

How do we understand and visualize uncertainty?

MALCOLM SAMBRIDGE, Australian National University, Canberra

CAROLINE BEGHEIN, Arizona State University, Tempe, USA

FREDERIK J. SIMONS, University College, London, U.K.

ROEL SNIEDER, Colorado School of Mines, Golden, USA

Geophysicists are often concerned with reconstructing subsurface properties using observations collected at or near the surface. For example, in seismic migration, we attempt to reconstruct subsurface geometry from surface seismic recordings, and in potential field inversion, observations are used to map electrical conductivity or density variations in geologic layers. The procedure of inferring information from indirect observations is called an inverse problem by mathematicians, and such problems are common in many areas of the physical sciences. The inverse problem of inferring the subsurface using surface observations has a corresponding forward problem, which consists of determining the data that would be recorded for a given subsurface configuration. In the seismic case, forward modeling involves a method for calculating a synthetic seismogram, for gravity data it consists of a computer code to compute gravity fields from an assumed subsurface density model. Note that forward modeling often involves assumptions about the appropriate physical relationship between unknowns (at depth) and observations on the surface, and all attempts to solve the problem at hand are limited by the accuracy of those assumptions. In the broadest sense then, exploration geophysicists have been engaged in inversion since the dawn of the profession and indeed algorithms often applied in processing centers can all be viewed as procedures to invert geophysical data.

Uncertainty in the information obtained from inverse problems is a central issue in many areas of science. In exploration and production geophysics in particular, the uncertainty surrounding the information retrieved from measurements that are inherently limited, incomplete, or redundant is the driver in making sound business decisions (Scales and Snieder, 2000), and as such assumes a central role in many studies. In global geophysical problems such as whole-earth seismology, the stakes are usually not monetary, but the implications of basing scientific judgment on the ambiguous or nonunique results of an inverse problem are not less important.

Rather than addressing uncertainty explicitly, however, geophysicists usually prefer to build models that fit the data to some acceptable level. In this case uncertainty is regarded as a measure of how far the preferred model is likely to be from the truth. A simple example where this is valid is linear regression (fitting straight lines through scattered x,y data). In regression problems a pair of error bars on each estimated unknown (e.g., the slope and gradient) represents their statistical uncertainty given the "known" errors in the data. Closely related to regression problems are linear discrete inverse problems, i.e., ones where the number of unknowns is fixed and the mathematical relationship between data and unknowns is linear. (Note that many exploration problems are either linear or are approximated as linear inverse problems. For example, techniques based on seismic ray theory fall into this category, e.g., traveltimes inversion, if the dependence of raypaths on the seismic wavespeeds is ignored.) For linear discrete inverse problems, well-known techniques can be used to obtain the standard deviations (or errors) of each parameter (Menke, 1989;

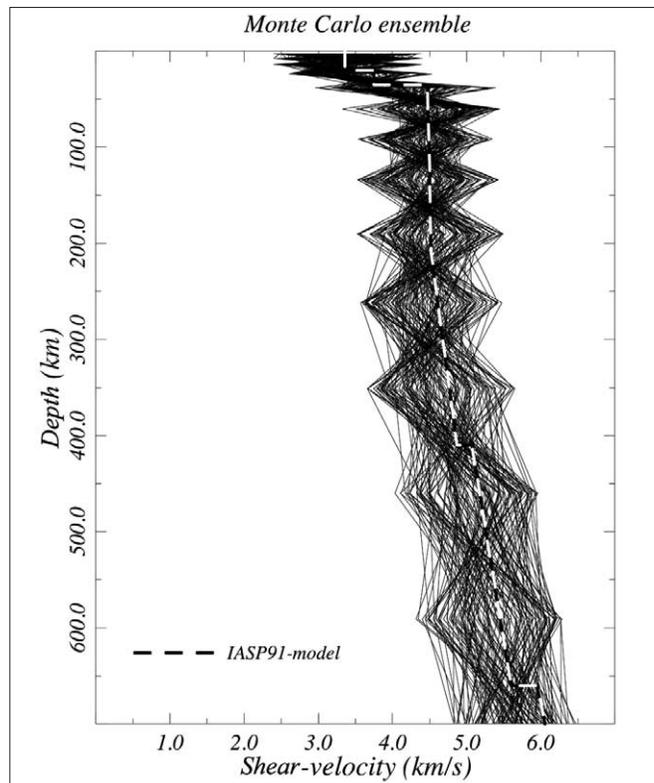


Figure 1. An ensemble of models obtained from a Monte Carlo search that fit measurements of the Rayleigh-wave group velocity within the error bounds.

