

## Inversion of seismic refraction data using genetic algorithms

Fabio Boschetti\*, Mike C. Dentith<sup>‡</sup>, and Ron D. List\*\*

### ABSTRACT

The use of genetic algorithms in geophysical inverse problems is a relatively recent development and offers many advantages in dealing with the nonlinearity inherent in such applications. However, in their application to specific problems, as with all algorithms, problems of implementation arise. After extensive numerical tests, we implemented a genetic algorithm to efficiently invert several sets of synthetic seismic refraction data. In particular, we aimed at overcoming one of the main problems in the application of genetic algorithms to geophysical problems: i.e., high dimensionality. The addition of a pseudo-subspace method to the genetic algorithm, whereby the complexity and dimensionality of a problem is progressively increased during the inversion, improves the convergence of the process. The method allows the region of the solution space containing the global minimum to be quickly found. The use of local optimization methods at the last stage of the search further improves the quality of the inversion. The genetic algorithm has been tested on a field data set to determine the structure and base of the weathered layer (regolith) overlaying a basement of granite and greenstones in an Archaean terrain of Western Australia.

### INTRODUCTION

Many geophysical optimization problems are nonlinear and result in irregular objective functions. Consequently, local optimization methods, e.g., matrix inversion, steepest descent, conjugate gradients, are prone to trapping in local minima, and their success depends heavily on the choice of starting model. Thus, global optimization methods that can avoid this limitation are particularly attractive for geophysical applications.

Genetic algorithms are a search method suitable for the global optimization of irregular, multimodal functions. Starting

with a set of initial solutions, these algorithms progressively modify the solution set by mimicking the evolutionary behavior of biological systems, until an acceptable result is achieved. For a more detailed description of genetic algorithms the reader is referred to Buckles and Petry (1992), Cavicchio (1970), Davis (1987, 1991), and Goldberg (1989). Because of their initial random and progressively more deterministic sampling of the function domain, these algorithms offer the possibility of locating, with relative efficiency, the most promising areas of the solution space. They are able to solve nonlinear, nonlocal optimization problems, without the need of curvature information and consequently without the need for derivative calculations. This feature can be particularly useful because it allows for the use of fast approximate forward modeling for which no exact derivative may be available, with consequent reduction of the computation effort (see for example the description of the ray-tracing routine employed in this work). Also, because genetic algorithms are based only on direct space sampling, any kind of linearization of the problem to be solved is avoided, with the consequent elimination of the errors involved in such an approximation.

These factors make genetic algorithms particularly attractive for addressing complex real-world problems, and they are being used increasingly to address geophysical problems (Billings et al., 1994; Jin and Madariaga, 1993; Mathias et al., 1993; Stoffa and Sen, 1991). However, the large dimensionality involved in most geophysical optimization problems can reduce the efficiency of genetic algorithm search, both in terms of quality of the result and computational cost.

In this paper, we describe the development of a method, which we refer to as the pseudo-subspace method, to facilitate the search in high-dimensional spaces by reducing the problem dimensionality in the first stage of the process and by progressively increasing it once promising areas in the solution space have been discovered. The method is applied to a set of 45-dimensional problems. This dimensionality is small compared to conventional seismic tomography inversion but, to our knowledge, it represents one of the highest-dimensional

Manuscript received by the Editor December 9, 1994; revised manuscript received February 20, 1996.

\*Department of Geology and Geophysics and Department of Mathematics, The University of Western Australia, Nedlands, Perth WA 6907, Australia.

<sup>‡</sup>Department of Geology and Geophysics, The University of Western Australia, Nedlands, Perth WA 6907, Australia.

\*\*Department of Mathematics, The University of Western Australia, Nedlands, Perth WA 6907, Australia.

© 1996 Society of Exploration Geophysicists. All rights reserved.

genetic algorithm applications to seismic data published so far. We hasten to add that 45 dimensions should not be considered as some maximum, and we believe that larger inversions are feasible using the method outlined. However, we suggest that an effort be made to keep the problem dimensionality as low as possible, e.g., by subdividing the search domain into a number of smaller subdomains, to facilitate and speed up the genetic algorithm process. This concept will be discussed extensively in the application to the regolith problem in the section "Application to Field Data" at the end of the paper.

We also show that by coupling the genetic algorithm search with local optimization in the last stages of the inversion, solutions very close to the global minimum in the solution space can be found more readily. This in turn shows that genetic algorithms are effective in locating the valley in the solution space containing the global minimum, enabling the local search to descend such a valley toward the correct solution.

The potential of the method has been tested and compared with traditional methods on a set of seismic refraction synthetic data modeling different refractor geometries. In particular, a horizontal refractor at different depths in the calculation domain (Figure 1a and Figure 7a), a refractor incorporating a large vertical step to model a vertical fault (Figure 8a), and a horizontal layering including an isolated body characterized by low velocity (Figure 9a) have been modeled. These tests are discussed in the section "Synthetic Examples."

#### SEISMIC REFRACTION DATA

Synthetic first-arrival traveltimes data were generated using a line length of 8000 m with a group interval of 200 m and shots spaced at every 500 m, that results in a total of 615 rays/traveltimes that are inverted to define the velocity structure in the first few hundred meters below the surface. For the purposes of inversion, the slowness distribution in the subsurface is defined by a  $9 \times 5$  grid whose spacing is 1000 m in the horizontal direction and 100 m in the vertical direction. Thus, the model has a 45-dimensional solution space, formed from the slowness values at each of the grid nodes.

The ray-tracing routine used in this study is a modified version of the routine written by Asakawa and Kawanaka (1993). Their ray-tracing algorithm is a numerical approximation but it gave excellent results when compared to other widely used algorithms, e.g., SEIS83. The routine has the advantage of being extremely fast; being able to calculate the first-arrival time for the 615 rays in the test data set in approximately 2 s using a 33 MHz 486 PC. In the original code the slowness is assumed to be constant inside the cells defined by the grid nodes. The calculation is performed only at the border of the cells, resulting in a very fast computation. We modified the routine by introducing linear variation of the slowness between the grid nodes. Slowness at the edge of a cell is determined by linear interpolation using the slownesses at the adjacent nodes. Within a cell, the ray is assumed to travel with a slowness equal to the average of the slowness values at the two points at which the ray intersects the borders of the cell. The advantage of using a linear slowness gradient is that it allows the description of the subsurface slowness field with fewer nodes than required by a constant slowness cell, with a consequent reduction in the dimension of the problem to be solved.

#### GENETIC ALGORITHMS

In setting up a genetic algorithm, choices relating to the parameter representation and genetic operators must be made. Such choices are generally problem dependent and are crucial to the success of the algorithm since they can seriously affect its performance. A general outline of genetic algorithms will now be given, and we will discuss the choices we made to develop an algorithm to successfully invert seismic refraction data.

An introduction to the pseudo-subspace method, which we added to the genetic algorithm to facilitate the search in high-dimensional spaces will also be given.

#### Parameter representation

Most of the research and almost all the theoretical analysis of genetic algorithms has been applied to binary-coded algorithms, i.e., each individual member of the population is represented in binary form. However, recent experimental results show that real coded genetic algorithms outperform binary ones in most applications (Davis, 1991; Wright, 1991). For this reason, a real-coded genetic algorithm was implemented, i.e., an individual is represented by an array of real values, which in the application presented here corresponds to the seismic slowness values at the nodes of a regular grid.

A second aspect of parameter representation relevant to the seismic refraction problem is that it is sensible to constrain the slowness values in the solution set to lie within reasonable limits. The first generation of individuals, i.e., the members of the starting models, are randomly generated between these values, and during the evolutionary process they are constrained to remain within these limits. The method described here was developed as part of a project to define the near surface structure of Precambrian rocks in the Western Australian shield. Thus, our models (see the section "Synthetic Examples") have slownesses of around 0.18 s/km and the values of slownesses in the solution set were limited to lay between 0.14 s/km and 0.25 s/km, which is equivalent to velocities between 7.14 km/s and 4.0 km/s.

When detailed a priori information is available, it may be incorporated easily into the inversion process at this stage. Nodes for which the seismic velocity is known with accuracy may be kept fixed during the entire process, or allowed to vary between close constraints, thereby reducing the complexity of the search domain.

#### Genetic operators

A genetic algorithm involves three basic operators corresponding to the biological processes of selection, crossover, and mutation. Selection involves the choice of the individuals for the generation of offspring. Crossover is the method of combining (mating) two individuals to produce an offspring. Mutation is the random changing of some individual within the population. Each operator can be implemented in different ways. Choosing the correct combination is vital to the effectiveness of the algorithm.

