Geophys. J. Int. (1994) **118**, 693–706

Hypocentre location: genetic algorithms incorporating problem-specific information

S. D. Billings, B. L. N. Kennett and M. S. Sambridge
Research School of Earth Sciences, Institute of Advanced Studies, Australian National University, Canberra ACT 0200, *Australia*

Accepted 1994 February 2. Received 1994 February 2; in original form 1993 May 10

SUMMARY

This paper shows how the performance of a fully non-linear earthquake location scheme can be improved by taking advantage of problem-specific information in the location procedure. The genetic algorithm is best viewed as a method of parameter space sampling that can be used for optimization problems. It has been applied successfully in regional and teleseismic earthquake location when the network geometry is favourable. However, on a series of test events with unfavourable network geometries the performance of the genetic algorithm is found to be poor.

We introduce a method to separate the spatial and temporal parameters in such a way that problems related to the strong trade-off between depth and origin time are avoided. Our modified algorithm has been applied to several test events. Performance over the unmodified algorithm is improved substantially and the computational cost is reduced. The algorithm is better suited to the determination of hypocentral location whether using arrival times, array information (slowness and azimuth) or a combination of both.

A second type of modification is introduced which exploits the weak correlation between the epicentral parameters and depth. This algorithm also improves performance over the standard genetic algorithm search, except in circumstances where the depth and epicentre are not weakly correlated, which occurs when the azimuthal coverage is very poor, or when azimuth and slowness information are incorporated. On a shallow nuclear explosion with only teleseismic P arrivals available, the algorithm consistently converged to a depth very close to the true depth, indicating superior depth estimation for shallow earthquake locations over the unmodified algorithm.

Key words: earthquake location, genetic algorithms, non-linear optimization.

1 INTRODUCTION

An optimization problem is any problem that involves a minimization or maximization of some function over a set of model parameters. When the functional relationship is linear, or weakly non-linear, gradient methods can be used. Using local curvature information one can iteratively improve a random initial model, until the adjustments become less than some specified tolerance. Examples include least squares, conjugate gradients and steepest descent. For non-linear functional relationships, or in situations where the computation of derivatives is inexact or expensive, global optimization methods such as nested grid search (Sambridge & Kennett 1986; Kennett 1992), Monte Carlo, simulated annealing (Kirkpatrick, Gelatt & Vecchi 1982, 1983) and genetic algorithms (Holland 1975) are preferable. These methods do not require derivatives and some rely on random processes to search the model space.

Earthquake location can be expressed as a problem of estimating four model parameters (latitude, longitude, depth and origin time) that best fit a set of arrival times of seismic waves at a number of different seismic stations, and possibly other information such as azimuth and slowness. Usually one minimizes the discrepancy between the observed and predicted arrival times using a maximum likelihood criterion. In this work we use a robust Lp-norm to measure the data misfit. For n arrival times this involves finding a four-vector (x, y, z, t) so that;

$$\sum_{i=1}^{n} |t_i^{\text{calc}}(x, y, z) + t - t_i^{\text{obs}}|^p / \sigma_i^p \qquad (1)$$

is minimal, where $t_i^{\text{calc}}(x, y, z)$ is the calculated traveltime from a location (x, y, z) to station i, t_i^{obs} is the observed arrival time of the seismic wave at station i, and σ_i is the standard deviation of the error in the ith arrival time. In the examples used here, traveltimes are calculated in the *iasp*91 velocity model of Kennett & Engdahl (1991). The value of p (usually $1 \leq p \leq 2$) determines the form of the statistical function used to describe the earthquake residuals (i.e. the distribution of errors). As p approaches a value of 1 the distribution becomes more robust and the misfit measure in eq. (1) becomes less sensitive to the occasional outlier. For $p = 2$ the earthquake residuals are assumed to follow a Gaussian distribution, which is non-robust and often a poor description of arrival-time errors (Buland 1986). A more practical choice is to use $p = 1$ (Sambridge & Kennett 1986; Kennett 1992) or $p = 1.25$ as in Billings (1994) and the present work.

Equation (1) is non-linear in the hypocentre coordinates, creating potential instabilities, especially when the recording network is sparse or has poor geometry (Smith 1976; Buland 1976; Herrmann 1979; Lee & Stewart 1981; Anderson 1981; Lienert & Frazer 1983). In addition, the computation of derivatives is complicated and costly when multiple arrivals, array observations or a 3-D velocity model are used. In these cases the earthquake hypocentre can be located using global optimization techniques such as the nested grid search algorithms of Sambridge & Kennett (1986) and Kennett (1992) or stochastic methods like the simulated-annealing algorithm of Billings (1994) and the genetic algorithms of Kennett & Sambridge (1992) and Sambridge & Gallagher (1993).

In this paper we investigate how the performance of an optimization algorithm can be improved by modifying the model space. The genetic algorithm has been shown by Kennett & Sambridge (1992) and Sambridge & Gallagher (1993) to be effective for regional and teleseismic earthquake location when the network geometry is ideal. We find that the algorithm can be unreliable and inaccurate on events with poor network geometries. This can be overcome by exploiting the interaction between different parameters in the location problem. We say that two hypocentral parameters are strongly correlated if accurate estimation of one hypocentral parameter is strongly dependent on accurate determination of the other. Depth and origin time are strongly correlated, while the epicentral parameters and depth are weakly correlated. An algorithm is presented which avoids problems associated with the strong correlation between depth and origin time by separating the spatial and temporal components of the search. Billings (1994) has already shown that this improves the performance of a simulated-annealing algorithm. A second algorithm is presented that exploits the weak correlation that exists between epicentre and depth by performing separate searches on the two coordinates spaces.

2 GENETIC ALGORITHMS FOR EARTHQUAKE LOCATION

The genetic algorithm was developed by Holland (1975) and introduced to geophysics by Stoffa & Sen (1991) and Sambridge & Drijkoningen (1992). Both contain comprehensive descriptions of the method. Kennett & Sambridge (1992) and Sambridge & Gallagher (1993) used the method for earthquake location.

The genetic algorithm has an analogy with biological systems, whereas simulated annealing has an analogy with thermodynamics. A genetic algorithm always operates on a group or 'population' of Q sets of hypocentral parameters simultaneously. Initially, these may be generated randomly and subsequent *generations* of models are constructed by the action of three operators each with biological analogues. These processes are selection, crossover and mutation, the action of each being controlled by a separate probability distribution. In the *selection* step those hypocentral estimates with the lowest misfit between observed and predicted arrival times (from eq. 1) have the highest probability of being passed on to the next generation. This introduces a process akin to survival of the fittest. New hypocetral parameters are created in the *crossover* step by swapping of segments between pairs of bit strings which are the binary encodings of the hypocentral parmeters. The *mutation* step allows individual bits to flip from 0 to 1, and vice versa, with a low probability. As successive generations pass, the misfit is driven towards the minimum by the accumulation of information about the shape of the misfit surface in 4-D space.

The model space is discretized and the finite set of model parameters are represented as concatenated binary strings. The performance of the algorithm depends critically on the discretization of the search space. For an extremely fine grid spacing the potential accuracy is good, but the model space is large, requiring many iterations to locate the global minimum. In contrast, too coarse a discretization allows for rapid location of the minimum on the coarse grid, but with a subsequent loss in accuracy, as the minimum on the grid may be far from the true global minimum. In addition, the global minimum of a coarse grid does not necessarily lie close to the true global minimum. In earthquake location the model space is potentially of size $360° \times 180° \times 670\,\mathrm{km} \times 600\,\mathrm{s}$. To achieve an accurate location over this search space the genetic algorithm would have to use a very large number of bits in the bit string. In practice, rough bounds of $2° \times 2° \times 60\,\mathrm{km} \times 12\,\mathrm{s}$ can usually be obtained before application of the genetic algorithm. Using eight bits to represent latitude and longitude and seven bits for depth and origin time, this gives a bit spacing of 0.87 km in latitude and longitude, 0.47 km in depth and 0.1 s in origin time.

A possible method of delineating bounds on the four coordinates for local networks, that could easily be extended to the case of teleseismic networks, is discussed by Sambridge & Kennett (1986). They use the arrival-order method of Anderson (1981) to find epicentral bounds and, if both P and S arrivals are available, the depth and origin time can also be bounded. The three events considered in this paper each have a preliminary location found by reporting agencies and so useful bounds on each parameter can be obtained using loose tolerances about this hypocentre.

The performance of the genetic algorithm is dependent on several adjustable parameters. An extensive set of tests on several earthquakes were used to determine a reliable set of these parameters. One parameter is related to an improvement to the selection step known as 'tournament selection' discussed in Goldberg & Deb (1991). This modified form of selection improves performance in hypocentre location and avoids problems associated with *ad hoc* scaling as tournament selection only makes use of the rank of the misfits. Random pairs of hypocentres are selected from the population of Q individuals and the respective misfits computed. A random number is generated between 0 and 1 and if it is less than the tournament selection probability (the adjustable parameter) the hypocentre with the higher misfit is passed on to the next generation, otherwise the lower misfit hypocentre is passed on. Both models are then put back into the initial population and the procedure is repeated until there are Q models in the 'offspring' population. The tournament selection probability, P_T, controls the likelihood that better models will survive the selection process. For $P_T = 0.5$ the search is random with no discrimination between good and bad models, while for $P_T > 0.5$ better models are favoured. The genetic algorithm was found to perform best when $P_T = 0.9$, implying that nearly all better models should be passed on to the offspring population. In the context of earthquake location, tournament selection is discussed in more detail by Sambridge & Gallagher (1993). A similar approach called 'update' was proposed by Stoffa & Sen (1991) but involved comparing parents with children.