Tarantola, 2005; Aster et al., 2005). There is, however, no general theory to characterize the uncertainty in nonlinear inverse problems (Snieder, 1998). The latter arise when simplifying assumptions on the relationship between data and unknowns are dropped, e.g., if we recalculate seismic raypaths every time the wavespeed model is updated in traveltimes inversion.

Characterizing uncertainty. We present examples from global seismology that capture some essential features of the characterization of uncertainty in geophysical inverse problems. Figure 1 shows a set of models of shear velocity in the upper mantle that fit a certain set of group velocity measurements of the fundamental-mode Rayleigh wave (a surface seismic wave confined to the vertical plane between source and receiver and sensitive to earth structure in the outer few hundred kilometers of the crust and lithosphere) (Douma et al., 1996). One might be tempted to characterize the uncertainty in this model by computing the standard deviation of this population of models at every depth. This approach is, however, flawed, because the variance of the population changes as the model parameterization is changed; a model with thinner layers allows for larger variations in the velocity. Moreover, characterizing the uncertainty by the variance only does not account for the trade-offs that are present in this population of models; a high veloc-

ity at a certain depth is, in general, associated with a lower velocity at adjacent depths.

The characterization of uncertainty is even more difficult to handle for models that represent a two- or three-dimensional field, as arises in tomographic problems. In that case displaying the variance constitutes a nontrivial visualization problem. One strategy is to adopt a color scheme based on hue-saturation-brightness (HSB) values. Hue and saturation can be used to represent the model value itself, with the brightness controlled by the width of the confidence interval. This conveys a visual impression of uncertainty through the intensity of the image, which can be intuitive. Alternatively image-processing techniques might be used to blur an image with a chosen filter whose width depends on the uncertainty value (Johnson, 2003). The space-filling option would be to display scalar uncertainty estimates as a separate image to accompany each tomographic slice shown. All of these techniques have been used with varying success. Visualization of simple scalar estimates of uncertainty is, however, rather limited since the correlation of errors between estimated parameters (described by off-diagonal elements of a model covariance matrix) as well as the trade-off between the parameters themselves (described by a resolution matrix) are not well characterized by simple error bars. (See Figure 2 and Box 1 for an explanation of model covariance and model resolution matrices.)

Visualizing sensitivity. Box 2 together with Figure 3 explain the concept of nonuniqueness in an idealized three-parameter linear inverse problem in terms of how data misfit varies as a function of the model parameters. The idea of directly visualizing the sensitivity of data fit to model parameters can be useful in some real problems also, i.e., by examining how data fit varies in the vicinity of a preferred solution model. If a least-squares misfit criterion is used to measure data fit then it can (in principle) be “mapped out” by repeated forward solutions either along axes passing through the “best fit” solution, or across 2D planes centered about a pair of parameters in the solution. Figure 4 shows some examples. Each colored dot is the result of a forward calculation with one or two variables perturbed from the preferred solution in parameter space. The height and color of the dot represents the fit to data with higher values indicating a worse fit. If the data model relationship were linear, each “slice” would have a simple quadratic shape. The minimum in the data misfit surface indicates the best data-fit model. If this is at a single unique point, then the inverse problem is said to have a global unique solution. In this case the curvature of the surface at the solution determines the model covariance and resolution matrices routinely calculated in discrete linear inverse problems (Menke, 1989; Tarantola, 2005). However if the lowest data misfit value lies in an extended flat region, or perhaps in a long, narrow valley (as in panel D of Figure 4), then there is no unique solution and the problem is said to be “ill-posed.” Since the shape of the misfit surface controls the relative degree of constraint placed on each parameter, trade-offs between pairs of parameters can be detected by examining contours of the data misfit, with long narrow valleys suggesting strong trade-offs between corresponding variables. This process of mapping out the misfit function in the vicinity of a solution is not restricted to linear problems and can be illustrative in nonlinear cases as well. Figure 4 shows an example of the least-squares data misfit used in a problem of seismic receiver function inversion. (A process whereby seismic shear wave speeds in beneath a recording station are inverted for, using a “transformed” three-component seis-

