Nonlinear Processes in Geophysics (1995) 2: 222 - 227

Nonlinear Processes in Geophysics

Identifying sets of acceptable solutions to non-linear, geophysical inverse problems which have complicated misfit functions

A. Lomax and R. Snieder

Department of Geophysics, Utrecht University, PO Box 80.021, 3508 TA Utrecht, The Netherlands

Received 16 September 1994 - Accepted 13 January 1995 - Communicated by D. Sornette

Abstract. A goal of geophysical inversion is to identify all models which give an acceptable misfit between predicted and observed data. However, because of the complexity of Earth structure, the non-linearity of physical processes in the Earth, and the insufficiency of geophysical data, many geophysical inverse problems may have a large number of distinct, acceptable solutions. These problems may be characterized by a complicated surface for the misfit function in the solution parameter space. For exploring such a surface, direct inversion and simple random search methods are often inadequate. However, directed search methods such as the genetic algorithm can be configured to balance convergent and random processes to find large sets of solutions that span the acceptable regions of complicated misfit surfaces.

1 Introduction

A goal of geophysical inversion is to find all earth models which, when operated upon by some forward method, produce synthetic data that gives an acceptable agreement with a set of observed data (Keilis-Borok and Yanovskaya, 1967; Press, 1968; Aki and Richards, 1980; Tarantola, 1987). In terms of the misfit function, a measure of the difference between observed and synthetic data that varies as a function of solution parameters, this goal of inversion is to identify all regions of the misfit surface with values below an acceptable level. This acceptance level is determined by considering variance in the data set and the loss of generality due to all simplifications and assumptions in formulating the inversion.

However, because of the complexity of Earth structure, the non-linearity of many physical processes in the Earth, and often insufficient and noisy geophysical data, many geophysical inverse problems have a large number of acceptable solutions. For such problems, the misfit function may be a complicated, irregular surface of peaks, valleys, ridges and local minima (Fig. 1; *c.f.* Scales *et al.*, 1992). In general, little may be known about the continuity, differentiability and number of minima of a complicated misfit surface. In addition, the shape of the misfit surface itself is not absolute, it will shift and deform in response to changes to the data set, a priori constraints, model parameterization, inversion procedure, approxima-

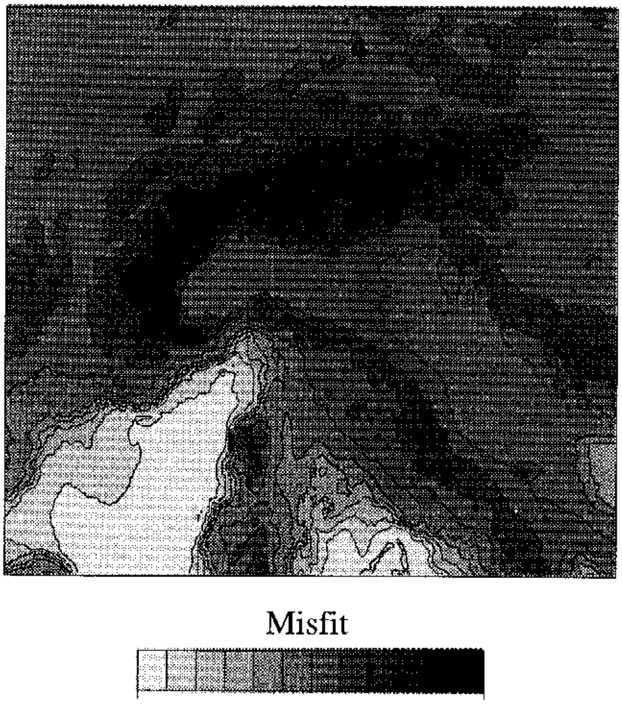

Fig. 1. Complicated misfit-surface in a 2-parameter solution-space. This surface representing the topography around the Alps has features that may be typical of the higher-dimensional misfit surfaces from geophysical inverse problems. Such features include: a broad, curving, elongated shape for the minimal region (related to poor constraint and trade-off between model parameters and uncertainty in the data), and multiple-minima and a fractal character (related to non-linearity).

Correspondence to: A. Lomax

tions in the physical theory of the forward problem and many other aspects of the inversion (*c.f.* Tarantola and Valette, 1982; Tarantola, 1987). This means, in particular, that the location of global and local minima of the misfit surface are dependent on the data set and formulation of problem as well as the properties of the "real" Earth.

