cylindrical surface into a 4×3 array.) Notice that the partitioning becomes transparent for the helix, $h = -1$. With the partitioning thus invisible, the matrix simply represents one-dimensional convolution and we have an alternative analytical approach, one-dimensional Fourier Transform. We often need to solve sets of simultaneous equations with a matrix similar to (12). The method we use is triangular factorization.

Although the autocorrelation \mathcal{R} has mostly zero values, the factored autocorrelation \mathcal{A} has a great number of nonzero terms, but fortunately they seem to be converging rapidly (in the middle) so truncation (of the middle coefficients) seems reasonable. I wish I could show you a larger matrix, but all I can do is to pack the signal \mathcal{A} into shifted columns of a lower triangular matrix \mathbf{A} like this:

$$\mathbf{A} = \left[\begin{array}{cccccccccccc} 1.8 & . & . & . & . & . & . & . & . & . & . & . \\ -.6 & 1.8 & . & . & . & . & . & . & . & . & . & . \\ 0.0 & -.6 & 1.8 & . & . & . & . & . & . & . & . & . \\ -.2 & 0.0 & -.6 & 1.8 & . & . & . & . & . & . & . & . \\ -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . & . & . & . & . & . \\ . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . & . & . & . & . \\ . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . & . & . & . \\ . & . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . & . & . \\ . & . & . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . & . \\ . & . & . & . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . & . \\ . & . & . & . & . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 & . \\ . & . & . & . & . & . & . & -.6 & -.2 & 0.0 & -.6 & 1.8 \end{array} \right] \quad (4.13)$$

If you will allow me some truncation approximations, I now claim that the laplacian represented by the matrix in equation (12) is factored into two parts $-\nabla^2 = \mathbf{A}'\mathbf{A}$ which are upper and lower triangular matrices whose product forms the autocorrelation seen in (12). Recall that triangular matrices allow quick solutions of simultaneous equations by backsubstitution. That is what we do with our deconvolution program.

4.3 CAUSALITY AND SPECTAL FACTORIZATION

Mathematics sometimes seems a mundane subject, like when it does the “accounting” for an engineer. Other times it brings unexpected amazing new concepts into our lives. This is the case with the study of causality and spectral factorization. There are many little-known, amazing, fundamental ideas here I would like to tell you about. We won’t get to the bottom of any of them but it’s fun and useful to see what they are and how to use them.

Start with an example. Consider a mechanical object. We can strain it and watch it stress or we can stress it and watch it strain. We feel knowledge of the present and past stress history is all we need to determine the present value of strain. Likewise, the converse, history of strain should tell us the stress. We could say there is a filter that takes us from stress to strain; likewise another filter takes us from strain to stress. What we have here is a pair of filters that are mutually inverse under convolution. In the Fourier domain, one is literally the inverse of the other. What is remarkable is that in the time domain, both are causal. They both vanish before zero lag $\tau = 0$.

Not all causal filters have a causal inverse. The best known name for one that does is “minimum-phase filter.” Unfortunately, this name is not suggestive of the fundamental property of interest, “causal with a causal (convolutional) inverse.” I could call it CwCI. An example of a causal filter without a causal inverse is the unit delay operator — with Z -transforms, the operator Z itself. If you delay something, you can’t get it back without seeing into the future, which you are not allowed to do. Mathematically, $1/Z$ cannot be expressed as a polynomial (actually, a convergent infinite series) in positive powers of Z .

Physics books don’t tell us where to expect to find transfer functions that are CwCI. I think I know why they don’t. Any causal filter has a “sharp edge” at zero time lag where it switches from nonresponsiveness to responsiveness. The sharp edge might cause the spectrum to be large at infinite frequency. If so, the inverse filter is small at infinite frequency. Either way, one of the two filters is unmanageable with Fourier transform theory which (you might have noticed in the mathematical fine print) requires signals (and spectra) to have finite energy which means the function must get real small in that immense space on the t -axis and the ω axis. It is impossible for a function to be small and its inverse be small. These imponderables get more manageable in the world of Time Series Analysis (discretized time axis).

4.3.1 The spectral factorization concept

Interesting questions arise when we are given a spectrum and find ourselves asking how to find a filter that has that spectrum. Is the answer unique? We’ll see not. Is there always an answer that is causal? Almost always, yes. Is there always an answer that is causal with a causal inverse (CwCI)? Almost always, yes.

Let us have an example. Consider a filter like the familiar time derivative $(1, -1)$ except let us downweight the -1 a tiny bit, say $(1, -\rho)$ where $0 < \rho < 1$. Now the filter $(1, -\rho)$ has a spectrum $(1 - \rho Z)(1 - \rho/Z)$ with autocorrelation coefficients $(-\rho, 1 + \rho^2, -\rho)$ that look a lot like a second derivative, but it is a tiny bit bigger in the middle. Two different waveforms, $(1, -\rho)$ and its time reverse both have the same autocorrelation. Spectral factorization could give us both $(1, -\rho)$ and $(\rho, -1)$ but we always want the one that is CwCI. The bad one is weaker on its first pulse. Its inverse is not causal. Below are two expressions for the filter inverse to $(\rho, -1)$, the first divergent (filter coefficients at infinite lag are infinitely strong), the second convergent but noncausal.

$$\frac{1}{\rho - Z} = \frac{1}{\rho} (1 + Z/\rho + Z^2/\rho^2 + \dots) \quad (4.14)$$

$$\frac{1}{\rho - Z} = \frac{-1}{Z} (1 + \rho/Z + \rho^2/Z^2 + \dots) \quad (4.15)$$

(Please multiply each equation by $\rho - Z$ and see it reduce to $1 = 1$).

So we start with a power spectrum and we should find a CwCI filter with that energy spectrum. If you input to the filter an infinite sequence of random numbers (white noise) you should output something with the original power spectrum.

We easily inverse Fourier transform the square root of the power spectrum getting a symmetrical time function, but we need a function that vanishes before $\tau = 0$. On the other hand, if we already had a causal filter with the correct spectrum we could manufacture many others. To do so all we need is a family of delay operators to convolve with. A pure delay filter does not change the spectrum of anything. Same for frequency-dependent delay operators. Here is an example of a frequency-dependent delay operator: First convolve with $(1, 2)$ and then deconvolve with $(2, 1)$. Both these have the same amplitude spectrum so their ratio has a unit amplitude (and nontrivial phase). If you multiply $(1 + 2Z)/(2 + Z)$ by its Fourier conjugate (replace Z by $1/Z$) the resulting spectrum is 1 for all ω .

Anything whose nature is delay is death to CwCI. The CwCI has its energy as close as possible to $\tau = 0$. More formally, my first book, FGDP, proves that the CwCI filter has for all time τ more energy between $t = 0$ and $t = \tau$ than any other filter with the same spectrum.

Spectra can be factorized by an amazingly wide variety of techniques, each of which gives you a different insight into this strange beast. They can be factorized by factoring polynomials, by inserting power series into other power series, by solving least squares problems, by taking logarithms and exponentials in the Fourier domain. I've coded most of them and still find them all somewhat mysterious.

Theorems in Fourier analysis can be interpreted physically in two different ways, one as given, the other with time and frequency reversed. For example, convolution in one domain amounts to multiplication in the other. If we were to express the CwCI concept with reversed domains, instead of saying the “energy comes as quick as possible after $\tau = 0$ ” we would say “the frequency function is as close to $\omega = 0$ as possible.” In other words, it is minimally wiggly with time. Most applications of spectral factorization begin with a spectrum, a real, positive function of frequency. I once achieved minor fame by starting with a real, positive function of space, a total magnetic field $\sqrt{H_x^2 + H_z^2}$ measured along the x -axis and I reconstructed the magnetic field components H_x and H_z that were minimally wiggly in space.

4.3.2 Cholesky decomposition

Conceptually the simplest computational method of spectral factorization might be “Cholesky decomposition.” For example, the matrix of (13) could have been found by Cholesky factorization of (12). The Cholesky algorithm takes a positive-definite matrix \mathbf{Q} and factors it into a triangular matrix times its transpose, say $\mathbf{Q} = \mathbf{T}'\mathbf{T}$.

It is easy to reinvent the Cholesky factorization algorithm. To do so, simply write all the components of a 3×3 triangular matrix \mathbf{T} and then explicitly multiply these elements times the transpose matrix \mathbf{T}' . You will find that you have everything you need to recursively build the elements of \mathbf{T} from the elements of \mathbf{Q} . Likewise for a 4×4 matrix, etc.

The 1×1 case shows that the Cholesky algorithm requires square roots. Matrix elements are not always numbers. Sometimes they are polynomials such as Z -transforms. To avoid square roots there is a variation of the Cholesky method. In this variation, we factor \mathbf{Q} into $\mathbf{Q} = \mathbf{T}'\mathbf{D}\mathbf{T}$ where \mathbf{D} is a diagonal matrix.

Once a matrix has been factored into upper and lower triangles, solving simultaneous equations is simply a matter of two backsubstitutions: (We looked at a special case of backsubstitution with equation (1.23).) For example, we often encounter simultaneous equations of the form $\mathbf{B}'\mathbf{B}\mathbf{m} = \mathbf{B}'\mathbf{d}$. Suppose the positive-definite matrix $\mathbf{B}'\mathbf{B}$ has been factored into triangle form $\mathbf{T}'\mathbf{T}\mathbf{m} = \mathbf{B}'\mathbf{d}$. To find \mathbf{m} we first backsolve $\mathbf{T}'\mathbf{x} = \mathbf{B}'\mathbf{d}$ for the vector \mathbf{x} . Then we backsolve $\mathbf{T}\mathbf{m} = \mathbf{x}$. When \mathbf{T} happens to be a band matrix, then the first backsubstitution is filtering down a helix and the second is filtering back up it. Polynomial division is a special case of back substitution.

Poisson’s equation $\nabla^2\mathbf{p} = -\mathbf{q}$ requires boundary conditions which we can honor when we filter starting from both ends. We cannot simply solve Poisson’s equation as an initial-value problem. We could insert the laplace operator into the polynomial division program, but the solution would diverge.

Being a matrix method, the Cholesky method of factorization has a cost proportional to the cube of the size of the matrix. Because our problems are very large and because the Cholesky method does not produce a useful result if we stop part way to completion, we look further. The Cholesky method is a powerful method but it does more than we require. The Cholesky method does not require band matrices, yet these matrices are what we very often find in applications, so we seek methods that take advantage of the special properties of band matrices.

4.3.3 Toeplitz methods

Band matrices are often called Toeplitz matrices. In the subject of Time Series Analysis are found spectral factorization methods that require computations proportional to the dimension of the matrix squared. They can often be terminated early with a reasonable partial result. Two Toeplitz methods, the Levinson method and the Burg method are described in my first textbook, FGDP. Our interest is multidimensional data sets so the matrices of interest are truly huge and the cost of Toeplitz methods is proportional to the square of the matrix size. Thus, before we find Toeplitz methods especially useful, we may need to find ways to take advantage of the sparsity of our filters.

4.3.4 Kolmogoroff spectral factorization

With Fourier analysis we find a method of spectral factorization that is as fast as Fourier transformation, namely $N \log N$ for a matrix of size N . This is very appealing. An earlier version of this book included such an algorithm. Pedagogically, I didn't like it in this book because it requires lengthy off-topic discussions of Fourier analysis which are already found in both my first book FGDP and my third book PVI.

The Kolmogoroff calculation is based on the logarithm of the spectrum. The logarithm of zero is minus infinity — an indicator that perhaps we cannot factorize a spectrum which becomes zero at any frequency. Actually, the logarithmic infinity is the gentlest kind. The logarithm of the smallest nonzero value in single precision arithmetic is about -36 which might not ruin your average calculation. Mathematicians have shown that the integral of the logarithm of the spectrum must be bounded so that some isolated zero values of the spectrum are not a disaster. In other words, we can factor the (negative) second derivative to get the first derivative. This suggests we will never find a causal bandpass filter. It is a contradiction to desire both causality and a spectral band of zero gain.

The weakness of the Kolmogoroff method is related to its strength. Fourier methods strictly require the matrix to be a band matrix. A problem many people would like to solve is how to handle a matrix that is “almost” a band matrix — a matrix where any band changes slowly with location.

4.4 WILSON-BURG SPECTRAL FACTORIZATION

(If you are new to this material, you should pass over this section.) Spectral factorization is the job of taking a power spectrum and from it finding a causal (zero before zero time) filter with that spectrum. Methods for this task (there are many) not only produce a causal wavelet, but they typically produce one whose convolutional inverse is also causal. (This is called the “minimum phase” property.) In other words, with such a filter we can do stable deconvolution. Here I introduce a new method of spectral factorization that looks particularly suitable for the task at hand. I learned this new method from John Parker Burg who attributes it to an old paper by Wilson (I find Burg's explanation, below, much clearer than Wilson's.)

To invoke the factorization subroutine, you need to supply one side of an autocorrelation function. For example, let us specify the negative of the 2-D Laplacian (an autocorrelation) in a vector $\mathbf{n} = 256 \times 256$ points long.

```
rr[0]    = 4.
rr[1]    = -1.
rr[256]  = -1.
```

4.4.1 Wilson-Burg theory

Newton's iteration for square roots

$$a_{t+1} = \frac{1}{2} \left(a_t + \frac{s}{a_t} \right) \quad (4.16)$$

converges quadratically starting from any real initial guess a_0 except zero. When a_0 is negative, Newton's iteration converges to the negative square root.

Quadratic convergence means that the square of the error $a_t - \sqrt{s}$ at one iteration is proportional to the error at the next iteration

$$a_{t+1} - \sqrt{s} \sim (a_t - \sqrt{s})^2 = a_t^2 - 2a_t\sqrt{s} + s > 0 \quad (4.17)$$

so, for example if the error is one significant digit at one iteration, at the next iteration it is two digits, then four, etc. We cannot use equation (17) in place of the Newton iteration itself, because it uses the answer \sqrt{s} to get the answer a_{t+1} , and also we need the factor of proportionality. Notice, however, if we take the factor to be $1/(2a_t)$, then \sqrt{s} cancels and equation (17) becomes itself the Newton iteration (16).

Another interesting feature of the Newton iteration is that all iterations (except possibly the initial guess) are above the ultimate square root. This is obvious from equation (17).

We can insert spectral functions in the Newton square-root iteration, for example $s(\omega)$ and $a(\omega)$. Where the first guess a_0 happens to match \sqrt{s} , it will match \sqrt{s} at all iterations. The Newton iteration is

$$2 \frac{a_{t+1}}{a_t} = 1 + \frac{s}{a_t^2} \quad (4.18)$$

Something inspires Wilson to express the spectrum $S = \bar{A}A$ as a Z -transform and then write the iteration

$$\frac{\bar{A}_{t+1}(1/Z)}{\bar{A}_t(1/Z)} + \frac{A_{t+1}(Z)}{A_t(Z)} = 1 + \frac{S(Z)}{\bar{A}_t(1/Z) A_t(Z)} \quad (4.19)$$

Now we are ready for the algorithm: Compute the right side of (19) by polynomial division forwards and backwards and then add 1. Then abandon negative lags and take half of the zero lag. Now you have $A_{t+1}(Z)/A_t(Z)$. Multiply out (convolve) the denominator $A_t(Z)$, and you have the desired result $A_{t+1}(Z)$. Iterate as long as you wish.

(Parenthetically, for those people familiar with the idea of minimum phase (if not, see FGDP or PVI), we show that $A_{t+1}(Z)$ is minimum phase: Both sides of (19) are positive, as noted earlier. Both terms on the right are positive. Since the Newton iteration always overestimates, the 1 dominates the rightmost term. After masking off the negative powers of Z (and half the zero power), the right side of (19) adds two wavelets. The $1/2$ is wholly real, and hence its real part always dominates the real part of the rightmost term. Thus (after masking negative powers) the wavelet on the right side of (19) has a positive real part, so the phase cannot loop about the origin. This wavelet multiplies $A_t(Z)$ to give the final wavelet $A_{t+1}(Z)$ and the product of two minimum-phase wavelets is minimum phase.)

The input of the program is the spectrum $S(Z)$ and the output is the factor $A(Z)$, a function with the spectrum $S(Z)$. I mention here that in later chapters of this book, the factor $A(Z)$ is known as the inverse Prediction-Error Filter (PEF). In the Wilson-Burg code below, $S(Z)$ and $A(Z)$ are Z -transform polynomials but their lead coefficients are extracted off, so for example, $A(z) = (a_0) + (a_1 Z + a_2 Z^2 + \dots)$ is broken into the two parts **a0** and **aa**.

user/gee/wilson.c

```

44 float wilson_factor(int niter      /* number of iterations */,
45                    float s0       /* zero-lag auto-correlation */,
46                    sf_filter ss    /* input auto-correlation */,
47                    sf_filter aa    /* output factor */,
48                    bool verb       /* verbosity flag */,
49                    float tol       /* tolerance */)
50 /*< Factor >*/
51 {
52     float eps;
53     int i, iter;
54
55     for(i=0; i < n2; i++) au[i] = 0.;
56     au[n-1] = s0;
57     b[0] = 1.;      /* initialize */
58
59     for(i=0; i < ss->nh; i++) { /* symmetrize input auto */
60         au[n-1+ss->lag[i]] = ss->flt[i];
61         au[n-1-ss->lag[i]] = ss->flt[i];
62     }
63
64     sf_helicon_init(aa);          /* multiply polynoms */
65     polydiv_init(n2, aa);        /* divide polynoms */
66     for(iter=0; iter < niter; iter++) {
67         polydiv_lop(false, false, n2, n2, au, bb); /* S/A */
68         polydiv_lop(true, false, n2, n2, cc, bb); /* S/(AA') */
69         eps = 0.;
70         for(i=1; i < n; i++) { /* b = plusside(1+cc) */
71             b[i] = 0.5*(cc[n-1+i] + cc[n-1-i]) / cc[n-1];
72             if (fabs(b[i]) > eps) eps = fabs(b[i]);
73         }
74         if (verb) sf_warning("wilson %d %f", iter, eps);
75         if (eps < tol) break;
76
77         sf_helicon_lop(false, false, n, n, b, c);
78     /* c = A b */
79
80     /* put on helix */
81     for(i=0; i < aa->nh; i++) aa->flt[i] = c[aa->lag[i]];
82 }
83 return sqrtf(cc[n-1]);
84 }

```

EXERCISES:

- 1 Fomel's factorization: A simple trick to avoid division in square root computation is to run Newton's method on the inverse square root instead. The iteration is then $R' = \frac{1}{2}R(3 - R^2X^2)$ where R converges (quadratically) to $1/\sqrt{X^2}$. To get the square root, just multiply R by X^2 . This leads to a reciprocal version of the Wilson-Burg algorithm. $A'/A + \bar{A}'/\bar{A} = 3 - A\bar{A}S$ Here is how it can work: Construct an inverse autocorrelation — for example, an ideal isotropic smoother; make a guess for A (min-phase roughener); iterate: (1) compute $3 - A\bar{A}S$, (2) take its causal part, (3) convolve with A to get A' . Each iteration involves just three convolutions (could be even done without helix).

4.5 HELIX LOW-CUT FILTER


If you want to see some tracks on the side of a hill, you want to subtract the hill and see only the tracks. Usually, however, you don't have a very good model for the hill. As an expedient you could apply a low-cut filter to remove all slowly variable functions of altitude. In chapter 1 we found the Sea of Galilee in Figure 1.3 to be too smooth for viewing pleasure so we made the roughened versions in Figure 1.6 using a filter based on equation (2.13), a one-dimensional filter that we could apply over the x -axis or the y -axis. In Fourier space such a filter has a response function of k_x or a function of k_y . The isotropy of physical space tells us it would be more logical to design a filter that is a function of $k_x^2 + k_y^2$. In Figure 7 we saw that the helix derivative \mathbf{H} does a nice job. The Fourier magnitude of its impulse response is $k_r = \sqrt{k_x^2 + k_y^2}$. There is a little anisotropy connected with phase (which way should we wind the helix, on x or y ?) but it is not nearly so severe as that of either component of the gradient, the two components having wholly different spectra, amplitude $|k_x|$ or $|k_y|$.

It is nice having the 2-D helix derivative, but we can imagine even nicer 2-D low-cut filters. In one dimension, equation (2.13) and (2.13) we designed a filters with an adjustable parameter, a cutoff frequency. We don't have such an object in 2-D so I set out to define one. It came out somewhat abstract and complicated, and didn't work very well, but along the way I found a simpler parameter that is very effective in practice. We'll look at it first.

First I had a problem preparing Figure 9. It shows shows the application of the helix derivative to a medical X-ray. The problem was that the original X-ray was all positive values of brightness so there was a massive amount of spatial low frequency present. Obviously an x -derivative or a y -derivative would eliminate the low frequency, but the helix derivative did not. This unpleasant surprise arises because the filter in equation (11) was truncated after a finite number of terms. Adding up the terms actually displayed in equation (11), they sum to .183 whereas theoretically the sum of all the terms should be zero. From the ratio of .183/1.791 we can say that the filter pushes zero frequency amplitude 90% of the way to zero value. When the image contains very much zero frequency amplitude, this is not good enough. Better results could be obtained with more coefficients, and I did use more coefficients, but simply removing the mean saved me from needing a costly number of filter coefficients.

We can visualize a plot of the magnitude of the 2-D Fourier transform of the filter (11).

Figure 4.9: Mammogram (medical X-ray). The cancer is the “spoked wheel.” (I apologize for the inability of paper publishing technology to exhibit a clear grey image.) The white circles are metal foil used for navigation. The little halo around a circle exhibits the impulse response of the helix derivative. hlx/mam mam



It is a 2-D function of k_x and k_y and it should resemble $k_r = \sqrt{k_x^2 + k_y^2}$. It does look like this even when the filter (11) has been truncated. The point of the cone $k_r = \sqrt{k_x^2 + k_y^2}$ becomes rounded and the truncated approximation of k_r does not reach zero at the origin of the (k_x, k_y) -plane. We can force it to vanish at zero frequency by subtracting .183 from the lead coefficient 1.791. I did not do that subtraction in Figure 10 which explains the whiteness in the middle of the lake.

Now let us return to my more logical but less effective approach. I prepared a half dozen medical X-rays like Figure 9. The doctor brought her young son to my office one evening to evaluate the results. In a dark room I would show the original X-ray on a big screen and then suddenly switch to the helix derivative. Every time I did this, her son would exclaim “Wow!” The doctor was not so easily impressed, however. She was not accustomed to the unfamiliar image. Fundamentally, the helix derivative applied to her data does compress the dynamic range making weaker features more readily discernable. We were sure of this from theory and from various geophysical examples. The subjective problem was her unfamiliarity with our display. I found that I could always spot anomalies more quickly on the filtered display, but then I would feel more comfortable when I would discover those same anomalies also present (though less evident) in the original data. Thinking this through, I decided the doctor would likely have been more impressed had I used a spatial lowcut filter instead of the helix derivative. That would have left the details of her image (above the cutoff frequency) unchanged altering only the low frequencies, thereby allowing me to increase the gain.

In 1-D we easily make a low-cut filter by compounding a first derivative (which destroys low frequencies) with a leaky integration (which undoes the derivative at all other frequencies). We can do likewise with a second derivative. In Z -transform notation, we would use something like $(-Z^{-1} + 2.00 - Z)/(-Z^{-1} + 2.01 - Z)$. (The numerical choice of the .01




Figure 4.10: Galilee roughened by gradient and by helical derivative. hlx/helgal helgal

controls the cutoff frequency.) We could use spectral factorization to break this spectrum into causal and anticausal factors. The analogous filter in 2-D is $-\nabla^2/(-\nabla^2 + k_0^2)$ which could also be factored as we did the helix derivative. I tried it. I ran into the problem that my helix derivative operator had a practical built-in parameter, the number of coefficients, which also behaves like a cutoff frequency. If I were to continue this project, I would use expressions for $-\nabla^2/(-\nabla^2 + k_0^2)$ directly in the Fourier domain where there is only one adjustable parameter, the cutoff frequency k_0 , and there is no filter length to confuse the issue and puff-up the costs.

A final word about the doctor. As she was about to leave my office she suddenly asked whether I had scratched one of her X-rays. We were looking at the helix derivative and it did seem to show a big scratch. What should have been a line was broken into a string of dots. I apologized in advance and handed her the original film negatives which she proceeded to inspect. “Oh,” she said, “Bad news. There are calcification nodules along the ducts.” So the scratch was not a scratch, but an important detail that had not been noticed on the original X-ray.

In preparing an illustration for here, I learned one more lesson. The scratch was small, so I enlarged a small portion of the mammogram for display. The very process of selecting a small portion followed by scaling the amplitude between maximum and minimum darkness of printer ink had the effect enhancing the visibility of the scratch on the mammogram itself. Now Figure 11 shows it to be perhaps even clearer than on the helix derivative.




Figure 4.11: Not a scratch. `hlx/mam scratch`

4.6 THE MULTIDIMENSIONAL HELIX

Till now the helix idea was discussed as if it were merely a two-dimensional concept. Here we explore its multidimensional nature. Our main goal is to do multidimensional convolution with a one-dimensional convolution program. This allows us to do multidimensional deconvolution with a one-dimensional deconvolutional program which is “magic”, i.e. many novel applications will follow.

We do multidimensional deconvolution with causal (one-sided) one-dimensional filters. Equation (7) shows such a one-sided filter as it appears at the end of a 2-D helix. Figure 12 shows it in three dimensions. The top plane in Figure 12 is the 2-D filter seen in equation (7). The top plane can be visualized as the area around the end of a helix. Above the top plane are zero-valued anticausal filter coefficients.




Figure 4.12: A 3-D causal filter at the starting end of a 3-D helix.

`hlx/XFig 3dpef`

It is natural to ask, “why not put the ‘1’ on a corner of the cube?” We could do that, but that is not the most general possible form. A special case of Figure 12, stuffing much of the volume with lots of zeros would amount to a ‘1’ on a corner. On the other hand, if we assert the basic form has a ‘1’ on a corner we cannot get Figure 12 as a special case. In a later chapter we’ll see that we often need as many coefficients as we can have near the ‘1’. In Figure 12 we lose only those neighboring coefficients that 1-D causality requires.

Geometrically, the three-dimensional generalization of a helix is like string on a spool, but that analogy does not illuminate our underlying conspiracy, which is to represent multidimensional convolution and deconvolution as one-dimensional.

4.7 SUBSCRIPTING A MULTIDIMENSIONAL HELIX

Basic utilities transform back and forth between multidimensional matrix coordinates and helix coordinates. The essential module used repeatedly in applications later in this book is `createhelix` on page 28. We begin here from its intricate underpinnings.

Fortran77 has a concept of a multidimensional array being equivalent to a one-dimensional array. Given that the hypercube specification `nd=(n1,n2,n3,...)` defines the storage dimension of a data array, we can refer to a data element as either `dd(i1,i2,i3,...)` or `dd(i1 +n1*(i2-1) +n1*n2*(i3-1) +...)`. The helix says to refer to the multidimensional data by its equivalent one-dimensional index (sometimes called its vector subscript or linear subscript).

The filter, however, is a much more complicated story than the data: First, we require all filters to be causal. In other words, the Laplacian doesn’t fit very well, since it is intrinsically noncausal. If you really want noncausal filters, you will need to provide your own time shifts outside the tools supplied here. Second, a filter is usually a small hypercube, say `aa(a1,a2,a3,...)` and would often be stored as such. For the helix we must store it in a special one-dimensional form. Either way, the numbers `na= (a1,a2,a3,...)` specify the dimension of the hypercube. In cube form, the entire cube could be indexed multidimensionally as `aa(i1,i2,...)` or it could be indexed one-dimensionally as `aa(ia,1,1,...)` or sometimes `aa[ia]` by letting `ia` cover a large range. When a filter cube is stored in its normal “tightly packed” form the formula for computing its one-dimensional index `ia` is

$$ia = i1 + a1*i2 + a1*a2*i3 + \dots$$

When the filter cube is stored in an array with the same dimensions as the data, `data...[n3][n2][n1]`, the formula for `ia` is

$$ia = i1 + n1*i2 + n1*n2*i3 + \dots$$

Module `decart` below contains two subroutines that explicitly provide us the transformations between the linear index `i` and the multidimensional indices `ii= (i1,i2,...)`. The two subroutines have the logical names `cart2line` and `line2cart`.

The fortran linear index is closely related to the helix. There is one major difference, however, and that is the origin of the coordinates. To convert from the linear index to

flt/lib/decart.c

```

22 void sf_line2cart(int dim      /* number of dimensions */,
23                  const int* nn /* box size [dim] */,
24                  int i         /* line coordinate */,
25                  int* ii       /* cartesian coordinates [dim] */)
26 /*< Convert line to Cartesian >*/
27 {
28     int axis;
29
30     for (axis = 0; axis < dim; axis++) {
31         ii[axis] = i%nn[axis];
32         i /= nn[axis];
33     }
34 }
35
36 int sf_cart2line(int dim      /* number of dimensions */,
37                 const int* nn /* box size [dim] */,
38                 const int* ii /* cartesian coordinates [dim] */)
39 /*< Convert Cartesian to line >*/
40 {
41     int i, axis;
42
43     if (dim < 1) return 0;
44
45     i = ii[dim-1];
46     for (axis = dim-2; axis >= 0; axis--) {
47         i = i*nn[axis] + ii[axis];
48     }
49     return i;
50 }

```

the helix lag coordinate, we need to subtract the fortran linear index of the “1.0” which is usually taken at `center= (1+a1/2, 1+a2/2, ..., 1)`. (On the last dimension, there is no shift because nobody stores the volume of zero values that would occur before the 1.0.) The `decart` module fails for negative subscripts. Thus we need to be careful to avoid thinking of the filter’s 1.0 (shown in Figure 12) as the origin of the multidimensional coordinate system although the 1.0 is the origin in the one-dimensional coordinate system.

Even in one dimension (see the matrix in equation (1.4)), to define a filter *operator* we need to know not only filter coefficients and a filter length, but we also need to know the data length. To define a multidimensional filter using the helix idea, besides the properties intrinsic to the filter, we also need to know the circumference of the helix, i.e., the length on the 1-axis of the data’s hypercube as well as the other dimensions `nd=(n1,n2,...)` of the data’s hypercube.

Thinking about convolution on the helix, it is natural to think about the filter and data being stored in the same way, that is, by reference to the data size. This would waste so much space, however, that our helix filter module `helix` on page 7 instead stores the filter coefficients in one vector and their lags in another. The *i*-th coefficient value of the filter goes in `aa->flt[i]` and the *i*-th lag `ia[i]` goes in `aa->lag[i]`. The lags are the same as the fortran linear index except for the overall shift of the 1.0 of a cube of data dimension `nd`. Our module for convolution on a helix, `helicon`, has already an implicit “1.0” at the filter’s zero lag so we do not store it. (It is an error to do so.)

Module `createhelix` on page 28 allocates memory for a helix filter and builds filter lags along the helix from the hypercube description. The hypercube description is not the literal cube seen in Figure 12 but some integers specifying that cube: the data cube dimensions `nd`, likewise the filter cube dimensions `na`, the parameter `center` identifying the location of the filter’s “1.0”, and a `gap` parameter used in a later chapter. To find the lag table, module `createhelix` first finds the fortran linear index of the `center` point on the filter hypercube. Everything before that has negative lag on the helix and can be ignored. (Likewise, in a later chapter we see a `gap` parameter that effectively sets even more filter coefficients to zero so their lags can be ignored too.) Then it sweeps from the center point over the rest of the filter hypercube calculating for a data-sized cube `nd`, the fortran linear index of each filter element. Near the end of the code you see the calculation of a parameter `lag0d`. This is the count of the number of zeros that a data-sized fortran array would store in a filter cube before the filter’s 1.0. We need to subtract this shift from the filter’s fortran linear index to get the lag on the helix.

A filter can be represented literally as a multidimensional cube like equation (7) shows us in two dimensions or like Figure 12 shows us in three dimensions. Unlike the helical form, in literal cube form, the zeros preceding the “1.0” are explicitly present so `lag0` needs to be added back in to get the fortran subscript. To convert a helix filter `aa` to fortran’s multidimensional hypercube `cube(n1,n2,...)` is module `boxfilter`: The `boxfilter` module is normally used to display or manipulate a filter that was estimated in helical form (usually estimated by the least-squares method).

