4

NA source inversion I - development of method

4.1 Introduction

The source parameters of moderate size earthquakes have commonly been determined using least-squares inversion of long period regional seismograms (e.g., Wallace et al., 1981; Wallace and Helmberger, 1982; Lay et al., 1994). This provides good constraint on the source mechanism, however the use of regional waveforms has a number of limitations, one of which is poor depth resolution. The use of teleseismic data allows good depth resolution and mechanism constraints, as well as having a number of other advantages such as being relatively insensitive to the Earth model used. However, until recently most methods for determining the source parameters of shallow teleseismic earthquakes have only been applied to large magnitude events (e.g., Langston, 1976; Langston et al., 1982; Nabelek, 1985). Such methods use a least-squares inversion of long period *P* and *SH* waveforms to determine the source depth, time function and fault orientation. Recent work on obtaining depth and mechanism estimates at far regional and teleseismic distances for smaller events includes that of Goldstein and Dodge (1999) who use *P*-waveform modelling.

In this chapter we outline a method to determine the source parameters (depth, mechanism and time function) of shallow, moderate size events using nonlinear waveform inversion of teleseismic data. We have shown that accurate synthetic seismograms at teleseismic distances can be calculated by using a modified reflectivity approach with a slowness integral adapted to the epicentral distance to the receiver (Section 2.5.2). This approach allows a full treatment of conversions and crustal reverberations but requires significant computational effort. In order to evaluate seismograms for varying source depth and mechanism rapidly, so that a fully nonlinear inversion method can be employed, we use a simple generalised ray scheme to model *P* and *S* and their corresponding surface reflections. We follow the method of Langston and Helmberger (1975) and adapt it to allow for different structures at the source and receiver ends. This provides a reasonable approximation without high computational cost, however, it is possible to use more accurate synthetics in the inversion when necessary; as for example in the case of problematic events where the simple synthetics are inadequate.

Unlike most previous source inversion methods, we make use of *SV* information in addition to *P* and *SH* information, and use short period or broadband velocity records instead of long period waveforms. For the inversion we use a direct search method called the neighbourhood algorithm (NA) (Sambridge, 1999a, b), which up until now has not been applied to the problem of source inversion. A relative amplitude approach is used for a number of reasons: firstly, the use of high frequency data requires a detailed knowledge of structure in order to accurately model absolute amplitudes, secondly, array beams are more likely to preserve relative amplitudes rather than absolute amplitudes. The source mechanism is represented in two different ways; the superposition of a double couple with an isotropic component, and a general moment tensor with six independent components. For both mechanism representations we can test the significance of any isotropic estimate by applying restrictions to the amount of isotropic component allowed in the inversion. We verify the inversion scheme with a number of synthetic tests. The use of a simple generalised ray scheme proves to be effective for inversion when compared with more accurate synthetic seismograms. Tests with synthetics for a perturbed source location demonstrate the relative robustness of the inversion scheme. We also examine the sensitivity of the inversion to variations in velocity structure.

4.2 The neighbourhood algorithm

We perform inversion for source depth, mechanism and time function using the neighbourhood algorithm (NA) (Sambridge, 1999a, b), which is a derivative-free, direct search method of inversion. The use of a fully nonlinear method of inversion has a number of advantages over a linearised inversion; it avoids possible dependence on the starting parameters, avoids potential numerical instability that can occur in iteratively applied linearised schemes, and also provides greater flexibility in the modelling. Since we are using fairly high frequency data it may be necessary to allow for a broad suite of crustal phases in order to accurately model the observed seismograms. The use of more complex synthetics in a linearised inversion scheme would make the partial derivatives difficult to calculate. However, the use of a derivative-free method such

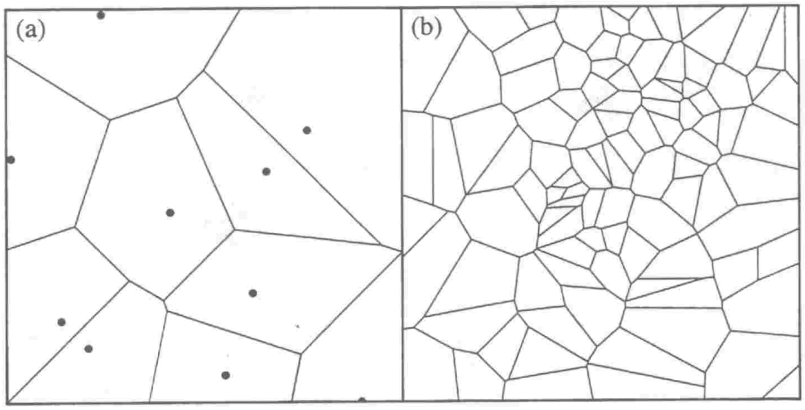

Fig. 4.1. (a) 10 quasi-uniform random points and their Voronoi cells. (b) The Voronoi cells about the first 100 points generated by the NA.

as the NA overcomes this problem. The NA is conceptually simple and its ease of use encourages us to apply it to the problem of waveform inversion for source parameters. We are then able to easily test various assumptions and explore the use of different earth structure models and seismogram calculation schemes if necessary, which would be more difficult if the inversion method relied on the calculation of partial derivatives.

4.2.1 Searching a parameter space

The NA is a direct search method of inversion which preferentially samples those regions of a multidimensional parameter space which have acceptable data fit. It has the ability to search efficiently by sampling simultaneously in different regions of parameter space, and is able to cope with multiple minima (Sambridge, 1998). The NA shares some characteristics with other nonlinear methods such as simulated annealing and genetic algorithms. Such methods often have a number of 'control' parameters which have to be tuned for each problem, whereas the NA requires just two control parameters which guide the algorithm in a conceptually simple manner. A further advantage of the NA over other direct search methods is that only the rank of the misfit function is used to compare models. This is of particular significance for seismic waveform inversion, as it avoids problems associated with scaling of the misfit function, and allows any type of user-defined misfit measure to be employed. These points are discussed in more detail in Sambridge (1999a).

The NA makes use of simple geometrical constructs known as Voronoi cells, which are nearest neighbour regions defined by a suitable distance norm (Voronoi, 1908).

Figure 4.1 provides an illustration of Voronoi cells for a two parameter problem. Figure 4.1(a) shows a set of 10 quasi-randomly distributed points and their corresponding Voronoi cells which define a neighbourhood about each point. This figure illustrates that the size and shape of each Voronoi cell is uniquely determined by the distribution of points. This geometrical property of Voronoi cells is exploited in the NA. At each stage of the inversion the parameter space is partioned into Voronoi cells, each of which is assigned a misfit value by performing forward modelling and calculating the misfit between the observations and predictions for each set of model parameters. Thus the Voronoi cells are used to create an approximate misfit surface at each iteration. The cells are then used to drive the search in parameter space by guiding the sampling towards those regions in which the current misfit is least. In the example shown in Figure 4.1 the NA distributes 20 new points randomly in the best 20 Voronoi cells at each iteration. The Voronoi cells are updated at each iteration, which forces sampling to be concentrated in the 'promising' regions identified from previous iterations. Figure 4.1(b) displays the Voronoi cells after five iterations of the NA, and we can see that the sampling is concentrated in certain regions of parameter space. These regions have a high density of points and so have Voronoi cells which are reduced in size. All the previous information obtained in the inversion is exploited in order to focus the search in on the regions with acceptable data fit. In this sense the algorithm is 'self-adaptive'.

The behaviour of the search algorithm is controlled by two parameters, n_s and n_r, and can be summarised by the following:

(i) First, an initial set of n_s models are generated randomly, and a misfit measure is calculated for each model.

(ii) Next, the n_r models with the lowest misfit are determined, and a uniform random walk is performed inside their Voronoi cells in order to generate a new set of n_s models.