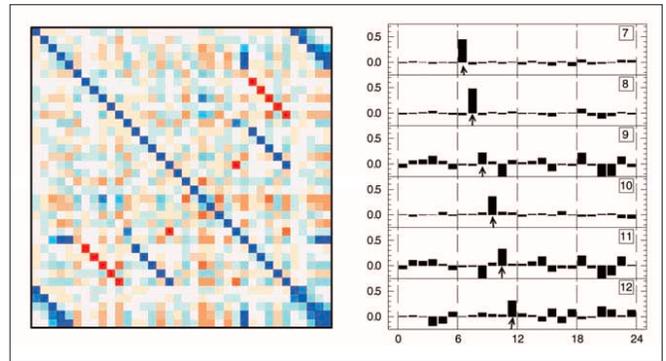


Figure 2. An example of a model covariance matrix (left) and six rows of a model resolution matrix (right) for an inverse problem with 24 parameters. On the left, blue indicates positive covariance values, and red indicates negative ones. On the right, the number in the box identifies the parameter index. Each column height represents the resolution value. The arrow shows where the row crosses the diagonal of the resolution matrix. For rows corresponding to well-resolved parameters, we expect to see a dominant column at the position of the corresponding arrow. (See Box 2).

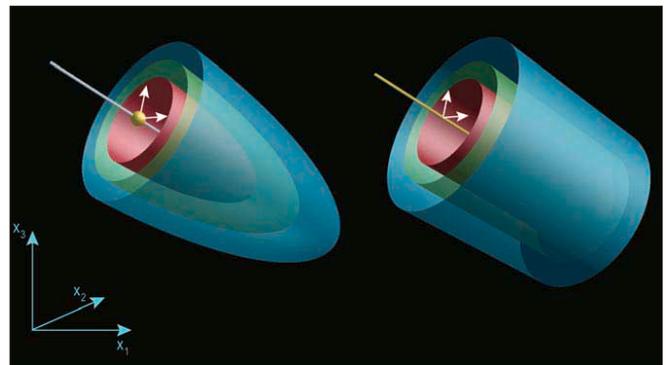


Figure 3. A visualization of the concept of nonuniqueness in a three-parameter linear inverse problem. See Box 2 for details.

mogram at the surface.) In this particular case the 24 unknowns represent a 1D shear wave velocity profile as a function of depth beneath a seismic station. The data misfit measures the difference between predicted and observed seismic waveforms at the surface. This is a nonlinear problem, which is clearly seen in the complex nature of the misfit function near the solution.

If the computational burden of the forward problem is high, the mapping out process may become rather costly and even impractical in some cases. However, since each forward solution is independent, one can take full advantage of modern cluster computing hardware, or large parallel computing facilities. Indeed for moderately sized problems and with appropriate software it is quite possible to interactively visualize such misfit surfaces in real time. (The software which generated Figure 4 did this across a networked set of desktop PCs, using a freely available software toolkit written in the Java programming language.) In this case repeated solutions of the forward problem in parallel show the limitations of linearized estimates of uncertainty (e.g., see Aster, 2005), which require that the misfit in the immediate vicinity of the solution is representative globally.

Characterizing the trade-off. Douma et al. (1996) used principal-component analysis to compute the patterns of variability that are hidden in the population of earth models in Figure 1. This approach is useful when the population models form one cluster, as in the example of Figure 1.

