

INVERSION FOR APPLIED GEOPHYSICS: A TUTORIAL

Douglas W. Oldenburg* and Yaoguo Li**

** UBC-Geophysical Inversion Facility*

Department of Earth and Ocean Sciences, University of British Columbia, Vancouver, Canada, V6T 1Z4

*** Department of Geophysics, Colorado School of Mines, Golden, Colorado, 80401*

INTRODUCTION

Throughout this book there are numerous cases where geophysics has been used to help solve practical environmental, geotechnical and exploration problems. The typical scenario is first to identify the physical property that is diagnostic of the sought geologic structure or buried object, for example density, seismic velocity, electrical conductivity, or magnetic susceptibility. The appropriate geophysical survey is then designed and field data are acquired and plotted. In some cases the information needed to solve the problem may be obtained directly from these plots, but in most cases more information about the subsurface is required. As an example, consider the magnetic field anomaly map presented in Figure 2. The existence of a buried object, and also approximate horizontal locations, can be inferred directly from that image. The map however, presents no information about the depth of the object or details regarding its shape. To obtain that information the data need to be inverted to generate a 3D subsurface distribution of the magnetic material.

Geophysicists have long recognized the need for inverting their data, but the state of computing resources a few decades ago meant that even the best inversion algorithms could recover estimates for only a small number of parameters. Consequently, the subsurface was modelled as a few layers, or as simple prismatic objects, and parameters were adjusted so that numerically modeled data fit the observations as well as possible. Unfortunately, earth structures are generally much more complicated and hence incorrect conclusions were often drawn. In the intervening years, there have been considerable advances in computing power and in the mathematics of inversion. As a result, more realistic inversions are being carried out and the usefulness of geophysics in solving applied problems is increasing. There is an increasing realization that geophysics, which is the only remote sensing discipline available for subsurface problems, must be an integral part of environmental, engineering and exploration problems. There is also a realization that extracting the information about the subsurface from the geophysical data is not a turn-key operation, and that cost-effective solutions generally involve a team of experts because the quality of images obtained from geophysical inversion depends critically upon other information. There is a bright future for the use of geophysics but it requires that inversion be applied with care and understanding. Anyone using inversion algorithms, or interpreting resultant images, must be familiar with the basic principles. Our goal is to present those principles, discuss in detail the essential elements of the inverse problem, and illustrate the techniques through synthetic and field examples.

There are two types of problems that are regularly referred to as “inverse” problems. The most common of these is concerned with obtaining estimates for a few parameters when there are many more data than unknowns; we

refer to this as the parameter estimation problem. These problems are generally solved with least squares approaches and there are a substantial number of books and papers available on this topic (e.g., Lawson and Hanson, 1974; Lines and Treitel, 1984). We treat this problem here but our main emphasis is on a second type of problem where our goal is to find a function, usually a physical property distribution. This will be our generic inverse problem.

Inverse theory is an exceedingly large topic and we cannot cover all aspects in depth in a limited document. We intend to outline the important elements for solving practical inverse problems. It is crucial to understand that achieving an interpretable solution from an inversion algorithm is not a “black-box” or “turn-key” operation. It is often easy to obtain an image from an inversion algorithm, but the challenge is to obtain the “best” image possible so that geophysics contributes maximally toward the solution of the particular problem. Quality interpretations require that the interpreter understand fundamentals regarding nonuniqueness of the solution, how prior information is used to “regularize” the problem, and the how well data should be fit. In this paper, the term “fitting” the data refers to how well the simulated data, corresponding to a recovered model, match the observations.

Of foremost importance is the understanding of nonuniqueness of the inverse problem. In Section 1 we use a simple seismic example to illustrate the nonuniqueness that occurs when we try to find a function. This clearly shows that additional information is always required if we are to compute a single answer to our problem. In Section 2 we outline a general methodology for solving the inverse problem as an optimization problem. A tutorial example for solving a linear inverse problem, and an interactive JAVA applet, are provided to help consolidate the concepts. The deleterious effects of overfitting or underfitting the data are explored and related to the Tikhonov curve, which plays a fundamental role in our inversion methodology. Examples for linear inversion include 2D cross-well tomography, 2D induced polarization (IP), and the 3D gravity problem. The more common, nonlinear inverse problem, is treated in Section 3 and the 2D DC resistivity problem is used as a prototypical example. The quality of the final inversion image depends greatly upon the geologic and geophysical prior information. Section 4 presents some typical prior information and illustrates how it can be incorporated. The 3D magnetic problem is used as an illustration. In Section 5 we confront the problem of estimating the noise level in the data, or equivalently, choosing a value of the regularization parameter. This is a crucial component of the inversion because it controls the amount of structural detail available in the final image. Section 6 treats the parameter estimation problem. We explore both linearized techniques and direct search algorithms that attempt to find a global minimum of a misfit function. As an example we choose the problem of using time domain electromagnetics to determine if a buried target is likely to be an unexploded ordnance (UXO). In Section 7 we look at the problem of estimating uncertainties in the recovered models. The section is logically divided into two parts to treat the parameter estimation problem, and the underdetermined inverse problem. Our chapter concludes with a commentary section in which we look at the inverse problem within the wider context of solving an applied problem and discuss elements of cost effectiveness and the necessary ingredients for success.

The focus of our chapter is on understanding the fundamentals, and we have tried to make the chapter useful to an audience with a wide range of mathematical backgrounds. Even for those with very limited mathematical background we hope the verbal explanations, computer applet, synthetic and field examples, and overview comments will be useful.

Before proceeding, we want to emphasize that this is a tutorial chapter and not a review. We have made no attempt to incorporate, and reference, all of the great papers and books that have been written on this subject. Our examples have primarily been drawn from our own work. This is not because we think our work is better than others, but it reflects the pragmatic availability of figures and background information about the application. We apologize at the outset for the extensive self-referencing and for omission of explicit references to other important

literary works.

Section 1.0: FUNDAMENTALS OF INVERSION

A geophysical remote sensing experiment is typified by the illustration in Figure 1. It consists of an energy source, an earth model, and a response. The energy source can be natural, as is the case for gravity, magnetics, and magnetotelluric surveys, or it can be a man-made device, such as a vibrator or explosive source used in seismic studies, or an electric current used in electromagnetic surveys. The earth model is characterized by 3D distributions of physical properties and the responses can be fields that are measured on or above the surface of the earth or in boreholes. The propagation of energy through the earth depends upon the 3D distribution of one or more physical properties, i.e., each datum is sensitive to a property variation in a volume. This is good news, because the datum therefore contains information about the earth away from the receiver location. This is what makes geophysics a valuable companion to direct sampling in a borehole which provides information only about a pencil-thin line into 3D space. It is also a complicating factor because there is not a one-to-one correspondence between a datum and the value of a physical property at a particular location in the earth. Because each datum is sensitive to what is happening in a volume, it should not be expected that a data image can directly provide localized information about the subsurface geology. Any data image represents a convoluted picture of the earth, and direct interpretation is therefore usually qualitative, difficult, and highly ambiguous. It is both desirable and necessary to carry out the interpretation based upon a more direct representation of the subsurface physical properties. Geophysical inversion is required to extract that information from the data.




Figure 1 The main components of a geophysical remote sensing experiment. Input, which is the energy source; earth filter, which is the distribution of physical properties; and the output, which is the measured value of physical fields.

Before tackling the inverse problem we need to be able to solve the “forward” problem. With reference to Figure 1, the forward problem involves calculating the responses under the assumption that the sources and the earth model are known. As an example, in Figure 2, a 3D susceptibility structure, magnetized by the earth’s field, generates magnetic field data at the surface. The total field anomaly is shown.

Data from a geophysical experiment can be generically written as

$$F_j[m] = d_j + n_j \quad j = 1, \dots, N \quad (1)$$

where F is a forward modelling operator that incorporates details of the survey design and the relevant physical equations, m is a generic symbol for a physical property distribution, and the right hand side represents the observed data d^{obs} which consist of the true data d plus additive noise n . F comes in a variety of forms, but most often it is an integral or differential operator, and the resultant equations must be solved numerically. Although this may entail

significant effort in itself, for example, computing electromagnetic responses for a 3D earth, forward modelling is a well-posed problem and there is a unique answer.

In a geophysical survey we acquire N data, d^{obs} , and some knowledge about their uncertainties. The goal of the inverse problem is to find the model m (susceptibility in the case of Figure 2) that produced the noisy observations. Shown schematically in Figure 3, this is a much more difficult process than forward modelling. The primary reason is that each geophysical datum depends upon a volumetric distribution of a physical property and information about the property is encoded in the data in a complex way. Secondly, it is unrealistic to expect that we can determine a 3D physical property distribution uniquely when we have only a few data. Even if each datum provided the precise value of a physical property at some location, then we have only N such values. That is not enough to provide the solution everywhere. It follows that any solution we obtain from an inverse algorithm must be nonunique, and if we find one solution that “acceptably” fits the data, there are infinitely many others that will fit just as well. Selection of a single (best guess) answer will require additional information.




Figure 2 In the forward problem, we compute the synthetic data that arise from a particular set of sources and receivers when the distribution of physical properties is known. In this example the earth’s main magnetic field $\vec{B}_0$ is the source, the model is the 3D distribution of magnetic susceptibility and the data are the total-field anomaly.




Figure 3 The goal of the inverse problem is to recover the physical property distribution that gave rise to the data. Here, we would like to recover the earth's magnetic susceptibility from the noise contaminated total field anomaly data measured at the surface.

The practical importance of nonuniqueness requires that we look at this issue directly. The physical property models of interest can be 1D, 2D, or 3D functions of space. Since any function has infinitely many degrees of freedom, or equivalently it requires infinitely many Fourier coefficients to specify it, then infinitely many data are required. The last step follows from the idea that each datum provides one piece of information and, at best, it could supply one numerical value of a Fourier coefficient. In practise, an infinite amount of perfectly accurate data never exist, but if it did, then a unique solution may exist. Some examples are: (1) 1D seismic refraction problem where the velocity increases continuously with depth. Knowledge of travel time $t(x)$ for all offsets x is sufficient to recover $v(z)$ down to a critical depth dependent upon the largest offset using Herglotz-Wiechert formula (see Aki and Richards, 2002); (2) In a DC resistivity experiment above a 1D electric conductivity structure, if the potential $V(x)$ due to a point source is known at all offsets x from a current source, then the 1D conductivity can be recovered uniquely from the potential (Langer, 1933). When the conductivity structure is 3D, it can be uniquely reconstructed from the continuous measurements of the electric potentials over the surface of the earth for every possible location of current source along a curve spanning the surface (Stevenson, 1934); (3) For an electromagnetic experiment above a 1D earth, if one knows the response at one location for all frequencies, or at one frequency but at all locations from the source, then it is possible to uniquely recover the electrical conductivity structure (Weidelt, 1972).

The above examples are a lead-in to the general principle that, in order to get a unique solution, the dimensionality of the data must match the dimensionality of the physical property one is trying to recover. For example, if the physical property varies with three independent variables, then the required data must also vary as a function of three independent variables. Note that for the magnetic survey, even knowing data everywhere on a plane would not be sufficient to generate a unique 3D distribution of susceptibility since the data only vary with two independent variables.

The theoretical results provide some comfort that, in principle, geophysical surveys have the ability to discern important information about the earth structure. Although algorithms based upon analytic inversions are rarely of practical use, and the theoretical uniqueness doesn't appear to play any role in practice, the uniqueness statement regarding dimensionality of data impacts upon how much information must be supplied by the user to obtain a geologically meaningful model for interpretation. For example, if a theoretical uniqueness exists, then by acquiring

sufficient data, and by imposing certain continuity conditions, such as requiring the model to be smooth to some extent, then inversion may be able to recover larger scale features of the earth. As we will see, this is true for 2D DC resistivity inversion. On the other hand, the lack of theoretical uniqueness in potential-field problems means that we must impose further conditions in order to produce meaningful models.

There is often much to be learned from a simple example, in which theoretical uniqueness has been established, because we can use that as a starting position to evolve to the geophysical reality where we have only a finite number of inaccurate data. Studying this yields fundamental understanding about non-uniqueness that applies in most of our problems.

1.1: Tutorial example: Recovering velocity from RMS data

We consider a 1D velocity structure $v(z)$. After conversion to time, the velocity $v(t)$ is related to the RMS velocity $V(t)$ by

$$V(t) = \left(\frac{1}{t} \int_0^t v^2(u) du \right)^{1/2} \quad (2)$$

This equation has a unique inverse

$$v(t) = V(t) \left(1 + \frac{2tV'(t)}{V(t)} \right)^{1/2} \quad (3)$$

where $V'(t) = dV(t)/dt$. So, if exact data $V(t)$ are known for all t then $v(t)$ can be found uniquely. Eq.(2) occurs in reflection seismic processing. Stacking velocities are obtained, and under suitable conditions, are equated to the RMS velocities. Discretization of eq.(3) yields Dix's formula.

The above two equations provide considerable insight about the nature of the forward and inverse problems that are generic to most problems encountered in geophysics. First we consider eq.(2). The RMS velocity at time t is an average of the interval velocity (squared) over a time interval $[0, t]$. Taking averages is a smoothing operation. We can see how the forward operator is a “smoother” by considering a velocity $v(t) = v_0 + v_1 \sin kt$. The plot of model $v(t)$ and data $V(t)$ are provided in Figure 4. The amplitude of the oscillations in the data decreases with increasing time. This is due to the increasing window width over which the average value is taken. For very large time the RMS velocity approaches the average velocity of $2000m/s$.

Eq.(3) states that the inverse operator can recover the model from the data. But if the forward problem has “smoothed”, then the inverse problem must “roughen”. The roughening in eq.(3) comes from the term $tV'(t)$. Differentiation is intrinsically a roughening process. This differentiation, plus amplification by the value of t , is sufficient to magnify the small fluctuations in the data at large time so that the original model is recovered. This is illustrated in Figure 4 where the input data in Figure 4(b) have been inverted to produce the velocity in Figure 4(c). That velocity is identical with the true velocity in Figure 4(a).

The exact inverse requires that data $V(t)$ be supplied for all times t . It follows immediately that if only a finite number of data, $V_j, j = 1, \dots, N$, are provided then there must be other velocities that fit the data exactly. These can be discovered by interpolating the data in different ways and using the exact inverse. In Figure 5 we show two models that fit 21 data exactly. A smooth velocity is constructed using a cubic spline interpolation, and a blocky model with linear gradients is found by using a linear interpolation. Since there are infinitely many ways to interpolate data, there must be an infinite number of solutions to the problem. This nonuniqueness is exacerbated




Figure 4 Inversion with complete and accurate data. (a) Velocity structure (b) RMS data $V(t)$, and (c) recovered velocity.

when the data are inaccurate, because now there are even more ways to interpolate the data. This example therefore shows that the nonuniqueness arises for two reasons: the data are finite in number and inaccurate.

There is another important aspect about the inverse problem that is easily seen from the RMS velocity example. The inverse problem is ill-conditioned or unstable. That is, arbitrarily small errors in the data can still generate arbitrarily large errors in the recovered model. This is illustrated in Figure 5. Gaussian random noise with a standard deviation of 20 m/s has been added to the data and the noisy data in Figure 5(c) have been interpolated with a cubic spline. The resultant velocity model in Figure 5(d) displays increasing error with time. In fact, the inversion result can be so inaccurate that it is physically unrealistic. This is observed at 1.7 s, where the recovered velocity is negative. A question that arises is: “how can relatively small perturbations to the data produce such disastrous consequences?” The answer lies in the fact that the inverse operator is a roughening operator. To examine this more quantitatively, suppose that the earth had a constant velocity $v(t) = v_0$ so that the data are $V(t) = v_0$. Allow a perturbation to the data in the form $a \cos(\omega_0 t)$ where $a \ll 1$ and ω_0 is the frequency of the perturbation. Using the exact inverse yields

$$v(t) = v_0 \left(1 + \frac{a t \omega_0 \cos \omega_0 t}{v_0} \right)^{1/2} \quad (4)$$

Thus the recovered model has an error amplitude that depends upon $a t \omega_0 / v_0$. Irrespective of how small a is, the error in the recovered model can be made arbitrarily large by increasing the frequency or time. The procedure is inherently unstable.




Figure 5 (a) Finite data set with two interpolations providing complete data, (b) subsequent velocities, (c) finite noisy data with interpolation fitting the points and misfitting the data, (d) resultant velocity model.

The ideas presented for the RMS inverse problem are generic to problems encountered in applied geophysics. To summarize, all geophysical forward modellings are “smoothing” operators and each datum is affected by a spatial distribution of a physical property. Each datum thus provides one “average” of the model, or one piece of information about the model. Yet each physical property distribution is a function and, as such, it has infinitely many degrees of freedom, and requires infinitely many pieces of information to specify it completely. In practise, the inverse problem will be nonunique because the data are finite in number and they are inaccurate. Moreover, the inverse problem is unstable in that small changes in the data can manifest themselves as large structures in the recovered model. These two issues are central to the development of all practical solutions of the inverse problem.

The fact that the inverse problem is nonunique and unstable means that it is an ill-posed problem (Tikhonov and Arsenin, 1977). A reformulation is required and all reformulations require that additional information, or prior knowledge about the model, be supplied. The final methodology depends upon what information is available and how it is represented. If the information about the model is characterized by known probabilities, then statistical formulations, as detailed in books by Tarantola (1987), Gelman et al. (1995), and Scales and Smith (1996), provide a good framework for solving the inverse problem. This approach treats both data and model parameters as random variables. Probability density functions for these quantities must be provided, and a theorem by Bayes is used to generate the *a posteriori* probability density function (PPD) for the model parameters. The model that maximizes the PPD is computed, and this maximum a posteriori probability (MAP) model, found by solving an optimization problem, is used for interpretation purposes.

In our approach, the incorporation of prior information is done by designing a model objective function (often called a regularization functional) that is directly related to the structural characteristics of the model. Different types of prior information can be included by altering the details of the model objective function. Out of all possible models that fit the data, we then select that one whose structural measure is minimum. The inverse problem is

solved by casting it as an optimization problem. There has been a tremendous amount of work done on solving the inverse problem in this manner. Good books include Tikhonov and Arsenin (1977), Parker (1994), Menke (1989), and Hansen (1998). In what follows, we present the basic ideas, but more detail and mathematical rigor is provided in those books, and in other literature regarding optimization, such as Nocedal and Wright (1999), Kelley (1999), and Vogel (2001).

Section 2: LINEAR INVERSION: AN OPTIMIZATION APPROACH

The data provide N inaccurate constraints upon the earth model and there are infinitely many models that acceptably fit those data. To find a particular solution we need some way of further evaluating each of these models so that we can distinguish between them. We need to define a “ruler”, or “norm” for the functions we are working with. This allows us to measure the “length” of each element and thereby select the smallest, or in some cases, the largest. The procedure is completely analogous to the challenge of selecting an individual from a group of persons. Some criterion for measurement is required, and it could be physical height, but it could be IQ, foot size, or a myriad of other quantities. Having decided upon the “ruler”, then the candidates can be ranked, and the smallest one can be selected. Clearly, the choice of “ruler”, or norm, is critical to the result.

In most problems we are interested in how the physical property varies in the spatial directions or how much “structure” it has. From an interpretation viewpoint it is deviations of the model away from a background earth that are likely of interest. Because of the ill-posedness of the inverse problem we can find models that are arbitrarily complicated, but it is not possible to make them arbitrarily simple. For instance a uniform halfspace will not generally fit a set of field observations. This motivates the goal of designing a norm that is associated with structure and finding the model that minimizes this quantity. It is hoped that features seen on this “simplest” model represent those in the earth.

A main consideration in designing the norm is that this is the vehicle by which prior information is included. For instance, if the earth is “blocky”, and thus characterized by boundaries across which there are large changes in the physical property, or if the distributions are expected to change smoothly, then the norm should be tailored to produce that type of model. Also, in most cases, some information about a background model is available through geophysical surveys or drill hole information, and that should be incorporated. Our goal is to find a solution that has the right character, is “close” to a background value, has a minimum of structure, and also satisfies any other prior information we might have about the model. We’ll elaborate more upon prior information and how to include it, but for now we introduce a flexible model objective function for 1D problems,

$$\phi_m(m) = \alpha_s \int (m - m_{ref})^2 dz + \alpha_x \int \left(\frac{dm}{dx} \right)^2 dx \quad (5)$$

and use this quantity as our norm for measuring the size of any m . The first term penalizes discrepancies between the constructed model and a reference model m_{ref} . If this were the only term in ϕ_m then the solution to the inverse problem would be a model that is as close as possible (using a squared, or l_2 norm) to the reference model. This is especially valuable when a good estimate of the background exists. The second term penalizes first order roughness in the x -direction and its role is to keep the constructed model as flat as possible, (ie. close to a constant value). The coefficients α_s, α_x can be altered to provide greater emphasis on closeness to a reference model or flatness.

Our inverse problem will be solved by finding the model that minimizes the model objective function and also “fits” the data. The concept of “fitting” the data means that some estimate of the “noise” be available.

Unfortunately “noise” within the context of inversion is everything that cannot be accounted for by a compatible relationship between model and data. That is $d^{obs} = F[m] + n + \delta$ where n might be additive noise such as that related to measurement errors, but δ represents other errors associated with discrepancies between the mathematical representation of the earth and the real earth. These arise when representing a 3D earth by a 2D approximation, or representing a continuous variation of physical parameters by piecewise constant functions. n might also include the effects of inaccurate locations, such as the electrode positions in a DC survey, or of other errors in source/receiver parameters. Generally none of the above mentioned errors are known so we proceed by assuming they are Gaussian. Sometimes covariance estimates are available but even this is not common. In the following analysis, we will assume that the data errors are Gaussian, independent and that the standard deviation of the j' th datum is ϵ_j . An associated measure of misfit is

$$\phi_d = \sum_{j=1}^N \left(\frac{d_j^{obs} - F[m]}{\epsilon_j} \right)^2 = ||W_d(d^{obs} - F[m])||^2 \quad (6)$$

where $W_d = \text{diag}(1/\epsilon)$.

Having defined a measure of misfit, the next task is to determine a tolerance value, such that if the misfit is about equal to that value, then we have an acceptable fit. Suppose that the standard deviations are known and that errors are Gaussian. Then ϕ_d becomes the chi-squared variable χ_N^2 with N degrees of freedom. This is a well known statistical quantity whose expected value is $\mathcal{E}[\chi_N^2] = N$ and standard deviation is $\sqrt{2N}$ (Parker, 1994). The inverse problem can now be formulated as follows: Find a model m that

$$\begin{aligned} &\text{minimize} && \phi_m \\ &\text{subject to} && \phi_d = \phi_d^* \end{aligned} \quad (7)$$

where ϕ_d^* is a target misfit. Setting $\phi_d^* \sim N$ should provide a good solution providing that our assumptions about the errors are valid.

The inverse problem in eq.(7) is solved as an unconstrained optimization problem. We minimize

$$\phi(m) = \phi_d + \beta \phi_m \quad (8)$$

where $0 < \beta < \infty$ is a constant and is generally known as the regularization parameter or tradeoff parameter. In eq.(8) it controls the relative weighting between ϕ_m and ϕ_d . As β approaches zero the function to be minimized consists only of the misfit ϕ_d . The model constructed will fit the data as well as possible, irrespective of the amount of structure that is built into the model. Conversely, if β becomes large, then the recovered model will make ϕ_m as small as possible (often equal to zero) and the misfit will be large. Clearly neither of these extremes is useful, and the role of β is to balance the relative importance of the two terms. This can be achieved by finding a value of β such that $\phi_d(\beta) = \phi_d^*$, which is the original constraint in eq.(7). That is, we have found a model that minimizes our model norm and also fits the data to a reasonable amount. This approach of choosing the β based upon the expected data misfit is called the discrepancy principle. We note that no other model will fit the data to this amount and have a smaller norm (Parker, 1994).

