UCRL-JC-135513 Pt 1 Version date: September 7, 1999

Analysis of Approximate Inverses in Tomography
I. Resolution Analysis of Common Inverses

James G. Berryman
Lawrence Livermore National Laboratory
P. O. Box 808 L-200
Livermore, CA 94551-9900

Abstract

The process of using physical data to produce images of important physical parameters is an inversion problem, and these are often called tomographic inverse problems when the arrangement of sources and receivers makes an analogy to x-ray tomographic methods used in medical imaging possible. Examples of these methods in geophysics include seismic tomography, ocean acoustic tomography, electrical resistance tomography, etc., and many other examples could be given in nondestructive evaluation and other applications. All these imaging methods have two stages: First, the data are operated upon in some fashion to produce the image of the desired physical quantity. Second, the resulting image must be evaluated in essentially the same timeframe as the image is being used as a diagnostic tool. If the resolution provided by the image is good enough, then a reliable diagnosis may ensue. If the resolution is not good enough, then a reliable diagnosis is probably not possible. But the first question in this second stage is always "How good is the resolution?" The concept of resolution operators and resolution matrices has permeated the geophysics literature since the work of Backus and Gilbert in the late 1960s. But measures of resolution have not always been computed as often as they should be because, for very data rich problems, these computations can actually be significantly more difficult/expensive than computing the image itself.

It is the purpose of this paper and its companion (Part II) to show how resolution operators/matrices can be computed economically in almost all cases, and to provide a means of comparing the resolution characteristics of many of the common approximate inverse methods. Part I will introduce the main ideas and analyze the behavior of standard methods such as damped least-squares, truncated singular value decomposition, the adjoint method, backprojection formulas, etc. Part II will treat many of the standard iterative inversion methods including conjugate gradients, Lanczos, LSQR, etc.

1 Introduction

The primary goal of tomography is to take a set of data and produce an image, in either two or three dimensions, of some physical quantity of interest. For example, x-ray tomography uses data on attenuation of x-rays to image density contrasts in a body, while seismic tomography uses measured traveltimes of soundwave propagation to deduce and image velocity structure of the earth. A secondary, but often equally important, goal of tomography is to make some statement about how reliable the resulting image is, or about how much more reliable certain parts of the image are compared to other parts. This secondary goal of tomography will be the main topic of the present work. A short name for the issue in question is analysis of "resolution" in tomography, but the resolution to which we refer is a technical quantity that may not agree with the reader's notion of what the term "resolution" should mean. To avoid this possible conflict between common and technical meanings of the term "resolution," we will refer to our current enterprise as "analysis of approximate inverses in tomography." By this phrase we mean that we intend to proceed towards the secondary goal of tomography by analyzing many of the types of procedures used to produce images, providing whenever possible a straightforward computational means of assessing the reliability of the images produced, and also gaining some insight into the practical issues of "resolution" in both the common and technical senses of the term.

The first part of this article focuses on qualitative aspects of the analysis, while the second part [Berryman, 1999] focuses on more quantitative aspects. The qualitative aspects of the analysis entail trying to seek common and desirable features in the approximate inverses and their resolution operators. For example, one desirable feature for problems involving real operators/matrices is a symmetric resolution operator, since it is very dificult to interpret and therefore to make any practical use of a nonsymmetric one. The quantitative aspects of the analysis entail construction of practical computational methods that allow direct computation of the resolution operators for any of the methods considered and in particular for the iterative methods that tend to be used for the larger problems occuring in modern imaging problems.

Before beginning the analysis, we will motivate our choice of problem formulation in the next section. Then, the seismic traveltime tomography problem is introduced as our generic example of the inverse problems to be studied. Section 4 introduces the Moore-Penrose pseudoinverse of the pertinent matrices, and emphasizes the role of the resolution projection operator both in the defining relations for the pseudoinverse and in its analysis. The l_p-norms are introduced in Section 5 and the significance of the absolute value norm (l_1) is discussed in Section 6. Regularization of least-squares solutions is discussed in Section 7 and other approximate methods of generating pseudoinverses are presented in Section 8. Section 9 provides an overview of the analysis of approximate inverses.

2 Continuous Versus Discrete Operators

The methods to be treated here can be applied to many different problems in imaging and inversion. We will take Earth imaging as our primary example, but the reader who has other interests should view the methods presented as just one concrete example of the very

general linear inversion problem of the form $Ax = y$, where A is some linear operator, y is data function or vector, and x is a model function or vector to be determined.

When studying the Earth, geophysicists often like to remind us that the physical properties of the Earth are highly variable from location to location and from scale to scale. Therefore, the student of the Earth should always keep in mind that whatever model of the Earth is being constructed, by whatever combination of experiment and analysis, will always be inadequate in some sense. For the present purposes, we will say such models are underparameterized. We could perhaps imagine the ideal Earth model would actually be a continuum model, while any model we can construct in practice is merely a crude cartoon of that continuum. Discretely parameterized models are, of course, very commonly used in computer codes. To believe that discretely parameterized models have meaning, we must have enough faith in the chosen parameters to think that they represent the "true" Earth locally at least in some average sense. Presumably it will always be hopeless to parameterize the full Earth continuum, but we hope it will also be unnecessary to do so. Processes taking place on a scale smaller than the grid spacing (call these the subgrid processes) can (we hope) be usefully approximated by treating a model with relatively few well-chosen parameters.

The methods we will describe in this paper could be formulated in at least two different ways: (1) in terms of an operator language that would apply to virtually any means of computing the desired quantities appearing in the analysis, or (2) in terms of matrices, which we may view as a discrete version of the equivalent operator. It is somewhat easier and more concrete to discuss the methods in terms of matrices and that is what we choose to do here.

3 Example: Seismic Traveltime Tomography

Various examples of the methods to be discussed in this article could be provided, but to be keep the analysis both simple and concrete we will limit most of our discussion to physical problems based on seismic traveltime tomography [Aki and Richards, 1980; Nolet, 1987]. Other choices such as ocean acoustic tomography [Munk *et al.*, 1995] or electrical impedance tomography [Yorkey *et al.*, 1987], among many others, could have been made, but an understanding of seismic traveltime tomography is sufficient to introduce all the main ideas and yet complex enough to be a good introduction to the main difficulties with these methods as well. We will introduce the problem of seismic traveltime tomography in this section and then return to it when we provide some worked examples at the end of the article.

A typical problem in seismic traveltime tomography is to infer the (isotropic) compressional-wave velocity distribution of an inhomogeneous medium, given a set of observed first-arrival traveltimes between sources and receivers of known location within the medium. This problem is common for borehole-to-borehole seismic tomography in oil field applications [Bording *et al.*, 1987; Rector, 1995]. We might also consider the problem of inverting for wave slowness when the absolute traveltimes are not known, as is normally the case in earthquake seismology.

3.1 Slowness Models

We can consider three kinds of slowness (reciprocal of velocity) models. If we permit the slowness to be a general function of position, $s(\mathbf{x})$, then this is a continuum model, which is explicitly excluded from consideration here. We can instead make one of two more restrictive assumptions that (i) the model comprises homogeneous blocks, or cells, with s_j denoting the slowness value of the jth cell, or (ii) the model is composed of a grid with values of slowness assigned at the grid points with some interpolation scheme to assign the values between grid points.

When it is not important which type of slowness model is involved, we will refer to the model abstractly as a vector $\mathbf{s}$ in a vector space. For a block model with n blocks, the model becomes a vector in R^n, the n-dimensional Euclidean vector space.

3.2 Fermat's Principle and Traveltime Functionals

The traveltime of a seismic wave is the integral of slowness along a ray path connecting the source and receiver. To make this more precise, we will define two functionals for traveltime.

Let P denote an arbitrary path connecting a given source and receiver in a slowness model $\mathbf{s}$. We will refer to P as a *trial ray path*. We define a functional τ^P which yields the traveltime along path P. Letting $\mathbf{s}$ be the continuous slowness distribution $s(\mathbf{x})$, we have

$$\tau^P(\mathbf{s}) = \int_P s(\mathbf{x})\, dl^P, \tag{1}$$

where dl^P denotes the infinitesimal distance along the path P.

Fermat's principle states that the correct ray path between two points is the one of least overall traveltime, *i.e.*, it minimizes $\tau^P(\mathbf{s})$ with respect to path P. [Actually, Fermat's principle is the weaker condition that the traveltime integral is *stationary* with respect to variations in the ray path, but for traveltime tomography using measured first arrivals it follows that the traveltimes must truly be *minima*.]

Let us define τ^* to be the functional that yields the traveltime along the Fermat (least-time) ray path. Fermat's principle then states

$$\tau^*(\mathbf{s}) = \min_{P \in Paths} \tau^P(\mathbf{s}), \tag{2}$$

where $Paths$ denotes the set of all continuous paths connecting the given source and receiver. The particular path that produces the minimum in (2) is denoted P^*. If more than one path produces the same minimum traveltime value, then P^* denotes any particular member in this set of minimizing paths.

To summarize, we have

$$\tau^*(\mathbf{s}) = \int_{P^*} s(\mathbf{x})\, dl^{P^*} = \min_P \int_P s(\mathbf{x})\, dl^P = \min_P \tau^P(\mathbf{s}). \tag{3}$$

The traveltime functional $\tau^*(\mathbf{s})$ is stationary with respect to small deviations from the path $P^*(s)$.

Snell's law is a well-known and easily derived consequence of the stationarity of the traveltime functional [Feynman *et al.*, 1963; Whitham, 1974].