The probability of crossover, which controls the amount of information exchange between hypocentres in the population, was set at 0.9. This means that 90 per cent of paired hypocentres exchange information via crossover. The probability of mutation, which controls the level of randomness in the search, was set at 0.04. Any individual bit has a 4 per cent chance of changing its value from 0 to 1 or vice versa. The population size, Q, was set at 24 models. In the absence of a better criterion, the search is terminated after a fixed number of iterations have been performed. By investigating performance on a large number of events it was found that performing about 50 iterations (we settled for 54) was ideal. When more iterations are performed, further improvements to the hypocentre are generally small. The model parameters are encoded using bit strings of length eight, for latitude and longitude, and length seven for depth and origin times. This means that there are a total of $2^{30} \approx 1 \times 10^9$ possible models. The genetic algorithm attempts to find an estimation of the global minimum by sampling only $54 \times 24 = 1024$ models.

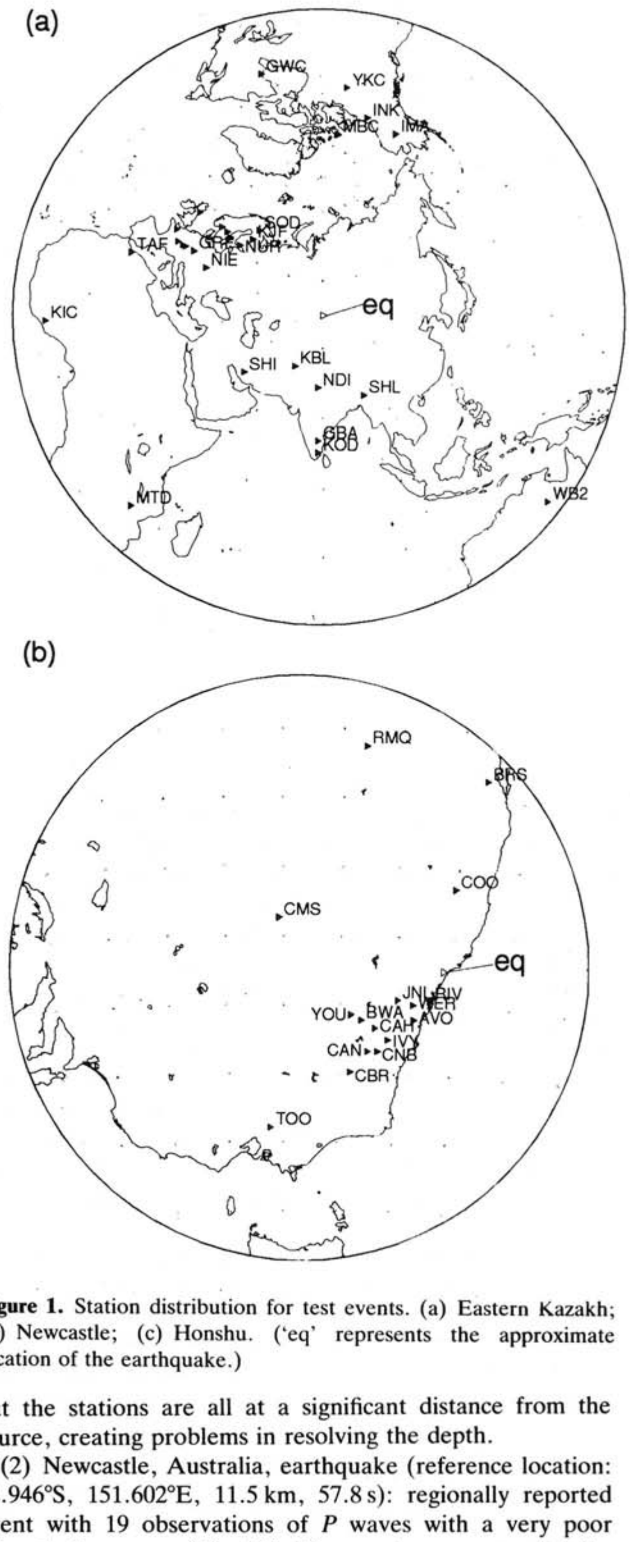

Figure 1. Station distribution for test events. (a) Eastern Kazakh; (b) Newcastle; (c) Honshu. ('eq' represents the approximate location of the earthquake.)

3 PERFORMANCE OF THE STANDARD GENETIC ALGORITHM

The performance of the genetic algorithm was investigated on three test events with differing characters of misfit surface. Each has the common feature of the absence of observations in a large azimuth zone. The three events were a nuclear explosion in Eastern Kazakh, and earthquakes in Newcastle (Australia) and Honshu (Japan). These events are the same as those used to test the simulated-annealing algorithm of Billings (1994).

(1) Eastern Kazakh nuclear explosion (reference location: 49.765 °N, 78.059 °E, 0.139 km, 59.8 s): a shallow event with 29 stations at teleseismic distances recording P arrivals only. The azimuthal coverage is moderately good (Fig. 1a) but the stations are all at a significant distance from the source, creating problems in resolving the depth.

(2) Newcastle, Australia, earthquake (reference location: 32.946°S, 151.602°E, 11.5 km, 57.8 s): regionally reported event with 19 observations of P waves with a very poor azimuthal coverage (Fig. 1b). The event was responsible for the first recorded deaths in Australia due to an earthquake.

(3) Honshu, Japan, earthquake (reference location: 35.260°N, 138.58°E, 166 km, 59.5 s): a predominantly teleseismically recorded event with several array observations of azimuth and slowness (Fig. 1c). Observations are absent from the entire Pacific sector.

Once hypocentral estimates have been generated for an event, it is difficult to determine how accurate these

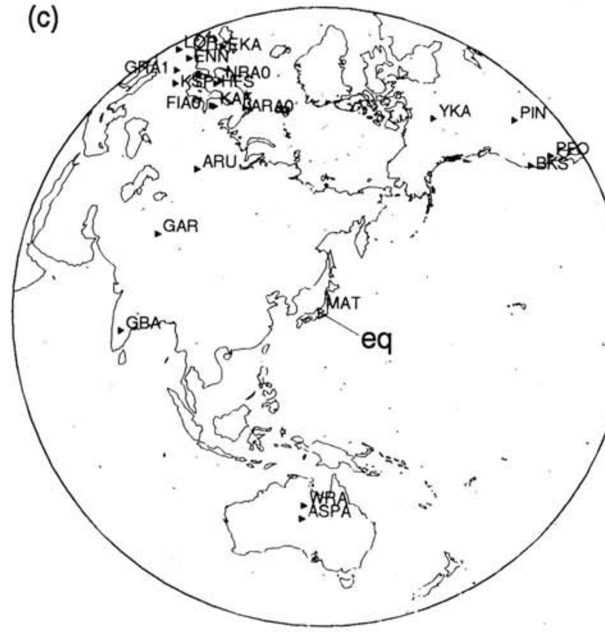

Figure 1. (*Continued.*)

locations are. If the hypocentre is already accurately known (as for nuclear explosions) the relative locations can be compared. However, the true hypocentre is not necessarily the best location under the criterion used to locate the event. The true arrival times are not known with perfect accuracy, and the calculation of traveltimes occurs in a 1-D approximation to the real earth, so the global minimum will *not* necessarily correspond to the true hypocentre. In addition, Kennett (1992) has shown that changing the statistical distribution of residuals can have as large an effect on the location as changing the velocity model used to calculate the traveltimes. Therefore, exact locations cannot be used to test the accuracy of an optimization algorithm, although they *can* be used to test the validity of a location criterion. In addition, the value of the misfit at the global minimum is not known in advance, so this measure cannot be used.

The performance of a stochastic algorithm is dependent on the random numbers used in the course of the optimization. Therefore, just one run of an algorithm is *insufficient* to determine the accuracy and reliability of the algorithm. As a means of testing reliability a total of 1000 separate locations were performed for each event. The performance of the algorithm can be analysed by considering the spread of locations and misfits, without prior knowledge of the location or misfit of the global minimum. There are three distinct outcomes from which conclusions can be drawn.

(1) The locations are all tightly clustered (say <2 km in space, <0.5 s in time) about a single region in the parameter space.

(2) The locations are spread throughout the parameter space and the misfits are variable.

(3) The locations are well spread but the misfits are similar.

In general if case 2 occurs then we would infer that the algorithm is unreliable and inaccurate. If case 3 occurs the misfit surface may have multiple minima indicating that the solution is non-unique, or at least not well constrained by the data. Billings (1994) has shown that on a long-wavelength scale the misfit surface is unimodal (at least for the events considered here), and consequently widely separated multiple minima are unlikely. If case 1 occurs the algorithm is either accurate and reliable, or has consistently converged to a local minimum or suboptimal solution. The long-wavelength unimodal structure indicates that this latter possibility is unlikely. Even though the performance of a single algorithm cannot be fully separated from effects due to the misfit structure, the relative performance of two competing algorithms can be contrasted by comparing the two sets of 1000 locations. This is how we shall compare the performance of the standard and modified genetic algorithms to be introduced in the next section.