A reasonable arrangement for a small 3-D filter is `na={5,3,2}` and `center={3,2,1}`. Using these arguments, I used `createhelix` on page 28 to create a filter. I set all the helix filter coefficients to 2. Then I used module `boxfilter` on page 29 to put it in a convenient form for display. Finally, I printed it:

```

                                user/gee/createhelix.c
36  sf_filter createhelix(int ndim      /* number of dimensions */,
37                        int* nd      /* data size [ndim] */,
38                        int* center  /* filter center [ndim] */,
39                        int* gap     /* filter gap [ndim] */,
40                        int* na      /* filter size [ndim] */)
41  /*< allocate and output a helix filter >*/
42  {
43      sf_filter aa;
44      int ii[SF_MAX_DIM], na123, ia, nh, lag0a, lag0d, *lag, i;
45      bool skip;
46
47      for (na123 = 1, i=0; i < ndim; i++) na123 *= na[i];
48      lag = (int*) alloca(na123*sizeof(int));
49
50      /* index pointing to the "1.0" */
51      lag0a = sf_cart2line (ndim, na, center);
52
53      nh=0;
54      /* loop over linear index. */
55      for (ia = 1+lag0a; ia < na123; ia++) {
56          sf_line2cart(ndim, na, ia, ii);
57
58          skip = false;
59          for (i=0; i < ndim; i++) {
60              if (ii[i] < gap[i]) {
61                  skip = true;
62                  break;
63              }
64          }
65          if (skip) continue;
66
67          lag[nh] = sf_cart2line(ndim, nd, ii);
68          nh++;                                /* got another live one */
69      }
70      /* center shift for nd cube */
71      lag0d = sf_cart2line(ndim, nd, center);
72      aa = sf_allocatehelix(nh); /* nh becomes size of filter */
73
74      for (ia=0; ia < nh; ia++) {
75          aa->lag[ia] = lag[ia] - lag0d;
76          aa->flt[ia] = 0.;
77      }
78
79      return aa;
80  }

```

user/gee/boxfilter.c

```

24 void box (int dim          /* number of dimaneions */,
25           const int *nd    /* data size [dim] */,
26           const int *center /* filter center [dim] */,
27           const int *na     /* filter size [dim] */,
28           const sf_filter aa /* input filter */,
29           int nc            /* box size */,
30           float* cube       /* output box [nc] */)
31 /*< box it >*/
32 {
33     int ii [SF_MAX_DIM];
34     int j, lag0a, lag0d, id, ia;
35
36     for (ia=0; ia < nc; ia++) {
37         cube[ia] = 0.;
38     }
39     lag0a = sf_cart2line(dim, na, center); /* 1.0 in na. */
40     cube[lag0a] = 1.;                      /* place it. */
41     lag0d = sf_cart2line(dim, nd, center); /* 1.0 in nd. */
42     for (j=0; j < aa->nh; j++) { /* inspect the entire helix */
43         id = aa->lag[j] + lag0d;
44         sf_line2cart(dim, nd, id, ii); /* ii = cartesian */
45         ia = sf_cart2line(dim, na, ii); /* ia = linear in aa */
46         cube[ia] = aa->flt[j];          /* copy the filter */
47     }
48 }

```

0.000	0.000	0.000	0.000	0.000
0.000	0.000	1.000	2.000	2.000
2.000	2.000	2.000	2.000	2.000

2.000	2.000	2.000	2.000	2.000
2.000	2.000	2.000	2.000	2.000
2.000	2.000	2.000	2.000	2.000

Different data sets have different sizes. To convert a helix filter from one data size to another, we could drop the filter into a cube with module `cube`. Then we could extract it with module `unbox` specifying any data set size we wish. Instead we use module `regrid` prepared by Sergey Fomel which does the job without reference to an underlying filter cube. He explains his `regrid` module thus:

Imagine a filter being cut out of a piece of paper and glued on another paper, which is then rolled to form a helix.

We start by picking a random point (let's call it `rand`) in the cartesian grid and placing the filter so that its center (the leading 1.0) is on top of that point. `rand` should be larger than (or equal to) `center` and smaller than `min (nold, nnew)`, otherwise the filter might stick outside the grid (our piece of paper.) `rand=nold/2` will do (assuming the filter is small), although nothing should change if you replace `nold/2` with a random integer array between `center` and `nold - na`.

The linear coordinate of `rand` is `h0` on the old helix and `h1` on the new helix. Recall that the helix lags `aa->lag` are relative to the center. Therefore, we need to add `h0` to get the absolute helix coordinate (`h`). Likewise, we need to subtract `h1` to return to a relative coordinate system.

user/gee/regrid.c

```

24 void regrid( int dim          /* number of dimensions */,
25              const int* nold  /* old data size [dim] */,
26              const int* nnew  /* new data size [dim] */,
27              sf_filter aa     /* modified filter */)
28 /*< change data size >*/
29 {
30     int i, h0, h1, h, ii[SF_MAX_DIM];
31
32     for (i=0; i < dim; i++) {
33         ii[i] = nold[i]/2-1;
34     }
35
36     h0 = sf_cart2line( dim, nold, ii ); /* midpoint lag on nold */
37     h1 = sf_cart2line( dim, nnew, ii ); /*                on nnew */
38     for (i=0; i < aa->nh; i++) { /* for all filter coefficients */
39         h = aa->lag[i] + h0;
40         sf_line2cart( dim, nold, h, ii );
41         aa->lag[i] = sf_cart2line( dim, nnew, ii ) - h1;
42     }
43 }

```


Chapter 5

Preconditioning

When I first realized that practical imaging methods in widespread industrial use amounted merely to the adjoint of forward modeling, I (and others) thought an easy way to achieve fame and fortune would be to introduce the first steps towards inversion along the lines of Chapter 2. Although inversion generally requires a prohibitive number of steps, I felt that moving in the gradient direction, the direction of steepest descent, would move us rapidly in the direction of practical improvements. This turned out to be optimistic. It was too slow. But then I learned about the conjugate gradient method that spectacularly overcomes a well-known speed problem with the method of steepest descents. I came to realize that it was still too slow. I learned this by watching the convergence in Figure 5.6. This led me to the helix method in Chapter . Here we'll see how it speeds many applications.

We'll also come to understand why the gradient is such a poor direction both for steepest descent and for conjugate gradients. An indication of our path is found in the contrast between and exact solution $\mathbf{m} = (\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{d}$ and the gradient $\Delta\mathbf{m} = \mathbf{A}'\mathbf{d}$ (which is the first step starting from $\mathbf{m} = \mathbf{0}$). Notice that $\Delta\mathbf{m}$ differs from \mathbf{m} by the factor $(\mathbf{A}'\mathbf{A})^{-1}$. This factor is sometimes called a spectrum and in some situations it literally is a frequency spectrum. In these cases, $\Delta\mathbf{m}$ simply gets a different spectrum from \mathbf{m} and many iterations are required to fix it. Here we'll find that for many problems, “preconditioning” with the helix is a better way.

5.1 PRECONDITIONED DATA FITTING

Iterative methods (like conjugate-directions) can sometimes be accelerated by a change of variables. The simplest change of variable is called a “trial solution”. Formally, we write the solution as

$$\mathbf{m} = \mathbf{S}\mathbf{p} \quad (5.1)$$

where \mathbf{m} is the map we seek, columns of the matrix \mathbf{S} are “shapes” that we like, and coefficients in \mathbf{p} are unknown coefficients to select amounts of the favored shapes. The variables \mathbf{p} are often called the “preconditioned variables”. It is not necessary that \mathbf{S} be an invertible matrix, but we'll see later that invertibility is helpful. Take this trial solution and insert it into a typical fitting goal

$$\mathbf{0} \approx \mathbf{F}\mathbf{m} - \mathbf{d} \quad (5.2)$$

and get

$$\mathbf{0} \approx \mathbf{F}\mathbf{S}\mathbf{p} - \mathbf{d} \quad (5.3)$$

We pass the operator $\mathbf{F}\mathbf{S}$ to our iterative solver. After finding the best fitting \mathbf{p} , we merely evaluate $\mathbf{m} = \mathbf{S}\mathbf{p}$ to get the solution to the original problem.

We hope this change of variables has saved effort. For each iteration, there is a little more work: Instead of the iterative application of \mathbf{F} and \mathbf{F}' we have iterative application of $\mathbf{F}\mathbf{S}$ and $\mathbf{S}'\mathbf{F}'$.

Our hope is that the number of iterations decreases because we are clever, or because we have been lucky in our choice of \mathbf{S} . Hopefully, the extra work of the preconditioner operator \mathbf{S} is not large compared to \mathbf{F} . If we should be so lucky that $\mathbf{S} = \mathbf{F}^{-1}$, then we get the solution immediately. Obviously we would try any guess with $\mathbf{S} \approx \mathbf{F}^{-1}$. Where I have known such \mathbf{S} matrices, I have often found that convergence is accelerated, but not by much. Sometimes it is worth using $\mathbf{F}\mathbf{S}$ for a while in the beginning, but later it is cheaper and faster to use only \mathbf{F} . A practitioner might regard the guess of \mathbf{S} as prior information, like the guess of the initial model \mathbf{m}_0 .

For a square matrix \mathbf{S} , the use of a preconditioner should not change the ultimate solution. Taking \mathbf{S} to be a tall rectangular matrix, reduces the number of adjustable parameters, changes the solution, gets it quicker, but lower resolution.

5.1.1 Preconditioner with a starting guess

In many applications, for many reasons, we have a starting guess \mathbf{m}_0 of the solution. You might worry that you could not find the starting preconditioned variable $\mathbf{p}_0 = \mathbf{S}^{-1}\mathbf{m}_0$ because you did not know the inverse of \mathbf{S} . The way to avoid this problem is to reformulate the problem in terms of a new variable $\tilde{\mathbf{m}}$ where $\mathbf{m} = \tilde{\mathbf{m}} + \mathbf{m}_0$. Then $\mathbf{0} \approx \mathbf{F}\mathbf{m} - \mathbf{d}$ becomes $\mathbf{0} \approx \mathbf{F}\tilde{\mathbf{m}} - (\mathbf{d} - \mathbf{F}\mathbf{m}_0)$ or $\mathbf{0} \approx \mathbf{F}\tilde{\mathbf{m}} - \tilde{\mathbf{d}}$. Thus we have accomplished the goal of taking a problem with a nonzero starting model and converting it a problem of the same type with a zero starting model. Thus we do not need the inverse of \mathbf{S} because the iteration starts from $\tilde{\mathbf{m}} = \mathbf{0}$ so $\mathbf{p}_0 = \mathbf{0}$.

5.2 PRECONDITIONING THE REGULARIZATION

The basic formulation of a geophysical estimation problem consists of setting up *two* goals, one for data fitting, and the other for model shaping. With two goals, preconditioning is somewhat different. The two goals may be written as:

$$\mathbf{0} \approx \mathbf{F}\mathbf{m} - \mathbf{d} \quad (5.4)$$

$$\mathbf{0} \approx \mathbf{A}\mathbf{m} \quad (5.5)$$

which defines two residuals, a so-called “data residual” and a “model residual” that are usually minimized by conjugate-gradient, least-squares methods.

To fix ideas, let us examine a toy example. The data and the first three rows of the matrix below are random numbers truncated to integers. The model roughening operator \mathbf{A} is a first differencing operator times 100.

d(m)	F(m,n)										iter	Norm
---	-----										---	-----
41.	-55.	-90.	-24.	-13.	-73.	61.	-27.	-19.	23.	-55.	1	20.00396538
33.	8.	-86.	72.	87.	-41.	-3.	-29.	29.	-66.	50.	2	12.14780140
-58.	84.	-49.	80.	44.	-52.	-51.	8.	86.	77.	50.	3	8.94393635
0.	100.	0.	0.	0.	0.	0.	0.	0.	0.	0.	4	6.04517126
0.	-100.	100.	0.	0.	0.	0.	0.	0.	0.	0.	5	2.64737511
0.	0.	-100.	100.	0.	0.	0.	0.	0.	0.	0.	6	0.79238468
0.	0.	0.	-100.	100.	0.	0.	0.	0.	0.	0.	7	0.46083349
0.	0.	0.	0.	-100.	100.	0.	0.	0.	0.	0.	8	0.08301232
0.	0.	0.	0.	0.	-100.	100.	0.	0.	0.	0.	9	0.00542009
0.	0.	0.	0.	0.	0.	-100.	100.	0.	0.	0.	10	0.00000565
0.	0.	0.	0.	0.	0.	0.	-100.	100.	0.	0.	11	0.00000026
0.	0.	0.	0.	0.	0.	0.	0.	-100.	100.	0.	12	0.00000012
0.	0.	0.	0.	0.	0.	0.	0.	0.	-100.	100.	13	0.00000000

Notice at the tenth iteration, the residual suddenly plunges 4 significant digits. Since there are ten unknowns and the matrix is obviously full-rank, conjugate-gradient theory tells us to expect the exact solution at the tenth iteration. This is the first miracle of conjugate gradients. (The residual actually does not drop to zero. What is printed in the `Norm` column is the square root of the sum of the squares of the residual components at the `iter`-th iteration minus that at the last iteration.)

5.2.1 The second miracle of conjugate gradients

The second miracle of conjugate gradients is exhibited below. The data and data fitting matrix are the same, but the model damping is simplified.

d(m)	F(m,n)										iter	Norm
---	-----										---	-----
41.	-55.	-90.	-24.	-13.	-73.	61.	-27.	-19.	23.	-55.	1	3.64410686
33.	8.	-86.	72.	87.	-41.	-3.	-29.	29.	-66.	50.	2	0.31269890
-58.	84.	-49.	80.	44.	-52.	-51.	8.	86.	77.	50.	3	-0.00000021
0.	100.	0.	0.	0.	0.	0.	0.	0.	0.	0.	4	-0.00000066
0.	0.	100.	0.	0.	0.	0.	0.	0.	0.	0.	5	-0.00000080
0.	0.	0.	100.	0.	0.	0.	0.	0.	0.	0.	6	-0.00000065
0.	0.	0.	0.	100.	0.	0.	0.	0.	0.	0.	7	-0.00000088
0.	0.	0.	0.	0.	100.	0.	0.	0.	0.	0.	8	-0.00000074
0.	0.	0.	0.	0.	0.	100.	0.	0.	0.	0.	9	-0.00000035
0.	0.	0.	0.	0.	0.	0.	100.	0.	0.	0.	10	-0.00000037
0.	0.	0.	0.	0.	0.	0.	0.	100.	0.	0.	11	-0.00000018
0.	0.	0.	0.	0.	0.	0.	0.	0.	100.	0.	12	0.00000000
0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	100.	13	0.00000000

Even though the matrix is full-rank, we see the residual drop about 6 decimal places after the third iteration! This convergence behavior is well known in the computational mathematics literature. Despite its practical importance, it doesn't seem to have a name or identified discoverer. So I call it the "second miracle."

Practitioners usually don't like the identity operator for model-shaping. Generally they prefer to penalize wiggleness. For practitioners, the lesson of the second miracle of conjugate gradients is that we have a choice of many iterations, or learning to transform independent variables so that the regularization operator becomes an identity matrix. Basically, such a

transformation reduces the iteration count from something about the size of the model space to something about the size of the data space. Such a transformation is called preconditioning. In practice, data is often accumulated in bins. Then the iteration count is reduced (in principle) to the count of full bins and should be independent of the count of the empty bins. This allows refining the bins, enhancing the resolution.

More generally, the model goal $\mathbf{0} \approx \mathbf{A}\mathbf{m}$ introduces a roughening operator like a gradient, Laplacian (and in chapter 6 a Prediction-Error Filter (PEF)). Thus the model goal is usually a filter, unlike the data-fitting goal which involves all manner of geometry and physics. When the model goal is a filter its inverse is also a filter. Of course this includes multidimensional filters with a helix.

The preconditioning transformation $\mathbf{m} = \mathbf{S}\mathbf{p}$ gives us

$$\begin{aligned}\mathbf{0} &\approx \mathbf{F}\mathbf{S}\mathbf{p} - \mathbf{d} \\ \mathbf{0} &\approx \mathbf{A}\mathbf{S}\mathbf{p}\end{aligned}\tag{5.6}$$

The operator \mathbf{A} is a roughener while \mathbf{S} is a smoother. The choices of both \mathbf{A} and \mathbf{S} are somewhat subjective. This suggests that we eliminate \mathbf{A} altogether by *defining* it to be proportional to the inverse of \mathbf{S} , thus $\mathbf{A}\mathbf{S} = \mathbf{I}$. The fitting goals become

$$\begin{aligned}\mathbf{0} &\approx \mathbf{F}\mathbf{S}\mathbf{p} - \mathbf{d} \\ \mathbf{0} &\approx \epsilon \mathbf{p}\end{aligned}\tag{5.7}$$

which enables us to benefit from the “second miracle”. After finding \mathbf{p} , we obtain the final model with $\mathbf{m} = \mathbf{S}\mathbf{p}$.

5.2.2 Importance of scaling

Another simple toy example shows us the importance of scaling. We use the same example as above except that the i -th column is multiplied by $i/10$ which means the i -th model variable has been divided by $i/10$.

d(m)	F(m,n)										iter	Norm
---	-----										----	-----
41.	-6.	-18.	-7.	-5.	-36.	37.	-19.	-15.	21.	-55.	1	11.59544849
33.	1.	-17.	22.	35.	-20.	-2.	-20.	23.	-59.	50.	2	6.97337770
-58.	8.	-10.	24.	18.	-26.	-31.	6.	69.	69.	50.	3	5.64414406
0.	10.	0.	0.	0.	0.	0.	0.	0.	0.	0.	4	4.32118177
0.	0.	20.	0.	0.	0.	0.	0.	0.	0.	0.	5	2.64755201
0.	0.	0.	30.	0.	0.	0.	0.	0.	0.	0.	6	2.01631355
0.	0.	0.	0.	40.	0.	0.	0.	0.	0.	0.	7	1.23219979
0.	0.	0.	0.	0.	50.	0.	0.	0.	0.	0.	8	0.36649203
0.	0.	0.	0.	0.	0.	60.	0.	0.	0.	0.	9	0.28528941
0.	0.	0.	0.	0.	0.	0.	70.	0.	0.	0.	10	0.06712411
0.	0.	0.	0.	0.	0.	0.	0.	80.	0.	0.	11	0.00374284
0.	0.	0.	0.	0.	0.	0.	0.	0.	90.	0.	12	-0.00000040
0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	100.	13	0.00000000

We observe that solving the same problem for the scaled variables has required a severe increase in the number of iterations required to get the solution. We lost the benefit of the second CG miracle. Even the rapid convergence predicted for the 10-th iteration is delayed until the 12-th.

5.2.3 Statistical interpretation

This book is not a statistics book. Never-the-less, many of you have some statistical knowledge that allows you a statistical interpretation of these views of preconditioning.

A statistical concept is that we can combine many streams of random numbers into a composite model. Each stream of random numbers is generally taken to be uncorrelated with the others, to have zero mean, and to have the same variance as all the others. This is often abbreviated as IID, denoting Independent, Identically Distributed. Linear combinations like filtering and weighting operations of these IID random streams can build correlated random functions much like those observed in geophysics. A geophysical practitioner seeks to do the inverse, to operate on the correlated unequal random variables and create the statistical ideal random streams. The identity matrix required for the “second miracle”, and our search for a good preconditioning transformation are related ideas. The relationship will become more clear in chapter 6 when we learn how to estimate the best roughening operator \mathbf{A} as a prediction-error filter.

Two philosophies to find a preconditioner:

1. Dream up a smoothing operator \mathbf{S} .
2. Estimate a prediction-error filter \mathbf{A} , and then use its inverse $\mathbf{S} = \mathbf{A}^{-1}$.

Deconvolution on a helix is an all-purpose preconditioning strategy for multidimensional model regularization.

The outstanding acceleration of convergence by preconditioning suggests that the philosophy of image creation by optimization has a dual orthonormality: First, Gauss (and common sense) tells us that the data residuals should be roughly equal in size. Likewise in Fourier space they should be roughly equal in size, which means they should be roughly white, i.e. orthonormal. (I use the word “orthonormal” because white means the autocorrelation is an impulse, which means the signal is statistically orthogonal to shifted versions of itself.) Second, to speed convergence of iterative methods, we need a whiteness, another orthonormality, in the solution. The map image, the physical function that we seek, might not be itself white, so we should solve first for another variable, the whitened map image, and as a final step, transform it to the “natural colored” map.

5.2.4 The preconditioned solver

Summing up the ideas above, we start from fitting goals

$$\begin{aligned} \mathbf{0} &\approx \mathbf{Fm} - \mathbf{d} \\ \mathbf{0} &\approx \mathbf{Am} \end{aligned} \tag{5.8}$$

and we change variables from \mathbf{m} to \mathbf{p} using $\mathbf{m} = \mathbf{A}^{-1}\mathbf{p}$

$$\begin{aligned} \mathbf{0} &\approx \mathbf{Fm} - \mathbf{d} = \mathbf{FA}^{-1} \mathbf{p} - \mathbf{d} \\ \mathbf{0} &\approx \mathbf{Am} = \mathbf{I} \mathbf{p} \end{aligned} \tag{5.9}$$

Preconditioning means iteratively fitting by adjusting the \mathbf{p} variables and then finding the model by using $\mathbf{m} = \mathbf{A}^{-1}\mathbf{p}$.

5.3 OPPORTUNITIES FOR SMART DIRECTIONS

Recall the fitting goals (5.10)

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{F}\mathbf{m} - \mathbf{d} = \mathbf{F}\mathbf{A}^{-1} \mathbf{p} - \mathbf{d} \\ \mathbf{0} &\approx \mathbf{r}_m = \mathbf{A}\mathbf{m} = \mathbf{I} \mathbf{p} \end{aligned} \quad (5.10)$$

Without preconditioning we have the search direction

$$\Delta\mathbf{m}_{\text{bad}} = \begin{bmatrix} \mathbf{F}' & \mathbf{A}' \end{bmatrix} \begin{bmatrix} \mathbf{r}_d \\ \mathbf{r}_m \end{bmatrix} \quad (5.11)$$

and with preconditioning we have the search direction

$$\Delta\mathbf{p}_{\text{good}} = \begin{bmatrix} (\mathbf{F}\mathbf{A}^{-1})' & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{r}_d \\ \mathbf{r}_m \end{bmatrix} \quad (5.12)$$

The essential feature of preconditioning is not that we perform the iterative optimization in terms of the variable \mathbf{p} . The essential feature is that we use a search direction that is a gradient with respect to \mathbf{p}' not \mathbf{m}' . Using $\mathbf{A}\mathbf{m} = \mathbf{p}$ we have $\mathbf{A}\Delta\mathbf{m} = \Delta\mathbf{p}$. This enables us to define a good search direction in model space.

$$\Delta\mathbf{m}_{\text{good}} = \mathbf{A}^{-1}\Delta\mathbf{p}_{\text{good}} = \mathbf{A}^{-1}(\mathbf{A}^{-1})'\mathbf{F}'\mathbf{r}_d + \mathbf{A}^{-1}\mathbf{r}_m \quad (5.13)$$

Define the gradient by $\mathbf{g} = \mathbf{F}'\mathbf{r}_d$ and notice that $\mathbf{r}_m = \mathbf{p}$.

$$\Delta\mathbf{m}_{\text{good}} = \mathbf{A}^{-1}(\mathbf{A}^{-1})' \mathbf{g} + \mathbf{m} \quad (5.14)$$

The search direction (5.14) shows a positive-definite operator scaling the gradient. Each component of any gradient vector is independent of each other. All independently point a direction for descent. Obviously, each can be scaled by any positive number. Now we have found that we can also scale a gradient vector by a positive definite matrix and we can still expect the conjugate-direction algorithm to descend, as always, to the “exact” answer in a finite number of steps. This is because modifying the search direction with $\mathbf{A}^{-1}(\mathbf{A}^{-1})'$ is equivalent to solving a conjugate-gradient problem in \mathbf{p} .

5.4 NULL SPACE AND INTERVAL VELOCITY

A bread-and-butter problem in seismology is building the velocity as a function of depth (or vertical travel time) starting from certain measurements. The measurements are described elsewhere (BEI for example). They amount to measuring the integral of the velocity squared from the surface down to the reflector. It is known as the RMS (root-mean-square) velocity. Although good quality echos may arrive often, they rarely arrive continuously for all depths. Good information is interspersed unpredictably with poor information. Luckily we can also

estimate the data quality by the “coherency” or the “stack energy”. In summary, what we get from observations and preprocessing are two functions of travel-time depth, (1) the integrated (from the surface) squared velocity, and (2) a measure of the quality of the integrated velocity measurement. Some definitions:

\mathbf{d} is a data vector whose components range over the vertical traveltime depth τ , and whose component values contain the scaled RMS velocity squared $\tau v_{\text{RMS}}^2 / \Delta\tau$ where $\tau / \Delta\tau$ is the index on the time axis.

\mathbf{W} is a diagonal matrix along which we lay the given measure of data quality. We will use it as a weighting function.

\mathbf{C} is the matrix of causal integration, a lower triangular matrix of ones.

\mathbf{D} is the matrix of causal differentiation, namely, $\mathbf{D} = \mathbf{C}^{-1}$.

\mathbf{u} is a vector whose components range over the vertical traveltime depth τ , and whose component values contain the interval velocity squared v_{interval}^2 .

From these definitions, under the assumption of a stratified earth with horizontal reflectors (and no multiple reflections) the theoretical (squared) interval velocities enable us to define the theoretical (squared) RMS velocities by

$$\mathbf{C}\mathbf{u} = \mathbf{d} \quad (5.15)$$

With imperfect data, our data fitting goal is to minimize the residual

$$\mathbf{0} \approx \mathbf{W}[\mathbf{C}\mathbf{u} - \mathbf{d}] \quad (5.16)$$

To find the interval velocity where there is no data (where the stack power theoretically vanishes) we have the “model damping” goal to minimize the wiggleness \mathbf{p} of the squared interval velocity \mathbf{u} .

$$\mathbf{0} \approx \mathbf{D}\mathbf{u} = \mathbf{p} \quad (5.17)$$

We precondition these two goals by changing the optimization variable from interval velocity squared \mathbf{u} to its wiggleness \mathbf{p} . Substituting $\mathbf{u} = \mathbf{C}\mathbf{p}$ gives the two goals expressed as a function of wiggleness \mathbf{p} .

$$\mathbf{0} \approx \mathbf{W}[\mathbf{C}^2\mathbf{p} - \mathbf{d}] \quad (5.18)$$

$$\mathbf{0} \approx \epsilon \mathbf{p} \quad (5.19)$$

5.4.1 Balancing good data with bad

Choosing the size of ϵ chooses the stiffness of the curve that connects regions of good data. Our first test cases gave solutions that we interpreted to be too stiff at early times and too flexible at later times. This leads to two possible ways to deal with the problem. One way modifies the model shaping and the other modifies the data fitting. The program below weakens the data fitting weight with time. This has the same effect as stiffening the model shaping with time.




Figure 5.1: Raw CMP gather (left), Semblance scan (middle), and semblance value used for weighting function (right). (Clapp) `prc/bob clapp`




Figure 5.2: Observed RMS velocity and that predicted by a stiff model with $\epsilon = 4$. (Clapp) `prc/bob stiff`




Figure 5.3: Observed RMS velocity and that predicted by a flexible model with $\epsilon = .25$ (Clapp) `prc/bob flex`

5.4.2 Lateral variations

The analysis above appears one dimensional in depth. Conventional interval velocity estimation builds a velocity-depth model independently at each lateral location. Here we have a logical path for combining measurements from various lateral locations. We can change the regularization to something like $\mathbf{0} \approx \nabla \mathbf{u}$. Instead of merely minimizing the vertical gradient of velocity we minimize its spatial gradient. Luckily we have preconditioning and the helix to speed the solution.


5.4.3 Blocky models

Sometimes we seek a velocity model that increases smoothly with depth through our scattered measurements of good-quality RMS velocities. Other times, we seek a blocky model. (Where seismic data is poor, a well log could tell us whether to choose smooth or blocky.) Here we see an estimation method that can choose the blocky alternative, or some combination of smooth and blocky.

Consider the five layer model in Figure 5.4. Each layer has unit travelttime thickness (so integration is simply summation). Let the squared interval velocities be (a, b, c, d, e) with strong reliable reflections at the base of layer c and layer e , and weak, incoherent, “bad” reflections at bases of (a, b, d) . Thus we measure V_c^2 the RMS velocity squared of the top three layers and V_e^2 that for all five layers. Since we have no reflection from at the base of the fourth layer, the velocity in the fourth layer is not measured but a matter for choice. In a smooth linear fit we would want $d = (c + e)/2$. In a blocky fit we would want $d = e$.

Our screen for good reflections looks like $(0, 0, 1, 0, 1)$ and our screen for bad ones looks like the complement $(1, 1, 0, 1, 0)$. We put these screens on the diagonals of diagonal matrices

Figure 5.4: A layered earth model. The layer interfaces cause reflections. Each layer has a constant velocity in its interior. prc/. rosales



G and **B**. Our fitting goals are:

$$3V_c^2 \approx a + b + c \quad (5.20)$$

$$5V_e^2 \approx a + b + c + d + e \quad (5.21)$$

$$u_0 \approx a \quad (5.22)$$

$$0 \approx -a + b \quad (5.23)$$

$$0 \approx -b + c \quad (5.24)$$


$$0 \approx -c + d \quad (5.25)$$

$$0 \approx -d + e \quad (5.26)$$

For the blocky solution, we do not want the fitting goal (5.25). Further explanations await completion of examples.

5.5 INVERSE LINEAR INTERPOLATION

Figure 5.5: The input data are irregularly sampled. prc/sep94 data



The first example is a simple synthetic test for 1-D inverse interpolation. The input data were randomly subsampled (with decreasing density) from a sinusoid (Figure 5.5). The forward operator **L** in this case is linear interpolation. We seek a regularly sampled model that could predict the data with a forward linear interpolation. Sparse irregular distribution of the input data makes the regularization enforcement a necessity. I applied convolution with the simple $(1, -1)$ difference filter as the operator **D** that forces model continuity (the first-order spline). An appropriate preconditioner **S** in this case is recursive causal integration.





Figure 5.6: Convergence history of inverse linear interpolation. Left: regularization, right: preconditioning. The regularization operator \mathbf{A} is the derivative operator (convolution with $(1, -1)$). The preconditioning operator \mathbf{S} is causal integration. `prc/sep94 conv`

As expected, preconditioning provides a much faster rate of convergence. Since iteration to the exact solution is never achieved in large-scale problems, the results of iterative optimization may turn out quite differently. Bill Harlan points out that the two goals in (5.8) conflict with each other: the first one enforces “details” in the model, while the second one tries to smooth them out. Typically, regularized optimization creates a complicated model at early iterations. At first, the data fitting goal (5.8) plays a more important role. Later, the regularization goal (5.8) comes into play and simplifies (smooths) the model as much as needed. Preconditioning acts differently. The very first iterations create a simplified (smooth) model. Later, the data fitting goal adds more details into the model. If we stop the iterative process early, we end up with an insufficiently complex model, not in an insufficiently simplified one. Figure 5.6 provides a clear illustration of Harlan’s observation.

Figure 5.7 measures the rate of convergence by the model residual, which is a distance from the current model to the final solution. It shows that preconditioning saves many iterations. Since the cost of each iteration for each method is roughly equal, the efficiency of preconditioning is evident.

Figure 5.7: Convergence of the iterative optimization, measured in terms of the model residual. The “p” points stand for preconditioning; the “r” points, regularization.

```
prc/sep94 schwab1
```



5.6 EMPTY BINS AND PRECONDITIONING

There are at least three ways to fill empty bins. Two require a roughening operator \mathbf{A} while the third requires a smoothing operator which (for comparison purposes) we denote \mathbf{A}^{-1} . The three methods are generally equivalent though they differ in important details.

The original way in Chapter 3 is to restore missing data by ensuring that the restored data, after specified filtering, has minimum energy, say $\mathbf{A}\mathbf{m} \approx \mathbf{0}$. Introduce the selection mask operator \mathbf{K} , a diagonal matrix with ones on the known data and zeros elsewhere (on the missing data). Thus $\mathbf{0} \approx \mathbf{A}(\mathbf{I} - \mathbf{K} + \mathbf{K})\mathbf{m}$ or

$$\mathbf{0} \approx \mathbf{A}(\mathbf{I} - \mathbf{K})\mathbf{m} + \mathbf{A}\mathbf{m}_k, \quad (5.27)$$

where we define \mathbf{m}_k to be the data with missing values set to zero by $\mathbf{m}_k = \mathbf{K}\mathbf{m}$.

A second way to find missing data is with the set of goals

$$\begin{aligned} \mathbf{0} &\approx \mathbf{K}\mathbf{m} - \mathbf{m}_k \\ \mathbf{0} &\approx \epsilon\mathbf{A}\mathbf{m} \end{aligned} \quad (5.28)$$

and take the limit as the scalar $\epsilon \rightarrow 0$. At that limit, we should have the same result as equation (5.27).

There is an important philosophical difference between the first method and the second. The first method strictly honors the known data. The second method acknowledges that when data misfits the regularization theory, it might be the fault of the data so the data need not be strictly honored. Just what balance is proper falls to the numerical choice of ϵ , a nontrivial topic.