These considerations imply that a comprehensive and meaningful solution to many inverse problems requires estimates of the acceptable regions of the misfit surface and of how the shape of the misfit surface depends on all assumptions in the problem. In this paper we discuss the first of these goals. We show how guided random-search techniques such as a genetic algorithm can be configured to identify the regions of a complex misfit surface that lie below a significance level. We call the solutions that lie within these regions *acceptable solutions*. We apply the genetic algorithm to a many-parameter geophysical inversion problem and show how the set of acceptable solutions indicates uncertainty and trade-offs in the results.

2 Existing Inversion Methods and Genetic Algorithm

Many geophysical problems are addressed through the use of direct, calculus based inversions. These methods require knowledge of an adequate starting solution, and that perturbations of the model are linearly related to changes in the data (Aki and Richards, 1980). They operate through a single or an iterative perturbation of the solution using locally determined gradients of the misfit surface; this requires a smooth and differentiable misfit function. Also, only a single final solution is identified, and, in general, this solution will be in the neighbourhood of and strongly dependent on the location of the starting solution (Fig. 2). Thus these methods are inherently unable to define the topography of a complicated misfit surface over a large solution-space.

Random, trail-and-error and enumerative or grid-search techniques are applicable to problems with complex misfit surfaces because these methods work directly with non-linear forward calculations and require no gradient information. Also they allow a large solution-space to be explored and produce multiple solutions. However commonly used techniques such as simple Monte Carlo and hedgehog (Keilis-Borok and Yanovskaya, 1967) become inefficient or impractical in very large solution-spaces. More recently a number of guided search techniques from the field of "artificial intelligence" have been developed which can sample efficiently a large solution-space. The genetic algorithm (GA) method (Goldberg, 1989; Davis, 1991, Holland, 1992) is one such technique that has been applied to geophysical problems (e.g. Stoffa and Sen, 1991; Sen and Stoffa, 1992; Sambridge and Drijkoningen, 1992; Sambridge and Gallagher, 1993; Jin and Madariaga, 1993; King, 1993; Nolte and Frazer, 1994; Lomax and Snieder, 1994).

The GA method is an iterative, guided search, which

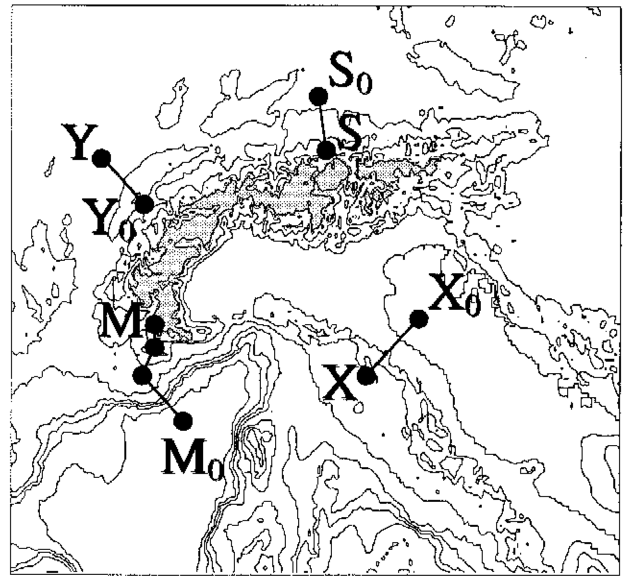

Fig. 2. Direct inversion on the complex misfit surface from Fig. 1. A region of acceptable misfit is shaded in grey. A single-step direct inversion (S_0 -> S) makes a perturbation to the starting solution S_0 that depends on the local gradient of the misfit surface at location S_0. A multiple-step direct inversion (M_0 -> M) makes successive perturbations to the solution that depend on the local gradient of the misfit surface at each solution location. For some starting solutions near to the acceptable region the final solutions S or M may be close to or within the acceptable region. But for most starting solutions, the complexity of the misfit surface and the distance of the starting solution from the acceptable region will prevent a solution within the acceptable region from being found (X_0 -> X). Even with starting solutions close to the acceptable region, the inversion may diverge (Y_0 -> Y).

applies stochastic operations to populations of trial solutions to find new solutions with smaller misfit. Beginning with a random population of solutions, succeeding populations are created by 1) *selection* - saving solutions with smaller misfits, 2) *crossover* - combining parts of two solutions to form new trial solutions, and 3) *mutation* - changing the values of some of the parameters of a solution. New populations are created for a set number of iterations or until some criteria for misfit reduction has been achieved. The GA search produces a large set of solutions which sample the solution-space globally but which can rapidly converge to a local or global minimum. As with other Monte Carlo methods, a primary limitation of the GA method is the speed of the forward calculation which typically must be done many thousands of times. Also, the construction and tuning of the selection, crossover and mutation operations are done in a primarily ad-hoc manner.