A histogram of the misfits of the final 1000 locations for each event is shown in Fig. 2. The horizontal axis represents the misfit of the final location found, with better models to the left, while the vertical axis represents the number of times that the misfit occurred in 1000 trials. The additional peak to the right of each diagram represents the number of locations produced with a higher misfit than the maximum shown on the plot. Two measures of the algorithms performance are also shown; the mean and median. The median is a more robust measure of performance than the mean, which can be distorted by a few locations with a very large misfit.

On the Eastern Kazakh event the genetic algorithm locates hypocentres most often with a misfit of just over 0.460, and finds fewer hypocentres with a lower misfit. There is a range in misfit values between 0.47 and 0.485 where almost no locations are found. Then at $\sim$0.490 there is a secondary peak, which may correspond to the presence of a local minimum. On the Newcastle earthquake, there is a more even misfit frequency distribution, but with a definite bias towards better models. For the Honshu earthquake there is a concentration in locations with a low misfit, which dies off quite rapidly with increasing misfit. However, there is a small secondary peak at a misfit value of about 0.62.

If all 1000 locations for each event in Fig. 2 are clustered *close* to the global minimum, and yet the misfit values are highly variable, then it can be inferred that the genetic algorithm has performed as well as it can for a highly irregular misfit surface (case 1). The variation in misfit is then predominantly due to systematic modelling errors and random noise in the data. However, if locations with relatively large misfit values are inadequate (i.e. far from the global minimum) then the inference is that the genetic algorithm is not a reliable method for earthquake location (case 2). To determine the quality of locations, it is necessary to determine the distribution of the hypocentres within parameter space. Fig. 3 represents the frequency of final locations for the Newcastle earthquake in four planes; latitude–longitude, latitude–depth, longitude–depth and depth–origin time. Each peak represents the number of locations that occur in a particular parameter range. The dimensions of the diagram are 22 km in latitude, 44 km in longitude, 50 km in depth and 2 s origin time,. In the latitude–longitude plane the genetic algorithm converges to a particular region in approximately 22 per cent of the trials.

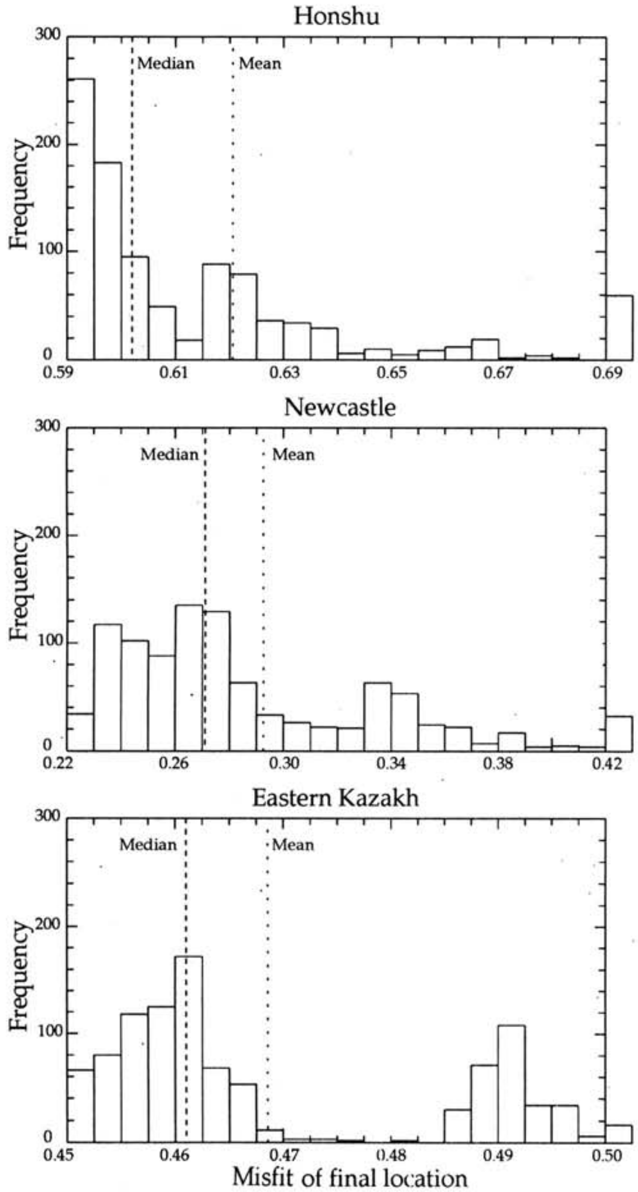

Figure 2. Frequency of misfits of hypocentral locations generated by the genetic algorithm in 1000 separate optimizations for each of the three test events. Lower misfit values indicate better performance, but cross comparison of misfit values between plots is not possible, as the misfit functions and minimum misfits are different for different events.

However, there is no consistency as the algorithm converges to locations with approximately a linear trend in latitude–longitude, with the epicentres differing by as much as 50 km. This is clearly an unacceptably large variation, especially when superimposed on the other uncertainties involved in earthquake location. The sections through the other three planes show that the standard genetic algorithm is unable to reliably resolve these parameters. Final depths differ by as much as 50 km and the origin time by up to 2 s. Similar poor performances are found for the other two events. Consequently, for events with poor network geometries, the genetic algorithm's performance corresponds to case 2; i.e. the misfit variation (seen in Fig. 2) is not due to systematic modelling errors and random noise in the data, but rather because accurate solution to the optimization problem has not been obtained.

The genetic algorithm used in these examples does not take advantage of any special features of the location problem, which makes it a very versatile and robust method of sampling a parameter space, but often results in it being out-performed by specialized optimization routines. The primary purpose of this paper is to investigate the effect of incorporating problem-specific information into a genetic algorithm. In earthquake location there are several correlations between hypocentral parameters that can be exploited, and these are discussed in the next section.

4 MODIFIED GENETIC ALGORITHMS EXPLOITING CORRELATIONS BETWEEN HYPOCENTRAL COORDINATES

There are two parameter correlations that can be exploited to improve the genetic algorithm's performance for earthquake location. The first is the strong correlation between depth and origin time and the second is the weak correlation between the epicentre and depth parameters.

4.1 Modified algorithm 1: exploiting the strong correlation between depth and origin time

Billings (1994) presents an earthquake location algorithm which separates the model space into spatial and temporal components, with a simulated-annealing search applied to space, and a golden section search (Whittle 1971) to time. The algorithm presented here is identical except the spatial simulated annealing search is replaced with a genetic algorithm search. For completeness the reasoning and methodology will be briefly reviewed.

It is well known that in the earthquake location problem there exists a strong trade-off between depth and origin time. Changing the depth of focus will change the traveltime between the source and the recording station, but the effect on the calculated arrival time at that station can be compensated by varying the origin time of the earthquake. This can best be seen in an example using the arrival times of the nuclear explosion at Eastern Kazakh. First, we define an 'optimum origin time' as that origin time which produces the lowest misfit at a particular spatial location. Fig. 4 shows how the origin time trades off with depth for a fixed epicentral location. In the two side panels we have plotted the misfit against origin time for depths of 1.04 and 9.06 km, respectively. The single well-defined minimum, in both plots, represents the optimum origin time at the corresponding depth. The central panel shows how the optimum origin time varies with depth. Notice how the optimum origin time changes by almost 2 s in a depth interval of just 10 km. For a teleseismic event our parameter ranges are about 60 km for depth and 12 s in origin time. Therefore, we can expect a variation of up to 10–12 s in origin time. We see that the two parameters are strongly correlated, and an accurate estimate of one parameter will be dependent on an accurate estimate of the other.

The strong correlation between depth and origin time means that, in general, the depth parameter will be difficult to resolve independently. However, the plot in Fig. 4