A third way to find missing data is to precondition equation (5.28), namely, try the substitution $\mathbf{m} = \mathbf{A}^{-1}\mathbf{p}$.

$$\begin{aligned} \mathbf{0} &\approx \mathbf{K}\mathbf{A}^{-1}\mathbf{p} - \mathbf{m}_k \\ \mathbf{0} &\approx \epsilon\mathbf{p} \end{aligned} \tag{5.29}$$

There is no simple way of knowing beforehand what is the best value of ϵ . Practitioners like to see solutions for various values of ϵ . Of course that can cost a lot of computational effort. Practical exploratory data analysis is more pragmatic. Without a simple clear theoretical basis, analysts generally begin from $\mathbf{p} = \mathbf{0}$ and abandon the fitting goal $\epsilon\mathbf{I}\mathbf{p} \approx \mathbf{0}$. Implicitly, they take $\epsilon = 0$. Then they examine the solution as a function of iteration, imagining that the solution at larger iterations corresponds to smaller ϵ . There is an eigenvector analysis indicating some kind of basis for this approach, but I believe there is no firm guidance.

Before we look at coding details for the three methods of filling the empty bins, we'll compare results of trying all three methods. For the roughening operator \mathbf{A} , we'll take the helix derivative \mathbf{H} . This is logically equivalent to roughening with the gradient ∇ because the (negative) laplacian operator is $\nabla'\nabla = \mathbf{H}'\mathbf{H}$.

5.6.1 SEABEAM: Filling the empty bins with a laplacian

Figure 5.8 shows a day's worth of data¹ collected at sea by SeaBeam, an apparatus for measuring water depth both directly under a ship, and somewhat off to the sides of the ship's track. The data is measurements of depth $h(x, y)$ at miscellaneous locations in the (x, y) -plane. The locations are scattered about, according to various aspects of the ship's navigation and the geometry of the SeaBeam sonic antenna. Figure 5.8 was made by binning with `bin2()` on page 12 and equation (1.15). The spatial spectra of the noise in the data could be estimated where tracks cross over themselves. This might be worth while, but we do not pursue it now.

Here we focus on the empty mesh locations where no data is recorded (displayed with the value of the mean depth \bar{h}). These empty bins were filled with module `mis2` on page 133. Results are in Figure 5.9. In Figure 5.9 the left column results from 20 iterations while the right column results from 100 iterations.


The top row in Figure 5.9 shows that more iterations spreads information further into the region of missing data.

It turned out that the original method strictly honoring known data gave results so similar to the second method (regularizing) that the plots could not be visually distinguished. The middle row in Figure 5.9 therefore shows the difference in the result of the two methods.

¹ I'd like to thank Alistair Harding for this interesting data set named April 18.

Figure 5.8: Depth of the ocean under ship tracks. Empty bins are displayed with an average depth \bar{h} .

prc/seab seabin



April 18 Binned

We see an outline of the transition between known and unknown regions. Obviously, the missing data is pulling known data towards zero.

The bottom row in Figure 5.9 shows that preconditioning spreads information to great distances much quicker but early iterations make little effort to honor the data. (Even though these results are for $\epsilon = 0$.) Later iterations make little change at long distance but begin to restore sharp details on the small features of the known topography.

What if we can only afford 100 iterations? Perhaps we should first do 50 iterations with preconditioning to develop the remote part of the solution and then do 50 iterations by one of the other two methods to be sure we attended to the details near the known data. A more unified approach (not yet tried, as far as I know) would be to unify the techniques. The conjugate direction method searches two directions, the gradient and the previous step. We could add a third direction, the smart direction of equation (5.14). Instead of having a 2×2 matrix solution like equation (2.70) for two distances, we would need to solve a 3×3 matrix for three.

Figure 5.9 has a few artifacts connected with the use of the helix derivative. Examine equation (11) to notice the shape of the helix derivative. In principle, it is infinitely long in the horizontal axis in both equation (11) and Figure 5.9. In practice, it is truncated. The truncation is visible as bands along the sides of Figure 5.9.

As a practical matter, no one would use the first two bin filling methods with helix derivative for the roughener because it is theoretically equivalent to the gradient operator ∇ which has many fewer coefficients. Later, in Chapter 6 we'll find a much smarter roughening operator \mathbf{A} called the Prediction Error Filter (PEF) which gives better results.

5.6.2 Three codes for inverse masking

The selection (or masking) operator \mathbf{K} is implemented in `mask()` on page 132.




Figure 5.9: The ocean bottom restoring missing data with a helix derivative.

```
prc/seabeam prcfill
```

```

                                filt/lib/mask.c
45     sf_adjnull (adj , add , nx , ny , x , y );
46
47     for (ix=0; ix < nx; ix++) {
48         if (m[ix]) {
49             if (adj) x[ix] += y[ix];
50             else    y[ix] += x[ix];

```

All the results shown in Figure 5.9 were created with the module `mis2` on the facing page. Code locations with `style=0,1,2` correspond to the fitting goals (5.27), (5.28), (5.29).

5.7 THEORY OF UNDERDETERMINED LEAST-SQUARES

Construct theoretical data with

$$\mathbf{d} = \mathbf{F}\mathbf{m} \quad (5.30)$$

Assume there are fewer data points than model points and that the matrix $\mathbf{F}\mathbf{F}'$ is invertible. From the theoretical data we estimate a model \mathbf{m}_0 with

$$\mathbf{m}_0 = \mathbf{F}'(\mathbf{F}\mathbf{F}')^{-1}\mathbf{d} \quad (5.31)$$

To verify the validity of the estimate, insert the estimate (5.31) into the data modeling equation (5.30) and notice that the estimate \mathbf{m}_0 predicts the correct data. Notice that equation (5.31) is not the same as equation (2.30) which we derived much earlier. What's the difference? The central issue is which matrix of $\mathbf{F}\mathbf{F}'$ and $\mathbf{F}'\mathbf{F}$ actually has an inverse. If \mathbf{F} is a rectangular matrix, then it is certain that one of the two is not invertible. (There are plenty of real cases where neither matrix is invertible. That's one reason we use iterative solvers.) Here we are dealing with the case with more model points than data points.

Now we will show that of all possible models \mathbf{m} that predict the correct data, \mathbf{m}_0 has the least energy. (I'd like to thank Sergey Fomel for this clear and simple proof that does *not* use Lagrange multipliers.) First split (5.31) into an intermediate result \mathbf{d}_0 and final result:

$$\mathbf{d}_0 = (\mathbf{F}\mathbf{F}')^{-1}\mathbf{d} \quad (5.32)$$

$$\mathbf{m}_0 = \mathbf{F}'\mathbf{d}_0 \quad (5.33)$$

Consider another model (\mathbf{x} not equal to zero)

$$\mathbf{m} = \mathbf{m}_0 + \mathbf{x} \quad (5.34)$$

which fits the theoretical data $\mathbf{d} = \mathbf{F}(\mathbf{m}_0 + \mathbf{x})$. Since $\mathbf{d} = \mathbf{F}\mathbf{m}_0$, we see that \mathbf{x} is a null space vector.

$$\mathbf{F}\mathbf{x} = \mathbf{0} \quad (5.35)$$

```

user/gee/mis2.c
25 void mis2(int niter      /* number of iterations */,
26           int nx        /* model size */,
27           float *xx      /* model */,
28           sf_filter aa   /* helix filter */,
29           const bool *known /* mask for known data */,
30           float eps      /* regularization parameter */,
31           bool doprec    /* to apply preconditioning */)
32 /*< interpolate >*/
33 {
34     int ix;
35     float *dd;
36
37     if (doprec) { /* preconditioned */
38         sf_mask_init(known);
39         polydiv_init(nx, aa);
40         sf_solver_prec(sf_mask_lob, sf_cgstep, polydiv_lob,
41                       nx, nx, nx, xx, xx, niter, eps, "end");
42         polydiv_close();
43     } else { /* regularized */
44         dd = sf_floatalloc(nx);
45         for (ix=0; ix < nx; ix++) {
46             dd[ix]=0.;
47         }
48
49         sf_helicon_init(aa);
50         sf_solver (sf_helicon_lob, sf_cgstep, nx, nx, xx, dd, niter,
51                  "known", known, "x0", xx, "end");
52         free(dd);
53     }
54     sf_cgstep_close();
55 }

```

First we see that \mathbf{m}_0 is orthogonal to \mathbf{x} because

$$\mathbf{m}_0' \mathbf{x} = (\mathbf{F}' \mathbf{d}_0)' \mathbf{x} = \mathbf{d}_0' \mathbf{F} \mathbf{x} = \mathbf{d}_0' \mathbf{0} = 0 \quad (5.36)$$

Therefore,

$$\mathbf{m}' \mathbf{m} = \mathbf{m}_0' \mathbf{m}_0 + \mathbf{x}' \mathbf{x} + 2 \mathbf{x}' \mathbf{m}_0 = \mathbf{m}_0' \mathbf{m}_0 + \mathbf{x}' \mathbf{x} \geq \mathbf{m}_0' \mathbf{m}_0 \quad (5.37)$$

so adding null space to \mathbf{m}_0 can only increase its energy. In summary, the solution $\mathbf{m}_0 = \mathbf{F}'(\mathbf{F}\mathbf{F}')^{-1}\mathbf{d}$ has less energy than any other model that satisfies the data.

Not only does the theoretical solution $\mathbf{m}_0 = \mathbf{F}'(\mathbf{F}\mathbf{F}')^{-1}\mathbf{d}$ have minimum energy, but the result of iterative descent will too, provided that we begin iterations from $\mathbf{m}_0 = 0$ or any \mathbf{m}_0 with no null-space component. In (5.36) we see that the orthogonality $\mathbf{m}_0' \mathbf{x} = 0$ does not arise because \mathbf{d}_0 has any particular value. It arises because \mathbf{m}_0 is of the form $\mathbf{F}' \mathbf{d}_0$. Gradient methods contribute $\Delta \mathbf{m} = \mathbf{F}' \mathbf{r}$ which is of the required form.

5.8 SCALING THE ADJOINT

First I remind you of a rarely used little bit of mathematical notation. Given a vector \mathbf{m} with components (m_1, m_2, m_3) , the notation **diag** \mathbf{m} means

$$\mathbf{diag} \mathbf{m} = \begin{bmatrix} m_1 & 0 & 0 \\ 0 & m_2 & 0 \\ 0 & 0 & m_3 \end{bmatrix} \quad (5.38)$$

Given the usual linearized fitting goal between data space and model space, $\mathbf{d} \approx \mathbf{F}\mathbf{m}$, the simplest image of the model space results from application of the adjoint operator $\hat{\mathbf{m}} = \mathbf{F}'\mathbf{d}$. Unless \mathbf{F} has no physical units, however, the physical units of $\hat{\mathbf{m}}$ do not match those of \mathbf{m} , so we need a scaling factor. The theoretical solution $\mathbf{m}_{\text{theor}} = (\mathbf{F}'\mathbf{F})^{-1}\mathbf{F}'\mathbf{d}$ tells us that the scaling units should be those of $(\mathbf{F}'\mathbf{F})^{-1}$. We are going to approximate $(\mathbf{F}'\mathbf{F})^{-1}$ by a diagonal matrix \mathbf{W}^2 with the correct units so $\hat{\mathbf{m}} = \mathbf{W}^2 \mathbf{F}'\mathbf{d}$.

What we use for \mathbf{W} will be a guess, simply a guess. If it works better than nothing, we'll be happy, and if it doesn't we'll forget about it. Experience shows it is a good idea to try. Common sense tells us to insist that all elements of \mathbf{W}^2 are positive. \mathbf{W}^2 is a square matrix of size of model space. From any vector $\tilde{\mathbf{m}}$ in model space with all positive components, we could guess that \mathbf{W}^2 be **diag** $\tilde{\mathbf{m}}$ to any power. To get the right physical dimensions we choose $\tilde{\mathbf{m}} = \mathbf{1}$, a vector of all ones and choose

$$\mathbf{W}^2 = \frac{1}{\mathbf{diag} \mathbf{F}'\mathbf{F}\mathbf{1}} \quad (5.39)$$

A problem with the choice (5.39) is that some components might be zero or negative. Well, we can take the square root of the squares of components and/or smooth the result.

To go beyond the scaled adjoint we can use \mathbf{W} as a preconditioner. To use \mathbf{W} as a preconditioner we define implicitly a new set of variables \mathbf{p} by the substitution $\mathbf{m} = \mathbf{W}\mathbf{p}$. Then $\mathbf{d} \approx \mathbf{F}\mathbf{m} = \mathbf{F}\mathbf{W}\mathbf{p}$. To find \mathbf{p} instead of \mathbf{m} , we iterate with the operator $\mathbf{F}\mathbf{W}$ instead of with \mathbf{F} . As usual, the first step of the iteration is to use the adjoint of $\mathbf{d} \approx \mathbf{F}\mathbf{W}\mathbf{p}$ to form

the image $\hat{\mathbf{p}} = (\mathbf{FW})'\mathbf{d}$. At the end of the iterations, we convert from \mathbf{p} back to \mathbf{m} with $\mathbf{m} = \mathbf{W}\mathbf{p}$. The result after the first iteration $\hat{\mathbf{m}} = \mathbf{W}\hat{\mathbf{p}} = \mathbf{W}(\mathbf{FW})'\mathbf{d} = \mathbf{W}^2\mathbf{F}'\mathbf{d}$ turns out to be the same as scaling.

By (5.39), \mathbf{W} has physical units inverse to \mathbf{F} . Thus the transformation \mathbf{FW} has no units so the \mathbf{p} variables have physical units of data space. Experimentalists might enjoy seeing the solution \mathbf{p} with its data units more than viewing the solution \mathbf{m} with its more theoretical model units.

The theoretical solution for underdetermined systems $\mathbf{m} = \mathbf{F}'(\mathbf{FF}')^{-1}\mathbf{d}$ suggests an alternate approach using instead $\hat{\mathbf{m}} = \mathbf{F}'\mathbf{W}_d^2\mathbf{d}$. This diagonal weighting matrix \mathbf{W}_d^2 must be drawn from vectors in data space. Again I chose a vector of all 1's getting the weight

$$\mathbf{W}_d^2 = \frac{1}{\text{diag } \mathbf{FF}'\mathbf{1}} \quad (5.40)$$

My choice of a vector of 1's is quite arbitrary. I might as well have chosen a vector of random numbers. Bill Symes, who suggested this approach to me, suggests using an observed data vector \mathbf{d} for the data space weight, and $\mathbf{F}'\mathbf{d}$ for the model space weight. This requires an additional step, dividing out the units of the data \mathbf{d} .

Experience tells me that a broader methodology than all above is needed. Appropriate scaling is required in both data space and model space. We need two other weights \mathbf{W}_m and \mathbf{W}_d where $\hat{\mathbf{m}} = \mathbf{W}_m\mathbf{F}'\mathbf{W}_d\mathbf{d}$.

I have a useful practical example (stacking in $v(z)$ media) in another of my electronic books (BEI), where I found both \mathbf{W}_m and \mathbf{W}_d by iterative guessing. First assume $\mathbf{W}_d = \mathbf{I}$ and estimate \mathbf{W}_m as above. Then assume you have the correct \mathbf{W}_m and estimate \mathbf{W}_d as above. Iterate. (Perhaps some theorist can find a noniterative solution.) I believe this iterative procedure leads us to the best diagonal pre- and post- multipliers for any operator \mathbf{F} . By this I mean that the modified operator $(\mathbf{W}_d\mathbf{FW}_m)$ is as close to being unitary as we will be able to obtain with diagonal transformation. Unitary means it is energy conserving and that the inverse is simply the conjugate transpose.

What good is it that $(\mathbf{W}_d\mathbf{FW}_m)'(\mathbf{W}_d\mathbf{FW}_m) \approx \mathbf{I}$? It gives us the most rapid convergence of least squares problems of the form

$$\mathbf{0} \approx \mathbf{W}_d(\mathbf{F}\mathbf{m} - \mathbf{d}) = \mathbf{W}_d(\mathbf{FW}_m\mathbf{p} - \mathbf{d}) \quad (5.41)$$

Thus it defines for us the best diagonal transform to a preconditioning variable $\mathbf{p} = \mathbf{W}_m^{-1}\mathbf{m}$ to use during iteration, and suggests to us what residual weighting function we need to use if rapid convergence is a high priority. Suppose we are not satisfied with \mathbf{W}_d being the weighting function for residuals. Equation (5.41) could still help us speed iteration. Instead of beginning iteration with $\mathbf{p} = \mathbf{0}$, we could begin from the solution \mathbf{p} to the regression (5.41).

The PhD thesis of James Rickett experiments extensively with data space and model space weighting functions in the context of seismic velocity estimation.

5.9 A FORMAL DEFINITION FOR ADJOINTS

In mathematics, adjoints are defined a little differently than we have defined them here (as matrix transposes).² The mathematician begins by telling us that we cannot simply form any dot product we want. We are not allowed to take the dot product of any two vectors in model space $\mathbf{m}_1 \cdot \mathbf{m}_2$ or data space $\mathbf{d}_1 \cdot \mathbf{d}_2$. Instead, we must first transform them to a preferred coordinate system. Say $\tilde{\mathbf{m}}_1 = \mathbf{M}\mathbf{m}_1$ and $\tilde{\mathbf{d}}_1 = \mathbf{D}\mathbf{d}_1$, etc for other vectors. We complain we do not know \mathbf{M} and \mathbf{D} . They reply that we do not really need to know them but we do need to have the inverses (aack!) of $\mathbf{M}'\mathbf{M}$ and $\mathbf{D}'\mathbf{D}$. A pre-existing common notation is $\sigma_m^{-2} = \mathbf{M}'\mathbf{M}$ and $\sigma_d^{-2} = \mathbf{D}'\mathbf{D}$. Now the mathematician buries the mysterious new positive-definite matrix inverses in the definition of dot product $\langle \mathbf{m}_1, \mathbf{m}_2 \rangle = \mathbf{m}_1' \mathbf{M}' \mathbf{M} \mathbf{m}_2 = \mathbf{m}_1' \sigma_m^{-2} \mathbf{m}_2$ and likewise with $\langle \mathbf{d}_1, \mathbf{d}_2 \rangle$. This suggests a total reorganization of our programs. Instead of computing $(\mathbf{m}_1' \mathbf{M}')(\mathbf{M} \mathbf{m}_2)$ we could compute $\mathbf{m}_1'(\sigma_m^{-2} \mathbf{m}_2)$. Indeed, this is the “conventional” approach. This definition of dot product would be buried in the solver code. The other thing that would change would be the search direction $\Delta \mathbf{m}$. Instead of being the gradient as we have defined it $\Delta \mathbf{m} = \mathbf{L}' \mathbf{r}$, it would be $\Delta \mathbf{m} = \sigma_m^{-2} \mathbf{L}' \sigma_d^{-2} \mathbf{r}$. A mathematician would *define* the adjoint of \mathbf{L} to be $\sigma_m^{-2} \mathbf{L}' \sigma_d^{-2}$. (Here \mathbf{L}' remains matrix transpose.) You might notice this approach nicely incorporates both residual weighting and preconditioning while yet evading the issue of where we get the matrices σ_m^2 and σ_d^2 or how we invert them. Fortunately, upcoming chapter 6 suggests how, in image estimation problems, to obtain sensible estimates of the elusive operators \mathbf{M} and \mathbf{D} . Paranthetically, modeling calculations in physics and engineering often use similar mathematics in which the role of $\mathbf{M}'\mathbf{M}$ is not so mysterious. Kinetic energy is mass times velocity squared. Mass can play the role of $\mathbf{M}'\mathbf{M}$.

So, should we continue to use $(\mathbf{m}_1' \mathbf{M}')(\mathbf{M} \mathbf{m}_2)$ or should we take the conventional route and go with $\mathbf{m}_1'(\sigma_m^{-2} \mathbf{m}_2)$? One day while benchmarking a wide variety of computers I was shocked to see some widely differing numerical results. Now I know why. Consider adding 10^7 identical positive floating point numbers, say 1.0's, in an arithmetic with precision of 10^{-6} . After you have added in the first 10^6 numbers, the rest will all truncate in the roundoff and your sum will be wrong by a factor of ten. If the numbers were added in pairs, and then the pairs added, etc, there would be no difficulty. Precision is scary stuff!

It is my understanding and belief that there is nothing wrong with the approach of this book, in fact, it seems to have some definite advantages. While the conventional approach requires one to compute the adjoint correctly, we do not. The method of this book (which I believe is properly called conjugate directions) has a robustness that, I'm told, has been superior in some important geophysical applications. The conventional approach seems to get in trouble when transpose operators are computed with insufficient precision.

² I would like to thank Albert Tarantola for suggesting this topic.

Chapter 6

Multidimensional autoregression

The many applications of least squares to the one-dimensional convolution operator constitute the subject known as “time-series analysis.” The *autoregression* filter, also known as the prediction-error filter (PEF), gathers statistics for us, not the autocorrelation or the spectrum directly but it gathers them indirectly as the inverse of the amplitude spectrum of its input. The PEF plays the role of the so-called “inverse-covariance matrix” in statistical estimation theory. Given the PEF, we use it to find missing portions of signals.

6.0.1 Time domain versus frequency domain

In the simplest applications, solutions can be most easily found in the frequency domain. When complications arise, it is better to use the time domain, to directly apply the convolution operator and the method of least squares.

A first complicating factor in the frequency domain is a required boundary in the time domain, such as that between past and future, or requirements that a filter be nonzero in a stated time interval. Another factor that attracts us to the time domain rather than the frequency domain is weighting functions.

Weighting functions are appropriate whenever a signal or image amplitude varies from place to place. Much of the literature on time-series analysis applies to the limited case of uniform weighting functions. Such time series are said to be “stationary.” This means that their statistical properties do not change in time. In real life, particularly in the analysis of echos, signals are never stationary in time and space. A stationarity assumption is a reasonable starting assumption, but we should know how to go beyond it so we can take advantage of the many opportunities that do arise. In order of increasing difficulty in the frequency domain are the following complications:

1. A time boundary such as between past and future.
2. More time boundaries such as delimiting a filter.
3. More time boundaries such as erratic locations of missing data.
4. Nonstationary signal, i.e., time-variable weighting.

5. Time-axis stretching such as normal moveout.

We will not have difficulty with any of these complications here, because we will stay in the time domain and set up and solve optimization problems by use of the conjugate-direction method. Thus we will be able to cope with great complexity in goal formulation and get the right answer without approximations. By contrast, analytic or partly analytic methods can be more economical, but they generally solve somewhat different problems than those given to us by nature.

6.1 SOURCE WAVEFORM, MULTIPLE REFLECTIONS

Here we devise a simple mathematical model for deep water bottom multiple reflections.¹ There are two unknown waveforms, the source waveform $S(\omega)$ and the ocean-floor reflection $F(\omega)$. The water-bottom primary reflection $P(\omega)$ is the convolution of the source waveform with the water-bottom response; so $P(\omega) = S(\omega)F(\omega)$. The first multiple reflection $M(\omega)$ sees the same source waveform, the ocean floor, a minus one for the free surface, and the ocean floor again. Thus the observations $P(\omega)$ and $M(\omega)$ as functions of the physical parameters are

$$P(\omega) = S(\omega)F(\omega) \quad (6.1)$$

$$M(\omega) = -S(\omega)F(\omega)^2 \quad (6.2)$$

Algebraically the solutions of equations (6.1) and (6.2) are

$$F(\omega) = -M(\omega)/P(\omega) \quad (6.3)$$

$$S(\omega) = -P(\omega)^2/M(\omega) \quad (6.4)$$

These solutions can be computed in the Fourier domain by simple division. The difficulty is that the divisors in equations (6.3) and (6.4) can be zero, or small. This difficulty can be attacked by use of a positive number ϵ to stabilize it. For example, multiply equation (6.3) on top and bottom by $P'(\omega)$ and add $\epsilon > 0$ to the denominator. This gives

$$F(\omega) = -\frac{M(\omega)P'(\omega)}{P(\omega)P'(\omega) + \epsilon} \quad (6.5)$$

where $P'(\omega)$ is the complex conjugate of $P(\omega)$. Although the ϵ stabilization seems nice, it apparently produces a nonphysical model. For ϵ large or small, the time-domain response could turn out to be of much greater duration than is physically reasonable. This should not happen with perfect data, but in real life, data always has a limited spectral band of good quality.

Functions that are rough in the frequency domain will be long in the time domain. This suggests making a short function in the time domain by local smoothing in the frequency domain. Let the notation $\langle \cdots \rangle$ denote smoothing by local averaging. Thus, to specify filters whose time duration is not unreasonably long, we can revise equation (6.5) to

$$F(\omega) = -\frac{\langle M(\omega)P'(\omega) \rangle}{\langle P(\omega)P'(\omega) \rangle} \quad (6.6)$$

¹ For this short course I am omitting here many interesting examples of multiple reflections shown in my 1992 book, PVI.

where instead of deciding a size for ϵ we need to decide how much smoothing. I find that smoothing has a simpler physical interpretation than choosing ϵ . The goal of finding the filters $F(\omega)$ and $S(\omega)$ is to best model the multiple reflections so that they can be subtracted from the data, and thus enable us to see what primary reflections have been hidden by the multiples.

These frequency-duration difficulties do not arise in a time-domain formulation. Unlike in the frequency domain, in the time domain it is easy and natural to limit the duration and location of the nonzero time range of $F(\omega)$ and $S(\omega)$. First express (6.3) as

$$0 = P(\omega)F(\omega) + M(\omega) \quad (6.7)$$

To imagine equation (6.7) as a fitting goal in the time domain, instead of scalar functions of ω , think of vectors with components as a function of time. Thus \mathbf{f} is a column vector containing the unknown sea-floor filter, \mathbf{m} contains the “multiple” portion of a seismogram, and \mathbf{P} is a matrix of down-shifted columns, each column being the “primary”.

$$\mathbf{0} \approx \mathbf{r} = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \end{bmatrix} = \begin{bmatrix} p_1 & 0 & 0 \\ p_2 & p_1 & 0 \\ p_3 & p_2 & p_1 \\ p_4 & p_3 & p_2 \\ p_5 & p_4 & p_3 \\ p_6 & p_5 & p_4 \\ 0 & p_6 & p_5 \\ 0 & 0 & p_6 \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} + \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \\ m_6 \\ m_7 \\ m_8 \end{bmatrix} \quad (6.8)$$

6.2 TIME-SERIES AUTOREGRESSION

Given y_t and y_{t-1} , you might like to predict y_{t+1} . The prediction could be a scaled sum or difference of y_t and y_{t-1} . This is called “autoregression” because a signal is regressed on itself. To find the scale factors you would optimize the fitting goal below, for the prediction filter (f_1, f_2) :

$$\mathbf{0} \approx \mathbf{r} = \begin{bmatrix} y_1 & y_0 \\ y_2 & y_1 \\ y_3 & y_2 \\ y_4 & y_3 \\ y_5 & y_4 \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \end{bmatrix} - \begin{bmatrix} y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} \quad (6.9)$$

(In practice, of course the system of equations would be much taller, and perhaps somewhat wider.) A typical row in the matrix (6.9) says that $y_{t+1} \approx y_t f_1 + y_{t-1} f_2$ hence the description of f as a “prediction” filter. The error in the prediction is simply the residual. Define the residual to have opposite polarity and merge the column vector into the matrix, so you get

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \approx \mathbf{r} = \begin{bmatrix} y_2 & y_1 & y_0 \\ y_3 & y_2 & y_1 \\ y_4 & y_3 & y_2 \\ y_5 & y_4 & y_3 \\ y_6 & y_5 & y_4 \end{bmatrix} \begin{bmatrix} 1 \\ -f_1 \\ -f_2 \end{bmatrix} \quad (6.10)$$

which is a standard form for autoregressions and prediction error.

Multiple reflections are predictable. It is the unpredictable part of a signal, the prediction residual, that contains the primary information. The output of the filter $(1, -f_1, -f_2) = (a_0, a_1, a_2)$ is the unpredictable part of the input. This filter is a simple example of a “prediction-error” (PE) filter. It is one member of a family of filters called “error filters.”

The error-filter family are filters with one coefficient constrained to be unity and various other coefficients constrained to be zero. Otherwise, the filter coefficients are chosen to have minimum power output. Names for various error filters follow:

$(1, a_1, a_2, a_3, \dots, a_n)$	prediction-error (PE) filter
$(1, 0, 0, a_3, a_4, \dots, a_n)$	gapped PE filter with a gap
$(a_{-m}, \dots, a_{-2}, a_{-1}, 1, a_1, a_2, a_3, \dots, a_n)$	interpolation-error (IE) filter

We introduce a free-mask matrix \mathbf{K} which “passes” the freely variable coefficients in the filter and “rejects” the constrained coefficients (which in this first example is merely the first coefficient $a_0 = 1$).

$$\mathbf{K} = \begin{bmatrix} 0 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{bmatrix} \quad (6.11)$$

To compute a simple prediction error filter $\mathbf{a} = (1, a_1, a_2)$ with the CD method, we write (6.9) or (6.10) as

$$\mathbf{0} \approx \mathbf{r} = \begin{bmatrix} y_2 & y_1 & y_0 \\ y_3 & y_2 & y_1 \\ y_4 & y_3 & y_2 \\ y_5 & y_4 & y_3 \\ y_6 & y_5 & y_4 \end{bmatrix} \begin{bmatrix} 0 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \end{bmatrix} + \begin{bmatrix} y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} \quad (6.12)$$

Let us move from this specific fitting goal to the general case. (Notice the similarity of the free-mask matrix \mathbf{K} in this filter estimation problem with the free-mask matrix \mathbf{J} in missing data goal (3.3).) The fitting goal is,

$$\mathbf{0} \approx \mathbf{Y}\mathbf{a} \quad (6.13)$$

$$\mathbf{0} \approx \mathbf{Y}(\mathbf{I} - \mathbf{K} + \mathbf{K})\mathbf{a} \quad (6.14)$$

$$\mathbf{0} \approx \mathbf{Y}\mathbf{K}\mathbf{a} + \mathbf{Y}(\mathbf{I} - \mathbf{K})\mathbf{a} \quad (6.15)$$

$$\mathbf{0} \approx \mathbf{Y}\mathbf{K}\mathbf{a} + \mathbf{Y}\mathbf{a}_0 \quad (6.16)$$

$$\mathbf{0} \approx \mathbf{Y}\mathbf{K}\mathbf{a} + \mathbf{y} \quad (6.17)$$

$$\mathbf{0} \approx \mathbf{r} = \mathbf{Y}\mathbf{K}\mathbf{a} + \mathbf{r}_0 \quad (6.18)$$

which means we initialize the residual with $\mathbf{r}_0 = \mathbf{y}$. and then iterate with

$$\Delta\mathbf{a} \leftarrow \mathbf{K}'\mathbf{Y}'\mathbf{r} \quad (6.19)$$

$$\Delta\mathbf{r} \leftarrow \mathbf{Y}\mathbf{K}\Delta\mathbf{a} \quad (6.20)$$

6.3 PREDICTION-ERROR FILTER OUTPUT IS WHITE

The relationship between spectrum and PEF


Knowledge of an autocorrelation function is equivalent to knowledge of a spectrum. The two are simply related by Fourier transform. A spectrum or an autocorrelation function encapsulates an important characteristic of a signal or an image. Generally the spectrum changes slowly from place to place although it could change rapidly. Of all the assumptions we could make to fill empty bins, one that people usually find easiest to agree with is that the spectrum should be the same in the empty-bin regions as where bins are filled. In practice we deal with neither the spectrum nor its autocorrelation but with a third object. This third object is the Prediction Error Filter (PEF), the filter in equation (6.10).

Take equation (6.10) for \mathbf{r} and multiply it by the adjoint \mathbf{r}' getting a quadratic form in the PEF coefficients. Minimizing this quadratic form determines the PEF. This quadratic form depends only on the autocorrelation of the original data y_t , not on the data y_t itself. Clearly the PEF is unchanged if the data has its polarity reversed or its time axis reversed. Indeed, we'll see here that knowledge of the PEF is equivalent to knowledge of the autocorrelation or the spectrum.

Undoing convolution in nature

Prediction-error filtering is also called “deconvolution”. This word goes back to very basic models and concepts. In this model one envisions a random white-spectrum excitation function \mathbf{x} existing in nature, and this excitation function is somehow filtered by unknown natural processes, with a filter operator \mathbf{B} producing an *output* \mathbf{y} in nature that becomes the *input* \mathbf{y} to our computer programs. This is sketched in Figure 6.1. Then we design a

Figure 6.1: Flow of information from nature, to observation, into computer. mda/XFig systems



prediction-error filter \mathbf{A} on \mathbf{y} , which yields a white-spectrumed output residual \mathbf{r} . Because \mathbf{r} and \mathbf{x} theoretically have the same spectrum, the tantalizing prospect is that maybe \mathbf{r} equals \mathbf{x} , meaning that the PEF \mathbf{A} has *deconvolved* the unknown convolution \mathbf{B} .