**Selection.**—Selection is the principal driving force that pushes the genetic algorithm population toward a final solution. It works by first assigning a measure of fitness to each

individual, as related to the value of the objective function at the corresponding point in the solution space. Then, using this assigned measure of fitness, a rule for selecting which individuals to use to create the next generation must be chosen. Thus, the selection operator can be implemented in many different ways and its influence can be varied. Implementations resulting in a strong selection pressure allow for fast convergence, but have the drawback of concentrating most of the population (and in few generations all of it) in a small part of the solution space. Because of such limited space sampling there is a greater risk of becoming trapped in a local minima (a phenomenon known as premature convergence). Implementations resulting in a low selection pressure will tend to decrease such a risk, but at the cost of a much slower convergence. For the algorithm to be effective, a balance between these two effects must be found.

A number of different selection implementations have been proposed in the literature (Davis, 1991; Goldberg, 1989) and two of the most common choices have been implemented: linear normalization selection, which has a high selection pressure (Davis, 1991; Goldberg, 1989), and parent selection, which has a low selection pressure (Cavichio, 1970).

In linear normalization selection, an individual is ranked according to its fitness, and then it is allowed to generate a number of offspring proportional to its rank position. Using the rank position rather than the actual fitness values avoids problems that occur when fitness values are very close to each other (in which case no individual would be favored) or when an extremely fit individual is present in the population (in such a case it would generate most of the offspring in the next generation). This selection technique pushes the population toward the solution in a reasonably fast manner, avoiding the risk of a single individual dominating the population in the space of one or two generations.

In parent selection, an individual is allowed to generate only one offspring regardless of its fitness. All the individuals are mated randomly and through crossover each couple creates two offspring. If the fitness of offsprings is better than those of the parents, they are substituted for them in the population. If not the parents keep on living. With this method the selection pressure is low and diversity is kept in the population because, as has been empirically demonstrated (Mahfoud, 1992), the offspring are statistically the closest points currently present in the population to the parents they have been generated from. This means that an individual is substituted by the (statistically) closest point if its fitness is better, a process that suggests the idea of a "stochastic parallel optimizer." A good summary of the advantages of such an operator can be found in Mahfoud (1992).

**Crossover.**—Different kinds of crossover have been discussed in the literature. We implemented and tested four of them: two-point crossover, multipoint crossover, uniform crossover, and uniform crossover with averaging of the parameters. These different methods basically describe how the two parents are combined. For a more complete discussion, the reader is referred to Davis (1991) and Goldberg (1989). Among these, uniform crossover was the most successful, and it has been incorporated in all the implementations tested in this study. We implemented uniform crossover in the following way: two individuals are chosen randomly and a random

number  $n$  between 1 and the number of parameters is selected; then  $n$  random gene locations are chosen and the floating point values of the parameters at such locations are swapped between the two individuals.

The proportion of a population that is mated after each generation is defined by the crossover rate. The experiments performed suggested an optimum tuning if the crossover rate is 0.8.

**Mutation.**—This operator parallels the biological process of mutation. Here it has been implemented in a traditional way, i.e., by randomly changing some parameter values in selected individuals. The probability of this event must be kept very low to reduce the chance of removing good individuals currently present in the population. In the tests discussed below, the mutation rate has been fixed at 0.01. Mutation has been implemented by assigning a parameter that is a random value between the minimum and the maximum velocity slowness limits.

### Pseudosubspace method

In a recent paper Williamson (1990), describes the inversion of seismic reflection data using a multistaged approach to local optimization, in which the subsurface is represented by a series of constant velocity cells. During the course of the inversion, a progressively smaller cell size is used, with a consequent increase in the resolution of the model. We have used a similar approach in defining the slowness grid that determines the model. However, in our models the subsurface distribution of slowness is defined at the nodes of a grid between which linear variation is assumed. Because of the linear slowness approximation, even a very coarse grid with only four nodes at the corners of the model still allows for a reasonable description of the slowness field and hence a reasonably accurate travel-time computation. With just four nodes, a genetic algorithm can rapidly explore the resulting low-dimensional solution space, while still being able to obtain a rough picture of the overall solution. Such a simplification would not be possible with constant velocity cells.

In the very first generations of our algorithms, we used a slowness grid defined by only a small number of points. Even in the corresponding low-dimensional solution space a genetic algorithm, rather than a traditional local optimizer, was used to avoid the need of any initial model in the inversion. Within a few generations, the algorithm successfully defined an approximately correct slownesses at these nodes. There is no need to run the algorithm for many generations at this stage because finding an exact solution is very unlikely. Thus, only a small population is needed at this stage.

For subsequent generations, the grid interval is halved, and each new grid point is assigned a value equal to the mean of the adjacent nodes. Again the algorithm is run, in what is now a higher dimensional solution space, with an initially larger population for a few generations. This process is repeated until the size of the grid reaches some predetermined limit, whereupon the algorithm is run for a larger number of generations, and with a larger population, until an acceptable convergence is reached. We found that an initial population of 20 individuals and a final one of 100 gave optimal results.

This "zooming" technique resembles a subspace method used in large scale optimization problems (Kennett et al., 1988;

Oldenburg et al., 1993). In this paper, we refer to such a technique when applied to genetic algorithms as the pseudosubspace method. Further details about the technique can be found in Boschetti et al. (1995a).

### COMPARISON OF DIFFERENT GENETIC ALGORITHMS

To explore the effects of the different choices described above, a series of genetic algorithms have been compared on a data set obtained as described above in the section "Seismic Refraction Data" using a horizontally layered synthetic model with linearly varying slowness in the vertical direction (model details shown in Figure 1a). The model is characterized by slowness values that allow for the occurrence of both diving and refracted rays and a uniform distribution of raypaths (Figure 1b).

Three different specific choices were tested: linear normalization selection, linear normalization selection with pseudosubspace method, and parent selection with pseudosubspace method. Each test consisted of five runs with different random seeds, i.e., with different randomly chosen initial populations. For each implementation, the best solution found by the algorithm in these five runs, together with the convergence curve is presented (Figures 1c–1h).

A comparison of Figures 1c and 1e shows an improvement in the performance of the linear normalization selection when the pseudosubspace method is used. The essential features of the test model are reproduced with only significant differences occurring at the top right of the model in an area of poor ray coverage. This suggests that the algorithm benefits from the inclusion of the pseudosubspace method.

The quality of the solution obtained with the parent technique (Figure 1g), even though it is obtained at a much higher cost in terms of the number of required function evaluations (Figure 1h) is not the best of the three solutions. This result suggests that efforts to maintain diversity within the population will not necessarily improve the quality of the final solution.

Clearly, the genetic algorithm with linear normalization selection including the pseudosubspace method outperformed the other implementations, and it has been used in all the experiments described in the rest of the paper. In the initial generations, using a coarse grid, the algorithm still appears to be able to locate the favorable area of the solution space in which to concentrate the more detailed and time consuming subsequent analysis.

### LOCAL SEARCH METHOD

The convergence curves in Figure 1 illustrate a well known characteristic of genetic algorithms—they are poor optimizers (DeJong, 1993). They show a very fast initial convergence, followed by progressively slower improvements. In fact such behavior is common to many optimization techniques, but is of particular concern in genetic algorithms. The form of the curve suggests that the algorithm should be stopped when an approximate solution has been found, because further improvements may be very costly. Improvements to the genetic algorithm solution can only develop through crossover or mutation, i.e., random events. Depending on the exact form of the solution space, further improvements to high-fitness solutions can be a rare event.

In the case of the best solution found by the genetic algorithm with the pseudosubspace method, small errors are present in

each parameter. This is because of the fact that the traveltime misfit caused by slightly wrong parameters may be reduced by introducing small errors at adjacent nodes. Improvement in the solution at this stage may be obtained only by the accurate tuning of most of the parameters at the same time, which cannot be achieved efficiently by a genetic algorithm process. However, further improvements to the inversion method can be obtained by combining the genetic algorithm with a local optimizing method. In the initial stages of the inversion, the space-sampling properties of the genetic algorithm is used to direct the search to the region close to the global solution. The solution so obtained can be further improved using a local search method such as a hill-climber algorithm.

Other hybrid genetic algorithm implementations have already been applied successfully to geophysical problems. A related approach has been used in Mathias et al. (1993), although in that case, a pseudosubspace method was not used, and local search was applied periodically to more than one solution obtained by the genetic algorithm. Also, Sen and Stoffa (1992) showed how the performance of genetic algorithms can be enhanced greatly by importing some elements from simulated annealing process.