(iii) The above steps are then repeated by calculating the misfit function for the most recently generated n_s models, and sampling inside the new Voronoi cells of the n_r models with lowest misfit. At each stage the size and shape of the Voronoi cells automatically adapt to the previously sampled models. This allows each successive iteration to concentrate sampling in regions of parameter space which have low data misfit.

The two control parameters, n_s and n_r, need to be tuned for each specific problem. For small (n_s, n_r) the algorithm is fairly localised in nature, whereas for larger (n_s,

n_r) the algorithm is more exploratory in nature. Note by definition $n_r \leq n_s$. For this application to waveform inversion we have found that broad sampling with a large n_r value gives the most satisfactory results. In practice, we also specify a maximum number of iterations and take the best fitting solution found as our preferred model.

4.2.2 Appraising the ensemble

Unlike most other inversion methods, the main objective of the NA is not to seek a single optimal model, but rather to generate an ensemble of models with acceptable data fit. The idea behind this is that the entire ensemble of models generated during the search stage contains useful information which can subsequently be analysed. The information acquired in the course of the inversion is used to guide a resampling of the parameter space to obtain measures of resolution and trade-off in the model parameters (see Sambridge (1999b) for more detail). This is performed within a Bayesian framework, and we refer the reader to Tarantola (1987) and Mosegaard and Tarantola (1995) for further information on this topic, rather than repeating material here. The Bayesian solution to the inverse problem is the posterior probability density function, which depends on the data, any prior information, and the statistics of all noise processes present. As yet we do not have adequate knowledge of the statistics of all noise processes present, due to the combination of model parameters being used in this case. Therefore we are currently unable to use such an approach to error analysis.

An alternative method has been developed by Farmer (2000), which attempts to map the complete region of acceptable solutions found by the NA. Some measure of whether a model is acceptable or not is needed, and this is represented by a certain misfit cut-off value. The models which satisfactorily fit the data (i.e. those with a misfit value below the misfit cut-off) are assigned equal misfit so that all acceptable models are treated equally by the algorithm. This results in concentrated sampling in and around the region of acceptable models. This method is discussed in more detail in Section 5.3.4, where we apply it to one of the events for which we have performed source inversion.

Even though global optimisation is not the primary objective of the NA, it has been shown to work well in this respect for both receiver function inversion (Sambridge, 1999a) and seismic event location (Sambridge and Kennett, 2001). Here we use the NA to seek a good match to seismograms recorded at teleseismic distances by varying the source parameters. We simply take the best fitting model found within the specified number of iterations as our estimate of the source depth, time function and mechanism.

4.3 Seismogram calculation scheme

Generalised ray contributions are generated for the direct waves (*P* and *S*) and their surface reflected phases (*pP, sP* and *pS, sS*) following the method of Langston and Helmberger (1975), which we have adapted to allow for different source and receiver structures. Below is a brief description of the method.

For the *P* wave response we first calculate the following expression for the near-source effects, including the free-surface reflections,

$$[\mathbf{\Sigma}]_P = [\Sigma_D^S]_P \mathrm{e}^{\mathrm{i}\omega T_P} + [\Sigma_U^S]_P [\mathbf{R}_F]_{PP} \mathrm{e}^{\mathrm{i}\omega T_{pP}} + [\Sigma_U^S]_S [\mathbf{R}_F]_{PS} \mathrm{e}^{\mathrm{i}\omega T_{sP}} \tag{4.1}$$

where Σ_U^S and Σ_D^S are the upward and downward source radiation terms (see Section 2.3), $\mathbf{R}_F$ is the free-surface reflection matrix near the source (Kennett, 1983, eqn 5.66), and T_P, T_{pP} and T_{sP} are the travel times for *P, pP, sP* respectively.

For the *S* wave response the following expressions are calculated,

$$[\mathbf{\Sigma}]_S = [\Sigma_D^S]_S \mathrm{e}^{\mathrm{i}\omega T_S} + [\Sigma_U^S]_P [\mathbf{R}_F]_{SP} \mathrm{e}^{\mathrm{i}\omega T_{pS}} + [\Sigma_U^S]_S [\mathbf{R}_F]_{SS} \mathrm{e}^{\mathrm{i}\omega T_{sS}} \tag{4.2}$$

$$[\mathbf{\Sigma}]_H = [\Sigma_D^S]_H \mathrm{e}^{\mathrm{i}\omega T_S} + [\Sigma_U^S]_H [\mathbf{R}_F]_{HH} \mathrm{e}^{\mathrm{i}\omega T_{sS}} \tag{4.3}$$

where T_S, T_{pS} and T_{sS} are the travel times for *S, pS, sS* respectively. These equations differ from those in Langston and Helmberger (1975), as we use wave functions normalised to unit energy transport in the vertical direction.

The displacement at the receiver corresponding to either the *P* or *S* response is then calculated using the following expression,

$$\bar{\mathbf{u}}(r, 0, \omega) = -\mathbf{W}_F\, \mathbf{\Sigma}(p, \omega)\, \mathrm{Q}(\omega)\, \mathrm{i}\omega M(\omega)\, \mathrm{e}^{-\mathrm{i}\frac{\pi}{2}} \tag{4.4}$$

which allows for a phase shift of $\pi/2$ due to complete reflection. The free-surface amplification factor at the receiver is given by $\mathbf{W}_F$, and the term $-\mathrm{i}\omega M(\omega)$ represents the farfield source time function, which we specify to be a trapezoid. At teleseismic distances and for shallow sources a single ray parameter is used, where p is the geometric slowness for the direct wave. The effect of attenuation is given by the following,

$$\mathrm{Q}(\omega) \approx \exp\{-\mathrm{i}\omega t_m^* \frac{1}{\pi} \ln\left(\frac{\omega}{2\pi}\right)\} \exp\{-\tfrac{1}{2}|\omega|(t_s^* + t_m^* + t_r^*)\} \tag{4.5}$$

where the effects of velocity dispersion are included in the mantle. The loss factors for the source, mantle and receiver structures (given by t_s^*, t_m^*, t_r^* respectively) are calculated using an attenuation model. In this thesis we use the *ak135* velocity model of Kennett et. al. (1995), with the corresponding attenuation profile of Montagner and Kennett (1996). We make an allowance for different source and receiver structures with a common mantle structure beneath a separation level at around 200 km depth.

4.4 Waveform inversion scheme

We perform waveform inversion using the NA for the following model parameters: source depth, source time function and source mechanism. The source time function is specified as a trapezoidal function in time, and is defined by the initial rise time, with the three time segments in the ratio 1:3:1. This is suitable for small events, but more complex forms with more parameters can be used for larger events. For the source mechanism we have experimented with two different styles of representation. The first is in terms of the superposition of a double-couple with an isotropic (explosive) component. The double-couple is specified by strike angle, dip angle and rake angle and the seismograms depend nonlinearly on each of these angles. The second representation is through a general moment tensor with six independent components (see e.g., Aki and Richards, 1980) that allows for a possible compensated linear vector dipole component in addition to double-couple and isotropic components. At fixed depth the seismograms depend linearly on the moment tensor components, but nonlinearity is introduced through variable source depth.

When we use the double-couple plus isotropic component mechanism we have six parameters to be determined, and are therefore working in a six-dimensional parameter space. For each of these parameters we have to specify upper and lower bounds to set up the search volume for the NA. For the strike angle we search between 0° and 360°, for the dip angle we search between 0° and 90°, and for the rake angle we search between 0° and 360°. The isotropic component is represented by an isotropic moment tensor, with a weighting factor between 0 and 5 to allow for a large range of values. In contrast, the moment tensor representation requires eight parameters and the six moment tensor components need to be able to take both positive and negative values. The search in this case will lie in an eight-dimensional parameter space.