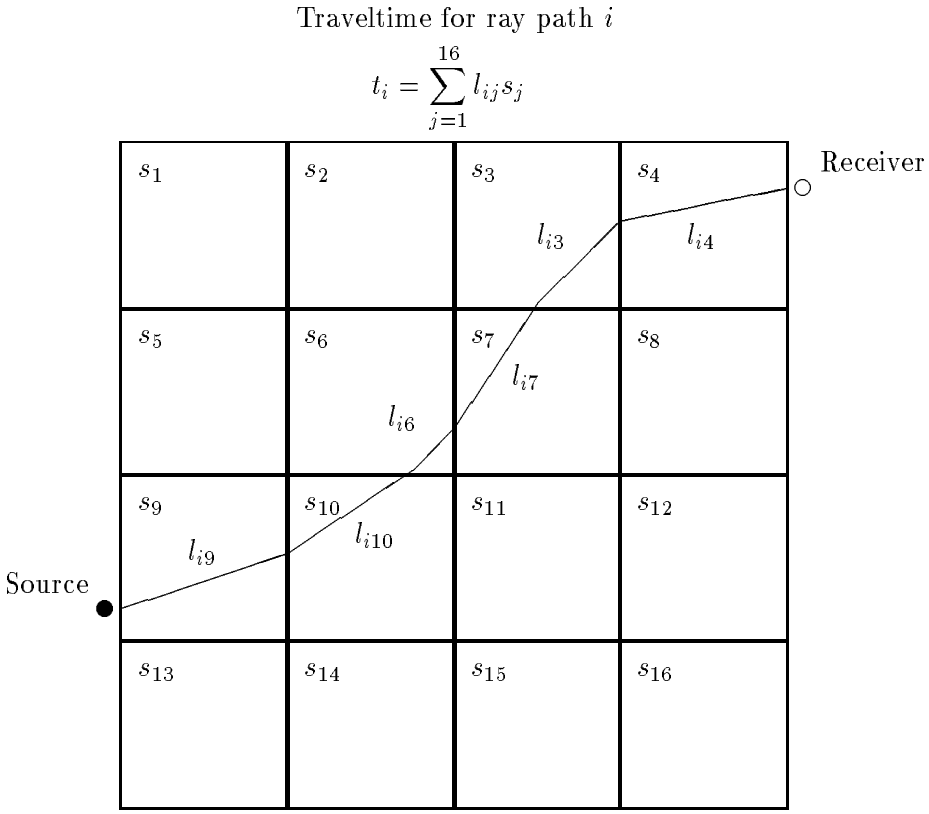

Figure 1: Schematic illustration of acoustic ray paths through a discrete, constant velocity/slowness model.

3.3 Seismic Inversion

Suppose we have a set of observed traveltimes, $t_1, \ldots, t_m$, from m source-receiver pairs in a medium of slowness $s(\mathbf{x})$. Let P_i be the Fermat ray path connecting the ith source-receiver pair. Neglecting observational errors, we can write

$$\int_{P_i} s(\mathbf{x})\, dl^{P_i} = t_i, \quad i = 1, \ldots, m. \tag{4}$$

Given a block model of slowness, let l_{ij} be the length of the ith ray path through the jth cell:

$$l_{ij} = \int_{P_i \cap \text{cell}_j} dl^{P_i}. \tag{5}$$

Given a model with n cells with the jth cell having constant slowness s_j, (4) can then be written

$$\sum_{j=1}^{n} l_{ij} s_j = t_i, \quad i = 1, \ldots, m. \tag{6}$$

Note that for any given i, the ray-path lengths l_{ij} are zero for most cells j, since a given ray path will in general intersect only a few of the cells in the model. Figure 1 illustrates the ray path intersections for a 2-D block model.

We can rewrite (6) in matrix notation by defining the column vectors $\mathbf{s}$ and $\mathbf{t}$ and the matrix $\mathbf{M}$ as follows:

$$\mathbf{s} = \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix}, \qquad \mathbf{t} = \begin{pmatrix} t_1 \\ t_2 \\ \vdots \\ t_m \end{pmatrix}, \qquad \mathbf{M} = \begin{pmatrix} l_{11} & l_{12} & \cdots & l_{1n} \\ l_{21} & l_{22} & \cdots & l_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ l_{m1} & l_{m2} & \cdots & l_{mn} \end{pmatrix}. \tag{7}$$

Equation (6) then becomes

$$\mathbf{M}\mathbf{s} = \mathbf{t}. \tag{8}$$

Equation (8) is just the discretized form of (3).

Equation (8) will be become our generic inversion problem in the remainder of this article. The right hand side is the *data* m-vector $\mathbf{t}$, whose elements are the m traveltime measurements in this application. The left hand side has as one of its factors an $m \times n$ matrix $\mathbf{M}$, which for this article will be assumed known (but for real applications is often another unknown in the problem [Berryman, 1990]). The remaining factor is the *model* n-vector $\mathbf{s}$, which is the model we seek in the inversion process. In practice, either $m < n$ or $m > n$ may occur in real problems. We refer to the case $m < n$ as underdetermined and the case $m > n$ as overdetermined, although this terminology is an oversimplification of the situation usually encountered in practice [Jackson, 1972]. Most of the remainder of the article will concentrate on solving this equation approximately for $\mathbf{s}$, and simultaneously providing some means of analyzing how well the method of approximate solution has worked. We return to the specific application to seismic traveltime tomography at the end of Part II.

4 Matrix Analysis and the Pseudoinverse

Various peculiar properties of a *rectangular matrix* $\mathbf{M}$ of size $m \times n$ (with $m \neq n$) can cause the innocent looking equation

$$\mathbf{M}\mathbf{s} = \mathbf{t} \tag{9}$$

to be very difficult to solve for the model vector $\mathbf{s}$ in terms of the data vector $\mathbf{t}$. Although generalized matrix inverses may be defined, each definition requires the choice of some figure of merit. The most common choices are based on error norms: the resulting model vector

should minimize an objective functional of the errors $\mathbf{t} - \mathbf{Ms}$. Although each objective functional generally leads to a unique approximate inverse, the choice of the objective functional itself is somewhat arbitrary. Therefore, it is important to analyze the range of useful choices and to determine general characteristics that help to distinguish among these choices.

Least-squares objective functionals are the ones most commonly used for several reasons:

1. The solution of the least-squares minimization problem is linear and therefore requires only linear algebraic computations.

2. A well-developed theory of generalized inverses exists and the most natural one of these (the Moore-Penrose pseudoinverse [Moore, 1920; Penrose, 1955a]) is intimately related to the least-squares solution [Penrose, 1955b].

3. Resolution properties of the least-squares solutions are well-understood.

Other types of approximate inverses may be and often are considered. For example, elementary backprojection formulas are themselves approximate inverses. More sophisticated objective functionals based on other measures of the error — instead of the squared error — are most often considered when the data are found and/or expected to be erratic. Robust inversion techniques sometimes use the l_1-norm (absolute error) either instead of or in combination with the l_2-norm (squared error) to reduce and/or eliminate the effects of outliers, which tend to dominate least-squares analyses but have little or no effect on the solution for l_1-based analysis. The biggest problems with these approximate inverses is that the minimum of an objective functional based on absolute values is found by linear programming (rather than linear algebra) and, furthermore, the minimum may not be unique. (More about this in Section 6.) Another important problem is that the resolution properties of these other approximate inverses (including the commonly used one-shot backprojection formulas) are not well-understood.

We examine these questions in this section. Our main purpose is to find qualitative means of comparing the performance of various approximate inversion techniques. Since the Moore-Penrose pseudoinverse is probably the most thoroughly analyzed of the generalized inverses, we check each approximate inverse for satisfaction of the four Moore-Penrose uniqueness conditions. Because of the uniqueness result, any inverse that satisfies all four is equivalent to the Moore-Penrose inverse. Most of the approximate inverses mentioned violate one or more of these conditions.

The more quantitative aspects of this problem will be examined in some detail in Part II. Important practical results of this analysis are explicit formulas for computing both model and data resolution matrices for many of the more common iterative methods, including conjugate directions/gradients [Hestenes and Stiefel, 1952], Lanczos [Lanczos, 1950], and LSQR [Paige and Saunders, 1982].

4.1 Moore-Penrose conditions

The Moore-Penrose pseudoinverse is very well-known, but we want to approach it from an unusual angle and so give a different derivation in this subsection. The advantage of this

approach will be seen when we are able to give a definite interpretation to each of the four Moore-Penrose uniqueness conditions.

We need the following fact to prove our main results in this section:

PROPOSITION 4.1 *For any real matrix* $\mathbf{A}$,

$$\mathbf{A}^T\mathbf{A} \equiv 0 \quad \textit{implies} \quad \mathbf{A} = 0. \tag{10}$$

Proof: If $\mathbf{A}$ is an arbitrary $m \times n$ real matrix whose column m-vectors are $\mathbf{a}_1$, $\mathbf{a}_2$, $\mathbf{a}_3$, ..., $\mathbf{a}_n$, then

$$\begin{aligned}\mathbf{A}^T\mathbf{A} &= \begin{pmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \vdots \\ \mathbf{a}_n^T \end{pmatrix} \begin{pmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \dots & \mathbf{a}_n \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{a}_1^T\mathbf{a}_1 & \mathbf{a}_1^T\mathbf{a}_2 & \dots & \mathbf{a}_1^T\mathbf{a}_n \\ \mathbf{a}_2^T\mathbf{a}_1 & \mathbf{a}_2^T\mathbf{a}_2 & \dots & \mathbf{a}_2^T\mathbf{a}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_n^T\mathbf{a}_1 & \mathbf{a}_n^T\mathbf{a}_2 & \dots & \mathbf{a}_n^T\mathbf{a}_n \end{pmatrix}. \end{aligned} \tag{11}$$

The diagonal components of $\mathbf{A}^T\mathbf{A}$ are the squares of the magnitudes of the column vectors. Then clearly, if $\mathbf{A}^T\mathbf{A} \equiv 0$, all the column vectors $\mathbf{a}_i \equiv 0$ and therefore $\mathbf{A} \equiv 0$. ∎

THEOREM 4.1 *The three equations*

$$\boxed{\begin{aligned} \mathbf{MXM} &= \mathbf{M}, \\ (\mathbf{MX})^T &= \mathbf{MX}, \\ (\mathbf{XM})^T &= \mathbf{XM}, \end{aligned}}$$

have a unique solution $\mathbf{X}$ *for any real matrix* $\mathbf{M}$ *of full rank.*

Proof: Suppose there are two solutions $\mathbf{X}$ and $\mathbf{Y}$, both satisfying the three conditions of the theorem. The first condition of the theorem applied to both $\mathbf{X}$ and $\mathbf{Y}$ gives

$$\mathbf{MXM} = \mathbf{M} = \mathbf{MYM} \tag{12}$$

and, therefore, implies that

$$\mathbf{M}(\mathbf{X} - \mathbf{Y})\mathbf{M} = 0. \tag{13}$$

Then, using $\mathbf{X}^T\mathbf{M}^T = (\mathbf{MX})^T = \mathbf{MX}$,

$$(\mathbf{X} - \mathbf{Y})^T\,\mathbf{M}^T\mathbf{M}(\mathbf{X} - \mathbf{Y}) = \mathbf{M}(\mathbf{X} - \mathbf{Y})\mathbf{M}(\mathbf{X} - \mathbf{Y}) = 0. \tag{14}$$

Because all the diagonal elements of the left-hand-side of (14) are sums of squares, it can only vanish (see PROPOSITION 4.1) if

$$\mathbf{M}\,(\mathbf{X}-\mathbf{Y})=0, \tag{15}$$

showing that the columns of $(\mathbf{X}-\mathbf{Y})$ must all lie in the right null space of $\mathbf{M}$.