A number of algorithms for the local optimization of functions in multi-dimensional spaces are described in the literature (Dixon and Szego, 1978; Press et al., 1992). The choice of the algorithm to use is problem-specific and often experimental trials are required. We tested four local search methods: the downhill SIMPLEX method, Powell's method, the conjugate gradient method (Press et al., 1992), and the local search routine from the Hill-climber method described in De la Maza and Yurez (1994). Of these, the SIMPLEX method, Powell's method, and the local search routine from the Hill-climber method do not use gradient information, while the conjugate gradient method does. The approximate calculation performed by the ray-tracing routine does not give gradient information. Consequently, the numerical calculation of the required derivatives is time consuming and the conjugate gradient method is considerably slower than the other methods.

After extensive tests, the SIMPLEX code was found to be the most reliable and stable algorithm and accordingly was used in the successive tests.

Figure 2 illustrates the result from the local optimization of the best individual from the genetic algorithm with the pseudosubspace method, together with its convergence curve. The SIMPLEX algorithm was able to improve the solution, minimizing the squared error up to a level of approximately zero misfit. The solution is almost indistinguishable from the original model (Figure 1a). All the layers are well reconstructed both in terms of velocity values and vertical position, with only a minor departure from the original at the middle left-hand side of the solution. Clearly, the two-staged procedure linking the genetic algorithms' global search with the local optimization has been very effective.

### EFFICIENCY

The performance of an inversion process is not only evaluated in terms of the accuracy of the solution it produces, but also in terms of its cost. In the case of a genetic algorithm, such cost is measured basically in terms of the number of function evaluations (ray tracings in this case) required to obtain

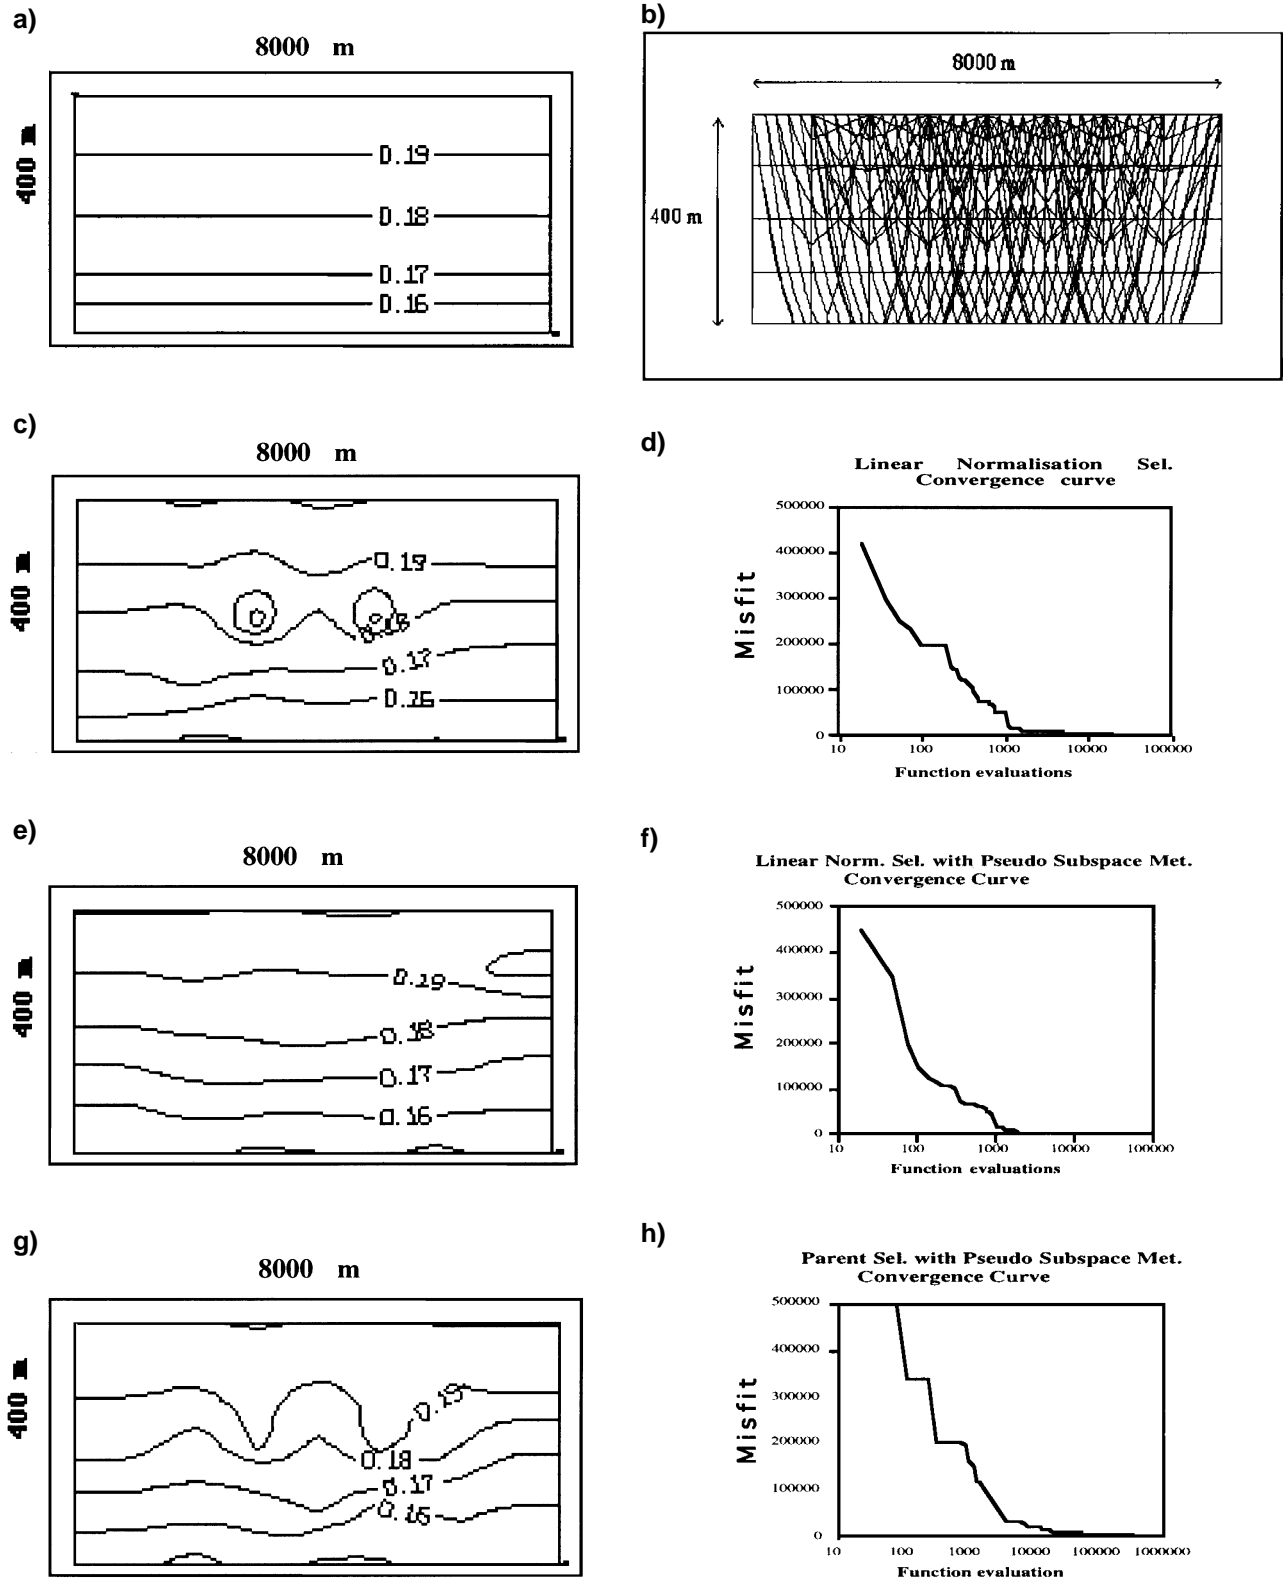


FIG. 1. Results from the comparison of different genetic algorithm implementations on a horizontally layered synthetic test with linearly varying slowness in the vertical direction. The (a) synthetic model together with (b) the ray diagram are presented. Both final models and convergence curves are shown for linear normalization selection with genetic algorithm (c) and (d), linear normalization selection with the pseudosubspace method (e) and (f), and parent selection with the pseudosubspace method (g) and (h). The contours represent the value of the seismic slowness in the calculation domain.

the final image. In general, genetic algorithms are considered a relatively expensive method, to be used only in cases where local procedures proved to be unsuitable. Raiche (1994) considers this to be particularly true for geophysical applications. However, our results suggest that this is not always the case.