Causal with causal inverse

Theoretically, a PEF is a causal filter with a causal inverse. This adds confidence to the likelihood that deconvolution of natural processes with a PEF might get the correct phase spectrum as well as the correct amplitude spectrum. Naturally, the PEF does not give the correct phase to an “all-pass” filter. That is a filter with a phase shift but a constant amplitude spectrum. (I think most migration operators are in this category.)

Theoretically we should be able to use a PEF in either convolution or polynomial division. There are some dangers though, mainly connected with dealing with data in small windows. Truncation phenomena might give us PEF estimates that are causal, but whose inverse is not, so they cannot be used in polynomial division. This is a lengthy topic in the classic literature. This old, fascinating subject is examined in my books, FGDP and PVI. A classic solution is one by John Parker Burg. We should revisit the Burg method in light of the helix.

PEF output tends to whiteness

The most important property of a prediction-error filter or PEF is that its output tends to a white spectrum (to be proven here). No matter what the input to this filter, its output tends to whiteness as the number of the coefficients $n \rightarrow \infty$ tends to infinity. Thus, the PE filter adapts itself to the input by absorbing all its color. This has important statistical implications and important geophysical implications.

Spectral estimation

The PEF's output being white leads to an important consequence: To specify a spectrum, we can give the spectrum (of an input) itself, give its autocorrelation, or give its PEF coefficients. Each is transformable to the other two. Indeed, an effective mechanism of spectral estimation, developed by John P. Burg and described in FGDP, is to compute a PE filter and look at the inverse of its spectrum.

Short windows

The power of a PE filter is that a short filter can often extinguish, and thereby represent, the information in a long resonant filter. If the input to the PE filter is a sinusoid, it is exactly predictable by a three-term recurrence relation, and all the color is absorbed by a three-term PE filter (see exercises). Burg's spectral estimation is especially effective in short windows.

Weathered layer resonance

That the output spectrum of a PE filter is white is also useful geophysically. Imagine the reverberation of the soil layer, highly variable from place to place, as the resonance between the surface and shallow more-consolidated soil layers varies rapidly with surface location because of geologically recent fluvial activity. The spectral color of this erratic variation on surface-recorded seismograms is compensated for by a PE filter. Usually we do not want PE-filtered seismograms to be white, but once they all have the same spectrum, it is easy to postfilter them to any desired spectrum.

6.3.1 PEF whiteness proof in 1-D

The basic idea of least-squares fitting is that the residual is orthogonal to the fitting functions. Applied to the PE filter, this idea means that the output of a PE filter is orthogonal to lagged inputs. The orthogonality applies only for lags in the past, because prediction knows only the past while it aims to the future. What we want to show here is different, namely, that the output is uncorrelated with *itself* (as opposed to the input) for lags in *both* directions; hence the output spectrum is white.

In (6.21) are two separate and independent autoregressions, $\mathbf{0} \approx \mathbf{Y}_a \mathbf{a}$ for finding the filter \mathbf{a} , and $\mathbf{0} \approx \mathbf{Y}_b \mathbf{b}$ for finding the filter \mathbf{b} . By noticing that the two matrices are really the same (except a row of zeros on the bottom of \mathbf{Y}_a is a row in the top of \mathbf{Y}_b) we realize that the two regressions must result in the same filters $\mathbf{a} = \mathbf{b}$, and the residual \mathbf{r}_b is a shifted version of \mathbf{r}_a . In practice, I visualize the matrix being a thousand components tall (or a million) and a hundred components wide.

$$\mathbf{0} \approx \mathbf{r}_a = \begin{bmatrix} y_1 & 0 & 0 \\ y_2 & y_1 & 0 \\ y_3 & y_2 & y_1 \\ y_4 & y_3 & y_2 \\ y_5 & y_4 & y_3 \\ y_6 & y_5 & y_4 \\ 0 & y_6 & y_5 \\ 0 & 0 & y_6 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \end{bmatrix} ; \quad \mathbf{0} \approx \mathbf{r}_b = \begin{bmatrix} 0 & 0 & 0 \\ y_1 & 0 & 0 \\ y_2 & y_1 & 0 \\ y_3 & y_2 & y_1 \\ y_4 & y_3 & y_2 \\ y_5 & y_4 & y_3 \\ y_6 & y_5 & y_4 \\ 0 & y_6 & y_5 \\ 0 & 0 & y_6 \end{bmatrix} \begin{bmatrix} 1 \\ b_1 \\ b_2 \end{bmatrix} \quad (6.21)$$

When the energy $\mathbf{r}'\mathbf{r}$ of a residual has been minimized, the residual \mathbf{r} is orthogonal to the fitting functions. For example, choosing a_2 to minimize $\mathbf{r}'\mathbf{r}$ gives $0 = \partial \mathbf{r}'\mathbf{r} / \partial a_2 = 2\mathbf{r}'\partial \mathbf{r} / \partial a_2$. This shows that \mathbf{r}' is perpendicular to $\partial \mathbf{r} / \partial a_2$ which is the rightmost column of the \mathbf{Y}_a matrix. Thus the vector \mathbf{r}_a is orthogonal to all the columns in the \mathbf{Y}_a matrix except the first (because we do not minimize with respect to a_0).

Our goal is a different theorem that is imprecise when applied to the three coefficient filters displayed in (6.21), but becomes valid as the filter length tends to infinity $\mathbf{a} = (1, a_1, a_2, a_3, \dots)$ and the matrices become infinitely wide. Actually, all we require is the last component in \mathbf{b} , namely b_n tend to zero. This generally happens because as n increases, y_{t-n} becomes a weaker and weaker predictor of y_t .

The matrix \mathbf{Y}_a contains all of the columns that are found in \mathbf{Y}_b except the last (and the last one is not important). This means that \mathbf{r}_a is not only orthogonal to all of \mathbf{Y}_a 's columns (except the first) but \mathbf{r}_a is also orthogonal to all of \mathbf{Y}_b 's columns except the last. Although \mathbf{r}_a isn't really perpendicular to the last column of \mathbf{Y}_b , it doesn't matter because that column has hardly any contribution to \mathbf{r}_b since $|b_n| \ll 1$. Because \mathbf{r}_a is (effectively) orthogonal to all the components of \mathbf{r}_b , \mathbf{r}_a is also orthogonal to \mathbf{r}_b itself. (For any \mathbf{u} and \mathbf{v} , if $\mathbf{r} \cdot \mathbf{u} = 0$ and $\mathbf{r} \cdot \mathbf{v} = 0$ then $\mathbf{r} \cdot (\mathbf{u} + \mathbf{v}) = 0$ and also $\mathbf{r} \cdot (a_1 \mathbf{u} + a_2 \mathbf{v}) = 0$).

Here is a detail: In choosing the example of equation (6.21), I have shifted the two fitting problems by only one lag. We would like to shift by more lags and get the same result. For this we need more filter coefficients. By adding many more filter coefficients we are adding many more columns to the right side of \mathbf{Y}_b . That's good because we'll be

needing to neglect more columns as we shift \mathbf{r}_b further from \mathbf{r}_a . Neglecting these columns is commonly justified by the experience that “after short range regressors have had their effect, long range regressors generally find little remaining to predict.” (Recall that the damped harmonic oscillator from physics, the finite difference equation that predicts the future from the past, uses only two lags.)

Here is the main point: Since \mathbf{r}_b and \mathbf{r}_a both contain the same signal \mathbf{r} but time-shifted, the orthogonality at all shifts means that the autocorrelation of \mathbf{r} vanishes at all lags. An exception, of course, is at zero lag. The autocorrelation does not vanish there because \mathbf{r}_a is not orthogonal to its first column (because we did not minimize with respect to a_0).

As we redraw $\mathbf{0} \approx \mathbf{r}_b = \mathbf{Y}_b \mathbf{b}$ for various lags, we may shift the columns only downward because shifting them upward would bring in the first column of \mathbf{Y}_a and the residual \mathbf{r}_a is not orthogonal to that. Thus we have only proven that one side of the autocorrelation of \mathbf{r} vanishes. That is enough however, because autocorrelation functions are symmetric, so if one side vanishes, the other must also.

If \mathbf{a} and \mathbf{b} were two-sided filters like $(\dots, b_{-2}, b_{-1}, 1, b_1, b_2, \dots)$ the proof would break. If \mathbf{b} were two-sided, \mathbf{Y}_b would catch the nonorthogonal column of \mathbf{Y}_a . Not only is \mathbf{r}_a not proven to be perpendicular to the first column of \mathbf{Y}_a , but it cannot be orthogonal to it because a signal cannot be orthogonal to itself.

The implications of this theorem are far reaching. The residual \mathbf{r} , a convolution of \mathbf{y} with \mathbf{a} has an autocorrelation that is an impulse function. The Fourier transform of an impulse is a constant. Thus the spectrum of the residual is “white”. Thus \mathbf{y} and \mathbf{a} have mutually inverse spectra.

Since the output of a PEF is white, the PEF itself has a spectrum inverse to its input.

An important application of the PEF is in missing data interpolation. We’ll see examples later in this chapter. My third book, PVI² has many examples³ in one dimension with both synthetic data and field data including the `gap` parameter. Here we next extend these ideas to two (or more) dimensions.

6.3.2 Simple dip filters

Convolution in two dimensions is just like convolution in one dimension except that convolution is done on two axes. The input and output data are planes of numbers and the filter is also a plane. A two-dimensional filter is a small plane of numbers that is convolved over a big data plane of numbers.

Suppose the data set is a collection of seismograms uniformly sampled in space. In other words, the data is numbers in a (t, x) -plane. For example, the following filter destroys any wavefront aligned along the direction of a line containing both the “+1” and the “−1”.

$$\begin{array}{ccc} -1 & \cdot & \\ \cdot & \cdot & \\ \cdot & 1 & \end{array} \quad (6.22)$$

² <http://sepwww.stanford.edu/sep/prof/pvi/toc.html/index.html>

³ <http://sepwww.stanford.edu/sep/prof/pvi/tsa/paper.html/node1.html>

The next filter destroys a wave with a slope in the opposite direction:

$$\begin{array}{ccc} & & 1 \\ & \cdot & \\ -1 & & \cdot \end{array} \quad (6.23)$$

To convolve the above two filters, we can reverse either on (on both axes) and correlate them, so that you can get

$$\begin{array}{ccccc} & & & & \\ & & \cdot & -1 & \cdot \\ & 1 & & \cdot & \cdot \\ & \cdot & & \cdot & 1 \\ & \cdot & -1 & & \cdot \end{array} \quad (6.24)$$

which destroys waves of both slopes.

A two-dimensional filter that can be a dip-rejection filter like (6.22) or (6.23) is

$$\begin{array}{ccc} a & \cdot & \\ b & \cdot & \\ c & 1 & \\ d & \cdot & \\ e & \cdot & \end{array} \quad (6.25)$$

where the coefficients (a, b, c, d, e) are to be estimated by least squares in order to minimize the power out of the filter. (In the filter table, the time axis runs vertically.)

Fitting the filter to two neighboring traces that are identical but for a time shift, we see that the filter coefficients (a, b, c, d, e) should turn out to be something like $(-1, 0, 0, 0, 0)$ or $(0, 0, -.5, -.5, 0)$, depending on the dip (stepout) of the data. But if the two channels are not fully coherent, we expect to see something like $(-.9, 0, 0, 0, 0)$ or $(0, 0, -.4, -.4, 0)$. To find filters such as (6.24), we adjust coefficients to minimize the power out of filter shapes, as in

$$\begin{array}{ccccc} v & a & \cdot & & \\ w & b & \cdot & & \\ x & c & 1 & & \\ y & d & \cdot & & \\ z & e & \cdot & & \end{array} \quad (6.26)$$

With 1-dimensional filters, we think mainly of power spectra, and with 2-dimensional filters we can think of temporal spectra and spatial spectra. What is new, however, is that in two dimensions we can think of dip spectra (which is when a 2-dimensional spectrum has a particularly common form, namely when energy organizes on radial lines in the (ω, k_x) -plane). As a short (three-term) 1-dimensional filter can devour a sinusoid, we have seen that simple 2-dimensional filters can devour a small number of dips.

6.3.3 PEF whiteness proof in 2-D

A well-known property (see FGDP or PVI) of a 1-D PEF is that its energy clusters immediately after the impulse at zero delay time. Applying this idea to the helix in Figure shows us that we can consider a 2-D PEF to be a small halfplane like with an impulse along a side. These shapes are what we see here in Figure 6.2.

Figure 6.2: A 2-D whitening filter template, and itself lagged. At output locations “A” and “B,” the filter coefficient is constrained to be “1”. When the semicircles are viewed as having infinite radius, the B filter is contained in the A filter. Because the output at A is orthogonal to all its inputs, which include all inputs of B, the output at A is orthogonal to the output of B.




Figure 6.2 shows the input plane with a 2-D filter on top of it at two possible locations. The filter shape is a semidisk, which you should imagine being of infinitely large radius. Notice that semidisk A includes all the points in B. The output of disk A will be shown to be orthogonal to the output of disk B. Conventional least squares theory says that the coefficients of the filter are designed so that the output of the filter is orthogonal to each of the inputs to that filter (except for the input under the “1,” because any nonzero signal cannot be orthogonal to itself). Recall that if a given signal is orthogonal to each in a given group of signals, then the given signal is orthogonal to all linear combinations within that group. The output at B is a linear combination of members of its input group, which is included in the input group of A, which are already orthogonal to A. Therefore the output at B is orthogonal to the output at A. In summary,

residual	\perp	fitting function
output at A	\perp	each input to A
output at A	\perp	each input to B
output at A	\perp	linear combination of each input to B
output at A	\perp	output at B

The essential meaning is that a particular lag of the output autocorrelation function vanishes.

Study Figure 6.2 to see for what lags all the elements of the B filter are wholly contained in the A filter. These are the lags where we have shown the output autocorrelation to be vanishing. Notice another set of lags where we have proven nothing (where B is moved to the right of A). Autocorrelations are centrosymmetric, which means that the value at any lag is the same as the value at the negative of that lag, even in 2-D and 3-D where the lag is a vector quantity. Above we have shown that a halfplane of autocorrelation values vanishes. By the centrosymmetry, the other half must vanish too. Thus the autocorrelation of the PEF output is an impulse function, so its 2-D spectrum is white.

The helix tells us why the proper filter form is not a square with the “1” on the corner. Before I discovered the helix, I understood it another way (that I learned from John P. Burg): For a spectrum to be white, *all* nonzero autocorrelation lags must be zero-valued. If the filter were a quarter-plane, then the symmetry of autocorrelations would only give us vanishing in another quarter, so there would be two remaining quarter-planes where the

autocorrelation was not zero.

Fundamentally, the white-output theorem requires a one-dimensional ordering to the values in a plane or volume. The filter must contain a halfplane of values so that symmetry gives the other half.

You will notice some nonuniqueness. We could embed the helix with a 90° rotation in the original physical application. Besides the difference in side boundaries, the 2-D PEF would have a different orientation. Both PEFs should have an output that tends to whiteness as the filter is enlarged. It seems that we could design whitening autoregression filters for 45° rotations also, and we could also design them for hexagonal coordinate systems. In some physical problems, you might find the nonuniqueness unsettling. Does it mean the “final solution” is nonunique? Usually not, or not seriously so. Recall even in one dimension, the time reverse of a PEF has the same spectrum as the original PEF. When a PEF is used for regularizing a fitting problem, it is worth noticing that the quadratic form minimized is the PEF times its adjoint so the phase drops out. Likewise, a missing data restoration also amounts to minimizing a quadratic form so the phase again drops out.

6.3.4 Examples of modeling and deconvolving with a 2-D PEF

Here we examine elementary signal-processing applications of 2-D prediction-error filters (PEFs) on both everyday 2-D textures and on seismic data. Some of these textures are easily modeled with prediction-error filters (PEFs) while others are not. All figures used the same 10×10 filter shape. No attempt was made to optimize filter size or shape or any other parameters.

Results in Figures 6.3-6.9 are shown with various familiar textures⁴ on the left as training data sets. From these training data sets, a prediction-error filter (PEF) is estimated using module `pef` on page 160. The center frame is simulated data made by deconvolving (polynomial division) random numbers by the estimated PEF. The right frame is the more familiar process, convolving the estimated PEF on the training data set.




Figure 6.3: Synthetic granite matches the training image quite well. The prediction error (PE) is large at grain boundaries so it almost seems to outline the grains.
mda/morgan granite

Theoretically, the right frame tends towards a white spectrum. Earlier you could notice

⁴ I thank Morgan Brown for finding these textures.

the filter size by knowing that the output was taken to be zero where the filter is only partially on the data. This was annoying on real data where we didn't want to throw away any data around the sides. Now the filtering is done without a call to the boundary module so we have typical helix wraparound.




Figure 6.4: Synthetic wood grain has too little white. This is because of the nonsymmetric brightness histogram of natural wood. Again, the PEF output looks random as expected.
mda/morgan wood




Figure 6.5: A banker's suit (left). A student's suit (center). My suit (right). The prediction error is large where the weave changes direction.
mda/morgan herr

Since a PEF tends to the inverse of the spectrum of its input, results similar to these could probably be found using Fourier transforms, smoothing spectra, etc. We used PEFs because of their flexibility. The filters can be any shape. They can dodge around missing data, or we can use them to estimate missing data. We avoid periodic boundary assumptions inherent to FT. The PEF's are designed only internal to known data, not off edges so they are readily adaptable to nonstationarity. Thinking of these textures as seismic time slices, the textures could easily be required to pass thru specific values at well locations.

6.3.5 Seismic field data examples

Figures 6.10-6.13 are based on exploration seismic data from the Gulf of Mexico deep water. A ship carries an air gun and tows a streamer with some hundreds of geophones. First we look at a single pop of the gun. We use all the geophone signals to create a single 1-D PEF for the time axis. This changes the average temporal frequency spectrum as shown




Figure 6.6: Basket weave. The simulated data fails to segregate the two dips into a checkerboard pattern. The PEF output looks structured perhaps because the filter is too small. mda/morgan fabric




Figure 6.7: Brick. Synthetic brick edges are everywhere and do not enclose blocks containing a fixed color. PEF output highlights the mortar. mda/morgan brick




Figure 6.8: Ridges. A spectacular failure of the stationarity assumption. All dips are present but in different locations. Never-the-less, the ridges have been sharpened by the deconvolution. mda/morgan ridges




Figure 6.9: Gulf of Mexico seismic section, modeled, and deconvolved. Do you see any drilling prospects in the simulated data? In the deconvolution, the strong horizontal layering is suppressed giving a better view of the hyperbolas. The decon filter is the same 10×10 used on the everyday textures. mda/morgan WGstack




Figure 6.10: ω spectrum of a shot gather of Figure 6.11 before and after 1-D decon with a 30 point filter. mda/pefex specdecon

in Figure 6.10. Signals from 60 Hz to 120 Hz are boosted substantially. The raw data has evidently been prepared with strong filtering against signals below about 8 Hz. The PEF attempts to recover these signals, mostly unsuccessfully, but it does boost some energy near the 8 Hz cutoff. Choosing a longer filter would flatten the spectrum further. The big question is, “Has the PEF improved the appearance of the data?”

The data itself from the single pop, both before and after PE-filtering is shown in Figure 6.11. For reasons of esthetics of human perception I have chosen to display a mirror image of the PEF’ed data. To see a blink movie of superposition of before-and-after images you need the electronic book. We notice that signals of high temporal frequencies indeed have the expected hyperbolic behavior in space. Thus, these high-frequency signals are wavefields, not mere random noise.

Given that all visual (or audio) displays have a bounded range of amplitudes, increasing the frequency content (bandwidth) means that we will need to turn down the amplification so we do not wish to increase the bandwidth unless we are adding signal.

Increasing the spectral bandwidth always requires us to diminish the gain.

The same ideas but with a two-dimensional PEF are in Figure 6.12 (the same data but with more of it squeezed onto the page.) As usual, the raw data is dominated by events arriving later at greater distances. After the PEF, we tend to see equal energy in dips in all directions. We have strongly enhanced the “backscattered” energy, those events that arrive later at *shorter* distances.

Figure 6.13 shows echos from the all shots, the nearest receiver on each shot. This picture of the earth is called a “near-trace section.” This earth picture shows us why there is so much backscattered energy in Figure 6.12 (which is located at the left side of Figure 6.13). The backscatter comes from any of the many of near-vertical faults.

We have been thinking of the PEF as a tool for shaping the spectrum of a display. But does it have a physical meaning? What might it be? Referring back to the beginning of the chapter we are inclined to regard the PEF as the convolution of the source waveform with some kind of water-bottom response. In Figure 6.12 we used many different shot-receiver separations. Since each different separation has a different response (due to differing moveouts) the water bottom reverberation might average out to be roughly an impulse. Figure 6.12 is a different story. Here for each shot location, the distance to the receiver is constant. Designing a single channel PEF we can expect the PEF to contain both the shot waveform and the water bottom layers because both are nearly identical in all the shots. We would rather have a PEF that represents only the shot waveform (and perhaps a radiation pattern).

Let us consider how we might work to push the water-bottom reverberation out of the PEF. This data is recorded in water 600 meters deep. A consequence is that the sea bottom is made of fine-grained sediments that settled very slowly and rather similarly from place to place. In shallow water the situation is different. The sands near estuaries are always shifting. Sedimentary layers thicken and thin. They are said to “on-lap and off-lap.” Here I do notice where the water bottom is sloped the layers do thin a little. To push the water bottom layers out of the PEF our idea is to base its calculation not on the raw data, but on the spatial prediction error of the raw data. On a perfectly layered earth a perfect spatial




Figure 6.11: Raw data with its mirror. Mirror had 1-D PEF applied, 30 point filter.
`mda/pefex decon0`




Figure 6.12: A 2-D filter (here 20×5) brings out the backscattered energy.
 mda/pefex decon1




Figure 6.13: Raw data, near-trace section (top). Filtered with a two-channel PEF (bottom). The movie has other shaped filters. `mda/pefex zof`

prediction error filter would zero all traces but the first one. Since a 2-D PEF includes spatial prediction as well as temporal prediction, we can expect it to contain much less of the sea-floor layers than the 1-D PEF. If you have access to the electronic book, you can blink the figure back and forth with various filter shapes.

6.4 PEF ESTIMATION WITH MISSING DATA

If we are not careful, our calculation of the PEF could have the pitfall that it would try to use the missing data to find the PEF, and hence it would get the wrong PEF. To avoid this pitfall, imagine a PEF finder that uses weighted least squares where the weighting function vanishes on those fitting equations that involve missing data. The weighting would be unity elsewhere. Instead of weighting bad results by zero, we simply will not compute them. The residual there will be initialized to zero and never changed. Likewise for the adjoint, these components of the residual will never contribute to a gradient. So now we need a convolution program that produces no outputs where missing inputs would spoil it.

Recall there are two ways of writing convolution, equation (1.4) when we are interested in finding the filter *inputs*, and equation (1.5) when we are interested in finding the *filter itself*. We have already coded equation (1.4), operator `helicon` on page 8. That operator was useful in missing data problems. Now we want to find a prediction-error filter so we need the other case, equation (1.5), and we need to ignore the outputs that will be broken because of missing inputs. The operator module `hconest` does the job.

```

                                user/gee/hconest.c
44      for (ia = 0; ia < na; ia++) {
45          for (iy = aa->lag[ia]; iy < ny; iy++) {
46              if(aa->mis[iy]) continue;
47
48              ix = iy - aa->lag[ia];
49
50              if( adj) a[ia] -= y[iy] * x[ix];
51              else    y[iy] -= a[ia] * x[ix];
52          }
53      }
```

We are seeking a prediction error filter $(1, a_1, a_2)$ but some of the data is missing. The data is denoted \mathbf{y} or y_i above and x_i below. Because some of the x_i are missing, some of the regression equations in (6.27) are worthless. When we figure out which are broken, we will put zero weights on those equations.

$$\mathbf{0} \approx \mathbf{r} = \mathbf{W}\mathbf{X}\mathbf{a} = \begin{bmatrix} w_1 & . & . & . & . & . & . & . \\ . & w_2 & . & . & . & . & . & . \\ . & . & w_3 & . & . & . & . & . \\ . & . & . & w_4 & . & . & . & . \\ . & . & . & . & w_5 & . & . & . \\ . & . & . & . & . & w_6 & . & . \\ . & . & . & . & . & . & w_7 & . \\ . & . & . & . & . & . & . & w_8 \end{bmatrix} \begin{bmatrix} x_1 & 0 & 0 \\ x_2 & x_1 & 0 \\ x_3 & x_2 & x_1 \\ x_4 & x_3 & x_2 \\ x_5 & x_4 & x_3 \\ x_6 & x_5 & x_4 \\ 0 & x_6 & x_5 \\ 0 & 0 & x_6 \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \end{bmatrix} \quad (6.27)$$

Suppose that x_2 and x_3 were missing or known bad. That would spoil the 2nd, 3rd, 4th, and 5th fitting equations in (6.27). In principle, we want w_2 , w_3 , w_4 and w_5 to be zero. In practice, we simply want those components of \mathbf{r} to be zero.

What algorithm will enable us to identify the regression equations that have become defective, now that x_2 and x_3 are missing? Take filter coefficients (a_0, a_1, a_2, \dots) to be all ones. Let \mathbf{d}_{free} be a vector like \mathbf{x} but containing 1's for the missing (or “freely adjustable”) data values and 0's for the known data values. Recall our very first definition of filtering showed we can put the filter in a vector and the data in a matrix or vice versa. Thus $\mathbf{X}\mathbf{a}$ above gives the same result as $\mathbf{A}\mathbf{x}$ below.

$$\begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ m_5 \\ m_6 \\ m_7 \\ m_8 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 2 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \mathbf{A}\mathbf{d}_{\text{free}} \quad (6.28)$$

The numeric value of each m_i tells us how many of its inputs are missing. Where none are missing, we want unit weights $w_i = 1$. Where any are missing, we want zero weights $w_i = 0$. The desired residual under partially missing inputs is computed by module `misinput` on the next page.

6.4.1 Internal boundaries to multidimensional convolution

Sometimes we deal with small patches of data. In order that boundary phenomena not dominate the calculation intended in the central region, we need to take care that input data is not assumed to be zero beyond the interval that the data is given.

The two little triangular patches of zeros in the convolution matrix in equation (6.27) describe end conditions where it is assumed that the data y_t vanishes before $t = 1$ and after $t = 6$. Alternately we might not wish to make that assumption. Thus the triangles filled with zeros could be regarded as missing data. In this one-dimensional example, it is easy to see that the filter, say `yy->mis` should be set to `true` at the ends so no output would ever

user/gee/misinput.c

```


23 void find_mask(int n          /* data size */,
24               const int *known /* mask for known data [n] */,
25               sf_filter aa     /* helical filter */)
26 /*< create a filter mask >*/
27 {
28     int i, ih;
29     float *rr, *dfre;
30
31     rr = sf_floatalloc(n);
32     dfre = sf_floatalloc(n);
33
34     for (i=0; i < n; i++) {
35         dfre[i] = known[i]? 0.:1.;
36     }
37
38     sf_helicon_init(aa);
39
40     for (ih=0; ih < aa->nh; ih++) {
41         aa->flt[ih] = 1.;
42     }
43
44     sf_helicon_lop(false, false, n, n, dfre, rr);
45
46     for (ih=0; ih < aa->nh; ih++) {
47         aa->flt[ih] = 0.;
48     }
49
50     for (i=0; i < n; i++) {
51         if ( rr[i] > 0.) aa->mis[i] = true;
52     }
53
54     free (rr);
55     free (dfre);
56 }

```

be computed there. We would like to find a general multidimensional algorithm to correctly specify `yy->mis` around the multidimensional boundaries. This proceeds like the missing data algorithm, i.e. we apply a filter of all ones to a data space template that is taken all zeros except ones at the locations of missing data, in this case y_0, y_{-1} and y_7, y_8 . This amounts to surrounding the original data set with some missing data. We need padding the size of the filter on all sides. The padded region would be filled with ones (designating missing inputs). Where the convolution output is nonzero, there `yy->mis` is set to `true` denoting an output with missing inputs.

The two-dimensional case is a little more cluttered than the 1-D case but the principle is about the same. Figure 4 shows a larger input domain, a 5×3 filter, and a smaller output domain. There are two things to notice. First, sliding the filter everywhere inside the outer

Figure 6.14: Domain of inputs and outputs of a two-dimensional filter like a PEF.



box, we get outputs (under the 1 location) only in the inner box. Second, (the adjoint idea) crosscorrelating the inner and outer boxes gives us the 3×5 patch of information we use to build the filter coefficients. We need to be careful not to assume that signals vanish outside the region where they are defined. In a later chapter we will break data spaces into overlapping patches, separately analyze the patches, and put everything back together. We do this because crosscorrelations change with time and they are handled as constant in short time windows. There we must be particularly careful that zero signal values not be presumed outside of the small volumes; otherwise the many edges and faces of the many small volumes can overwhelm the interior that we want to study.

In practice, the input and output are allocated equal memory, but the output residual is initialized to zero everywhere and then not computed except where shown in figure 4. Below is module `bound` to build a selector for filter outputs that should never be examined or even computed (because they need input data from outside the given data space). Inputs are a filter `aa` and the size of its cube `na = (na(1), na(2), ...)`. Also input are two cube dimensions, that of the data last used by the filter `nold` and that of the filter's next intended use `nd`. (`nold` and `nd` are often the same). Module `bound` begins by defining a bigger data space with room for a filter surrounding the original data space `nd` on all sides. It does this by the line `nb=nd+2*na`. Then we allocate two data spaces `xx` and `yy` of the bigger size `nb` and pack many ones in a frame of width `na` around the outside of `xx`. The filter `aa` is also filled with ones. The filter `aa` must be regridded for the bigger `nb` data space (regridding merely changes the lag values of the ones). Now we filter the input `xx` with `aa` getting `yy`. Wherever the output is nonzero, we have an output that has been affected by the boundary. Such an output should not be computed. Thus we allocate the logical mask `aa->mis` (a part of the helix filter definition in module `helix` on page 7) and wherever we see a nonzero value of `yy` in the output, we designate the output as depending on missing inputs by setting `aa->mis` to `true`.

```

user/gee/bound.c
28 void bound (int dim          /* number of dimensions */,
29             const int *nold /* old data coordinates [dim] */,
30             const int *nd   /* new data coordinates [dim] */,
31             const int *na   /* filter box size [dim] */,
32             const sf_filter aa /* helix filter */)
33 /*< Mark helix filter outputs where input is off data. >*/
34 {
35     int iy, my, ib, mb, i, nb[SF_MAX_DIM], ii[SF_MAX_DIM];
36     float *xx, *yy;
37
38     my = mb = 1;
39     for (i=0; i < dim; i++) {
40         nb[i] = nd[i] + 2*na[i]; /* nb is a bigger space. */
41         mb *= nb[i];
42         my *= nd[i];
43     }
44
45     xx = sf_floatalloc(mb); yy = sf_floatalloc(mb);
46
47     for (ib=0; ib < mb; ib++) {
48         sf_line2cart(dim, nb, ib, ii);
49         xx[ib] = 0.;
50         for (i=0; i < dim; i++)
51             if(ii[i]+1 <= na[i] || ii[i]+1 > nb[i]-na[i]) {
52                 xx[ib] = 1.;
53                 break;
54             }
55     }
56     sf_helicon_init( aa);
57     regrid(dim, nold, nb, aa);
58     for (i=0; i < aa->nh; i++) aa->flt[i] = 1.;
59     /* apply filter */
60     sf_helicon_lop(false, false, mb, mb, xx, yy);
61     regrid(dim, nb, nd, aa);
62     for (i=0; i < aa->nh; i++) aa->flt[i] = 0.;
63
64     aa->mis = sf_boolalloc(my); /* attach missing designation */
65     for (iy = 0; iy < my; iy++) { /* map to padded space */
66         sf_line2cart(dim, nd, iy, ii);
67         for (i=0; i < dim; i++) ii[i] += na[i];
68         ib = sf_cart2line(dim, nb, ii);
69         aa->mis[iy] = (bool) (yy[ib] > 0.);
70     }
71
72     free (xx); free (yy);
73 }

```

In reality one would set up the boundary conditions with module `bound` before identifying locations of missing data with module `misinput`. Both modules are based on the same concept, but the boundaries are more cluttered and confusing which is why we examined them later.