Using $\mathbf{M}^T\mathbf{X}^T=(\mathbf{X}\mathbf{M})^T=\mathbf{X}\mathbf{M}$, we also find that

$$\mathbf{M}^T\,(\mathbf{X}-\mathbf{Y})^T\,(\mathbf{X}-\mathbf{Y})\,\mathbf{M}=(\mathbf{X}-\mathbf{Y})\,\mathbf{M}\,(\mathbf{X}-\mathbf{Y})\,\mathbf{M}=0 \tag{16}$$

implies

$$(\mathbf{X}-\mathbf{Y})\,\mathbf{M}=0, \tag{17}$$

showing that the rows of $(\mathbf{X}-\mathbf{Y})$ must all lie in the left null space of $\mathbf{M}$. [Note that (15) and (17) were obtained without using the assumption that $\mathbf{M}$ is of full rank.] However, since $\mathbf{M}$ is assumed to be of full rank, one or both of these equations — either (15) or (17) or both — can only be satisfied if $(\mathbf{X}-\mathbf{Y})\equiv 0$. Thus, $\mathbf{X}=\mathbf{Y}$, and the solution is unique. ∎

THEOREM 4.2 *(Moore-Penrose) The four equations*

$$\boxed{\begin{aligned}\mathbf{M}\mathbf{X}\mathbf{M}&=\mathbf{M},\\ \mathbf{X}\mathbf{M}\mathbf{X}&=\mathbf{X},\\ (\mathbf{M}\mathbf{X})^T&=\mathbf{M}\mathbf{X},\\ (\mathbf{X}\mathbf{M})^T&=\mathbf{X}\mathbf{M},\end{aligned}}$$

have a unique solution $\mathbf{X}$ *for any real matrix* $\mathbf{M}$.

Proof: Suppose there are two solutions $\mathbf{X}$ and $\mathbf{Y}$, both satisfying all four conditions of the theorem. Then, the proof of the preceding theorem requires only the three boxed relations of THEOREM 4.1 to derive (15) and (17) and therefore shows that the first, third, and fourth conditions of the present Theorem imply

$$\mathbf{M}\mathbf{X}=\mathbf{M}\mathbf{Y} \tag{18}$$

and

$$\mathbf{X}\mathbf{M}=\mathbf{Y}\mathbf{M}. \tag{19}$$

Multiplying (18) on the left by $\mathbf{X}$ and (19) on the right by $\mathbf{Y}$, we find

$$\mathbf{X}\mathbf{M}\mathbf{X}=\mathbf{X}\mathbf{M}\mathbf{Y}=\mathbf{Y}\mathbf{M}\mathbf{Y}. \tag{20}$$

Substituting the second condition from the statement of the present theorem into (20) on both the left and the right shows that

$$\mathbf{X}=\mathbf{X}\mathbf{M}\mathbf{Y}=\mathbf{Y}. \tag{21}$$

Thus, the solution $\mathbf{X}$ is unique. ∎

The unique generalized inverse satisfying the conditions of THEOREM 4.2 is the Moore-Penrose pseudoinverse $\mathbf{X} \equiv \mathbf{M}^\dagger$, where the dagger notation means simply that $\mathbf{M}^\dagger$ is the Moore-Penrose pseudoinverse of $\mathbf{M}$. The unique generalized inverse found in THEOREM 4.1 is also the Moore-Penrose inverse, but we see that the additional condition $\mathbf{XMX} = \mathbf{X}$ is a consequence of the other three conditions when $\mathbf{M}$ is of full rank — since then either $\mathbf{MX} = \mathbf{I}_m$ or $\mathbf{XM} = \mathbf{I}_n$ must hold, when $m \neq n$.

4.2 Significance of the Moore-Penrose conditions

Recalling the physical interpretation of $\mathbf{Ms} = \mathbf{t}$ given after (8), we will consistently refer to $\mathbf{s}$ as the model vector and to $\mathbf{t}$ as the data vector. Following this convention, any n-vector that might be multiplied on the right hand side of $\mathbf{M}$ is a vector in the "model space," while any m-vector whose tranpose might be applied to the left hand side of $\mathbf{M}$ (or equivalently the vector itself to the right hand side of $\mathbf{M}^T$) is a vector in the "data space."

To interpret the Moore-Penrose conditions, we must first define two types of *ideal* resolution operator in terms of the Moore-Penrose pseudoinverse $\mathbf{M}^\dagger$. The *model resolution* operator is defined as

$$\mathcal{R}_{model} = \mathbf{M}^\dagger\mathbf{M}, \tag{22}$$

while the data resolution operator is defined as

$$\mathcal{R}_{data} = \mathbf{M}\mathbf{M}^\dagger. \tag{23}$$

By considering the singular-value decompositions of $\mathbf{M}$ and $\mathbf{M}^\dagger$, it is easy to see that these two operators are projection operators respectively on the parts of the model and data spaces spanned by the left and right singular vectors of $\mathbf{M}$ corresponding to nonzero singular values. Precisely the same resolution operators arise if we consider the normal matrix $\mathbf{M}^T\mathbf{M}$ for overdetermined systems or the normal matrix $\mathbf{M}\mathbf{M}^T$ for underdetermined systems. The unique resolution matrix for the symmetric matrix $\mathbf{M}^T\mathbf{M}$ is $\mathcal{R}_{model}$ as defined in (22), and similarly $\mathcal{R}_{data}$ as defined in (23) is the unique resolution matrix for $\mathbf{M}\mathbf{M}^T$.

The resolution matrices defined in this way are ideal resolution matrices that define the subspaces spanned by the singular vectors having nonzero singular values. Knowledge of these matrices/operators helps us to gain insight into the degree of trustworthiness of the model $\mathbf{s}$ constructed by using the exact $\mathbf{M}^\dagger$ and also the degree to which the information in the data vector $\mathbf{t}$ has been used in the reconstruction. If both resolution matrices are identity matrices, then a perfect inverse has been achieved. But in realistic inverse problems, this essentially never happens. If the matrix is rank deficient or if it is so large that it is impractical to compute $\mathbf{M}^\dagger$, then one or the other and often both $\mathcal{R}_{model}$ and $\mathcal{R}_{data}$ differ from the identity matrices of the ideal case.

In particular, we will define analogs of these matrices for problems in which an approximate inverse $\mathbf{X}$ is used in place of $\mathbf{M}^\dagger$. In such circumstances, we will call the resulting matrices the "effective resolution matrices." However, we want to be clear that effective resolution matrices do not necessarily provide good approximations to the ideal resolution matrices. In fact, the ideal resolution matrices may be completely irrelevant to the problem

obtained using some choices of the approximate inverse. Instead we want to define the "effective resolution matrices" so they correctly compute the appropriate projection operator for the approximate inverse $\mathbf{X}$ no matter how crude it may be.

Each of the four Moore-Penrose conditions can now be given a simple interpretation. For example, the last two conditions are constraints on the form of what we will call "effective resolution matrices." In particular, the effective model resolution [Backus and Gilbert, 1968; Jackson, 1972] is

$$\mathcal{E}_{model} \equiv \mathbf{XM}, \tag{24}$$

so the last condition

$$(\mathbf{XM})^T = \mathbf{XM} \tag{25}$$

is a statement that *the effective model resolution should be symmetric.* Similarly, the third condition

$$(\mathbf{MX})^T = \mathbf{MX} \tag{26}$$

states that *the effective data resolution* [Wiggins, 1972; Jackson, 1972]

$$\mathcal{E}_{data} \equiv \mathbf{MX} \tag{27}$$

should be symmetric. If these symmetry conditions are violated by an approximate inverse, we will generally find it difficult to interpret the off-diagonal terms in these resolution matrices. Thus, symmetry of the resolution matrices is a desirable quality and one we should therefore seek to enforce in any approximate inverse.

An understanding of the significance of the second condition of Moore-Penrose follows easily from the analysis of the last section. We see that the condition

$$\mathbf{XMX} = \mathbf{X}. \tag{28}$$

is *a uniqueness condition on the approximate inverse* $\mathbf{X}$. Lack of uniqueness for the matrix $\mathbf{X}$ may arise if it contains contributions from the null space of $\mathbf{M}$. To see that (28) explicitly excludes the possibility of having such contributions in $\mathbf{X}$, consider the singular-value decompositions (SVDs) [Lanczos, 1960; Golub and Van Loan, 1989]

$$\mathbf{M} = \sum_i \lambda_i \mathbf{u}_i \mathbf{v}_i^T \quad \text{and} \quad \mathbf{X} = \sum_i \lambda_i^{-1} \mathbf{v}_i \mathbf{u}_i^T + \epsilon_0 \mathbf{v}_0 \mathbf{u}_0^T, \tag{29}$$

where the vectors $\mathbf{u}_i$ and $\mathbf{v}_i$ are the singular vectors of $\mathbf{M}$ as well as the eigenvectors of $\mathbf{MM}^T$ and $\mathbf{M}^T\mathbf{M}$ respectively, the λ_i are the corresponding singular values with λ_i^2 being the nonvanishing eigenvalues of both $\mathbf{MM}^T$ and $\mathbf{M}^T\mathbf{M}$. The vectors $\mathbf{v}_0$ and $\mathbf{u}_0$ are any vectors from the left and right null spaces of $\mathbf{M}$, *i.e.*, satisfying $\mathbf{Mv}_0 = 0$ and $\mathbf{u}_0^T\mathbf{M} = 0$. Then, since the presence of $\mathbf{M}$ in (28) guarantees that the left hand side of (28) cannot have any contribution of the form $\mathbf{v}_0\mathbf{u}_0^T$ while the right hand side does have such a term if $\epsilon_0 \neq 0$, equality in (28) clearly implies $\epsilon_0 \equiv 0$, which is a uniqueness (minimum norm) condition.