Clearly, an important decision in terms of the efficiency of the two-stage optimization techniques described above, is when to terminate the genetic algorithm and implement the local search. We now implemented the genetic algorithm to output the best solution every 50 generations after the final grid configuration has been achieved. The best solutions after 100, 150, 200, and 300 generations for the genetic algorithm with linear normalization selection incorporating the pseudosubspace method are presented in Figure 3. These pictures should be compared with the result at the end of the process, i.e., after 400 generations, already shown in Figure 1e. It is clear that an acceptable solution is obtained after only 100 generations, with only minor improvements occurring up to 300 generations after which variations are negligible.

An even more interesting test is to improve, by local search, such solutions and compare the final results. In Figure 4, the results obtained using the local search, initialized with the solutions after 100 and 150 generations, are shown. The differences between the results are minimal, even when compared with the result from the solution locally improved after 400 generations (Figure 2a). This shows that genetic algorithms have discovered the "good" valley in the solution space very rapidly. Furthermore, the computation effort, i.e., the number of function evaluations, for the local improvement of the 100, 150, and 400 generations solutions by the SIMPLEX code is almost the same (Table 1). Clearly, the space sampling of the genetic algorithm after the 100th generation has not been very productive. After a solution located in a good valley has been found, there is little advantage in using the genetic algorithm for further space sampling because more effective results may be obtained by the local search.

It is very difficult to generalize about when to stop the genetic algorithm and begin the local search. A number of alternatives are available, for instance, waiting for no improvements to happen for a few generations, or waiting for variations in the misfit to fall below a predetermined value. However, choosing such a threshold can be very difficult and problem specific. Because the genetic algorithm process is nondeterministic, it is possible

to have many generations without any improvement, followed by very rapid improvements as new domains are discovered by crossover or mutation. Research to define effective criteria suitable for controlling the inversion of the seismic data is on-going.

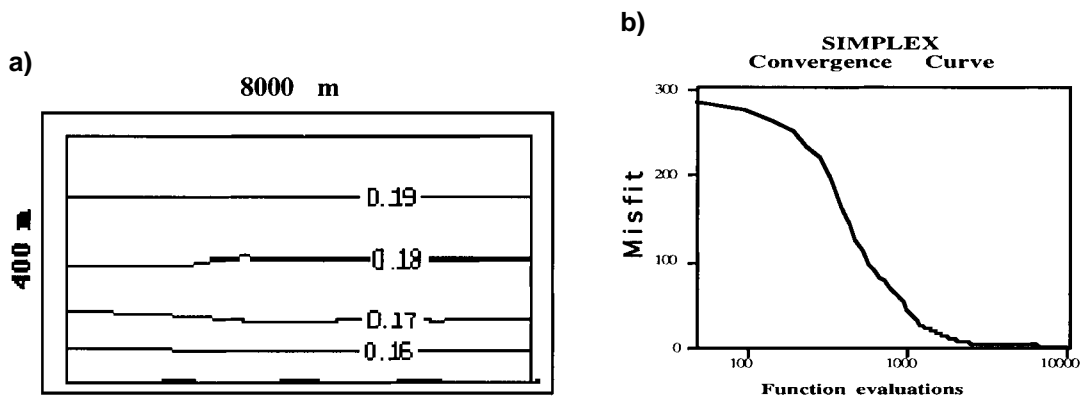
As shown in Table 1, in 100 generations the genetic algorithm performs approximately 6000 function evaluations to locate the region in the solution space containing the global minima. Such a computation effort is comparable to that required by the local search to further improve the solution. Thus, a genetic algorithm applied to this problem should be considered good not only in terms of accuracy but also in terms of the computation effort. Inversion of the synthetic data employing only the local optimizer starting from a random point in the solution space (i.e., under the same conditions used to initialize the genetic algorithm process) were attempted and resulted in the process getting trapped in local minima very far from the global solution. Even in these cases, the process required a number of iterations larger than required by the genetic algorithm to find an acceptable solution.

### STABILITY

As described above, each implementation of the genetic algorithm was tested on five different random seeds (initial populations). The stability of the proposed inversion procedure can be illustrated by comparing the results of the five different tests. Since the genetic algorithm process is nondeterministic, we expect the solutions to differ. The best solution is illustrated in Figure 1e. In Figure 5, the four other results obtained after 150 generations using the linear normalization genetic algorithm with the pseudosubspace method are presented. In two cases, there has been good reconstruction of the layering with minor anomalous features (Figure 5a and 5b), while in the remaining

**Table 1. The number of function evaluations performed in 100, 150, and 400 generations by the genetic algorithm and the number of evaluations required by the local search to optimize these models.**

Generations	GAs	Local search
100	6316	10403
150	10814	10620
100	24286	10355



**FIG. 2.** Result from the optimization of the genetic algorithm solution with the (a) SIMPLEX algorithm and (b) convergence curve.

two examples major anomalies are present (Figure 5c and 5d). However, the results obtained after local optimization using the SIMPLEX code are in acceptable agreement with the synthetic image in every case (Figure 6).

This result suggests that in all five tests the genetic algorithm with the pseudosubspace method has been able to find good valleys in the solution space. All these solutions were sufficiently close to the global minimum for the local optimizer to further reduce the error. All the tests were initialized randomly but converged to the same solution. This suggests that only one global solution is present in the application, i.e., no cases of clear ambiguity are manifested, although obviously the number of tests is small.

Given the relative low cost of the genetic algorithm process, as a general rule it might be advisable to perform more than one genetic algorithm inversion with different, randomly chosen, initial populations and locally improve the best solutions found. These runs would help identify where more than one solution to the problem is possible.

Similarly, the solution obtained from the local search does not necessarily have zero misfit. In such a case it is impossible to discriminate if such solution represents a local minimum or a point whose surrounding space topography is so complicated that it cannot be improved further even by the local search. In one of the tests we performed, this has actually been the case. Calculating the function values along the direction connecting