6.4.2 Finding the prediction-error filter

The first stage of the least-squares estimation is computing the prediction-error filter. The second stage will be using it to find the missing data. The input data space contains a mixture of known data values and missing unknown ones. For the first stage of finding the filter, we generally have many more fitting equations than we need so we can proceed by ignoring the fitting equations that involve missing data values. We ignore them everywhere that the missing inputs hit the filter.

The codes here do not address the difficulty that maybe too much data is missing so that all weights are zero. To add stabilization we could supplement the data volume with a “training dataset” or by a “prior filter”. With things as they are, if there is not enough data to specify a prediction-error filter, you should encounter the error exit from `cgstep()` on page 51.

```

user/gee/pef.c
25 void find_pef(int nd      /* data size */,
26               float* dd   /* data [nd] */,
27               sf_filter aa /* estimated filter */,
28               int niter    /* number of iterations */)
29 /*< find PEF >*/
30 {
31     hconest_init( dd, aa);
32     sf_solver(hconest_lob, sf_cgstep, aa->nh, nd, aa->flt, dd,
33              niter, "x0", aa->flt, "end");
34     sf_cgstep_close();
35 }
```

6.5 TWO-STAGE LINEAR LEAST SQUARES

In Chapter 3 and Chapter 5 we filled empty bins by minimizing the energy output from the filtered mesh. In each case there was arbitrariness in the choice of the filter. Here we find and use the optimum filter, the PEF.

The first stage is that of the previous section, finding the optimal PEF while carefully avoiding using any regression equations that involve boundaries or missing data. For the second stage, we take the PEF as known and find values for the empty bins so that the power out of the prediction-error filter is minimized. To do this we find missing data with module `mis2()` on page 133.

This two-stage method avoids the nonlinear problem we would otherwise face if we included the fitting equations containing both free data values and free filter values. Presumably, after two stages of linear least squares we are close enough to the final solution that we could switch over to the full nonlinear setup described near the end of this chapter.

The synthetic data in Figure 6.15 is a superposition of two plane waves of different directions, each with a random (but low-passed) waveform. After punching a hole in the data, we find that the lost data is pleasingly restored, though a bit weak near the side boundary. This imperfection could result from the side-boundary behavior of the operator or from an insufficient number of missing-data iterations.




Figure 6.15: Original data (left), with a zeroed hole, restored, residual selector (right).
mda/hole hole

The residual selector in Figure 6.15 shows where the filter output has valid inputs. From it you can deduce the size and shape of the filter, namely that it matches up with Figure 4. The ellipsoidal hole in the residual selector is larger than that in the data because we lose regression equations not only at the hole, but where any part of the filter overlaps the hole.

The results in Figure 6.15 are essentially perfect representing the fact that that synthetic example fits the conceptual model perfectly. Before we look at the many examples in Figures 6.16-6.19 we will examine another gap-filling strategy.

6.5.1 Adding noise (Geostat)

In chapter 3 we restored missing data by adopting the philosophy of minimizing the energy in filtered output. In this chapter we learned about an optimum filter for this task, the prediction-error filter (PEF). Let us name this method the “minimum noise” method of finding missing data.

A practical problem with the minimum-noise method is evident in a large empty hole such as in Figures 6.16- 6.17. In such a void the interpolated data diminishes greatly. Thus we have not totally succeeded in the goal of “hiding our data acquisition footprint” which

we would like to do if we are trying to make pictures of the earth and not pictures of our data acquisition footprint.

What we will do next is useful in some applications but not in others. Misunderstood or misused it is rightly controversial. We are going to fill the empty holes with something that looks like the original data but really isn't. I will distinguish the words "synthetic data" (that derived from a physical model) from "simulated data" (that manufactured from a statistical model). We will fill the empty holes with simulated data like what you see in the center panels of Figures 6.3-6.9. We will add just enough of that "wall paper noise" to keep the variance constant as we move into the void.

Given some data \mathbf{d} , we use it in a filter operator \mathbf{D} , and as described with equation (6.27) we build a weighting function \mathbf{W} that throws out the broken regression equations (ones that involve missing inputs). Then we find a PEF \mathbf{a} by using this regression.

$$\mathbf{0} \approx \mathbf{r} = \mathbf{W}\mathbf{D}\mathbf{a} \quad (6.29)$$

Because of the way we defined \mathbf{W} , the "broken" components of \mathbf{r} vanish. We need to know the variance σ of the nonzero terms. It can be expressed mathematically in a couple different ways. Let $\mathbf{1}$ be a vector filled with ones and let \mathbf{r}^2 be a vector containing the squares of the components of \mathbf{r} .

$$\sigma = \sqrt{\frac{1}{N} \sum_i r_i^2} = \sqrt{\frac{\mathbf{1}'\mathbf{W}\mathbf{r}^2}{\mathbf{1}'\mathbf{W}\mathbf{1}}} \quad (6.30)$$

Let us go to a random number generator and get a noise vector \mathbf{n} filled with random numbers of variance σ . We'll call this the "added random noise". Now we solve this new regression for the data space \mathbf{d} (both known and missing)

$$\mathbf{0} \approx \mathbf{r} = \mathbf{A}\mathbf{d} - \mathbf{n} \quad (6.31)$$

keeping in mind that known data is constrained (as detailed in chapter 3).

To understand why this works, consider first the training image, a region of known data. Although we might think that the data defines the white noise residual by $\mathbf{r} = \mathbf{A}\mathbf{d}$, we can also imagine that the white noise determines the data by $\mathbf{d} = \mathbf{A}^{-1}\mathbf{r}$. Then consider a region of wholly missing data. This data is determined by $\mathbf{d} = \mathbf{A}^{-1}\mathbf{n}$. Since we want the data variance to be the same in known and unknown locations, naturally we require the variance of \mathbf{n} to match that of \mathbf{r} .

A very minor issue remains. Regression equations may have all of their required input data, some of it, or none of it. Should the \mathbf{n} vector add noise to every regression equation? First, if a regression equation has all its input data that means there are no free variables so it doesn't matter if we add noise to that regression equation because the constraints will overcome that noise. I don't know if I should worry about how *many* inputs are missing for each regression equation.

It is fun making all this interesting "wall paper" noticing where it is successful and where it isn't. We cannot help but notice that it seems to work better with the genuine geophysical data than it does with many of the highly structured patterns. Geophysical data is expensive to acquire. Regrettably, we have uncovered a technology that makes counterfeiting much easier.

Examples are in Figures 6.16-6.19. In the electronic book, the right-side panel of each figure is a movie, each panel being derived from different random numbers.




Figure 6.16: The herringbone texture is a patchwork of two textures. We notice that data missing from the hole tends to fill with the texture at the edge of the hole. The spine of the herring fish, however, is not modeled at all. `mda/morgan herr-hole-fillr`




Figure 6.17: The brick texture has a mortar part (both vertical and horizontal joins) and a brick surface part. These three parts enter the empty area but do not end where they should. `mda/morgan brick-hole-fillr`

The seismic data in Figure 6.19 illustrates a fundamental principle: In the restored hole we do not see the same spectrum as we do on the other panels. This is because the hole is filled, not with all frequencies (or all slopes) but with those that are most predictable. The filled hole is devoid of the unpredictable noise that is a part of all real data.

6.5.2 Inversions with geostat

In geophysical estimation (inversion) we use model styling (regularization) to handle the portion of the model that is not determined by the data. This results in the addition of minimal noise. Alternately, like in Geostatistics, we could make an assumption of statistical stationarity and add much more noise so the signal variance in poorly determined regions matches that in well determined regions. Here is how to do this. Given the usual data fitting and model styling goals

$$\mathbf{0} \approx \mathbf{Lm} - \mathbf{d} \quad (6.32)$$

$$\mathbf{0} \approx \mathbf{Am} \quad (6.33)$$




Figure 6.18: The theoretical model is a poor fit to the ridge data since the prediction must try to match ridges of all possible orientations. This data requires a broader theory which incorporates the possibility of nonstationarity (space variable slope).

mda/morgan ridges-hole-fillr




Figure 6.19: Filling the missing seismic data. The imaging process known as “migration” would suffer diffraction artifacts in the gapped data that it would not suffer on the restored data.

mda/morgan WGstack-hole-fillr

We introduce a sample of random noise \mathbf{n} and fit instead these regressions

$$\mathbf{0} \approx \mathbf{Lm} - \mathbf{d} \quad (6.34)$$

$$\mathbf{0} \approx \mathbf{Am} - \mathbf{n} \quad (6.35)$$

Of course you get a different solution for each different realization of the random noise. You also need to be a little careful to use noise \mathbf{n} of the appropriate variance. Figure 6.20 shows a result on the SeaBeam data. Bob Clapp developed this idea at SEP and also applied it




Figure 6.20: Top left is binned data. Top right extends the data with a PEF. The bottom two panels add appropriately colored random noise in the regions of missing data.

mda/geostat bobsea

to interval velocity estimation, the example of Figures 5.1-5.3.

6.5.3 Infill of 3-D seismic data from a quarry blast

Finding missing data (filling empty bins) requires use of a filter. Because of the helix, the codes work in spaces of all dimensions.

An open question is how many conjugate-direction iterations are needed in missing-data programs. When estimating filters, I set the iteration count `niter` at the number of free filter parameters. Theoretically, this gives me the exact solution but sometimes I run double the number of iterations to be sure. The missing-data estimation, however is a completely different story. The number of free parameters in the missing-data estimation, could be very large. This often implies impractically long compute times for the exact solution. In practice I experiment carefully with `niter` and hope for the best. I find that where gaps are small, they fill in quickly. Where the gaps are large, they don't, and more iterations are required. Where the gaps are large is where we must experiment with preconditioning.

Figure 6.21 shows an example of replacing missing data by values predicted from a 3-D PEF. The data was recorded at Stanford University with a 13×13 array of independent recorders. The figure shows 12 of the 13 lines each of length 13. Our main goal was to measure the ambient night-time noise. By morning about half the recorders had dead batteries but the other half recorded a wave from a quarry blast. The raw data was distracting to look at because of the many missing traces so I interpolated it with a small 3-D filter. That filter was a PEF.




Figure 6.21: The left 12 panels are the inputs. The right 12 panels are outputs.
mda/miss3 passfill

6.5.4 Imposing prior knowledge of symmetry

Reversing a signal in time does not change its autocorrelation. In the analysis of stationary time series, it is well known (FGDP) that the filter for predicting forward in time should

be the same as that for “predicting” backward in time (except for time reversal). When the data samples are short, however, a different filter may be found for predicting forward than for backward. Rather than average the two filters directly, the better procedure is to find the filter that minimizes the sum of power in two residuals. One is a filtering of the original signal, and the other is a filtering of a time-reversed signal, as in equation (6.36), where the top half of the equations represent prediction-error predicting forward in time and the second half is prediction backward.

$$\begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \end{bmatrix} = \begin{bmatrix} y_3 & y_2 & y_1 \\ y_4 & y_3 & y_2 \\ y_5 & y_4 & y_3 \\ y_6 & y_5 & y_4 \\ y_1 & y_2 & y_3 \\ y_2 & y_3 & y_4 \\ y_3 & y_4 & y_5 \\ y_4 & y_5 & y_6 \end{bmatrix} \begin{bmatrix} 1 \\ a_1 \\ a_2 \end{bmatrix} \quad (6.36)$$

To get the bottom rows from the top rows, we simply reverse the order of all the components within each row. That reverses the input time function. (Reversing the order within a column would reverse the output time function.) Instead of the matrix being diagonals tipping 45° to the right, they tip to the left. We could make this matrix from our old familiar convolution matrix and a time-reversal matrix

$$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

It is interesting to notice how time-reversal symmetry applies to Figure 6.15. First of all, with time going both forward and backward the residual space gets twice as big. The time-reversal part gives a selector for Figure 6.15 with a gap along the right edge instead of the left edge. Thus, we have acquired a few new regression equations.

Some of my research codes include these symmetries, but I excluded them here. Nowhere did I see that the reversal symmetry made noticeable difference in results, but in coding, it makes a noticeable clutter by expanding the residual to a two-component *residual array*.


Where a data sample grows exponentially towards the boundary, I expect that extrapolated data would diverge too. You can force it to go to zero (or any specified value) at some distance from the body of the known data. To do so, surround the body of data by missing data and surround that by specification of “enough” zeros. “Enough” is defined by the filter length.

6.5.5 Hexagonal coordinates

In a two-dimensional plane it seems that the one-sidedness of the PEF could point in any direction. Since we usually have a rectangular mesh, however, we can only do the calculations along the axes so we have only two possibilities, the helix can wrap around the 1-axis, or it can wrap around the 2-axis.

Figure 6.22: Top is known data. Middle includes the interpolated values. Bottom is the filter with the leftmost point constrained to be unity and other points chosen to minimize output power.


mda/misif man1



Curiously, constraining the filter to be of the form $(a_{-2}, a_{-1}, 1)$ in Figure 6.23 yields the same interpolated missing data as in Figure 6.22. I understand that the sum squared of the coefficients of $A(Z)P(Z)$ is the same as that of $A(1/Z)P(Z)$, but I do not see why that would imply the same interpolated data; never the less, it seems to.

Figure 6.23: The filter here had its rightmost point constrained to be unity—i.e., this filtering amounts to backward prediction. The interpolated data seems to be identical to that of forward prediction.

mda/misif man3



6.6.1 Objections to interpolation error

In any data interpolation or extrapolation, we want the extended data to behave like the original data. And, in regions where there is no observed data, the extrapolated data should drop away in a fashion consistent with its spectrum determined from the known region.


My basic idea is that the spectrum of the missing data should match that of the known data. This is the idea that the spectrum should be unchanging from a known region to an unknown region. A technical word to express the idea of spectra not changing is “stationary.” This happens with the PEF (one-sided filter) because its spectrum tends to the inverse of that of the known data while that of the unknown data tends to the inverse of that of the PEF. Thus the spectrum of the missing data is the “inverse of the inverse” of the spectrum of the known. The PEF enables us to fill in the missing area with the spectral shape of the known area. (In regions far away or unpredictable, the spectral shape may be the same, but the energy drops to zero.)

On the other hand, the interpolation-error filter, a filter like $(a_{-2}, a_{-1}, 1, a_1, a_2)$, should fail to do the job because it has the wrong spectrum. (I am stating this fact without proof).

To confirm and show these concepts, I prepared synthetic data consisting of a fragment of a damped exponential, and off to one side of it an impulse function. Most of the energy is in the damped exponential. Figure 6.24 shows that the spectrum and the extended data


are about what we would expect. From the extrapolated data, it is impossible to see where the given data ends.

Figure 6.24: Top is synthetic data with missing portions. Middle includes the interpolated values. Bottom is the filter, a *prediction-error* filter which may look symmetric but is not quite. mda/misif syn1



For comparison, I prepared Figure 6.25. It is the same as Figure 6.24, except that the filter is constrained in the middle. Notice that the extended data does *not* have the spectrum of the given data—the wavelength is much shorter. The boundary between real data and extended data is not nearly as well hidden as in Figure 6.24.

Figure 6.25: Top is the same synthetic data. Middle includes the interpolated values. Bottom is the filter, an *interpolation-error* filter. mda/misif syn3



6.6.2 Packing both missing data and filter into a vector

Now let us examine the theory and coding behind the above examples. Define a roughening filter $A(\omega)$ and a data signal $Y(\omega)$ at some stage of interpolation. The fitting goal is $0 \approx A(\omega)Y(\omega)$ where the filter $A(\omega)$ has at least one time-domain coefficient constrained to be nonzero and the data contains both known and missing values. Think of perturbations ΔA and ΔY . We neglect the nonlinear term $\Delta A \Delta Y$ as follows:

$$0 \approx (A + \Delta A)(Y + \Delta Y) \quad (6.37)$$

$$0 \approx A\Delta Y + Y\Delta A + AY + \Delta A\Delta Y \quad (6.38)$$

$$0 \approx A\Delta Y + Y\Delta A + AY \quad (6.39)$$

Let us use matrix algebraic notation to rewrite the fitting goals (6.39). For this we need mask matrices (diagonal matrices with ones on the diagonal where variables are free and zeros where they are constrained i.e., where $\Delta a_i = 0$ and $\Delta y_i = 0$). The free-mask matrix for missing data is denoted \mathbf{J} and that for the PE filter is \mathbf{K} . The fitting goal (6.39) becomes

$$\mathbf{0} \approx \mathbf{AJ}\Delta\mathbf{y} + \mathbf{YK}\Delta\mathbf{a} + (\mathbf{Ay} \text{ or } \mathbf{Ya}) \quad (6.40)$$

Defining the original residual as $\bar{\mathbf{r}} = \mathbf{A}\mathbf{y}$ this becomes

$$\mathbf{0} \approx \begin{bmatrix} \mathbf{AJ} & \mathbf{YK} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{y} \\ \Delta\mathbf{a} \end{bmatrix} + \bar{\mathbf{r}} \quad (6.41)$$

For a 3-term filter and a 7-point data signal, the fitting goal (6.40) becomes

$$\begin{bmatrix} a_0 & . & . & . & . & . & . & y_0 & . & . \\ a_1 & a_0 & . & . & . & . & . & y_1 & y_0 & . \\ a_2 & a_1 & a_0 & . & . & . & . & y_2 & y_1 & y_0 \\ . & a_2 & a_1 & a_0 & . & . & . & y_3 & y_2 & y_1 \\ . & . & a_2 & a_1 & a_0 & . & . & y_4 & y_3 & y_2 \\ . & . & . & a_2 & a_1 & a_0 & . & y_5 & y_4 & y_3 \\ . & . & . & . & a_2 & a_1 & a_0 & y_6 & y_5 & y_4 \\ . & . & . & . & . & a_2 & a_1 & . & y_6 & y_5 \\ . & . & . & . & . & . & a_2 & . & . & y_6 \end{bmatrix} \begin{bmatrix} \mathbf{J} & \mathbf{0} \\ \mathbf{0} & \mathbf{K} \end{bmatrix} \begin{bmatrix} \Delta y_0 \\ \Delta y_1 \\ \Delta y_2 \\ \Delta y_3 \\ \Delta y_4 \\ \Delta y_5 \\ \Delta y_6 \\ \hline \Delta a_0 \\ \Delta a_1 \\ \Delta a_2 \end{bmatrix} + \begin{bmatrix} \bar{r}_0 \\ \bar{r}_1 \\ \bar{r}_2 \\ \bar{r}_3 \\ \bar{r}_4 \\ \bar{r}_5 \\ \bar{r}_6 \\ \bar{r}_7 \\ \bar{r}_8 \end{bmatrix} \approx \mathbf{0} \quad (6.42)$$

Recall that \bar{r}_t is the convolution of a_t with y_t , namely, $\bar{r}_0 = y_0 a_0$ and $\bar{r}_1 = y_0 a_1 + y_1 a_0$, etc. To optimize this fitting goal we first initialize $\mathbf{a} = (1, 0, 0, \dots)$ and then put zeros in for missing data in \mathbf{y} . Then we iterate over equations (6.43) to (6.47).

$$\mathbf{r} \leftarrow \mathbf{A}\mathbf{y} \quad (6.43)$$

$$\begin{bmatrix} \Delta\mathbf{y} \\ \Delta\mathbf{a} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{J}'\mathbf{A}' \\ \mathbf{K}'\mathbf{Y}' \end{bmatrix} \mathbf{r} \quad (6.44)$$

$$\Delta\mathbf{r} \leftarrow \begin{bmatrix} \mathbf{AJ} & \mathbf{YK} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{y} \\ \Delta\mathbf{a} \end{bmatrix} \quad (6.45)$$

$$\mathbf{y} \leftarrow \text{cgstep}(\mathbf{y}, \Delta\mathbf{y}) \quad (6.46)$$

$$\mathbf{a} \leftarrow \text{cgstep}(\mathbf{a}, \Delta\mathbf{a}) \quad (6.47)$$

This is the same idea as all the linear fitting goals we have been solving, except that now we recompute the residual \mathbf{r} inside the iteration loop so that as convergence is achieved (*if* it is achieved), the neglected nonlinear term $\Delta A \Delta Y$ tends to zero.

My initial research proceeded by linearization like (6.39). Although I ultimately succeeded, I had enough difficulties that I came to realize that linearization is dangerous. When you start “far enough” from the correct solution the term $\Delta A \Delta Y$ might not actually be small enough. You don’t know how small is small, because these are not scalars but operators. Then the solution may not converge to the minimum you want. Your solution will depend on where you start from. I no longer exhibit the nonlinear solver `missif` until I find a real data example where it produces noticeably better results than multistage linear-least squares.

The alternative to linearization is two-stage linear least squares. In the first stage you estimate the PEF; in the second you estimate the missing data. If need be, you can re-estimate the PEF using all the data both known and missing (downweighted if you prefer).

If you don't have enough regression equations because your data is irregularly distributed, then you can use binning. Still not enough? Try coarser bins. The point is that nonlinear solvers will not work unless you begin close enough to the solution, and the way to get close is by arranging first to solve a sensible (though approximate) linearized problem. Only as a last resort, after you have gotten as near as you can, should you use the nonlinear least-squares techniques.

6.7 LEVELED INVERSE INTERPOLATION

Eighteenth- and nineteenth- century mathematics literature gives us many methods of interpolating functions. These classical methods are generally based on polynomials. The user specifies some order of polynomial and the theory works out the coefficients. Today our interest is in both interpolating and extrapolating wavefields (which are solutions to low order differential equations) and we use methods that are much better behaved than polynomials when extrapolating data, methods which behave acceptably when faced with contradictory data values, and methods which also apply in two and three dimensions.

In Chapter 3, subroutine `invint1()` on page ?? solved the problem of inverse linear interpolation, which is, given scattered data points, to find a function on a uniform mesh from which linear interpolation gives the scattered data points. To cope with regions having no data points, the subroutine requires an input roughening filter. This is a bit like specifying a differential equation to be satisfied between the data points. The question is, how should we choose a roughening filter? The importance of the roughening filter grows as the data gets sparser or as the mesh is refined.

Figures 6.22-6.25 suggest that the choice of the roughening filter need not be subjective, nor a priori, but that the prediction-error filter (PEF) is the ideal roughening filter. Spectrally, the PEF tends to the inverse of its input hence its output tends to be “level”. Missing data that is interpolated with this “leveler” tends to have the spectrum of given data.

6.7.1 Test results for leveled inverse interpolation

Figures 6.26 and 6.27 show the same example as in Figures 3.13 and 3.14. What is new here is that the proper PEF is not given but is determined from the data. Figure 6.26 was made with a three-coefficient filter $(1, a_1, a_2)$ and Figure 6.27 was made with a five-coefficient filter $(1, a_1, a_2, a_3, a_4)$. The main difference in the figures is where the data is sparse. The data points in Figures ??, 6.26 and 6.27 are samples from a sinusoid.

Comparing Figures 3.13 and 3.14 to Figures 6.26 and 6.27 we conclude that by finding and imposing the prediction-error filter while finding the model space, we have interpolated beyond aliasing in data space.

Figure 6.26: Interpolating with a three-term filter. The interpolated signal is fairly monofrequency.
mda/levint subsine3





Figure 6.27: Interpolating with a five term filter.
mda/levint subsine5



6.7.2 Analysis for leveled inverse interpolation

Here we see how the interpolation beyond aliasing was done. The first “statement of wishes” is that the observational data \mathbf{d} should result from a linear interpolation \mathbf{L} of the uniformly sampled model space \mathbf{m} ; that is, $\mathbf{0} \approx \mathbf{Lm} - \mathbf{d}$. Expressing this as a change $\Delta\mathbf{m}$ gives the fitting goal in terms of the model change, $\mathbf{0} \approx \mathbf{L}\Delta\mathbf{m} + (\mathbf{Lm} - \mathbf{d}) = \mathbf{L}\Delta\mathbf{m} + \mathbf{r}$. The second wish is really an assertion that a good way to find missing parts of a function (the model space) is to solve for the function and its PEF at the same time. We are merging the fitting goal (3.15) for irregularly sampled data with the fitting goal (6.42) for finding the prediction-error filter.

$$\mathbf{0} \approx \mathbf{r}_d = \mathbf{L}\Delta\mathbf{m} + (\mathbf{Lm} - \mathbf{d}) \quad (6.48)$$

$$\mathbf{0} \approx \mathbf{r}_m = \mathbf{A}\Delta\mathbf{m} + \mathbf{MK}\Delta\mathbf{a} + (\mathbf{Am} \text{ or } \mathbf{Ma}) \quad (6.49)$$

Writing this out in full for 3 data points and 6 model values on a uniform mesh and a PEF of 3 terms, we have

$$\left[\begin{array}{cccccc|cccc} .8 & .2 & . & . & . & . & & & & \\ . & . & 1 & . & . & . & & & & \\ . & . & . & . & .5 & .5 & & & & \\ \hline a_0 & . & . & . & . & . & m_0 & . & . & \\ a_1 & a_0 & . & . & . & . & m_1 & m_0 & . & \\ a_2 & a_1 & a_0 & . & . & . & m_2 & m_1 & m_0 & \\ . & a_2 & a_1 & a_0 & . & . & m_3 & m_2 & m_1 & \\ . & . & a_2 & a_1 & a_0 & . & m_4 & m_3 & m_2 & \\ . & . & . & a_2 & a_1 & a_0 & m_5 & m_4 & m_3 & \\ . & . & . & . & a_2 & a_1 & . & m_5 & m_4 & \\ . & . & . & . & . & a_2 & . & . & m_5 & \end{array} \right] \left[\begin{array}{c} \Delta m_0 \\ \Delta m_1 \\ \Delta m_2 \\ \Delta m_3 \\ \Delta m_4 \\ \Delta m_5 \\ \Delta m_6 \\ \hline \Delta a_0 \\ \Delta a_1 \\ \Delta a_2 \end{array} \right] + \left[\begin{array}{c} r_{d0} \\ r_{d1} \\ r_{d2} \\ \hline r_{m0} \\ r_{m1} \\ r_{m2} \\ r_{m3} \\ r_{m4} \\ r_{m5} \\ r_{m6} \\ r_{m7} \end{array} \right] \approx \mathbf{0} \quad (6.50)$$

where r_m is the convolution of the filter a_t and the model m_t , where r_d is the data misfit $\mathbf{r} = \mathbf{L}\mathbf{m} - \mathbf{d}$, and where \mathbf{K} was defined in equation (6.11).

Before you begin to use this nonlinear fitting goal, you need some starting guesses for \mathbf{m} and \mathbf{a} . The guess $\mathbf{m} = 0$ is satisfactory (as explained later). For the first guess of the filter, I suggest you load it up with $\mathbf{a} = (1, -2, 1)$ as I did for the examples here.

6.7.3 Seabeam: theory to practice

I provide here a more fundamental theory for dealing with the Seabeam data. I originally approached the data in this more fundamental way, but with time, I realized that I paid a high price in code complexity, computational speed, and reliability. The basic problem is that the elegant theory requires a good starting model which can only come from the linearized theory. I briefly recount the experience here, because the fundamental theory is interesting and because in other applications, you will face the challenge of sorting out the fundamental features from the essential features.

The linear-interpolation operator carries us from a uniform mesh to irregularly distributed data. Fundamentally we seek to solve the inverse problem to go the other direction. A nonlinear approach to filling in the missing data is suggested by the one-dimensional examples in Figures 6.26–6.27, where the PEF and the missing data are estimated simultaneously. The nonlinear approach has the advantage that it allows for completely arbitrary data positioning, whereas the two-stage linear approach forces the data to be on a uniform mesh and requires there not be too many empty mesh locations.

For the 2-D nonlinear problem, we follow the same approach we used in one dimension, equations (6.48) and (6.49), except that the filtering and the linear interpolations are two dimensional.

I have had considerable experience with this problem on this data set and I can report that bin filling is easier and works much more quickly and reliably. Eventually I realized that the best way to start the nonlinear iteration (6.48-6.49) is with the final result of bin filling. Then I learned that the extra complexity of the nonlinear iteration (6.48-6.49) offers little apparent improvement to the quality of the SeaBeam result. (This is not to say that we should not try more variations on the idea).

Not only did I find the binning method faster, but I found it to be *much* faster (compare a minute to an hour). The reasons for being faster (most important first) are,

1. Binning reduces the amount of data handled in each iteration by a factor of the average number of points per bin.
2. The 2-D linear interpolation operator adds many operations per data point.
3. Using two fitting goals seems to require more iterations.

(Parenthetically, I later found that helix preconditioning speeds the Seabeam interpolation from minutes to seconds.)

The most serious criticism of the nonlinear approach is that it does not free us from the linearized approaches. We need them to get a “close enough” starting solution to the

nonlinear problem. I learned that the iteration (6.48-6.49), like most nonlinear sequences, behaves unexpectedly and badly when you start too far from the desired solution. For example, I often began from the assumed PEF being a Laplacian and the original map being fit from that. Oddly, from this starting location I sometimes found myself stuck. The iteration (6.48-6.49) would not move towards the map we humans consider a better one.

Having said all those bad things about iteration (6.48-6.49), I must hasten to add that with a different type of data set, you might find the results of (6.48-6.49) to be significantly better.

6.7.4 Risky ways to do nonlinear optimization

I have noticed that some geophysicists have adopted a risky method of nonlinear optimization, which is not advocated in the professional optimization literature. This risky method is to linearize a goal (with a multiparameter model space), then optimize the linearized goal, then relinearize, etc. The safer method is to relinearize after each step of CD.

An instructive example I learned about many years ago was earthquake epicenter location. Model space is latitude, longitude, and origin time. When people added a new variable, the depth, the solutions went wild until they learned to restrict the depth to zero until the other three parameters were stabilized. Apparently the instability stems from the fact that depth and origin time affect distant receivers in a similar way.

6.7.5 The bane of PEF estimation

This is the place where I would like to pat myself on the back for having “solved” the problem of missing data. Actually, an important practical problem remains. I’ve been trying to coax younger, more energetic people to think about it. The problem arises when there is too much missing data.

The bane of PEF estimation is too much missing data

Then *all* the regression equations disappear. The nonlinear methods are particularly bad because if they don’t have a good enough starting location, they can and do go crazy. My only suggestion is to begin with a linear PEF estimator. Shrink the PEF and coarsen the mesh in model space until you do have enough equations. Starting from there, hopefully you can refine this crude solution without dropping into a local minimum.

Another important practical problem remains, that of nonstationarity. We’ll see the beginnings of the solution to that problem in chapter .

6.8 MULTIVARIATE SPECTRUM

A common spectrum is the Fourier spectrum. More fundamentally, a spectrum is a decomposition of a model space or data space into components. The components are in some sense independent; more specifically, the components are orthogonal to one another. An-

other well-known spectrum is provided by eigenvectors and eigenvalues. In statistical signal processing we handle a third type of spectrum, the multivariate spectrum.

Working in an optimization problem, we begin from residuals between theory and practice. These residuals can be scaled to make new optimization residuals before we start minimizing their energy. What scaling should we use? The scaling can be a simple weighting function or a filter. A filter is simply a weighting function in Fourier space.

The basic idea of common sense, which also comes to us as results proven by Gauss or from the theory of statistical signal processing, is this: The optimization residuals should be roughly of equal scale. This makes sense because squaring magnifies scale, and anything small will be ignored while anything large will dominate. Scaling optimization residuals to be in a common range makes them all equally influential on the final solution. Not only should optimization residuals be of like scale in physical space, they should be of like scale in Fourier space or eigenvector space, or any other space that we might use to represent the optimization residuals. This implies that the optimization residuals should be uncorrelated. If the optimization residuals were correlated, they would have a spectrum that was not white. Not white means of differing sizes in Fourier space. Residuals should be the same size as one another in physical space, likewise in Fourier space. Thus the optimization residuals should be orthogonal and of unit scale, much like Fourier components or as eigenvectors are orthonormal.

Let us approach the problem backwards. Suppose we have two random variables that we take to be the ideal optimization residuals x_1 and x_2 . In reality the two may be few or trillions. In the language of statistics, the optimization residuals are expected to have zero mean, an idea that is formalized by writing $E(x_1) = 0$ and $E(x_2) = 0$. Likewise these ideal optimization residuals have equal energy, $E(x_1^2) = 1$ and $E(x_2^2) = 1$. Finally, these two optimization residuals are uncorrelated, a condition which is written as $E(x_1 x_2) = 0$. The expectation symbol $E()$ is like a summation over many instances of the random variable.