Similarly, the first condition of Moore-Penrose

$$\mathbf{MXM} = \mathbf{M} \tag{30}$$

is a type of *uniqueness condition on the matrix* $\mathbf{M}$. It explicitly excludes the possibility that terms from the null space of the approximate inverse $\mathbf{X}$ contribute to the matrix $\mathbf{M}$. Thus, (28) requires that a given $\mathbf{M}$ has only one $\mathbf{X}$, while (12) requires that $\mathbf{X}$ is the approximate inverse for only one $\mathbf{M}$. Of course, strict uniqueness holds only if all four conditions are satisfied simultaneously. In fact, we will see this point elucidated further in the following discussion of the various approximate inverses considered. We find that these approximate inverses may be conveniently characterized by which subset of the Moore-Penrose conditions they violate.

Finally, we note that (28) and (30) can be rewritten to emphasize the importance of the resolution operators as

$$\mathcal{R}_{model}\mathbf{X} = \mathbf{X}\mathcal{R}_{data} = \mathbf{X} \tag{31}$$

and

$$\mathbf{M}\mathcal{R}_{model} = \mathcal{R}_{data}\mathbf{M} = \mathbf{M}. \tag{32}$$

Thus, $\mathbf{X}$ is uniquely determined by its projections on the resolved parts of the data and model spaces, and $\mathbf{M}$ uniquely determines what these resolved parts are when the resolution operators are both symmetric.

All these interpretations may be understood very easily by considering the singular-value decompositions of the matrices $\mathbf{M}$ and $\mathbf{X}$, as we did to demonstrate our interpretation of (28) as another type of uniqueness condition.

5 Significance of l_p Norms

The most often used objective function in optimization theory is the squared-error functional, but a more general objective function commonly considered is

$$\Psi_p(\mathbf{s}) = \sum_{i=1}^{m} w_i \left| \sum_{j=1}^{n} l_{ij}s_j - t_i \right|^p , \tag{33}$$

where $1 \leq p \leq \infty$ [Claerbout and Muir, 1973; Claerbout, 1976; Dines and Lytle, 1979; Scales and Gersztenkorn, 1988; Scales *et al.*, 1988; Scales *et al.*, 1990]. This functional is the sum of the absolute values of the errors taken to the power p. When $p = 2$, the functional reduces to the least-squares functional. Nonnegative weights $w_i \neq 1$ may also be used in this sum if desired, but we will take $w_i \equiv 1$ here and concentrate instead on the major differences between these functionals for $p \neq 2$ and for the most common exponent $p = 2$. (We discuss weighted least squares in Section 8.3.)

To gain some understanding of the effect of the exponent p in the objective functional, it will be helpful to consider a simpler problem first. If we want to estimate a variable x

from a series of measurements x_i for $i = 1, \ldots, m$, we may also consider

$$\Omega_p(x) = \sum_{i=1}^{m} |x - x_i|^p \,. \tag{34}$$

Differentiating (34) with respect to x and setting the result equal to zero gives

$$\frac{d\Omega_p}{dx} = p \sum_{i=1}^{m} |x - x_i|^{p-1} \operatorname{sgn}(x - x_i) = 0, \tag{35}$$

where

$$\begin{aligned} \operatorname{sgn}(r) &= r/|r| \quad \text{if} \quad r \neq 0, \\ &= 0 \quad \text{if} \quad r = 0. \end{aligned} \tag{36}$$

If an x can be found that satisfies (35), then that x corresponds to the location of the minimum of Ω_p. Equation (35) shows that ps less than unity would have the undesirable property of making the equation singular at points where variable x approaches one of the data points $x = x_i$, and furthermore one stationary point of the equation occurs at infinity; thus, to avoid these difficulties, the previously stated requirement that $1 \leq p$ follows.

Rewriting (35) using the definition of the sign function, we find the condition for the minimum is

$$\sum_{i=1}^{m} |x - x_i|^{p-2} (x - x_i) = 0. \tag{37}$$

When $p = 2$, Equation (37) has the simple solution

$$x = \frac{1}{m} \sum_{i=1}^{m} x_i, \tag{38}$$

showing that the minimum for the squared-error choice of exponent occurs when x is the arithmetic mean of the data. When $p = 1$, (37) shows that

$$\sum_{i=1}^{m} \operatorname{sgn}(x - x_i) = 0, \tag{39}$$

which can be seen to be a formal definition of the data median for the set $\{x_i\}$ if there is a unique solution to (39). Otherwise, the median need not be unique. (See Figure 2.) The most significant advantage of the median over the mean is that it is independent of the values of any outliers, whereas the mean depends linearly on these values and, therefore, can be quite sensitive to the presence of very bad (extreme) data.

The other fairly common choice of exponent is $p = \infty$. However, it is a bit easier to understand the significance of this limit by first considering finite p's that are large and even. Then, (37) becomes

$$\sum_{i=1}^{m} (x - x_i)^{p-1} = 0, \tag{40}$$

where exponent $p-1$ is large and odd. For the sake of argument, we assume that the smallest (x_{min}) and largest (x_{max}) values are unique and satisfy $x_{min} \ll x_{max}$. Then, if p is very large and x takes (as it must for satisfaction of the equality) some intermediate value (*i.e.*, $x_{min} \ll x \ll x_{max}$) so that both positive and negative terms occur on the left hand side of (40), the only terms that make a significant contribution to (40) are the extreme ones in the data sequence $\{x_i\}$. Thus, (40) reduces (with very small corrections) to

$$(x - x_{min})^{p-1} + (x - x_{max})^{p-1} \simeq 0, \tag{41}$$

or equivalently

$$(x - x_{min})^{p-1} \simeq (x_{max} - x)^{p-1}. \tag{42}$$

Taking the $(p-1)$-th root and solving for x gives

$$x \simeq \tfrac{1}{2}(x_{min} + x_{max}), \tag{43}$$

showing that the minimum of the objective functional for large p is the average of the smallest and largest values in the sequence x_i, *i.e.*, it depends *only* on the outliers. In the limit of $p \to \infty$, $\simeq$ is replaced by $=$ in (43). Such behavior clearly tends to overemphasize extreme values and to produce results hypersensitive to errors in the data, so values of $p > 2$ are seldom considered for use with real, noisy data.

6 Absolute Value Norm (l_1)

Now that we understand why $p > 2$ is not a desirable alternative for use with real data, we limit consideration to the main alternative to least-squares, *i.e.*, the absolute value norm with $p = 1$ in (33). To find the minimum of (33), we take the derivative with respect to s_j and find the set of conditions for a minimum is

$$\sum_{i=1}^{m} l_{ij} \mathrm{sgn}\left(\sum_{k=1}^{n} l_{ik} s_k - t_i\right) = 0 \quad \text{for} \quad j = 1, \ldots, n. \tag{44}$$

There are n conditions in (44) — one for each of the n cells in the model. It is not hard to see that all n conditions in (44) cannot be satisfied in general for an arbitrary set of ray-path lengths l_{ij}. For example, if a particular cell has only three ray paths passing through it and if the lengths of the paths are comparable, there is no combination of pluses and minuses such that $\pm l_{1j} \pm l_{2j} \pm l_{3j} = 0$. This difficulty is not at all exceptional for our problem. Rather it is expected to be the rule that $\sum_{i=1}^{m} (\pm) l_{ij} \neq 0$ for almost all cells j. So the minimum generally cannot be found in this manner — the technical reason being that the function to be minimized is simply nonanalytic (nondifferentiable) at the actual minimum, or minima. Since the only points of nonanalyticity for this problem are those for which the argument of the sign function vanishes, the minimum of Ψ_1 (or possibly *minima* since they may not be unique) is found by solving a linear programming problem (rather than a linear algebra problem). An example of how this works is presented in Figure 3.

We assume the mathematical inversion problem has been formulated so the number of data (m) exceeds the number of unknowns (n) in the discretized model. (This choice is not

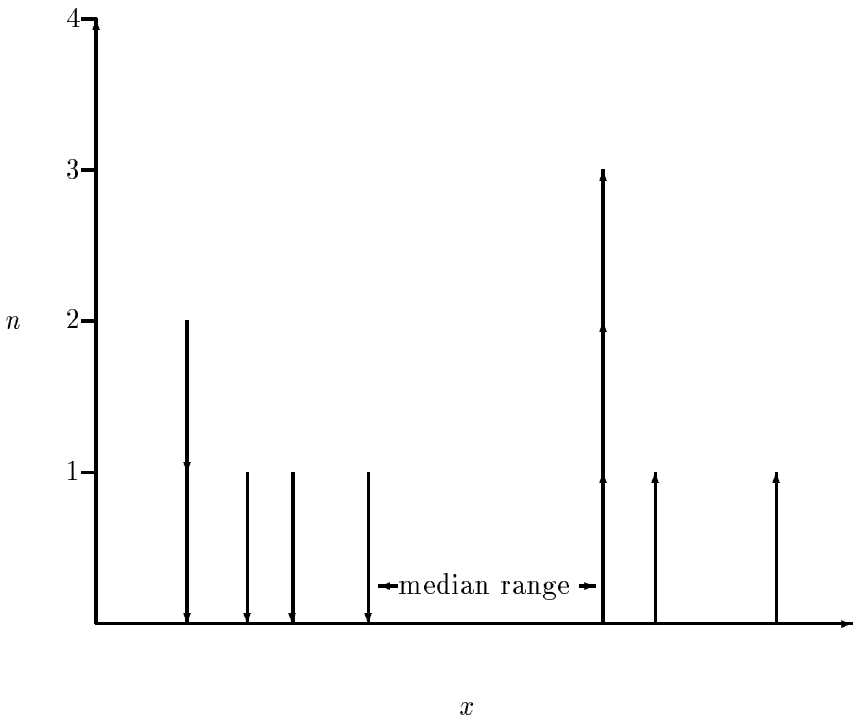

Figure 2: A median is found by varying x in (1.38) until the number of pluses and minuses balances. An arrow pointing upward corresponds to a positive contribution and an arrow pointing downwards corresponds to a negative contribution. As in the case illustrated, the median is not necessarily unique if the total number of data is even, or if the total number is odd but there are multiple (n) occurrences of some values.