In summary, the important points are the following: Solution of the inverse problem is nonunique because the data are finite in number and inaccurate. Selection of a single answer, from the infinitely many that fit the data, requires additional information about the model. We will attempt to include all of the information that is available. In addition, we adopt the philosophy that we want to have a model that is as simple as possible so that we don't over interpret the data. Our goals can be accomplished by designing a model objective function ϕ_m and finding

a model that minimizes this subject to adequately fitting the data (i.e. having the misfit ϕ_d be acceptably small). To compute our solution we minimize $\phi_d + \beta\phi_m$ where β is an adjustable regularization constant that determines the degree to which the data are fit.

The ideas outlined above are applicable to all inverse problems. However, the practical solution depends upon many factors, for example: whether the data are linear or nonlinear functionals of the model, the size of the problem, the computations required to do the forward modelling and so on. To crystallize the ideas, and to set the foundation for all inversions, we begin with a simple linear inverse problem as an illustration.

2.1: The linear inverse problem

Problems in geophysics can be linear or nonlinear. Linear problems are the easiest to solve. For nonlinear problems we usually linearize the equations and iterate. Either way, the solution of the linear inverse problem forms the heart of most analysis.

How do we know if a problem is linear? Let m_1 and m_2 be two models, and let c_1 and c_2 be coefficients. A problem is linear if and only if the following is true:

$$F[c_1m_1 + c_2m_2] = c_1F[m_1] + c_2F[m_2]. \quad (9)$$

Essentially linear problems are those in which superposition holds. The response from composite structures is equal to the sum of responses from the individual structures and there is no interaction between the two structures. Some geophysical problems are linear. Gravity and magnetic problems are linear; straight ray tomography is linear if the model is slowness rather than velocity, the RMS problem considered above is linear if the model stands for v^2 .

Linear problems can be written as an integral equation with the integrand being the product of the model and kernel functions that are independent of the model. For a 1D problem this takes the form

$$d_j = \int_a^b g_j(x)m(x)dx, \quad j = 1, \dots, N \quad (10)$$

where d_j is the j 'th datum, g_j is the j 'th kernel function and m is the model. The limits on the integral arise from the physical experiment (the definition of the domain of the model) and the kernel functions depend upon the physics of the problem, boundary conditions, and locations of sources and receivers. They play a fundamental role in the inverse problem. Each kernel is effectively the window through which each datum sees the model, since each datum is given by the inner product of the model with the kernel function over the model domain.

As an example, consider a 1D problem shown in Figure 6. The kernel functions are chosen to be a combination of exponentially decaying sinusoids.

$$g_j(x) = e^{jpx} \cos(2\pi jqx), \quad j = 0, \dots, 19 \quad (11)$$

These kernels typify those arising in surface geophysical surveys. For example, they emulate an electromagnetic experiment carried out at different frequencies. In such surveys, higher frequency waves are attenuated more rapidly as they penetrate the earth. In Figure 7 we show one of the kernels overlaid with the model, the product of those functions, and the final datum. This process explicitly shows that we obtain information about the model only in those regions where the kernel has significant amplitude, and that the final datum is a weighted average, or moment,




Figure 6 A synthetic 1D example illustrating the inverse methodology. The top panel shows the model consisting of a boxcar and a Gaussian bump. The lower panel displays twenty kernel functions for the data. The kernels are exponentially decaying sinusoids.

of the model. An understanding of this process is essential. The kernel functions are the crucial elements in our ability to obtain information about the model. In a good experiment we hope to adjust the survey parameters to obtain as many independent moments of the model as possible. Unfortunately, we can't make the kernels any shape we want (unless we are just working in a mathematical exercise) because the shape of the kernels depend upon the physical equations.

Demo applet: hands on exercises

The main text in this section explains the basic ideas with necessary figures. More insight can be obtained by working with the applet that is provided in the accompanying CD-ROM. The text in this and other boxes provides instructions for working with the applet.

Upon starting the applet, it will display the model shown in Figure 6 and a set of data generated by the 20 kernels. Noise contaminated data are also displayed.

1. You can keep the standard demo model or create a new model of your choice. To create a new model, place the cursor at the point you want to start changing the model and draw the desired curve while holding the left mouse button.
2. The button "kernel" provides access to a pop-up window for examining kernel functions and altering their parameters. You can keep the default choice of kernels or change them.

2.1.1: Generating a numerical solution

The models we are attempting to find are functions but we need to discretize the problem so that a numerical solution can be computed. There are two routes. We can carry through all of the mathematics using functions and

then discretize the problem at the last stage; or we can perform the discretization at the outset. We choose the latter option. One justification for this is that the forward problem of most geophysical methods can only be solved numerically and discretization is required from the beginning anyway.

There are various ways to discretize any problem (Parker, 1994; Hansen, 1998). In the work done here, we introduce a model mesh and divide our earth into “rectangular cells”, which are horizontal layers in 1D, rectangular blocks in 2D, and cuboidal cells in 3D. These are illustrated in Figure 8.

For a 1D problem, the nodal points of the mesh are designated as $x = (x_0, x_1, \dots, x_M)$. We shall also assume that the physical property is constant in each cell, that is, our model is a piecewise constant function. The model space is therefore an M -dimensional vector space $m = (m_1, m_2, \dots, m_M)$. Parameterizing in this manner is not




Figure 7 The numerical value for each datum is given by the integration of the product of the model and the kernel function associated with that datum. The top panel shows the model in Figure 6 overlaid with the fourth kernel function. The lower panel displays the product of the kernel and the model, and the corresponding datum is equal to the area under the curve (shaded region). Note that this datum has a negative value even though the model is completely positive.




Figure 8 Illustration of discretization used in 1D, 2D, and 3D geophysical inverse problems. Within each cell the physical property is assumed to be constant. This is not an undue restriction so long as the cells are sufficiently small.

an undue restriction so long as the cells are sufficiently small. Within the context of our inverse problems, the discretization should be sufficiently fine that it does not affect the outcome of the inversion. That is, the same solution for the optimization problem should arise if the cell size is decreased by a factor of two. In such cases the discrete solution we are calculating is effectively the same as the continuous solution.

With the above discretization, the forward modelling can be written as $Gm = d$ where G is an $N \times M$ matrix whose elements

$$G_{jk} = \int_{x_{k-1}}^{x_k} g_j(x) dx, \quad (12)$$

are the integrals of the j 'th kernel function over the k 'th cell. We note that the rows of G are effectively the kernel functions (imagine that we used a midpoint rule for evaluation of eq.(12); then each element of G is the value of the kernel multiplied by δx_k , the discretization interval). The j 'th datum is the inner product of the j 'th row of G with the model vector, that is, each datum is a weighted average of the model.

The system of equations $Gm = d$ is underdetermined when $N < M$. Our model space has dimension M but we have only N activated vectors. So even if the vectors are linearly independent then there are $M - N$ degrees of freedom and hence infinitely many solutions.

To generate a particular solution we introduce the misfit and model objective functions, which when evaluated with our cellularized representation of the earth, become:

$$\phi_d = \sum \left(\frac{d_j^{obs} - Gm}{\epsilon_j} \right)^2 = ||W_d(Gm - d^{obs})||^2 \quad (13)$$

and

$$\begin{aligned} \phi_m(m) &= \alpha_s \int (m - m_{ref})^2 dz + \alpha_x \int \left(\frac{d}{dx} (m - m_{ref}) \right)^2 dx \\ &= \alpha_s ||W_s(m - m_{ref})||^2 + \alpha_x ||W_x(m - m_{ref})||^2 \\ &= ||W_m(m - m_{ref})||^2 \end{aligned} \quad (14)$$

where matrices W_s and W_x are the smallest and flattest model weighting matrices, and W_m is the combined weighting matrix given by $W_m^T W_m = \alpha_s W_s^T W_s + \alpha_x W_x^T W_x$.

The final objective function to be minimized is

$$\begin{aligned} \phi(m) &= \phi_d + \beta \phi_m \\ &= ||W_d(Gm - d^{obs})||^2 + \beta ||W_m(m - m_{ref})||^2 \end{aligned} \quad (15)$$

This function depends upon M variables. Its minimum is found by differentiating with respect to each of the variables and setting the resulting equations to zero, that is, the solution model produces zero gradient of the objective function in eq.(15). Setting the gradient $\nabla \phi_m = 0$ yields

$$(G^T W_d^T W_d G + \beta W_m^T W_m) m = G^T W_d^T W_d b^{obs} + \beta W_m^T W_m m_{ref} \quad (16)$$

which is an $M \times M$ invertable system of equations provided that $\beta > 0$. The recovered solution is

$$m_{rec} = (G^T W_d^T W_d G + \beta W_m^T W_m)^{-1} (G^T W_d^T W_d b^{obs} + \beta W_m^T W_m m_{ref}) \quad (17)$$

To illustrate this process, we generate twenty data associated with the model and kernels shown in Figure 6 and subsequently contaminate them with Gaussian random noise whose standard deviation is equal to a minimum

value plus a percentage of the datum magnitude. This simulates practical observations where large-valued data are typically characterized by a relative error whereas small-valued data are characterized by an absolute error estimate. Figure 9 displays both accurate data and noise-contaminated data for our simulation. When these data are inverted using the above-discussed algorithm, we recover a model shown in Figure 10. The value of the regularization parameter β was chosen so that a target misfit of $\phi_d^* = 20$ was achieved. Comparison with the true model shows that the constructed model has recovered the major features but high-frequency variation is lost. This is expected since the data are finite in number and inaccurate and since the model norm was designed to find a smooth model. (The reader can refer to the previous boxes on the use of the inversion applet to reproduce these results and experiment with other scenarios).




Figure 9 The accurate and noise-contaminated data for the synthetic example shown in Figure 6. The plus (+) indicates the accurate data, and the circle (o) indicates the noisy data.




Figure 10 The top panel shows the recovered model, which is overlaid with the true model (dashed line) for comparison. Major features in the true model are well-represented. This model achieves the expected misfit value of 20. The predicted data are shown with the observed data in the lower panel. The two data sets match each other well, but they are not identical.

Demo applet: Example of linear inversion

Continuing with the inversion demo applet, we now generate the data, add noise to simulated observations and invert them to construct a model:

1. Once a model is drawn, the data are automatically computed and displayed below the model plot. The “noise” button allows you to add different amounts of noise to the data that are generated whenever the true model is changed. This button will bring up a dialog box where you can input the percentage and minimum threshold of error to be added. The example in the text uses 2% of each datum value and a minimum of 2% of the maximum value. You can also change the seed for the random number generator to obtain different random sequences or make sure you are working with the same realization of noise.
2. The next step is to design a model objective function. The basic form in eq.(14) is implemented. The default reference model is $m_{ref} = 0$. You can make another reference model using the right-hand button on your mouse. Set the cursor down at the left most endpoint, click the button, and draw your reference model. You can choose different coefficients α_s and α_x to look for models with different smoothness. Increasing the ratio α_x/α_s makes the model smoother. The example in text uses a value of $\alpha_x/\alpha_s = 100$.
3. You then need to decide on the target misfit value. This is done by selecting the option *chifact* and setting the value to 1 in the lower-right corner of the applet. The target misfit is given by the number of the data multiplied by the *chifact*. (The other options offer different approaches to the choice of regularization parameter, but they will be discussed later.)
4. Press the “invert” button to initiate the inversion process. An appropriate value of regularization parameter is found and the inverse solution is obtained by solving eq.(16). The recovered model will be displayed on top of the true model. The predicted data are displayed along with the noisy data. Below the data display are a number of lines summarizing information about the inversion such as the chosen regularization parameter, target and achieved data misfit values, and value of the model objective function.

2.1.2: The effects of under-fitting or over-fitting the data

In the above example, the true misfit evaluated with the particular realization of Gaussian noise, was $\phi_d^{true} = 19.5$. This is slightly different than our target value of $\phi_d^* = 20 (= N)$, but still we have fit the data to about the right global value with our choice of β . We now examine how the model is affected when our misfit is substantially above or below the true value of the noise. This is best done by using the applet but an example is given in Figure 11 where we show models for two other misfits. In Figure 11(a), $\phi_d^* = 100$, (obtained by setting *chifact* = 5 and running the applet) which is substantially larger than ϕ_d^{true} . Fitting the data too poorly (underfitting the data) has resulted in a model that has less structure. Resolution is lost and information about the model that is in the signal, is not being extracted. Alternatively, misfitting the data to values significantly smaller than ϕ_d^{true} (overfitting the data) generates excessive structure. The model in Figure 11(b) has a misfit of 8.2, which is substantially smaller than the true value. The excessive structure is the result of attempting to fit the noise that, generally, requires a great deal of complexity in the model. In almost all problems, attempting to fit data to values less than the true noise level, will require excessive amounts of structure and will result in only a modest decrease in the misfit.

This behavior is further illustrated by examining the continuous change of data misfit and model norm as functions of the regularization parameter. Figure 12(a) shows that the data misfit increases monotonically with

the regularization parameter. Figure 12(b) shows the monotonic decrease of the model norm with the regularization parameter. One notable feature of the curves is that the data misfit remains nearly constant for a wide range of small β , but the model norm changes rapidly in that range. At these regularization values, the inversion is attempting to fit the noise in the data. The result is that a great deal of structure is being put into the model but little misfit reduction is achieved. The lower panel of the figure displays the data misfit as a function of the model norm. This curve is known as the Tikhonov curve and it plays an important role in inversion theory. Each point on this curve corresponds to the results of minimizing eq.(15) for a different value of β . When β is large (the left hand portion of the curve), the solutions correspond to small model norm but large misfits. Conversely, when β is small (right hand portion of the curve) the misfits are small but the model norm is large. The transition point seems to be a compromise and it is also close to the solution based on $\phi_d^* = N$. This observation has been used in the L-curve criterion to estimate an optimal regularization, which we will discuss later.

Demo applet: Effect of different data misfit

Explore the effect of underfitting and overfitting the data. During the inversion, the applet automatically generates a complete Tikhonov curve (data misfit as a function of model norm), which is shown in the top-right portion of the applet. The point on the curve that corresponds to the current solution is marked by a brown dot.

1. To explore the effect of underfitting the data, click the left mouse button on the portion of the Tikhonov curve to the left of the current solution. As you move to the left end of the curve, each point on the curve corresponds to an inversion with increased data misfit. You will observe that the model becomes smoother, and the predicted data deviate more, from the observed data.
2. Moving towards the right end of the Tikhonov curve, the solution corresponds to smaller regularization and the data misfit decreases slowly while the model norm increases rapidly. The effect can be seen clearly in the recovered model which accumulates a large amount of structure and ultimately bears little resemblance to the true model.

2.1.3: Understanding the Tikhonov Curve

The mathematical understanding of the behavior of the Tikhonov curve is achieved through the use of singular value decomposition (SVD) (Lanczos, 1958). This is elaborated upon in many references (e.g., Parker, 1994; Hansen, 1998), but it is so fundamental to our problem that we reiterate the essential points here. An undergraduate level background in linear algebra is required for this section.

It suffices to consider the following problem. We minimize

$$\phi(m) = \|Gm - d^{obs}\|^2 + \beta \|m\|^2 \quad (18)$$

to obtain the sought solution. This system arises when the data errors are unity (that is, the data equations have been normalized by the standard deviation so that $W_d = I$), and only a smallest model component ($W_m = I$) is used in the objective function. It is a simplified form of the general objective function in eq.(14), but there is no loss of generality since that equation can be reduced to this form by a simple transformation called a reduction to standard form (e.g., Golub and Von Matt, 1997).

Fundamental understanding of the behavior of the Tikhonov curve is obtained by considering a solution of a system of equations using SVD. For the moment consider solving our data system $Gm = d$ without a regularization




Figure 11 Two inversions illustrating the effect of underfitting and overfitting the observations. The top panel shows the recovered model when the data are grossly underfit; for that model $\phi_d = 100$ while the true misfit is 19.5. Resolution of model features is lost. The lower panel shows the model obtained when data are over-fit; $\phi_d = 8.4$. Excessive structure is required to reproduce the noise in data.

term. The matrix G is $N \times M$ and therefore is not directly invertable. However, it has a singular value decomposition as a product of three matrices, $G = U\Lambda V^T$, where $U = (u_1, \dots, u_p)$ is an $N \times p$ matrix, Λ is a diagonal matrix with nonzero positive elements $\lambda_1 > \lambda_2 > \dots > \lambda_p$ and $V = (v_1, v_2, \dots, v_p)$ is an $M \times p$ matrix. p is number of nonzero singular values. Vectors in data space and model space are connected via the relationships $Gv_i = \lambda_i u_i$ and $G^T u_i = \lambda_i v_i$. The vectors u_i and v_i are orthonormal, and $U^T U = I_p$ and $V^T V = I_p$. The result of this decomposition is that the M -dimensional model space, and the N -dimensional data space, are divided into activated and unactivated parts. The p -dimensional subspaces of the activated portions are connected through the three quantities u_i, λ_i, v_i . Figure 13 illustrates the decomposition.

The activated portion of model space $\mathcal{R}^M$ is given by $\mathcal{G} = \text{asp}\{g_1, \dots, g_N\}$, that is all linear combinations of the rows of G . Assuming that the rows are independent, then the dimension of the activated space is N . The unactivated portion, or annihilator space, $\mathcal{A}$, is made up of all vectors that are orthogonal to the rows of G . Thus $\mathcal{A}$ contains all of the vectors such that $Gm = 0$. There are $M - N$ vectors required to describe this subspace.

Any vector in $\mathcal{R}^M$ can be written as $m = m^{\parallel} + m^{\perp}$ where $m^{\parallel} \in \mathcal{G}$ and $m^{\perp} \in \mathcal{A}$. Each element in $\mathcal{A}$ is an annihilator for the problem and, if our only objective is to fit the data, then we can add as much or as little of these vectors without affecting the data fit. That is, if m^* is a solution to the problem then $m^* + c_1 g_a$, where $g_a \in \mathcal{A}$, is also a solution.

The SVD solution of $Gm = d$ is

$$m_c = VV^T m = V\Lambda^{-1}U^T d. \quad (19)$$

In component form this becomes

$$m_c = \sum_{i=1}^p \frac{(u_i^T d)}{\lambda_i} v_i. \quad (20)$$

Eq.(20) shows that the solution is progressively built up from the vectors v_i . In most problems encountered in geophysics in which only one physical property is sought through inversion, these vectors have the characteristic that as their order number increases so does their frequency content. This is illustrated in Figure 15 for the example problem. The first vector v_1 has no zero crossings, the second has two, and so on. So the vectors are somewhat like sinusoids and thus continuing the summation in eq.(20) is similar to building up a solution by using sinusoids with progressively increasing frequency. The amplitude for v_i depends upon two terms. The first is the inner product between the data and the vector u_i . The second, and most important factor, is the size of the singular value λ_i . Figure 14 displays the singular values for our example problem. The observed logarithmic decay is characteristic of




Figure 12 The top and middle panels respectively show the change of data misfit and model norm as function of the regularization parameter. The lower panel is the Tikhonov curve which shows the data misfit value as a function of the model norm. The star indicates the solution achieving the expected misfit of 20.




Figure 13 Decomposition of the M-dimensional model space and the N-dimensional data space.




Figure 14 Singular values of the matrix in the example problem. The top panel shows the 20 kernel functions which form the rows of the matrix G . The bottom panel shows the singular values which decrease rapidly and span 10 orders of magnitude.

most problems in geophysics, and this reflects the fact that some of the data are almost linearly dependent. Dividing by a small singular value would not necessarily be a problem if the numerator $(u_i^T d)$ was sufficiently small. In practise however, the data $d^{obs} = d^{true} + n$. So the constructed model has the form

$$m_c = \sum \frac{(u_i^T d^{obs})}{\lambda_i} v_i = \sum \frac{(u_i^T d^{true})}{\lambda_i} v_i + \sum \frac{(u_i^T n)}{\lambda_i} v_i \quad (21)$$

Even if the error is small, the term $(u_i^T n)/\lambda_i$ can be large since λ_i can be very small. This amplifies the effect of high order singular vectors in the constructed model. In fact, it can be large enough to make the solution incomprehensible.

The above analysis illustrates that the small singular values are the progenitors of the instability of the inverse problem. Any methodology that attempts to solve the inverse problem must take this into account. Two approaches are common. The first winnows out as many higher-order singular vectors as needed to obtain a solution that is reasonably simple and still fits the data. This is called the truncated SVD (TSVD). The second approach is to dampen the effect of the higher-order singular vectors so that they do not affect the solution adversely. This is effectively the way our minimization solution in eq.(18) has stabilized the problem.

The minimization of

$$\phi(m) = \|Gm - d^{obs}\|^2 + \beta \|m\|^2 \quad (22)$$

yields

$$(G^T G + \beta I)m = G^T d^{obs} \quad (23)$$

Substituting the singular value decomposition $G = U\Lambda V^T$ into this yields a solution

$$\begin{aligned} m_c &= V^T V m = V(\Lambda^2 + \beta I)^{-1} \Lambda U^T d^{obs} \\ &= V T \Lambda^{-1} U^T d^{obs} \end{aligned} \quad (24)$$

where the matrix T is diagonal and has components

$$t_i = \frac{\lambda_i^2}{\lambda_i^2 + \beta}. \quad (25)$$

The t_i are referred to as filter coefficients in the Tikhonov regularization and these are plotted in Figure 16 for the example problem considered here. In component form the solution is

$$m_c = \sum \frac{t_i (u_i^T d)}{\lambda_i} v_i \quad (26)$$

So the regularization term acts to dampen the amplitude of the basis vectors associated with smaller singular values. If the regularization parameter $\beta \gg \lambda_i^2$ then $t_i \approx 0$ and the associated basis function will not be added. This is exactly what we want because this basis vector doesn't reduce the misfit much anyway and yet it contributes greatly to the structural complexity in the model. This is quantified by evaluating the misfit and model objective function terms

$$\phi_d = \sum_{i=1}^p (1 - t_i)^2 (u_i^T d)^2 \quad (27)$$

$$\phi_m = \sum_{i=1}^p \left(\frac{t_i}{\lambda_i} \right)^2 (u_i^T d)^2 \quad (28)$$

To summarize, the first few singular vectors have many desirable properties. They add a large smooth component to the structure, they greatly reduce the misfit and they don't contribute substantially to the model norm. For these vectors we want $t_i \sim 1$ and so β should be chosen to allow this, i.e., $\beta \ll \lambda_i^2$. Conversely, components associated with small singular values add a high structural component, don't significantly reduce the misfit and greatly increase the model norm. These contributions are undesirable and we want $\beta \gg \lambda_i^2$ so that $t_i \approx 0$.

All practical algorithms make use of the above fundamental ideas although implementation varies. In small-scale problems, where the SVD of a matrix can be computed, then using the filter approach discussed above, or imposing a truncation (TSVD), is common. In large-scale problems it is not possible to carry out the decomposition and the matrix systems are solved using iterative techniques. Irrespective of how the final system is solved, the roll of β and its effect on overfitting or underfitting the data remains the same.