A popular approach in tomographic problems are check-

erboard or spike tests whereby a synthetic input model (i.e., forward-modeled predictions based on a known set of parameters) is reinverted for. Comparison of input and output models gives an impression of likely blurring effects and length scales of resolvability. This technique has its critics, but interestingly it is essentially a discrete version of one advocated by Parker (1994) as a way of estimating resolving widths in a 1D continuous inverse problem (see Figure 4.02a of Parker, 1994).

In general, uncertainty estimation is complicated when the inverse problem is nonlinear, data noise characteristics are poorly known, model parameterization is inadequate, and approximations in the forward theory are present. Seismologists have to grapple with all of these issues and are expected to convey meaningful uncertainty estimates to their peers through visualization.

Changing the parameters. A commonly used heuristic procedure to estimate the effect of uncertainty arising from choices made in parameterization and regularization is to perform repeated inversions with these factors varied between runs. A qualitative estimate of the robustness of constructed models is made by comparing the results of different inversions. An example, taken from the field of global seismology (Beghein and Trampert, 2003), shows how this can fail even if the problem is linear. Here three dependent variables (α , β , and η), which characterize seismic anisotropy in earth's inner core, are constrained by normal mode data. The upper and middle panels of Figure 5 show the result of two linearized inversions using various levels of damping. (Mathematicians know damping as regularization and use it to stabilize inversion of ill-conditioned matrices. In geophysics, this often takes the form of imposing additional a priori information, such as smoothness constraints.) The model shown in the upper panel of Figure 5 is heavily damped, the model in the lower panel less so. The lack of change in the constructed models might tempt a conclusion that the model uncertainty is small in the deepest layers. This is, however, not so. It turns out the data have little sensitivity to structure in the deepest layers (none at the center of the earth) and the regularization has forced the variation to nil. The lowest panel in Figure 5 shows the results of a model space search using the Neighborhood Algorithm, a new automated parameter search technique which seeks out regions of model space where data fit is high and concentrates sampling there. The solid lines in Figure 5 show variability spanned by twice the standard error in the family of models found which were compatible with the data. In the upper regions of the model the results are similar to the lightly damped linear inversion but at the base the spread is large, confirming that the normal-mode data poorly resolve earth structure there, and uncertainty should be large, not small as suggested by the "tweak it and see" strategy. The models in the new ensemble are also compatible with independent traveltime data which are sensitive to the deepest part of the alpha model. This example shows the problems that can occur in using simple "bootstrap" type estimates of uncertainty on regularized solutions of ill-posed inverse problems.