Now suppose there exists a transformation \mathbf{B} from these ideal optimization residuals to two experimental residuals y_1 and y_2 , say $\mathbf{y} = \mathbf{B}\mathbf{x}$ where

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (6.51)$$

The experimental residuals y_1 and y_2 are likely to be neither orthogonal nor equal in energy. From the column vector \mathbf{y} , the experimenter can form a square matrix. Let us also allow the experimenter to write the symbol $E()$ to denote summation over many trials or over many sections of data, ranges over time or space, over soundings or over receiver locations. The experimenter writes

$$\mathbf{R} = E(\mathbf{y}\mathbf{y}') \quad (6.52)$$

$$\mathbf{R} = E(\mathbf{B}\mathbf{x}\mathbf{x}'\mathbf{B}') \quad (6.53)$$

Given a random variable r , the expectation of $2r$ is simply $E(2r) = 2E(r)$. The $E()$ symbol is a summation on random variables, but constants like the coefficients of \mathbf{B} pass right

through it. Thus,

$$\mathbf{R} = \mathbf{B} E(\mathbf{x}\mathbf{x}') \mathbf{B}' \quad (6.54)$$

$$\mathbf{R} = \mathbf{B} E \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \begin{bmatrix} x_1 & x_2 \end{bmatrix} \right) \mathbf{B}' \quad (6.55)$$

$$\mathbf{R} = \mathbf{B} \begin{bmatrix} E(x_1x_1) & E(x_1x_2) \\ E(x_2x_1) & E(x_2x_2) \end{bmatrix} \mathbf{B}' \quad (6.56)$$

$$\mathbf{R} = \mathbf{B}\mathbf{B}' \quad (6.57)$$

Given a matrix \mathbf{R} , there is a simple well-known method called the Cholesky factorization method that will factor \mathbf{R} into two parts like \mathbf{B} and \mathbf{B}' . The method creates for us either an upper or a lower triangular matrix (our choice) for \mathbf{B} . You can easily reinvent the Cholesky method if you multiply the symbols for two triangular matrices like \mathbf{B} and \mathbf{B}' and notice the procedure that works backwards from \mathbf{R} to \mathbf{B} . The experimenter seeks not \mathbf{B} , however, but its inverse, the matrix that takes us from the experimental residuals to the ideal optimization residuals that are uncorrelated and of equal energies. The Cholesky factorization costs N^3 computations, which is about the same as the cost of the matrix inversion of \mathbf{R} or \mathbf{B} . For geophysical maps and other functions on Cartesian spaces, the Prediction Error Filter (PEF) accomplishes the same general goal and has the advantage that we have already learned how to perform the operation using operators instead of matrices.

The multivariate spectrum of experimental residuals \mathbf{y} is the matrix $\mathbf{R} = E(\mathbf{y}\mathbf{y}')$. For optimum model finding, the experimental residuals (squared) should be weighted inversely (matrix inverse) by their multivariate spectrum.

If I were a little stronger at analysis (or rhetoric) I would tell you that the optimizers pre-conditioned variable \mathbf{p} is the statisticians IID (Independent Identically Distributed) random variable. For stationary (statistically constant) signals and images, \mathbf{A}_m is the model-space PEF. Echo soundings and interval velocity have statistical properties that change with depth. There \mathbf{A}_m is a diagonal weighting matrix (perhaps before or after a PEF).

6.8.1 What should we optimize?

Least-squares problems often present themselves as fitting goals such as

$$\mathbf{0} \approx \mathbf{F}\mathbf{m} - \mathbf{d} \quad (6.58)$$

$$\mathbf{0} \approx \mathbf{m} \quad (6.59)$$

To balance our possibly contradictory goals we need weighting functions. The quadratic form that we should minimize is

$$\min_m (\mathbf{F}\mathbf{m} - \mathbf{d})' \mathbf{A}_n' \mathbf{A}_n (\mathbf{F}\mathbf{m} - \mathbf{d}) + \mathbf{m}' \mathbf{A}_m' \mathbf{A}_m \mathbf{m} \quad (6.60)$$

where $\mathbf{A}_n' \mathbf{A}_n$ is the inverse multivariate spectrum of the noise (data-space residuals) and $\mathbf{A}_m' \mathbf{A}_m$ is the inverse multivariate spectrum of the model. In other words, \mathbf{A}_n is a leveler on the data fitting error and \mathbf{A}_m is a leveler on the model. There is a curious unresolved issue: What is the most suitable constant scaling ratio of \mathbf{A}_n to \mathbf{A}_m ?

6.8.2 Confusing terminology for data covariance

Confusion often stems from the mean of the data $E(\mathbf{d})$.

An experimentalist would naturally believe that the expectation of the data is solely a function of the data, that it can be estimated by averaging data.

On the other hand, a theoretician's idea of the expectation of the observational data $E(\mathbf{d})$ is that it is the theoretical data \mathbf{Fm} , that the expectation of the data $E(\mathbf{d}) = \mathbf{Fm}$ is a function of the model. The theoretician thinks this way because of the idea of noise $\mathbf{n} = \mathbf{Fm} - \mathbf{d}$ as having zero mean.

Seismological data is highly complex but also highly reproducible. In studies like seismology, the world is deterministic but more complicated than our ability to model. Thus, as a practical matter, the discrepancy between observational data and theoretical data is more realistically attributed to the theoretical data. It is not adequately modeled and computed.

This superficial difference in viewpoint becomes submerged to a more subtle level by statistical textbooks that usually define weighting functions in terms of variances instead of spectra. This is particularly confusing with the noise spectrum $(\mathbf{A}'_n \mathbf{A}_n)^{-1}$. It is often referred to as the “data covariance” defined as $E[(\mathbf{d} - E(\mathbf{d}))(\mathbf{d} - E(\mathbf{d}))']$. Clearly, the noise spectrum is the same as the data covariance only if we accept the theoretician's definition that $E(\mathbf{d}) = \mathbf{Fm}$.

There is no ambiguity and no argument if we drop the word “variance” and use the word “spectrum”. Thus, (1) the “inverse noise spectrum” is the appropriate weighting for data-space residuals; and (2) the “inverse model spectrum” is the appropriate model-space weighting. Theoretical expositions generally require these spectra to be given as “prior information.” In this book we see how, when the model space is a map, we can solve for the “prior information” along with everything else.

The statistical words “covariance matrix” are suggestive and appealing, but I propose not to use them because of the ambiguity of $E(\mathbf{d})$. For example, we understand that people who say “data covariance” intend the “multivariate noise spectrum” but we cannot understand their meaning of “model covariance”. They should intend the “multivariate model spectrum” but that implies that $E(\mathbf{m}) = \mathbf{0}$, which seems wrong. Avoiding the word “covariance” avoids the problem.

6.8.3 Hermeneutics

Hermeneutics is the study of the methodological principles of interpretation. Historically, it refers to bible study. Never-the-less, it seems entirely appropriate for Geophysical Estimation. If Albert's book is “Inverse Problem Theory” and mine is “Inverse Problem Practice”, and if the difference between theory and practice is smaller in theory than it is in practice, then there are two fundamental questions:

1. In theory, what is the difference between theory and practice? In theory, the difference is data error.
2. In practice, what is the difference between theory and practice? One suggestion is that the discrepancy is entirely due to inadequate modeling. It is well known that

geophysical data is highly repeatable. The problem is that the modeling neglects far too much.

Here is a perspective drawn from analysis of the human genome: "The problem is that it is possible to use empirical data to calibrate a model that generates simulated data that is similar to the empirical data. The point of using such a calibrated model is to be able to show how strange certain regions are if they don't fit the simulated distribution, which is based on the empirical distribution." In other words, "inversion" is just the process of calibrating a model. To learn something new we study the *failures* of such models.

Chapter 7

Spatial aliasing and scale invariance

Landforms are not especially predictable. Therefore, crude PEF approximations are often satisfactory. Wavefields are another matter. Consider the “shape” of the acoustic wavefronts at this moment in the room you are in. The acoustic wavefield has statistical order in many senses. If the 3-D volume is filled with waves emitted from a few point sources, then (with some simplifications) what could be a volume of information is actually a few 1-D signals. When we work with wavefronts we can hope for more dramatic, even astounding, results from estimating properly.

The plane-wave model links an axis that is not aliased (time) with axes (space) that often are.

We often characterize data from any region of (t, x) -space as “good” or “noisy” when we really mean it contains “few” or “many” plane-wave events in that region. Where regions are noisy, there is no escaping the simple form of the Nyquist limitation. Where regions are good we may escape it. Real data typically contains both kinds of regions. Undersampled data with a broad distribution of plane waves is nearly hopeless. Undersampled data with a sparse distribution of plane waves offer us the opportunity to resample without aliasing. Consider data containing a spherical wave. The angular bandwidth in a plane-wave decomposition appears huge *until we restrict attention to a small region* of the data. (Actually a spherical wave contains very little information compared to an arbitrary wave field.) It can be very helpful in reducing the local angular bandwidth if we can deal effectively with tiny pieces of data. If we can deal with tiny pieces of data, then we can adapt to rapid spatial and temporal variations. This chapter shows such tiny windows of data.

7.1 INTERPOLATION BEYOND ALIASING

A traditional method of data interpolation on a regular mesh is a four-step procedure: (1) Set zero values at the points to be interpolated; (2) Fourier transform; (3) Set to zero the high frequencies; and (4) Inverse transform. This is a fine method and is suitable for many applications in both one dimension and higher dimensions. However, this method fails to

take advantage of our prior knowledge that seismic data has abundant fragments of plane waves that link an axis that is not aliased (time) to axes that often are (space).

7.1.1 Interlacing a filter

The filter below can be designed despite alternate missing traces. This filter destroys plane waves. If the plane wave should happen to pass halfway between the “d” and the “e”, those two points could interpolate the halfway point, at least for well-sampled temporal frequencies, and the time axis should always be well sampled. For example, $d = e = -.5$ would almost destroy the plane wave and it is an aliased planewave for its higher frequencies.

$$\begin{array}{ccccccccc} a & \cdot & b & \cdot & c & \cdot & d & \cdot & e \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \end{array} \quad (7.1)$$

We could use module `pef` on page 160 to find the filter (1), if we set up the lag table `lag` appropriately. Then we could throw away alternate zeroed rows and columns (rescale the lag) to get the filter

$$\begin{array}{ccccc} a & b & c & d & e \\ \cdot & \cdot & 1 & \cdot & \cdot \end{array} \quad (7.2)$$

which could be used with subroutine `mis1()` on page 70, to find the interleaved data because both the filters (1) and (2) have the same dip characteristics.

Figure 1 shows three plane waves recorded on five channels and the interpolated data. Both the original data and the interpolated data can be described as “beyond aliasing,”




Figure 7.1: Left is five signals, each showing three arrivals. With the data shown on the left (and no more), the signals have been interpolated. Three new traces appear between each given trace, as shown on the right. lal/lace lace3

because on the input data the signal shifts exceed the signal duration. The calculation requires only a few seconds of a two-stage least-squares method, in which the first stage estimates a PEF (inverse spectrum) of the known data, and the second uses the PEF to estimate the missing traces. Figure 1 comes from PVI which introduces the clever method described above. We will review how that was done and examine the F90 codes

that generalize it to N -dimensions. Then we'll go on to more general methods that allow missing data in any location. Before the methods of this section are applied to field data for migration, data must be broken into many overlapping tiles of size about like those shown here and the results from each tile pieced together. That is described later in chapter .

A PEF is like a differential equation. The more plane-wave solutions you expect, the more lags you need on the data. Returning to Figure 1, the filter must cover four traces (or more) to enable it to predict three plane waves. In this case, $\mathbf{na}=(9,4)$. As usual, the spike on the 2-D PEF is at $\mathbf{center}=(5,1)$. We see the filter is expanded by a factor of $\mathbf{jump}=4$. The data size is $\mathbf{nd}=(75,5)$ and $\mathbf{gap}=0$. Before looking at the code `lace` on page 4 for estimating the PEF, it might be helpful to recall the basic utilities `line2cart` and `cart2line` on the current page for conversion between a multidimensional space and the helix filter lag. We need to sweep across the whole filter and “stretch” its lags on the 1-axis. We do not need to stretch its lags on the 2-axis because the data has not yet been interlaced by zero traces. The line `ii[0] *= jump` means we interlace the 1-axis but not the 2-axis because the data has not yet been interlaced with zero traces.

After the PEF has been found, we can get missing data in the usual way with with module `mis2` on page 133.

7.2 MULTISCALE, SELF-SIMILAR FITTING

Large objects often resemble small objects. To express this idea we use *axis scaling* and we apply it to the basic theory of prediction-error filter (PEF) fitting and missing-data estimation.

Equations (3) and (4) compute the same thing by two different methods, $\mathbf{r} = \mathbf{Y}\mathbf{a}$ and $\mathbf{r} = \mathbf{A}\mathbf{y}$. When it is viewed as fitting goals minimizing $\|\mathbf{r}\|$ and used along with suitable constraints, (3) leads to finding filters and spectra, while (4) leads to finding missing data.

$$\begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \\ r_9 \end{bmatrix} = \begin{bmatrix} y_2 & y_1 \\ y_3 & y_2 \\ y_4 & y_3 \\ y_5 & y_4 \\ y_6 & y_5 \\ y_3 & y_1 \\ y_4 & y_2 \\ y_5 & y_3 \\ y_6 & y_4 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{bmatrix} \mathbf{a} \quad (7.3)$$

user/gee/lace.c

```

27 sf_filter lace_pef(int dim      /* number of dimensions */,
28                  float *dd     /* data */,
29                  int jump      /* filter stretch */,
30                  int n         /* data size */,
31                  int *nd       /* data dimensions [dim] */,
32                  int *center   /* filter center [dim] */,
33                  int *gap      /* filter gap [dim] */,
34                  int *na       /* filter size [dim] */)
35 /*< estimate PEF >*/
36 {
37     int *savelags, ii[SF_MAX_DIM]; /* holding place */
38     int ih, nh, lag0, j;
39     sf_filter aa;
40
41     aa = createhelix(dim, nd, center, gap, na);
42     savelags = aa->lag;
43     nh = aa->nh;
44
45     aa->lag = sf_intalloc(nh); /* prepare interlaced helix */
46     lag0 = sf_cart2line(dim, na, center);
47
48     for (ih=0; ih < nh; ih++) { /* sweep through the filter */
49         sf_line2cart(dim, na, ih+lag0+1, ii);
50         for (j=0; j < dim; j++) {
51             ii[j] -= center[j];
52         }
53         ii[0] *= jump; /* interlace on 1-axis */
54         aa->lag[ih] = sf_cart2line(dim, nd, ii);
55     }
56     na[0] *= jump;
57     bound(dim, nd, nd, na, aa); /* define aa->mis */
58     na[0] /= jump;
59
60     find_pef(n, dd, aa, nh*2); /* estimate aa coefficients */
61     free(aa->lag);
62     aa->lag = savelags; /* restore filter lags */
63
64     return aa;
65 }

```

$$\begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \\ r_9 \end{bmatrix} = \begin{bmatrix} a_2 & a_1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & a_2 & a_1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & a_2 & a_1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & a_2 & a_1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & a_2 & a_1 \\ \hline a_2 & \cdot & a_1 & \cdot & \cdot & \cdot \\ \cdot & a_2 & \cdot & a_1 & \cdot & \cdot \\ \cdot & \cdot & a_2 & \cdot & a_1 & \cdot \\ \cdot & \cdot & \cdot & a_2 & \cdot & a_1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \mathbf{y} \quad (7.4)$$

A new concept embedded in (3) and (4) is that one filter can be applicable for different stretchings of the filter's time axis. One wonders, "Of all classes of filters, what subset remains appropriate for stretchings of the axes?"

7.2.1 Examples of scale-invariant filtering

When we consider all functions with vanishing gradient, we notice that the gradient vanishes whether it is represented as $(1, -1)/\Delta x$ or as $(1, 0, -1)/2\Delta x$. Likewise for the Laplacian, in one dimension or more. Likewise for the wave equation, as long as there is no viscosity and as long as the time axis and space axes are stretched by the same amount. The notion of "dip filter" seems to have no formal definition, but the idea that the spectrum should depend mainly on slope in Fourier space implies a filter that is scale-invariant. I expect the most fruitful applications to be with dip filters.

Resonance or viscosity or damping easily spoils scale-invariance. The resonant frequency of a filter shifts if we stretch the time axis. The difference equations

$$y_t - \alpha y_{t-1} = 0 \quad (7.5)$$

$$y_t - \alpha^2 y_{t-2} = 0 \quad (7.6)$$

both have the same solution $y_t = y_0 \alpha^{-t}$. One difference equation has the filter $(1, -\alpha)$, while the other has the filter $(1, 0, -\alpha^2)$, and α is not equal to α^2 . Although these operators differ, when $\alpha \approx 1$ they might provide the same general utility, say as a roughening operator in a fitting goal.

Another aspect to scale-invariance work is the presence of "parasitic" solutions, which exist but are not desired. For example, another solution to $y_t - y_{t-2} = 0$ is the one that oscillates at the Nyquist frequency.

(Viscosity does not necessarily introduce an inherent length and thereby spoil scale-invariance. The approximate frequency independence of sound absorption per wavelength typical in real rocks is a consequence of physical inhomogeneity at all scales. See for example Kjartansson's constant Q viscosity, described in IEI. Kjartansson teaches that the decaying solutions $t^{-\gamma}$ are scale-invariant. There is no "decay time" for the function $t^{-\gamma}$. Differential equations of finite order and difference equations of finite order cannot produce $t^{-\gamma}$ damping, yet we know that such damping is important in observations. It is easy to manufacture $t^{-\gamma}$ damping in Fourier space; $\exp[(-i\omega)^{\gamma+1}]$ is used. Presumably, difference equations can make reasonable approximations over a reasonable frequency range.)

7.2.2 Scale-invariance introduces more fitting equations

The fitting goals (3) and (4) have about double the usual number of fitting equations. Scale-invariance *introduces extra equations*. If the range of scale-invariance is wide, there will be more equations. Now we begin to see the big picture.

1. Refining a model mesh improves accuracy.
2. Refining a model mesh makes empty bins.
3. Empty bins spoil analysis.
4. If there are not too many empty bins we can find a PEF.
5. With a PEF we can fill the empty bins.
6. To get the PEF and to fill bins we need enough equations.
7. Scale-invariance introduces more equations.

An example of these concepts is shown in Figure 2.




Figure 7.2: Overcoming aliasing with multiscale fitting. lal/multiscale mshole


Additionally, when we have a PEF, often we still cannot find missing data because conjugate-direction iterations do not converge fast enough (to fill large holes). Multiscale convolutions should converge quicker because they are like mesh-refinement, which is quick. An example of these concepts is shown in Figure 3.

7.2.3 Coding the multiscale filter operator

Equation (3) shows an example where the first output signal is the ordinary one and the second output signal used a filter interlaced with zeros. We prepare subroutines that allow

Figure 7.3: Large holes are filled faster with multiscale operators.

lal/multiscale msiter



for more output signals, each with its own filter interlace parameter given in the table `jump[ns]`. Each entry in the jump table corresponds to a particular scaling of a filter axis. The number of output signals is `ns` and the number of zeros interlaced between filter points for the `j`-th signal is `jump[j]-1`.

The multiscale helix filter is defined in module `mshelix` on page 7, analogous to the single-scale module `helix` on page 7. A new function `onescale` extracts our usual helix filter of one particular scale from the multiscale filter.

We create a multiscale helix with module `createmshelix` on page 8. An expanded scale helix filter is like an ordinary helix filter except that the lags are scaled according to a `jump`.

First we examine code for estimating a prediction-error filter that is applicable at many scales. We simply invoke the usual filter operator `hconest` on page 155 for each scale.

The multiscale prediction-error filter finding subroutine is nearly identical to the usual subroutine `find_pef()` on page 160. (That routine cleverly ignores missing data while estimating a PEF.) To easily extend `pef` to multiscale filters we replace its call to the ordinary helix filter module `hconest` on page 155 by a call to `mshconest`. The purpose of `pack(dd,.true.)` is to produce the one-dimensional array expected by our solver routines.

Similar code applies to the operator in (4) which is needed for missing data problems. This is like `mshconest` on page 9 except the adjoint is not the filter but the input.

The multiscale missing-data module `msmis2` is just like the usual missing-data module `mis2` on page 133 except that the filtering is done with the multiscale filter `mshelicon`.

user/gee/mshelix.c

```

25 #ifndef _mshelix_h
26
27 typedef struct mshelixfilter {
28     int nh, ns;
29     float* flt;
30     int** lag;
31     bool** mis;
32     sf_filter one;
33 } *msfilter;
34 /* ^ */
35
36 #endif
37
38 void onescale(int i, msfilter aa)
39 /*< select one scale from multiple scales >*/
40 {
41     aa->one->lag = aa->lag[i];
42     if (NULL != aa->mis) aa->one->mis = aa->mis[i];
43 }

```

7.3 References

Canales, L.L., 1984, Random noise reduction: 54th Ann. Internat. Mtg., Soc. Explor. Geophys., Expanded Abstracts, 525-527.

Rothman, D., 1985, Nonlinear inversion, statistical mechanics, and residual statics estimation: Geophysics, **50**, 2784-2798

Spitz, S., 1991, Seismic trace interpolation in the F-X domain: Geophysics, **56**, 785-794.

```

user/gee/createms helix.c
41 msfilter createms helix(int ndim      /* number of dimensions */,
42                        int* nd       /* data size [ndim] */,
43                        int* center   /* filter center [ndim] */,
44                        int* gap      /* filter gap [ndim] */,
45                        int ns        /* number of scales */,
46                        int* jump     /* filter scaling [ns] */,
47                        int* na       /* filter size [ndim] */)
48 /*< allocate and output a multiscale helix filter >*/
49 {
50     msfilter msaa;
51     sf_filter aa;
52     int is , ih , nh , id , n123 , nb[SF_MAX_DIM];
53
54     aa = createhelix(ndim, nd, center, gap, na);
55     nh = aa->nh;
56
57     msaa = msallocate(nh, ns);
58     for (is=0; is < ns; is++) {
59         for (ih=0; ih < nh; ih++) /* expanded scale lags */
60             msaa->lag[is][ih] = aa->lag[ih]*jump[is];
61     }
62     sf_deallocatehelix(aa);
63
64     n123=1;
65     for (id=0; id < ndim; id++) n123 *= nd[id];
66     msaa->mis = sf_boolalloc2(n123, ns);
67
68     aa = msaa->one;
69     for (is=0; is < ns; is++) { /* for all scales */
70         onescale(is, msaa); /* extract a filter */
71         aa->mis = NULL;
72         for (id=0; id < ndim; id++) nb[id] = na[id]*jump[is];
73         bound(ndim, nd, nd, nb, aa); /* set up its boundaries */
74         for (id=0; id < n123; id++)
75             msaa->mis[is][id] = aa->mis[id]; /* save them */
76         free(aa->mis);
77     }
78
79     return msaa;
80 }
81

```

user/gee/mshconest.c

```

48     for (is=0; is < msaa->ns; is++) {
49         onescale(is, msaa);
50         hconest_lop(adj, true, na, nx, a, y+is*nx);
51     }

```

user/gee/mspef.c

```

26 void find_pef(int nd      /* data size */,
27               float* dd   /* data */,
28               msfilter aa /* estimated filter */,
29               int niter   /* number of iterations */)
30 /*< estimate PEF >*/
31 {
32     float *ee;
33     int is, id, ns;
34
35     ns = aa->ns;
36     ee = sf_floatalloc(nd*ns);
37     for (is=0; is < ns; is++) {
38         for (id=0; id < nd; id++) {
39             ee[id+is*nd] = dd[id];
40         }
41     }
42
43     mshconest_init(dd, aa);
44     sf_solver(mshconest_lop, sf_cgstep, aa->nh, nd*ns, aa->flt,
45             ee, niter, "x0", aa->flt, "end");
46     sf_cgstep_close();
47
48     free(ee);
49 }

```

user/gee/mshelicon.c

```

26     for (is=0; is < aa->ns; is++) {
27         onescale(is, aa);
28         sf_helicon_init(aa->one);
29         sf_helicon_lop(adj, true, nx, nx, xx, yy+is*nx);
30     }

```

user/gee/msmis2.c

```

29 void msmis2(int niter      /* number of iterations */,
30             int nx        /* data size */,
31             int ns        /* number of scales */,
32             float *xx      /* data */,
33             msfilter aa    /* filter */,
34             const bool *known /* mask for known data */)
35 /*< interpolate >*/
36 {
37     int ix, nxs;
38     float *dd;
39
40     nxs = ns*nx;
41
42     dd = sf_floatalloc(nxs);
43     for (ix=0; ix < nxs; ix++) {
44         dd[ix]=0.;
45     }
46
47     mshelicon_init(aa);
48     sf_solver (mshelicon_lob, sf_cgstep, nx, nxs, xx, dd, niter,
49               "known", known, "x0", xx, "end");
50
51     sf_cgstep_close();
52     free(dd);
53 }

```


Chapter 8

Nonstationarity: patching

There are many reasons for cutting data planes or image planes into overlapping pieces (patches), operating on the pieces, and then putting them back together again, as depicted in Figure 1. The earth’s dip varies with lateral location and depth. The dip spectrum and spatial spectrum thus also varies. The dip itself is the essence of almost all earth mapping, and its spectrum plays an important role in the estimation any earth properties. In statistical estimation theory, the word to describe changing statistical properties is “nonstationary”.




Figure 8.1: Decomposing a wall of information into windows (also called patches). Left is an example of a 2-D space input to module `patch`. Right shows a close-up of the output (top left corner). `pch/XFig antoine`

We begin this chapter with basic patching concepts along with supporting utility code. The language of this chapter, *patch*, *overlap*, *window*, *wall*, is two-dimensional, but it may

as well be three-dimensional, *cube*, *subcube*, *brick*, or one-dimensional, *line*, *interval*. We sometimes use the language of windows on a wall. But since we usually want to have overlapping windows, better imagery would be to say we assemble a quilt from patches.

The codes are designed to work in any number of dimensions. After developing the infrastructure, we examine some two-dimensional, time- and space-variable applications: adaptive steep-dip rejection, noise reduction by prediction, and segregation of signals and noises.

8.1 PATCHING TECHNOLOGY

A plane of information, either data or an image, say `wall(nwall1, nwall2)`, will be divided up into an array of overlapping windows each window of size `(nwind1, nwind2)`. To choose the number of windows, you specify `(npatch1, npatch2)`. Overlap on the 2-axis is measured by the fraction `(nwind2*npatch2)/nwall2`. We turn to the language of F90 which allows us to discuss N -dimensional hypercubes almost as easily as two-dimensional spaces. We define an N -dimensional volume (like the wall) with the vector `nwall= (nwall1, nwall2, ...)`. We define subvolume size (like a 2-D window) with the vector `nwind=(nwind1, nwind2, ...)`. The number of subvolumes on each axis is `npatch=(npatch1, npatch2, ...)`. The operator `patch` on page 3 simply grabs one patch from the wall, or when used in adjoint form, it puts the patch back on the wall. The number of patches on the wall is `product(npatch)`. Getting and putting all the patches is shown later in module `patching` on page 6.

The i -th patch is denoted by the scalar counter `ipatch`. Typical patch extraction begins by taking `ipatch`, a C linear index, and converting it to a multidimensional subscript `jj` each component of which is less than `npatch`. The patches cover all edges and corners of the given data plane (actually the hypervolume) even where `nwall/npatch` is not an integer, even for axes whose length is not an integer number of the patch length. Where there are noninteger ratios, the spacing of patches is slightly uneven, but we'll see later that it is easy to reassemble seamlessly the full plane from the patches, so the unevenness does not matter. You might wish to review the utilities `line2cart` and `cart2line` on the current page which convert between multidimensional array subscripts and the linear memory subscript before looking at the patch extraction-putback code: The cartesian vector `jj` points to the beginning of a patch, where on the wall the `(1,1,...)` coordinate of the patch lies. Obviously this begins at the beginning edge of the wall. Then we pick `jj` so that the last patch on any axis has its last point exactly abutting the end of the axis. The formula for doing this would divide by zero for a wall with only one patch on it. This case arises legitimately where an axis has length one. Thus we handle the case `npatch=1` by abutting the patch to the beginning of the wall and forgetting about its end. As in any code mixing integers with floats, to guard against having a floating-point number, say 99.9999, rounding down to 99 instead of up to 100, the rule is to always add .5 to a floating point number the moment before converting it to an integer. Now we are ready to sweep a window to or from the wall. The number of points in a window is `size(wind)` or equivalently `product(nwind)`.


Figure 2 shows an example with five nonoverlapping patches on the 1-axis and many overlapping patches on the 2-axis.


```

                                user/gee/patch.c
41 void patch_lop (bool adj, bool add, int nx, int ny,
42                float* wall, float* wind)
43 /*< patch operator >*/
44 {
45     int i, j, shift;
46
47     sf_adjnull (adj, add, nx, ny, wall, wind);
48     sf_line2cart(dim, npatch, ipatch, jj);
49     for(i = 0; i < dim; i++) {
50         if(npatch[i] == 1) {
51             jj[i] = 0;
52         } else if (jj[i] == npatch[i]-1) {
53             jj[i] = nwall[i] - nwind[i];
54         } else {
55             jj[i] = jj[i]*(nwall[i] - nwind[i])/(npatch[i] - 1.0);
56         }
57     }
58
59     /* shift till the patch start */
60     shift = sf_cart2line(dim, nwall, jj);
61     for(i = 0; i < ny; i++) {
62         sf_line2cart(dim, nwind, i, ii);
63         j = sf_cart2line(dim, nwall, ii) + shift;
64         if (adj) wall[j] += wind[i];
65         else     wind[i] += wall[j];
66     }
67 }

```

Figure 8.2: A plane of identical values after patches have been cut and then added back. Results are shown for $nwall=(100,30)$, $nwind=(17,6)$, $npatch=(5,11)$. For these parameters, there is gapping on the horizontal axis and overlap on the depth axis. pch/patch parcel



8.1.1 Weighting and reconstructing

The adjoint of extracting all the patches is adding them back. Because of the overlaps, the adjoint is not the inverse. In many applications, inverse patching is required; i.e. patching things back together seamlessly. This can be done with weighting functions. You can have any weighting function you wish and I will provide you the patching reconstruction operator $\tilde{\mathbf{I}}_p$ in

$$\tilde{\mathbf{d}} = [\mathbf{W}_{\text{wall}} \mathbf{P}' \mathbf{W}_{\text{wind}} \mathbf{P}] \mathbf{d} = \tilde{\mathbf{I}}_p \mathbf{d} \quad (8.1)$$


where \mathbf{d} is your initial data, $\tilde{\mathbf{d}}$ is the reconstructed data, \mathbf{P} is the patching operator, \mathbf{P}' is adjoint patching (adding the patches). \mathbf{W}_{wind} is your chosen weighting function in the window, and \mathbf{W}_{wall} is the weighting function for the whole wall. You specify any \mathbf{W}_{wind} you like, and module `mkwallwt` below builds the weighting function \mathbf{W}_{wall} that you need to apply to your wall of reconstructed data, so it will undo the nasty effects of the overlap of windows and the shape of your window-weighting function. You do not need to change your window weighting function when you increase or decrease the amount of overlap between windows because \mathbf{W}_{wall} takes care of it. The method is to use adjoint `patch` on page 3 to add the weights of each window onto the wall and finally to invert the sum wherever it is non-zero. (You lose data wherever the sum is zero).

No matrices are needed to show that this method succeeds, because data values are never mixed with one another. An equation for any reconstructed data value \tilde{d} as a function of the original value d and the weights w_i that hit d is $\tilde{d} = (\sum_i w_i d) / \sum_i w_i = d$. Thus, our process is simply a “partition of unity.”

To demonstrate the program, I made a random weighting function to use in each window with positive random numbers. The general strategy allows us to use different weights in different windows. That flexibility adds clutter, however, so here we simply use the same weighting function in each window.

The operator $\tilde{\mathbf{I}}_p$ is called “idempotent.” The word “idempotent” means “self-power,” because for any N , $0^N = 0$ and $1^N = 1$, thus the numbers 0 and 1 share the property that raised to any power they remain themselves. Likewise, the patching reconstruction operator multiplies every data value by either one or zero. Figure 3 shows the result obtained when a plane of identical constant values \mathbf{d} is passed into the patching reconstruction operator $\tilde{\mathbf{I}}_p$. The result is constant on the 2-axis, which confirms that there is adequate sampling on the 2-axis, and although the weighting function is made of random numbers, all trace of random numbers has disappeared from the output. On the 1-axis the output is constant, except for being zero in gaps, because the windows do not overlap on the 1-axis.