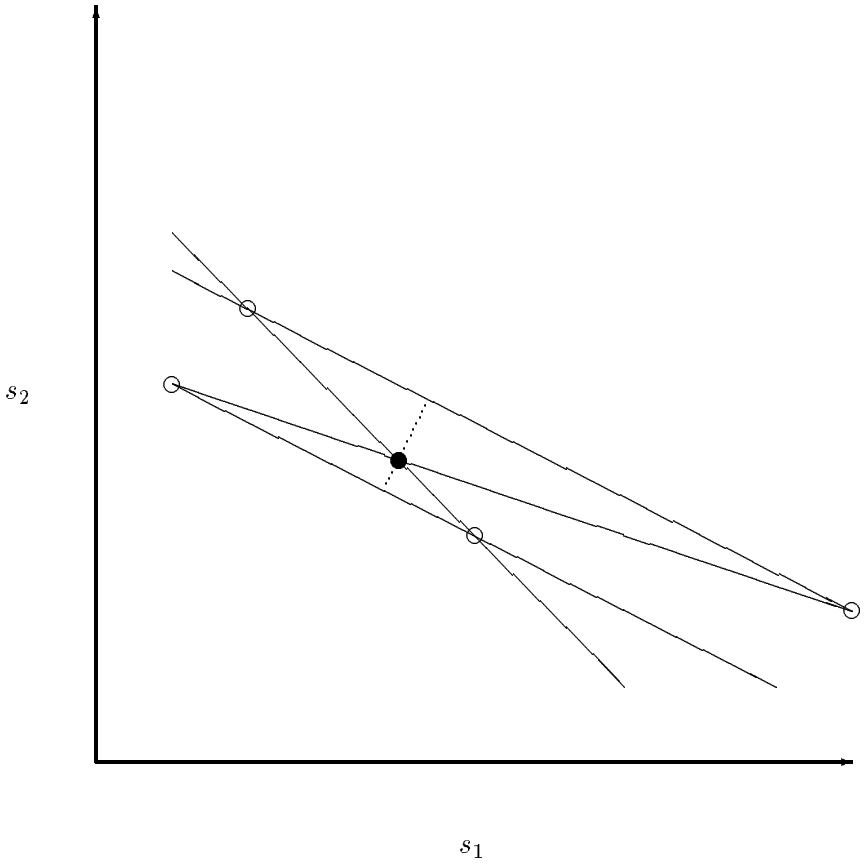

Figure 3: The minimum of the objective functional generally occurs at some point that solves a determined subset of the overdetermined set exactly.

universal. Some authors prefer to deal with underdetermined systems such that $m < n$.) In general, the solution $\mathbf{s}$ minimizing the l_1-norm of the errors may then be expressed as the solution of the rearrangement of (8) given by

$$\begin{pmatrix} \mathbf{M}' \\ \mathbf{M}'' \end{pmatrix} \mathbf{s} \simeq \begin{pmatrix} \mathbf{t}' \\ \mathbf{t}'' \end{pmatrix}, \tag{45}$$

where the m rows of the ray-path matrix have been reordered so that $\mathbf{M}'$ is an $n \times n$ (generally) *invertible submatrix* of $\mathbf{M}$ and the data vector $\mathbf{t}$ has also been reordered so that

$$\mathbf{M}'\mathbf{s} = \mathbf{t}' \quad \text{or} \quad \mathbf{s} = (\mathbf{M}')^{-1}\,\mathbf{t}'. \tag{46}$$

The remaining ray-paths contained in $\mathbf{M}''$ and the corresponding data $\mathbf{t}''$ are used only in the process of selecting a particular rearrangement of $\mathbf{M}$ out of the total of $m!/n!(m-n)!$ possible rearrangements. The final choice is made by finding a *segmentation* of the matrix such that

$$\sum_{i=n+1}^{m} \left| \sum_{j=1}^{n} l_{ij} s_j - t_i \right| \tag{47}$$

is minimized. The values for $i = 1, \ldots, n$ do not appear in (47) because, by assumption, these terms vanish identically due to (46). We will call any rearrangement of the rows and columns of $\mathbf{M}$ into parts to be satisfied ($\mathbf{M}'$) and parts to be ignored ($\mathbf{M}''$) a *segmentation* of $\mathbf{M}$; a particular choice that minimizes (47) is *an optimal segmentation*. If only one segmentation optimizes (47), then it is the unique optimal segmentation.

The approximate inverse resulting from this l_1-norm analysis is easily seen to be

$$\mathbf{X} = \begin{pmatrix} (\mathbf{M}')^{-1} & \mathbf{O}_{n,m-n} \end{pmatrix}, \tag{48}$$

where $\mathbf{O}_{n,m-n}$ is an $n \times (n-m)$ matrix of zeroes. Then, the effective model resolution matrix is trivial for this problem, since

$$\mathcal{E}_{model} = \mathbf{X}\mathbf{M} = \mathbf{I}_n. \tag{49}$$

The model is apparently perfectly resolved for this subset of the ray-path data, assuming only that $\mathbf{M}'$ is invertible. The effective data resolution matrix is not so simple however, since

$$\mathcal{E}_{data} = \mathbf{M}\mathbf{X} = \begin{pmatrix} \mathbf{I}_n & \mathbf{O}_{n,m-n} \\ \mathbf{M}''\,(\mathbf{M}')^{-1} & \mathbf{O}_{m-n,m-n} \end{pmatrix}, \tag{50}$$

showing that the first n traveltime constraints are satisfied exactly while the other $m-n$ traveltime constraints are essentially ignored. This interpretation is not entirely correct however, since the apparently ignored data were actually used in (47) while making the choice of optimal matrix segmentation.

Note that the data resolution matrix for the l_1-norm is not symmetric, so $(\mathbf{M}\mathbf{X})^T \neq \mathbf{M}\mathbf{X}$. Checking the other Moore-Penrose conditions by multiplying (50) on the right by $\mathbf{M}$, we find that

$$\mathbf{M}\mathbf{X}\mathbf{M} = \mathbf{M}, \tag{51}$$

and by multiplying (49) on the right by $\mathbf{X}$ that

$$\mathbf{XMX} = \mathbf{X}. \tag{52}$$

Since the identity matrix $\mathbf{I}_n$ is symmetric, (49) also shows directly that

$$(\mathbf{XM})^T = \mathbf{XM}. \tag{53}$$

Thus, $\mathbf{X}$ for the absolute value norm satisfies three out of four of the Moore-Penrose conditions — failing only to have a symmetric data resolution matrix. For most applications, this is a rather minor failing since the data resolution is seldom computed — or even considered.

This analysis of the resolution for the absolute value norm is somewhat unsatisfactory, since it suggests that the model resolution is not only symmetric but also always perfect. A more sensible alternative to this overly optimistic result is desirable, but so far has not been found. A statistical measure of resolution constructed by considering various nonoptimal matrix segmentations might be useful, but we will not pursue this possibility further here.

7 Regularized Least-Squares

There are two main approaches to regularizing least-squares methods: (*i*) damping [Levenberg, 1944; Marquardt, 1963, 1970; Tikhonov and Arsenin, 1977] and (*ii*) truncated singular-value decomposition (SVD) [Eckart and Young, 1936; Lanczos, 1960; Golub and Kahan, 1965; Golub and Van Loan, 1989]. Damping typically involves adding a positive diagonal term to the singular matrix $\mathbf{M}^T\mathbf{M}$ appearing in the normal equations for the least-squares solution. Truncated singular-value decomposition involves computing the SVD of the ray-path matrix $\mathbf{M}$ and then, when computing the approximate inverse of $\mathbf{M}$, discarding not only the singular vectors with zero singular value as is normally done in the Moore-Penrose pseudoinverse, but also discarding other singular vectors if the corresponding singular values lie below some threshold ε. Penrose [1955b] has shown that the full SVD actually gives the generalized inverse required for solution of the least-squares problem. So the truncated SVD is a type of regularized least-squares method; it eliminates terms that would introduce very noisy behavior into the final solution due to the presence of small denominators (*i.e.*, the inverses of the smallest eigenvalues).

To illustrate these concepts, consider an approximate inverse for $\mathbf{M}$ using damped least-squares of the form

$$\mathbf{X}_\mu = \left(\mathbf{M}^T\mathbf{M} + \mu\mathbf{I}\right)^{-1}\mathbf{M}^T, \tag{54}$$

where μ is a small positive constant. (Various other choices of damping term are possible if we substitute for the term $\mu\mathbf{I}$, which is just the simplest and most common form used.) In contrast, the truncated SVD approach gives

$$\mathbf{X}_q = \sum_{i=1}^{q} \frac{1}{\lambda_i}\mathbf{z}_i\mathbf{y}_i^T, \tag{55}$$

where λ_i is the ith singular value of $\mathbf{M}$ and the SVD of the ray-path matrix is

$$\mathbf{M} = \sum_{i=1}^{r} \lambda_i \mathbf{y}_i \mathbf{z}_i^T. \tag{56}$$

The upper limit of the sum in (56) satisfies $q \leq r$ such that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_q \geq \varepsilon$.