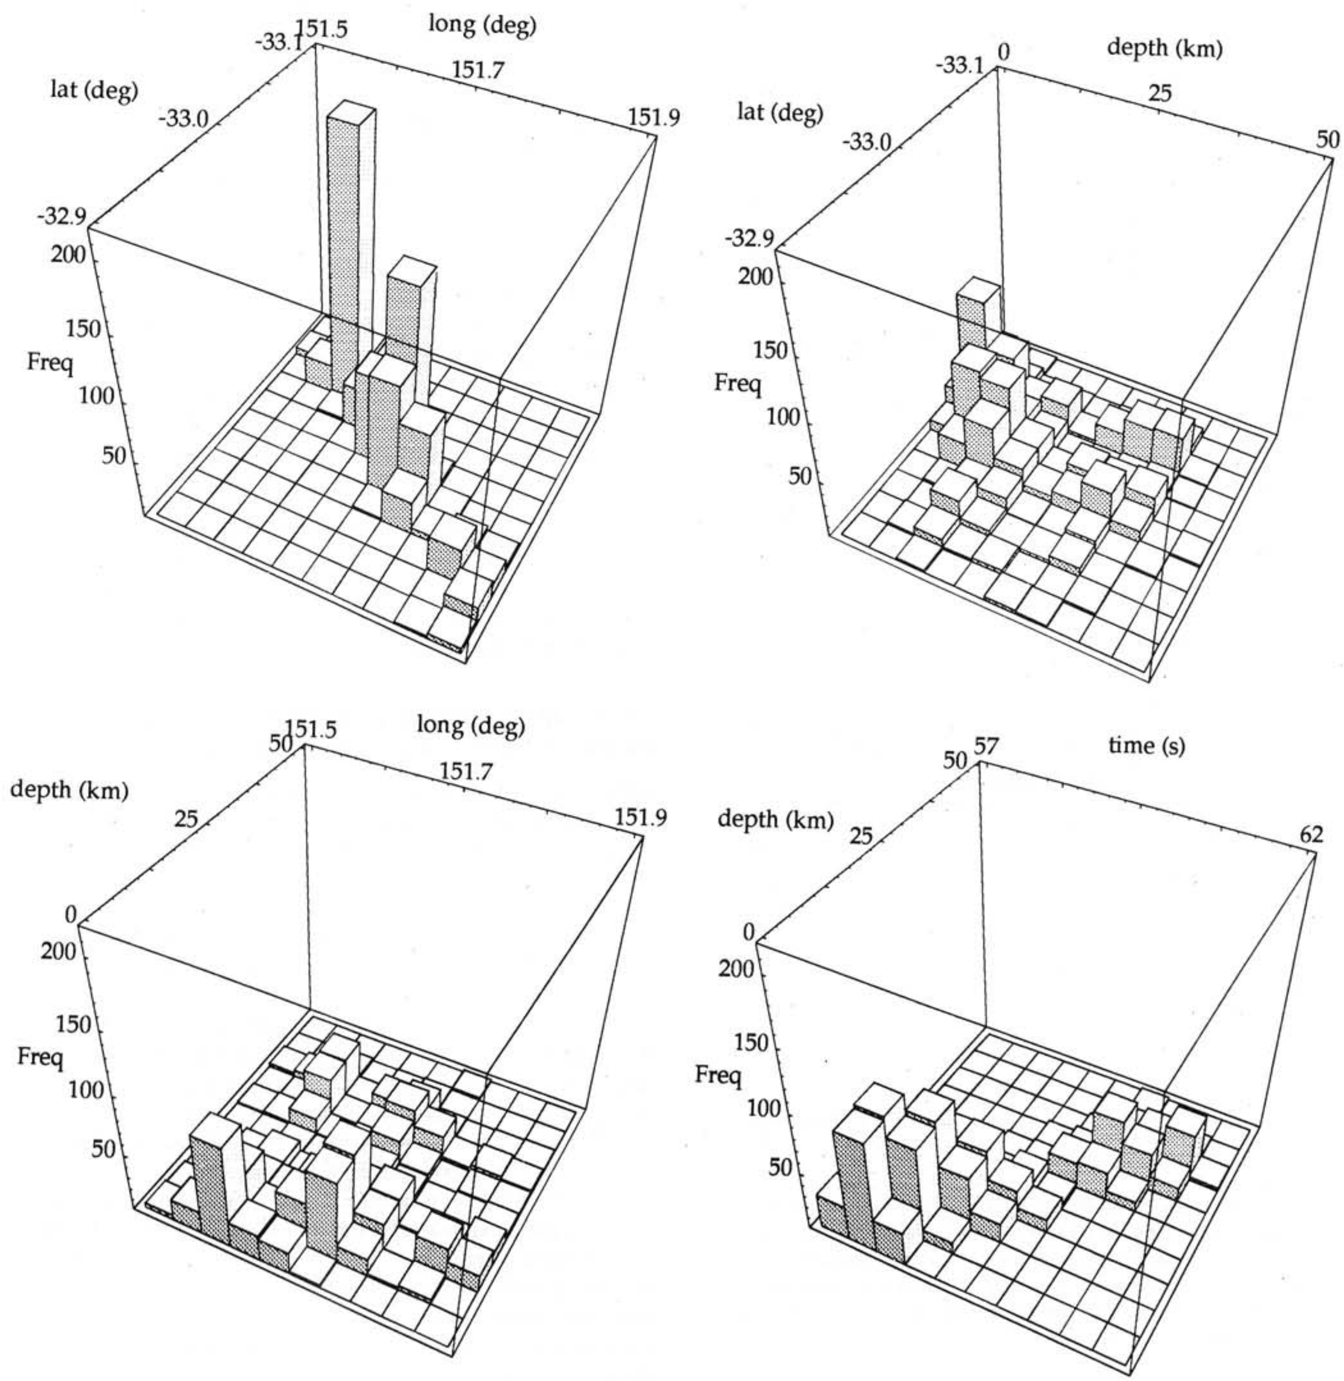

Figure 3. The frequency with which, after 1000 trials, the final hypocentral location is in the range defined by the projection of locations onto the plane in question for the Newcastle earthquake. The diffuse pattern of hypocentral locations shown here indicates that the genetic algorithm is not reliable.

suggests that by separating the spatial and temporal searches and computing the optimum origin time for each spatial location, estimation of the depth parameter will be much simpler. Since the traveltime between source and receiver is fixed regardless of the origin time, then arrival times, and hence misfits, may be calculated for different origin times with only a few simple arithmetic operations. Consequently, the optimum origin time can be quickly found using a golden section search (Whittle 1971) over the depth parameter. The golden section search for finding the minimum value of a function in one dimension is similar to the bisection method of root finding. It is appropriate because it does not require the computation of derivatives. For the special case of Gaussian statistics there exists a simple analytical formula for the optimum origin time (see Billings 1994).

With this modification the genetic algorithm need only be applied to the spatial part of the problem, because in effect

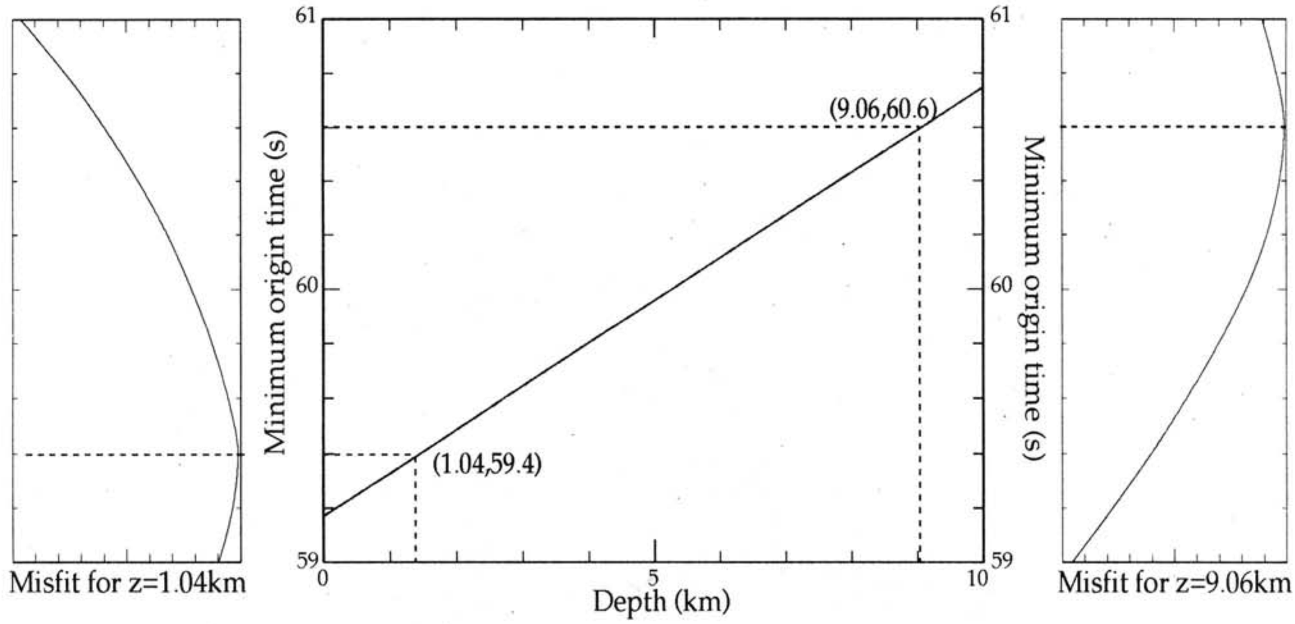

Figure 4. The variation of optimum origin time with depth for a fixed latitude and longitude for the Eastern Kazakh nuclear explosion. The central box represents the variation of the optimum origin time with depth. The smaller plots either side display the variation in misfit value with origin time for a fixed depth ($z = 1.04$ km and $z = 9.06$ km) and illustrate how the optimum origin time is defined. The large variation in optimum origin time with depth indicates that these two coordinates are strongly correlated.

the misfit in eq. (1) is always evaluated at the optimum origin time for that depth. This means that a smaller population size of around 20 models may be used. The probabilities of tournament selection, crossover and mutation were set at the same values as for a genetic algorithm search over all four hypocentral coordinates (0.9, 0.9 and 0.04 respectively). The decreased complexity of the search space results in the algorithm converging more rapidly, and 40 iterations were usually found to be sufficient.

4.2 Modified algorithm 2: exploiting the weak correlation between epicentre and depth

For events with a good azimuthal coverage the epicentre is largely constrained by the distribution of stations, and a reasonably accurate estimation can be made by using Anderson's (1981) arrival order method. Ray tracing through an assumed earth model is required primarily for determining the depth of the earthquake. Additonally, in the neighbourhood of the global minimum, the optimum depth at a particular epicentre is little affected by small changes in the epicentre, implying that latitude and longitude are weakly correlated with depth. An algorithm can be developed that exploits this weak correlation by performing separate searches in epicentre and depth-time.