Figure 15 Singular vectors for the example problem. The first singular vector, which is associated with the largest singular value, is the simplest and has no zero crossing. Subsequent singular vectors become more complicated and exhibit an increased number of zero crossings.




Figure 16 The filter coefficients (+) in the Tikhonov regularization are compared with the singular values (*). Note that the filter coefficients decrease more rapidly than the singular values. As a result, the singular vectors of smaller singular values are heavily damped in the constructed model.

Demo applet: Further exploration

The main text and the preceding exercises with the applet illustrate the basics of linear inversion. However, the outcome of an inversion depends on a number of factors such as the data available, the noise in the data, and the type of model objective function used. The reader is encouraged to explore these aspects and gain further understanding by using the applet.

1. Examine the effect of a reduced data set by using two different data types generated from the default model.
 - a. Test-I: select only the first six kernels with default parameters. This is analogous to acquiring fewer data. Generate the data and invert them to construct a model using the default model objective function. Note the loss of resolution in the constructed model. This is because the ten data associated with high-frequency kernels are omitted.
 - b. Test-II: select the first 20 kernels as given but increase the decay of the kernels. For example, you can change the decay parameter to $p = -0.75$. This is analogous to working with a survey system that has a shallower penetration depth. Inverting this set of data will produce a model that defines the boxcar anomaly well but loses the definition of the Gaussian bump at depth.
2. Examine the effect of data error by starting with the default model and kernels. You can increase the amount of noise added to the data by clicking the noise button and increase both the percentage error and the minimum noise level (for example: 20% and a minimum of 5. Inverting the noisier data using the option of $chifact=1$ will produce a model that only shows two broad anomalies with little resolution.
3. Examine the model objective function by varying the parameters α_s and α_x . These two parameters can be changed by clicking the alpha button. In order to see the effect clearly in the applet, you need to increase the noise level in the data.
 - a. Test-I: Choose $\alpha_s = 1$ and $\alpha_x = 0$ to produce a smallest model.
 - b. Test-II: Choose $\alpha_s = 0$ and $\alpha_x = 1$ to produce a flattest model.

Note that the smallest model has more structure in it while the flattest model has a much smoother variation.

2.2: Examples of Linear Inversion

In reality there are few linear inverse problems in geophysics, and even those that are linear, are often supplied with additional constraints such as positivity, or more general bound constraints on the physical properties. These extended problems require that solution strategies for nonlinear problems, which we will discuss later, be employed. Nevertheless gravity, magnetics, induced polarization (IP) and cross-well tomography that assumes straight rays, can be solved with linear analysis. Here we consider three examples: (1) a cross-well tomography problem, (2) IP inversion in 2 dimensions without positivity, and (3) a 3D gravity inversion without bound constraints. We have selected these examples to show how the inversion methodology, outlined thus far, works. The examples have been chosen for a number of reasons. Firstly, each is commonly used in applied geophysics and so is of practical relevance. Secondly, the associated kernel functions for each method are substantially different and it is of interest to see the effects of this. Thirdly, the problems are solved in two and three dimensions. Fourthly, inverting these data sets brings to light other issues that become important in practical inversions of these and other types of geophysical data. At this stage it is sufficient to work with synthetic models.

2.2.1: Cross-Well Tomography

In a typical cross-well seismic tomography survey, first arrival travel times from sources in one borehole are recorded by receivers in the other borehole. The travel time depends upon the velocity structure between the two holes and the goal of inversion is to recover the velocity. The travel time for the j 'th datum is given by

$$t_j = \int_{\Gamma_j} \frac{1}{v(l)} dl \quad (29)$$

where Γ_j is the ray path, and l is the distance along the path. In general the ray path depends upon the velocity structure, but if the velocity variation is small, then the rays travel in approximately straight lines from the source to receiver. By adopting this straight-ray assumption, and by working with slowness $1/v$, the problem is linear.

As an example we consider the velocity model in Figure 17. Twenty receivers in the borehole at $x = 100$ meters measure arrival times from ten sources in the borehole at $x = 0$ meters. The data, contaminated with Gaussian noise of 0.02 seconds, are shown in Figure 17(b).




Figure 17 (a) A high velocity (small slowness) block is embedded in a uniform medium. (b) Travel-time data contaminated with noise. The source location is plotted on the vertical axis, the receiver location is on the horizontal axis. (c) The recovered model. (d) Predicted travel times from the recovered model.

To proceed with the inversion we first define a 2D model objective function that is the discretized version of

$$\phi_m(m) = \alpha_s \int_A (m - m_{ref})^2 dx dz + \alpha_x \int_A \left(\frac{dm}{dx} \right)^2 dx dz + \alpha_z \int_A \left(\frac{dm}{dz} \right)^2 dx dz \quad (30)$$

where $m = 1/v$ is the slowness. To discretize we introduce a 2D rectangular mesh with M elements so that the model objective function can be written as eq (14). With the same mesh, the data can be written as

$$d_j = \sum_{k=1}^M G_{jk} m_k \quad j = 1, N \quad (31)$$

where G_{jk} is the length of the j 'th raypath in the k 'th cell, and m_k denotes the slowness in the cell, and d_j is the travel time for the j 'th source-receiver pair. The matrix G is very sparse. Each row has non-zero elements only for those cells through which the raypath crosses. That is, each datum is providing information about those cells but not about cells away from the raypath.

The data equations in eq.(31) are in the form $d = Gm$ and the misfit function ϕ_d is defined by eq.(13), where, in this synthetic example, we have the luxury of knowing the exact values of the standard deviations. We specify the reference model to be a constant equal to that of the background slowness. Eq.(16) is solved and a value of β is found such that $\phi_d^* = 200$, which is the number of data.

The recovered model is shown in Figure 17(c) along with the predicted data. The inversion has been moderately successful. A zone of high velocity, coinciding with the approximate location of the initial block, is reasonably well defined. Because of the smoothing components in the model norm we see boundaries of the block are smeared, especially in the X-shaped region where the raypaths are concentrated. The cells in the two V-shaped regions on the top and bottom have no rays going through them and the final solution is entirely controlled by the model norm. The model therefore coincides with the reference model.

2.2.2: Inversion of Induced Polarization Data

IP data are routinely acquired in mineral exploration, geotechnical and environmental problems. For background on the physics of the IP problem, data collection and case histories, the reader is referred to Sumner (1976), Fink *et al.* (1990), Wait (1959), or other papers in this volume. Although the complete IP problem is nonlinear, for many cases of practical interest the chargeability η is small and a linearized relationship exists between the observed data and chargeability. For our analysis, we assume that the earth model has been cellularized in the manner outlined earlier, and that the chargeability for the k 'th cell is η_k . The data, referred to as apparent chargeability, are expressed as

$$\eta_{aj} = \sum_{k=1}^M J_{jk} \eta_k, \quad j = 1, \dots, N \quad (32)$$

The apparent chargeabilities can be dimensionless, as in Seigel's (1959) definition, or have units of msec, mrad, or pfe (percent frequency effect) depending upon the system with which they were acquired (Oldenburg and Li, 1994). The matrix elements J_{ij} are the sensitivities for IP data and they are dependent upon the electrode geometry used for acquisition, and upon the electrical conductivity. The conductivity is usually estimated by inverting DC resistivity data acquired at the same time as the IP data. It is important to note that, unlike the cross-well tomography example above, the matrix J is full, and thus each datum is affected by the presence of chargeability at every point in the

subsurface. Consequently it is generally not possible to make geologic inferences directly from a pseudo-section plot of the data.

For a numerical example, we consider a 2D chargeability model with a uniform background conductivity. We first investigate the kernel functions associated with these data. A plot of the IP sensitivity for two n-spacings of a dipole-dipole electrode geometry is shown in Figure 18. The sensitivity has large amplitude near the surface, peaks near the electrodes, and decays with distance from the electrode. The corresponding IP datum is obtained by integrating the product of the chargeability with the sensitivity shown in the figure. As a result, the near surface chargeability has the strongest effect on the measured apparent chargeability and hence near-surface inhomogeneities often mask the signal we are interested in. We also note that the sensitivity changes sign with spatial location, and it becomes negative in the region between the electrodes near the surface. Therefore, some subsurface regions contribute positively to the measured data while others contribute negatively. As a result, although the intrinsic chargeability is positive, the apparent chargeability can be negative or positive depending where the chargeable region is in the subsurface. Another characteristic of the sensitivity is its more rapid change near the surface and broad diffused appearance at depth. This reflects the fact that the resolving power of surface arrays is better near the surface and worsens with depth.

To illustrate IP inversion, we carry out a dipole-dipole survey over the structure in Figure 19(a). The simulated data, with added Gaussian noise, are displayed in Figure 19(b). The apparent chargeability pseudo-section is dominated by responses from the near-surface chargeability variations, and the deeper targets are masked. Inversion is required. To proceed, we first define a 2D model objective function that is the discretized version of eq.(30). The




Figure 18 The sensitivities of IP data from dipole-dipole arrays with a dipole spacing of 10 m. The top panel shows the sensitivity for n-spacing of 1, and the lower panel shows the sensitivity for n-spacing of 5. Both sections are plotted using four gray levels to show the polarity of the sensitivities. They are characterized by a central region of negative values surrounded by a semi-annular positive region. Farther out, they become negative again. Although the intrinsic chargeability is always positive, any chargeable body located within the negative zones of sensitivity will produce negative apparent chargeability.

data equations in eq.(32) are already in the form $d = Gm$ and the misfit function ϕ_d is defined by eq. (13), where, again in this synthetic example, we have the luxury of knowing the exact values of the standard deviations. Eq.(16) is solved and a value of β is found such that $\phi_d^* = 93$, which is the number of data. The 2D model was described with 2500 cells. The recovered model, shown in Figure (19c), images the major anomalous chargeability zone reasonably well. Panel-d displays the predicted apparent chargeability pseudo-section. Note that it reproduces the observed data well, but without the high-frequency variations due to the noise. To further examine the difference, panel-e shows the data misfit map, which is the difference between the observed and predicted data normalized by the error standard deviation. The optimal data fit should result in a sufficiently random-looking misfit map without any long-wavelength features in this misfit map. This example is a good illustration.

Overall, the inversion has done a moderately good job in delineating the surface and subsurface chargeability zones. However, there is one major aspect of the image that is not satisfactory. The recovered model has regions of negative chargeability. This is not physically plausible since chargeability is, by definition, a positive quantity. Positivity is important additional information and we will see that incorporating it not only makes the image physically viable, but it also sharpens up the signature of the buried object. Incorporating positivity requires that the inversion be recast as a nonlinear problem and so the details of how to accomplish that are postponed until next section. However, to motivate the need for carrying this out, we present the model obtained by imposing positivity, in Figure 19(f). It is a superior solution.

2.2.3: Gravity Inversion

In our third example we invert surface gravity data to recover a 3D anomalous (either excess or deficient) density structure. The vertical component of the gravity field is given by

$$g_z(\vec{r}_0) = \gamma \int_V \rho(\vec{r}) \frac{z - z_0}{|\vec{r} - \vec{r}_0|^3} dv \quad (33)$$

where ρ is the density contrast, $\vec{r}_0$ is the observation location and $\vec{r}$ is the source location. V represents the volume of the anomalous mass, and γ is the gravitational constant. Here we have adopted a Cartesian coordinate system having its origin on the earth's surface and the z-axis pointing vertically downward.

Our density model is chosen to simulate two dissolution tunnels that have different strike lengths and dip. Figure 20(a) shows a volume rendered image of the density contrast from a viewpoint looking northwest. The tunnel at the left has a depth of burial varying from 50 m at the south end to 350 m at the north end, while the tunnel on the right has a constant depth to the top of 200 m. Two negative gravity anomalies corresponding to the tunnels are shown as contour map Figure 20(b). The data set consists of 254 points gathered at a station interval of 25 m along 7 parallel lines spaced 150 meter apart and Gaussian noise has been added to simulate observational errors. The data locations are marked by the symbol +.

To perform gravity inversion, we again use cell-based discretization. We divide the region of interest into a set of 3D prismatic cells ($M=32000$) and assume a constant density contrast within each cell. The gravity field at the j 'th location is,

$$d_j = \sum_{k=1}^M G_{jk} \rho_k \quad j = 1 \cdots N \quad (34)$$

where the sensitivity G_{jk} quantifies the contribution of the k 'th cell to the j 'th datum. It is given by

$$G_{jk} = \gamma \int_{\Delta V_k} \frac{z - z_{0j}}{|\vec{r} - \vec{r}_{0j}|^3} dv, \quad (35)$$




Figure 19 Inversion of 2D IP data. Panels: (a) the true chargeability model, (b) simulated data with added uncorrelated Gaussian noise, (c) chargeability model recovered from a 2D IP inversion without positivity, (d) predicted responses from the above recovered model, (e) the difference between the observed and predicted data normalized by the error standard deviation, and (f) chargeability model recovered with a positivity constraint.

where ΔV_k is the volume of the k 'th cell. Eq.(34) is expressed in matrix form $Gm = d$ as in eq.(13).

To apply the inversion algorithm we define a model norm identical to that in the 2D IP inversion except that now we also include the derivative term in the third dimension. The misfit is defined by eq.(13), and again, we have used the known standard deviations of the errors. Eq.(16) is solved, along with a value of β , such that $\phi_d^* = N$. Figure 20(c) shows recovered density contrast as a volume-rendered image. The predicted data from this inversion are shown in Figure 20(d), which is a good match to the observations. The difference map between the observations and predicted data is shown in panel-f. It is sufficiently random as desired. The model in Figure 20(c) identifies two regions of negative density and locates their horizontal extents. However, the major density contrast is concentrated at the surface and there is no indication that the tunnel on the left is dipping. We conclude that, although the inversion has been a technical success, we have failed to obtain geologic information regarding the depth of the tunnels.

This result highlights difficulties associated with the inversion of potential field data: we face not only the non-uniqueness caused by the availability of finite number of inaccurate data, but also the fundamental ambiguity of potential fields that obey Laplace's equation. According to Gauss' theorem, there will be infinitely many different subsurface density distributions that will reproduce a set of data on the surface even if the data are continuous and infinitely accurate. In terms of dimensionality, we have sought to recover the density distribution as a function of spatial position in 3D while we have at most a single 2D map of independent data available to us. Therefore, there is not enough information in the data and theoretically, we do not expect to recover reliable depth information unless we further restrict the admissible models.

In response to this inherent non-uniqueness, there are two routes that can be chosen. The first is to give up on 3D inversion and restrict the gravity interpretation to plan maps or equivalent layer densities. The other alternative is to modify the inversion algorithm so that it produces the type of model that is desired. In this case we want the density distributed with depth and not concentrated at the surface. This can be accomplished by using a depth weighting function that will be discussed in a future section. Also, for this particular gravity problem, there is the possibility of incorporating negativity, which is also discussed later. As motivation for those sections, and also to show that the basic inversion methodology outlined is still useful and only needs some alterations, we present our final inversion of these gravity data. Figure 20(e) shows the inversion result obtained after the algorithm has been modified so that a depth weighting, and bound constraints restricting density to the range $[-1.5, 0]$, have been applied. The tunnels are well located with clear indication of a dipping structure in one of them.

2.2.4 Summary Comments for the Linear Inverse Examples

The three examples presented here show that the basic inversion methodology can work well in some cases but there are other cases where more information, such as positivity, or altering the inverse problem to include a depth weighting, are needed. In the following sections we address these issues. As a starting point, however we look at strategies needed to solve the general nonlinear inverse problem.




Figure 20 Example of gravity inversion. Panel-a shows a volume rendered image of the density contrast corresponding to the tunnels. Panel-b shows the simulated gravity anomaly, which has been contaminated by Gaussian noise. Panel-c is the result of inverting the observed data using a generic model objective function, and panel-d displays the predicted data. The difference between the observed and predicted data are shown in panel-f. Panel-e shows the inversion result obtained when depth weighting and negativity are incorporated into the inversion.

Section 3.0: NON-LINEAR INVERSION

For most geophysical problems the forward modelling equation $F[m] = d$ is nonlinear, that is, the relationship eq.(9) does not hold. As an example consider the dc resistivity equation

$$\nabla \cdot (\sigma \nabla \phi) = -I\delta(r - r_s) \quad (36)$$

where the forward modelling operator involves the product of model, σ (electrical conductivity), with the gradient of the response ϕ (electrical potential). Physically, the nonlinearity arises because electrical charges, accumulated in regions where the conductivity is changing, have fields that affect the magnitude of charge buildup in other areas. This interaction means that the principle of superposition with respect to the conductivities is not valid. Another example is in refraction tomography where the travel time is

$$t_j = \int_{\Gamma_j} c(l) dl \quad (37)$$

where c denotes the slowness, Γ_j is the ray path, and l is the distance along the path. When a straight-ray assumption is used, as we did earlier, then the travel times depend linearly upon the slowness. However, that is an approximation, because the true ray path depends upon the velocity structure. This dependence makes the refraction problem non-linear. Without major simplifying assumptions, data from all electromagnetic and seismic problems are intrinsically nonlinear functionals of the physical properties. Even the IP problem that we treated earlier, is really nonlinear; the IP equation $J\eta = d$ is an approximation that is valid only when the chargeability η is small. In summary, nonlinear problems in geophysics are the rule, rather than the exception, and we now turn our attention to solving them.

The issues for solving the nonlinear inverse problem are fundamentally the same as those in the linear problem. We still desire to find a function from a limited number of data and hence fundamental problems regarding nonuniqueness and data misfit must be addressed at the outset. As a result, the same optimization methodology for solving the linear inverse problem can be used to solve nonlinear problem. We specify a misfit function ϕ_d and a model norm ϕ_m and minimize $\phi(m) = \phi_d + \beta\phi_m$ where β is the regularization parameter. For the present we will continue with our methodology of choosing β such that the final misfit is equal to ϕ_d^* .

We haven't stressed this point previously, but we are able to minimize $\phi(m) = \phi_d + \beta\phi_m$ in one step only if $\phi(m)$ is a quadratic function. Only under that condition does taking a gradient yield a linear system of equations to be solved. The optimization problem becomes nonlinear when either ϕ_d or ϕ_m is non-quadratic. That can occur because the forward mapping for the data is nonlinear, because the model objective function is non-quadratic, or because additional constraints, such as positivity or bound constraints, are imposed upon the solution. In fact, the optimization problem of minimizing eq.(15) and finding a β such that $\phi_d = \phi_d^*$ is a nonlinear problem. We have made it into a "linear" problem by fixing β so that objective function is quadratic. This leads to a linear system of equations that is repeatedly solved with different β 's, and ultimately an acceptable β is selected.

The literature regarding solving constrained and unconstrained nonlinear optimization problems is vast. Excellent books are Nocedal and Wright (1999), Kelley (1999), Gill et al. (1981), Dennis and Schnabel (1996). It is not our intention to attempt to provide an in-depth summary of that material. Rather we shall outline some general principles behind the Newton or Gauss-Newton procedures that are typically used so that essential elements are discussed.

Let

$$\phi(m) = \phi_d + \beta\phi_m(m) = ||W_d(d^{obs} - F[m])||^2 + \beta||W_m(m - m_{ref})||^2 \quad (38)$$

and assume that β is fixed. Our goal is to find m that minimizes this functional. An iterative procedure is required and so we let $m^{(n)}$ be the current model and δm be a perturbation. Expanding eq.(38) in a Taylor series yields

$$\phi(m^{(n)} + \delta m) = \phi(m^{(n)}) + g^T \delta m + \frac{1}{2} \delta m^T H \delta m + \dots \quad (39)$$

where $g = \nabla \phi(m)$ is the gradient and $H = \nabla \nabla \phi(m)$ with components

$$g_i = \frac{\partial \phi(m)}{\partial m_i} \quad H_{ij} = \frac{\partial^2 \phi(m)}{\partial m_i \partial m_j} \quad (40)$$

These are evaluated at the current model $m^{(n)}$. We want to find a perturbation such that eq.(39) is minimized. Neglecting the higher order terms and taking the derivative of eq.(39) with respect to δm , and setting the resultant equal to zero, yields Newton's equation for the perturbation

$$H \delta m = -g \quad (41)$$

The solution is updated by setting $m^{(n+1)} = m^{(n)} + \delta m$ and the process is continued until convergence, namely that the gradient is acceptably close to zero.

To evaluate eq.(41) we need to compute the gradient and Hessian. The gradient of eq.(38) is

$$g(m) = J^T W_d^T W_d (F[m] - d^{obs}) + \beta W_m^T W_m (m - m_{ref}) = 0 \quad (42)$$

where the matrix $J(m) = \frac{\partial F}{\partial m}$ is the sensitivity matrix. Its elements $J_{jk} = \frac{\partial F_j}{\partial m_k}$ quantify the change in the j 'th datum that arises from a unit change in the k 'th model parameter. The Hessian has the form

$$H = (\nabla J^T) W_d^T W_d (F[m] - d^{obs}) + J^T W_d^T W_d J + \beta W_m^T W_m \quad (43)$$

The first terms accounts for the change in the sensitivity when the model is perturbed. It is computationally burdensome to compute and also its importance decreases as the optimization proceeds because the difference between the predicted and observed data become smaller. For these reasons this term is generally neglected and we use an approximate Hessian $\tilde{H}$ in the computations:

$$\tilde{H} = J^T W_d^T W_d J + \beta W_m^T W_m \quad (44)$$

The matrix $\tilde{H}$ to be inverted is an $M \times M$ positive definite symmetric matrix so its inverse exists. The resultant equation to be solved is

$$(J^T W_d^T W_d J + \beta W_m^T W_m) \delta m = J^T W_d^T W_d (d^{obs} - F[m^{(n)}]) - \beta W_m^T W_m (m^{(n)} - m_{ref}) \quad (45)$$

This is often referred to as the Gauss-Newton equation.

The local quadratic representation of the true multi-dimensional quadratic surface described by eq.(39) is given by

$$f(\delta m) = \phi(m) + g^T \delta m + \frac{1}{2} \delta m^T \tilde{H} \delta m \quad (46)$$

where $\tilde{H}$ is an approximate Hessian matrix. Minimizing this with respect to δm yields the $\tilde{H} \delta m = -g$. If $f(\delta m)$ is an adequate approximation to $\phi(m^{(n)} + \delta m)$ then the perturbation will be good and the updated model can be $m^{(n+1)} = m^{(n)} + \delta m$. However, if $f(\delta m)$ is a poor approximation, then the recovered δm may have the wrong direction and/or be the wrong size. There are two general strategies for dealing with such situations.

In trust region approaches, the approximate Hessian is modified and the size of a potential step is restricted so that f is a good approximation. Then δm has (ideally) both the correct direction and step length and it can be added to the existing model to generate the updated solution. The other strategy is to accept that δm has the right direction but its magnitude is incorrect. Invariably the magnitude is too large so the updated model is given by $m^{(n+1)} = m^{(n)} + \gamma \delta m$ where $\gamma \in [0, 1]$ is a damping factor that reduces the step-length. γ is chosen so that $\phi^{(n+1)} < \phi^{(n)}$, that is, some progress has been made to reduce the value of the objective function. The combination of using the Gauss-Newton eq.(45) and step-length control produces the damped Gauss-Newton methodology.

Any of the above strategies for computing a step can be continued until convergence has been achieved, that is, until the gradient is sufficiently close to zero and the objective function no longer decreases. This statement trivializes some difficult issues connected with implementing the procedure numerically for highly nonlinear problems. Nevertheless, in principle, for a good implementation, the behavior of ϕ_m and ϕ_d as a function of iteration might look like that represented in Figure 21(a). At convergence, the solution corresponds to a point on the Tikhonov curve. Where it ends up on that curve depends upon the value of β that was selected.