For discrimination purposes we need to provide accurate constraints on the significance of any isotropic component in the source mechanism estimate. The source inversion can be performed with no restriction on the allowable isotropic component, or alternatively with a restriction imposed. If the source mechanism is represented as the superposition of a double-couple and an isotropic component, it is simple to impose the restriction that the source mechanism is a pure double-couple by limiting the range of values the isotropic component is allowed to take. When using a moment tensor representation for the source mechanism we can impose the restriction of zero trace (i.e. $M_{xx}+M_{yy}+M_{zz} = 0$), and consequently have a seven-dimensional parameter space. The fit of the two types of inversion can then be compared in order to determine the significance of any estimated isotropic component. For the case of an earthquake, if a large isotropic estimate is obtained for an unconstrained inversion then we would expect a lower misfit to be obtained when the isotropic component is restricted. On the other hand, for an explosion we would expect to obtain a lower misfit when no restriction is imposed on the isotropic component.

The forward modelling for each set of model parameters is performed using generalised ray theory as described in Section 4.3. We need to quantify the misfit between the observed and synthetic seismograms for each model. The model is sought that most closely matches the synthetics to the observed data, but we do not know the exact probability distribution for the data errors, therefore we cannot define an absolute measure of data misfit. Since we are using a derivative-free inversion scheme, based on only the rank of the misfit function, we can investigate the use of different misfit measures. We are not restricted to the conventional squared residual measure (i.e. a L_2 norm) which is commonly used for linear inversion. We perform tests with synthetic data to determine which misfit measure gives the most stable results, and initially consider the following misfit measures: normalised cross-correlation, and L_p norm measures with $p = 1, 1.5, 2$. An advantage of the L_1 and $L_{1.5}$ measures over the conventional squared residual measure is that they are robust and tolerant to outliers. In each case we apply a weighting factor to each station, depending on the signal to noise ratio for the particular record.

The normalised cross-correlation measure we use is given by,

$$M = \frac{1}{N_s \langle S \rangle} \sum_{n=1}^{N_s} (1 - C(u_n^{obs}, u_n^{syn})) S_n \tag{4.6}$$

where N_s is the number of stations, $C(u_n^{obs}, u_n^{syn})$ is the normalised cross-correlation between the observed and predicted seismograms at the nth station, calculated for an appropriate window, S_n is the signal to noise ratio at the nth station, and $\langle S \rangle = N_s^{-1} \sum_n S_n$ is the averaged signal to noise ratio across all the stations.

The L_p norm misfit measure we use is given by the following expression,

$$M_p = \left[\frac{1}{N_s \langle S \rangle} \sum_{n=1}^{N_s} S_n \int_{t_1}^{t_2} \mathrm{d}t \, |u_n^{obs}(t) - u_n^{syn}(t)|^p \right]^{1/p} . \tag{4.7}$$

Tests with synthetic data show that the most stable results are obtained when using a L_2 norm misfit measure. With a normalised cross-correlation measure it is possible to fit the waveforms at all but one station, for example, and thereby obtain a low misfit value. The normalised cross-correlation measure tends to favour those models which fit the data well, whereas the L_2 norm measure tends to penalise those models which do not fit the data. The L_2 norm measure appears to converge to the correct result in a moderate number of iterations, even when multiple minima are present. In contrast, the L_1 and $L_{1.5}$ norms tend to lead to local minima instead, in this case the more robust statistics are too tolerant of misfits in the seismograms and do not provide strong enough constraints. In light of these results with synthetic data, further investigations were performed on real data using L_p measures with $p > 2$ which are less tolerant to outliers. The results of these investigations are presented later in Section 5.3.3. We have used an L_2 norm misfit measure in the synthetic tests presented in this chapter and the illustrations with real data presented in Chapter 5.

Since we are using a relative amplitude approach in this study both the observed and synthetic seismograms are normalised to have unit maximum amplitude prior to inversion. This means that factors such as the scalar moment do not need to be included, and sensitivity to site amplification effects is minimised. In order to overcome the problem of travel time variability due to factors such as lateral heterogeneity and source location errors, the data and synthetics are aligned prior to inversion to avoid spurious fits. Each observed and synthetic seismogram is aligned to an arbitrary reference time using a cross correlation technique. The time window used in the inversion is 51.2 s for both P and S, and the window starts 20 s before the time of arrival for the observed examples, and 5 s before the time of arrival for the synthetic examples. A longer time window may be needed for deeper sources, which have a greater time separation between the various phases. The synthetic seismograms are calculated for a broad range of frequencies, between 0.01 Hz and 2 Hz. No filtering was performed for the synthetic tests, however the real data used in the observed examples were filtered using parameters discussed later. For the observational examples we use short period or broadband velocity records, and do not make any correction for instrument response.

The control parameters (n_s and n_r) used in this application of the NA were obtained after trials over a limited range of values. Those parameters which give consistently stable results with good convergence were chosen. We are able to achieve a good match to the seismic waveforms, with good correspondence between the best-fit model and the input model, even though exhaustive testing of the control parameters has not been undertaken. We have chosen a value of 16 for the parameter n_s, with a value of 8 for the parameter n_r. The NA is then initiated by generating 16 random models in parameter space, and the same number of models are generated at each subsequent iteration by resampling the 8 Voronoi cells within which the lowest current misfits are attained. The algorithm continues for 40 iterations, and tests show that performing more iterations does not significantly improve the best-fit model.

The following steps are performed in the inversion for source depth, mechanism and time function:

(i) First, an estimate of the source depth is needed to initiate the inversion and this can be obtained from a variety of sources. The depth estimate is usually obtained from bulletin sources if these are available, or alternatively, from a source location procedure which uses the NA for available arrival-time information (Sambridge and Kennett, 2001). We also need to take the estimates of distance, azimuth and expected arrival time at each station from the initial location estimate.

(ii) This hypocentral estimate is then linked to the waveform inversion, which uses the NA to search a range of source depths around the depth estimate, as well as searching a range of source time function parameters and source mechanisms.

(iii) Forward modelling is performed for each set of model parameters using generalised ray theory, and the misfit between the observed and predicted seismograms is calculated.

(iv) Once the algorithm has completed the specified number of iterations, the best fitting source depth, time function, and mechanism is taken as our preferred model.

4.5 Synthetic tests

In order to investigate the accuracy of the depth and source mechanism inversion scheme we perform a number of synthetic tests. We simulate an event for which a source mechanism determined from long-period recordings was reported in the Bulletin of the International Seismological Centre (ISC). The sample event occurred in western Iran (Figure 4.2), on 20 June 1990. We use the National Earthquake Information Centre (NEIC) moment tensor solution to calculate the synthetics, with a source depth of 17 km. A trapezoidal source time function is used with a rise time of 1.5 s. The 'observed' noise-free seismograms are calculated using both generalised ray theory, and the full representation described in Chapter 2, which includes the effects of shallow reverberations and conversions. We also perform tests where we introduce errors into the source location estimate, to determine how robust the inversion is in the presence of noise. In each case the best-fit model is compared to the input model to determine the accuracy of the inversion scheme. The inversions are performed using nine stations in the teleseismic distance range, whose locations are given in Table 4.1.

Fig. 4.2. Location of the western Iran event and station distribution used in the inversion.