7.1 Damped least-squares

To check satisfaction of the Moore-Penrose conditions for a damped least-squares inverse, we first compute the effective resolution matrices

$$\mathcal{E}_{data} = \mathbf{M}\mathbf{X}_\mu = \sum_{i=1}^{r} \frac{\lambda_i^2}{\lambda_i^2 + \mu} \mathbf{y}_i \mathbf{y}_i^T \tag{57}$$

and

$$\mathcal{E}_{model} = \mathbf{X}_\mu \mathbf{M} = \sum_{i=1}^{r} \frac{\lambda_i^2}{\lambda_i^2 + \mu} \mathbf{z}_i \mathbf{z}_i^T, \tag{58}$$

showing that both resolution matrices are symmetric. It is characteristic of damped least-squares that — unlike the truncated SVD approach — all the eigenvalues λ_i^2 get perturbed additively. If we think of the multiplicative factor

$$\frac{\lambda_i^2}{\lambda_i^2 + \mu} \tag{59}$$

as a measure of our ability to resolve the ith eigenvector, then the resolution of $\mathbf{y}_i$ and $\mathbf{z}_i$ is high (unity) when $\mu = 0$ and very low when μ is a large number. So this method degrades the resolution of all eigenvectors simultaneously, though not at the same rate, since the rate is dependent on the ratio μ/λ_i^2. This reduces the impact of the smallest (noise-like) eigenvalues on the final reconstruction, but simultaneously reduces the resolvability of the largest (signal-like) eigenvalues. Checking the other Moore-Penrose conditions, we find that

$$\mathbf{M}\mathbf{X}_\mu\mathbf{M} = \sum_{i=1}^{r} \frac{\lambda_i^3}{\lambda_i^2 + \mu} \mathbf{y}_i \mathbf{z}_i^T \neq \mathbf{M} \tag{60}$$

for any value of $\mu > 0$, and that

$$\mathbf{X}_\mu\mathbf{M}\mathbf{X}_\mu = \sum_{i=1}^{r} \frac{\lambda_i^3}{(\lambda_i^2 + \mu)^2} \mathbf{z}_i \mathbf{y}_i^T \neq \mathbf{X}_\mu = \sum_{i=1}^{r} \frac{\lambda_i}{(\lambda_i^2 + \mu)} \mathbf{z}_i \mathbf{y}_i^T, \tag{61}$$

again for any positive μ. Thus, the damped least-squares method produces an inverse that satisfies only the two symmetry conditions of Moore-Penrose for the resolution matrices.

7.2 Truncated SVD

By contrast, it is easy to see that the truncated SVD inverse satisfies all but one of the Moore-Penrose conditions. Clearly, $\mathbf{X}_q\mathbf{M}\mathbf{X}_q = \mathbf{X}_q$ and both $\mathbf{M}\mathbf{X}_q$ and $\mathbf{X}_q\mathbf{M}$ are symmetric. However, $\mathbf{M}\mathbf{X}_q\mathbf{M} \neq \mathbf{M}$; instead we have

$$\mathbf{M}\mathbf{X}_q\mathbf{M} = \sum_{i=1}^{q} \lambda_i \mathbf{y}_i \mathbf{z}_i^T \equiv \mathbf{X}_q^\dagger. \tag{62}$$

We see that the pseudoinverse of $\mathbf{X}_q$ given by $\mathbf{X}_q^\dagger$ is just a truncated version of $\mathbf{M}$, where a set of the smallest $(r-q)$ nonzero eigenvalues have been discarded. The same truncated SVD may serve as the approximate inverse for many $\mathbf{M}$s, since $\mathbf{X}_t$ does not satisfy the first Moore-Penrose condition. Thus, one of the Moore-Penrose uniqueness conditions is violated.

It should already be clear from the foregoing analysis that the truncated SVD is preferred over the damped least-squares method when accuracy of the inversion is the only concern. Damping shifts all the singular values/eigenvalues away from zero, while truncated SVD preserves the largest singular values and preferentially sets the smallest ones to zero. However, when the matrices involved are very large, it is usually not practical to compute the full SVD of the (ray-path) matrix $\mathbf{M}$; then it may be preferable to sacrifice some accuracy for the sake of computational economy.

8 Other Methods and Alternative Criteria

Many alternative approximate inversion criteria may be considered. We present four of the more important ones in this section: (1) backprojection, (2) the adjoint method, (3) weighted least-squares, and (4) the Backus-Gilbert inverse. None of these methods (except for weighted least-squares with trivial weighting) is equivalent to least-squares or the Moore-Penrose inverse, as we shall see.

8.1 Backprojection formulas

The concept of backprojection is very common in tomographic imaging of all types. First, we measure some kind of line integral (traveltime or decay of intensity) through a body. Then we try to image a relevant property of the body (wave speed or attenuation coefficient) by guessing the distribution of that property (the first guess might be a constant value) and then predicting the values of the various line integrals. The difference between the predicted value and the measured value of each line integral is information we have about values of the body's properties along the line associated with the line integral. If we assume (for lack of other information) that the differences are equally probable to have arisen at all points along the line, then we can "backproject" these differences to obtain an update. These backprojections are done separately for each line and then summed (averaged) to obtain the image from the backprojection technique. This approach is very easy to implement in any application, and gives useful and qualitatively correct images. Many more sophisticated

versions of this idea may also be considered to obtain improved quantitative imaging and inversion of the data.

Let $\mathbf{C}$ be a diagonal $n \times n$ matrix whose diagonal elements are the column sums of $\mathbf{M}$ given by

$$C_{jj} = \sum_{i=1}^{m} l_{ij}, \tag{63}$$

and let $\mathbf{L}$ be a diagonal $m \times m$ matrix whose diagonal elements are the row sums of $\mathbf{M}$ given by

$$L_{ii} = \sum_{j=1}^{n} l_{ij} \tag{64}$$

In seismic tomography, we call $\mathbf{C}$ the coverage matrix and $\mathbf{L}$ the length matrix, since C_{jj} is the total length of those segments of all the ray paths that cross through cell j and L_{ii} is the total length of the ith ray-path through all cells. With these definitions, we have

$$\begin{aligned} \mathbf{M}\mathbf{v} &= \mathbf{L}\mathbf{u} \\ \mathbf{M}^T\mathbf{u} &= \mathbf{C}\mathbf{v}, \end{aligned} \tag{65}$$

with $\mathbf{u}^T = (1, 1, 1, \ldots, 1)$ an m-vector of ones and $\mathbf{v}^T = (1, 1, 1, \ldots, 1)$ an n-vector of ones.

One convenient backprojection formula is

$$s_j = \frac{1}{C_{jj}} \sum_{i=1}^{m} l_{ij} \left(\frac{t_i}{L_{ii}} \right), \tag{66}$$

since the value of the tomographic integral is t_i, which must be spread equally over a ray path of length L_{ii}. The total coverage of the jth cell is C_{jj} and is made up of the m contributions l_{ij}, for $i = 1, \ldots, m$. This is a "path-length weighted" backprojection formula with the weights adding properly to unity, since $\sum_i l_{ij}/C_{jj} = 1$, for $i = 1, \ldots, m$. Then, using the notation introduced in (63) and (64), (66) can be rewritten as

$$\mathbf{s} \simeq \mathbf{C}^{-1}\mathbf{M}^T\mathbf{L}^{-1}\mathbf{t}, \tag{67}$$

where we may assume that the diagonal matrices $\mathbf{C}$ and $\mathbf{L}$ are both invertible, for without loss of generality we may simply remove uncovered cells or vanishing ray paths from the problem until the resulting matrices $\mathbf{C}$ and $\mathbf{L}$ are invertible. From (67), we deduce that the effective (approximate) inverse matrix for this backprojection formula is

$$\mathbf{X} = \mathbf{C}^{-1}\mathbf{M}^T\mathbf{L}^{-1}. \tag{68}$$

The effective resolution matrices associated with this approximate inverse are

$$\mathcal{E}_{data} = \mathbf{M}\mathbf{X} = \mathbf{M}\mathbf{C}^{-1}\mathbf{M}^T\mathbf{L}^{-1} \tag{69}$$

for data resolution and

$$\mathcal{E}_{model} = \mathbf{X}\mathbf{M} = \mathbf{C}^{-1}\mathbf{M}^T\mathbf{L}^{-1}\mathbf{M} \tag{70}$$

for model resolution.

Recalling our definition of $\mathbf{O}$ as a matrix of zeroes, the eigenvalue problem for the matrix $\mathbf{M}$ may be formulated as

$$\begin{pmatrix} \mathbf{O}_{m,m} & \mathbf{M} \\ \mathbf{M}^T & \mathbf{O}_{n,n} \end{pmatrix} \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} = \lambda_i \begin{pmatrix} \mathbf{L} & \mathbf{O}_{m,n} \\ \mathbf{O}_{n,m} & \mathbf{C} \end{pmatrix} \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} \tag{71}$$

and can then be manipulated into the form

$$\begin{pmatrix} \mathbf{L}^{-1}\mathbf{M}\mathbf{C}^{-1}\mathbf{M}^T & \mathbf{O}_{m,n} \\ \mathbf{O}_{n,m} & \mathbf{C}^{-1}\mathbf{M}^T\mathbf{L}^{-1}\mathbf{M} \end{pmatrix} \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} = \lambda_i^2 \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix}. \tag{72}$$

From (72), it is not hard to show that the singular-value decomposition of the effective resolution matrices may then be expressed as

$$\mathcal{E}_{data} = \mathbf{M}\mathbf{X} = \sum_{i=1}^{r} \lambda_i^2 \frac{\mathbf{L}\mathbf{u}_i\mathbf{u}_i^T}{\mathbf{u}_i^T\mathbf{L}\mathbf{u}_i}, \tag{73}$$

and

$$\mathcal{E}_{model} = \mathbf{X}\mathbf{M} = \sum_{i=1}^{r} \lambda_i^2 \frac{\mathbf{v}_i\mathbf{v}_i^T\mathbf{C}}{\mathbf{v}_i^T\mathbf{C}\mathbf{v}_i}, \tag{74}$$

where $r \leq \min(m,n)$ is the rank of $\mathbf{M}$. Recalling (65), we see that one singular value λ_i of (71) is $\lambda = 1$. Since this eigenvalue belongs to an eigenvector having all positive elements, the Perron-Frobenius theorem for nonnegative matrices [Bapat and Raghavan, 1997] says that this eigenvalue is the largest eigenvalue and that the eigenvector associated with this eigenvalue is unique. Therefore, the eigenvalues of (71) are bounded above by unity so that $0 \leq \lambda_i^2 \leq 1$. Using λ_i^2 as a measure of the resolution of the i-th eigenvector, the effective resolution matrices perfectly resolve only homogeneous background eigenvectors (which have the only eigenvalues exactly equal to unity). Other eigenvectors' contributions to the resolution are degraded (from the optimal value of unity) in proportion to the square of their eigenvalues.