It is unlikely that a first guess for β yields a solution where the misfit is equal to a desired value. So, just as in the linear inverse problem, a strategy is required to find the a value for the regularization parameter.

There are numerous strategies for estimating an acceptable value of β . Here we outline two of them. The first makes use of a cooling schedule for β and some variant of a damped Gauss-Newton algorithm. Cooling strategies are part of the more general framework of continuation methods (Nocedal and Wright, 1999). The inversion begins by initializing the regularization parameter to β_1 which is large enough so that the model term ($W_m^T W_m$) in the matrix in eq.(45) dominates the data sensitivity term ($J^T W_d^T W_d J$). The minimization is performed, with fixed β_1 to arrive at a point $m^{(1)}$ on the Tikhonov curve. Convergence of this inner iteration is governed by requiring the gradient to be sufficiently small. The value of β is then reduced and a new minimization is carried out beginning with the previously found model. This generates $m^{(2)}$, $m^{(3)}$, and so on. The procedure is continued until a solution with an acceptably small misfit has been found. In practise, there are questions about how quickly β can be cooled and there is a need to guard against poor steps, but the above procedure, as sketched in Figure 21(b) illustrates the basic convergence path.

An alternative way to reach the desired target misfit involves dynamically updating β at each iteration. The method, outlined in Parker (1994), has two stages. The first concentrates on finding a model that fits to the desired amount, and the second phase continues minimization of the objective function while keeping the misfit fixed. At each iteration the Gauss-Newton eq.(45) is solved to generate a perturbation $\delta m(\beta)$ for a range of β values. The nonlinear misfit $\phi_d(m^{(n)} + \delta m(\beta))$ is evaluated and, for the first phase, the value of β that produces the greatest decrease in this objective function is chosen. In the second stage, when the misfit can be achieved, there are two values of β which generate perturbations such that $\phi_d(m^{(n)} + \delta m) = \phi_d^*$. In this case, the largest value of β is the desired choice since the model associated with that will have the smaller model norm. Typical curves corresponding to this approach are shown in Figure 21(c). In some of the following examples in this chapter we adopt this strategy but impose an additional restriction that limits the amount that the misfit is reduced at each iteration. That is, the target misfit for iteration $(n + 1)$ is $\phi_d^{*(n+1)} = \max[\theta \phi_d^{*(n)}, \phi_d^*]$ where ϕ_d^* is the final target misfit and θ is typically in the range $(0.1, 0.5)$. This is a cautionary step and is introduced to help prevent excess structure being added to the model at any iteration.

3.1: Example for Nonlinear Inversion

As an example of the above optimization strategy we invert DC resistivity data taken over a 2D earth structure




Figure 21 (a) Convergence plot for minimization with a fixed β . The solution terminates at a point on the Tikhonov curve. (b) Convergence plot for a “cooling” methodology to find β that produces an acceptably small misfit. (c) Data misfits vs β for selected iterations in Parker’s two-stage algorithm.

shown in Figure 22. On the left, one resistive and two conductive prisms are buried beneath highly conductive surface blocks in a homogenous earth. One prism extends to infinite depth. A vertical contact separates media of highly different conductivities, and on the right, a circular shaped Gaussian conductor is buried beneath a layer of random conductivity. A pole-dipole survey with $n = 1, 8$ and $a = 10$ m was simulated, and the data, after being contaminated with 5% Gaussian noise, are shown in Figure 22(c). Simulation of the data was performed by solving the differential eq.(36) with appropriate boundary conditions. We used a nodal-based finite volume technique based upon the work of Dey and Morrison (1979) and McGillivray (1992).

For the inversion, the model variable was chosen to be $m = \log \sigma$. This ensures that conductivity will be positive and it also handles the situation where conductivity varies by many orders of magnitude. In these situations it is the relative change in conductivity that is of interest and logarithmic variables are appropriate. The objective function ϕ_m is the same as that of the IP inversion eq.(30) except there is a reference model incorporated. The reference model was specified to be a halfspace of $400\Omega m$ and $(\alpha_s, \alpha_x, \alpha_z) = (.001, 1, 1)$. Minimizing eq.(15) subject to $\phi_d^* = N$ produces the model in Figure 22(b). The predicted data, shown in Figure 22(d), are a good match to the observations. The recovered model bears considerable likeness to the true model; the surface inhomogeneities are well delineated and the underlying conductive prisms are clearly visible. The vertical contact is sharply imaged




Figure 22 Example of 2D DC resistivity inversion. Panel (a) is the resistivity model. (b) the recovered model, (c) synthetic dipole-dipole pseudo-section showing the data contaminated with noise, and (d) predicted data from the model in (b).




Figure 23 Convergence curves: (a) misfit ϕ_d , (b) model norm ϕ_m and (c) regularization parameter β as a function of iteration

at the surface but it extends only to limited depth. The variable surface layer is well defined and the conductor rises distinctly above the resistive background. Differences between the images in (a) and (c) are that amplitudes of anomalies are reduced, the boundaries of the recovered prisms are smooth and the constructed model merges at depth with the reference model. The characteristic smoothness of features in the reconstructed model arises from using the l_2 -norm of the horizontal and vertical derivatives in the model objective function. The fact that the recovered model becomes equal to the reference model at depth is a manifestation of the limited depth of investigation of the surface DC resistivity data.

The convergence curves for the above minimization is shown in Figure 23. We have used Parker's (1994) strategy to estimate β , but at each iteration we have restricted the decrease in misfit to be a factor of two. This can be seen in the misfit plot. As the misfit decreases, the model norm, and hence structure on the model, increases. In the last few iterations the misfit holds its target value and the model norm is somewhat reduced. The algorithm converges when the desired misfit is achieved and there is no longer significant change in the model norm.

Field inversion of DC resistivity data, and IP data are routinely carried out to recover 2D conductivity structures. Examples for mineral exploration can be found in Oldenburg *et al.* (1997), Mutton (1997), Kowalczyk (2000), and

Phillips *et al.* (2002) and there is no need to reproduce them here. Also, the same methodology can be applied for 3D problems (e.g., Ellis and Oldenburg, 1994; Park and Van, 1991; Sasaki, 1994).

Section 4.0: INCORPORATING PRIOR INFORMATION

We have previously discussed that the data provide only limited constraints upon the model and that there are infinitely many models that will fit the data to the same degree. In order to find a particular solution that contains meaningful information about the earth, we need to include more information about the likely earth model. In applied geophysics problems this prior information takes numerous forms. It can be geological, geophysical, and or philosophical. Whatever the case, the goal of a good algorithm is to incorporate this information into the inversion.

Geologic or structural information is often available and it can take many different forms. In mineral exploration problems for example, there is usually a generic geologic model that outlines likely rock units and their geometry, scale lengths of features, and mineralization. Attaching physical property values to this image provides a geophysicist with a representation of the likely earth. Combining geologic descriptions with physical property estimates, to create a potential geophysical model, is an essential part of the analysis. From this, we can then attempt to extract information about the character of the expected physical property distribution. Particular questions might be: Are the boundaries smooth or sharp? Is there preferential smoothness in a particular directions? Are some units expected to be dipping in a particular direction? The task is then to design a model norm such that the norm-minimizing solution has these properties. In new exploration areas, the details about the expected models can be scant. However, in more mature investigations, there may be considerable information that comes from boreholes or from other geophysical surveys and subsequent interpretation. Importantly, there is usually some information regarding the variation of the physical property value for each rock unit that can be obtained.

By compositing the above information one can generate a reference model that typifies the area of investigation. Sometimes the detail on the reference model is minimal (uniform halfspaces are a common default option), but in other cases, where considerable information is available, the reference model can be quite complicated.

The last form of information falls into the category of being philosophical and we have discussed this earlier. Because solutions can be made arbitrarily complex, and thus completely uninterpretable, we choose to find that solution that minimizes the model norm. If we find a solution that has the right character, is as “simple” as possible, and satisfies all available physical constraints, then we hope that features observed on that solution emulate the large scale features of the earth. This is a conservative strategy. By adopting it we may miss some structures but hopefully we don’t introduce false structures.

In this section we illustrate how the above types of information can be incorporated into our inversion. The section has three parts. In the first we introduce a flexible generic model objective and show how the various components are related to information the user might have available. The second part concerns additional “*tuning*” or revision of the algorithm that is sometimes required to eliminate particular artifacts. The depth weighting needed for inverting potential field data is an example of this tuning. The third part discusses the incorporation of bounds on the model parameters.

4.1: Prior information and model objective function

The prior information available might be: (a) knowledge of a background or reference model and (b) a general sense that the structures should be smooth or that they have sharp edges. There are many possible ways to design

objective functions that can incorporate this information. We offer one that is generic and has been useful to us. It has terms like

$$\phi_m(m) = \alpha_s \int_V w_s |m - m_{ref}|^p dv + \alpha_x \int_V w_x \left| \frac{\partial(m - m_{ref})}{\partial x} \right|^p dv + \dots \quad (47)$$

where p is a real number, usually greater than unity. The differences between this objective function and that considered earlier, are the generalization to represent the exponent by p and the inclusion of additional weighting functions w_s, w_x .

An important quantity in eq.(47) is the reference model m_{ref} . How much is known about the reference model, and its degree of complexity varies greatly between problems. In reconnaissance surveys where little is known, the reference model might be a uniform halfspace, and for some problems just the zero model. In other surveys, the knowledge of the physical property distribution, built up through previous analyses or direct measurement, might be quite detailed. In generating such reference models, we are often more certain about its value in some locations than in others. This is especially true when a reference model has been derived from borehole measurements. Our confidence in the reference model is likely high around the boreholes but decreases away from there. To incorporate this variation in certainty, the model norm includes a positive weighting w_s that has relatively high values in regions where we want the constructed model to be closer to the reference. These weightings, with or without a reference model, are also useful in the derivative terms. For instance, a low value for a weighting w_x in

$$\phi_x = \alpha_x \int_{vol} w_x(x, y, z) \left(\frac{dm}{dx} \right)^2 dv \quad (48)$$

allows the inversion algorithm to produce large gradients in the x -direction in a particular region of the model.

Reference models can play an important role in the inversion of gravity and magnetic data where the data have no inherent depth resolution, and they also can be very useful in cooperative inversions where it is hypothesized that different physical properties are correlated. For instance susceptibility and chargeability are sometimes correlated and if a good model exists for one of these, then it can be used to design a weighting function for the inversion of the other quantity. We'll illustrate later the use of the reference model and specific weighting functions in model construction.

The α coefficients in eq.(47) determine the relative weighting of the terms in the objective function and hence they can be used to adjust the character of the final model. Their effect can be dramatic. To illustrate this we invert the DC resistivity data in the preceding section after adjusting the coefficients. The models in Figure 24 are designed respectively to be close to a reference halfspace, have minimal variation in the x -direction, and minimal variation in the z -direction. Judged from the criterion of fitting the data, these models are all equally acceptable. The predicted data from each model is shown in the figure and the final misfits are identical. It is highly unlikely however that models (b) and (c) are good candidates for the local geology. So these models are mathematical solutions but they are not reasonable candidates for earth structure. We comment here that if this conclusion has been drawn, then some prior information has been used, although it was applied after the inversion was completed. The decision about whether the recovered model has an acceptable structural character is a "soft" or subjective constraint that is difficult to quantify. Nevertheless, interpreters who have worked extensively in particular geologic areas often have strong opinions regarding the possible acceptable character of the model. This can play an important role in many practical inversions.




Figure 24 Three resistivity models which reproduce the data in Figure 22. The models are generated by altering the coefficients α in the model objective function. (a) the model which is close to a background resistivity of 400 $\Omega\text{-m}$, (b) the model which has minimum structure in the horizontal direction, (c) the model which has minimum structure in the vertical direction. The predicted data for each inverted model is in the right hand column. All models fit the data to the same degree. The models and predicted data can be compared with the generic inversion in Figure 22

If models (b) and (c) in Figure 24 are regarded as being too extreme then the α 's can be readjusted and the inversion rerun. Setting them equal to $(\alpha_s, \alpha_x, \alpha_z) = (.001, 1, 1)$ yields the model in Figure 22(b). Earth models that are intermediate between that, and models (b)-(c) in Figure 24, can easily be generated by adjusting the coefficients. The final values of the coefficients needed to generate an optimum image are generally obtained through a trial process. However, length scales, obtained by evaluating ratios of the α coefficients have proved useful. Since $\phi(m)$ is a scalar that is generated by summing up a number of terms, the terms in the objective must have the same units. Hence the coefficients $\alpha_x, \alpha_y, \alpha_z$ have dimensions of $(length)^2$ compared to α_s . Thus a length scale in the x-direction is defined by

$$L_x = \sqrt{\frac{\alpha_x}{\alpha_s}} \quad (49)$$

As an approximate guideline, structure in the horizontal direction that has a scale length much less than L_x will be heavily penalized in the inversion, while structure that has a larger scale length will be admitted. In the DC resistivity example L_x was chosen to be 30m which corresponded to a ratio of α coefficients of about 10^{-3} .

In the above example, changing the ratio of the coefficients allowed structure to be preferentially smooth in the horizontal or vertical directions. Effectively these directions are tied to the cartesian framework in which the inversion is carried out. Only a minor modification is required so that preferential smoothness in a particular part of the model is orientated in an arbitrarily assigned direction (Li and Oldenburg, 2000). This can be beneficial when we have prior dip and strike direction associated with geologic structure and an inversion that reflects that character is sought.

Structural character can also be controlled by the exponent p in eq.(47). In principle p can be any positive number but generally it is chosen so $1 < p < 2$ because even that range can result in very different solutions. The value of the term in vertical braces increases with the size of p , so large values of the model, or large derivatives, carry a higher penalty weight when $p = 2$ compared to when $p = 1$. When $p = 2$, which is the value used in our objective functions thus far, sharp boundaries are manifested as smooth transition zones in the constructed model. Reducing the value of p allows more blocky models to be obtained.

To illustrate the ideas, we return to the applet example (although the supplied applet does not have the functionality to generate the following results). Figure 25 shows the result of minimizing the discrete form of

$$\phi_m(m) = \int |m|^p dx \quad (50)$$

for $p = 1$ and $p = 2$. This general measure is used to define the l_p norm of the model. The l_1 -norm solution is characterized by a set of impulses. The solution is “spiky”. That is, the final model is comprised of a few impulses interjected onto an otherwise zero model. The solution for $p = 1$ can be obtained through linear programming techniques, and a fundamental theorem of linear programming assures us that the maximum number of non-zero elements is equal to the number of data constraints. In this case, that is twenty. However, because the data are fit only to within a tolerance, the final number of impulses is reduced from this maximum. More general problems when $p > 1$ can be solved by using Newton methods for minimization, as outlined here, or solving a series of weighted least squares problems (e.g. Scales and Gerztenkorn, 1988). Irrespective of the method of solution, as p increases from 1 to 2, the character of the solution will change from the spiky solution in panel-b to the smoother solution in panel-a.

Minimizing the l_1 -norm of the model does not yield a desirable solution for our problem, but it is ideal for others. For instance, the earth model in reflection seismology is often assumed to be comprised of uniform layers, and seismic energy recorded at the earth’s surface is due to reflections at those boundaries. The processed normal incidence seismic data are ideally a bandlimited representation of the subsurface reflectivity as a function of time. These data can be inverted to recover a sparse spike representation of the reflectivity (Levy and Fullagar, 1981). The amplitude of the reflections are related to the acoustic impedance contrast between the layers and the time delay between the impulses is related to the thickness of the layers.

A better solution for our applet problem can be obtained by applying the norm to the derivative. Figure 25 shows the result of minimizing the discrete form of

$$\phi_m(m) = \int \left| \frac{dm}{dx} \right|^p dx \quad (51)$$

for $p = 2$ and $p = 1$. The $p = 1$ solution is characterized by sharp transitions zones and a nonzero gradient exists only on a few cells. This matches the character of the “block” in the true earth model. However, the Gaussian feature at depth is less well represented by this solution than it was for the l_2 solution in panel-a. The important point about the four solutions in Figure 25 is that they are all equally valid if judged solely from the viewpoint of fitting the data. However geological inferences drawn from these models would be considerably different. This stresses again the importance of trying to design a model objective function that generates a model that has the “right” character as determined from prior knowledge of the model.

We close this section by showing two examples. The first is a 2D cross-well straight-ray tomography problem where the data are linear functionals of the model. The second is a time-domain electromagnetic problem where the relationship between data and model is nonlinear.




Figure 25 l_2 and l_1 norm models for the applet example. Minimizing: (a) $|m|^2$, (b) $|m|^1$, (c) $|dm/dx|^2$, and (d) $|dm/dx|^1$.

For the tomography problem the 2D objective function we minimize

$$\phi_m(m) = \int \left| \frac{dm}{dx} \right|^p dx dz + \int \left| \frac{dm}{dz} \right|^p dx dz \quad (52)$$

with $p = 2$ and $p = 1$. The results are shown in Figure 26. The boundaries are smeared and there is a significant overshoot in the recovered amplitude for the $p = 2$ solution. The $p = 1$ solution is a much better representation of the block.

As a second example, we invert time domain electromagnetic data, in which the relationship between the data and the physical property is nonlinear. The example below is from Farquharson and Oldenburg, (1998). The layered earth model is shown in Figure 27. The transmitter was a $5m \times 5m$ square loop and the receiver measured the voltage in a horizontal coil at a distance 20 m from the source. The data, contaminated with Gaussian errors, are shown in Figure 27. Details regarding how to compute the responses can be found in Ward and Hohmann (1987). The earth is divided into 50 layers, and two inversions are carried out. α_s was set equal to zero in both, so the objective function is concerned only with variation of conductivity with depth. Panels (b) and (c) in Figure 27 show the models obtained when $p = 2$ and $p = 1$, respectively. The l_2 inversion smooths out the boundaries while the l_1 solution achieves a blocky model. The width and conductivity of the conductive layer are reasonably well reproduced and overall, this is a better model to interpret.




Figure 26 The l_2 – and l_1 –norm inversion models for the cross-well tomography. (a) true model, (b) data, (c) l_2 –norm model, (d) predicted data from l_2 model, (e) l_1 –norm model, (f) predicted data from the l_1 model.

4.2: “Tuning” the formulation of the inverse problem

The previous examples illustrate the effects of using different definitions for the model objective function in an effort to construct a model that has the right character. But despite the flexibility afforded by these objective functions, the result may still not be adequate. This is the case for inverting gravity or magnetic data. We saw earlier that surface gravity data, when inverted, produced a causative density anomaly appearing near the surface rather than at depth (see Figure 20). Changing to an l_1 –norm does not alter that fundamental characteristic. The basic understanding of this result lies in the nature of the kernel functions. They decay with distance from the source and thus changes in a physical property at locations close to observational point will have the greatest effect. In the case of gravity, the concentration of the physical property will be at the surface if the observations were above the surface. If the observations were in a borehole, then physical property variation will tend to be located along the borehole. The same conclusions apply to magnetic data. This is an unacceptable result and means that the algorithm must be modified. Here we introduce the magnetic problem and work with a synthetic example of a susceptible cube in a halfspace to illustrate how an inversion algorithm can be altered or “tuned” so that it alleviates artifacts and produces models that are more realistic. This tuning is really another way of incorporating prior information




Figure 27 The true conductivity model is shown by the dashed line in (b) and (c). The observed voltages in a loop receiver are in (a). The l_2 and l_1 -norm inversions of the data are in (b) and (c) respectively.

because it is implemented by adjusting the algorithm so that the solution looks like the synthetic input models that were initially used to generate the data, and that the synthetic models are considered to be representative of the earth.

If the effects of self-demagnetization and remanent magnetization are neglected, the anomalous magnetic flux intensity is related to the magnetic susceptibility, κ , by the integral equation

$$\vec{B}(\vec{r}_j) = \frac{\mu_0}{4\pi} \int_V \kappa(r) \vec{H}_0 \cdot \nabla \nabla \left(\frac{1}{|\vec{r}_j - \vec{r}|} \right) dv \quad (53)$$

where $\vec{H}_0$ is the Earth's local magnetic field and can be assumed to be constant over the volumes typically used in exploration. Measured data can be any component of the anomalous field or, as in total field surveys, the projection of the anomalous field onto the direction of the earth's field. Data may be acquired in the air, on the surface, or in boreholes. Discretizing eq.(53) by using a 3D cellularization in which the magnetic susceptibility is constant in each cell, and letting d_j denote an arbitrary datum yields $d_j = \sum_{k=1}^M G_{jk} \kappa_k$ where G_{jk} is the sensitivity that quantifies the effect of a unit susceptibility in the k 'th cell on the j 'th datum. Efficient formulae for carrying out the needed volume integrals over a prismatic cube are prevalent in the literature. We have used Sharma's (1966) formula.




Figure 28 Example of 3D magnetic inversion. (a) The 3D magnetic susceptibility, (b) total field data contaminated with Gaussian noise, (c) 3D susceptibility generated with a generic inversion, (d) predicted data from the model in (c).

A susceptibility model is shown in Figure 28(a) and the total field data, contaminated with 2 nT of Gaussian noise are shown in Figure 28(b). A generic inversion, using the model objective function with $(\alpha_s, \alpha_x, \alpha_y, \alpha_z) = (10^{-4}, 1, 1, 1)$, is provided in that figure along with the predicted data. The white square indicates the true boundaries of the susceptible block. The recovered susceptibility is concentrated at the surface.

The magnetic kernels in eq.(53) behave as $1/r^3$ where r is the distance between the source cell and observation location. For surface data, kernels for all data therefore have a $1/z^3$ intrinsic decay. This is illustrated in Figure 29 where a linear and logarithmic plot for the sensitivity along a column of cells below a surface observation point are shown in panels (a) and (b) respectively. The decay of the kernel is so severe that little structure can be seen on a linear scale. A log plot of the absolute value of the sensitivity must be used.

To allow the algorithm to generate models that have structure at depth, we introduce a depth weighting that counteracts this general decay of the kernels. One possibility (Li and Oldenburg, 1996) is to define a weighting

$$w(z) = \frac{1}{(z + z_0)^{\nu/2}} \quad (54)$$

where the parameters ν and z_0 can be chosen to maximize the match between $w^2(z)$ and the depth decay of the sensitivities. For magnetic data $\nu \approx 3$, and z_0 depends upon the observation height and cell size. The effect of the weighting is to alter the kernels, linear combinations of which make up the final model. In Figure 29(c) we show a




Figure 29 Depth decay of magnetic kernels. Panel (a) displays the decay of the kernel with depth plotted on a linear scale. Little variation can be seen beyond the depth immediately below the surface. Panel (b) shows the same decay plotted on a logarithmic scale. It changes in amplitude by three orders of magnitude. Panel (c) shows the kernels in one cross-section after a depth weighting has been applied. The amplitude is now much more balanced with respect to depth but there is enough structure both horizontally and vertically in the weighted kernel for model construction.

cross-section of a 3D kernel with depth weighing. We note that although the weighting function has compensated for much of the decay, variations in the horizontal and vertical directions still remain in the weighted sensitivities. Therefore, they are still useful as basis functions for constructing a model.