Table 4.1. *Station locations relative to western Iran event.*

Station	Distance (°)	Azimuth (°)
KEV	34.97	347
TOL	41.41	291
SCP	88.72	323
SLR	65.47	201
COL	77.56	7
MAJO	68.41	60
BJI	51.26	65
LZH	43.29	74
KMI	46.70	89

4.5.1 Tests with simple seismograms

We undertake a varied set of tests with synthetic data to verify the effective operation of the neighbourhood algorithm. We first perform tests of the inversion scheme by

Table 4.2. *Comparison of inversion results with the input model, using a double couple plus isotropic term representation of the source mechanism. The 'observed' seismograms have been calculated using the simple synthetic scheme.*

Type	Depth (km)	Source function (s)	Strike (°)	Dip (°)	Rake (°)	Isotropic weighting	misfit
Model (fault plane 1)	17.0	1.5	202	38	156	0.0	
Model (fault plane 2)			311	76	54		
P inversion	16.7	1.6	120	88	254	0.8	0.091
Constrained P inversion	17.1	1.5	197	37	155	0.0	0.006
S inversion (T)	17.2	1.5	113	79	302	2.0	0.003
S inversion (R & T)	16.7	1.6	197	30	155	0.1	0.008

calculating 'observed' seismograms using the simplified approximations described in Section 4.3. We generate seismograms for each station using the epicentral distances and azimuths shown in Table 4.1 with the NEIC mechanism and the *ak135* reference model. We then attempt to match these 'observed' seismograms to the predicted seismograms with the aid of the neighbourhood algorithm. The source inversion tests with simple seismograms are obtained using a representation of the source mechanism in terms of a double-couple plus an isotropic component. The inversion results using *P* and *S* waveforms separately are given in Table 4.2, and are illustrated in Figures 4.3, 4.4, 4.5 and 4.6.

The first set of inversions are performed using just the *P* waveforms as recorded on the vertical component. For an unconstrained inversion quite a large isotropic component is obtained (Table 4.2 and Figure 4.3), due to the limited sampling of the source radiation pattern with the available station distribution. The fault plane parameters are poorly constrained due to the dominance of the isotropic component, however, the source depth and time function are well constrained. Although the predicted mechanism differs markedly from the input mechanism, the pattern sampled by the stations is well represented and there is a good waveform fit (Figure 4.3). If we impose the restriction of zero isotropic component and perform the inversion again, the resulting source mechanism is a much better representation of the NEIC mechanism (Figure 4.4), with fault plane parameters which are close to the input values (Table 4.2). The seismogram fit is improved, and a somewhat lower misfit value is obtained, which indicates that a pure slip dislocation (an "earthquake") is compatible with the data.

The *S* wave group provides a somewhat different sampling of the source radiation pattern and it is therefore useful to examine how an inversion based solely on *S* can

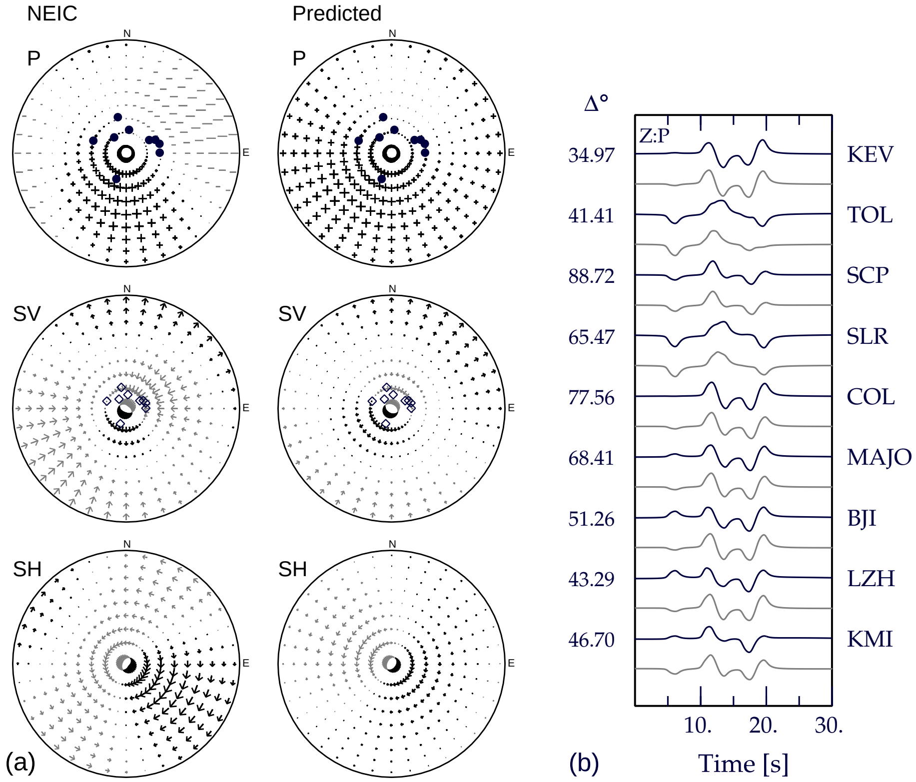

Fig. 4.3. Results of a *P* only waveform inversion for the simulated western Iran event, using a double couple plus isotropic component representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

perform. Commonly only the transverse component of *S* is used in source inversion (e.g., Nabelek, 1985), however, the use of *SH* waveforms alone cannot provide any constraints on the isotropic component. This is illustrated in Figure 4.5, where a very large isotropic component is obtained when we undertake an inversion using just the *SH* waveforms. Again, the fault plane parameters are poorly constrained due to the dominance of the isotropic component, however the source depth and time function are well constrained (Table 4.2). Even though the predicted mechanism differs significantly from the input model, the seismogram fit is excellent (Figure 4.5). When we use both the radial and transverse components of the *S* waveforms, and so sample both *SV* and *SH* information, we obtain a much better result from the inversion (Figure 4.6), since now we have much more information on the radiation pattern. The isotropic component is negligible, and the source mechanism is well constrained, along with the

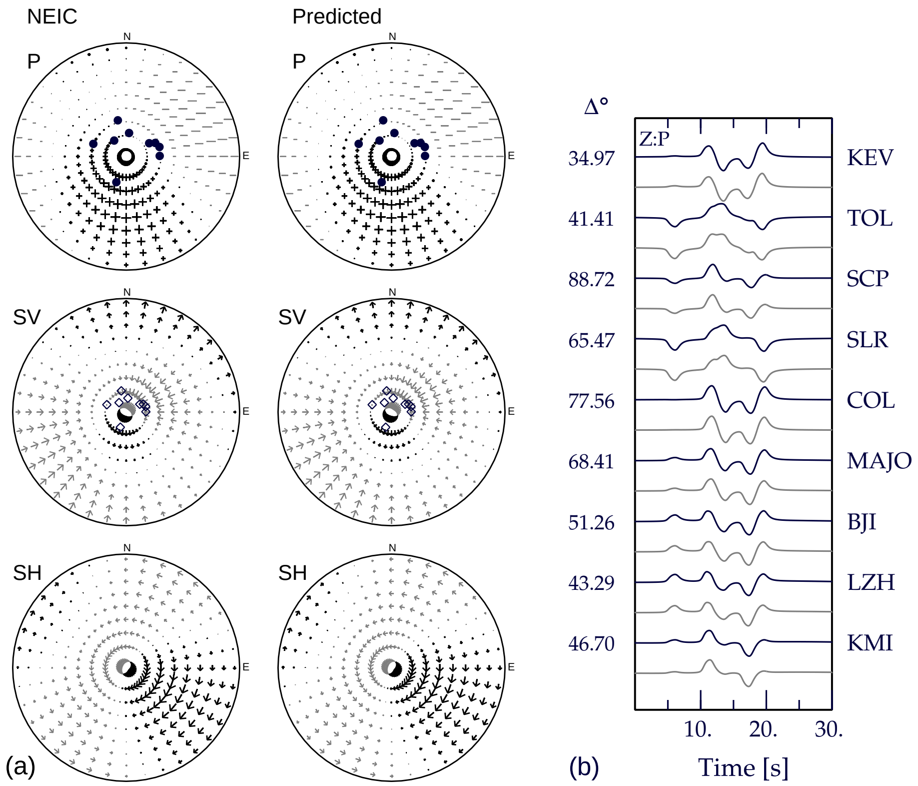

Fig. 4.4. Results of a *P* only waveform inversion for the simulated western Iran event, using a double couple representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

source depth and time function (Table 4.2). This shows that the inclusion of *SV* wave data can provide the extra information needed to constrain the source mechanism.