We introduce 'optimum depth' in a similar manner to optimum origin time and define it as that depth which produces the minimum misfit for a fixed epicentral location. For each depth we calculate the misfit using the optimum origin time at that depth. The optimum depth may therefore be plotted as a function of the epicentral parameters and for the Eastern Kazakh event this surface is shown in Fig. 5. The range in latitude and longitude shown spans 2° and corresponds to the size of the model space. The grid spacing is approximately 4.2 km. The depth varies from 0 to 35 km, and latitude–longitude locations with optimum depths below this depth are coloured black. The best approximation to the global minimum is shown by a cross. Notice in the neighbourhood of the global minimum, that the optimum depth is zero and does not vary with latitude and longitude. For this event it can be concluded that in the neighbourhood of the global minimum, the depth is only weakly dependent on the epicentre. This feature may not exist for events with poorer azimuthal coverage and an algorithm relying on a weak correlation between epicentre and depth may break down.

The nature of the tau-spline representation (Buland & Chapman 1983) used in conjunction with the *iasp*91 tables means that it is computationally more expensive to recompute misfit at different depths than at different latitudes or longitudes. The total number of repeat calculations of misfit for different depths can be reduced by coupling a genetic algorithm search on latitude–longitude with a nested golden section search in depth and origin time. The 1-D golden section search can be used in two dimensions by performing a step-wise minimization. Each time the misfit needs to be calculated for a particular depth, the optimum origin time is estimated by performing a golden section search on time. A reference origin time is then available to compute the misfit and advantage is taken of the ease of re-calculating misfit for different origin times. In addition, there is no need to minimize the depth for every epicentral location, because we are assuming that the depth is only weakly correlated with epicentre. A single, fixed depth (and origin time) can be used as a reference to calculate epicentral misfits.

The modified algorithm consists of the following steps.

(1) A large number of hypocentres are randomly generated with a uniform distribution, and their misfits are

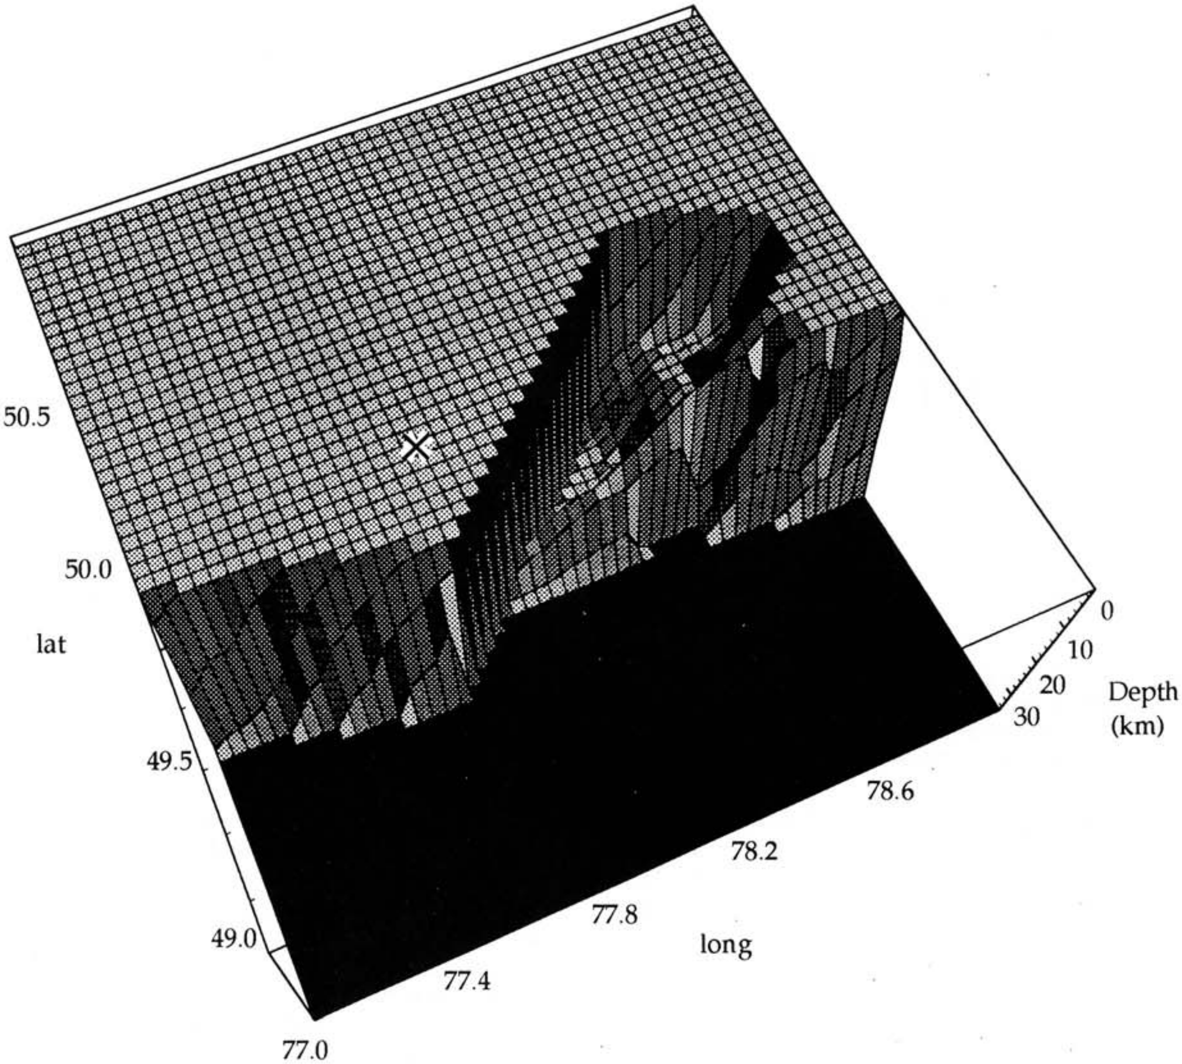

Figure 5. The variation in optimum depth with epicentre for the Eastern Kazakh nuclear explosion. The nearest approximation to the global minimum is shown by a cross. In the neighbourhood of the global minimum the depth is independent of epicentre. The particular range of latitude and longitude shown here represents the model space through which the genetic algorithm searches for the global minimum.

calculated according to eq. (1). The latitude, x_{op}, and longitude, y_{op}, from the hypocentre with lowest misfit is used as a reference to calculate an optimum depth and origin time pair, (z_{min}, t_{min}), by a nested golden section search as described in the previous paragraph. This depth and origin time pair is used as a reference for computing the misfits, F, of other members of the latitude–longitude population which we can index as (x_i, y_i);

$$F(x_i, y_i) = F(x_i, y_i, z_{min}, t_{min}). \tag{2}$$

(2) The latitude and longitude population is evolved using the genetic algorithm, with eq. (2) used to calculate misfits.

(3) Step (2) is repeated until the newly evolved population contains a latitude–longitude pair, (x_j, y_j), that decreases the misfit below that of the best epicentre, (x_{op}, y_{op}), found up until that point. This condition occurs when

$$F(x_j, y_j, z_{min}, t_{min}) < F(x_{op}, y_{op}, z_{min}, t_{min}). \tag{3}$$

The new latitude–longitude pair then becomes the optimum epicentral estimate. As the epicentre has altered, the optimum depth and origin time may also change, possibly only slightly because of the weak correlation between epicentre and depth. The nested golden section search is then used to find new estimates of the optimum depth, z'_{min}, and origin time, t'_{min}.

(4) Steps (2) and (3) are repeated, except the reference depth and origin time used to calculate misfits via eq. (2) are z'_{min}, t'_{min}.

The tournament selection number and the probabilities of crossover and mutation were chosen as before. A population size of 14 was found to give good results. The algorithm was found to be fairly slow to converge, and 100 iterations were necessary. Re-calculating the misfit for different latitude and longitudes is less expensive than for different depths, and so these 100 iterations can be performed relatively efficiently. Finally, better results were obtained if a large number of hypocentres (we chose 36) were randomly generated at the beginning of the search.

This ensures that a reasonably good latitude–longitude pair is chosen to compute the first minimal depth and origin time.

5 RESULTS

The two modified algorithms were tested on the same three events used to test the original version of the genetic algorithm. The events are particularly suitable because they display a range of characteristics that may affect relative performance. The Eastern Kazakh nuclear explosion has a relatively good azimuthal coverage (Fig. 1a) with the depth weakly dependent on the epicentre (Fig. 5). As the *P* arrivals are all at teleseismic distances there may be some difficulty in resolving the depth. The azimuthal coverage of the Newcastle event is poor (Fig. 1b), possibly resulting in a dependence of depth on epicentre, which may affect the performance of the epicentre–depth-separated algorithm. The inclusion of azimuth and slowness information in the Honshu event (Fig. 1c), changes the character of the misfit surface, in particuar the depth–origin time correlation. This may affect the performance of both modified algorithms.

For each event and each algorithm, 1000 separate locations were again performed using different strings of random numbers. The results are summarized in Fig. 6 where the algorithms are arranged horizontally and the events vertically. The diagrams represent the number of times in 1000 locations that the algorithm converged to a location with a particular misfit. Recall that misfit gives a measure of the performance of the algorithm, with low values indicative of good performance. Notice that the vertical scales differ between events. The space–time-separated algorithm (top row) improves performance over the standard search on four hypocentral coordinates (middle row) in all three events. The epicentre–depth-separated algorithm (bottom row) represents an improvement on the Eastern Kazakh event but not for Newcastle or Honshu, where the depth and epicentre are not weakly correlated. The dependence of optimum depth on epicentre arises because of the poor azimuthal coverage for the Newcastle earthquake and the inclusion of azimuth and slowness for the Honshu earthquake. We now consider each event separately.