Incorporating the weighting into all components of the objective function can be accomplished by defining

$$\begin{aligned} \phi_m(m) = & \alpha_s \int_V w^2(z)(m - m_{ref})^2 dv + \alpha_x \int_V \left(\frac{\partial w(z)(m - m_{ref})}{\partial x} \right)^2 dv \\ & + \alpha_y \int_V \left(\frac{\partial w(z)(m - m_{ref})}{\partial y} \right)^2 dv + \alpha_z \int_V \left(\frac{\partial w(z)(m - m_{ref})}{\partial z} \right)^2 dv. \end{aligned} \quad (55)$$

Carrying out the inversion with a depth weighting produces the model in Figure 30(c). The model is much improved and the major susceptibility is now concentrated near the true depth. We note that the depth weighting has introduced

features in the image that might be considered undesirable. Specifically, the prism has a tail extending to depth and there are also broad regions of negative susceptibility. These artifacts can be greatly reduced by incorporating the additional constraint of positivity, which we will delineate in the next section.

The depth weighting described above is also important when inverting gravity data. In that case the exponential in eq.(54) should be modified to $\nu \approx 2$ since gravity data decay as $1/r^2$. When working with borehole gravity or magnetic data the weighting is modified to be a “distance” based weighting, or for general data sets it can be cast as a “sensitivity-based” weighting function. In these cases $w(z)$ premultiplying the model term in eq.(55) is replaced by $w(x, y, z)$. We note that the use of sensitivity based weighting can be important in non-potential field problems such as wide-aperture seismic refraction data (Wang, 1995) and also in DC resistivity and IP inversions where there are few current sources. Such is the case for gradient array surveys, where without any counteraction, substantial variation in conductivity or chargeability will occur at the surface of the earth and especially in the vicinity of the current electrodes. Sensitivity based weighting can counteract these effects. As the number of sources increases and we begin to acquire sufficient data to delineate 3D structure, then the need for a sensitivity weighting becomes less important.

When implementing a feature like depth weighting, some degree of empiricism needs to be incorporated. That is, synthetic models that typify the expected geologic structure should be created and inversions carried out. The effects of the weighting parameters in the model objective function can be evaluated and their values altered. In effect, this type of “tuning” of an inversion algorithm is a way of incorporating additional information and designing an algorithm that gives the right type of structure for known synthetic scenarios. In a recent paper, Haber




Figure 30 (a) true model, (b) model from generic inversion, (c) model with depth weighting but without positivity, (d) model with depth weighting and positivity

and Tenorio (2003) add more mathematical structure to this approach and show how a best model objective function can be obtained from a suite of prior models or training models.

Having generated the algorithm, it is then essential to observe how the inversion behaves when applied to data from models that are different from that which was used to generate the modifications. As a case in point we consider the magnetic problem. The depth weighting was designed and tested for a buried prism. Inversion showed that the constructed model was at approximately the right depth. However, it would be completely unacceptable if the modifications to the algorithm were such that the constructed models always had features at depth irrespective of the input data. Subsequent tests of the inversion of magnetic data, arising from scattered blocks at the surface, resulted in models in which the susceptibility was concentrated at the surface and not put to large depths even though a depth weighting had been used. This provides some comfort that the inversion algorithm hasn't been over-designed.

4.3: Bound constraints

In the 3D magnetics example, application of depth weighting prevented the concentration of recovered susceptibility at the surface. However, the resulting model was unsatisfactory because the main susceptibility feature was flanked by negative zones. This was because the observed data contain negative anomalies due to the oblique angle of magnetization and anomaly projection, and it was much easier to reproduce these data by putting negative susceptibility in the model. This problem can be rectified by imposing a positivity constraint on the recovered susceptibility.

More generally, in many problems, hard-constraint prior information about physical property values is available from petrophysical logs or from other sources. For instance the physical property may be known to be positive (always the case for the electrical conductivity, generally the case for magnetic susceptibility, and sometimes the case for density contrast). At an even more informative level, the physical property may be known to lie within particular bounds, $[m^- < m < m^+]$. We want to include this information into the inversion and solve the optimization problem

$$\begin{aligned} \text{minimize : } & \phi = \phi_d + \beta\phi_m, \\ \text{subject to : } & m^- < m < m^+. \end{aligned} \quad (56)$$

There are numerous ways to attack the solution of eq.(56). In the examples presented in this paper we use a primal logarithmic barrier method. The logarithmic barrier method was originally developed for solving linear and quadratic programming problems with inequality constraints (e.g. Gill *et al.*, 1991; Wright, 1997). Basically it adds a barrier term in the objective function to implement the bound constraints. In general, if an upper and lower bounds are imposed on each cell then the objective function takes the form (e.g, Gill *et al.*, 1991)

$$\phi(m) = \phi_d + \beta\phi_m - 2\lambda \left\{ \sum_{j=1}^M \ln\left(\frac{m_j - m^-}{m^+ - m^-}\right) + \sum_{j=1}^M \ln\left(\frac{m^+ - m_j}{m^+ - m^-}\right) \right\} \quad (57)$$

where λ is the barrier parameter, and the tradeoff parameter β is fixed. As the name suggests, the logarithmic barrier term forms a barrier along the boundary of the feasible domain and keeps the solution within the user-imposed bounds. Minimizing eq.(57) is a nonlinear problem and hence linearization and iteration is required. The solution is obtained by solving a sequence of nonlinear minimizations with decreasing λ . As λ approaches zero, the sequence of solutions approaches the solution of eq.(56). The details of our implementation for the gravity and magnetic problems are provided in Li and Oldenburg (2003). If only positivity is imposed, it can be implemented by setting $m^- = 0$ and m^+ to be a large number.

The benefits of incorporating positivity in the magnetic inversion are illustrated in Figure 30(d). The negative zones have disappeared and so has the susceptibility tail. This is a much improved model of the susceptibility structure.

As final examples, we return to the 2D IP and 3D gravity problems illustrated earlier. In the case of IP inversion, physics dictates that the chargeability is positive. We imposed this constraint in the inversion of data shown in Figure 19(f). This generates a much sharper image of the buried chargeable block compared to that obtained without positivity. In the gravity problem, we suppose that the density contrast for the tunnels is $-1.5 < \rho < 0$. Solving the inverse problem with these constraints yields the model shown in Figure 20(e). The recovered density contrast is much more compact and delineates both the dipping and horizontal tunnels.

4.4: Field Example: 3D inversion of magnetic data

The synthetic model for magnetic susceptibility illustrates that a nonlinear inversion which includes positivity and depth weighting can work well. We have “tuned” the algorithm on a simple problem, but a realistic test is required. We close this section by inverting total field magnetic anomaly data acquired over a mineral deposit.

Total field magnetic data taken over the Raglan Deposit in Northern Quebec are shown in Figure 31. Two regions of high magnetic field are observed. These coexist with ultramafic outcrops and the geologic question was whether the outcrops were associated with a single flow unit or with isolated zones that extend to depth. For the 3D inversion, the earth was modelled with a 40x40x10 grid of cells (16,000 cells) and the horizontal dimensions of the cells was 100 x100 m. The data were assumed to have noise of 2% plus 5 nT. In defining the objective function we set $\alpha_s = 0.0001$, $\alpha_x = \alpha_y = \alpha_z = 1$ and the reference model equal to zero. Positivity was included in the inversion. The predicted data are shown in Figure 31.

An important aspect of 3D inversion is the ability to display a 3D model. A 3D volume of susceptibilities can be displayed in cross-sectional cuts or as a 3D-isometric image after volume rendering. Cross-sectional cuts have the advantage that values of the physical property in each cell are displayed. These are the numbers that reproduce the field data and from which quantitative questions can be answered. The disadvantage of the cross-sections is that it is sometimes difficult to get a sense of 3D structure. This is where 3D volume rendered images have an advantage, but care must be taken since a different cutoff for volume rendering will generate a different image. The 3D isometric image in Figure 31 has a threshold level of 0.04. Susceptibilities less than that are transparent. The image shows that the surface outcrops are connected by a flow unit at depth. This was verified by a subsequent drilling, so the inversion results contributed critical information for clarifying the geological understanding of this deposit.




Figure 31 3D magnetic inversion at Raglan Deposit. The top right diagram shows the total field aeromagnetic data. The inducing field has $I=83^\circ$ and $D=-32^\circ$. The data are contoured in nT on the scale shown by the grayscale bar. The volume rendered image of the inverted susceptibility model is shown at the left. The displayed surface provides a representation of the ultramafic flow. Indicated in this representation is the intersecting drillhole that was spotted based upon the inversion result. The bottom right image shows the predicted data. Note that these are an excellent representation of the observations. Large misfit occurs only near rapidly varying features corresponding to near-surface susceptibility variations.

Section 5.0: CHOOSING A REGULARIZATION PARAMETER

The inverse problem, when formulated as an unconstrained optimization problem, involves minimizing $\phi(m) = \phi_d + \beta\phi_m$. Thus far we have concentrated upon defining ϕ_m and showing its importance in generating a solution that is interpretable. Issues with respect to the noise have been disregarded by assuming that the noise was Gaussian, independent and uncorrelated and had known standard deviation. The inverse problem was solved by finding the model that minimized the model objective function and produced a target misfit ϕ_d^* where the numerical value $\phi_d^* \approx N$ was obtained using the discrepancy principle. That is, the final model was designed to fit the data to the degree justified by a well-defined statistical measure associated with Gaussian statistics. In practise however, the noise on the data does not conform to these Gaussian assumptions and, even if it did, the standard deviation of the data are not generally available.

“Errors” in the data constraints used in the inversions arises from different sources. The data equations can be written as

$$F_j[m] = d_j^t + n_j + \delta_j \quad (58)$$

where d_j^t is the true signal. The additive noise n_j can arise from natural or man-made sources such as the measuring instruments, wind, power lines, automobiles etc. The δ_j represent modelling errors. In reality, the forward modelling is mathematically expressed as $\mathcal{F}[m] = d$ where m is a function. For purposes of analysis, we work with a numerical approximation $F[m]$ where F is a discrete operator (finite difference modelling or integral equation representation) and m is a finite vector of parameters. Any discrepancy between $\mathcal{F}[m]$ and $F[m]$, that is, any difference between the true forward modelling and the approximate forward modelling that arises from numerical implementation, is encapsulated in the δ_j . This could result from using too few cells in a finite difference solution, or not solving a system of equations accurately enough. Differences also arise from more fundamental reasons. The numerical solution might assume the earth is 1D or 2D, instead of 3D. Or, some basic physics may not be accounted for in the forward modelling operator. For instance, when inverting IP data we work with the DC resistivity equations but data could be contaminated by EM coupling which requires that the complete Maxwell’s equations be used.

We therefore see that “errors” are complicated and they are not likely uncorrelated and Gaussian as we have assumed. Even if we tried to treat errors as correlated, we don’t know what the correlation is. So we’ll stay with our assumption about independent Gaussian statistics and do as well as we can. To proceed in practical problems we take the following route. Let $d_j^{obs} = d_j^t + n_j$ and assume that the noise can be represented as a Gaussian independent random variable. Initial estimates for the standard deviations ϵ_j are required. This is not a simple operation and it is likely survey dependent. A lower bound for ϵ_j can be obtained from repeated measurements if they are available. Other factors to be incorporated include effects of survey parameters such as mislocation and orientation of sources and receivers. What is important at this stage, is that ϵ_j approximately characterizes the correct relative errors. That is, if one datum has twice the standard deviation as another, then the estimated standard deviation should reflect this. If this can be accomplished then the normalized data in $||W_d d^{obs}||$ all have the same standard deviation but the absolute noise is indeterminate by a constant scaling factor. A useful estimator of the standard deviation is $\epsilon = h_1 + h_2 * |d^{obs}|$; that is, a base level plus a percentage value. Percentage errors, in themselves, are not appropriate when there are zero crossings of the data.

Assignment of incorrect standard errors in the inverse problem, and finding a model with $\phi_d^* \sim N$, results in overfitting or underfitting the data. The consequences of this can be examined by using the applet. If data are overfit, excessive structure that is purely an artefact of the noise will be observed. If the data are underfit, then there is information about the model that has not been extracted from data. Ideally, the generated model should

reproduce the signal exactly, and misfit entirely the additive noise. That is, if $d_j^{obs} = d_j^t + n_j$ then the constructed model ideally has the property that $F[m] = d^t$ so the misfit should be the expected value of

$$\phi_d = \sum \left(\frac{d_j^{obs} - F[m]}{\epsilon_j} \right)^2 = \sum \left(\frac{d_j^t + n_j - F[m]}{\epsilon_j} \right)^2 = \sum \left(\frac{n_j}{\epsilon_j} \right)^2 \quad (59)$$

where ϵ_j is an estimate of the true standard deviation of the datum. But if the estimated standard deviations are in error, say by a constant scaling factor, then so too is the expected misfit. This uncertainty in size of ϕ_d can be accommodated by adjusting the regularization parameter β . In this section we present two ways, the L-curve and GCV methods, for automatically estimating β . Both techniques are widely used in practical problems.

Demo applet: The effect of incorrectly specifying the data error in the inversion

The easiest way to investigate incorrect assignments of standard deviations is to use the parameter “*chifact*”. Let ϵ_j be the estimated standard error for the j 'th datum. It is related to the true standard error ϵ_j^t by $\epsilon_j = c\epsilon_j^t$ where c is a constant. Then the evaluated misfit

$$\phi_d = \sum \left(\frac{d_j^{obs} - F[m]}{\epsilon_j} \right)^2 = \sum \left(\frac{d_j^{obs} - F[m]}{c\epsilon_j^t} \right)^2 = \frac{\phi_d^t}{c^2} \quad (60)$$

So the expected value of the misfit is N/c^2 . That is, if we had underestimated the standard deviation of the data by a factor of two ($c = 0.5$) then the misfit we should be asking for is $4N$. In the applet, the parameter *chifact* performs this role.

1. Start the applet and choose a default model, default kernels, and default noise level. Then perform the inversion using *chifact*=1. This is the appropriate level of misfit.
2. Rerun the inversion with *chifact*<1 (e.g., 0.1). Note the increase in structure in the recovered model.
3. Invert the data again with *chifact*>1 (e.g., 10). Note the loss of details in the model.

This is the same exercise as a previous one on the effect of underfitting and overfitting the data. In fact the earlier exercise is designed to show the importance of fitting the observations to the correct level.

5.1: Automatic estimation of regularization parameter

As discussed above, the purpose of specifying a correct value for the regularization parameter is to achieve the correct level of data misfit. Therefore, estimating the regularization parameter is equivalent to estimating the noise in the data. In the case of unknown noise level in the data, we have a couple of choices. One option is to first supply an estimate of the noise, fit the data according to a discrepancy principle, assess the recovered model visually, make a subjective decision about whether the data have been overfit or underfit, and then rerun the inversion asking for a different misfit. Equivalently, the same procedure is invoked by trial and error runs using different values of β . This is an interactive and subjective process. An alternate procedure is to estimate the noise level as an integral part of the inversion process. Since we usually do not have enough independent information to estimate noise level for individual data, we opt to estimate the overall level of noise in the data. We first assume that the data have been normalized by their relative standard errors so that the scaled data have errors that are specified by a Gaussian distribution with an unknown variance. The goal of estimating this unknown variance of the noise is equivalent to finding the regularization parameter β . There are two approaches for estimating β . The first is the L-curve criterion.

This is a somewhat heuristic approach and uses the general character of the Tikhonov curve. The second is the generalized cross-validation (GCV), which has a more rigorous mathematical basis.

5.1.1: L-curve

The use of the L-curve for solving linear inverse problems was first introduced by Lawson and Hanson (1974) and has been championed by Hansen (1998) who presents extensive background and numerous examples. The reader can refer to these two books for details. Here we present only the essential ideas.

The heuristic process for determining β follows directly from the analysis of the Tikhonov curve done in Section 2. Figure 32 shows an illustration of a typical Tikhonov curve. In region-I of the Tikhonov curve where β is large, the misfit can be greatly reduced without requiring large amounts of structure to be added to the model. In this regime, the model is being built up using basis vectors that have long wavelength. Correspondingly we are fitting the long wavelength structure in the data which has a greater likelihood of representing the physical model. In region-II, however, there is very little reduction in misfit even though a large amount of structure is being added to the model. In this regime we are fitting the high frequency component of noise in the data and additional structure in the model is not physical. The Tikhonov curve generally exhibits two straight regions with a sharp corner between them on a log-log plot. This general behavior is described as an L-curve. The previous commentary argues that the desired β is away from the two straight regions and Hansen's suggestion is that an optimum value corresponds to the "knee" of the L-curve. At this point, the coherent signal in the data has been reproduced, yet the recovered model is not adversely affected by the noise in the data. The corner point can be evaluated numerically by plotting (ϕ_d, ϕ_m) on a log-log plot and looking for the point of maximum curvature. In Figure 33, we plot the curvature as a function of β for the applet example. The L-curve estimate for β is 1700, which is slightly smaller than that obtained from the discrepancy principle. The resulting data misfit is also slightly lower. The corresponding model is shown in panel-c and it compares well with the recovered model shown in Figure 10.

The estimation of the regularization parameter in nonlinear problems can also be done using the L-curve. Simplistically, one can carry out a series of inversions to produce the full Tikhonov curve, find the β that corresponds to point of maximum curvature, and then carry out the final inversion. This would be an expensive proposition. Alternatively, we can estimate an optimal regularization parameter for each iteration of the non-linear inversion. It




Figure 32 Tikhonov curve: In region-I the data are underfit and small amounts of additional structure can greatly reduce the misfit. In region-II the data are overfit. Large amounts of structure added to the model produce only a marginal decrease in misfit.

is employed as follows. At each iteration a linear system of equations, such as eq.(45), needs to be solved. This system is solved for a sequence of β 's and a Tikhonov curve is generated. The value of β corresponding to the maximum curvature is chosen and used in eq.(45) to solve for a perturbation δm . The underlying assumption is the following. There are two kinds of error at each iteration. The first is the noise in the observed data, and the second is the linearization error resulted from neglecting the high order terms in the Taylor expansion of $F[m]$. Fitting either of these errors generally requires excessive structure in the model. The L-curve estimate of regularization parameter is therefore an appropriate value to use at each iteration. As the iterations progress and a final solution is approached, the linearization error becomes smaller and the required regularization parameter becomes smaller. At convergence, the linearization error diminishes and the L-curve criterion effectively estimates the original data error. Our experience with the L-curve in many practical experiments has borne this out well. (Li and Oldenburg, 1999; Farquharson and Oldenburg., 2000).




Figure 33 Determining the optimal regularization parameter using an L-curve criterion. Panel (a) shows the Tikhonov curve of the 1D example with L-curve estimate of β marked by a *. (b) is the curvature of the L-curve as a function of β . (c) is the constructed (dashed) and true (solid) model.

Demo applet: Explore the utility of the L-curve and GCV criteria

1. Start the applet and use the default model and other parameters. Invert the data by setting $chifact=1$ in the lower right portion of the applet, and press the “invert” button. Note the characteristics of the inversion result including the recovered model, values of the regularization parameter and data misfit. The applet also displays the Tikhonov curve and GCV curve that were constructed during the inversion.
2. Choose the L-curve option for regularization parameter. The applet will automatically update the inversion to display the result of using the L-curve criterion. Note that the brown dot indicating the choice of regularization parameter moves down to the corner of the Tikhonov curve, and the data misfit is slightly smaller, and the amplitude of the recovered model increases slightly. Over all, however, the result is very similar to that obtained with $chifact=1$.
3. Next, choose the GCV option. Now, the applet will display the GCV inversion result. The GCV estimate of regularization should be close to the L-curve value, but it can be either slightly greater or smaller depending on the particular realization of noise added to the data.
4. One can cycle through the three options ($chifact=1$, L-curve, and GCV) to examine the change in the regularization parameter, the resulting model, the data misfit, and model objective function value. To further explore the behavior of the two automatic estimation methods, press the “noise” button and alter either the magnitude of noise added to the data, or change the seed to get a different realization of noise. This allows one to gain intuition about the variation of the L-curve and GCV estimates compared with the result from discrepancy principle

5.1.2: Generalized Cross Validation (GCV)

Estimating the unknown noise level has been a much studied problem in statistics. One popular technique is General Cross Validation (GCV) (Wahba, 1990; Golub *et al.*, 1979; Golub and Von Matt, 1997). The basic idea for GCV is that a good solution to the inverse problem is one that is not unduly sensitive to any particular datum. This means that the solution will predict a datum even if that datum is not used when calculating the model. This accounts for the name cross-validation. The desired value of β is that which best predicts all of the data.

Conceptually the computations could proceed as follows. We first pick a value of β . For an arbitrary n 'th point, discard it from the data set to generate a reduced data set $\tilde{d}^{obs(n)}$. The inversion of the new data set is obtained by minimizing

$$\phi_n(m) = \left\| Gm - \tilde{d}^{obs(n)} \right\|^2 + \beta \|W_m m\|^2 \quad (61)$$

Let $m_n(\beta)$ denote the solution. Next carry out the forward modelling to predict d_n^{pred} , the value of the missing datum. The difference between the predicted and the observed datum is accumulated in the Cross Validation (CV) function. When the above procedure is carried out for all values of $n = 1, N$ we obtain

$$CV(\beta) = \sum_{n=1}^N (d_n^{pred} - d_n^{obs})^2 \quad (62)$$

We now repeat the above procedure with other values of β and compute the CV curve as a function of β . The minimum of the CV function identifies the β for which the data change the least if we discard arbitrary data values.




Figure 34 Tikhonov curve with GCV estimate marked by a *. (b) GCV curve (c) constructed and true model.

The CV as defined above is not invariant under orthogonal transformations of the equations and it was thus replaced by the Generalized Cross Validation (GCV) function (Golub et al., 1979). Subsequent improvements regarding how to compute this function (Wahba, 1990; Golub and Von Matt, 1997) have led to evaluating

$$GCV(\beta) = \frac{\|d(\beta) - d^{obs}\|^2}{(\text{trace}(I - C(\beta)))^2} \quad (63)$$

where

$$C(\beta) = G(G^T G + \beta W_m^T W_m)^{-1} G^T \quad (64)$$

In eq.(63) $d(\beta)$ is the predicted data obtained from inverting all of the observations.

The GCV curve for our template problem is shown in Figure 34. There is a well defined minimum and the corresponding value of β produces the model shown in that figure.

If you are exploring with the applet then you will notice that the β value from the GCV tends to overfit the data slightly, and sometimes there is no minimum, and hence GCV can fail. However, in most cases it works fairly well. As a practical matter, the GCV is a statistical estimator and works better as the number of data increases. In general the number of data should be 30 or more, but we have successfully used it in examples with fewer data.

The GCV methodology can be applied to nonlinear problems in the same manner as the L-curve. (Haber and Oldenburg, 2000). At each iteration the sensitivity matrix J replaces the matrix G above. The GCV estimate for the linearized equations is chosen and β is obtained by finding the minimum. The underlying philosophy is the same as that in the L-curve approach. That is, at each iteration, the right hand side of the equations to be solved

is contaminated with two types of errors. There are correlated errors arising from nonlinear terms and additive Gaussian noise. The β estimate from GCV tends to be associated with the additive random noise. The correlated nonlinear errors will then be reproduced by the constructed model. As such, some damping at each iteration is required. Haber and Oldenburg (2000) used a damped Gauss-Newton strategy. Farquharson and Oldenburg (2000) introduced a cooling to limit the rate at which β could proceed from its initial value to the final value generated by the GCV analysis.