4.5.2 Tests with complex seismograms

We have shown that good results are obtained using simple synthetics, therefore we proceed with a more realistic set of tests where we calculate the 'observed' seismograms using the full teleseismic representation described in Chapter 2. The use of the more complete calculation scheme introduces some noticeable changes in the character of some of the waveform segments, as we now include shallow reverberations and conversions. We continue to use the simplified approximations in the inversion, because of the speed of computation, but now the 'observations' and 'predictions' are generated using very different schemes which provides a more realistic test of the inversion algorithm.

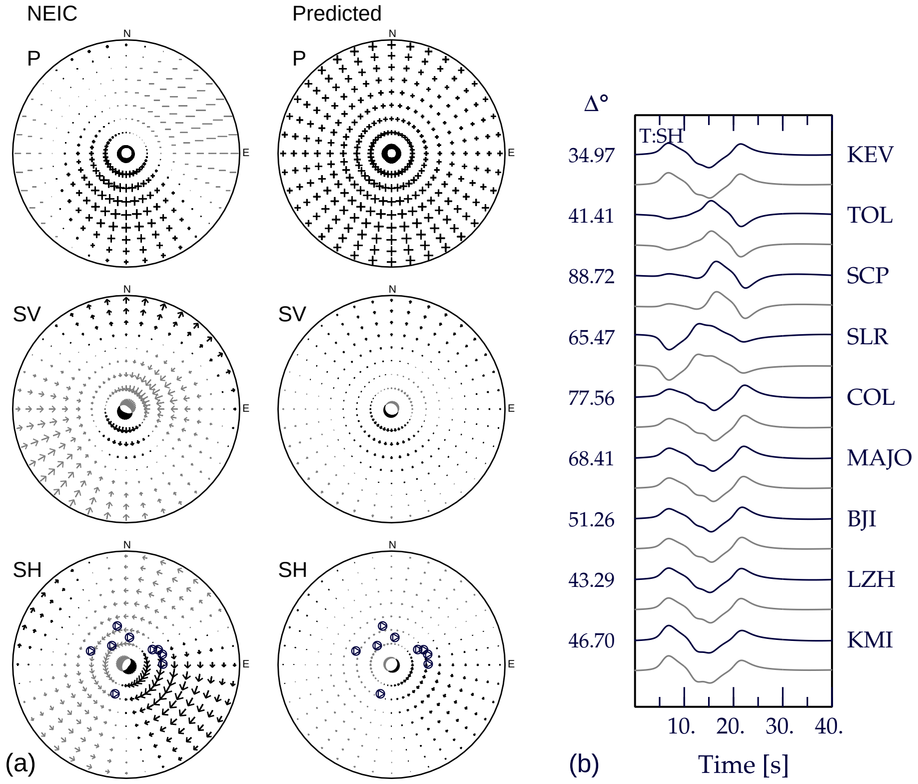

Fig. 4.5. Results of an inversion using just the *S* waveforms as recorded on the transverse component for the simulated western Iran event, using a double couple plus isotropic component representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms for the transverse component of *S*. The epicentral distances are displayed on the left, and the station names on the right.

Table 4.3. *Comparison of inversion results with the input model, using a double couple plus isotropic term representation of the source mechanism. The 'observed' seismograms have been calculated using the full teleseismic representation.*

Type	Depth (km)	Source function (s)	Strike (°)	Dip (°)	Rake (°)	Isotropic weighting	misfit
Model (fault plane 1)	17.0	1.5	202	38	156	0.0	
Model (fault plane 2)			311	76	54		
P and S inversion	17.8	1.5	119	87	301	0.5	0.17
Constrained P and S inversion	17.7	1.6	312	73	61	0.0	0.12

We perform joint inversions of both *P* and *S* waveforms for the new 'dataset'. We use the vertical component of *P*, and both the radial and transverse components of *S*. The combined *S* wave data were given half the weighting of the *P* wave data when

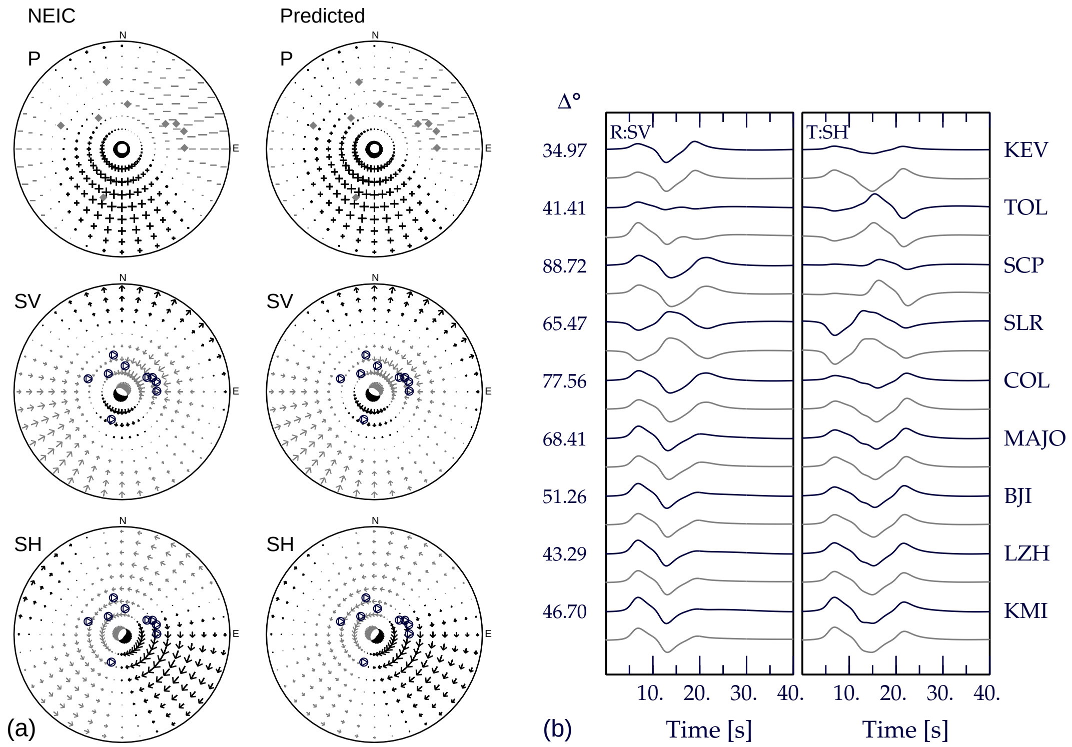

Fig. 4.6. Results of a *S* only waveform inversion for the simulated western Iran event, using a double couple plus isotropic component representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms. The radial (left) and transverse (right) components of *S* are shown. The epicentral distances are displayed on the left, and the station names on the right.

calculating the misfit function. We experiment with both representations of the seismic source in the inversions, with differing dimensionality of the parameter space. The inversion results for a double couple plus isotropic component representation are given in Table 4.3, and illustrated in Figures 4.7 and 4.8, and those for a moment tensor representation are given in Table 4.4 and Figure 4.9.