5.1 Eastern Kazakh

Consider the relative performances of the three algorithms on the Eastern Kazakh nuclear explosion. The genetic algorithm with no separation (middle) produces two peaks in misfit frequency, suggesting an interaction with a secondary minimum. The median, which gives a measure of the performance of an algorithm, is at a misfit value of 0.462. The space–time-separated algorithm is much more likely to locate the earthquake accurately with a median of 0.457 and a definite bias towards locating models with lower misfit. Notice, that there is still some concentration of locations with relatively high misfit. The epicentre–depth-separated algorithm produces locations intermediate in quality between the other two algorithms (median 0.459). Clearly, the two modified algorithms improve performance over the standard genetic algorithm search. The ability of the two modified algorithms to locate the depth of the global minimum is much better than that of the standard algorithm (Fig. 7). The epicentre–depth-separated algorithm gives the most consistent depth estimation, with all hypocentral locations at a depth of 0 km (true depth 0.139 km). The space–time algorithm gives only a slightly poorer estimate of depth with 98 per cent of depths between 0 and 1 km. In contrast, the standard genetic algorithm search has poor depth estimation with only a slight bias towards locating the event very shallow with only 15 per cent of trials between 0 and 1 km. The algorithm interacts with the valley in depth and origin time and in many trials is unable to successfully locate the more promising regions. Both modified algorithms effectively eliminate problems related to the depth–origin time trade-off and are consequently better able to estimate the depth parameter.

5.2 Newcastle

Referring back to Fig. 6, consider the relative performances on the Newcastle earthquake. The standard genetic algorithm search produces a fairly uniform spread of misfit values with only a slight bias towards better models. The median misfit value of hypocentral locations is 0.271. The space–time-separated algorithm improves performance substantially, reducing the median to 0.236. In terms of hypocentral locations the algorithm is better able to resolve coordinates than the standard search, as shown in Figs 3 and 8. The plot for the space–time-separated algorithm (Fig. 8) has the final locations clustered in a much smaller region, indicating a greater likelihood of locating the region of the global minimum. When contrasting performance of the two algorithms using Figs 3 and 8 note that the horizontal scales are the same but the vertical scale in Fig. 8 has been more than doubled from 220 to 500. This means that most of the peaks in Fig. 8 would go off the scale in Fig. 3. Instead of a fairly diverse spread in locations there is now a strong bias towards locating the earthquake in a fairly small region. This indicates that the space–time-separated algorithm produces more accurate and reliable locations than the standard search.

The epicentre–depth-separated algorithm does not perform as well in comparison to the space–time-separated algorithm. While more likely to produce a better result than the standard search (the median is lower at 0.267), it is also much more likely to produce a very bad result, with 25 per cent of locations having misfits larger than the largest values shown on the plot. There are two reasons for this poor performance. First, the poor azimuthal coverage evident in Fig. 1(b) results in a dependence of latitude and longitude on depth. Consequently, the optimum depth varies significantly with epicentral position (Fig. 9). There are two regions where the optimum depth is approximately constant (at 0 and 30 km) but neither of these are close to the depth at the global minimum ($\sim$10 km). The algorithm relies on a weak correlation between depth and origin time in the neighbourhood of the global minimum, which is clearly not present. Secondly, in $\sim$17 per cent of trials the algorithm converged to a depth of 20.6 km which is close to the 20 km discontinuity in the *iasp*91 velocity model. For depths below 20 km, waves are identified as *Pb* waves, while above 20 km they are identified as *Pg* waves. It is generally not possible to separately identify *Pg* and *Pb* waves on a real

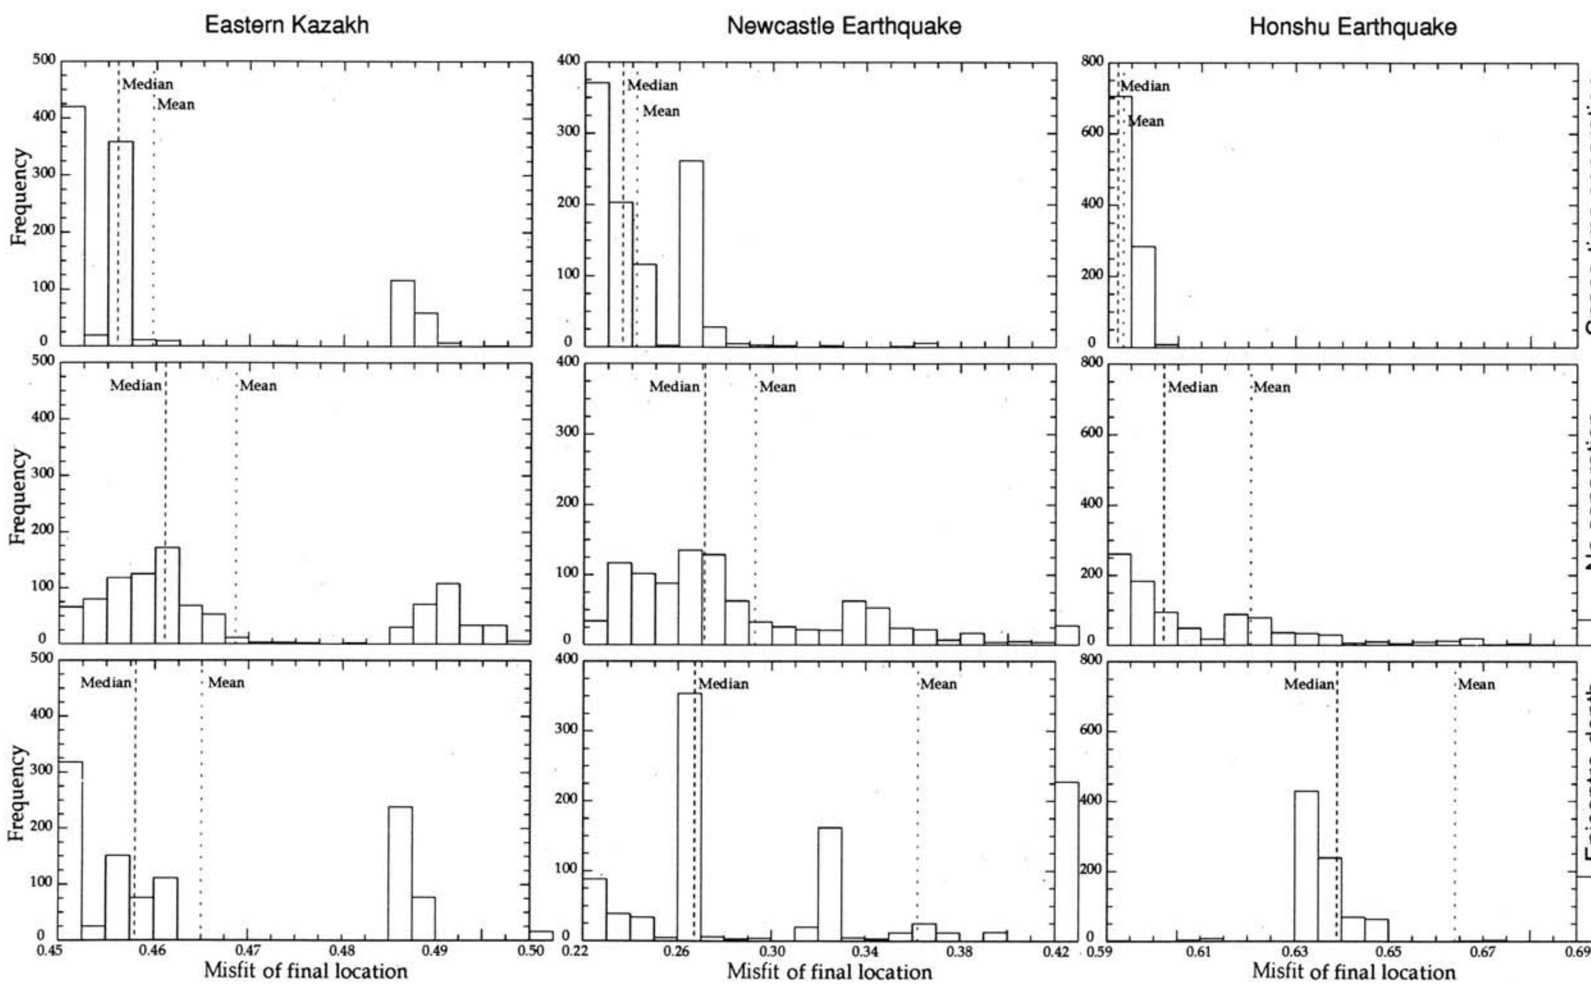

Figure 6. Frequency of misfits of hypocentral locations generated by the three algorithms (standard search, epicentre and depth–time separated) in 1000 separate optimizations for each of the three test events. The algorithms are arranged in rows and the events in columns. Lower misfit values indicate better performance. The space–time-separated algorithm improves performance substantially over the standard search. The epicentre–depth-separated algorithm improves performance on the Eastern Kazakh event, but tends to perform poorly on Newcastle and Honshu where depth and epicentre are no longer weakly correlated.