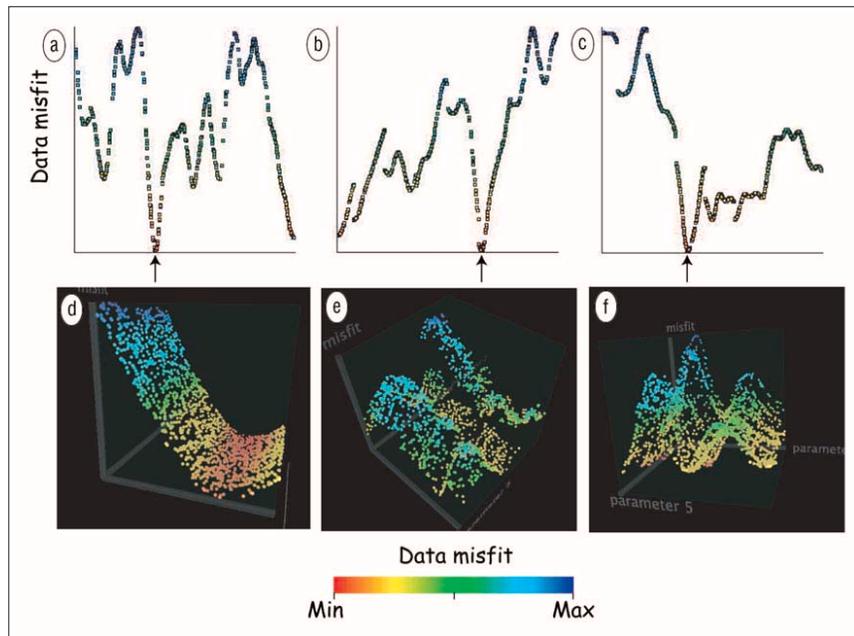


Figure 4. Exploring data fit about a single solution in the seismic receiver function problem. Here a 24-dimensional parameter space is explored in real time using interactive software. The top panels show the variation of data misfit along three axes through a solution. In each case the optimum model lies at the misfit minimum (marked with an arrow). Behavior can be multimodal and the data misfit may even contain discontinuities (as in cases b and c) indicating strong nonlinearities in the inverse problem. Panels d–f show 2D data misfit surfaces. Height and color represents increased data misfit. This type of repeated forward modeling combined with interactive visualization can help determine the degree of constraint placed on parameters, show trade-offs and assess the nonlinearity of the problem. Real time rotation of “dot” figures, such as those in d–f creates a clearer impression of the the misfit surface than can be displayed in the “snapshot” figures shown here.

Posing questions to data. In investigating a complex geologic problem, like the origin of seismic anisotropy below a continent, perhaps the best approach mixes the intuitive with the formal, and the graphically clever with the explicit. In our view the essence of an inverse problem is about asking questions of data, rather than constructing models per se. Hence uncertainty estimates are only useful if we know how they alter decisions or interpretations drawn from the data. Simons and van der Hilst (2003) take this view in a study of upper mantle anisotropy beneath Australia. Here the crucial question is: Are the directions of seismic anisotropy (the fast axes, shown as black sticks in Figure 6) aligned with the absolute plate motion or not? Oddly, “absolute” plate motion is a relative concept. Assuming hot spots to constitute a fixed reference frame for the movement of the earth’s lithospheric plates is not without detractors. But requiring the net rotation of all lithospheric plates to be zero reflects mantle strain under only certain restrictive assumptions. The (lack of) alignment of seismic anisotropy in either reference frame, hot spot (HS), or no-net-rotation (NNR), may help us “decide” which is most physically plausible and supported by the data.

The lower panels of Figure 6 show histograms of the angular difference between anisotropy and HS (gray bars) and NNR (red squares) plate motion directions, as a function of depth. The flat slopes above 140 km depth suggest anisotropy is not aligned with absolute plate motion here, in either reference frame; below it, the opposite appears to be true. The hot-spot reference frame (gray) is supported over the no-net-rotation frame (red): Its histograms are more peaked at the zero angle indicating perfect alignment. How are these conclusions affected by uncertainty? It is nearly impossible to assign a significance to the difference in steepness in the histograms. The directions of anisotropy are the result of a mul-