The first two tests use the representation of the source mechanism in terms of a double-couple plus an isotropic component. Firstly, an unconstrained joint inversion is performed, in which a significant isotropic component is introduced with a substantial departure from the NEIC solution (Figure 4.7). The match to the seismic waveforms is good and as a result the source depth and time function are well constrained (Table 4.3). We observe from Figure 4.7 that the main differences in the *P*, *SV* and *SH* radiation patterns occur well away from the limited portion sampled by the teleseismic 'observations' for this shallow source. We can check whether the data is compatible with a pure slip dislocation source represented by a pure double-couple,

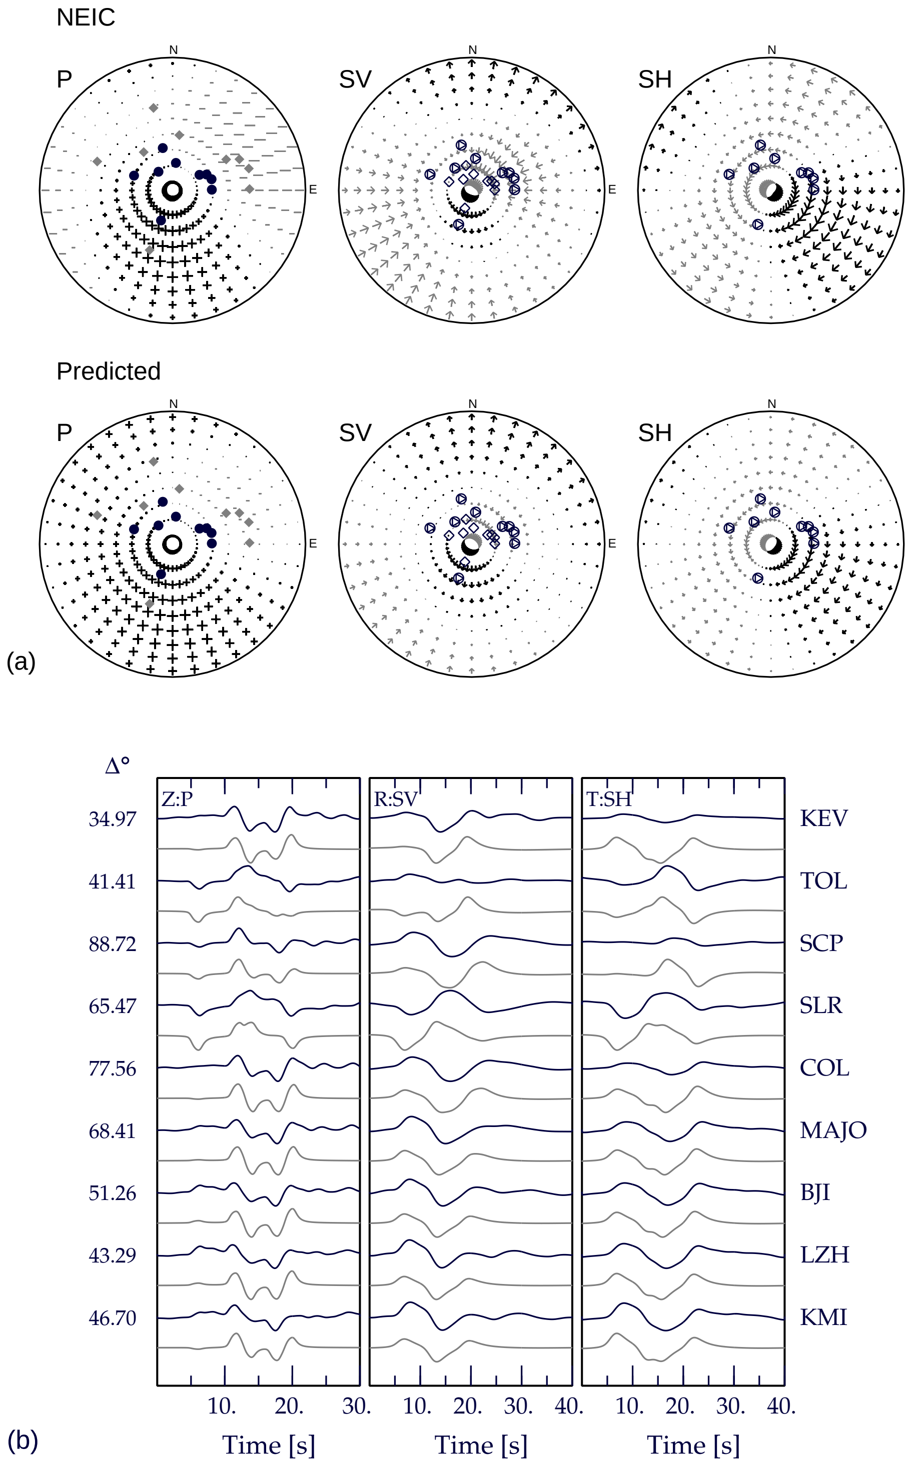

Fig. 4.7. Results of a joint *P* and *S* waveform inversion for the simulated western Iran event, using a double couple plus isotropic component representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (top), and the predicted source mechanism obtained from the inversion (bottom). (b) Comparison between the complex 'observed' (black traces) and predicted (grey traces) seismograms. The vertical component of *P* (left), and the radial (middle) and transverse (right) components of *S* are shown. The epicentral distances are displayed on the left, and the station names on the right.