3 Genetic Algorithm Inversion for Acceptable Solutions

Unfortunately, with a complicated misfit-surface, a GA may define only one or two local minima (Fig. 3a), and different local minima may be found depending on the parameters controlling the GA search (e.g. Goldberg and

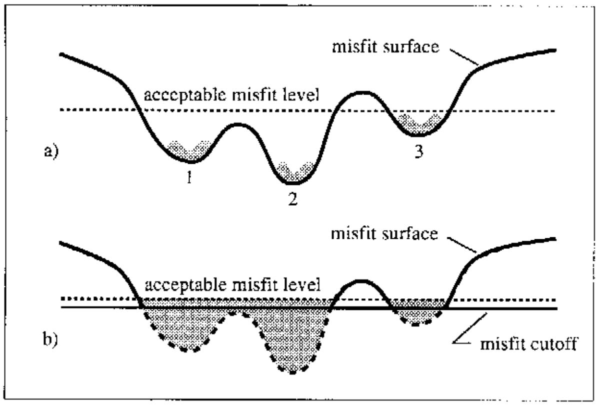

Fig. 3. a) A "typical" GA may find solutions located near only one of the minima, 1, 2 or 3 in the misfit surface. But a goal of geophysical inversion is to find all solutions with misfit below some acceptance level (e.g. data variance); all such solutions may give useful information about the problem. b) The identification of a representative sample of all acceptable solutions is a practical way to achieve the goal of geophysical inversion. A GA with a misfit cutoff value at or below the acceptable level can aid in identifying a wide range of acceptable solutions and avoiding deep local minima.

Richardson, 1987; Stoffa and Sen, 1991; Nolte and Frazer (1994). This behaviour occurs when a GA is configured for rapid minimization of the misfit; the results are strongly influenced by the best existing solutions in the population (Fig. 4, top). Such convergence to a local minimum not only prevents the identification of all acceptable solutions, but also may produce a localized, poorly distributed set of solutions that is not appropriate for later statistical analyses or for use to estimate the *a posteriori* probability density function for the inversion (Tarantola and Valette, 1987; Sen and Stoffa, 1992; Nolte and Frazer, 1994).

Here, we configure the GA to attempt to find sets of acceptable solutions - solutions representing all regions of the solution-space that give a misfit with the data below some acceptance level. We begin with a GA similar to that described by Sambridge and Drijkoningen (1992), but the rate of crossover is set lower relative to the rate of mutation, and the best solution of each generation is never explicitly carried over to the next generation (no *elitism*, see Goldberg, 1989). In addition, we define a minimum misfit "cutoff" value and reset lower misfits to this value; this helps to prevent the stalling of the GA in deep minima that are much lower than the acceptable level (Fig. 3b). This cutoff makes the method more of a *stochastic* search and less of an *optimization* method for regions of the solution space with misfit below the cutoff value. The resulting GA configured to find acceptable solutions is referred to as GA1.

Figure 4 shows a comparison of a "typical" GA, the GA1 and a random search in imaging an acceptable misfit region of the sample misfit surface in Fig. 1. A "typical" GA (Fig. 4, top row) converges rapidly towards a local minima, testing many solutions in the neighbourhood of this minima; it finds a large number of acceptable solutions, but they are highly clustered and do not form a good sample of the acceptable misfit region. With different starting populations (different columns in Fig. 4), different minima are be identified. The GA1 (Fig. 4, middle row) converges towards the acceptable misfit region, but tests many solutions throughout the solution-space; it finds fewer acceptable solutions that the "typical" GA, but these solutions are less clustered and give a better image of the acceptable misfit region. With different starting populations the GA1 shows more stability in the results than the GA. A random search or crude Monte Carlo (Fig 4, bottom row) tests solutions that are well distributed throughout the solution-space; but finds very few acceptable solutions and does not image the acceptable misfit region as well as the GA1.