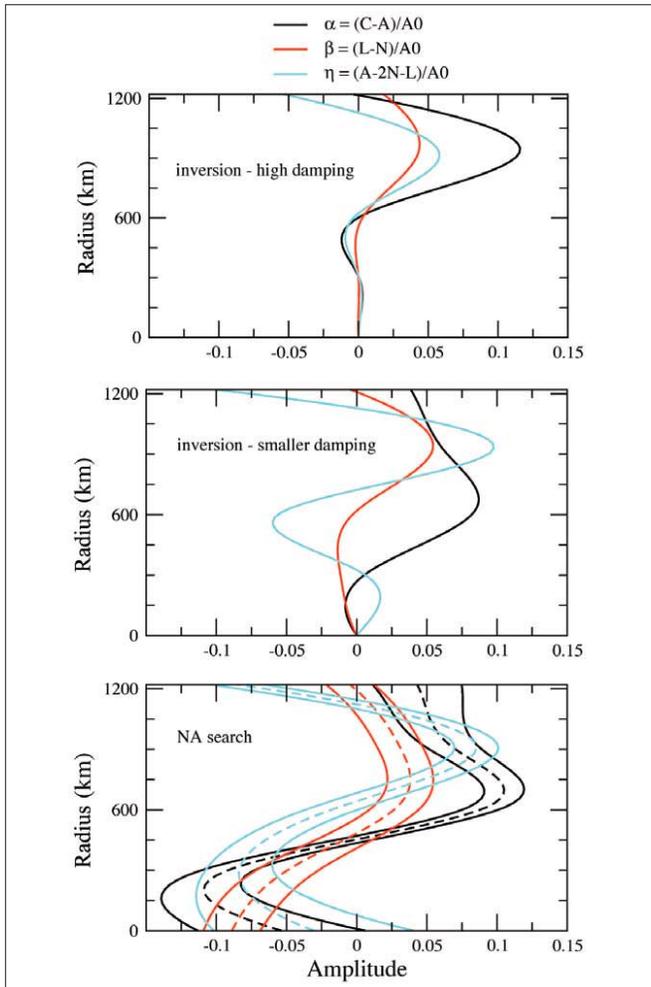


Figure 5. Models of inner core anisotropy resulting from the inversion of normal-mode data (upper two panels) and from the application of the Neighborhood Algorithm (lower panel) to these same data. (Normal modes are whole-earth oscillations excited by large earthquakes.) α , β , and η are the three anisotropic parameters to which the data are sensitive. A larger damping was applied to the models displayed in the top panel than to the models in the middle panel. The third panel displays the range of good data-fitting models obtained from a model-space search approach. The dotted line represents the mean model, and the thick surrounding lines correspond to two standard deviations.

titude of inversion steps, starting from the individual waveforms of many thousands of seismograms, inverted nonlinearly for path-average structure, culminating in a three-dimensional model which is a result of a linear inversion of all the path constraints. Along the way, choices about parameterization were made, damping factors chosen, subjective errors of “what fits” are assigned, etc. A veritable nightmare of uncertainty for any decision-maker.

Simons and van der Hilst (2003) address this question by experimenting with synthetically generated data which mimic the hypothetical truth as closely as possible (see their Figure 5). Working back from what they want to conclude (Is the decision sound? Can we tell the cases apart? Which is best supported by data?), they test out various what-if scenarios. They calculate forward models based on known input models, reinvert the synthetic data, and compare the effect on the histograms directly. This, they do for a variety of random noise levels on the input data. It appears that in Australia, the data do allow them to determine the difference between the two reference frames—and, what’s more, their conclusions stand up to a sea of noise. Such an approach is heuristic, incomplete, and imprecise, but in the absence of a tractable theory lack-

Box 1. Model covariance and resolution matrices.

A common way of describing uncertainty in linear inverse problems is by means of model covariance and resolution matrices. Figure 2 shows an example of a model covariance (left) and the rows of a model resolution matrix (right) for a 24-parameter inverse problem. In essence the model covariance matrix shows how errors in the data propagate into errors in the estimated model parameters. It is a square symmetric matrix with as many rows and columns as there are unknowns in the problem. Off-diagonal entries indicate how errors are correlated between pairs of model parameters. For example if the value is positive then the errors in the two corresponding parameters will tend to be positively correlated [e.g. when data errors force one parameter too high (low) they will also tend to force the other too high (low)]. A negative entry indicates the converse. Significant off-diagonal entries indicate strong trade-offs in the variables. Diagonal entries give information on the variance of corresponding parameters, i.e., how data noise affects them individually.

It is important to note that the uncertainty is described by the whole matrix and not just the diagonal entries. In practice off-diagonal entries are often ignored which means that the influence of data noise can be severely underestimated, even in linear inverse problems. For example this is often the case in published catalogs of earthquake locations, where usually only the diagonal entries of the calculated model covariance values are ever recorded. As the number of unknowns increases model covariance matrices quickly become difficult to calculate both due to the computational burden, and also because the number of entries grows quadratically with the number of unknowns.