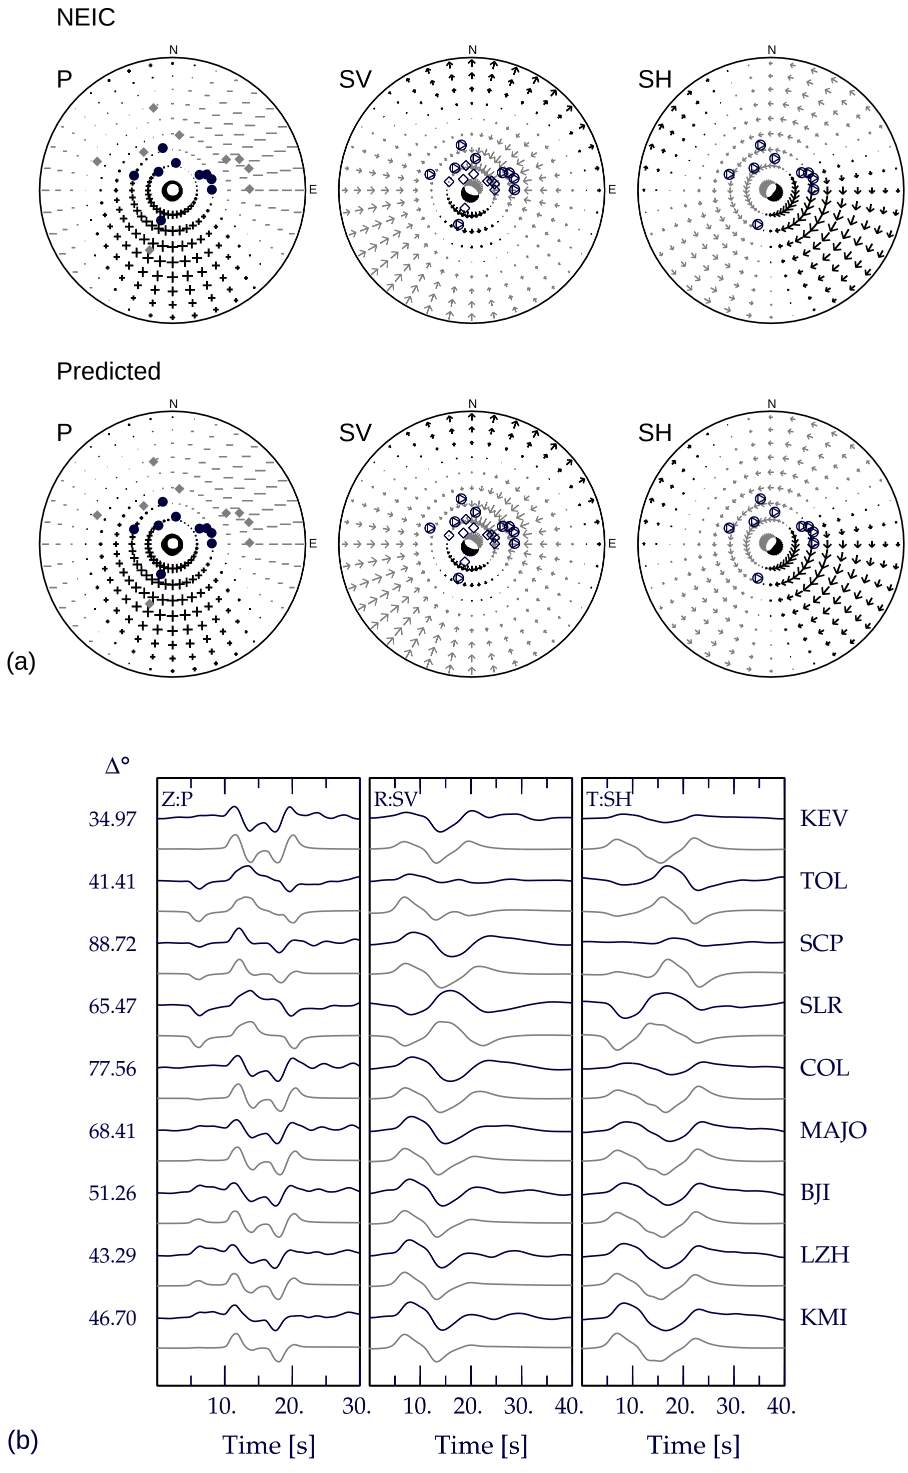

Fig. 4.8. Results of a joint *P* and *S* waveform inversion for the simulated western Iran event, using a double couple representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (top), and the predicted source mechanism obtained from the inversion (bottom). (b) Comparison between the complex 'observed' (black traces) and predicted (grey traces) seismograms. The vertical component of *P* (left), and the radial (middle) and transverse (right) components of *S* are shown. The epicentral distances are displayed on the left, and the station names on the right.

Table 4.4. *Comparison of inversion results with the input model, using a moment tensor representation of the source mechanism. The 'observed' seismograms have been calculated using the full representation.*

Type	Depth (km)	Source function (s)	M_{xx}	M_{yy}	M_{zz}	M_{xy}	M_{xz}	M_{yz}	misfit
Model	17.0	1.5	0.34	-0.73	0.39	-0.27	-0.63	-0.36	
P and S inversion	16.8	1.7	0.61	-0.62	0.67	-0.24	-0.85	-0.51	0.14

and undertake a second inversion with the isotropic component constrained to zero (effectively a five parameter space). The resulting source mechanism is a much better representation of the NEIC mechanism (Figure 4.8) and the strike, dip and rake angles are well constrained, with values which are very close to the input values (Table 4.3). The inversion results for the source depth and time function agree well with the original input values. A lower misfit is obtained, and the waveform match is improved (Figure 4.8) when the isotropic component is constrained to zero and thus the pure double-couple model is to be favoured.

Using the same 'dataset' as the last two tests we carry out a joint *P* and *S* waveform inversion using an eight parameter space where the source mechanism is represented in terms of six independent moment tensor components. The source depth and time function fit well when compared with the input model (Table 4.4). A comparison between the 'observed' and predicted seismograms is shown in Figure 4.9, and an excellent waveform match is achieved. It is difficult to judge immediately the comparison of the moment tensor components, however, plotting the radiation patterns (Figure 4.9) reveals that the recovered source mechanism is very similar to the 'observed' mechanism except for a small isotropic component. In order to compare the isotropic components from different inversions we calculate the normalised trace of the moment tensor. This gives a value of 0.4 for the isotropic component, which is smaller than that obtained using the alternative source mechanism representation (which gives a value of 1.0). This is probably because the moment tensor representation allows for a compensated linear vector dipole component, whereas the alternative representation does not, therefore requiring a larger isotropic component in order to fit the data. We note that the misfit level achieved is slightly better than for the unconstrained *P* and *S* inversion with the double-couple and isotropic component.

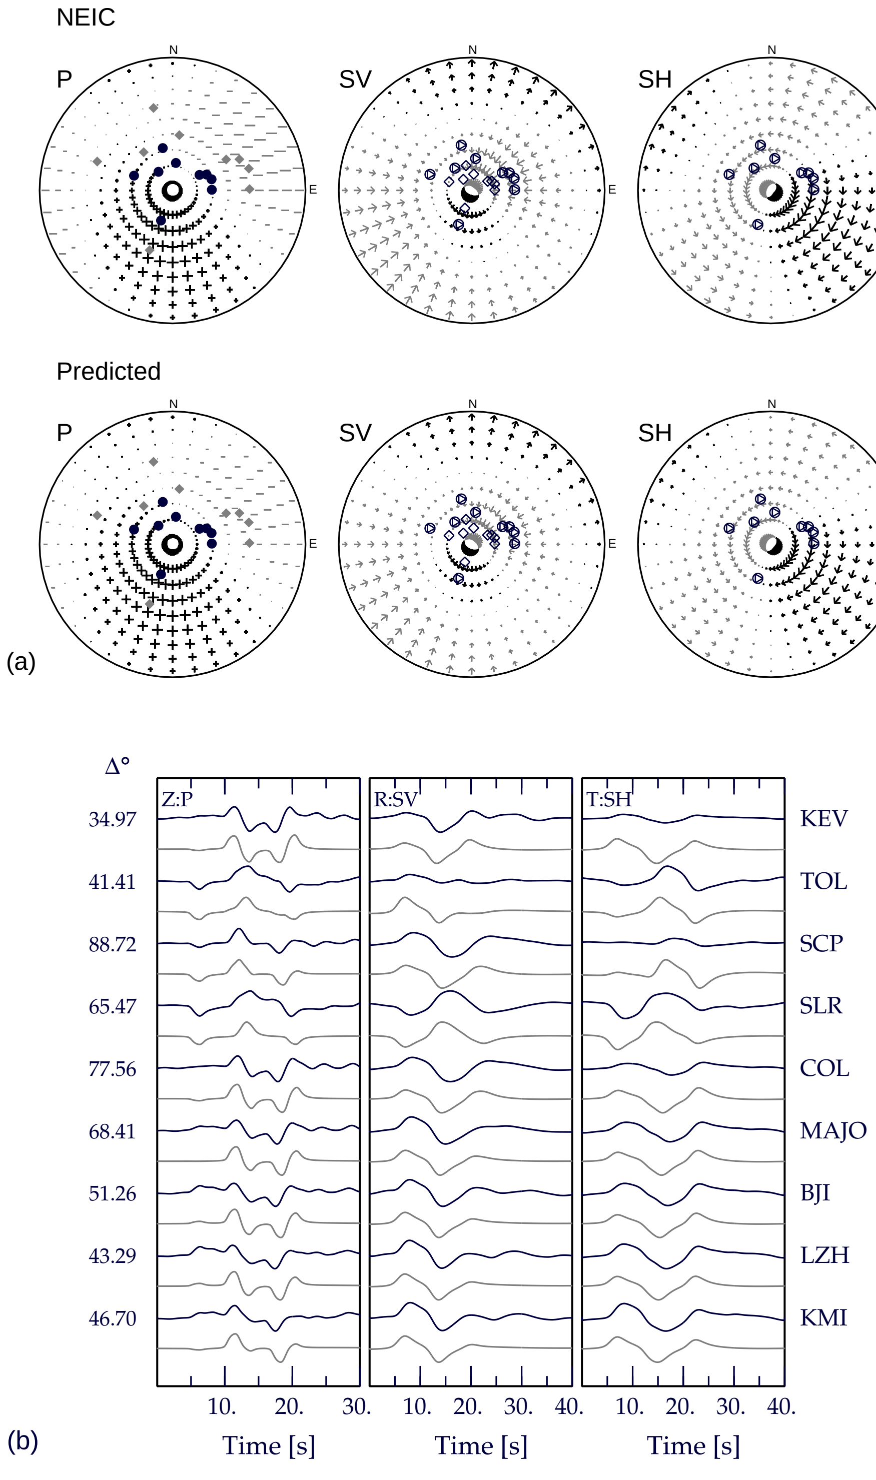

Fig. 4.9. Results of a joint *P* and *S* waveform inversion for the simulated western Iran event, using a moment tensor representation of the source. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the NEIC moment tensor solution (top), and the predicted source mechanism obtained from the inversion (bottom). (b) Comparison between the complex 'observed' (black traces) and predicted (grey traces) seismograms. The vertical component of *P* (left), and the radial (middle) and transverse (right) components of *S* are shown. The epicentral distances are displayed on the left, and the station names on the right.