As an example for the use of L-curve and GCV in a nonlinear problem we show the inversion of airborne electromagnetic data to simultaneously recover both a 1D conductivity and magnetic susceptibility structure. The earth model consists of two conductive and susceptible prisms buried in a homogeneous earth. The data are the vertical components of the secondary magnetic field measured at 10 frequencies from 110 Hz to 56KHz. The source is a vertical magnetic dipole and the receiver is at a distance of 10 metres away at the same height. Synthetic responses were generated along a line that crossed the two prisms and Gaussian noise was added. Prior to inversion, the data were weighted by their standard deviations. At each station, 1D inversions were carried out using the L-curve and GCV strategies. The inversions were stitched together to form 2D images shown in Figure 35. In the L-curve images, the two regions of high conductivity are outlined although there are side-wings and overshoots that are characteristic of fitting 3D structure with 1D models. The regions of imaged susceptibility have fewer artifacts and correspond quite well with the susceptible prisms. Of particular note however, is the behavior of the misfit functional. At the edges of the model, where the earth is 1D, the reconstructed model fits the data to a level that would be expected using chi-squared statistics. That is, the misfit values are about 20 which is equal to the number of data. Larger misfits are observed in the location of the anomalous bodies. In those regions the data are severely affected by 3D effects and it is not possible to fit those data with a 1D earth. The L-curve has recognized this. The same general statements apply to the GCV results. In fact, the GCV misfits and final models are similar to those of the L-curve. The GCV results however, show slightly more structure and it is noted that the misfit is somewhat reduced compared to the L-curve misfits.

We find the above result quite encouraging because of the difficulty of the example. Not only are the data nonlinear functionals of the physical properties, but there are “modelling” errors as a consequence of representing a 3D earth with a 1D approximation. Yet, both automatic estimators gave similar and reasonable results. There are differences in detail, but it must be remembered that the β obtained from the L-curve or GCV should only be regarded as a first estimate for the regularization. As with all inversions, if there appears to be too much or too little structure, then the regularization can be adjusted and the inversion rerun. What is of major importance here is that the β estimated from the GCV and L-curve provides a good starting value for any such experimentation.

In summary, the choice of regularization parameter is extremely important because it controls the amount of structure in the final image. In the above, we have provided basic strategies that are applicable for linear and nonlinear problems. None of these methods is a panacea, but each provides a possible outcome that can be interpretable, or it provides a ball-park number for the regularization, which the user can then adjust. This also re-enforces our earlier comment that inversion is not a blackbox and it should not be treated as a turn-key operation. In cases of unknown data error, the user’s judicious interaction in estimating data error and, therefore, the value of regularization parameter, is of crucial importance.

Section 6.0: PARAMETRIC INVERSION

The concern of the chapter thus far has been to recover a function, usually a physical property distribution, that gave rise to the observations. To proceed numerically, we discretized the problem and assumed that the model was




Figure 35 L-curve and GCV criteria for estimating the regularization parameter in nonlinear inverse problems. The data, not shown, are the in-phase and quadrature phase signals at 10 frequencies for a horizontal coplanar system flown 30 meters above a 3D earth containing two conductive and magnetic prisms. The results for the L-curve are shown on the left. (a) misfit, (b) recovered conductivity, (c) recovered magnetic susceptibility. The white rectangle shows the true location of the prisms. The corresponding GCV results are shown in the right-hand column.

constant in each cell. The number of cells was considerably larger than the number of data and hence a large under-determined system of equations was to be solved. A specific solution was obtained by minimizing $\phi = \phi_d + \beta \phi_m$ where ϕ_m incorporated prior information about the model and β was adjusted so that a desirable misfit was achieved.

An alternative strategy is to formulate the inverse problem in terms of a few parameters that need to be recovered and to work with data sets in which the number of data is greater than the number of parameters. This leads to a parameter estimation problem and the overdetermined problem can be solved by least-squares. This approach is commonly used in geophysical literature and there are a number of reasons for adopting this procedure. We list some below:

1. The earth structure, (or the solution of the posed inverse problem), is accurately specified by the values of some parameters. For instance, it might be known that the earth consists of three flat-lying layers and we seek to recover the velocities and thicknesses of the layers. This prior information is exceedingly powerful in limiting the space of acceptable models and, when it is available, there is no justification for treating the inversion as an under-determined problem. Another example is the earthquake location problem. If the background velocity is known, and if the earthquake is located at a point, then the inverse problem is formulated to recover the three spatial coordinates of the earthquake hypocenter (e.g. Stein and Wyssession, 2003).
2. On some occasions the computational complexities of the forward modelling are too difficult to be used in practise and hence the true forward modelling is replaced by an approximate forward modelling that is controlled by a few parameters. Practical work in finding unexploded ordnances (UXO) often resorts to this. A UXO

is an unexploded munition such as a hand grenade, mortar, or bomb. Its detailed structure is complicated and it can be composed of multiple materials with differing physical properties and have complex asymmetric structures such as tail fins and brass rings. If the electromagnetic sensor is close to such an object then full 3D electromagnetic forward modelling is required to simulate the responses. However, if the sensor is a scale length or so away from the UXO then its response is well approximated by that from electromagnetic dipoles. The goal of the practical UXO inverse algorithm is to locate the object and determine whether it is likely to be a UXO or uninteresting shrapnel that should be left in the ground. The parameter-based inversion in this case is sufficient and it has the added benefit of greatly reducing the computational burden (Pasion and Oldenburg, 2001).

3. A parameter-based model is introduced to reduce computational costs of solving the inverse problem and also to generate a “simple” solution that can be readily interpreted. For example, it is not uncommon to represent the earth by a few homogenous layers in which the physical property values and the layer thicknesses are unknowns. Since there are only a few parameters, and the number of data is considerably greater than the number of parameters, there may be no need to introduce additional regularization into the problem. By using only a few layers, the final model is “simple”. Simplicity of the result along with the reduced computation time makes this an attractive option. Unfortunately, an incorrect choice of parameters (e.g. the number of layers) can greatly alter the outcome.

Before turning to the solution of the parameter estimation problem, valuable insight can be obtained by examining the strategy for achieving simplicity by reducing the number of parameters in the problem. We will use our applet example to explore the effects of reducing the number of layers and thereby make a transition between underdetermined and overdetermined systems. This also allows us to discuss the effects of discretization on the inverse problem. (Note however that the supplied version of the applet does not allow the number of cells to be changed.)

In designing a discretization, we are governed by two conflicting ideas. On the one hand we want cells small enough to emulate complicated geologic structures. On the other hand, we want to work with a minimum number of cells so the computational cost is reduced. In practise, it is sufficient to work with cells whose size is significantly less than the resolving power of the data. A sufficient test is to verify that the recovered model does not significantly change when the cell dimensions are reduced. We illustrate this with the generic applet example. The constructed model, using 100 cells and corresponding to a misfit of $\phi_d = 20$ shown in Figure 10 is reproduced in 36(b). The model obtained when the number of cells is 200 is shown in Figure 36(a). There is no discernible difference. However, differences begin to arise as the number of cells is reduced. In the remaining panels of that figure we show the results of inverting with 50, 20, 10 and 5 cells. In each case we attempted to find a model that had a misfit of 20. When the number of cells was less than 10, this was no longer possible even when β was reduced to zero. In the case where $M = 5$, the misfit was $\phi_d = 40$.




Figure 36 Effects of model discretization on the inversion. The panels (a)-(f) show respectively the recovered models using 200, 100, 50, 20, 10, 5 cells. All models, except (f), have a misfit of 20.

This example illustrates a couple of important points. The first is that there is a cell size that is sufficiently small so that the discretization does not significantly affect the inversion result. Using cell sizes that are smaller than this will not improve results, it only serves to increase the computations. The second point is that the parameterization can act as a regularization. As the number of cells is reduced, model space becomes so restricted that no additional regularization is required and a solution can be obtained by minimizing the misfit only.

Formulating the inversion to minimize only the misfit, and regularizing by the parameterization, can be successful, but it can also be problematic. The following example illustrates the essential issues. Figure 37(b) shows the best fitting model when there are 10 cells. The fit is $\phi_d = 18.4$ and the recovered model is a good representation of the true model. It is also very similar to that obtained in Figure 36(e). With five cells the reconstructed model shows two regions of energy. The first block is recreated fairly well because the model discretization happened to coincide with the true width of the block (this is unlikely to happen in practice). The smooth feature is somewhat mislocated because of discretization. With 20 cells, the misfit is reduced to 15 and the model elements essentially reflect the algorithm's attempt to fit the noise. Note that the vertical scale has units of 10^7 . At this point the problem has been over-parameterized and additional regularization is required.




Figure 37 Best fitting models obtained by using different discretizations. Models recovered by using 5, 10, 20 cells are respectively shown in (a)-(c). Note that the scale in panel (c) is 10^7 .

6.1: Parameter Estimation

There is extensive literature, both in mathematical and geophysical texts, concerned with estimating values of parameters from observed data sets and it would not be constructive to repeat that information here. Our goal is limited to presenting some essential aspects about this problem and to contrast it with the general inversion problem we have treated thus far. Our short commentary should not be interpreted to imply that we don't think this problem is important. Rather, we can be short because of the material that exists on this subject, and also because some basic strategies for solving the parametric inverse problem have already been dealt with in this chapter: from a mathematical viewpoint, the parameter estimation problem is a subset of the general inversion methodology that we have been investigating.

In the parameter estimation problem the unknowns can be represented by a vector $m = (m_1, m_2, \dots, m_M)$ but the elements might denote physical property values, layer thicknesses, location coordinates of corners of a body etc. Since the parameters generally have different dimensions and sizes, most estimation procedures work with normalized quantities that have been scaled by typical values. The new parameters are dimensionless and have magnitudes in the order of unity. There must be more data than parameters so that the problem is overdetermined. The goal of the inverse problem is to find values of the parameters that produce a best fit to the observations, so the inverse problem is still formulated as an optimization problem, but the objective function consists only of a

misfit term. The least squares function

$$\phi_d = \sum \left(\frac{d_j^{obs} - F[m]}{\epsilon_j} \right)^2 = ||W_d(d_j^{obs} - F[m])||^2 \quad (65)$$

where ϵ is an estimated standard deviation of the data, is typically used but other representations of goodness of fit are viable. The immediate advantage of working with the overdetermined problem is that the problem seems simpler. Compared to solutions for the underdetermined problem, the model objective function has been omitted, as well as the difficult task of estimating a regularization parameter. Also, there are generally fewer parameters to estimate and the forward modelling is often faster to solve because of the reduced modelling. The tradeoff is that the relationship between the model parameters and the data can be highly nonlinear. This can translate into an objective function that is quite irregular, having multiple minima, valleys, or even cliffs (see, for example, Scales *et al.*, 1992; Billings *et al.*, 1994, and Deng and Scales, 1999). It is constructive to think about the objective function as being a geographic topographic surface. The elevation is $z = F(x, y)$ and so there are two variables. We desire to find the value of x and y that makes z minimum. Imagine standing at a particular location and wanting to find the lowest elevation point, z_{min} . Our natural instincts are to go “downhill”. We look for the direction of steepest descent and follow it. If the topography is a smooth bowl then we readily arrive at our destination. (see Figure 38(a)) Alternatively, if the topography is tortuous (Figure 38(b)) then the nearest local minimum may be far from our goal. When this occurs, the only way to make progress is to climb up the local topography and proceed in a direction that seems more promising. The Newton-type procedures employed thus far are guaranteed to take us to the nearest local minimum, but there is no inherent mechanism to recognize that the global minimum has been found. Confidence can be achieved if we continue to arrive at the same solution when we begin with different starting models, but we cannot circumvent the reality that our descent algorithm will be successful only if we start in the valley of the global minimum.




Figure 38 Illustration of a smoothly varying data misfit function with a single global minimum (panel-a), and a complex misfit function with multiple minima (panel b). In both cases, the global minimum corresponds to the best fitting model. The values of the function are plotted as a contour diagram at the bottom face of the cube. In (b) the contours are tortuous and there are multiple minima with valleys.

The above deficiency of descent algorithms has been one of the motivations for developing different direct search methods. Such methods aim to explore the entire parameter space to locate the parameter values corresponding to the global minimum. For our 2-parameter topographic surface, this means sampling the x and y directions, and selecting the location corresponding to the smallest elevation.

The simplest and most intuitive approach would be to sample each parameter uniformly within a given interval. Let us consider the objective function in Figure 38(b) and extend it to higher dimensions. Suppose that the topographic area is $1\text{ km} \times 1\text{ km}$ and we want to be within 10 meters of the global minimum. One way to accomplish this is to sample each direction uniformly at N points. If we chose to sample the x -direction every 10 meters then $N = 100$. Correspondingly, the total number of samples needed in two dimensions is N^2 . For higher dimensional problems, the number of samples required increases exponentially with dimensionality. If M parameters $(m_1, m_2, \dots, m_M)$ are to be found, then straightforward sampling requires N^M evaluations of the forward modelling $F(m)$. Even when $F(m)$ can be easily evaluated, this strategy limits the number of unknowns to be less than about 10. The enormous computational cost involved illustrates the need for more efficient sampling. This leads to the development of Monte Carlo techniques for global minimization.

In Monte Carlo techniques, random models are selected from an infinitude of possibilities, and forward modelling is carried out to evaluate data misfit. Models are accepted or rejected depending upon how well they fit the data and the best fitting model is found. The practical difficulty with such algorithms is the challenge of making them efficient by sampling as few models as possible. Various strategies, including genetic algorithms (Goldberg, 1989), simulated annealing (Kirkpatrick *et al.*, 1983), and neighborhood algorithms (Sambridge, 1999a), and others have been developed. Sambridge and Mosegaard (2002) provide an excellent introduction and tutorial for Monte

Carlo methods and their use in geophysical inverse problems. That paper also has an extensive reference literature that serves as a gateway to this area of research.

Significant improvement regarding the size of the problem can be made by using importance sampling, that is, attempting to isolate those regions of parameter space in which the minima are more likely located (again see Sambridge and Mosegaard, 2002). However, even these techniques become onerous as the number of variables increases, and practical limits with today's computations limit the problem size to a few hundreds of parameters, or perhaps a thousand if the objective function is particularly easy to evaluate.

A positive feature of the global search algorithm is that a change in the objective function is trivially handled. This is because only the value of the objective function, not its derivatives, needs to be evaluated. This flexibility enables one to utilize different misfit measures or statistics with little extra effort. For example, our measure of misfit has been a squared criterion (65), but this is only one choice. If the data errors have outliers, then a more robust misfit is the l_1 -norm

$$\phi_d = \sum \left| \frac{d_j^{obs} - F[m]}{\epsilon_j} \right| \quad (66)$$

In fact any l_p -norm, (Huber (1964), Eklom (1973, 1987)) might better represent the statistics of the errors and hence be a better misfit to minimize. These norms can easily be incorporated into a global search algorithm. On the other hand, the majority of descent algorithms require that $g = \nabla \phi_d$ be evaluated. This can be complicated by the functional form of ϕ_d and the need to compute $\partial F / \partial m$. We note that if the problem is such that this latter quantity cannot be evaluated, the user has no choice but to use a global search method.

Despite the potential problems of using gradient algorithms for parametric problems they are still successfully used. For example, Christensen (1996, 1997) use them in the interpretation of time domain EM data for ground water exploration; Pasion and Oldenburg (2001) and Billings *et al.* (2002) use them for UXO characterization problems. In the following we provide a few comments about their solutions. We then provide an example of using both Newton-type and direct search algorithms to determine if a buried object is likely a UXO.

A Newton-type procedure for solving a parameter estimation problem is essentially the same as that presented in Section 3. The principal difference is that in minimizing eq.(65) the gradient,

$$g(m) = J^T W_d^T W_d (F[m] - d^{obs}) \quad (67)$$

and the Hessian matrix $H = \nabla \nabla \phi_d$, no longer involve a term related to the model objective function. To make the gradient zero, a perturbation δm away from a model $m^{(n)}$ is found by solving.

$$H(m^{(n)}) \delta m = -g(m^{(n)}) \quad (68)$$

In eq.(68), H is either the full Hessian matrix, or some approximation to it, and this can yield Newton, quasi-Newton, Gauss-Newton, or steepest descent perturbations. Step sizes are limited by damping the perturbation after solving eq.(68) or by using a trust region approach which restricts the size of a perturbation so that the quadratic Taylor expansion is valid. The popular Marquardt-Levenberg approach is an example of this (Marquardt, 1963). Many algorithms also can incorporate additional bound constraints (e.g., Lawson and Hanson, 1974; Kelley, 1999). As an introduction to the plethora of algorithms that exist, the reader can visit the NEOS website (NEOS Guide, 2003).

As an example directly related to environmental geophysics, we look at the task of using time domain electromagnetic (TEM) data to identify a buried UXO. From an inversion perspective, this problem encompasses




Figure 39 TEM data for 4 time channels are given in the column on the left. Predicted data are shown in the center column and differences are on the right. The dotted lines show the actual locations at which data were recorded.

most of the usual elements encountered in parameter estimation problems. It also is a natural framework on which to pose a follow-up question regarding uncertainties in estimates and the need to invoke global optimization strategies.

Figure 39 shows the data at four time channels that were collected during a TEM survey over a known UXO target. The target was a 105 mm mortar that was hand-buried 44 cm beneath the surface. The survey instrument was a Geonics EM63 which has a $1m \times 1m$ source coil and the data are vertical components of the time derivative of the magnetic field measured in the center of the loop. The source and receiver are about 50 cm above the earth's surface. The goal of the survey is to locate the object and generate diagnostics to show that the object is likely a UXO and not a piece of scrap metal. Complete forward modelling would require solving Maxwell's equations but this is a very computationally demanding task. In such cases, some form of reduced modelling is sought.

If an object is approximately spherical, or if the distance between it and the observer is much larger than the size of the object, then the magnetic response can be modelled as that due to a dipole, and the TEM response modelled as that due to a decaying dipole. In the same limit, the TEM electromagnetic signature of a buried axi-symmetric body can be assumed to be represented by two decaying dipoles (one aligned with the axis of symmetry of the object, and the other perpendicular to the axis) and located at the center of the object. Each decay has the form

$$L(t) = k(t + \alpha)^{-\beta} e^{-t/\tau} \quad (69)$$

The data at any time are thus controlled by 13 parameters: the x, y, z locations of the object, two angles characterizing the orientation of the principal axis, and four parameters k, α, β, τ for each dipole. The reader is referred to Pasion and Oldenburg (2001) for details about the validity of this modelling.

The objective function in the inversion is given by eq.(65) and a Newton algorithm that allows upper and lower bounds on the parameters (Kelley, 1999) was used for the minimization. The predicted data, and difference from the observed, are given in Figure 39. The starting and recovered parameters are given in Table I and the final misfit was $\phi_d = 4514$. Because the ordnance was hand-placed, the location (x, y, z) and orientation ϕ, θ of the target are known. These parameter values are well recovered. The other parameters in the “known” column have been inferred from separate numerical experiments and we are not certain how correct they are. They are presented in brackets. Details about interpretation of parameters are given in Pasion and Oldenburg (2001), but the most important parameters are the values of β and k . The average of β is larger than 0.7 indicating that the object is ferrous. The ratio k_1/k_2 is greater than unity and this indicates that the object is rod-like. The ratio β_1/β_2 is less than unity, and this is also diagnostic of a rod. Hence, from the inversion, the object is estimated to be magnetic and rod-like which makes it a good candidate for being a UXO. It was, as we stated earlier, a 105 mm mortar.

A number of conclusions can be drawn from this example. Firstly, the two-dipole model seems to be a relatively good mathematical model for the UXO item. Secondly, the inversion algorithm has worked well because: (1) the data have been adequately fit, (2) good estimates have been obtained for the five known parameters, and (3) useful estimates of the other parameters have been recovered.

In the above Newton procedure for solving the parameter estimation problem, iteration is continued until the zero gradient condition is effectively attained. Unfortunately, the inter-connection between the parameters in the problem and the data is generally highly nonlinear and the misfit functional is often very rough and complicated with many local minima. The success of finding a desired global minimum therefore depends heavily upon the starting model. To illustrate this we rerun the example but begin with a very poor starting model. The result is given in Table II. The final misfit is $\phi_d = 7009$. The horizontal location of the target has been reasonably well recovered but the depth is incorrect and object is interpreted to be non-magnetic and plate-like.

The uncertainty of knowing whether we have found a sufficiently good minimum motivates the application of a global search algorithm. Here we use the Neighborhood Algorithm (NA) (Sambridge, 1999a) to search for a global minimum. The result is listed in Table 1. The minimum misfit achieved is $\phi_d = 4692$, which is slightly larger than that obtained with the Newton algorithm, and the parameter values obtained are similar to those from the linearized minimization when the first starting model was used.

As summary comments regarding parametric inversion, we have shown that a parameter set corresponding to a best fitting model can be found through linearized approaches. Because of the possible dependence upon the starting model, (that is, getting trapped in a local minimum) it is also advisable to begin the inversion from a number of different starting models and assess the results. If the number of parameters is not large then search algorithms can be used to find a global minimum. We note also, that the use of gradient methods will outperform global search methods if we are sufficiently close to the solution. This means that some difficult practical problems could be attacked with a two-stage methodology. A global search can be used to provide a starting model for a second phase Newton-type procedure that quickly gets to the local minimum.

As in the general inverse problem, however, it must be remembered that the recovered model is only one possible solution, and it depends upon details of the objective function, ϕ_d , that is being minimized. Changing the objective function in any way, even changing the estimated standard deviations of the data, will provide a different

answer. In practice, we don't know the true statistics of the errors, and estimating these is a continual challenge in every data set. This is often the most significant shortcoming in parametric inversions. Even if we have an estimate of the statistics of the noise, we should not become over-zealous about finding the model that makes this absolute minimum. The true model may have a realization of noise that does not correspond to this value. These comments motivate the idea that we need to obtain not only a value of the model parameters but also some assessment of their uncertainties. We address this issue in the next section.

Section 7.0: ASSESSING THE UNCERTAINTY IN INVERTED MODELS

For the underdetermined inverse problem in which we tried to find a function, and for the overdetermined problem where we attempted to find values of a few parameters, we found a single solution by minimizing a

Table 1 The known parameter values are given in column 2. Values in brackets are best estimates obtained through other analyses. Columns 3–5 pertain to a Newton-type solution for minimization. The uncertainties are linearized estimates. Column 6 is the global minimum produced by the NA algorithm. The means and standard deviations of the parameters estimated from marginal distributions using NA-Bayes are shown in the last two columns.

Parameter	Known values	starting parameters (Newton)	recovered parameters (Newton)	Standard deviations (Newton)	Global Search (NA-Bayes)	Mean (NA-Bayes)	standard deviations NA-Bayes
Northing x (meters)	1.83	2.0	1.78	0.001	1.78	1.78	.004
Easting y (meters)	2.00	2.0	2.05	0.001	2.04	2.04	.02
Depth z (meters)	0.44	0.52	0.45	0.001	0.47	0.47	.01
ϕ (degrees)	0	-45	9.71	0.29	9.4	9.5	1.5
θ (degrees)	90	45	85.0	0.19	84.2	84.3	0.8
k_1	(69.6)	88.9	74.2	0.75	76.7	76.7	1.5
beta_1	(0.64)	0.8	0.71	0.016	0.72	.72	.02
gamma_1	(20.4)	4.7	22.7	2.6	26.7	26.6	0.6
k_2	(20.1)	22.2	27.6	0.56	29.4	29.1	1.0
beta_2	(1.08)	1.15	1.06	0.024	1.08	1.08	0.01
gamma_2	(7.59)	3.13	5.73	0.63	6.74	6.8	1.0
beta_av			0.89			0.90	.05
k1/k2			2.69			2.64	0.2
beta/beta2			0.67			0.67	0.1

particular objective function. An immediate question arises as to how well features of the model, or values of the model parameters, represent the true solution. Effectively, the question is: “what uncertainty can be ascribed to the computed result?” This is a critical question and strategies used to answer this are problem dependent. In this section we briefly look at ways to gain quantitative insight about uncertainty. We make no pretenses about trying to cover this subject in depth because this is currently the subject of intensive research.