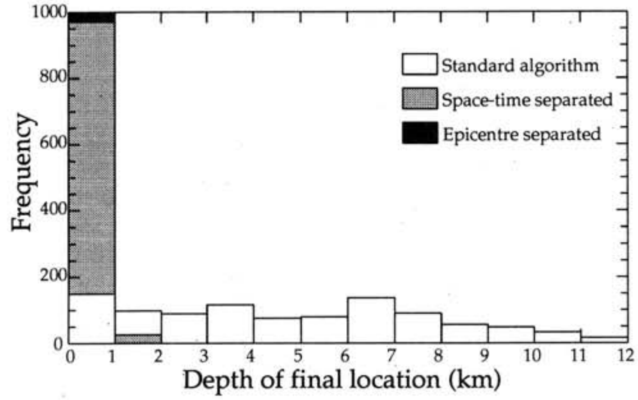

Figure 7. The number of times in 1000 trials that the final hypocentral location is at the depth shown on the horizontal scale. Lower peaks are drawn in front of the higher peaks. The two modified algorithms have much better depth resolution than the original algorithm.

seismogram. The *Pb* and *Pg* waves have different traveltimes, so instead of a smooth increase in velocity with depth, there is a discontinuous change, which creates a discontinuous jump in the misfit surface. The epicentre–depth algorithm has a strong interaction with this feature and is unable to accurately estimate the depth for this event. The 20 km discontinuity in the *iasp*91 model is between two layers of constant velocity. Gradient methods artificially concentrate locations on this type of discontinuity and it would appear that the epicentre–depth-separated algorithm is also capable of producing such artificial concentrations.

5.3 Honshu

The relative performance of the different algorithms on the Honshu earthquake can be compared by referring again to Fig. 6. The standard search is capable of producing poor locations but has a definite bias towards higher quality locations (median 0.602). The space–time-separated algorithm improves performance substantially, with ~70 per cent of trials in the best misfit bracket. The median is very low at 0.592 with almost no poor performances. The inclusion of array observations has changed the nature of the misfit surface, and in particular the trade-off between depth and origin time. However, the space–time-separated algorithm is still able to perform well, because it does not depend on a strong correlation between depth and origin time. It assumes a worst-case scenario of a strong correlation and will work equally well when the correlation is weak or non-existent. The epicentre–depth-separated algorithm performs very poorly on this event with almost no high-quality locations in the 1000 trails. In 20 per cent of trials the final misfit is too large to fit on the plot, while the algorithm is most likely to produce locations with a misfit in the mid-range of the plot (median 0.639). The epicentre–depth-separated algorithm performs so poorly because azimuth and slowness introduce a dependence of depth on epicentre. Azimuth gives the direction of the earthquake from the observing station, while slowness gives a take-off angle. Together azimuth and slowness define a curved path through the Earth along which the earthquake must have occurred, resulting in a dependence of depth on epicentre.

6 DISCUSSION

The space–time-separated algorithm improves performance on all three test events. On an event larger set of events, including events with ideal network geometry, the algorithm is never out-performed by the standard search over all four hypocentral coordinates. In terms of computational cost the modified algorithm is also cheaper to implement. It evaluates many more misfits (in the order of 4000) than the standard search, but a great many of them differ only in the origin time parameter and therefore may be formed efficiently. As the number of observations increases, the cost approaches that of the standard genetic algorithm, because the discrepancy between cost of re-calculating in time compared to space decreases with increasing numbers of observations. The algorithm relying on an epicentre–depth separation improves performance when the azimuthal coverage of the network is good, but can produce less reliable solutions when this is not the case. It is about 20 per cent cheaper to implement than either of the other two algorithms presented here.

The space–time-separated algorithm improves performance on all three events. By minimizing in time for every spatial location the genetic algorithm needs to search through a model space of decreased complexity. Instead of four unknowns the algorithm needs to find only three unknowns, reducing the number of possible models from 1.1×10^9 to 8.4×10^6, which is a substantial reduction. Of course, the origin time still needs to be determined, but this can be achieved using a simple and fast 1-D optimization routine. In addition, the variation of the misfit surface for the three spatial parameters is much less complicated than the misfit surface for all four parameters and represents a subspace of the original four-parameter space. A further reason for the improved convergence results from the elimination of the depth–time relationship. As far as the genetic algorithm is concerned there are no longer numerous depths and origin times with very similar misfits. This is a problem for the standard genetic algorithm over all four hypocentral coordinates as often accurate depth estimates are generated but the associated origin time is poor, giving a relatively high misfit. It is then the task of crossover and mutation to improve the origin time. In this modified algorithm accurate depth estimates will always have small misfits as the origin time is minimized for each depth, substantially decreasing the complexity of the problem.

The epicentre–depth-separated algorithm performs well when the depth is weakly dependent on the latitude and longitude. When this is not the case, the basic premise on which the algorithm was constructed breaks down, and the performance is generally poor. However, the algorithm is capable of producing accurate depth estimates for events which traditionally cause problems for standard location procedures, such as shallow earthquakes or explosions with only teleseismic observations of *P* arrivals available. In contrast to the epicentre–depth-separated algorithm, the space–time-separated algorithm is capable of producing reliable locations even when the correlation between depth