Thus, we conclude that this particular choice of backprojection formula does not satisfy any of the Moore-Penrose conditions. However, many other choices of backprojection formula might be considered. It is not hard to think of examples that satisfy some of the Moore-Penrose conditions, *e.g.*, choosing $\mathbf{X} = \alpha\mathbf{M}^T$ where α is some positive scalar produces an approximate inverse with symmetric resolution matrices (see the next subsection).

Since the backprojection formula performs badly in light of the Moore-Penrose conditions, we may choose to consider other criteria to evaluate its performance. For example,

$$\mathbf{X}\mathbf{M}\mathbf{v} = \mathbf{v}, \quad \mathbf{v}^T\mathbf{X}\mathbf{M} = \mathbf{v}^T, \tag{75}$$

and

$$\mathbf{M}\mathbf{X}\mathbf{u} = \mathbf{u}, \quad \mathbf{u}^T\mathbf{M}\mathbf{X} = \mathbf{u}^T, \tag{76}$$

(if satisfied) show that the rows and columns of the resolution matrices all sum to unity. These are all conditions that we would like for the resolution matrices to satisfy, since they all would be satisfied by a perfect inverse (assuming one exists). Such criteria serve to check the normalization of the resolution matrices.

8.2 The adjoint

A particularly simple variant of the backprojection formula is the adjoint approximation

$$\mathbf{X} = \alpha \mathbf{M}^T, \tag{77}$$

where α is a scalar with units chosen to guarantee that $\mathbf{X}$ has the correct dimensions, and its magnitude can be chosen to statisfy some additional criterion. This simple choice satisfies the third and fourth Moore-Penrose conditions, while violating the first two. The Moore-Penrose analysis of this case is elementary and is left to the reader.

8.3 Weighted least-squares inverse

When working with real data, we often find that data quality is far from uniform. If some data are better than other data, we may want to exclude the bad data, or we may to include it but give it less weight than the data we perceive to be more reliable. Such considerations lead to the concept of weighted least-squares inversion. Using our standard problem notation, the data may consist of m measurements t_i for $i = 1, \ldots, m$, the model may consist of n cells each having a constant physical parameter s_j for $j = 1, \ldots, n$ to be determined. We may decide to assign weights $0 \leq W_{ii} \leq 1$ to each term in a least-squares error functional, so we have

$$\Psi_{wls}(\mathbf{s}) = \sum_{i=1}^{m} W_{ii} \left[t_i - (\mathbf{Ms})_i \right]^2 . \tag{78}$$

If $\mathbf{W}$ is the diagonal matrix whose diagonal components are the W_{ii}'s, then it is straightforward to show that the minimum of Ψ_{wls} occurs when

$$\mathbf{M}^T \mathbf{W} \left(\mathbf{t} - \mathbf{Ms} \right) = 0. \tag{79}$$

Equation (79) is the normal equation for weighted least-squares, and the vector (or vectors) $\mathbf{s}$ that satisfy the equation are the solution (or solutions) of the weighted least-squares optimization problem.

Probably the simplest means of analyzing this problem is to note that, if we choose to look at it in a certain way, the problem reduces to standard least-squares analysis. Since the matrix $\mathbf{W}$ is diagonal with nonnegative real numbers along the diagonal, the meaning of its square root $\mathbf{W}^{1/2}$ is clear and we can define a new problem,

$$\tilde{\mathbf{M}}\mathbf{s} = \tilde{\mathbf{t}}, \tag{80}$$

with $\tilde{\mathbf{M}} = \mathbf{W}^{1/2}\mathbf{M}$ and $\tilde{\mathbf{t}} = \mathbf{W}^{1/2}\mathbf{t}$, which we want to satisfy in the least-squares sense. Then, Ψ_{wls} may be equivalently redefined as

$$\Psi_{wls}(\mathbf{s}) = \sum_{i=1}^{m} \left[\tilde{t}_i - \left(\tilde{\mathbf{M}}\mathbf{s} \right)_i \right]^2 , \tag{81}$$

which when minimized leads to the standard normal equations

$$\tilde{\mathbf{M}}^T \left(\tilde{t} - \tilde{\mathbf{M}}\mathbf{s} \right) = 0, \tag{82}$$

having the same solution (or solutions) $\mathbf{s}$ as (79).

Taking the approach of the preceding paragraph, we have reduced the weighted least-squares problem to a modified version of the original least-squares problem and all the previous analysis of least-squares carries over, so we do not need to repeat it here. We might ask a more subtle question however about how and to what extent the weighting scheme has modified the solution of our original problem $\mathbf{Ms} = \mathbf{t}$. Presumably, if the weights are all nonzero *and* we have found a solution such that the minimum value of $\Psi_{wls}(\mathbf{s}) \equiv 0$, then and only then will the solutions (or sets of solutions) be the same in weighted and the unweighted least-squares problems. This situation will virtually never arise in practice, so we may expect there will usually be significant differences between the solutions of the two problems. We will not attempt to analyze this question here, but point out that a careful analysis of this issue would be a valuable addition to the literature. In particular, there are many subtle and important issues yet to be addressed concerning the choice of such weighting schemes and relationships between these schemes and a variety of hybrid methods [Huber, 1977; Scales and Gersztenkorn, 1988] that try to make optimum use of both l_2 and l_1 objective functionals in the presence of very noisy data. Such methods are a part of the broad subject of nonlinear optimization (since we might then choose the weights in a data dependent way) and therefore beyond our present scope.

8.4 Backus-Gilbert inverse

Backus and Gilbert [1968; 1970] formulated an approach to the inverse problem that is especially well-suited to problems in geophysics and imaging of the Earth. They stress the finiteness of the number of data points available for any inversion process, the imprecision of those data, and the (virtually) infinite dimensionality of the continuum quantities to be estimated from the data. Backus and Gilbert make the further assumption in their approach that a "good" Earth model is nevertheless available so the inverse problem may be legitimately linearized with respect to that model.

The mathematical formulation of their approach assumes that the discrete data γ_i are related to an Earth model $E(r)$ and some set of functions $G_i(r)$ that characterize the interaction between the Earth model and the data collected according to an integral relation of the form

$$\int_0^1 G_i(r)E(r)dr = \gamma_i, \quad \text{for} \quad i = 1, \ldots, m. \tag{83}$$

The functions $G_i(r)$ are assumed to be known, resulting from a foward computation using the "good" Earth model. An example of such data might be the frequencies of vibration of either the Earth or the Sun [Lapwood and Usami, 1981; Morelli and Dziewonski, 1987; Vorontsov and Zharkov, 1989], where i is then an index over various modes of vibration.

To characterize the accuracy and resolution of the final result, Backus and Gilbert introduce a function $A(r; r_0)$ with the desired properties that

$$\int_0^1 A(r; r_0)dr = 1 \quad \text{and} \quad \int_0^1 A(r; r_0)E(r)dr \simeq E(r_0). \tag{84}$$

So A is intended to act much like a Dirac delta function $\delta(r - r_0)$. The goal of the method is to try to construct such an A from the known functions G_i in the form of an expansion

$$A(r; r_0) = \sum_{i=1}^{m} a_i(r_0) G_i(r). \tag{85}$$

Substituting (85) into (84) shows that, in terms of the measured data, we must have one constraint on the a_i's

$$\sum_{i=1}^{m} a_i(r_0) \int_0^1 G_i(r) dr = 1 \tag{86}$$

and an approximate formula for the Earth model in terms of the γ_i's

$$\sum_{i=1}^{m} a_i(r_0) \gamma_i \simeq E(r_0). \tag{87}$$

In the Backus-Gilbert theory, the set of equations (83) plays the same role as $\mathbf{Ms} = \mathbf{t}$ in our previous discussions, while equation (87) plays the same role as $\mathbf{XMs} = \mathbf{Xt}$. Thus, the $a_i(r_0)$'s play the role of a single row of an approximate inverse operator $\mathbf{X}$, such that $\sum X_{ji} \gamma_i = E_j \equiv E(r_0)$. We will therefore define the vector $\mathbf{a}^T = (\, a_1, a_2, \ldots, a_m \,)$ as a row vector of the Backus-Gilbert approximate inverse $\mathbf{X}$.

So far we do not have enough constraints to estimate the $a_i(r_0)$'s. Making use of the idea that A should resemble a delta function, Backus and Gilbert introduce a spread function

$$S(A; r_0) = 12 \int_0^1 (r - r_0)^2 A^2(r; r_0) dr. \tag{88}$$

If A is a delta function $\delta(r - r_0)$, this integral vanishes. So S is a measure of the deviation of A from a delta function. Substituting the expansion (85) of A in terms of the $G_i(r)$'s into (88), we see that S is given by

$$S(A; r_0) = \sum_{i=1}^{n} \sum_{j=1}^{m} a_i(r_0) a_j(r_0) N_{ij}(r_0), \tag{89}$$

where the matrix $N_{ij}(r_0)$ is given by

$$N_{ij}(r_0) = 12 \int_0^1 (r - r_0)^2 G_i(r) G_j(r) dr. \tag{90}$$

One row of the Backus and Gilbert inverse $\mathbf{X}$ is then found by minimizing (89) with respect to the a_i's, subject to the constraint (86). If we define

$$c_i = \int_0^1 G_i(r) dr, \tag{91}$$

then it is not difficult to show that

$$\mathbf{a}(r_0) = \left(\mathbf{c}^T \mathbf{N}^{-1}(r_0) \mathbf{c} \right)^{-1} \mathbf{N}^{-1}(r_0) \mathbf{c}, \tag{92}$$

where $\mathbf{N}$ is the matrix whose elements are the N_{ij}'s, $\mathbf{a}$ is the vector whose elements are the a_i's, $\mathbf{c}$ is the vector whose elements are the c_i's. We have assumed that $\mathbf{N}$ is invertible. When it is invertible, satisfaction of the constraint (86) by the result (92) is then straightforward to check. Otherwise, we must again invoke some approximation method to solve the Backus-Gilbert inversion problem.