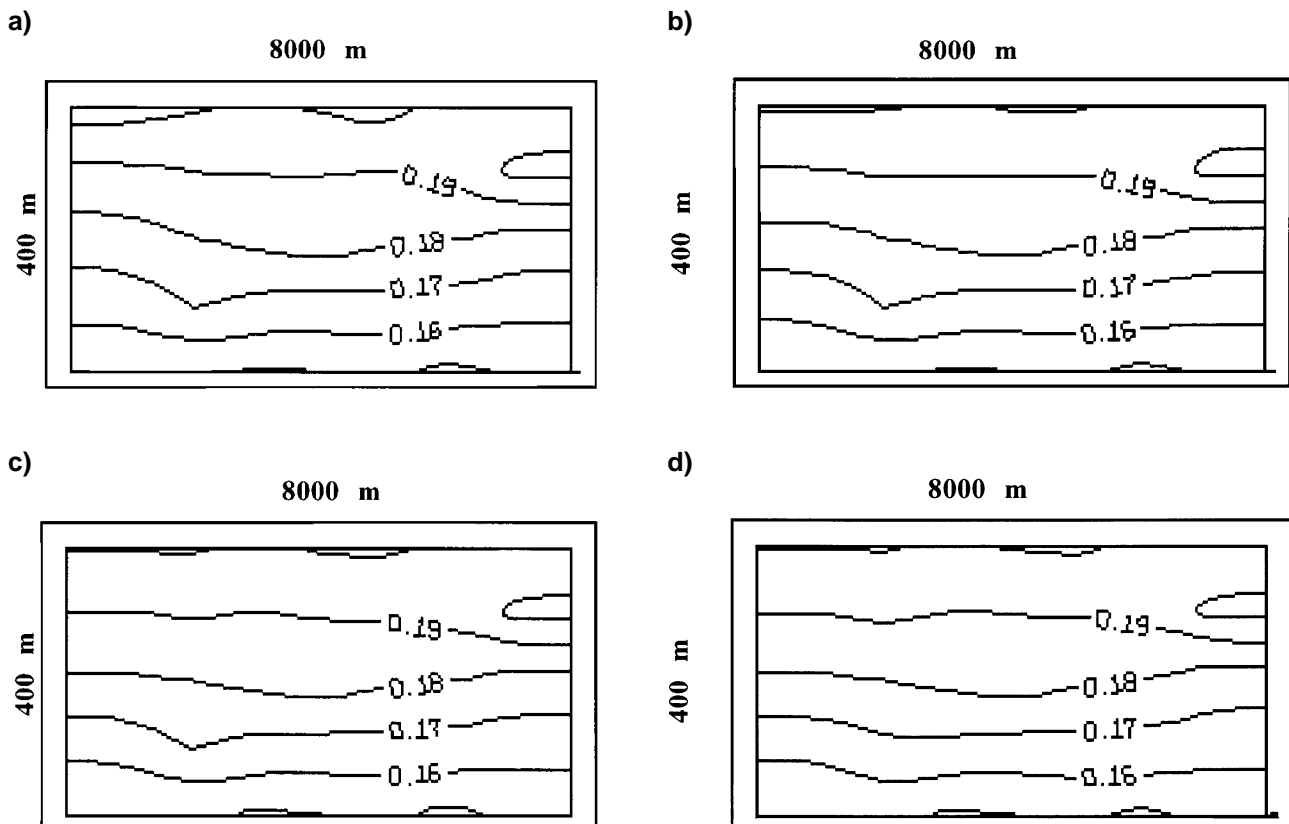


FIG. 3. Best solutions after (a) 100, (b) 150, (c) 200, and (d) 300 generations for the genetic algorithm with pseudosubspace method. An acceptable solution is already obtained after only 100 generations.




FIG. 4. Models resulting from local optimization of the best solutions obtained after (a) 100 and (b) 150 generations of genetic algorithm with pseudosubspace method. Only minor differences may be found between the two models.

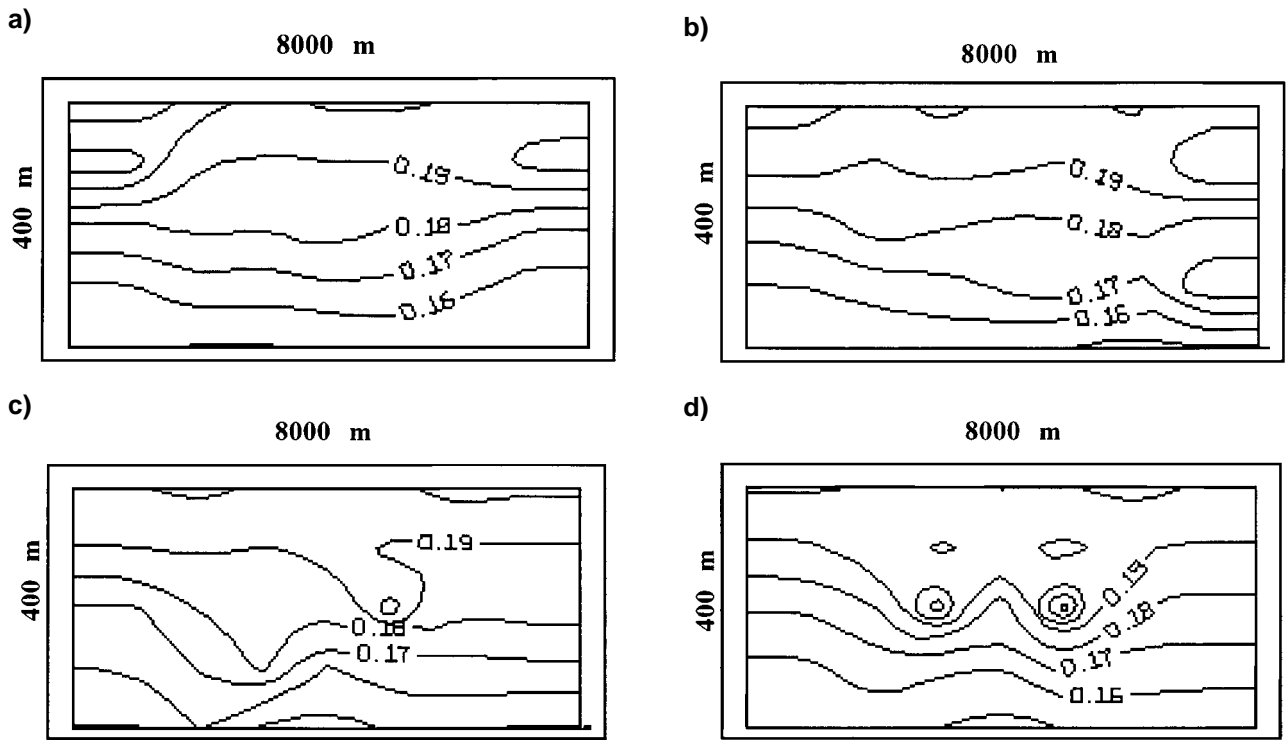


FIG. 5. Results obtained using different starting populations after 150 generations using genetic algorithm with pseudosubspace method. These results should be compared to the synthetic model in Figure 1a. Solutions in (a) and (b) are comparable to the synthetic image, however (c) and (d) contain discrepancies.

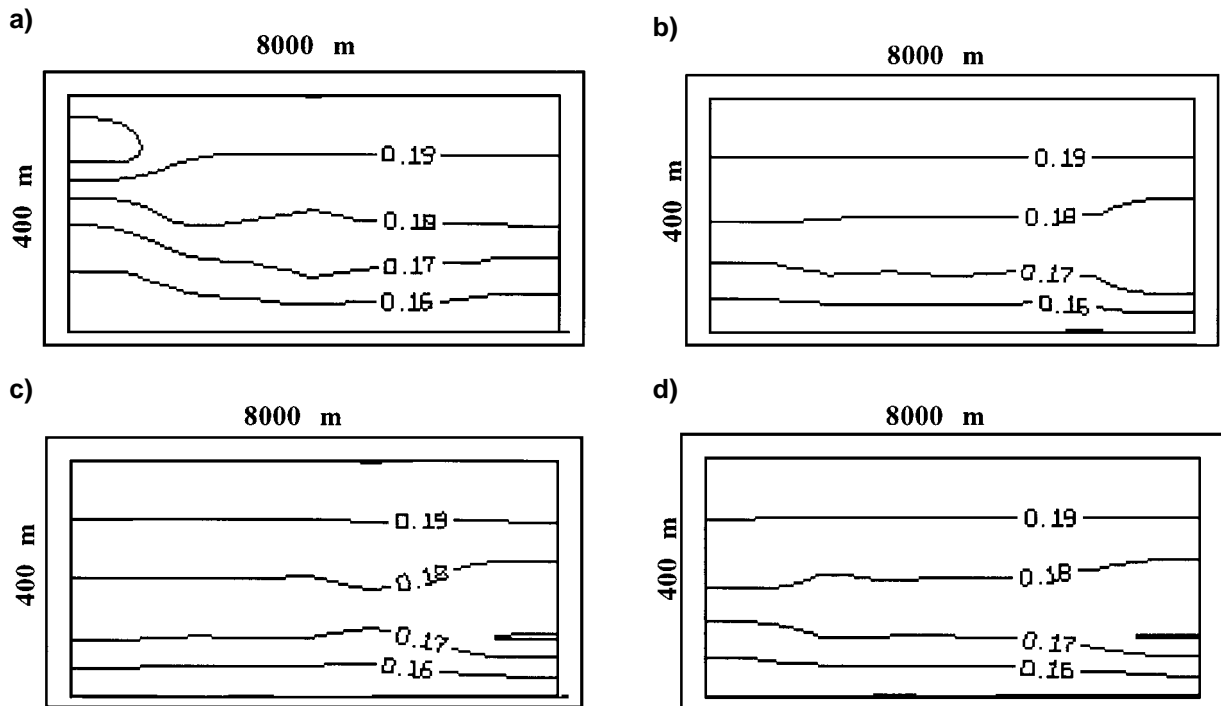


FIG. 6. Results from local optimization, using the SIMPLEX code, of the four solutions shown in Figure 5 from the genetic algorithm with pseudo-subspace method. All the reconstructions are in good agreement with the synthetic case (Figure 1a).



a solution obtained from a local search to the synthetic image, showed values to be constantly decreasing even at very small steps. Thus, the local search had failed to complete the exploration of the good valley. It was a lucky trial: not all the possible directions can be examined in a 45-dimensional space, and the existence of curved valleys connecting the local solution to the global one usually cannot be excluded. Thus, the assumption that the output from a local search must be a local minimum should not be taken for granted. However, in the test examples the difference between the local optimizer solution and the global minimum apparently is not large.