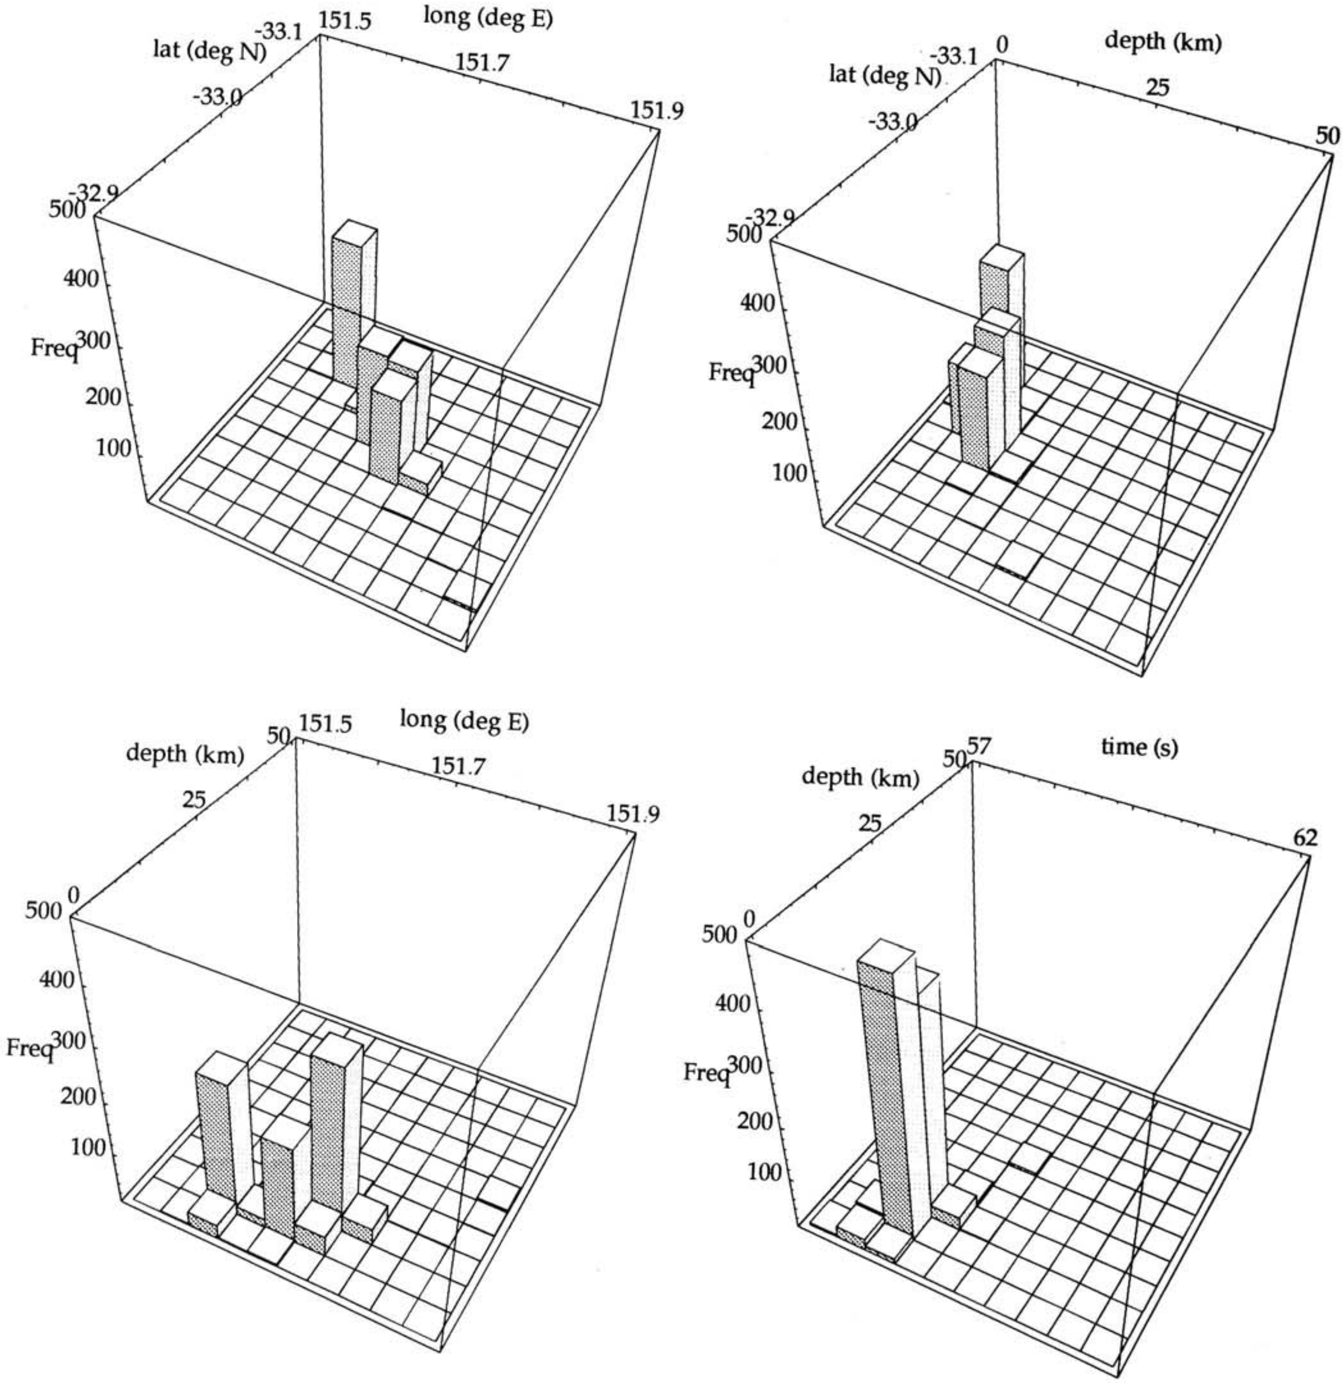

Figure 8. The frequency with which, after 1000 trials, the final hypocentral location found by the space–time-separated algorithm is in the range defined by the projection of locations onto the plane in question for the Newcastle earthquake. The strongly separated algorithms produce a tight clustering of hypocentres indicating a better performance that the standard genetic algorithm. Compare with Fig. 3 but note that the vertical scale here is twice that on Fig. 3.

and origin time is not strong. This arises because the epicentre–depth-separated algorithm *depends critically* on the weak correlation between depth and time, while the space–time-separated algorithm was designed to *avoid problems* associated with the strong dependence of depth on origin time. It assumes a worst-case scenario of strong correlation and is capable of producing good locations no matter how the depth and origin time are correlated.

In this paper we have demonstrated that incorporating problem-specific information into earthquake location with genetic algorithms will generally improve performance. This information can also be incorporated into a simulated-annealing algorithm (Billings 1994). The algorithm developed in Billings (1994) used a space–time separation together with information related to the long-wavelength unimodal structure of the misfit function. In that case also,

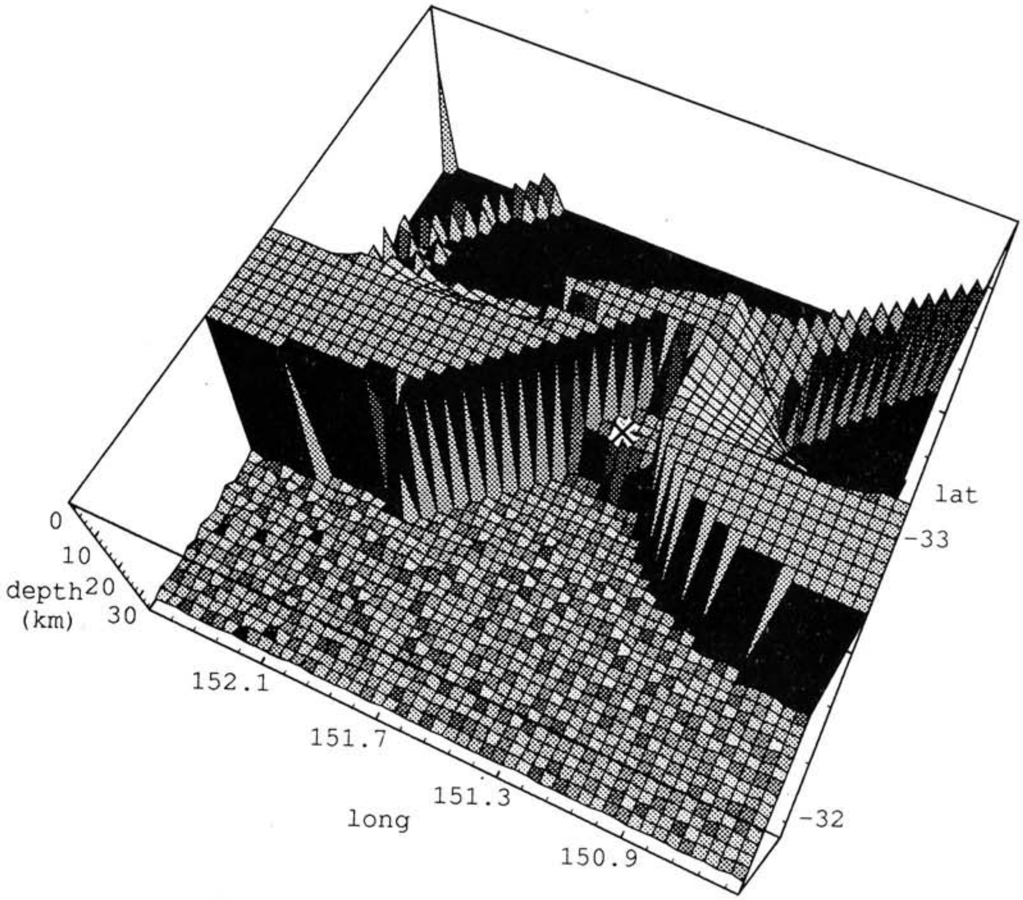

Figure 9. The variation in optimum depth with epicentre for the Newcastle earthquake. The nearest approximation to the global minimum is shown by a cross. There are two plateaus in the optimum depth, at 0 and 30 km. Neither are close to the depth of the global minimum which is at ~10 km. The particular range of latitude and longitude shown here represents the model space through which the genetic algorithm searches for the global minimum.

the modified algorithm was shown to improve performance substantially over a standard simulated-annealing search. The incorporation of problem-specific knowledge makes the genetic and simulated-annealing algorithms highly specialized and suited only to the particular problem at hand. However, the accurate estimation of earthquake hypocentres is so important that the development of specialized optimization routines is well justified.

ACKNOWLEDGMENTS

S.B. would like to thank the Research School of Earth Sciences for financial support during the preparation of this work, in the form of an A. L. Hales Honours Scholarship. The constructive comments of two anonymous referees, were much appreciated. This work was partially supported by DARPA/NMRO grant MDA972-91-J-1021.

REFERENCES

Anderson, K.R., 1981. Epicentral location using arrival time order, *Bull. seism. Soc. Am.*, **71**, 541–545.

Billings, S.D., 1994. Simulated annealing for earthquake location, *Geophys. J. Int.*, **118**, 680–692 (this issue).

Buland, R., 1976. The mechanics of locating earthquakes, *Bull. seism. Soc. Am.*, **66**, 173–187.

Buland, R., 1986. Uniform reduction error analysis, *Bull. seism. Soc. Am.*, **76**, 217–230.

Buland, R. & Chapman, C.H., 1983. The computation of seismic traveltimes, *Bull. seism. Soc. Am.*, **73**, 1271–1302.

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

Goldberg, D.E. & Deb, K., 1991. A comparative analysis of selection schemes used in genetic algorithms, in *Foundations of Genetic Algorithms*, pp. 230–241, ed. Rawlins, G.J.E. Morgan Kaufmann, San Mateo, CA.

Herrmann, R.B., 1979. FASTHYPO—A hypocentre location program, *Earthq. Notes*, **50**, 25–37.

Holland, J.H., 1975. *Adaption in Natural and Artificial Systems*, University of Michigan Press, Michigan.

Kennett, B.L.N., 1992. Locating oceanic earthquakes—the influence of regional models and location criteria, *Geophys. J. Int.*, **108**, 848–854.

Kennett, B.L.N. & Engdahl, E.R., 1991. Traveltimes for global earthquake location and phase identification, *Geophys. J. Int.*, **105**, 429–465.

Kennett, B.L.N. & Sambridge, M.S., 1992. Earthquake location: genetic algorithms for teleseisms, *Phys. Earth planet. Inter.*, **75**, 103–110.

Kirkpatrick, S.C., Gelatt, D. & Vecchi, M.P., 1982. Optimisation by simulated annealing, *IBM Research Report, RC* 9355, New York.

Kirkpatrick, S.C., Gelatt, D. & Vecchi, M.P., 1983. Optimisation by simulated annealing, *Science*, **220,** 671–680.

Lee, W.H.K. & Stewart, S.W., 1981. *Principles and Applications of Microearthquake Networks*, Academic Press, London.

Lienert, B.R. & Frazer, L.N., 1983. An improved earthquake location algorithm, *EOS, Trans. Am. geophys. Un.*, **64,** 267.

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

Sambridge, M.S. & Gallagher, K., 1993. Earthquake hypocentre location using genetic algorithms, *Bull. seism. Soc. Am.*, **83,** 1467–1491.

Sambridge, M.S. & Kennett, B.L.N., 1986. A novel method of hypocentre location, *Geophys. J. R. astr. Soc.*, **87,** 679–697.

Smith, E.G.C., 1976. Scaling the equations of condition to improve conditioning, *Bull. seism. Soc. Am.*, **66,** 2075–2081.

Stoffa, P.L. & Sen, M.K., 1991. Nonlinear multiparameter optimisation using genetic algorithms: Inversion of plane-wave seismograms, *Geophysics*, **11,** 1794–1810.

Whittle, P., 1971. *Optimisation under Constraints*, Wiley, London.