Different approaches are required for different types of problems and we thus divide this section into two parts and look respectively at the parameter estimation problem and at the under-determined inverse problem.

7.1: Uncertainty in parametric inversion

In parametric problems there are two approaches that can be used to estimate uncertainty and these are linked to how the problem is solved. If the solution was obtained from a Newton-type algorithm then a local analysis of the curvature of the objective function can be used to estimate the covariance of the parameters (e.g. Menke,1989).

Table 2 Illustrates the effect of a poor choice of starting model with a Newton-type algorithm. The first four columns are the same as in Table I. Note how the linearized standard deviations from second solution do not encompass the known values or those from the first inversion.

Parameter	Known values	Starting parameters	Recovered parameters	Standard deviations	Starting parameters	Recovered parameters	Standard deviations
Northing x (meters)	1.83	2.0	1.78	0.001	3.0	1.79	.001
Easting y (meters)	2.00	2.0	2.05	0.001	3.0	2.09	.001
Depth z (meters)	0.44	0.52	0.45	0.001	1.57	0.31	.003
ϕ (degrees)	0	-45	9.71	0.29	0	-101.7	0.58
θ (degrees)	90	45	85.0	0.19	0	14.5	0.19
k_1	(69.6)	88.9	74.2	0.75	1	11.1	0.87
beta_1	(0.64)	0.8	0.71	0.016	1	0.87	0.096
gamma_1	(20.4)	4.7	22.7	2.6	1	1.72	0.027
k_2	(20.1)	22.2	29.227.6	0.56	1	27.28	0.33
beta_2	(1.08)	1.15	1.06	0.024	1	0..98	0.013
gamma_2	(7.59)	3.13	5.73	0.63	1	75.44	25.80
beta_av			0.89			0.92	
k1/k2			2.69			0.41	
beta/beta2			0.67			0.89	

Let C be the covariance matrix for the data. To keep a consistent notation, let $C^{-1} = W_d^T W_d$. The least squares misfit objective function to be minimized is

$$\begin{aligned}\phi_d &= (d - F[m])^T C^{-1} (d - F[m]) \\ &\equiv ||W_d(d - F[m])||^2\end{aligned}\tag{70}$$

For a linear problem where $F[m] = Gm$ the covariance of the model estimates are given by

$$Cov[m] = (G^T W_d^T W_d G)^{-1}\tag{71}$$

For linear problems this information is globally correct. For non-linear problems, however, a localized linear analysis must be undertaken. If m_c is our computed model and δm is a perturbation, then we can write

$$F[m_c + \delta m] = F[m_c] + J[m_c]\delta m\tag{72}$$

where J is the sensitivity matrix evaluated at the computed model. The covariance of the model estimates becomes

$$Cov[m] = (J^T W_d^T W_d J)^{-1},\tag{73}$$

which is the inverse Hessian matrix. The square root of the diagonal elements of the covariance matrix defines the estimated standard deviation of computed parameters. If the true objective function is similar to a quadratic surface centered at our solution, then the above estimate of uncertainty should be useful. However, if the true objective function is substantially different, or worse, if we are at one of the nondesirable multiple minima, the above covariance estimates will not reflect the true uncertainty. For these reasons, eq.(73) usually underestimates the uncertainties. To illustrate the effects we return to the UXO problem. Table II shows the estimated standard deviations for the “good” and “poor” models. Standard deviations were obtained by taking the square roots of the diagonal elements of the covariance matrix. For the “good” model, the true values are within a few standard deviations from the calculated values. For the poor model, the true model values are many standard deviations away. The conclusion from this is that if we have been fortunate enough to construct a model that is close to the global minimum then a local uncertainty analysis may provide meaningful results. Alternatively, if we are far from the global solution then our covariance values do not reflect the true uncertainties.

The second approach to solving the parameter estimation problem was to use a global search algorithm. An important aspect of search algorithms is that they produce an ensemble of models (e.g., Lomax and Snieder, 1995) that can be analysed to generate estimates for the most likely values of the parameters and their uncertainty. The proper analytic framework lies in the statistical approach, Bayesian Inference, to inverse problems (Tarantola and Valette, 1982; Tarantola, 1987; Scales and Smith, 1996; Sambridge and Mosegaard, 2002 and references therein).

In a Bayesian formulation of the inverse problem we require data-independent prior information about the model parameters expressed through a probability density function $\rho(m)$ and also a likelihood function $L(m)$ which measures the likelihood of a model m through its misfit $\phi_d(m)$. Assuming Gaussian statistics, the likelihood function can be expressed as a probability density of the form

$$L(m) = ce^{-\phi_d(m)},\tag{74}$$

where c is a constant. The two probability density functions are combined through Bayes’ theorem to generate a posteriori probability density function $\sigma(m)$ over the model space $\mathcal{M}(m)$

$$\sigma(m) = L(m)\rho(m)\tag{75}$$

Any average of a function $f(m)$ can be obtained by

$$E_f = \int_{\mathcal{M}} \sigma(m) f(m) dm \quad (76)$$

This allows us to compute mean values, variances, and covariances of the parameters. For instance if $f(m) = m_i$ then the integration would yield the mean of that parameter. Unfortunately it is not generally possible to carry out the integrations in eq.(76) but approximations to the integral can be obtained if there are a large number of models sampled according to $\sigma(m)$. This is what the NA algorithm (Sambridge, 1999b) attempts to do and Monte Carlo integration techniques are used to evaluate the integrals. Of particular importance is the marginal probability density distribution, which is the probability density function for a particular parameter, irrespective of the values of the other parameters. For instance the marginal pdf for m_1 is the evaluation of eq.(76) over dimensions 2 to M . Visualizing these marginal pdf's reveals whether they are single or multi-modal, and means and standard deviations can be calculated directly.

We return again to the UXO problem. An ensemble of 22,020 models was used to generate marginal density distributions for individual parameters. The results were obtained from the computer programs supplied by Sambridge. Figure 40 displays the resultant marginal pdf's for different UXO parameters. A mean value and standard deviation of each parameter can be estimated from that distribution and those are listed in Table 1. We note in particular that the marginal pdf's show that β is likely greater than 0.7, which indicates that the object is magnetic; and, using the means of the distributions, $k_1/k_2 > 1$ and $\beta_1/\beta_2 < 1$. These indicate that the object is rod-like.

The above analysis can be made more quantitative. $f(m)$ in eq.(76) can be a function of multiple variables. In Figure 41 we show the marginal distributions for the average value of β and for ratios β_1/β_2 and k_1/k_2 . We arrive at an assessment that the object is magnetic and rod-like.

7.2: Assessing uncertainty in under-determined problems

Estimating uncertainties in inverse problems, where the goal is to recover a function, is more difficult than in parameter estimation problems. A function requires infinitely many independent data to determine it completely. With only a finite number of data, the nonuniqueness is such that the value of a model at any particular point is undetermined. This means that, as we decrease the cell size in our discrete solution to the inverse problem, the uncertainty of the value in each cell increases. It is therefore a futile effort to assign an uncertainty to each cell in a constructed model without providing some additional qualification. We can only obtain uncertainty about spatial averages of the model. Also, our uncertainty analysis should include our regularization functional and any prior information about the model. In the following we provide a few ways to understand and approach the problem of uncertainty estimation in under-determined problems. This area of inverse theory is sometimes regarded as “appraisal” and is complementary to the process of model construction where we attempted to find a single solution to interpret. We begin with ideas by Backus and Gilbert (1968, 1970) to show what information can be obtained from the data alone, and an extension of that to calculate bounds on spatial averages of the model. We briefly treat linearized methods to estimating localized uncertainties about a constructed model. The section ends with some practical approaches where insight about resolution and structure is obtained by carrying out additional inversions.




Figure 40 Marginal pdf's for some selected parameters in the UXO discrimination problem.




Figure 41 The marginal pdf for the average value of β is shown at the left. The observation that the mean is considerably greater than 0.7 indicates that object is magnetic. Histograms for ratios β_1/β_2 and k_1/k_2 are respectively given in middle and right panels. These indicate that the object is rod-like.

7.2.1: Averaging functions

To begin, consider the linear inverse problem of recovering the model $m(x)$ from a set of exact data

$$d_j = \int_0^1 g_j(x)m(x)dx \quad j = 1, N \quad (77)$$

We have shown earlier in this chapter that there are functions $m^*(x)$ such that

$$0 = \int_0^1 g_j(x) m^*(x) dx \quad j = 1, N \quad (78)$$

Clearly we can add as much of $m^*(x)$ to any solution $m(x)$ and still fit the data. It thus follows that the uncertainty about the value of $m(x)$ at any point is infinite. This was recognized by Backus and Gilbert (1968, 1970) and they abandoned the goal of assigning errors to points of a function and instead calculated errors for spatial averages of the model centered around a point of interest x_0 . For our linear problem the averages have the form

$$\langle m(x_0) \rangle = \int_0^1 m(x) A(x, x_0) dx \quad (79)$$

where $\langle m(x_0) \rangle$ is the average value, $A(x, x_0)$ is the averaging function. It has the form

$$A(x, x_0) = \sum_{j=1}^N \alpha_j g_j(x) \quad (80)$$

where α_j are coefficients that depend upon x_0 . The coefficients are computed so that $A(x, x_0)$ is designed to be a localized function (ideally a delta function) centered on x_0 . Importantly, the average value is unique. All models that fit the data have this same average. Substitution of eq.(80) into eq.(79) yields

$$\langle m(x_0) \rangle = \sum_{j=1}^N \alpha_j d_j \quad (81)$$

and therefore the average is a linear combination of the data. The averaging function quantifies the resolving power of the experiment and it is extremely important in understanding what level of detail about the model can be obtained purely from the data.

When the data are contaminated with noise, there is uncertainty in the average value. For most geophysical experiments even small amounts of noise yield average values with unacceptably large uncertainty. What is required is to tradeoff resolution for accuracy and a formalism for carrying that out is provided. Readers are referred to Parker (1977) and Oldenburg (1984) for concise summaries on the methodology. In the end, the interpreter has three quantities by which to assess what is known about the model near a point x_0 . They are: the average value $\langle m(x_0) \rangle$, a measure of its uncertainty, and the averaging function $A(x, x_0)$. The technique is applicable to nonlinear problems but the analysis looks only at models that are linearly close to the constructed model. This suffers from the same difficulty encountered earlier in the UXO problem. If the constructed model is close to the true model then a linearized appraisal analysis produces good results, but if we are far away, then the analysis is marginalized.

Some attempts have been made to extend the ideas of Backus and Gilbert. For example, Oldenburg (1983) used linear programming methods to compute upper and lower bounds for averages of the model

$$\bar{m} = \frac{1}{x_2 - x_1} \int_{x_1}^{x_2} m(x) dx \quad (82)$$

subject to the model adequately fitting the data. The motivation for this is practical. In some studies, we may not be too concerned about detailing the variation of the model in a particular depth range, but it might be of importance to know that the average value of a physical property over some depth range is bounded between two

values. For example, knowing that the average velocity between 250 and 275 meters is confined to the range [1250, 1300] m/sec might be sufficient. In an analogous manner to Backus-Gilbert appraisal, there is a tradeoff between uncertainty and resolution. The uncertainty in the average value decreases as the integration width increases and the interpreter is free to use a width and computed bounds which are useful. The above optimization problem can be solved by using only data as constraints or it can be extended to incorporate some types of prior information about the physical property. For example, positivity, or bounds on the value of particular cells are straightforwardly incorporated. The technique can be applied to nonlinear problems but this requires that a global minimum of an objective function is found. However, in the 1D magnetotelluric problem, Dosso and Oldenburg (1991) compared the results of an iterative linearized algorithm with those from simulated annealing, whose goal was to find a global minimum, and found that the bounds on the averages were quite comparable even though the models from which the bounds were computed were somewhat different. More extensive work on obtaining bounds on averages of the model is provided in Stark (1992).

7.2.2: Model resolution and covariance matrices

The other approaches to assessing nonuniqueness incorporate the regularization functional and other prior knowledge about the model. The analysis has been presented by numerous authors (e.g., Menke, 1989; Tarantola, 1987). For the linear problem where we minimized

$$\phi(m) = ||W_d(Gm - d^{obs})||^2 + \beta ||W_m m||^2 \quad (83)$$

the solution is

$$m = (G^T W_d^T W_d G + \beta W_m^T W_m)^{-1} G^T W_d^T W_d d^{obs} \quad (84)$$

$$m = H^{-1} G^T W_d^T W_d d^{obs} \quad (85)$$

where

$$H = G^T W_d^T W_d G + \beta W_m^T W_m \quad (86)$$

is the Hessian matrix.

If we write $d^{obs} = Gm_t + n$, where m_t is the true model, and n is the additive error, then the solution in eq.(84) can be written as

$$m = Rm_t + \Delta m, \quad (87)$$

where

$$R = H^{-1} G^T W_d^T W_d G \quad (88)$$

is called the resolution matrix. Δm is a bias term and it is the model obtained when inverting a data set consisting of pure noise. Neglecting the effects of noise for the moment, eq.(87) shows that the $i'th$ component of our solution is $m_i = r_i^T m_t$ where r_i is the $i'th$ row of R . Thus, m_i can be interpreted as a weighted average of the true model. If r_i is zero except for a value of unity in the $i'th$ cell, then the value of the model for that cell has been exactly recovered. In general this will not be the case, and usually the row function exhibits a peak at the $i'th$ cell and is smeared out away from there. Examples of these averaging functions are shown in Figure 42.




Figure 42 Averaging and point spread functions. The first column shows the averaging functions (corresponding to the rows of the resolution matrix) at locations $x=(0.05, 0.3, 0.5)$ respectively. The second column shows the point spread functions (corresponding to the columns of the resolution matrix) at the same three locations.

The rows of matrix R are thus quite akin to the Backus-Gilbert averaging functions and similarly, the average value has uncertainty associated with it because of data error. The covariances of the recovered values in eq.(84) are given by $\mathcal{E}(\Delta m \Delta m^T)$ where $\mathcal{E}$ is the expected value. Thus

$$Cov[m] = H^{-1} G^T W_d^T W_d G H^{-1} \quad (89)$$

The diagonal elements of the covariance matrix yields the estimated standard deviations of the computed model elements. These standard deviations should be used together with the averaging functions obtained from the resolution matrix.

Another way to gain insight about the resolving power of the experiment is to see how an impulse in the model appears in the inverted model. This is easily done by making m_i equal to zero everywhere except on the i^{th} cell where it can have a unit value. The inversion result $m = R m_i$ will likely be a smoothed feature approximately centered around the position of the impulse but will have structure away from that point. The extent to which the smoothing has happened provides a qualitative measure of the resolving power of the inversion and these functions are referred to as point spread functions (PSF). Clearly the point spread functions correspond to the columns of the resolution matrix R in eq.(87). Three of these have been plotted in Figure 42. We notice that the width of the PSF increases with the depth at which the impulse was applied. This indicates that our resolving power is worsening

with depth. A crude measure for resolving power can be assigned from examining the width of the PSF. We would not expect to see structure that has a scale-length less than this width. We notice, in this particular example, that the averaging functions (rows of R) and the point spread functions (columns of R) are very similar. This is not usually the case since the matrix R is not symmetric in general. It becomes symmetric when the regularization term, controlled by β , approaches zero.

For nonlinear problems, the analysis in this section follows through except that the matrix G in eqs.(87) is replaced by the sensitivity matrix J evaluated at the final model obtained in the construction process. The covariances and associated averaging functions pertain to models that are linearly close to the constructed model. Many authors have used this approach to appraise their inversion results. Alumbaugh and Newman (2000) provide an example connected with cross well tomograph using electromagnetic data and that paper provides references to the extensive use of this approach. In nonlinear problems the point spread functions refer to how impulse perturbations of the constructed model are manifested in the constructed model. An example of a nonlinear problem and further discussion is given in Parker (1994).

Finally, a potentially powerful mathematical basis for treating uncertainty in underdetermined problems lies in the framework of Bayesian analysis. We have presented the essence of this earlier within the context of the overdetermined problem, and the formalism is equally applicable to our problem here provided that the problem has been discretized. Calculation of the posterior probability density function requires knowledge of the likelihood function and the prior model density function. In principle this posterior pdf contains all of the information needed to assess uncertainty, but integration of the multi-dimensional integrals is computationally prohibitive. The usual procedure is to find a single model that maximizes the posterior density function and carry out a local analysis associated with that model. Readers are referred to Menke (1989) and Tarantola (1987) for more details.

7.2.3: Limited exploration of model space

The ideas presented in the previous sections have attempted to quantify the uncertainty in average values of the model or look at how impulse perturbations would manifest themselves in the constructed model. These approaches are useful, however, further insight can be obtained by carrying out other inversions that are designed to answer specific questions. For instance, examination of our constructed model undoubtedly raises questions regarding existence of particular features and it is important to know whether those features are demanded by the data (and thus are resolved) or if they are a consequence of our regularization. Such questions may be answered by carrying out a limited exploration of model space to generate other models that differ from our preferred model, but are still geologically plausible, and in accordance with our prior information.

One way to generate such models is to alter the model objective function ϕ_m by changing the coefficients, weighting functions, or the reference model. As an example we return to the inversion of DC resistivity data presented in Section 3. In Figure 43 we show five models obtained by carrying out the inversion using different reference models. The reference models correspond to different halfspaces or have linear gradients in resistivity. Viewing these models provides insight about which features might be real. In particular, the fact that all images show the four buried prisms provides some confidence that they exist. We can't however, say this with certainty because our search of model space has been limited, and there may be undiscovered geologically plausible models that do not have one of the prisms. Nevertheless, the results increase our confidence that the earth might have these features.

The models in Figure 43 are reasonably similar in the near-surface regions even though the reference model differs. At depth however, all models revert to the reference. Structure at depth is therefore controlled by prior




Figure 43 Five resistivity models obtained by using different reference models. The reference model in panels (a)-(c) are halfspaces with different resistivity. Panels (d)-(e) have resistivities that progressively increase and decrease with depth.

information built into the objective function, while the structure at the surface seems to be needed to satisfy the data. In comparing the models, there is a corridor above which the models are reasonably similar, and below which, the models are controlled by the reference model. This corridor defines a “depth of investigation” (DOI) for the survey and can be used to truncate parts of the recovered model that are controlled only by the model objective function. In Figure 44 we show the DOI index from Oldenburg and Li (1999) for the above models and use that to truncate (or successively dim) the image. Truncating the image using this DOI index, or some equivalent technique, is good practise when showing models to people who are not familiar with the details of inversion. Experience has taught us that it is not possible to tell someone to ignore parts of an image once it has been shown.

The images in Figures 43 and 22 can be used to discuss other generic questions that arise when inversions are carried out. Consider the two buried conductors at $x = -60$ and $x = 25\text{m}$ in the image in Figure 22. A question of practical interest might be whether these bodies really have limited depth extent. A first answer is obtained from the DOI investigation. In Figure 44(b) the DOI values are superposed upon the model obtained from the inversion. The body near $x = -60$ appears well within the region where the data have significant influence and we conjecture that it is likely closed. However, the body at $x = 25\text{m}$ has substantial conductivity at depths where the DOI index is large, and hence we are not seeing its bottom with the survey data provided. It may extend to greater depths than suggested by Figure 44(b). Supporting evidence for that arises from the model in Figure 43(e). That model, generated with a reference model whose conductivity increased with depth, shows that the conductor extends to depths of a 100 meters where it meets with a layer of higher conductivity. If this model is geologically plausible then we must conclude that the second conductor could extend to depth and thus we don’t know whether it is open or closed.




Figure 44 The DOI index is shown in (c). Values close to unity indicate locations where the recovered model is insensitive to the value of the supplied reference model, while values approaching zero indicate locations where the inversion result is controlled entirely by the reference model. This information is superposed upon the constructed model in (b). The contour interval is 0.1. In (a) the intensity of the recovered model is scaled by the DOI value. White regions correspond to areas where the resistivity has minimal effect on the data.

Finding a model that displays, or doesn't display, a specific feature is a powerful way to assess existence of structures. In effect, this is hypothesis testing, and the thought process goes as follows. A best inversion is carried out and particular features of the image are of interest. For the example above, it might be existence of the conductive block at $x = -60m$. We suspect that this feature is real because it exists in the image even though we have purposely tried to generate a model that is as smooth as possible. Also, we have observed it on all of the inversion images. However, we can focus more directly on the question of existence of this body by trying to construct a geologically reasonable counter-example that doesn't have the feature. To do this, let the reference model be the uniform halfspace of 400 ohm-m and introduce a weighting into the first component of the model objective function so that it becomes

$$\iint w(x, z)(m(x, z) - m_{ref})^2 dx dz \quad (90)$$

where the weighting $w(x, z)$ has large values in a region around the resistive region shown in Figure 45(d) and is unity away from there. The resulting weight function is shown in Figure 45(a). Setting the weighting to 100 and

carrying out the inversion produces the model in Figure 45(b). The inversion is able to find a model that doesn't have anomalous conductivity over the rectangular region, however, some extra conductivity at the surface, and also at depth, is required to reproduce the data. If this model is deemed to be geologically reasonable then a counter example has been found and one cannot conclude that the initial body must exist. Figure 45(c) shows the result obtained after setting the weighting to 10.




Figure 45 Resistivity models generated after penalizing conductivity variation in a region around the conductive block at $x = -75m$. The initial inversion produced the model in (d). Panel (a) shows the location of the non-unity portions of the weighting function. Panel (b) is the inversion result with a weighting value of 100. Panel (c) is the inversion result with a weighting value of 10.

The generic scenario for hypothesis testing is as follows. A principal inversion is carried out to produce an image to be interpreted. Questions about existence or nonexistence of structure, or questions regarding location of an observed structure, can be formulated. This often occurs when a geologist, who is familiar with the project area, is viewing the inversion results. Hypotheses are formed and the appropriate objective function is then designed to test the hypothesis. Enhanced, or reduced, confidence in the existence, nonexistence, location etc. is inferred from the outcome of the inversion. It must be remembered that model space is infinitely large, and therefore we can never sample it sufficiently to determine that a particular feature is observed on all models that adequately fit the

data. However, a single geologically plausible counter example is powerful information to show that the feature doesn't exist. If focused attempts to construct the counter example fail, then that provides increased confidence that the feature exists.

In summary, the assessment of uncertainties in the constructed model can be attacked through various procedures ranging from Backus-Gilbert appraisal, linearized analysis, full non-linear inference, or model construction. The variety of approaches indicates that this is still a fertile field for research and that the type of appraisal information needed is problem dependent.

Section 8.0: COMMENTS ON PRACTICAL IMPLEMENTATION OF INVERSION PROCEDURES

This completes our tutorial on inversion. However, since this chapter is included in a practical book for solving environmental, engineering and exploration problems, we feel that it is worthwhile to incorporate the work presented thus far, into a larger picture. There are two main aspects upon which we would like to comment. These pertain to issues of cost effectiveness of geophysical inversion and some components, particularly from the man-power side, that are crucial to using inversion techniques successfully.