The Moore-Penrose analysis for Backus-Gilbert is left to the reader, but a summary of the results is provided in the next section.

9 Comparisons of Approximate Inverses

We have seen that the various approximate inverses can be partially characterized by noting which and how many of the Moore-Penrose conditions they satisfy. We now present a summary of these results in Table 1.

TABLE 1. Qualitative comparison of various approximate inverses. Note that many choices of backprojection method are possible, so in general "backprojection" may not satisfy any of the Moore-Penrose conditions. It is possible to choose a backprojection method that guarantees satisfaction of some subset of the conditions; the adjoint is an example of this feature. Least-squares produces a solution equivalent to that of the pseudoinverse $\mathbf{M}^\dagger$ and is therefore represented by the first entry. Weighted least-squares methods produces the pseudoinverse solution of a slightly modified problem, which for our present purposes is also represented by the first entry.

Approximate Inverse	Moore-Penrose Conditions			
	$\mathbf{MXM} = \mathbf{M}$	$\mathbf{XMX} = \mathbf{X}$	$(\mathbf{XM})^T = \mathbf{XM}$	$(\mathbf{MX})^T = \mathbf{MX}$
$\mathbf{M}^\dagger$(SVD)	yes	yes	yes	yes
Truncated SVD	no	yes	yes	yes
Damped l_2	no	no	yes	yes
Adjoint	no	no	yes	yes
l_1	yes	yes	yes	no
Backus-Gilbert	no	yes	no	no
Backprojection	no	no	no	no

In Table 2 we preview results from the subsequent discussion in Part II [Berryman, 1999] on iterative methods. Note that all the iterative methods of matrix inversion lead directly to approximations of the Moore-Penrose pseudoinverse $\mathbf{M}^\dagger$. In general, carrying these iterative processes to completion will (at least in principle) produce a result closely approximating $\mathbf{M}^\dagger$. Incomplete or partially completed iteration leads to results that are

generally qualitatively similar to a truncated SVD, as we will show in Part II. The most startling result found in this analysis is the fact that the LSQR algorithm of Paige and Saunders [1982] does *not* satisfy the second Moore-Penrose condition, $\mathbf{XMX} = \mathbf{X}$, when the data vector $\mathbf{t}$ and the ray-path matrix $\mathbf{M}$ are inconsistent. We find that a relatively straightforward modification of LSQR (using a different choice of starting vector) may be introduced to correct this flaw.

TABLE 2. Qualitative comparison of various iterative inverses.

Iterative Inverse	Moore-Penrose Conditions			
	$\mathbf{MXM} = \mathbf{M}$	$\mathbf{XMX} = \mathbf{X}$	$(\mathbf{XM})^T = \mathbf{XM}$	$(\mathbf{MX})^T = \mathbf{MX}$
Fully iterated CD, CG, Lanczos, or Modified LSQR	yes	yes	yes	yes
Full LSQR Inverse	yes	no	yes	yes
Partial Lanczos, or Modified LSQR	no	yes	yes	yes
Partial Conjugate Directions/gradients	no	yes	no	yes
Partial LSQR Inverse	no	no	yes	yes
Damped Lanczos or LSQR	no	no	yes	yes

Acknowledgments

I thank S. F. Ashby, R. Burridge, J. F. Claerbout, S. Fomel, D. M. Goodman, J. B. Keller, F. Muir, G. Nolet, G. Papanicolaou, B. N. Parlett, S. Ronen, M. A. Saunders, W. W. Symes, H. A. van der Vorst, and J. J. Zucca for helpful discussions, lectures, collaborations, correspondence and/or criticism. This work was performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48 and supported by the Environmental Management Science Program with oversight by the Engineering Research Division of the Office of Basic Energy Sciences, Office of Energy Research, and by the Office of Science and Technology, Office of Environmental Management. Early work in this area was performed while the author was visiting the Geophysics Department at Stanford University. Partial support for this visit was provided by the Stanford Exploration Project, whose sponsors are hereby gratefully acknowledged.

References

Aki, K., and P. Richards, 1980, *Quantitative Seismology: Theory and Methods*, Freeman, San Francisco, California.

Backus, G. and F. Gilbert, 1968, The resolving power of gross earth data, *Geophys. J. R. Astron. Soc.* **16**, 169–205.

Backus, G., and F. Gilbert, 1970, Uniqueness in the inversion of inaccurate gross earth data, *Philos. Trans. R. Soc. London* **266A**, 123–192.

Bapat, R. B., and T. E. S. Raghavan, 1997, *Nonnegative Matrices and Applications*, Cambridge University Press, Cambridge.

Berryman, J. G., 1990, Stable iterative reconstruction algorithm for nonlienar traveltime tomography, *Inverse Problems* **6**, 21–42.

Berryman, J. G., 1999, Analysis of approximate inverses in tomography. II. Iterative inverses, *Optimization and Engineering*, submitted.

Bording, R. P., A. Gersztenkorn, L. R. Lines, J. A. Scales, and S. Treitel, 1987, Applications of seismic travel-time tomography, *Geophys. J. R. Astron. Soc.* **90**, pp. 285–303

Claerbout, J. F., 1976, *Fundamentals of Geophysical Data Processing with Applications to Petroleum Prospecting*, McGraw-Hill, New York, pp. 123–129.

Claerbout, J. F., and F. Muir, 1973, Robust modeling with erratic data, *Geophysics* **38**, 826–844.

Dines, K. A., and R. J. Lytle, 1979, Computerized geophysical tomography, *Proc. IEEE* **67**, 1065–1073.

Eckart, C., and G. Young, 1936, The approximation of one matrix by another of lower rank, *Psychometrika* **1**, 211–218.

Feynman, R. P, R. B. Leighton, and M. Sands, 1963, *The Feynman Lectures on Physics*, Vol. I, Addison-Wesley, Reading Massachusetts, Chapt. 26.

Golub, G. and W. Kahan, 1965, Calculating the singular values and pseudo-inverse of a matrix, *SIAM J. Numer. Anal.* **2**, 205–224.

Golub, G. H., and C. F. Van Loan, 1989, *Matrix Computations*, Johns Hopkins University Press, Baltimore, p. 341.

Hestenes, M. R., and E. Stiefel, 1952, Methods of conjugate gradients for solving linear systems, *J. Res. Nat. Bur. Stan.* **B 49**, 409–436.

Huber, P. J., 1977, *Robust Statistical Procedures*, SIAM, Philadelphia, Pennsylvania, 1977 and 1996.

Jackson, D. D., 1972, Interpretation of inaccurate, insufficient, and inconsistent data, *Geophys. J. Roy. Astron. Soc.* **28**, 97–109.

Lanczos, C., 1950, An iterative method for the solution of the eigenvalue problem of lienar differential and integral operators, *J. Res. Nat. Bur. Stand.* **45**, 255–282.

Lanczos, C., 1961, *Linear Differential Operators*, SIAM, Philadelphia, Pennsylvania, 1961 and 1996, pp. 120–127.

Lapwood, E. R., and T. Usami, 1981, *Free Oscillations of the Earth*, Cambridge University Press, Cambridge.

Levenberg, K., 1944, A method for the solution of certain non-linear problems in least-squares, *Quart. Appl. Math.* **2**, 164–168.

Marquardt, D. W., 1963, An algorithm for least-squares estimation of nonolinear parameters, *SIAM J. Appl. Math.* **11**, 431–441.

Marquardt, D. W., 1970, Generalized inverses, ridge regression, biased linear estimation, and nonlinear estimation, *Technometrics* **12**, 591–612.

Moore, E. H., 1920, *Bull. Amer. Math. Soc.* **26**, 394–395.

Morelli, A., and A. M. Dziewonski, 1987, The harmonic expansion approach to the retrieval of deep Earth structure, in *Seismic Tomography*, edited by G. Nolet, Reidel, Dordrecht, pp. 251–274.

Munk, W., P. Worcester, and C. Wunsch, 1995, *Ocean Acoustic Tomography*, Cambridge University Press, Cambridge, pp. 239–271.

Nolet, G. (ed.), 1987, *Seismic Tomography: With Application in Global Seismology and Exploration Geophysics*, Reidel, Dordrecht.

Paige, C. C., and M. A. Saunders, 1982, LSQR: An algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Softw.* **8**, 43–71.

Penrose, R., 1955a, A generalized inverse for matrices, *Proc. Cambridge Philos. Soc.* **51**, 406–413.

Penrose, R., 1955b, On best approximation solutions of linear matrix equations, *Proc. Cambridge Philos. Soc.* **52**, 17–19.

Rector, J., 1995, Crosswell methods: Where are we, where are we going? *Geophysics* **60**, 629–630.

Scales, J. A., P. Docherty, and A. Gersztenkorn, 1990, Regularisation of nonlinear inverse problems: Imaging the near-surface weathering layer, *Inverse Problems* **6**, 115–131.

Scales, J. A., and A. Gersztenkorn, 1988, Robust methods in inverse theory, *Inverse Problems* **4**, 1071–1091.

Scales, J. A., A. Gersztenkorn, and S. Treitel, 1988, Fast l_p solution of large, sparse, linear systems: Application to seismic travel time tomography, *J. Comput. Phys.* **75**, 314–333.

Tikhonov, A. N., and V. Y. Arsenin, *Solution of Ill-Posed Problems* (Winston, New York, 1977).

Vorontsov, S. V., and V. N. Zharkov, 1989, Helioseismology: Theory and interpretation ofexperimental data, *Sov. Sci. Rev. E Astrophys. Space Phys.* **7**, 1–103.

Whitham, G. B., 1974, *Linear and Nonlinear Waves*, Wiley, New York, Chapt. 7, pp. 247–250.

Wiggins, R. A., 1972, The general linear inverse problem: Implications of surface waves and free oscillations for Earth structure, *Rev. Geophys. Space Phys.* **10**, 251-285.

Yorkey, T. J., J. G. Webster, and W. J. Tompkins, 1987, Comparing reconstruction algorithms for electrical impedance tomography, *IEEE Trans. Biomed. Engng.* **BME-34**, 843–52.