Figure 8.3: A plane of identical values passed through the idempotent patching reconstruction operator. Results are shown for the same parameters as Figure 2. `pch/patch idempatch`



Module `patching` assists in reusing the patching technique. It takes a linear operator \mathbf{F} . as its argument and applies it in patches. Mathematically, this is $[\mathbf{W}_{\text{wall}} \mathbf{P}' \mathbf{W}_{\text{wind}} \mathbf{F} \mathbf{P}] \mathbf{d}$. It is assumed that the input and output sizes for the operator `oper` are equal.

user/gee/mkwallwt.c

```

25 void mkwallwt(int dim      /* number of dimensions */,
26               int* npatch /* number of patches [dim] */,
27               int* nwall  /* data size [dim] */,
28               int* nwind  /* patch size [dim] */,
29               float* windwt /* window weighting (input) */,
30               float* wallwt /* wall weighting (output) */)
31 /*< make wall weight >*/
32 {
33     int i, j, ip, np, n, nw;
34
35     np = 1;
36     n = 1;
37     nw = 1;
38
39     for (j=0; j < dim; j++) {
40         np *= npatch[j];
41         n *= nwall[j];
42         nw *= nwind[j];
43     }
44
45     for (i = 0; i < n; i++) {
46         wallwt[i] = 0.;
47     }
48
49     patch_init(dim, npatch, nwall, nwind);
50
51     for (ip=0; ip < np; ip++) {
52         patch_lop(true, true, n, nw, wallwt, windwt);
53         patch_close ();
54     }
55
56     for (i = 0; i < n; i++) {
57         if ( wallwt[i] != 0.) wallwt[i] = 1. / wallwt[i];
58     }
59 }

```

```

user/gee/patching.c
51  for (i=0; i < n; i++) data[i] = 0.;
52
53  patch_init(dim, npatch, nwall, nwind);
54  for (ip = 0; ip < np; ip++) {
55      /* modl -> winmodl */
56      patch_lop(false, false, n, nw, modl, winmodl);
57      /* winmodl -> windata */
58      oper(false, false, nw, nw, winmodl, windata);
59      /* apply window weighting */
60      for (iw=0; iw < nw; iw++) windata[iw] *= windwt[iw];
61      /* data <- windata */
62      patch_lop(true, true, n, nw, data, windata);
63      patch_close();
64  }
65
66  /* windwt -> wallwt */
67  mkwallwt(dim, npatch, nwall, nwind, windwt, wallwt);
68
69  /* apply wall weighting */
70  for (i=0; i < n; i++) data[i] *= wallwt[i];

```

8.1.2 2-D filtering in patches


A way to do time- and space-variable filtering is to do invariant filtering within each patch. Typically, we apply a filter, say \mathbf{F}_p , in each patch. The composite operator, filtering in patches, $\tilde{\mathbf{F}}$, is given by

$$\tilde{\mathbf{d}} = [\mathbf{W}_{\text{wall}} \mathbf{P}' \mathbf{W}_{\text{wind}} \mathbf{F}_p \mathbf{P}] \mathbf{d} = \tilde{\mathbf{F}} \mathbf{d} \quad (8.2)$$

I built a triangular weighting routine `tentn()` that tapers from the center of the patch of the filter's *outputs* towards the edges. Accomplishing this weighting is complicated by (1) the constraint that the filter must not move off the edge of the input patch and (2) the alignment of the input and the output. The layout for prediction-error filters is shown in Figure 4. We need a weighting function that vanishes where the filter has no outputs.

Figure 8.4: Domain of inputs and outputs of a two-dimensional prediction-error filter.


pch/XFig rabdomain



The amplitude of the weighting function is not very important because we have learned how to put signals back together properly for arbitrary weighting functions. We can use any

Figure 8.5: Placement of tent-like weighting function in the space of filter inputs and outputs.

pch/XFig rabtent



pyramidal or tent-like shape that drops to zero outside the domain of the filter output. The job is done by subroutine `tentn()`. A new parameter needed by `tentn` is `a`, the coordinate of the beginning of the tent.

user/gee/tent.c

```


45  /* loop in the window */
46  for (i=0; i < nw; i++) {
47      sf_line2cart(dim, nwind, i, x);
48
49      windwt[i] = 1.;
50      for (j=0; j < dim; j++) {
51          if (x[j] >= start[j] && x[j] <= end[j]) {
52              w = (x[j]-mid[j])/wid[j];
53              windwt[i] *= SF_MAX(0., 1.-fabs(w));
54          } else {
55              windwt[i] = 0.;
56          }
57      }
58  }

```

In applications where triangle weights are needed on the *inputs* (or where we can work on a patch without having interference with edges), we can get “triangle tent” weights from `tentn()` if we set filter dimensions and lags to unity, as shown in Figure ??.

Figure 8.6: Window weights from `tentn()` with `nwind=(61,19)`, `center=(31,1)`, `a=(1,1)`.


pch/patch wind1wt



Triangle weighting functions can sum to a constant if the spacing is such that the midpoint of one triangle is at the beginning of the next. I imagined in two dimensions that something similar would happen with shapes like Egyptian pyramids of Cheops, $2 - |x - y| + |x + y|$. Instead, the equation $(1 - |x|)(1 - |y|)$ which has the tent-like shape shown in Figure ?? adds up to the constant flat top shown in Figure ??. (To add interest to


Figure ??, I separated the windows by a little more than the precise matching distance.) In practice we may chose window shapes and overlaps for reasons other than the constancy of the sum of weights, because `mkwallwt` on page 5 accounts for that.

Figure 8.7: (Inverse) wall weights with `n1=100, w1=61, k1=2, n2=30, w2=19, k2=2`
`pch/patch wall1wt`



Finally is an example of filtering a plane of uniform constants with an impulse function. The impulse function is surrounded by zeros, so the filter output patches are smaller than the input patches back in Figure 3. Here in Figure 8, both axes need more window density.

Figure 8.8: Filtering in patches Mid with the same parameters as in Figures 2 and 3. Additionally, the filter parameters are `a1=11 a2=5 lag1=6 lag2=1`. Thus, windows are centered on the 1-axis and pushed back out the 2-axis. `pch/patch cinloip`



8.1.3 Designing a separate filter for each patch

Recall the prediction-error filter subroutine `find_pef()` on page 160. Given a data plane, this subroutine finds a filter that tends to whiten the spectrum of that data plane. The output is white residual. Now suppose we have a data plane where the dip spectrum is changing from place to place. Here it is natural to apply subroutine `find_pef()` in local patches. This is done by subroutine `find_lopef()`. The output of this subroutine is an array of helix-type filters, which can be used, for example, in a local convolution operator `loconvol`. We notice that when a patch has fewer regression equations than the filter has coefficients, then the filter is taken to be that of the previous patch.

8.1.4 Triangular patches

I have been running patching code for several years and my first general comment is that realistic applications often call for patches of different sizes and different shapes. (Tutorial, non-interactive C code is poorly suited to this need.) Raw seismic data in particular seems more suited to triangular shapes. It is worth noting that the basic concepts in this chapter have ready extension to other shapes. For example, a rectangular shape could be duplicated into two identical patches; then data in one could be zeroed above the diagonal and in the other below; you would have to allow, of course, for overlap the size of the filter. Module `pef` on page 160 automatically ignores the zeroed portion of the triangle, and it is irrelevant

user/gee/lopec.c

```

58  patch_init(dim, npatch, nwall, nwind);
59  for (ip=0; ip < np; ip++) {
60      bb = aa+ip;
61
62      patch_lop(false, false, n, nw, wall, windata);
63      if (NULL != mask) {
64          patch_lop(false, false, n, nw, mask, winmask);
65          for (iw=0; iw < nw; iw++) {
66              known[iw] = (winmask[iw] != 0.);
67          }
68          find_mask(nw, known, bb);
69      }
70      for (mis=iw=0; iw < nw; iw++) {
71          if (!bb->mis[iw]) mis++;
72      }
73      if (mis > nh) { /* enough equations */
74          find_pef(nw, windata, bb, nh);
75      } else if (ip > 1) { /* use last PEF */
76          for (ih=0; ih < nh; ih++) {
77              bb->flt[ih] = (bb-1)->flt[ih];
78          }
79      }
80      patch_close();
81  }

```

user/gee/loconvol.c

```

27  void loconvol_init(sf_filter aa_in)
28  /*< initialize with the first filter >*/
29  {
30      aa = aa_in;
31  }
32
33  void loconvol_lop(bool adj, bool add, int nx, int ny,
34                  float *xx, float *yy)
35  /*< convolve >*/
36  {
37      sf_helicon_init(aa);
38      aa++;
39
40      sf_helicon_lop(adj, add, nx, ny, xx, yy);
41  }

```

what `mis2()` on page 133 does with a zeroed portion of data, if a triangular footprint of weights is designed to ignore its output.

EXERCISES:

- 1 Code the linear operator $\mathbf{W}_{\text{wall}}\mathbf{P}'\mathbf{W}_{\text{wind}}\mathbf{P}$ including its adjoint.
- 2 **Smoothing program.** Some familiar operations can be seen in a new light when done in patches. Patch the data. In each patch, find the mean value. Replace each value by the mean value. Reconstruct the wall.
- 3 **Smoothing while filling missing data.** This is like smoothing, but you set window weights to zero where there is no data. Because there will be a different set of weights in each window, you will need to make a simple generalization to `mkwallwt` on page 5.
- 4 **Gain control.** Divide the data into patches. Compute the square root of the sum of the squares of all the data in each patch. Divide all values in that patch by this amount. Reassemble patches.

8.2 STEEP-DIP DECON

Normally, when an autoregression filter (PEF) predicts a value at a point it uses values at earlier points. In practice, a gap may also be set between the predicted value and the earlier values. What is not normally done is to supplement the fitting signals on nearby traces. That is what we do here. We allow the prediction of a signal to include nearby signals at earlier times. The times accepted in the goal are inside a triangle of velocity less than about the water velocity. The new information allowed in the prediction is extremely valuable for water-velocity events. Wavefronts are especially predictable when we can view them along the wavefront (compared to perpendicular or at some other angle from the wavefront). It is even better on land, where noises move more slowly at irregular velocities, and are more likely to be aliased.

Using `lopef` on page 9, the overall process proceeds independently in each of many overlapping windows. The most important practical aspect is the filter masks, described next.

8.2.1 Dip rejecting known-velocity waves

Consider the two-dimensional filter

$$\begin{array}{ccc} & +1 & \\ -1 & 0 & -1 \\ & +1 & \end{array} \quad (8.3)$$

When this filter is applied to a field profile with 4 ms time sampling and 6 m trace

spacing, it should perfectly extinguish 1.5 km/s water-velocity noises. Likewise, the filter

$$\begin{array}{ccccc}
 & & +1 & & \\
 & & 0 & & \\
 & & 0 & & \\
 -1 & 0 & -1 & & \\
 & & 0 & & \\
 & & 0 & & \\
 & & +1 & &
 \end{array} \tag{8.4}$$

should perfectly extinguish water noise when the trace spacing is 18 m. Such noise is, of course, spatially aliased for all temporal frequencies above 1/3 of Nyquist, but that does not matter. The filter extinguishes them perfectly anyway. Inevitably, the filter cannot both extinguish the noise and leave the signal untouched where the alias of one is equal to the other. So we expect the signal to be altered where it matches aliased noise. This simple filter does worse than that. On horizontal layers, for example, signal wavelets become filtered by $(1, 0, 0, -2, 0, 0, 1)$. If the noise is overwhelming, this signal distortion is a small price to pay for eliminating it. If the noise is tiny, however, the distortion is unforgivable. In the real world, data-adaptive deconvolution is usually a good compromise.

The two-dimensional deconvolutions filters we explore here look like this:

$$\begin{array}{cccccccccc}
 x & x & x & x & x & x & x & x & x & \\
 x & x & x & x & x & x & x & x & x & \\
 . & x & x & x & x & x & x & x & . & \\
 . & x & x & x & x & x & x & x & . & \\
 . & x & x & x & x & x & x & x & . & \\
 . & . & x & x & x & x & x & . & . & \\
 . & . & x & x & x & x & x & . & . & \\
 . & . & x & x & x & x & x & . & . & \\
 . & . & . & x & x & x & . & . & . & \\
 . & . & . & x & x & x & . & . & . & \\
 . & . & . & x & x & x & . & . & . & \\
 . & . & . & . & . & . & . & . & . & \\
 . & . & . & . & . & . & . & . & . & \\
 . & . & . & . & 1 & . & . & . & . &
 \end{array} \tag{8.5}$$

where each . denotes a zero and each x denotes a (different) adjustable filter coefficient that is chosen to minimize the power out.

You can easily imagine variations on this shape, such as a diamond instead of a triangle. I invite you to experiment with the various shapes that suggest themselves.

8.2.2 Tests of steep-dip decon on field data

Low-velocity noises on shot records are often not fully suppressed by stacking because the noises are spatially aliased. Routine field arrays are not perfect and the noise is often

extremely strong. An interesting, recently-arrived data set worth testing is shown in Figure 9.




Figure 8.9: Gravel plain ground roll (Middle East) Worth testing. pch/mideast gravel2D

I scanned the forty Yilmaz and Cumro shot profiles for strong low-velocity noises and I selected six examples. To each I applied an AGC that is a slow function of time and space (triangle smoothing windows with triangle half-widths of 200 time points and 4 channels). Because my process simultaneously does both low-velocity rejection and deconvolution, I prepared more traditional 1-D deconvolutions for comparison. This is done in windows of 250 time points and 25 channels, the same filter being used for each of the 25 channels in the window. In practice, of course, considerably more thought would be given to optimal window sizes as a function of the regional nature of the data. The windows were overlapped by about 50%. The same windows are used on the steep-dip deconvolution.

It turned out to be much easier than expected and on the first try I got good results on all all six field profiles tested. I have not yet tweaked the many adjustable parameters. As you inspect these deconvolved profiles from different areas of the world with different recording methods, land and marine, think about how the stacks should be improved by the deconvolution. Stanford Exploration Project report 77 (SEP-77) shows the full suite of results. Figure 10 is a sample of them.

Unexpectedly, results showed that 1-D deconvolution also suppresses low-velocity noises. An explanation can be that these noises are often either low-frequency or quasimonochromatic.

As a minor matter, fundamentally, my code cannot work ideally along the side bound-




Figure 8.10: Top is a North African vibrator shot profile (Y&C #10) after AGC. Middle is gapped 1-D decon. Bottom is steep-dip decon. pch/steep wz

aries because there is no output (so I replaced it by the variance scaled input). With a little extra coding, better outputs could be produced along the sides if we used spatially one-sided filters like

$$\begin{array}{ccccc}
 x & x & x & x & x \\
 . & x & x & x & x \\
 . & x & x & x & x \\
 . & . & x & x & x \\
 . & . & x & x & x \\
 . & . & . & x & x \\
 . & . & . & x & x \\
 . & . & . & . & . \\
 . & . & . & . & . \\
 . & . & . & . & 1
 \end{array} \tag{8.6}$$

These would be applied on one side of the shot and the opposite orientation would be applied on the other side. With many kinds of data sets, such as off-end marine recording in which a ship tows a hydrophone streamer, the above filter might be better in the interior too.

8.2.3 Are field arrays really needed?

Field arrays cancel random noise but their main function, I believe, is to cancel low-velocity coherent noises, something we now see is handled effectively by steep-dip deconvolution. While I do not advocate abandoning field arrays, it is pleasing to notice that with the arrival of steep-dip deconvolution, we are no longer so dependent on field arrays and perhaps coherent noises can be controlled where field arrays are impractical, as in certain 3-D geometries. A recently arrived 3-D shot profile from the sand dunes in the Middle East is Figure 11. The strong hyperbolas are ground roll seen in a line that does not include the shot. The open question here is, how should we formulate the problem of ground-roll removal in 3-D?

8.2.4 Which coefficients are really needed?

Steep-dip decon is a heavy consumer of computer time. Many small optimizations could be done, but more importantly, I feel there are some deeper issues that warrant further investigation. The first question is, how many filter coefficients should there be and where should they be? We would like to keep the number of nonzero filter coefficients to a minimum because it would speed the computation, but more importantly I fear the filter output might be defective in some insidious way (perhaps missing primaries) when too many filter coefficients are used. Perhaps if 1-D decon were done sequentially with steep-dip decon the number of free parameters (and hence the amount of computer time) could be dropped even further. I looked at some of the filters and they scatter wildly with the Nyquist frequency (particularly those coefficients on the trace with the “1” constraint). This suggests using a damping term on the filter coefficients, after which perhaps the magnitude of a filter coefficient will be a better measure of whether this practice is really helpful. Also, it would, of course, be fun to get some complete data sets (rather than a single shot profile) to see the difference in the final stack.




Figure 8.11: Sand dunes. One shot, six parallel receiver lines. pch/mideast dune3D

8.3 INVERSION AND NOISE REMOVAL

Here we relate the basic theoretical statement of geophysical inverse theory to the basic theoretical statement of separation of signals from noises.

A common form of linearized geophysical inverse theory is

$$\mathbf{0} \approx \mathbf{W}(\mathbf{Lm} - \mathbf{d}) \quad (8.7)$$

$$\mathbf{0} \approx \epsilon \mathbf{A} \mathbf{m} \quad (8.8)$$

We choose the operator $\mathbf{L} = \mathbf{I}$ to be an identity and we rename the model \mathbf{m} to be signal \mathbf{s} . Define noise by the decomposition of data into signal plus noise, so $\mathbf{n} = \mathbf{d} - \mathbf{s}$. Finally, let us rename the weighting (and filtering) operations $\mathbf{W} = \mathbf{N}$ on the noise and $\mathbf{A} = \mathbf{S}$ on the signal. Thus the usual model fitting becomes a fitting for signal-noise separation:

$$\mathbf{0} \approx \mathbf{N}(-\mathbf{n}) = \mathbf{N}(\mathbf{s} - \mathbf{d}) \quad (8.9)$$

$$\mathbf{0} \approx \epsilon \mathbf{S} \mathbf{s} \quad (8.10)$$

8.4 SIGNAL-NOISE DECOMPOSITION BY DIP

Choose noise \mathbf{n} to be energy that has no spatial correlation and signal \mathbf{s} to be energy with spatial correlation consistent with one, two, or possibly a few plane-wave segments.

(Another view of noise is that a huge number of plane waves is required to define the wavefield; in other words, with Fourier analysis you can make anything, signal or noise.) We know that a first-order differential equation can absorb (kill) a single plane wave, a second-order equation can absorb one or two plane waves, etc. In practice, we will choose the order of the wavefield and minimize power to absorb all we can, and call that the signal.

\mathbf{S} is the operator that absorbs (by prediction error) the plane waves and \mathbf{N} absorbs noises and $\epsilon > 0$ is a small scalar to be chosen. The difference between \mathbf{S} and \mathbf{N} is the spatial order of the filters. Because we regard the noise as spatially uncorrelated, \mathbf{N} has coefficients only on the time axis. Coefficients for \mathbf{S} are distributed over time and space. They have one space level, plus another level for each plane-wave segment slope that we deem to be locally present. In the examples here the number of slopes is taken to be two. Where a data field seems to require more than two slopes, it usually means the “patch” could be made smaller.

It would be nice if we could forget about the goal (10) but without it the goal (9), would simply set the signal \mathbf{s} equal to the data \mathbf{d} . Choosing the value of ϵ will determine in some way the amount of data energy partitioned into each. The last thing we will do is choose the value of ϵ , and if we do not find a theory for it, we will experiment.

The operators \mathbf{S} and \mathbf{N} can be thought of as “leveling” operators. The method of least-squares sees mainly big things, and spectral zeros in \mathbf{S} and \mathbf{N} tend to cancel spectral lines and plane waves in \mathbf{s} and \mathbf{n} . (Here we assume that power levels remain fairly level in time. Were power levels to fluctuate in time, the operators \mathbf{S} and \mathbf{N} should be designed to level them out too.)

None of this is new or exciting in one dimension, but I find it exciting in more dimensions. In seismology, quasisinusoidal signals and noises are quite rare, whereas local plane waves are abundant. Just as a short one-dimensional filter can absorb a sinusoid of any frequency, a compact two-dimensional filter can absorb a wavefront of any dip.

To review basic concepts, suppose we are in the one-dimensional frequency domain. Then the solution to the fitting goals (10) and (9) amounts to minimizing a quadratic form by setting to zero its derivative, say

$$0 = \frac{\partial}{\partial \mathbf{s}'} ((\mathbf{s}' - \mathbf{d}')\mathbf{N}'\mathbf{N}(\mathbf{s} - \mathbf{d}) + \epsilon^2 \mathbf{s}'\mathbf{S}'\mathbf{S}\mathbf{s}) \quad (8.11)$$

which gives the answer

$$\mathbf{s} = \left(\frac{\mathbf{N}'\mathbf{N}}{\mathbf{N}'\mathbf{N} + \epsilon^2 \mathbf{S}'\mathbf{S}} \right) \mathbf{d} \quad (8.12)$$

$$\mathbf{n} = \mathbf{d} - \mathbf{s} = \left(\frac{\epsilon^2 \mathbf{S}'\mathbf{S}}{\mathbf{N}'\mathbf{N} + \epsilon^2 \mathbf{S}'\mathbf{S}} \right) \mathbf{d} \quad (8.13)$$

To make this really concrete, consider its meaning in one dimension, where signal is white $\mathbf{S}'\mathbf{S} = 1$ and noise has the frequency ω_0 , which is killable with the multiplier $\mathbf{N}'\mathbf{N} = (\omega - \omega_0)^2$. Now we recognize that equation (12) is a notch filter and equation (13) is a narrow-band filter.

The analytic solutions in equations (12) and (13) are valid in 2-D Fourier space or dip space too. I prefer to compute them in the time and space domain to give me tighter control

on window boundaries, but the Fourier solutions give insight and offer a computational speed advantage.

Let us express the fitting goal in the form needed in computation.

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix} \approx \begin{bmatrix} \mathbf{N} \\ \epsilon \mathbf{S} \end{bmatrix} \mathbf{s} + \begin{bmatrix} -\mathbf{N} \mathbf{d} \\ \mathbf{0} \end{bmatrix} \quad (8.14)$$

user/gee/signoi.c

```

46 void signoi_lop (bool adj, bool add, int n1, int n2,
47                float *data, float *sign)
48 /*< linear operator >*/
49 {
50     sf_helicon_init (nn);
51     polydiv_init (nd, ss);
52
53     sf_adjnull (adj, add, n1, n2, data, sign);
54
55     sf_helicon_lop (false, false, n1, n1, data, dd);
56     sf_solver_prec (sf_helicon_lop, sf_cgstep, polydiv_lop,
57                   nd, nd, nd, sign, dd, niter, eps, "end");
58     sf_cgstep_close ();
59
60     nn++;
61     ss++;
62 }

```

As with the missing-data subroutines, the potential number of iterations is large, because the dimensionality of the space of unknowns is much larger than the number of iterations we would find acceptable. Thus, sometimes changing the number of iterations **niter** can create a larger change than changing **epsilon**. Experience shows that helix preconditioning saves the day.

8.4.1 Signal/noise decomposition examples

Figure 12 demonstrates the signal/noise decomposition concept on synthetic data. The signal and noise have similar frequency spectra but different dip spectra.

Before I discovered helix preconditioning, Ray Abma found that different results were obtained when the fitting goal was cast in terms of **n** instead of **s**. Theoretically it should not make any difference. Now I believe that with preconditioning, or even without it, if there are enough iterations, the solution should be independent of whether the fitting goal is cast with either **n** or **s**.

Figure 13 shows the result of experimenting with the choice of ϵ . As expected, increasing ϵ weakens **s** and increases **n**. When ϵ is too small, the noise is small and the signal is almost the original data. When ϵ is too large, the signal is small and coherent events are pushed into the noise. (Figure 13 rescales both signal and noise images for the clearest display.)




Figure 8.12: The input signal is on the left. Next is that signal with noise added. Next, for my favorite value of `epsilon=1.`, is the estimated signal and the estimated noise.
 pch/signoi signoi




Figure 8.13: Left is an estimated signal-noise pair where `epsilon=4` has improved the appearance of the estimated signal but some coherent events have been pushed into the noise. Right is a signal-noise pair where `epsilon=.25`, has improved the appearance of the estimated noise but the estimated signal looks no better than original data.
 pch/signoi signeps

Notice that the leveling operators \mathbf{S} and \mathbf{N} were both estimated from the original signal and noise mixture $\mathbf{d} = \mathbf{s} + \mathbf{n}$ shown in Figure 12. Presumably we could do even better if we were to reestimate \mathbf{S} and \mathbf{N} from the estimates \mathbf{s} and \mathbf{n} in Figure 13.

8.4.2 Spitz for variable covariances

Since signal and noise are uncorrelated, the spectrum of data is the spectrum of the signal plus that of the noise. An equation for this idea is

$$\sigma_d^2 = \sigma_s^2 + \sigma_n^2 \quad (8.15)$$

This says resonances in the signal and resonances in the noise will both be found in the data. When we are given σ_d^2 and σ_n^2 it seems a simple matter to subtract to get σ_s^2 . Actually it can be very tricky. We are never given σ_d^2 and σ_n^2 ; we must estimate them. Further, they can be a function of frequency, wave number, or dip, and these can be changing during measurements. We could easily find ourselves with a negative estimate for σ_s^2 which would ruin any attempt to segregate signal from noise. An idea of Simon Spitz can help here.

Let us reexpress equation (15) with prediction-error filters.

$$\frac{1}{\bar{A}_d A_d} = \frac{1}{\bar{A}_s A_s} + \frac{1}{\bar{A}_n A_n} = \frac{\bar{A}_s A_s + \bar{A}_n A_n}{(\bar{A}_s A_s)(\bar{A}_n A_n)} \quad (8.16)$$

Inverting

$$\bar{A}_d A_d = \frac{(\bar{A}_s A_s)(\bar{A}_n A_n)}{\bar{A}_s A_s + \bar{A}_n A_n} \quad (8.17)$$

The essential feature of a PEF is its zeros. Where a PEF approaches zero, its inverse is large and resonating. When we are concerned with the zeros of a mathematical function we tend to focus on numerators and ignore denominators. The zeros in $\bar{A}_s A_s$ compound with the zeros in $\bar{A}_n A_n$ to make the zeros in $\bar{A}_d A_d$. This motivates the “Spitz Approximation.”

$$\bar{A}_d A_d = (\bar{A}_s A_s)(\bar{A}_n A_n) \quad (8.18)$$

It usually happens that we can find a patch of data where no signal is present. That’s a good place to estimate the noise PEF A_n . It is usually much harder to find a patch of data where no noise is present. This motivates the Spitz approximation which by saying $A_d = A_s A_n$ tells us that the hard-to-estimate A_s is the ratio $A_s = A_d/A_n$ of two easy-to-estimate PEFs.

It would be computationally convenient if we had A_s expressed not as a ratio. For this, form the signal $\mathbf{u} = \mathbf{A}_n \mathbf{d}$ by applying the noise PEF A_n to the data \mathbf{d} . The spectral relation is

$$\sigma_u^2 = \sigma_d^2 / \sigma_n^2 \quad (8.19)$$

Inverting this expression and using the Spitz approximation we see that a PEF estimate on \mathbf{u} is the required A_s in numerator form because

$$A_u = A_d / A_n = A_s \quad (8.20)$$

8.4.3 Noise removal on Shearer's data

Professor Peter Shearer¹ gathered the earthquakes from the IDA network, an array of about 25 widely distributed gravimeters, donated by Cecil Green, and Shearer selected most of the shallow-depth earthquakes of magnitude greater than about 6 over the 1981-91 time interval, and sorted them by epicentral distance into bins 1° wide and stacked them. He generously shared his edited data with me and I have been restacking it, compensating for amplitude in various ways, and planning time and filtering compensations.

Figure 14 shows a test of noise subtraction by multidip narrow-pass filtering on the Shearer-IDA stack. As with prediction there is a general reduction of the noise. Unlike with prediction, weak events are preserved and noise is subtracted from them too.

Besides the difference in theory, the separation filters are much smaller because their size is determined by the concept that “two dips will fit anything locally” ($a_2=3$), versus the prediction filters “needing a sizeable window to do statistical averaging.” The same aspect ratio a_1/a_2 is kept and the page is now divided into 11 vertical patches and 24 horizontal patches (whereas previously the page was divided in 3×4 patches). In both cases the patches overlap about 50%. In both cases I chose to have about ten times as many equations as unknowns on each axis in the estimation. The ten degrees of freedom could be distributed differently along the two axes, but I saw no reason to do so.

8.4.4 The human eye as a dip filter

Although the filter seems to be performing as anticipated, no new events are apparent. I believe the reason that we see no new events is that the competition is too tough. We are competing with the human eye, which through aeons of survival has become is a highly skilled filter. Does this mean that there is no need for filter theory and filter subroutines because the eye can do it equally well? It would seem so. Why then pursue the subject matter of this book?

The answer is 3-D. The human eye is not a perfect filter. It has a limited (though impressive) dynamic range. A nonlinear display (such as wiggle traces) can prevent it from averaging. The eye is particularly good at dip filtering, because the paper can be looked at from a range of grazing angles and averaging window sizes miraculously adjust to the circumstances. The eye can be overwhelmed by too much data. The real problem with the human eye is that the retina is only two-dimensional. The world contains many three-dimensional data volumes. I don't mean the simple kind of 3-D where the contents of the room are nicely mapped onto your 2-D retina. I mean the kind of 3-D found inside a bowl of soup or inside a rock. A rock can be sliced and sliced and sliced again and each slice is a picture. The totality of these slices is a movie. The eye has a limited ability to deal with movies by optical persistence, an averaging of all pictures shown in about 1/10 second interval. Further, the eye can follow a moving object and perform the same averaging. I have learned, however, that the eye really cannot follow two objects at two different speeds

¹ I received the data for this stack from Peter Shearer at the Cecil and Ida Green Institute of Geophysics and Planetary Physics of the Scripps Oceanographic Institute. I also received his permission to redistribute it to friends and colleagues. Should you have occasion to copy it please reference ?? it properly. Examples of earlier versions of these stacks are found in the references. Professor Shearer may be willing to supply newer and better stacks. His electronic mail address is shearer@mahi.ucsd.edu.




Figure 8.14: Stack of Shearer's IDA data (left). Multidip filtered (right). It is pleasing that the noise is reduced while weak events are preserved. pch/ida sneps

and average them both over time. Now think of the third dimension in Figure 14. It is the dimension that I summed over to make the figure. It is the 1° range bin. If we were viewing the many earthquakes in each bin, we would no longer be able to see the out-of-plane information which is the in-plane information in Figure 14.

To view genuinely 3-D information we must see a movie, or we must compress the 3-D to 2-D. There are only a small number of ways to compress 3-D to 2-D. One is to select planes from the volume. One is to sum the volume over one of its axes, and the other is a compromise, a filtering over the axis we wish to abandon before subsampling on it. That filtering is a local smoothing. If the local smoothing has motion (out of plane dip) of various velocities (various dips), then the desired process of smoothing the out of plane direction is what we did in the in-plane direction in Figure 14. But Figure 14 amounts to more than that. It amounts to a kind of simultaneous smoothing in the *two* most coherent directions whereas in 3-D your eye can smooth in only *one* direction when you turn your head along with the motion.

If the purpose of data processing is to collapse 3-D data volumes to 2-D where they are comprehensible to the human eye, then perhaps data-slope adaptive, low-pass filtering in the out-of-plane direction is the best process we can invent.

My purpose in filtering the earthquake stacks is to form a guiding “pilot trace” to the analysis of the traces *within* the bin. Within each bin, each trace needs small time shifts and perhaps a small temporal filter to best compensate it to . . . to what? to the pilot trace, which in these figures was simply a stack of traces in the bin. Now that we have filtered in the range direction, however, the next stack can be made with a better quality pilot.

8.5 SPACE-VARIABLE DECONVOLUTION

Filters sometimes change with time and space. We sometimes observe signals whose spectrum changes with position. A filter that changes with position is called nonstationary. We need an extension of our usual convolution operator `hconv` on page 155. Conceptually, little needs to be changed besides changing `aa(ia)` to `aa(ia,iy)`. But there is a practical problem. Fomel and I have made the decision to clutter up the code somewhat to save a great deal of memory. This should be important to people interested in solving multidimensional problems with big data sets.

Normally, the number of filter coefficients is many fewer than the number of data points, but here we have very many more. Indeed, there are `na` times more. Variable filters require `na` times more memory than the data itself. To make the nonstationary helix code more practical, we now require the filters to be constant in patches. The data type for nonstationary filters (which are constant in patches) is introduced in module `nhelix`, which is a simple modification of module `helix` on page 7. What is new is the integer valued vector `pch(nd)` the size of the one-dimensional (helix) output data space. Every filter output point is to be assigned to a patch. All filters of a given patch number will be the same filter. Nonstationary helices are created with `createnhelix`, which is a simple modification of module `createhelix` on page 28. Notice that the user must define the `pch(product(nd))` vector before creating a nonstationary helix. For a simple 1-D time-variable filter, presumably

```

                                user/gee/nhelix.c
10 typedef struct nhelixfilter {
11     int np;
12     sf_filter* hlx;
13     bool* mis;
14     int *pch;
15 } *nfilter;

```

`pch` would be something like $(1, 1, 2, 2, 3, 3, \dots)$. For multidimensional patching we need to think a little more.