8.1: Cost effective use of geophysical inversion

This chapter has illustrated that important information about the subsurface can be obtained through mathematical inverse procedures. In practise however the usefulness of inversion depends upon the cost (computer, manpower, time) of inverting the data, compared to the dollar impact that the information makes on the project. This requires an evaluation of the cost effectiveness of carrying out the inversion. We do not advocate that all data be inverted. Rather, the decision to employ inversion, or different levels of sophistication in the inversion, depends upon the question that needs to be addressed in the particular environmental, engineering or exploration problem. When solving a specific problem there may be several stages at which questions are posed and for which geophysics can be of benefit. The flow chart in Figure 46 outlines some scenarios that might arise within the scope of a major project. Most individual problems will make use of one or more of these decision paths.

8.1.1: Stage I: Reconnaissance data and map interpretation

In initial stages of problem solving, ground or airborne geophysical data are acquired with the goals of obtaining geologic information and for identifying "targets". Standard processing can be applied, and the data are generally presented in map form. Good data maps are essential in any interpretation. On the large scale they provide information about the geology, such as predominant strike direction, faults, lineaments etc. This is information that is not easily obtained from inversions. In fact, strike direction, or locations of contacts, can be considered to be prior information for some inversion processes. The map data may also provide indication of targets. For example a magnetic survey done to find an underground storage tank, pipe, a UXO or mineral deposit, might produce an anomaly map that identifies the locations of these objects. Is inversion necessary? The answer is problem dependent and is guided by prior knowledge, risk, and cost. In the case of the storage tank which is likely close to the surface and reasonably well located in the horizontal directions, the digging might proceed straight away. The same is




Figure 46 Flow chart that puts geophysical inversion in the context of solving practical problems.

true for the buried pipe, although to reduce risk of damaging the pipe, an estimate of depth of burial would be helpful. This could be provided through simple rules of thumb applied to cross-sectional profiles, or obtained by solving a parameter estimation problem. However, when the objects are at greater depth, or if information about their location is not directly evident from the data maps, then it will be cost effective to invert the data. This is usually the case for mineral exploration.

8.1.2: Stage II: A first pass inversion

There is an immense difference in the amount of processing required to plot the data in some map form compared to inverting the data. There are also importance differences in the data quality that is needed in pursuing each of these paths. In providing map data, where the goal is to find “bumps”, various massagings of the data still yield interpretable results. For instance, the eye can detect “bumps”, or other features, even though the data are scaled, are biased by an additive constant or smooth trend, or have had ad hoc filtering applied to reduce specific forms of observed noise. Only qualitative information can be extracted from such approximate data. On the other hand, inversion procedures are generally unforgiving about ad hoc processing. In quantitative analysis, it is essential that we understand precisely what the data are. At the heart of the inversion procedure is a mathematical model of the process which relates the physical property distribution to the measured data. Survey parameters (locations and orientations of sources and receivers), units and normalizations of the observations, and information about topography must be known as accurately as possible. Even a single missing piece of information, like the normalization of data, can prevent an inversion from proceeding. The first cost in carrying out an inversion can be the cost of acquiring these relevant details. This is all too frequently neglected by some contractors.

The next costs pertain to manpower and computing. Our general procedure involved minimizing $\phi = \phi_d + \beta \phi_m$. The details of the recovered model depends upon the explicit form of ϕ_d and ϕ_m and the selection of the regularization parameter β . Specifying ϕ_d requires that the user estimate proper relative errors for the data. This requires some analysis and understanding of the field survey. Specifying ϕ_m also requires thought and effort

because there are many aspects of prior information that can be included. Some experimentation may be required to determine which objective function is satisfactory for the particular problem. This can be done with synthetic examples where data from a “typical” model are simulated and inverted. We note that there are thus manpower and computing costs in preparing for the first inversion.

The inversion can be performed once the data, their errors, and a model objective, have been decided upon. In many cases a few inversions are needed to explore the effects of different errors assignments or different regularization parameters. Also, with data sets like gravity or magnetics, it may be desirable to invert data with a different regional removed, since this is a pre-inversion step in which it is always uncertain as to whether an optimum job was done. Finally, having arrived at a “best” model for interpretation, it will be informative to carry out a few additional inversions by varying the model objective function to get a limited exploration of model space and to get some sense about the resolving capabilities of the data, or to see if certain features are likely artifacts of the inversion. The conclusion to be drawn is that any individual data set will likely be inverted a number of times. This requires that an experienced person, who understands the geologic problem, the geophysical survey, and the fundamentals of inversion, be in charge. Nevertheless, the procedure can be cost effective. An example for mineral exploration is presented in Phillips *et al.* (2001)

8.1.3: Stage III: Detailed investigations

Large scale investigations such as detailed mineral exploration or mine development, major contaminant site cleanup and site investigation for waste material, require that inversion be an integral part of the analysis. The highest resolution images possible are demanded, and information from different geophysical surveys, geology, and boreholes must be incorporated into a single interpretation. This requires a dedicated team, ideally including geophysicists, geologists, engineers and others from relevant disciplines. The currency of exchange is images, and the relationship between physical properties and buried objects or geologic units is essential. This demands that extensive petrophysical information be available. An important aspect of the interpretation is the compatibility of the generated image with prior information. If the image appears to contradict previously existing ideas about structure or location of an anomalous body, then a modified objective function can be generated and the inversion rerun to test a hypothesis. Visualization capabilities, and software that can manipulate different 3D images, keep track of drillhole information, store models on different grids, allow the user to build reference models, and generate correlations between models on different grids, are needed. These projects are generally long term and the inversion processing is iterative, interactive and essential.

8.2: Keys to successful geophysical inversion

Success can be judged only if clear goals for geophysical inversion have been previously established. It is essential to know whether the objective is to find a possible target or whether it was to determine details about its structure. Irrespective of the goal, a successful inversion of geophysical data requires two ingredients: (1) high quality survey design and data acquisition, (2) high quality inversion of the data. The first item ensures that the data contain the “signal” needed to answer the question, even though the signal may be encoded in the observations in a complicated manner. The second item ensures that the information is extracted from the data in such a way as to answer the question. These two components are intimately linked because the amount of structure in the final model depends upon the specifics of the inversion algorithm, and also upon which data are acquired and upon their

accuracy. “Success” can be enhanced by adopting a more holistic approach in which the survey design makes use of the inversion algorithm. A possible scenario to strive for includes the following steps:

- Generate a synthetic earth model and assign representative physical property values to the various elements. This model should contain the topography if it is likely to have important effects on the data.
- Specify a trial geophysical survey and carry out a forward modelling. Estimate realistic noise levels for the data and add a realization of this noise to generate the simulated field data.
- Design a model objective function and inversion algorithm that includes prior information appropriate for the survey area.
- Carry out the inversion of synthetic data.
- Assess whether the inverse result can adequately answer your question. (e.g. does it show the target body?). If not, then either more, or different data need to be collected, or the data needs to be collected to a higher accuracy. This last item is important. The resolving power of the data increases as the signal to noise level increases. Structures that were not visible at one noise level may be visible at a reduced noise level.

Completing the above steps will help ensure a successful survey and subsequent inversion. The list also underscores the connections between data collection, the inversion algorithm, and the ability to answer a final question.

A key ingredient to the success of any inversion is the personnel involved in the processing. It is relatively easy to obtain an image from any inversion algorithm, but obtaining a best interpretable image requires an experienced professional. This comment may be obvious, but it is so crucial to the long term growth of responsible use of geophysics in solving applied problems, that we comment upon it here.

There is a major difference between forward modelling and inversion that may not be sufficiently appreciated by those not familiar with inversion procedures. Forward modelling is a well-defined procedure, and numerical algorithms can be run by persons having minimal experience. There might be some rules with respect to designing finite difference or finite element meshes, or adjustment of tolerance parameters for monitoring how well a system of equations is solved, but the process is generally “turn-key”. No extensive training is required. This contrasts markedly with the inverse problem, whose solution cannot reliably be obtained by implementing a computing code in a “black box” manner. To achieve success, an inversionist must have skills and understanding about the data (both the field data and synthetically modelled data), about the prior information, and about the inversion process.

8.2.1: Understanding the field, and synthetically modelled, data:

The quantitative nature of inversion requires that the interpreter has a fundamental understanding about the geophysical data, how they were collected, and knowledge about any processing that has been carried out. At the data collection level, estimates regarding misorientations, or incorrect locations of sources and receivers, as well as topography, are needed. This information, when combined with data statistics from repeatability measurements, or discrepancies in field reciprocity checks, provide estimates about the uncertainties in the data that can be used when designing a misfit function ϕ_d . If the data are further processed prior to inversion, for instance, removal of specific types of noise, or removal of the effects of “regional” sources, then the interpreter must be familiar with these processes and understand how parameters in those algorithms affect the data. The quality of the inversion result is heavily dependent upon the quality of data supplied, and the adage, “garbage in — garbage out” is operative here.

The interpreter must also understand the forward modelling used to emulate the field survey. Discrepancies between the mathematical modelling and the true earth system enter as “noise” in the inverse problem and may

impact upon how ϕ_d is defined and how well the data should ultimately be fit. This may include neglecting the effects of EM coupling in IP surveys, assuming the physical properties are isotropic when they are anisotropic, the effects of small scale geologic noise on the data, (like a small highly conducting rock near an electrode), or the effects of dimensionality (working in 2D when the object is 3D). It is important to identify items that could realistically affect the discrepancy between the mathematical model of the earth and the true earth. Once identified, there is usually some way of forward modelling with a more complicated algorithm to estimate the magnitude of the approximations or neglected effects.

8.2.2: Understanding the prior information and the inversion process:

There should be a clear objective regarding what is expected from the inversion algorithm, and the interpreter should be connected with the survey design. Basic understanding about the geologic structure and expected physical property values is essential. The inversionist should also know what prior information is available and how to incorporate it into the inverse problem. This requires a good grasp of the fundamentals of inversion.

When an inversion is run, the result may not be satisfactory (by whatever criterion), or the program may have failed to converge. The process of tracking down the source of the problem must then begin. It could be a problem with the data, or with details of the inversion algorithm. If the inversion has been successful, then issues such as hypothesis testing, carrying out limited explorations of models space, and plotting of images arise. In all of these endeavors, experience and knowledge are required.

CONCLUSION

In this chapter we have concentrated on the fundamentals of inversion and have attempted, through synthetic and field examples, to illustrate the power of inversion. The process is nontrivial to implement and requires good algorithms and a good understanding of the inverse process on the part of the interpreter. We hope this chapter can contribute positively in both of these areas.

ACKNOWLEDGMENTS

We thank Dwain Butler for inviting us to write this paper and for his efforts in handling the manuscript. We also thank John Scales and an unknown referee for their comments and suggestions about the manuscript. We especially thank our colleagues at the Geophysical Inversion Facility. Roman Shekhtman, Eldad Haber, Jiuping Chen, Colin Farquharson, Len Pasion, Steven Billings and Francis Jones have supplied some of the figures presented here and have helped shape the material content of this chapter. We also thank Roman Shekhtman for generating the Java Applet used in this tutorial. We thank Malcolm Sambridge for making his Neighborhood Algorithms available.

REFERENCES

- Aki K. and Richards, P. G., 2002, *Quantitative seismology*, University Science Book.
- Alumbaugh, D., and Newman, G., 2000, Image appraisal for 2-D and 3-D electromagnetic inversion, *Geophysics*, **65**, 1455–1467.
- Backus, G. E., and Gilbert, J. F., 1968, The resolving power of gross earth data: *Geophys. J. Roy. Astr. Soc.*, **16**, 169–205.

- , 1970, Uniqueness in the inversion of inaccurate gross earth data, *Phil. Trans. R. Soc. Lond., Ser., A*, **266**, 123–192.
- Billings, S. D., Kennett, B.L.N., and Sambridge, M. S., 1994, Hypocenter location: Genetic algorithms incorporating problem specific information, *Geophys. J. R. Astron. Soc.*, **118**, 693–706.
- Billings, S. D., Pasion, L. R. & Oldenburg, D.W., 2002, UXO discrimination and identification using magnetometry, *Proceedings of SAGEEP*
- Christensen, N.B., 1996: Imaging of Transient Electromagnetic Soundings Using a Scaled Frequency Derivative. In: Jacobsen, B.H., Mosegaard, K., and Sibani, P., ed., *Inverse Methods. Interdisciplinary Elements of Methodology, Computation and Applications. Lecture Notes in Earth Sciences*, Springer-Verlag, Berlin-Heidelberg, 205-214.
- , 1997: Electromagnetic subsurface imaging - A case for an adaptive Born approximation. *Surveys in Geophysics* 18, 477-510,.
- Deng, L. H., and Scales, J. A., 1999, Estimating the topography of multi-dimensional fitness functions, *Cent. for Wave Phenomena Tech. Rep. 208*, Colorado School of Mines, Golden.
- Dennis, J. E., and Schnabel, R. B., 1996, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, SIAM Classics in Applied Mathematics.
- Dey, A. and Morrison, H. F., 1979. Resistivity modelling for arbitrarily shaped three-dimensional structures, *Geophysics*, **44**, 753–780.
- Dosso, S. E., and Oldenburg, D. W., 1991, Magnetotelluric appraisal using simulated annealing, *Geophys. J. Int.*, **106**, 379–385.
- Eklblom, H., 1973, Calculation of linear best L_p -approximations, *BIT*, **13**, 292–300.
- , 1987, The L_1 -estimate as limiting case of an L_p - or Huber-estimate, in *Statistical data analysis based on the L_1 -norm and related methods*, ed. Dodge, Y., Elsevier, Amsterdam..
- Ellis, R. G. and Oldenburg, D. W., 1994. The pole-pole 3-D DC-resistivity inverse problem: a conjugate-gradient approach: *Geophys. J. Int.*, **119**, 187–194.
- Farquharson C.G., and Oldenburg D.W., 1998, "Nonlinear inversion using general measures of data misfit and model structure", *Geophysical Journal International*, **134**, 213-227.
- , 2000, Automatic estimation of the trade-off parameter in nonlinear inverse problems using the GCV and L-curve criteria, SEG 70th Ann. Internat. Mtg, Soc. Expl. Geophys.
- Fink, J. B., Sternberg, B. K., McAlister, E. O., and Wieduwilt, W. G., Eds., 1990, *Induced polarization*, SEG Investigations in Geophysics No. 4, Soc. Expl. Geophys., Tulsa
- Gelman, A., Carlin, J. B., Stern, H. S. & Rubin, D. B., 1995. *Bayesian Data Analysis*, Chapman and Hall, London.
- Gill, P. E., Murray, W., and Wright, M. H., 1981, *Practical Optimization* Academic Press Inc.
- Gill, P. E., Murray, W., Ponceleon, D. B., and Saunders, M., 1991. Solving reduced KKT systems in barrier methods for linear and quadratic programming. Technical Report SOL 91–7, Stanford University.
- Goldberg, D. E., 1989, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, Mass.

- Golub, G. H., Heath, M., and Wahba, G., 1979, Generalized cross-validation as a method for choosing a good ridge parameter, *Technometrics*, **21**, 215–223.
- Golub, G. H., and Von Matt, U., 1997, Generalized cross-validation for large scale problems, *J. of Computational and Graphic Statistics*, **6**, 1–34.
- Haber, E., and Oldenburg, D.W., 2000. A GCV-based method for nonlinear ill-posed problems, *Computational Geosciences*, **4**, 41–63.
- Haber, E., and Tenorio, L., 2003, Learning regularization functionals-a supervised training approach. *Inverse Problems* **19**, No. 3., 611–626.
- Hansen, P. C., 1992, Analysis of discrete ill-posed problems by means of the L-curve, *SIAM Review*, **34**, 561, 580.
- , 1998, *Rank-deficient and discrete ill-posed problems: Numerical aspects of linear inversion*, SIAM
- Huber, P.J., 1964, Robust estimation of a location parameter, *Ann. math. Stat.*, **35**, 73–101.
- Kelley, C. T., 1999, *Iterative Methods for Optimization*, SIAM.
- Kirkpatrick, S. C., D. Gelatt, and M. P. Vecchi, 1983, Optimization by simulated annealing, *Science*, **220**, 671–680.
- Kowalczyk, P., 2000, Inversion of induced polarization data from Donlin Creek, Alaska, 70th Ann. Internat/ Mtg., Soc. of Expl. Geophys., 1067–1070.
- Lanczos, C., 1958, Linear systems in self-adjoint form, *Am. math. Monthly*, **65**, 665–679.
- Langer, R.E., 1933, An inverse problem in differential equations, *Am. Math. Soc Bull.*, ser,2 v29, 814-820
- Lawson, C. L. and Hanson, R. J., 1974. *Solving least-squares problems*, Prentice-Hall, Englewood Cliffs, NJ.
- Levy, S. and Fullagar, P. K., 1981, Reconstruction of a sparse spike train from a portion of its spectrum and application to high resolution deconvolution. *Geophysics*, **44**, 1235–1243.
- Li, Y. and Oldenburg, D. W., 1996. 3-D inversion of magnetic data, *Geophysics*, **61**, 394–408.
- , 1999, 3D inversion of DC resistivity data using an L-curve criterion, 69th Ann. Internat. Mtg., Soc.Expl. Geophys., Expanded Abstracts
- , 2000. Incorporating geologic dip information into geophysical inversion, *Geophysics*, **65**, 148–157.
- , 2003, Fast inversion of large scale magnetic data using wavelet transforms and a primal logarithmic barrier method, *Geophysical Journal International*, **152**, 251–265.
- Lines, L. R. and Treitel, S., 1984. A review of least-squares inversion and its application to geophysical problems: *Geophys. Prosp.*, Eur. Assn. Geosci. Eng., **32**, 159-186.
- Lomax, A., and Snieder, R., 1995, Identifying sets of acceptable solutions to nonlinear geophysical inverse problems which have complicated misfit functions, *Nonlinear Proc. Geophys.*, **2**, 222–227.
- Marquardt, D. W., 1963, An algorithm for least-squares estimation of non-linear parameters, *J. Siam*, **11**, 431–441.
- McGillivray, P.R. 1992, Forward modeling of DC resistivity and MMR Data, Ph.D. Dissertation, University of British Columbia, Vancouver.
- Menke, W., 1989. *Geophysical Data Analysis : Discrete Inverse Theory*, Academic Press
- Mutton, A. J., 1997, The application of geophysics during evaluation of the century zinc deposit, in Ed. Gubins, A. G., *Proceedings of the Exploration 97, Fourth Decennial International Conference on Mineral Exploration*, 599–614.

- NEOS Guide for optimization, 2003, <http://www-fp.mcs.anl.gov/otc/Guide>, last accessed June 25, 2003.
- Nocedal, J. and Wright, S. J., 1999, *Numerical Optimization*, Springer.
- Oldenburg, D. W., 1983, Funnel functions in linear and nonlinear appraisal, *J. Geophys. Res.*, **88**, 7387–7398.
- , 1984, An introduction to linear inverse theory, *IEEE Trans. Geosci. Remote Sensing*, **GE-22**, 665–674.
- Oldenburg, D. W., and Li, Y., 1994, Inversion of induced polarization data, *Geophysics*, **59**, 1327–1341.
- , 1999, Estimating depth of investigation in DCfind resistivity and IP surveys, *Geophysics*, **64**, 403–416.
- Oldenburg, D. W., Li, Y., and Ellis, R. G., 1997, Inversion of geophysical data over a copper-gold porphyry deposit: A case history for Mt. Milligan, *Geophysics*, **62**, 1419–1431
- Park, S. K. and Van, G. P., 1991. Inversion of pole-pole data for 3-D resistivity structure beneath arrays of electrodes: *Geophysics*, **56**, 951–960.
- Parker, R. L., 1977, Understanding inverse theory, *Ann. Rev. Earth Planet. Sci*, **5**, 35–64.
- 1994. *Geophysical Inverse theory*, Princeton Univ Press
- Pasion, L., and Oldenburg, D.W., 2001, A discrimination algorithm for UXO Using Time Domain Electromagnetic Induction, *Journal of Environmental and Engineering Geophysics*, **6**, 91–102.
- Phillips, N., Oldenburg, D. W., and Chen, J., Li, Y., and Routh, P., 2001, Cost effectiveness of geophysical inversions in mineral exploration: Applications at San Nicolas", *The Leading Edge*, **20**, 1351–1360.
- Sambridge, M., 1999a, Geophysical inversion with a neighbourhood algorithm, I: Searching a parameter space, *Geophys. J. Int.*, **138**, 479–494.
- , 1999b, Geophysical inversion with a neighbourhood algorithm, II: Appraising an ensemble, *Geophys. J. Int.*, **138**, 727–746..
- Sambridge, M., and Mosegaard, K., 2002, Monte Carlo methods in geophysical inverse problems, *Review of Geophysics*, **40**,
- Sasaki, Y., 1994. 3-D resistivity inversion using the finite-element method: *Geophysics*, **59**, 1839–1848.
- Scales, J., and Gersztenkorn, 1988, Robust methods in inverse theory, *Inverse problems*, **4**, 1071–1091.
- Scales, J., and Smith, M. L., 1996. *Introductory Geophysical Inverse Theory: Part I*, Samizdat Press.
- Scales, J. A., Smith, M. L. and Fischer, T. L. ,1992, Global optimization methods for multimodal inverse problems, *J. Comput. Phys.*, **103(2)**, 258–268.
- Seigel, H. O., 1959. Mathematical formulation and type curves for induced polarization: *Geophysics*, **24**, 547–565.
- Sharma, P. V., 1966, Rapid computation of magnetic anomalies and demagnetization effects caused by bodies of arbitrary shape: *Pure Appl. Geophys.*, **64**, 89–109.
- Stark, P.B., 1992, Inference in infinite-dimensional inverse problems: Discretization and duality, *J. Geophys. Res.* **97**, 14055–14082.
- Stein, S., and Wysession, M., 2003, *An introduction to seismology, earthquake, and earth structure*, Blackwell Publishing.
- Stevenson, A. F., 1934, On the theoretical determination of earth resistance from surface potential measurements, *Physics*, **V3** 114–124

- Sumner, J.S., 1976, Principles of induced polarization for geophysical exploration, Elsevier.
- Tarantola, A., 1987. *Inverse Problem Theory*, Elsevier, Amsterdam.
- Tarantola, A., and Valette, B., 1982, Inverse problems=quest for information, J. Geophys., 50, 159–170.
- Tikhonov, A. V. and Arsenin, V. Y., 1977. Solution of ill-posed problems, Ed. Fritz, J., John-Wiley & Sons, New York.
- Vogel, C. R., 2001, *Computational Methods for Inverse Problems*, Frontiers in Applied Mathematics, SIAM, , Philadelphia.
- Wahba, G., 1990. *Spline models for observational data*, SIAM, Philadelphia.
- Wait, J.,R. (editor), 1959, Overvoltage research and geophysical applications, Pergamon Press.
- Wang, B., 1995, Effective approaches to handling non-uniform data coverage problem for wide-aperture refraction/reflection profiling, 65th Ann. Internat. Mtg., Soc. Expl. Geophys., Expanded Abstracts, 659–662.
- Ward, S. H., and Hohmann, G.W., 1987, Electromagnetic theory for geophysical applications, in Nabighian, M. N., Ed., Electromagnetic methods in applied geophysics, **01**, Soc. of Expl. Geophys., 313–364.
- Weidelt, P., 1972, The inverse problem of Geomagnetic induction, Z. Geophys., **38**, 257–289.
- Wright, S. J., 1997. Primal-dual interior-point methods, SIAM, Philadelphia.