A measure of the relative efficiency of the three methods is given by the number of acceptable solutions identified (about 330/90/20 for GA, GA1 and random search, respectively); but, as indicated above, this measure does not indicate the quality of the distribution of the acceptable solutions. Also, if the dimension of the parameter space were increased, the relative efficiency of the random search would decrease rapidly.

The results in Fig. 4 show that the GA1 gives a better image of the acceptable region of the misfit surface than a "typical" GA, while maintaining the high efficiency of the GA relative to a random search. However, the GA1 still shows some tendency to produce clustered solutions in the acceptable misfit region, and the sampling of the remainder of the parameter space is not as well distributed as the random search. There are other modifications and parameter adjustments to the GA that may help to achieve the goal of finding a representative set of all acceptable models (e.g. Goldberg and Richardson, 1987), and much work remains to be done on this problem.

4 Application to a Geophysical Problem

We illustrate the use of the GA1 to find sets of acceptable solutions in a seismological inverse problem, summarizing some of the results of Lomax and Snieder (1994). We invert synthetic group-velocity estimates to determine layered S velocity models. The data consists of fundamental Rayleigh group-velocity dispersion estimates from 10 to 350 sec period for the layered earth model iasp91 (Kennett and Engdahl, 1991) with noise added by incorporating the scatter from real data (see Lomax and Snieder, 1994). This gives a noisy group velocity data set with a mean at each period that corresponds to the dispersion for the iasp91 model. The models are defined by 4 crustal and 14 mantle velocity-depth nodes and a Moho discontinuity at a depth between 15 and 70 km. The parameterization gives about 10^{45} possible models, though the number of significantly different models is the order of 10^{10} to 10^{20}. We invert the noisy synthetic data using this model parameterization and the GA1 procedure described above, with the expectation