The right side of Figure 2 shows six rows of the resolution matrix. These describe how independently each parameter can be recovered from the data. Each row corresponds to a particular parameter and the height of the columns gives the relative “leakage” or blurring of one parameter into another. For example in Figure 2 parameter 7 (first panel) is relatively well resolved since its own column (shown by the arrow beneath) is much higher than the others, whereas parameter 9 (third panel) is poorly resolved because its column is virtually indistinguishable from the other columns. Unlike the model covariance matrix the resolution matrix does not depend on the errors in the data but instead reflects the uncertainty arising from the nature of the forward problem as well as the amount and distribution of the available data. For example, in seismic traveltimes inversion it is influenced by the distribution of the sources and receivers (and hence raypaths) as well as the thinness of the rays implicit in an infinite-frequency ray theory approximation.

ing approximations (in both the physics and the mathematics), without access to infinite computer time, and in an imperfect world, directly focusing on how uncertainty affects decisions/interpretations is a practical way forward.

The way forward? Characterizing and communicating the uncertainty in geophysical inverse problems is still a problem of formidable proportions. Finding effective solutions is crucial for establishing a meaningful interface with those who make financial or scientific decisions based on the data. Increases in computing power make it possible to produce not just a single model, but populations of models, that can be used as input for stochastic modeling for the production of reservoirs. The example of Figure 1 shows, however, that even for simple one-dimensional earth models a large population is needed to characterize the uncertainty.

As discussed above, one way of looking at inversion is a process of asking questions of data. Trying to build a com-

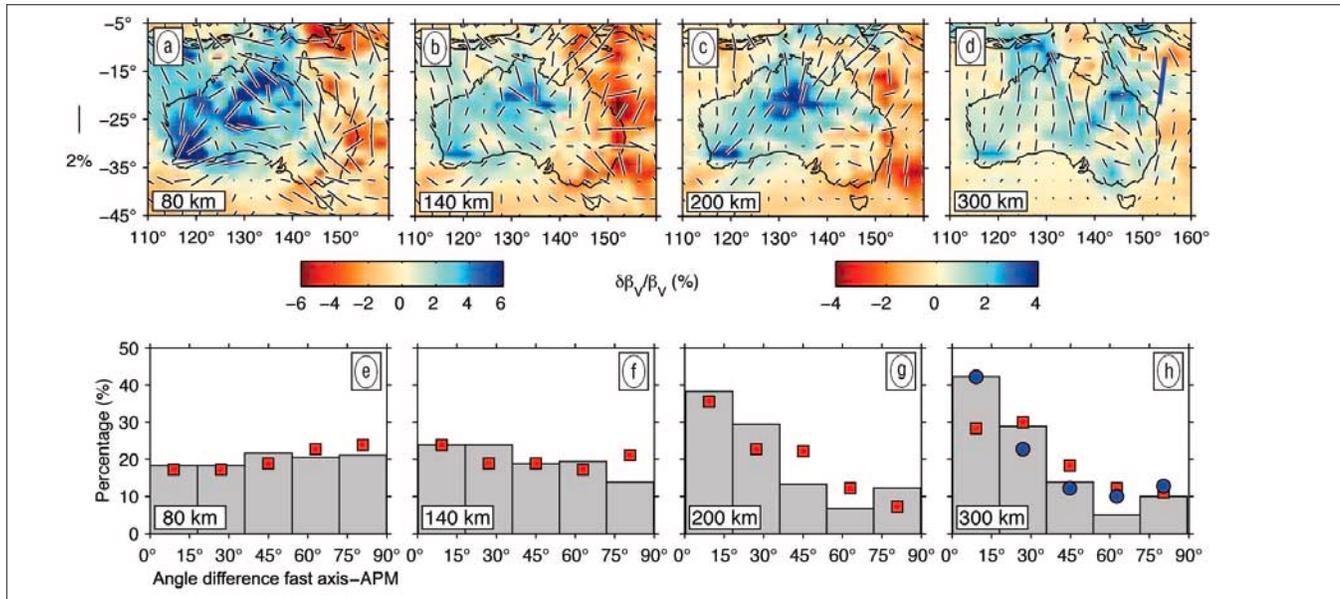


Figure 6. Are the anisotropic directions of fast shear-wave speed propagation in the Australian upper mantle related to the movement of the Indo-Australian plate, and in what frame of reference? (a–d) Isotropic shear-wave speed anomalies (faster, slower than average) and fast polarization directions (black bars). (e–h) Histograms of the angle between fast directions and plate motion in two reference frames (gray bars, red squares), and compared to an independent estimate derived from the anisotropy data themselves (blue).