### SYNTHETIC EXAMPLES

The two-stage inversion procedure described above has been tested further on three other synthetic data sets: a shallow horizontal refractor (see Figure 7), a refractor incorporating a step (Figure 8) and an isolated buried body with anomalous velocity (Figure 9). These noise-free data sets allow us to test the method without the complications caused by the presence of noise. The same conditions used in the flat horizontal layers inversion example have been applied in all these tests. All the synthetic images have been created with a  $9 \times 5$  slowness grid resulting in a 45-dimensional inversion problem.

Notice that the only a priori information used in these inversions is contained in the genetic algorithm implementation, i.e., slowness field parameterization and slowness value constraints. This information is the same for all the different tests, and no specific a priori information or starting model is needed in the individual runs.

The pseudosubspace method has been implemented in three stages. In the first stage, the inversion is performed on a  $3 \times 2$  grid, whose spacing is 4000 m in the horizontal and 400 m in the vertical direction. The node spacing is then halved at each stage. Consequently, in the second stage, the inversion is performed on a  $5 \times 3$  grid and in the last stage it reaches the final configuration of  $9 \times 5$  nodes. Each test consisted of five runs with different random seeds, and the best solution found by the genetic algorithm has been improved further by the local search.

**Shallow horizontal refractor.**—This configuration is similar to the horizontal flat layer case previously analyzed, except that the refractor is raised to the middle of the model and the slowness below it kept constant (Figure 7a). It approximates the case of a shallow layer whose velocity increases with depth overlying a layer with a constant and higher velocity. This configuration is particularly interesting because it tests the ability of the inversion process to recover the depth of the refractor correctly. Intuitively, it is expected that many possible solutions to the problem exist with increased depth to the refractor being compensated by a higher velocity for the overlying area and vice-versa. The solution obtained by the genetic algorithm is shown in Figure 7c and after local optimization in Figure 7d. The match between true model and the final solution is very good, with the refractor correctly located in the vertical direction and perfectly horizontal.

From the raypath's diagram (Figure 7b), we see that no rays penetrate to the lower part of the model, and hence no real solution can be expected from such area. The slowness values

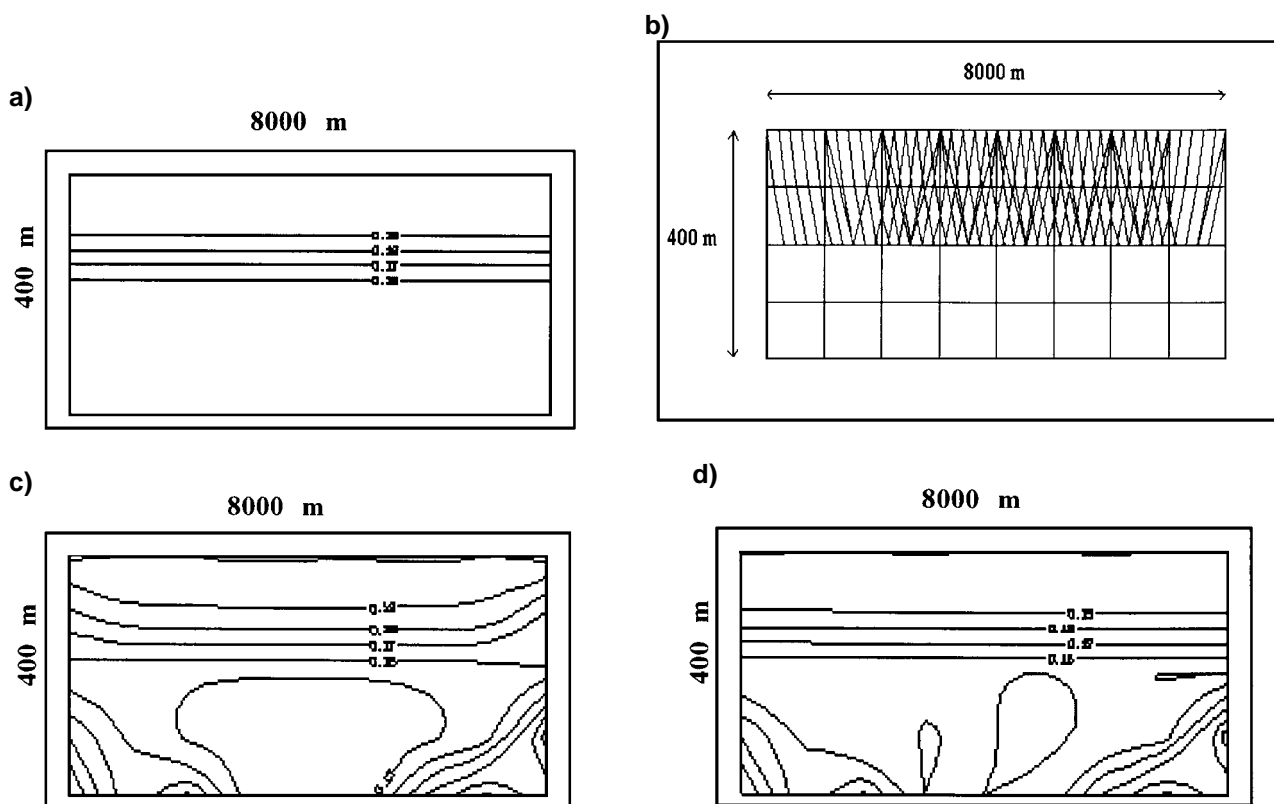


FIG. 7. Inversion of a synthetic data set simulating the presence of a shallow horizontal refractor. (a) Synthetic model, (b) ray diagram, (c) result obtained using the genetic algorithm with the pseudosubspace method, and (d) result obtained after local optimization.

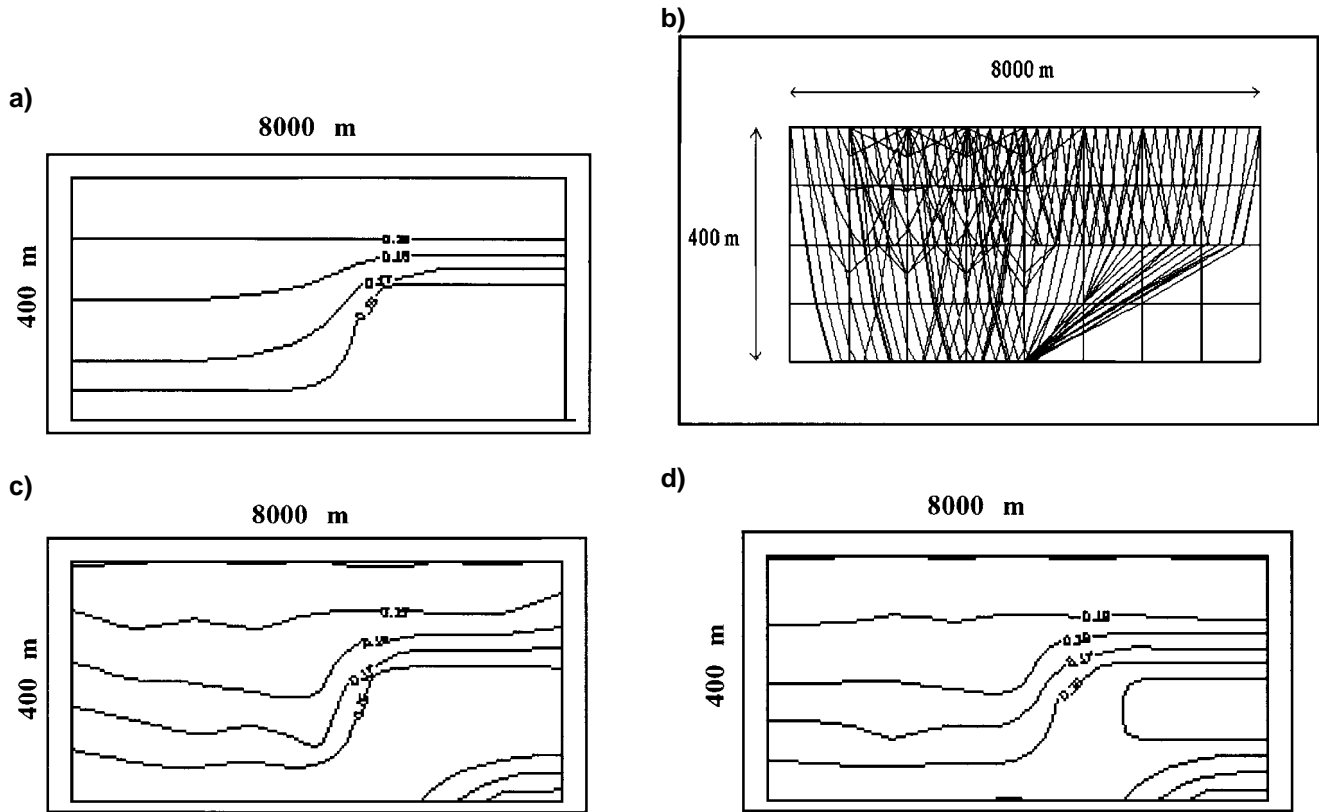


FIG. 8. Inversion of a synthetic data set simulating the presence of a step refractor. (a) Synthetic model, (b) ray diagram, (c) result obtained using the genetic algorithm with the pseudosubspace method, and (d) result obtained after local optimization.