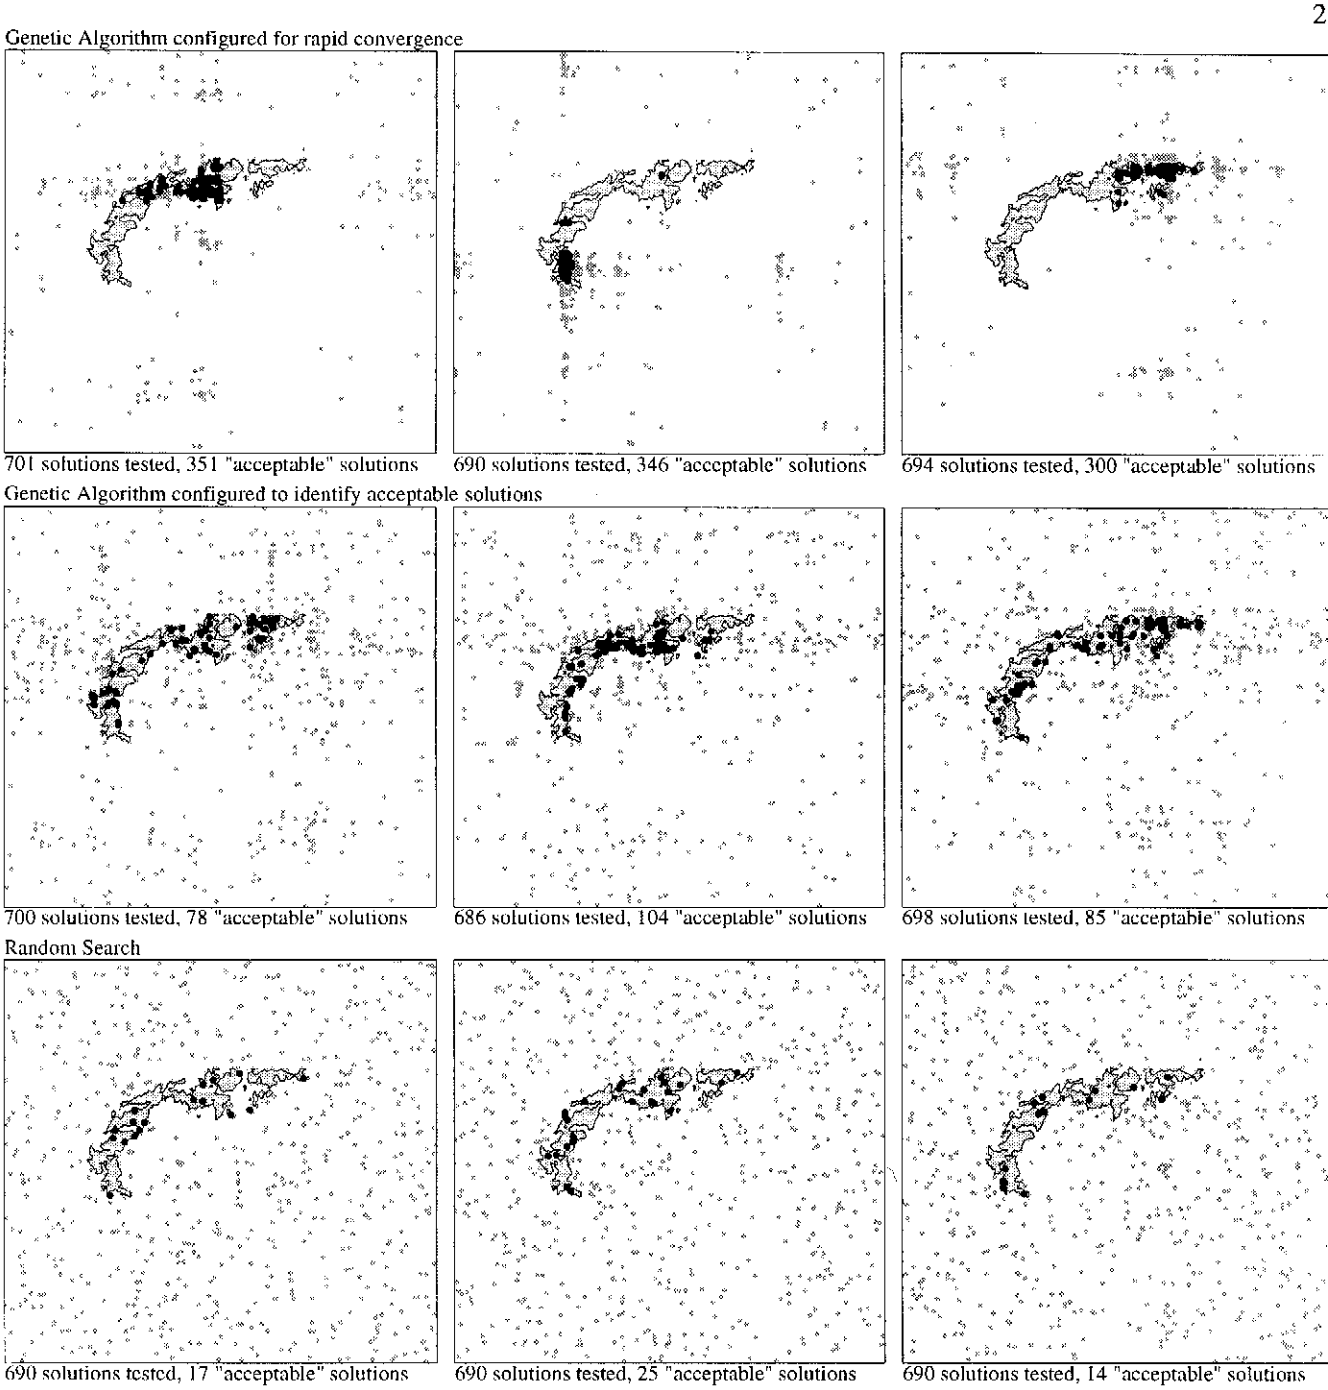

Fig. 4. Comparison of a "typical" GA tuned for rapid convergence (top row), a GA configured to find acceptable solutions (GA1, middle row) and a random search (bottom row) in imaging an acceptable misfit region (grey tone) of the misfit surface of Fig. 1. Each row shows three runs with different starting populations for a single method; all runs have about the same number of tested solutions and require about the same computation time.

that the iasp91 S velocity distribution will be contained within the scatter of acceptable models. Note that this test does not use exact parameterization - the input iasp91 model cannot be represented exactly by the nodal models used for inversion.