Box 2. Nonuniqueness and uncertainty.

The concept of nonuniqueness can be succinctly explained with the aid of visualization. Consider the case of a simple three-parameter inverse problem where the mathematical relationship between model parameters and data is linear. Let's assume the discrepancy between predictions and observations is measured with a standard least-squares criterion. For every point in the three-dimensional parameter/model space the fit to data can be calculated. Figure 3 shows two different cases. The left panel is for a "well-posed" inverse problem with a unique solution. We see isosurfaces of a least-squares data fit function. Every point on the isosurface has the same fit to data, and in this case they form ellipsoids closed about a unique best-fit set of parameters (represented by a yellow sphere). In the right panel the isosurfaces are cylindrical and correspond to an ill-posed problem. Here there is no unique solution to the inverse problem because every point along the yellow axis has minimum data misfit. In the left case, uncertainty is well described by the confidence ellipsoids. On the right side, uncertainty is defined only in directions perpendicular to the axis and the combination of parameters along the axis is completely unconstrained. The right inverse problem has nonunique solutions with a complete trade-off between parameters corresponding to the direction of the axis.

plete model of the subsurface from some particular data set is tantamount to asking a very broad-ranging question. For example "Give me your best estimate of the electrical conductivity at every position in a 3D volume within my petroleum reservoir from the available EM field measurements." Many of the difficulties in characterizing uncertainty arise from the imprecise nature of the questions we ask of data. For example, there may be many (or an infinite number) of electrical conductivity models that fit EM data accurately, and this nonuniqueness means that it may not be possible to determine uncertainty everywhere in the 3D volume at once. If a more precise question is asked the situation may be different. For example "What is the maximum depth (or volume) of that reservoir or ore body, given the available data and some reasonable bounds on material properties everywhere?" Often

uncertainty is both simpler to determine and easier to convey if the underlying data question is simpler or more precise. Perhaps then the way forward is to ask many detailed questions of our data rather than a few loose ones. In each case uncertainty may have a more straightforward interpretation and the answers to many simple questions may collectively convey reliable information. This is not to say that attempts at "complete" inversion via model building are not useful. As statistician Samuel Karlin put it: "The purpose of models is not to fit the data but to sharpen the questions."

It is clear that further research is needed. Returning to exploration geophysics, we hope this contribution will stimulate further research on the characterization of uncertainty in nonlinear inverse problems, the visualization of uncertainty, and on creating an effective interface between geophysicists who construct models and those who are responsible for business decisions.

Suggested reading. *Parameter Estimation and Inverse Problems* by Aster et al. (Elsevier, 2005). "Robust normal mode constraints on inner-core anisotropy from model space search" by Beghein and Trampert (*Science*, 2003). "Ensemble inference in terms of empirical orthogonal functions" by Douma et al. (*Geophysical Journal International*, 1996). "A next step: Visualizing errors and uncertainty" by Johnson (*Computer Graphics and Applications*, IEEE, 2003). *Geophysical Data Analysis: Discrete Inverse Theory* by Menke (Academic Press, 1989). *Geophysical Inverse Theory* by Parker (Princeton University Press, 1994). "The anatomy of inverse problems" by Scales and Snieder (*GEOPHYSICS*, 2000). "Seismic and mechanical anisotropy and the past and present deformation of the Australian lithosphere" by Simons and van der Hilst (*Earth and Planetary Science Letters*, 2003). "The role of nonlinearity in inverse problems" by Snieder (*Inverse Problems*, 1998). *Inverse Problem Theory and Methods for Model Parameter Estimation* by Tarantola (Siam, 2005). **TJE**

Acknowledgments: Peter Rickwood assisted with the production of Figure 4. The toolkit used for parameter search on a computational cluster is available from <http://rses.anu.edu.au/cadi/Software.html>.

Corresponding author: Malcolm.Sambridge@anu.edu.au