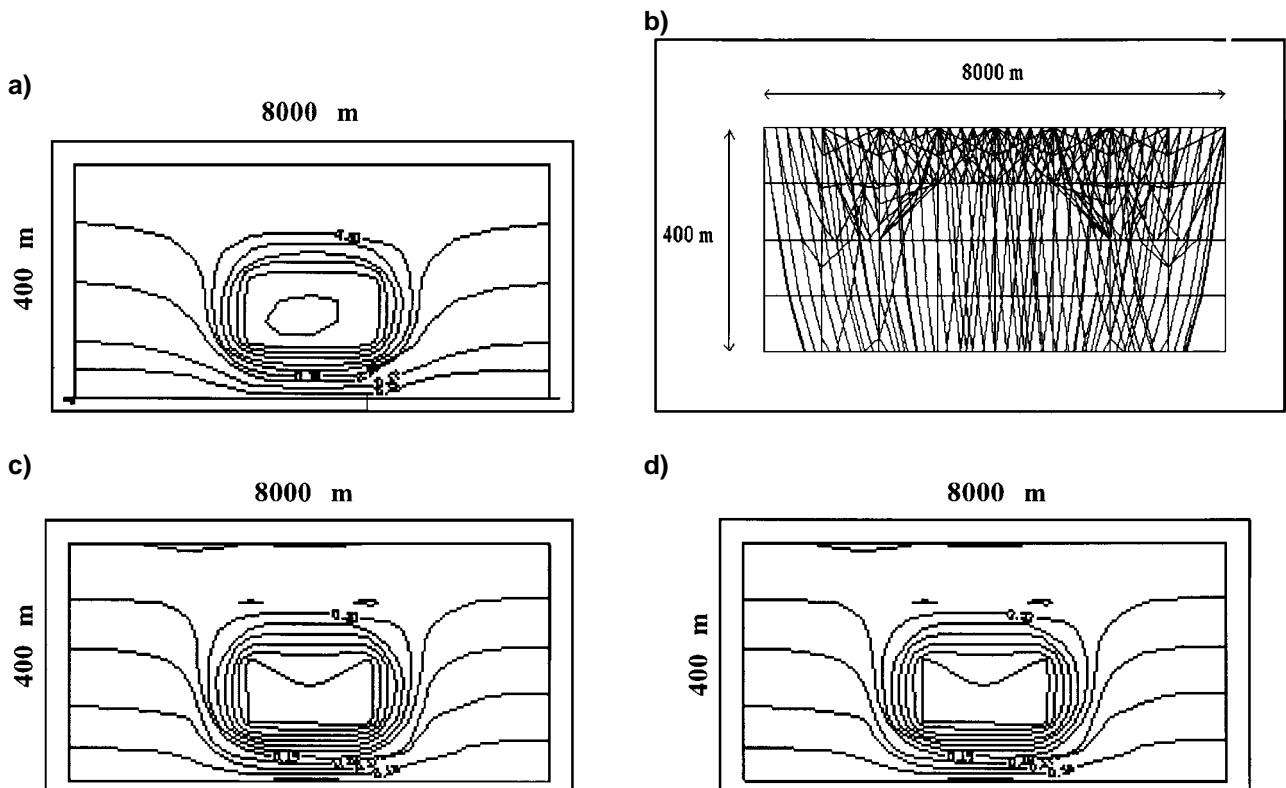


FIG. 9. Inversion of a synthetic data set simulating the presence of a concealed body. (a) Synthetic model, (b) ray diagram, (c) result obtained using the genetic algorithm with the pseudosubspace method, and (d) result obtained after local optimization.

are simply artifacts of the genetic algorithm. Note further that as the rays are concentrated in the area above the refractor it is not surprising that the slowness in this area is well recovered.

**Step refractor.**—This model is a development of the previous one with a step incorporated into the refractor (Figure 8a). As it can be seen in Figure 8b, the model allows rays to dive through the refractor. Again, the results after the two stages of the inversion procedure are presented in Figure 8c (solution from genetic algorithm) and Figure 8d (from the genetic algorithm plus local optimizer). The final solution is accurate in the right-hand and left-hand side of the model; however, small errors are present in the central region. Overall, the final model is very close to the original one, with the best fit in areas where there are most rays, especially when these vary in orientation.

**Concealed body.**—In the final model, an isolated body with low velocity is superimposed on a vertical velocity gradient (Figure 9a). As shown in Figure 9b, most of the rays run along the top and the borders of the body. In the resulting solution, the outline of the body is reconstructed almost perfectly (Figure 9d) and this allows, as a consequence, excellent recovery of the vertical velocity gradient. Note that the result obtained using the genetic algorithm (Figure 9c) is so close to the starting model that the local search was not able to improve it further.

#### APPLICATION TO FIELD DATA

The algorithm presented in this study has also been applied to a real data set to test its efficiency in practical situations. Refraction data from a seismic survey near the Nevoria Gold Mine, Southern Cross, Western Australia have been used. The area is characterized mainly by greenstones and granitoids overlain by a thick weathered cover (regolith). Such weathered profiles are of great exploration interest and they have recently been the target of various geophysical studies to deduce which is the most viable method to map their base and internal structure. For such a purpose, the efficiency of the seismic refraction method has been tested. During that survey, shots were fired and data recorded at 151 stations 25 m apart. The profile lies across an almost vertical contact between greenstone and granitoids that reaches the surface near shot  $n$ . 45 in Figure 10. The first arrivals were analyzed with the plus-minus method (Hagedoorn, 1959) in a study performed, in collaboration, by the University of Western Australia and Curtin University in Perth, Western Australia (see Dentith et al., 1992). The results are presented in Figure 10.

The same data set has then been inverted with the use of a genetic algorithm with the pseudosubspace method. However, in this case no local optimization has been used. The rationale behind such choice lies in the fact that improvements to the genetic algorithm solution through the use of local optimizers can be obtained only by a very small decrease in the error misfit. This can be seen from the convergence curve in Figure 2b. The success of such a process in different synthetic tests is very important from a theoretical perspective because it shows that a genetic algorithm search using the pseudosubspace method is successful in detecting the valley containing the global minimum in the solution space. However, on real data sets the refinements on the error misfit required by the local optimizer may fall well below the limitations imposed by the presence of

noise. In such circumstances, the local optimization of the genetic algorithm solution is obviously useless. A detailed study of the influence of noise in the tomographic inversion of refraction data, also aimed at establishing the maximum amount of noise the algorithm can tolerate, is currently on-going. However, the synthetic tests we presented in this study show that the quality of the genetic algorithm solution, even without local optimization, is satisfactory and accordingly its performance has been tested in a field data application.

In the experiment described here, the inversion has been carried out on a domain 3750 m long and 160 m deep, with a resolution of 75 m in the horizontal and 40 m in the vertical direction, resulting in a  $51 \times 5$  nodes grid. To reduce the problem dimensionality, the domain has been divided into seven small subdomains ( $9 \times 5$  nodes), each overlapping for two nodes at the lateral borders. Figure 11 helps to describe the process. In this way, the problem has been reduced to the same dimensionality as that of the synthetic tests above. Once all the seven subdomains were inverted, the nodes at the extreme of the single subdomains were disregarded because of the low ray density in such areas, and the remaining nodes linked to obtain the global  $51 \times 5$  node solution.