Finally, we are ready for the convolution operator. The operator `nhconest` on page 25 allows for a different filter in each patch. A filter output `y[iy]` has its filter from the patch `ip=aa->pch[iy]`.

Because of the massive increase in the number of filter coefficients, allowing these many filters takes us from overdetermined to very undetermined. We can estimate all these filter coefficients by the usual deconvolution fitting goal (6.18)

$$\mathbf{0} \approx \mathbf{r} = \mathbf{YK}\mathbf{a} + \mathbf{r}_0 \quad (8.21)$$

but we need to supplement it with some damping goals, say

$$\begin{aligned} \mathbf{0} &\approx \mathbf{YK}\mathbf{a} + \mathbf{r}_0 \\ \mathbf{0} &\approx \epsilon \mathbf{R}\mathbf{a} \end{aligned} \quad (8.22)$$

where \mathbf{R} is a roughening operator to be chosen.

Experience with missing data in Chapter 3 shows that when the roughening operator \mathbf{R} is a differential operator, the number of iterations can be large. We can speed the calculation immensely by “preconditioning”. Define a new variable \mathbf{m} by $\mathbf{a} = \mathbf{R}^{-1}\mathbf{m}$ and insert it into (22) to get the equivalent preconditioned system of goals.

$$\mathbf{0} \approx \mathbf{YKR}^{-1}\mathbf{m} \quad (8.23)$$

$$\mathbf{0} \approx \epsilon \mathbf{m} \quad (8.24)$$

The fitting (23) uses the operator \mathbf{YKR}^{-1} . For \mathbf{Y} we can use subroutine `nhconest()` on page 25; for the smoothing operator \mathbf{R}^{-1} we can use nonstationary polynomial division with operator `npolydiv()`:

Now we have all the pieces we need. As we previously estimated stationary filters with the module `pef` on page 160, now we can estimate nonstationary PEFs with the module `npof` on page 27. The steps are hardly any different. Near the end of module `npof` is a filter `reshape` from a 1-D array to a 2-D array.

Figure 15 shows a synthetic data example using these programs. As we hope for deconvolution, events are compressed. The compression is fairly good, even though each event has a different spectrum. What is especially pleasing is that satisfactory results are obtained

user/gee/createnhelix.c

```

29 nfilter createnhelix(int dim      /* number of dimensions */,
30                      int *nd      /* data size [dim] */,
31                      int *center  /* filter center [dim] */,
32                      int *gap     /* filter gap [dim] */,
33                      int *na      /* filter size [dim] */,
34                      int *pch     /* patching [product(nd)] */)
35 /*< allocate and output a non-stationary filter >*/
36 {
37     nfilter nsaa;
38     sf_filter aa;
39     int n123, np, ip, *nh, i;
40
41     aa = createhelix(dim, nd, center, gap, na);
42
43     n123=1;
44     for (i=0; i < dim; i++) {
45         n123 *= nd[i];
46     }
47     np = pch[0];
48     for (i=0; i < n123; i++) {
49         if (pch[i] > np) np=pch[i];
50     }
51     np++;
52
53     nh = sf_intalloc(np);
54     for (ip=0; ip < np; ip++) {
55         nh[ip] = aa->nh;
56     }
57     nsaa = nallocate(np, n123, nh, pch);
58     for (ip=0; ip < np; ip++) {
59         for (i=0; i < aa->nh; i++) {
60             nsaa->hlx[ip]->lag[i] = aa->lag[i];
61         }
62         nbound(ip, dim, nd, na, nsaa);
63     }
64
65     sf_deallocatehelix(aa);
66
67     return nsaa;
68 }

```

```

user/gee/nhconest.c
51  for (iy=0; iy < ny; iy++) {
52      if (aa->mis[iy]) continue;
53
54      ip = aa->pch[iy];
55      lag = aa->hlx[ip]->lag;
56      na = aa->hlx[ip]->nh;
57
58      for (ia=0; ia < na; ia++) {
59          ix = iy - lag[ia];
60          if (ix < 0) continue;
61
62          if (adj) {
63              a[ia+nhmax*ip] += y[iy] * x[ix];
64          } else {
65              y[iy] += a[ia+nhmax*ip] * x[ix];
66          }
67      }
68  }

```




Figure 8.15: Time variable deconvolution with two free filter coefficients and a gap of 6.

pch/tvdecon tvdecon

user/gee/npolydiv.c

```

56     for (id=0; id < nd; id++) {
57         tt[id] = adj? yy[id]: xx[id];
58     }
59
60     if (adj) {
61         for (iy=nd-1; iy >= 0; iy--) {
62             ip = aa->pch[iy];
63             lag = aa->hlx[ip]->lag;
64             flt = aa->hlx[ip]->flt;
65             na = aa->hlx[ip]->nh;
66             for (ia=0; ia < na; ia++) {
67                 ix = iy - lag[ia];
68                 if (ix < 0) continue;
69                 tt[ix] -= flt[ia] * tt[iy];
70             }
71         }
72         for (id=0; id < nd; id++) {
73             xx[id] += tt[id];
74         }
75     } else {
76         for (iy=0; iy < nd; iy++) {
77             ip = aa->pch[iy];
78             lag = aa->hlx[ip]->lag;
79             flt = aa->hlx[ip]->flt;
80             na = aa->hlx[ip]->nh;
81             for (ia=0; ia < na; ia++) {
82                 ix = iy - lag[ia];
83                 if (ix < 0) continue;
84                 tt[iy] -= flt[ia] * tt[ix];
85             }
86         }
87         for (id=0; id < nd; id++) {
88             yy[id] += tt[id];
89         }
90     }

```



```

                                user/gee/npef.c
29 void find_pef(int nd      /* data size */,
30               float *dd   /* data */,
31               nfilter aa  /* estimated filter */,
32               nfilter rr  /* regularization filter */,
33               int niter   /* number of iterations */,
34               float eps   /* regularization parameter */,
35               int nh      /* filter size */)
36 /*< estimate non-stationary PEF >*/
37 {
38     int ip, ih, na, np, nr;
39     float *flt;
40
41     np = aa->np;
42     nr = np*nh;
43     flt = sf_floatalloc(nr);
44
45     nhconest_init(dd, aa, nh);
46     npolydiv2_init(nr, rr);
47
48     sf_solver_prec(nhconest_lop, sf_cgstep, npolydiv2_lop,
49                    nr, nr, nd, flt, dd, niter, eps, "end");
50     sf_cgstep_close();
51     npolydiv2_close();
52
53     for (ip=0; ip < np; ip++) {
54         na = aa->hlx[ip]->nh;
55         for (ih=0; ih < na; ih++) {
56             aa->hlx[ip]->flt[ih] = -flt[ip*nh + ih];
57         }
58     }
59
60     free(flt);
61 }

```

in truly small numbers of iterations (about three). The example is for two free filter coefficients $(1, a_1, a_2)$ per output point. The roughening operator \mathbf{R} was taken to be $(1, -2, 1)$ which was factored into causal and anticausal finite difference.

I hope also to find a test case with field data, but experience in seismology is that spectral changes are slow, which implies unexciting results. Many interesting examples should exist in two- and three-dimensional filtering, however, because reflector dip is always changing and that changes the *spatial* spectrum.

In multidimensional space, the smoothing filter \mathbf{R}^{-1} can be chosen with interesting directional properties. Sergey, Bob, Sean and I have joked about this code being the “double helix” program because there are two multidimensional helixes in it, one the smoothing filter, the other the deconvolution filter. Unlike the biological helixes, however, these two helixes do not seem to form a symmetrical pair.

EXERCISES:

- 1 Is `nhconest` on page 25 the inverse operator to `npolydiv` on page 26? Do they commute?
- 2 Sketch the matrix corresponding to operator `nhconest` on page 25. HINTS: Do not try to write all the matrix elements. Instead draw short lines to indicate rows or columns. As a “warm up” consider a simpler case where one filter is used on the first half of the data and another filter for the other half. Then upgrade that solution from two to about ten filters.


Chapter 9

Plane waves in three dimensions

In this chapter we seek a deeper understanding of plane waves in *three* dimensions, where the examples and theory typically refer to functions of time t and two space coordinates (x, y) , or to 3-D migration images where the t coordinate is depth or travelttime depth. As in Chapter , we need to decompose data volumes into subcubes, shown in Figure 9.1.

Figure 9.1: Left is space of inputs and outputs. Right is their separation during analysis.

lmn/XFig rayab3D



In this chapter we will see that the wave model implies the 3-D whitener is not a cube filter but two planar filters. The wave model allows us to determine the scale factor of a signal, even where signals fluctuate in strength because of interference. Finally, we examine the local-monoplane concept that uses the superposition principle to distinguish a sedimentary model cube from a data cube.

9.1 THE LEVELER: A VOLUME OR TWO PLANES?

In *two* dimensions, levelers were taken to be PEFs, small rectangular planes of numbers in which the time axis included enough points to include reasonable stepouts were included and the space axis contained one level plus another space level, for each plane-wave slope supposed to be present.

We saw that a whitening filter in *three* dimensions is a small volume with shape defined by subroutine `createhelix()`. It might seem natural that the number of points on the x - and y -axes be related to the number of plane waves present. Instead, I assert that if the volume contains plane waves, we don't want a *volume* filter to whiten it; we can use a *pair* of *planar* filters to do so and the order of those filters is the number of planes thought to be simultaneously present. I have no firm mathematical proofs, but I offer you

some interesting discussions, examples, and computer tools for you to experiment with. It seems that some applications call for the volume filter while others call for the two planes. Because two planes of numbers generally contain many fewer adjustable values than a volume, statistical-estimation reasons also favor the planes.

What is the lowest-order filter that, when applied to a volume, will destroy one and only one slope of plane wave?

First we seek the answer to the question, “What is the lowest order filter that will destroy one and only one plane?” To begin with, we consider that plane to be horizontal so the volume of numbers is $f(t, x, y) = b(t)$ where $b(t)$ is an arbitrary function of time. One filter that has zero-valued output (destroys the plane) is $\partial_x \equiv \partial/\partial x$. Another is the operator $\partial_y \equiv \partial/\partial y$. Still another is the Laplacian operator which is $\partial_{xx} + \partial_{yy} \equiv \partial^2/\partial x^2 + \partial^2/\partial y^2$.

The problem with $\partial/\partial x$ is that although it destroys the required plane, it also destroys $f(t, x, y) = a(t, y)$ where $a(t, y)$ is an *arbitrary* function of (t, y) such as a cylinder with axis parallel to the x -axis. The operator $\partial/\partial y$ has the same problem but with the axes rotated. The Laplacian operator not only destroys our desired plane, but it also destroys the well known nonplanar function $e^{ax} \cos(ay)$, which is just one example of the many other interesting shapes that constitute solutions to Laplace’s equation.

I remind you of a basic fact: When we set up the fitting goal $\mathbf{0} \approx \mathbf{A}\mathbf{f}$, the quadratic form minimized is $\mathbf{f}'\mathbf{A}'\mathbf{A}\mathbf{f}$, which by differentiation with respect to \mathbf{f}' gives us (in a constraint-free region) $\mathbf{A}'\mathbf{A}\mathbf{f} = \mathbf{0}$. So, minimizing the volume integral (actually the sum) of the squares of the components of the gradient implies that Laplace’s equation is satisfied.

In any volume, the lowest-order filter that will destroy level planes and no other wave slope is a filter that has one input and *two outputs*. That filter is the gradient, $(\partial/\partial x, \partial/\partial y)$. Both outputs vanish if and only if the plane has the proper horizontal orientation. Other objects and functions are not extinguished (except for the non-wave-like function $f(t, x, y) = \text{const}$). It is annoying that we must deal with *two* outputs and that will be the topic of further discussion.

A wavefield of tilted parallel planes is $f(t, x, y) = g(\tau - p_x x - p_y y)$, where $g()$ is an arbitrary one-dimensional function. The operator that destroys these tilted planes is the two-component operator $(\partial_x + p_x \partial_t, \partial_y + p_y \partial_t)$.

The operator that destroys a family of dipping planes

$$f(t, x, y) = g(\tau - p_x x - p_y y)$$

is

$$\left[\begin{array}{c} \frac{\partial}{\partial x} + p_x \frac{\partial}{\partial t} \\ \frac{\partial}{\partial y} + p_y \frac{\partial}{\partial t} \end{array} \right]$$

9.1.1 PEFs overcome spatial aliasing of difference operators

The problem I found with finite-difference representations of differential operators is that they are susceptible to spatial aliasing. Even before they encounter spatial aliasing, they are susceptible to accuracy problems known in finite-difference wave propagation as “frequency dispersion.” The aliasing problem can be avoided by the use of spatial prediction operators such as

$$\begin{array}{c} \cdot \quad a \\ \cdot \quad b \\ 1 \quad c \\ \cdot \quad d \\ \cdot \quad e \end{array} \quad (9.1)$$

where the vertical axis is time; the horizontal axis is space; and the “.”s are zeros. Another possibility is the 2-D whitening filter


$$\begin{array}{c} f \quad a \\ g \quad b \\ 1 \quad c \\ \cdot \quad d \\ \cdot \quad e \end{array} \quad (9.2)$$

Imagine all the coefficients vanished but $d = -1$ and the given 1. Such filters would annihilate the appropriately sloping plane wave. Slopes that are not exact integers are also approximately extinguishable, because the adjustable filter coefficients can interpolate in time. Filters like (9.2) do the operation $\partial_x + p_x \partial_t$, which is a component of the gradient in the plane of the wavefront, *and* they include a temporal deconvolution aspect and a spatial coherency aspect. My experience shows that the operators (9.1) and (9.2) behave significantly differently in practice, and I am not prepared to fully explain the difference, but it seems to be similar to the gapping of one-dimensional filters.

You might find it alarming that your teacher is not fully prepared to explain the difference between a volume and two planes, but please remember that we are talking about the factorization of the volumetric spectrum. Spectral matrices are well known to have structure, but books on theory typically handle them as simply $\lambda \mathbf{I}$. Anyway, wherever you see an \mathbf{A} in a three-dimensional context, you may wonder whether it should be interpreted as a cubic filter that takes one volume to another, or as two planar filters that take one volume to two volumes as shown in Figure 9.2.

Figure 9.2: An inline 2-D PEF and a crossline 2-D PEF both applied throughout the volume. To find each filter, minimize each output power independently.

lmn/XFig rayab3Doper



9.1.2 My view of the world

I start from the idea that the four-dimensional world (t, x, y, z) is filled with expanding spherical waves and with quasispherical waves that result from reflection from quasispherical objects and refraction through quasihomogeneous materials. We rarely, if ever see in an observational data cube, an entire expanding spherical wave, but we normally have a two- or three-dimensional slice with some wavefront curvature. We analyze data subcubes that I call bricks. In any brick we see only local patches of apparent plane waves. I call them platelets. From the microview of this brick, the platelets come from the “great random-point-generator in the sky,” which then somehow convolves the random points with a platelike impulse response. If we could deconvolve these platelets back to their random source points, there would be nothing left inside the brick because the energy would have gone outside. We would have destroyed the energy inside the brick. If the platelets were coin shaped, then the gradient magnitude would convert each coin to its circular rim. The plate sizes and shapes are all different and they damp with distance from their centers, as do Gaussian beams. If we observed *rays* instead of wavefront platelets then we might think of the world as being filled with noodles, and then. . . .

How is it possible that in a small brick we can do something realistic about deconvolving a spheroidal impulse response that is much bigger than the brick? The same way as in one dimension, where in a small time interval we can estimate the correct deconvolution filter of a long resonant signal. A three-point filter destroys a sinusoid.


The inverse filter to the expanding spherical wave might be a huge cube. Good approximations to this inverse at the brick level might be two small planes. Their time extent would be chosen to encompass the slowest waves, and their spatial extent could be two or three points, representing the idea that normally we can listen to only one person at a time, occasionally we can listen to two, and we can never listen to three people talking at the same time.

9.2 WAVE INTERFERENCE AND TRACE SCALING

Although neighboring seismometers tend to show equal powers, the energy on one seismometer can differ greatly from that of a neighbor for both theoretical reasons and practical ones. Should a trace ever be rescaled to give it the same energy as its neighbors? Here we review the strong theoretical arguments against rescaling. In practice, however, especially on land where coupling is irregular, scaling seems a necessity. The question is, what can go wrong if we scale traces to have equal energy, and more basically, where the proper scale factor cannot be recorded, what should we do to get the best scale factor? A related question is how to make good measurements of amplitude versus offset. To understand these issues we review the fundamentals of wave interference.

Theoretically, a scale-factor problem arises because *locally*, wavefields, not energies, add. Nodes on standing waves are familiar from theory, but they could give you the wrong idea that the concept of node is one that applies only with sinusoids. Actually, destructive interference arises anytime a polarity-reversed waveform bounces back and crosses itself. Figure 9.3 shows two waves of opposite polarity crossing each other. Observe that one seismogram has a zero-valued signal, while its neighbors have anomalously higher amplitudes

Figure 9.3: Superposition of plane waves of opposite polarity.
lmn/scale super



and higher energies than are found far away from the interference. The situation shown in Figure 9.3 does not occur easily in nature. Reflection naturally comes to mind, but usually the reflected wave crosses the incident wave at a later time and then they don't extinguish. Approximate extinguishing occurs rather easily when waves are quasi-monochromatic. We will soon see, however, that methodologies for finding scales all begin with deconvolution and that eliminates the monochromatic waves.

9.2.1 Computing the proper scale factor for a seismogram

With data like Figure 9.3, rescaling traces to have equal energy would obviously be wrong. The question is, “How can we determine the proper scale factor?” As we have seen, a superposition of N plane waves exactly satisfies an N -th order (in x) difference equation. Given a 2-D wave field, we can find its PEF by minimizing output power. Then we ask the question, could rescaling the traces give a lower output power? To answer this, we set up an optimization goal: Given the leveler (be it a cubic PEF or two planar ones), find the best trace scales. (After solving this, we could return to re-estimate the leveler, and iterate.) To solve for the scales, we need a subroutine that scales traces and the only tricky part is that the adjoint should bring us back to the space of scale factors. This is done by `scaletrace`. Notice that to estimate scales, the adjoint forms an inner product of the raw data on the

user/gee/scaletrace.c

```

40   for (i=i2=0; i2 < n2; i2++) {
41       for (i1=0; i1 < n1; i1++, i++) {
42           if( adj) scale[i2] += sdata[i] * data[i];
43           else    sdata[i] += scale[i2] * data[i];
44       }
45   }

```

previously scaled data. Let the operator implemented by `scaletrace` be denoted by \mathbf{D} , which is mnemonic for “data” and for “diagonal matrix,” and let the vector of scale factors be denoted by \mathbf{s} and the leveler by \mathbf{A} . Now we consider the fitting goal $\mathbf{0} \approx \mathbf{A}\mathbf{D}\mathbf{s}$. The trouble with this fitting goal is that the solution is obviously $\mathbf{s} = \mathbf{0}$. To avoid the trivial solution $\mathbf{s} = \mathbf{0}$, we can choose from a variety of supplemental fitting goals. One possibility is that for the i -th scale factor we could add the fitting goal $s_i \approx 1$. Another possibility, perhaps better if some of the signals have the opposite of the correct polarity, is that the sum of the scales should be approximately unity. I regret that time has not yet allowed me to identify some interesting examples and work them through.

9.3 LOCAL MONOPLANE ANNIHILATOR

LOMOPLAN (LOcal MONo PLane ANnihilator) is a data-adaptive filter that extinguishes a local monoplane, but cannot extinguish a superposition of several planes. We presume an ideal sedimentary model as made of (possibly curved) parallel layers. Because of the superposition principle, data can be a superposition of several plane waves, but the ideal model should consist locally of only a single plane. Thus, LOMOPLAN extinguishes an ideal model, but not typical data. I conceived of LOMOPLAN as the “ultimate” optimization criterion for inversion problems in reflection seismology (Claerbout, 1992b) but it has not yet demonstrated that it can attain that lofty goal. Instead, however, working in two dimensions, it is useful in data interpretation and in data quality inspection.

The main way we estimate parameters in reflection seismology is that we maximize the coherence of theoretically redundant measurements. Thus, to estimate velocity and statics shifts, we maximize something like the power in the stacked data. Here I propose another optimization criterion for estimating model parameters and missing data. An interpreter looking at a migrated section containing two dips in the same place suspects wave superposition more likely than bedding texture superposition. To minimize the presence of multiple dipping events in the same place, we should use the mono plane annihilator (MOPLAN) filter as the weighting operator for any fitting goal. Because the filter is intended for use on images or migrated data, not on data directly, I call it a *plane* annihilator, not a *planewave* annihilator. (A time-migration or merely a stack, however, might qualify as an image.) We should avoid using the word “wavefront” because waves readily satisfy the superposition principle, whereas images do not, and it is this aspect of images that I advocate and formulate as “prior information.”

An example of a MOPLAN in two dimensions, $(\partial_x + p_x \partial_\tau)$, is explored in Chapter 4 of PVI Claerbout (1992a), where the main goal is to estimate the (τ, x) -variation of p_x . Another family of MOPLANs arise from multidimensional prediction-error filtering described earlier in this book and in PVI, Chapter 8.

Here I hypothesize that a MOPLAN may be a valuable weighting function for many estimation problems in seismology. Perhaps we can estimate statics, interval velocity, and missing data if we use the principle of minimizing the power out of a LOcal MONo PLane ANnihilator (LOMOPLAN) on a migrated section. Thus, those embarrassing semicircles that we have seen for years on our migrated sections may hold one of the keys for unlocking the secrets of statics and lateral velocity variation. I do not claim that this concept is as powerful as our traditional methods. I merely claim that we have not yet exploited this

concept in a systematic way and that it might prove useful where traditional methods break.

For an image model of nonoverlapping curved planes, a suitable choice of weighting function for fitting problems is the local filter that destroys the best fitting local plane.

9.3.1 Mono-plane deconvolution

The coefficients of a 2-D monoplane annihilator filter are defined to be the same as those of a 2-D PEF of spatial order unity; in other words, those defined by either (9.1) or (9.2).

The filter can be lengthened in time but not in space. The choice of exactly two columns is a choice to have an analytic form that can exactly destroy a single plane, but cannot destroy two. Applied to two signals that are statistically independent, the filter (9.2) reduces to the well-known prediction-error filter in the left column and zeros in the right column. If the filter coefficients were extended in both directions on t and to the right on x , the two-dimensional spectrum of the input would be flattened.

9.3.2 Monoplanes in local windows

The earth dip changes rapidly with location. In a small region there is a local dip and dip bandwidth that determines the best LOMOPLAN (LOcal MOPLAN). To see how to cope with the edge effects of filtering in a small region, and to see how to patch together these small regions, recall subroutine `patchn()` on page ?? and the weighting subroutines that work with it.

Figure 9.4 shows a synthetic model that illustrates local variation in bedding. Notice dipping bedding, curved bedding, unconformity between them, and a fault in the curved bedding. Also, notice that the image has its amplitude tapered to zero on the left and right sides. After local monoplane annihilation (LOMOPLAN), the continuous bedding is essentially gone. The fault and unconformity remain.

The local spatial prediction-error filters contain the essence of a factored form of the inverse spectrum of the model.

Because the plane waves are local, the illustrations were made with module `lopef` on page ??.

9.3.3 Crossing dips

Figure 9.5 deserves careful study. The input frame is dipping events with amplitudes slowly changing as they cross the frame. The dip of the events is not commensurate with the mesh, so we use linear interpolation that accounts for the irregularity along an event. The output panel tends to be small where there is only a single dip present. Where two dips cross, they tend to be equal in magnitude. Studying the output more carefully, we notice that of the two dips, the one that is strongest on the input becomes irregular and noisy on the output, whereas the other dip tends to remain phase-coherent.




Figure 9.4: Left is a synthetic reflectivity model. Right is the result of local monoplane annihilation. `lmn/sep73 sigmoid`




Figure 9.5: Conflicting dips before and after application of a local monoplane annihilator. `lmn/sep73 conflict`

I could rebuild Figure 9.5 to do a better job of suppressing monodip areas if I passed the image through a lowpass filter, and then designed a gapped deconvolution operator. Instead, I preferred to show you high-frequency noise in the place of an attenuated wavefront.

The residual of prediction-error deconvolution tends to have a white spectrum in time. This aspect of deconvolution is somewhat irritating and in practice it requires us to post-filter for display, to regain continuity of signals. As is well known (PVI, for example), an alternative to postfiltering is to put a gap in the filter. A gapped filter should work with 2-D filters too, but it is too early to describe how experimenters will ultimately choose to arrange gaps, if any, in 2-D filters. There are some interesting possibilities. (Inserting a gap also reduces the required number of CD iterations.)

9.3.4 Tests of 2-D LOMOPLAN on field data

Although the LOMOPLAN concept was developed for geophysical *models*, not raw *data*, initial experience showed that the LOMOPLAN program is effective for quality testing data and data interpretation.

Some field-data examples are in Figures 9.6 and 9.7. These results are not surprising. A dominant local plane is removed, and noise or the second-from-strongest local plane is left. These data sets fit the local plane model so well that subtracting the residual noise from the data made little improvement. These figures are clearer on a video screen. To facilitate examination of the residual on Figure 9.6 on paper (which has a lesser dynamic range than video), I recolored the white residual with a short triangle filter on the time axis.




Figure 9.6: Data section from the Gulf of Mexico (left) and after LOMOPLAN (right)
 Press button for movie. [lmn/sep73 dgulf](#)

The residual in Figure 9.7 is large at the dead trace and wherever the data contains

crossing events. Also, closer examination showed that the strong residual trace near 1.1 km offset is apparently slightly time-shifted, almost certainly a cable problem, perhaps resulting from a combination of the stepout and a few dead pickups. Overall, the local-plane residual shows a low-frequency water-velocity wave seeming to originate from the ship.




Figure 9.7: Portion of Yilmaz and Cumro data set 27 (left) and after LOMOPLAN (right).
Press button for movie. [lmn/sep73 yc27](#)

9.4 GRADIENT ALONG THE BEDDING PLANE

The LOMOPLAN (LOcal MONoPLane ANnihilator) filter in three dimensions is a deconvolution filter that takes a volume in and produces two volumes out. The x -output volume results from a first order prediction-error filter on the x -axis, and the y -output volume is likewise on the y -axis.

Although I conceived of 2-D LOMOPLAN as the “ultimate” optimization criterion for inversion problems in reflection seismology of sedimentary sections, it turned out that it was more useful in data interpretation and in data-quality inspection. In this study, I sought to evaluate usefulness with *three*-dimensional data such as 3-D stacks or migrated volumes, or 2-D prestack data.

In experimenting with 3-D LOMOPLAN, I came upon a conceptual oversimplification, which although it is not precisely correct, gives a suitable feeling of the meaning of the operator. Imagine that the earth was flat horizontal layers, except for occasional faults. Then, to find the faults you might invoke the horizontal gradient of the 3-D continuum of data. The horizontal components of gradient vanish except at a fault, where their relative magnitudes tell you the orientation of the fault. Instead of using the gradient vector, you could use prediction-error filters of first order (two components) along x and y directions.

3-D LOMOPLAN is like this, but the flat horizontal bedding may be dipping or curved. No output is produced (ideally) except at faults. The 3-D LOMOPLAN is like the gradient *along the plane of the bedding*. It is nonzero where the bedding has an intrinsic change.

LOMOPLAN flags the bedding where there is an intrinsic change.

9.4.1 Definition of LOMOPLAN in 3-D

Three-dimensional LOMOPLAN is somewhat like multiple passes of two-dimensional LOMOPLAN; i.e., we first LOMOPLAN the (t, x) -plane for each y , and then we LOMOPLAN the (t, y) -plane for each x . Actually, 3-D LOMOPLAN is a little more complicated than this. Each LOMOPLAN filter is designed on all the data in a small (t, x, y) volume.

To put the L^Ocal in LOMOPLAN we use subcubes (bricks). Recall that we can do 2-D LOMOPLAN with the prediction-error subroutine `find_lopef()` on page ?? . To do 3-D LOMOPLAN we need to make two calls to subroutine `find_lopef()`, one for the x -axis in-line planar filters and one for the y -axis crossline filters. That is what I will try next time I install this book on a computer with a bigger memory.

9.4.2 The quarterdome 3-D synthetic (qdome)

Figure 9.4 used a model called “Sigmoid.” Using the same modeling concepts, I set out to make a three-dimensional model. The model has horizontal layers near the top, a Gaussian appearance in the middle, and dipping layers on the bottom, with horizontal unconformities between the three regions. Figure 9.8 shows a vertical slice through the 3-D “qdome” model and components of its LOMOPLAN. There is also a fault that will be described later. The most interesting part of the qdome model is the Gaussian center. I started from the equation of a Gaussian

$$z(x, y, t) = e^{-(x^2+y^2)/t^2} \quad (9.3)$$

and backsolved for t

$$t(x, y, z) = \sqrt{\frac{x^2 + y^2}{-\ln z}} \quad (9.4)$$

Then I used a random-number generator to make a blocky one-dimensional impedance function of t . At each (x, y, z) location in the model I used the impedance at time $t(x, y, z)$, and finally defined reflectivity as the logarithmic derivative of the impedance. Without careful interpolation (particularly where the beds pinch out) a variety of curious artifacts appear. I hope to find time to use the experience of making the qdome model to make a tutorial lesson on interpolation. A refinement to the model is that within a certain subvolume the time $t(x, y, z)$ is given a small additive constant. This gives a fault along the edge of the subvolume. Ray Abma defined the subvolume for me in the qdome model. The fault looks quite realistic, and it is easy to make faults of any shape, though I wonder how they would relate to realistic fault dynamics. Figure 9.9 shows a top view of the 3-D qdome model and components of its LOMOPLAN. Notice that the cross-line spacing has been chosen to be double the in-line spacing. Evidently a consequence of this, in both Figure 9.8 and Figure 9.9, is that the Gaussian dome is not so well suppressed on the crossline cut as on the in-line cut. By comparison, notice that the horizontal bedding




Figure 9.8: Left is a vertical slice through the 3-D “qdome” model. Center is the in-line component of the LOMOPLAN. Right is the cross-line component of the LOMOPLAN.

lmn/sep77 qdomesico

above the dome is perfectly suppressed, whereas the dipping bedding below the dome is imperfectly suppressed.

Finally, I became irritated at the need to look at *two* output volumes. Because I rarely if ever interpreted the polarity of the LOMOPLAN components, I formed their sum of squares and show the single square root volume in Figure 9.10.

9.5 3-D SPECTRAL FACTORIZATION

Hi Sergey, Matt, and Sean, Here are my latest speculations, plans:

The 3-D Lomoplan resembles a gradient, one field in, two or three out. Lomoplan times its adjoint is like a generalized laplacian. Factorizing it yields a lomoplan generalization of the helix derivative, i.e. a one-to-one operator with the same spectral characteristic as the original lomoplan. It will probably not come out to be a juxtaposition of planes, will be more cube like.

The advantage of being one-to-one is that it can be used as a preconditioner. The application, naturally enough, is estimating things with a prescribed dip spectrum. Things like missing data and velocities.

Why use multiplanar lomoplan estimates if they will then be converted by this complicated process into a cube? Why not estimate the cube directly? Maybe to impose the “pancake” model instead of the noodle model of covariance. Maybe to reduce the number




Figure 9.9: Left is a horizontal slice through the 3-D qdome model. Center is the in-line component of the LOMOPLAN. Right is the cross-line component of the LOMOPLAN.
 lmn/sep77 qdometoco




Figure 9.10: Left is the model. Right is the magnitude of the LOMOPLAN components in Figure 9.9.
 lmn/sep77 qdometora

of coefficients to estimate.

I haven't figured out yet how to convert this speculation into an example leading to some figures. If you like the idea, feel free to beat me to it :)

REFERENCES

- Claerbout, J. F., 1992a, *Earth Soundings Analysis: Processing Versus Inversion*: Blackwell Scientific Publications.
- , 1992b, Information from smiles: Mono-plane-annihilator weighted regression, *in* SEP-73, 409–420, Stanford Exploration Project.

Chapter 10

Some research examples

SEP students and researchers have extended the work described in this book. A few of their results are summarized here without the details and working codes.

10.1 GULF OF MEXICO CUBE

David Lumley from Chevron gave James Rickett some nice 3-D data from the Gulf of Mexico. There movie shows time slices at intervals of about 6ms. These slices are about 18 feet apart. That is about 7,000 years of deposition in the Gulf of Mexico. Altogether it is about a million years (about the age of the human species). Figure 10.1 shows some nice time slices.




Figure 10.1: Some time slices show a salt dome, some river channels, a dendritic drainage canyon, and a fault. rez/canyon canyon