Figure 5 shows the acceptable models and their $\pm 1\sigma$ and $\pm 2\sigma$ spread in velocity at each depth from three GA1 runs with the noisy synthetic data. Here, acceptable models are defined as those models giving a synthetic dispersion curve with an r.m.s misfit with the data (E_s) less than 0.85 times the r.m.s scatter of the data values (E_d); this insures that the predicted data from acceptable solutions will fall within the scatter of the real data. The distribution of acceptable models is an estimate of the shape of the better-fitting region of the misfit surface in the solution space. Allowing for differences in parameterization in the crust and at the "400" km discontinuity between the iasp91 model and the models used for inversion, and the noise added to the data,

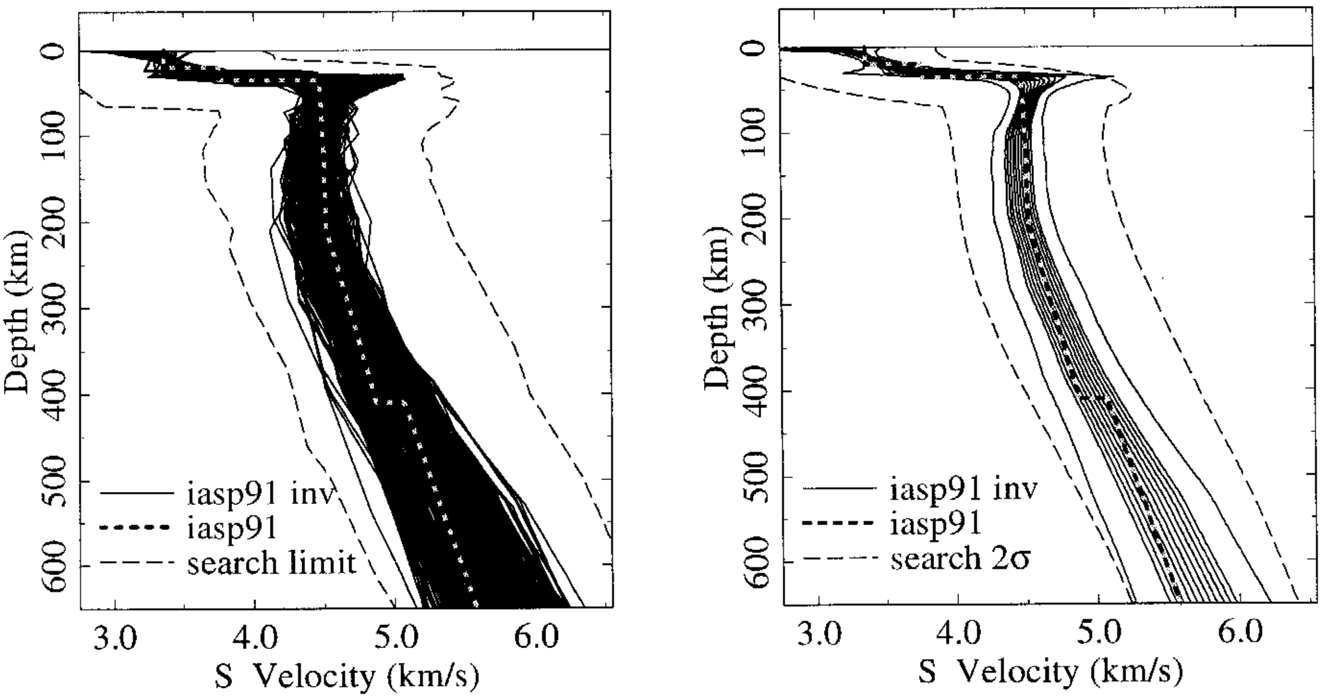

Fig. 5. Combined results from three runs of the GA1 for the synthetic iasp91 data. The left panel shows all acceptable models (misfit $E_s \leq 0.85E_0$) and the search limits (dashed lines). The right panel shows the $\pm 1\sigma$ spread (closely spaced solid lines) and the $\pm 2\sigma$ spread (outer solid lines) in S velocity at each depth for the acceptable models and the $\pm 2\sigma$ spread of the search (dashed lines). The iasp91 model used to generate the synthetic data is indicated in both panels (thick dotted line). The search limits and the $\pm 2\sigma$ spread of the search are estimated using 1000 randomly generated models with the same model parameterization employed in the inversion. The results from three runs are combined because of the tendency of the GA1 to exhibit some clustering of solutions. Note that the characteristics of the sampling by the GA of the parameter space are such that the σ spreads shown here do not give formal uncertainty estimates.

the $\pm 1\sigma$ spread of acceptable models in Fig. 5 shows excellent recovery of the original iasp91 model.