The solution is shown in Figure 12. The result below the refractor is not reliable because of the limited amount of rays diving into such area. To facilitate the interpretation, the position of the refractor has been marked by a thick line and the slowness values below the line have been muted (notice that the refractor position is discretized at the nodes location). Such operation has been performed with the use of the ray-tracing




FIG. 10. Result of the analysis of the real data set recorded close to the Nevoria Gold Mine, Southern Cross, Western Australia, with the plus-minus method. A schematic description of the geology of the area (obtained through lithologic analysis of RAB chips) is also given: a weathered layer (c) overlays a basement formed by (a) greenstones and (b) granitoids divided by an almost vertical contact.




FIG. 11. Schematic description of the process performed to subdivide the search domain into seven small subdomains. Notice that such subdomains overlap at the lateral borders to disregard the results obtained in areas of poor ray coverage.

routine: each node slowness value is altered by a small amount both in negative and positive signs, if the error misfit does not undergo any change it means that no rays dive in an area close to the node and such node should be muted.

The result in Figure 12 agrees well with the previous analysis obtained by the more traditional method presented in Figure 10. Notice that a detailed comparison of the two pictures is not possible because of both the approximations involved in the plus-minus method and the coarse parameterization in the genetic algorithm inversion. Nevertheless, the two images are characterized by the same main features—the steep contact toward the left part of the picture as well as the acute undulation in the refractor position in the center, and the smoother undulation at the right-hand side of the picture. Some disagreement is present immediately to the left-hand side of the area where the refractor reaches the surface (shots 35–40 in Figure 12). This may be caused by the coarse parameterization in the genetic algorithm inversion.

Eventually, an attempt to invert the overall domain in a single process has been performed. The pseudosubspace method has been implemented in such a way as to divide the search into three stages of different dimensionality. In the first stage, the search is performed on a grid with 13 nodes in the horizontal direction and two in the vertical. Then the dimensionality is increased to  $25 \times 3$  and eventually to the final configuration of  $49 \times 5$  nodes, approximating the  $51 \times 5$  dimensionality of the

final image in Figure 12. The result is presented in Figure 13. It can be noticed that the solution lacks definition in the reconstruction of the refractor position. The solution space is too large to be searched accurately, and some details cannot be resolved. Furthermore, errors are left in some parts of the picture. The deep undulation in the middle part of the image is substituted by a flat refractor with a low velocity anomaly over it. This kind of error was also present in the synthetic test inverted with a genetic algorithm without subspace search (see Figure 1c) and is caused by details that cannot be resolved in search domains that are too large. It is evident that the efforts to keep the problem dimensionality as low as possible by subdividing the search domain into small subdomains have been particularly beneficial.

Nevertheless, it should be noted that the solution in Figure 13 is not totally unsatisfactory, as the contact between greenstones and granitoids is well recovered as well as reasonably good in the reconstruction of the average refractor position. Also, the stratification overlying the refractor resembles the one presented in Figure 12. This result suggests that problems whose dimensionality is larger than the  $9 \times 5$  nodes synthetic tests presented above should not be considered completely beyond the methods potential. This confirms the results obtained with physical model data in which tomographic problems up to 105 dimensions were inverted satisfactorily (see Boschetti et al., 1995b).




FIG. 12. Inversion of the real data set with the genetic algorithm with the pseudosubspace method. The image has been obtained by subdividing the search domain into seven small subdomains and by linking the individual results. The steep contact and the undulations in the refractor position agree well with the previous analysis obtained with the plus-minus method presented in Figure 10.




FIG. 13. Inversion of the real data set with the genetic algorithm with the pseudosubspace method without the subdivision of the search domain. The solution should be compared with Figure 12. Errors are present in the refractor position and in the presence of a low velocity anomaly in the middle part of the picture.

## CONCLUSIONS

Geophysical problems tend to have larger dimensionality than most of the optimization problems where genetic algorithms have been applied traditionally. This means that genetic algorithms are not usually applied to large complex geophysical problem.

This study shows that by reducing the number of parameters to be inverted in the first stage of the search and by the progressive increase of the complexity and dimensionality of the problem genetic algorithm's performances can be enhanced greatly and larger problems could be tackled. Also, any further possibility of reducing the problem dimensionality, such as subdividing the search into small subdomains, should be pursued.

The method potentiality cannot be compared at this time with traditional seismic processing techniques in very large problems, and it should be considered a useful tool to obtain a relatively fast and accurate preliminary analysis. However, the good results obtained in the reconstruction of both the refractor position and the slowness field on synthetic and field data are particularly promising and suggest that further analysis of the method should be pursued.

## ACKNOWLEDGMENTS

We express special thanks to P. Yurez and Dr. M. De la Maza for allowing the local search routine from their Hill-climbing algorithm to be tested in this work. F. Boschetti is supported by an Overseas Postgraduate Research Scholarship funded by the Department of Employment, Education, and Training and by a University Research Scholarship from The University of Western Australia.

## REFERENCES

- Asakawa, K., and Kawanaka, T., 1993, Seismic ray tracing using linear travelttime interpolation: *Geophys. Prosp.*, **41**, 99–111.
- Billings, S., Kennett, B., and Sambridge, M., 1994, Hypocentre location: Genetic algorithms incorporating problem specific information: *Geophys. J. Internat.*, **118**, 693–706.
- Boschetti, F., Dentith, M., and List, R., 1995a, Genetic algorithms incorporating a pseudosubspace method: Presented at the Internat. Conf. on Evolutionary Computing, ICEC 95, 557–560.
- 1995b, A staged genetic algorithm for tomographic inversion of seismic refraction data: *Expl. Geophys.*, **25**, 173–178.
- Buckles, B. P., and Petry, F. E., 1992, Genetic algorithms: IEEE Computer Soc. Press Tech. Series.
- Cavicchio, D. J., 1970, Adaptive search using simulated evolution: Ph.D., University of Michigan.
- Davis, L., 1987, Genetic algorithms and simulated annealing: Morgan Kaufmann Publishers, Inc.
- 1991, Handbook on genetic algorithms: Van Nostrand Reinhold.
- De la Maza, M., and Yurez, D., 1994, Dynamic hill climbing: *AI Expert*, March, 26–31.
- DeJong, K., 1993, Genetic algorithms are NOT function optimizers: Presented at the Foundation of Genetic Algorithms 2.
- Dentith, M. C., Jones, M. C., and Trench, A., 1992, Exploration for gold-bearing iron formation in the Burbidge area of the Southern Cross Greenstone Belt, W.A.: *Expl. Geophys.*, **23**, 111–116.
- Dixon, L., and Szego, G., 1978, Towards global optimization: North-Holland Publ. Co.
- Goldberg, D. E., 1989, Genetic algorithms in search, optimization, and machine learning: Addison-Wesley Publ. Co., Inc.
- Hagedoorn, J. G., 1959, The plus-minus method of interpolating seismic refraction sections: *Geophys. Prosp.*, **7**, 158–182.
- Jin, S., and Madariaga, R., 1993, Background velocity inversion with a genetic algorithm: *Geophys. Res. Lett.*, **20**, 93–96.
- Kennett, B. L. N., Sambridge, M. S., and Williamson, P. R., 1988, Subspace methods for large inverse problems with multiple parameter classes: *Geophys. J. Internat.*, **94**, 237–247.
- Mahfoud, S. W., 1992, Crowding and preselection revisited: *Parallel Problem Solving from Nature*, **2**, 27–35.
- Mathias, K., Whitley, D., Stork, C., and Kusuma, T., 1993, Staged hybrid genetic algorithm search for seismic data imaging: Presented at the Internat. Conf. on Evolutionary Computing, ICEC 1994.
- Oldenburg, D. W., McGillivray, P. R., and Ellis, R. S., 1993, Generalized subspace methods for large-scale inverse problems: *Geophys. J. Internat.*, **114**, 12–20.
- Press, W. H., Vetterling, W. T., Teukolsky, S. A., and Flannery, B. P., 1992, Numerical recipes: Cambridge Univ. Press.
- Raiche, A., 1994, Progress, problems, and challenges: *Surveys in Geophys.*, **15**, 159–207.
- Sen, M. K., and Stoffa, P. L., 1992, Rapid sampling of model space using genetic algorithms: Examples from seismic waveform inversion: *Geophys. J. Internat.*, **108**, 281–292.
- Stoffa, P. L., and Sen, M. K., 1991, Nonlinear multiparameter optimization using genetic algorithms: Inversion of plane-wave seismograms: *Geophysics*, **56**, 1794–1810.
- Williamson, P. R., 1990, Tomographic inversion in reflection seismology: *Geophysics*, **56**, 255–274.
- Wright, A. H., 1991, Genetic algorithms for real parameter optimization: Presented at the Foundation of Genetic Algorithms.