The scatter of acceptable models indicates the constraint and trade-offs of model parameters; this is analogous to the information given by the curving, elongated shape of the sample misfit surface in Fig. 1. The large scatter of acceptable solutions at the top of the crust, the top of the mantle, and below about 350 km in the mantle indicate regions of model space that are not well constrained by the inversion (Fig. 5). This lack of constraint has several causes. First, the lack of group velocity data at periods less than 10 seconds and greater that 350 seconds leads to the fanning of solutions at the top and bottom of the model; below about 550 km the scatter and 2σ spread of acceptable solutions is nearly as broad as for the search limits which indicates that the data imposes almost no constraint in this region. Secondly, the physics of the forward problem smooths and obscures information about the model. This process results in the scatter in models around the Moho, which reflects the physical limitations in resolving a first order discontinuity with fundamental mode dispersion data alone. A third cause of scatter in the models is the scatter in the data. This effect is present at all depths in the model and at all periods in the data and overlaps with the other causes of scatter mentioned above.

The diversity in acceptable models near the Moho shows that, in general, as the depth of the Moho discontinuity increases, the velocities in the lower crust and uppermost mantle also increase (Fig. 6). In consequence, most of the models with a shallower Moho do not have a high-velocity mantle lid. Also, the models with relatively low crustal velocities near the Moho tend to have relatively high velocity in the uppermost mantle. These correlations and trade-offs indicate the inability of the dispersion data to uniquely define the structure of the lowermost crust, the Moho and the uppermost mantle. However, the liberal parameterization and large number of solutions produced with the GA1 inversion gives an illustration of possible structures near the Moho and their trade-offs. In contrast, a method such as direct inversion may require, for example, a fixed crustal structure and Moho depth and will define only one solution for the deeper structure; a feature of this structure, such as a strong mantle lid, may be given a stronger emphasis in interpretation than if the correlations and trade-offs between such features in many acceptable models were known.

5 Discussion

In geophysical inversion the identification of all classes of acceptable solutions is important for a comprehensive and meaningful interpretation of the results. Achieving this goal requires information about the topography of the misfit

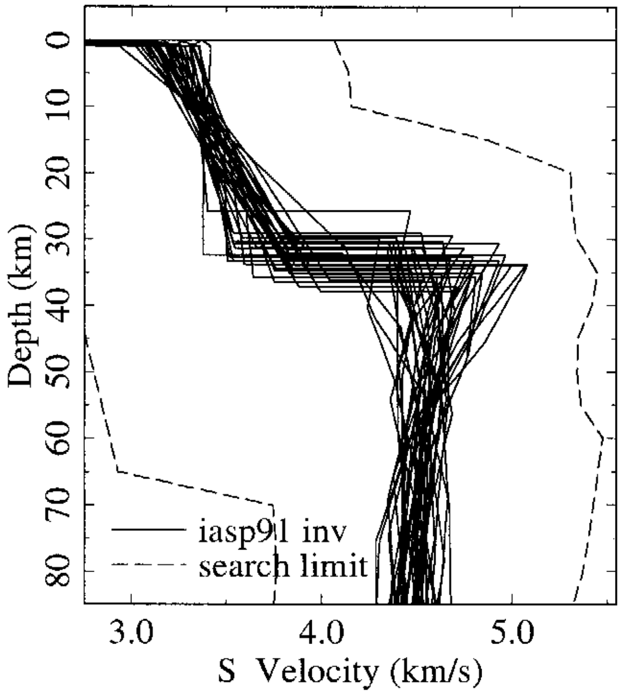

Fig. 6. Upper 85 km of the GA1 results for the synthetic iasp91 data from Fig. 5 showing a subset of acceptable models (misfit $E_s \leq 0.85E_d$) and the search limits (dashed lines).

surface, which may be rough and irregular for many interesting problems. We have presented a genetic algorithm tuned to find the regions of acceptable misfit on complicated misfit surfaces and have shown that this method performs better that a genetic algorithm tuned for rapid convergence or a random search for a complicated, 2-parameter misfit surface. We apply this algorithm to obtain many acceptable solutions in a multi-parameter geophysical inversion problem and show how this set of solutions indicates uncertainty and trade-off in the results.

The genetic algorithm is a powerful tool for problems for which there exists a single, clear, "optimum" solution (e.g. many geophysical problems *after* the use of many assumptions and physical approximations), or for problems where the identification of any one acceptable solution is adequate (e.g. process optimization in a factory), and, with careful tuning, it can identify multiple minima in complicated misfit surfaces. But methods based on the genetic algorithm may not be the most efficient for the goal of obtaining a representative set of acceptable solutions on complicated misfit surfaces. In addition, the genetic algorithm does not sample from the *a posteriori* probability density function (Nolte and Frazer, 1994) and so the solutions obtained cannot be formally incorporated in a general non-linear inverse theory (e.g. Tarantola and Valette, 1982). Other techniques from the field of "artificial intelligence" (e.g. simulated annealing (Kirkpatrick et al., 1983; Tarantola, 1987; Koren et al., 1991), or perhaps new algorithms developed with these goals on mind, may prove superior. A careful balancing of random global searching, convergent local searching, and the retention of multiple acceptable solutions may be necessary in any successful approach.

Acknowledgements. This research was supported by the Netherlands Organization for Scientific Research (NWO) through the Pionier project PGS 76-144.

References

Aki, K. and Richards, P.G., *Quantitative seismology*, W.H. Freeman, San Francisco, 1980.

Davis, L. (Ed.), *Handbook of genetic algorithms*, Van Nostrand Reinhold, New York, 1992.

Goldberg, D.E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA, 1989.

Goldberg, D.E. and Richardson, J., Genetic algorithms with sharing for multimodal function optimization, in Grefenstette, J.J. (Ed.) *Genetic Algorithms and their Applications, Proceedings of the Second International Conference on Genetic Algorithms and their applications*, Lawrence Erlbaum Associates, Hillsdale, NJ, 41-49, 1987.

Holland, J.H., *Adaptation in natural and artificial systems*, MIT Press, Cambridge, MA, 1992.

Jin, S. and Madariaga, R., Background velocity inversion with a genetic algorithm, *Geophys. Res. Lett.*, *20*, 93-96, 1993.

Keilis-Borok, V.I. and Yanovskaya, T.B., Inverse problems in seismology (structural review), *Geophys. J. R. Astr. Soc*, *13*, 223-234, 1967.

Kennett, B.L.N. and Engdahl, E.R., Travel times for global earthquake location and phase identification, *Geophys. J. Int.*, *105*, 429-466, 1991.

King, S., The genetics of mantle viscosity, abstract, in: Cadek, O. and Yuen, D.A. (Eds.), Geodynamics workshop in the Czech Republic, *Terra Nova*, *5*, 581, 1993.

Kirkpatrick, S., Gelatt Jr., C.D., and Vecchi, M.P., Optimization by simulated annealing, *Science*, *220*, 671-680, 1983.

Koren, Z., Mosegaard, K., Landa, E., Thore, P., and Tarantola, A., Monte Carlo estimation and resolution analysis of seismic background velocities, *J. Geophys. Res.*, *96*, 20289-20299, 1991.

Lomax, A. and Snieder, R., The contrast in upper mantle shear-wave velocity between the East European Platform and Tectonic Europe obtained with genetic algorithm inversion of Rayleigh wave group dispersion, *Geophys. J. Int.*, (in press), 1995.

Nolte, B. and Frazer, L.N., Vertical seismic profile inversion with genetic algorithms, *Geophys. J. Int.*, *117*, 162-178, 1994.

Press, F., Earth models obtained by Monte Carlo inversions, *J. Geophys. Res.*, *73*, 5223-5234, 1968.

Sambridge, M. and Drijkoningen, G., Genetic algorithms in seismic waveform inversion, *Geophys. J. Int.*, *109*, 323-342, 1992.

Sambridge, M. and Gallagher, K., Earthquake hypocenter locations using genetic algorithms, *Bull. seism. soc. Am.*, *83*, 1467-1491, 1993.

Scales, J.A., Smith, M.L., and Fischer, T.L., Global optimization methods for multimodal inverse problems, *J. Comp. Phys.*, *103*, 258-268, 1992.

Sen, M.K. and Stoffa, P.L., Rapid sampling of model space using genetic algorithms: examples from seismic waveform inversion, *Geophys. J. Int.*, *108*, 281-292, 1992.

Stoffa, P.L. and Sen, M.K., Nonlinear multiparameter optimization using genetic algorithms: Inversion of plane-wave seismograms, *Geophysics*, *56*, 1794-1810, 1991.

Tarantola, A., *Inverse problem Theory*, Elsevier, Amsterdam, 1987.

Tarantola, A. and Valette, B., Inverse problems = quest for information, *J. Geophys.*, *50*, 159-170, 1982.