4.5.3 Perturbed source location

In this set of synthetic tests we perturb the source location, thus providing a more realistic test of the inversion scheme in the presence of errors and noise. We use three test cases, where the latitude and longitude of the source location are perturbed by random values in the range ±0.05°, and the source depth is perturbed by random values in the range ±5 km. These values are chosen to reflect the errors in the source location for this well constrained event, as reported in the Bulletin of the ISC.

The perturbed source location is used to calculate the 'observed' seismograms, while the unperturbed source location (Table 4.1) is used to calculate the predicted seismograms. In this way we can simulate the effect of errors in the source location, where the actual hypocenter differs from the estimated hypocenter that is used in the waveform inversion. The 'observed' seismograms are calculated using the full teleseismic representation. The inversion is performed in two ways; firstly with the waveforms shifted in time in order to align the direct arrivals on both the observed and predicted seismograms, and secondly without correcting for the time shifts introduced by the source location perturbation.

In the first set of tests we represent the source mechanism in terms of a double-couple plus an isotropic component, and perform inversions using P wave data as recorded on the vertical component. Inversions are performed with three different source perturbations, for both aligned and unaligned data. The source depths and time functions recovered from the various inversions match the input values well. For the aligned data the recovered source depths are within ±0.9 km of the input values, and the source time functions are within ±0.06 s. For the unaligned data we match the source depths to within ±2.3 km, and the source time functions to within ±0.06 s. In most cases a large isotropic component is obtained due to the lack of constraint on this parameter from the use of P wave data alone, with the given station distribution. Therefore the strike, dip and rake angles are poorly constrained due to the dominance of the isotropic component in these cases. One test case with unaligned data obtained a fairly small isotropic component, and so the strike, dip and rake angles are well constrained. Two out of the three test cases obtained lower misfit values with unaligned data than with aligned data.

We use one of our test cases to provide an illustration of inversion with a perturbed source location. The source location has been perturbed by -0.05° in latitude, 0.02° in longitude, and 1.8 km in source depth (giving a source depth of 18.8 km). We use a moment tensor representation of the source mechanism, and perform a joint inversion

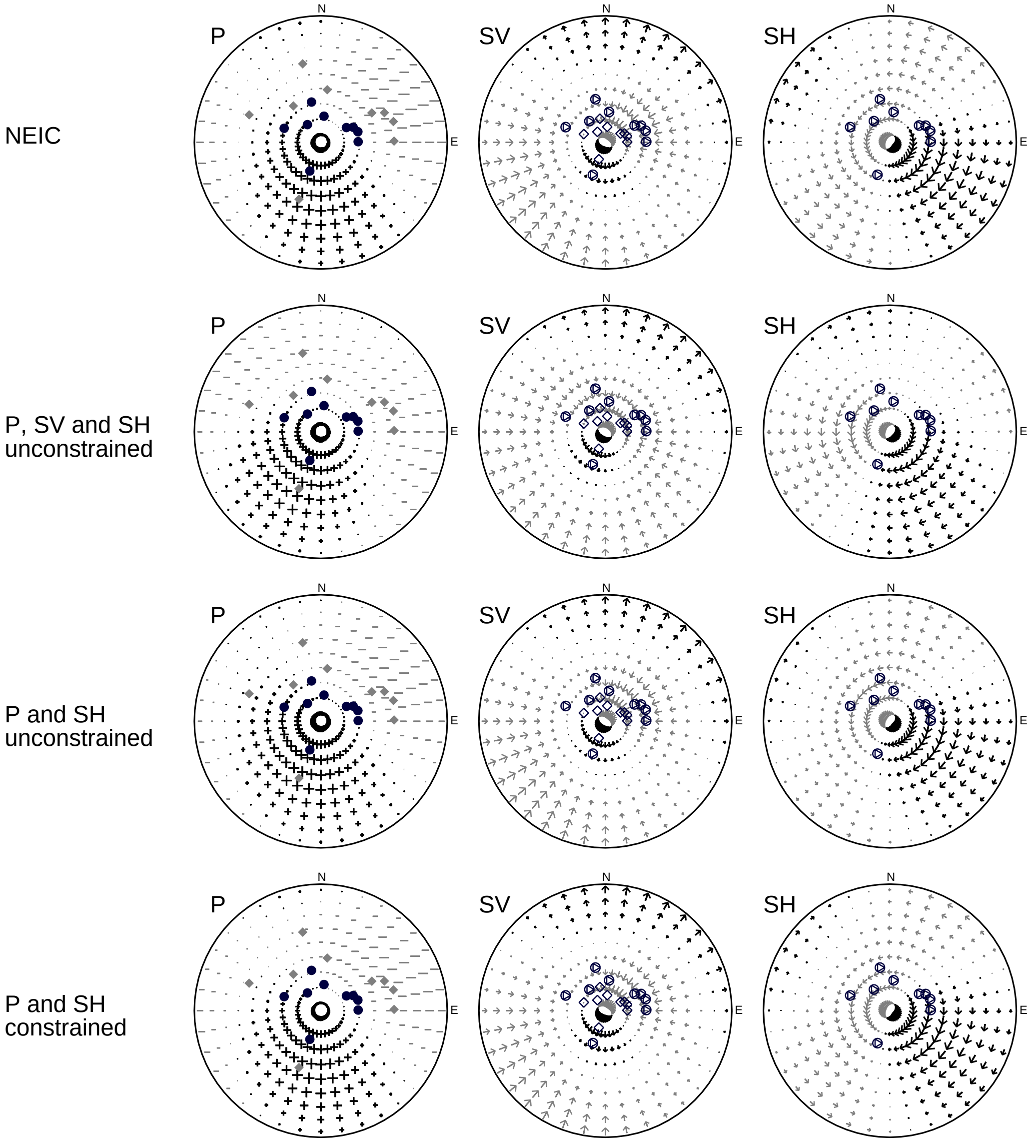

Fig. 4.10. (top) Source radiation patterns for *P*, *SV* and *SH* waves for the western Iran event using the NEIC moment tensor solution. Lower three sets of radiation patterns are the result of various inversions (described in the text). The plotted symbols represent the average take-off angle and azimuth for various phases; solid circles for *P*, solid triangles for *pP*, solid (grey) diamonds for *pS*; open circles for *S*, open triangles for *sS*, and open diamonds for *sP*.

of both *P* and *S* wave data. We do not correct for the time shifts introduced by the perturbation, so that timing errors of up to 0.5 s are present.

We first perform an inversion using *P*, *SV* and *SH* wave data with a moment tensor representation and an unconstrained isotropic component. The resulting source depth and time function (19.4 km and 1.5 s respectively) match the input values well. Plotting the radiation patterns reveals a very close match to the input source mechanism (Figure

4.10), with a negligible isotropic component. This is an ideal case, assuming we have *P*, *SV* and *SH* information available, however this may not always be so. We perform another inversion using only *P* and *SH* information to simulate a more realistic case. The main difference now is that we have less constraint on the isotropic component due to the loss of *SV* information. In fact, we now obtain a significant isotropic component, with a normalised trace of 0.2. The recovered source mechanism is a fairly good representation of the input mechanism (Figure 4.10), even though it contains a significant isotropic component. The resulting source depth of 19.2 km matches the model well, as does the recovered source time function of 1.6 s. A comparison between the 'observed' seismograms and the predicted seismograms calculated from the best fit model is shown in Figure 4.11. A good match is obtained for both the vertical component of *P* and the transverse component of *S*, as would be expected due to the close match between the recovered model parameters and the input parameters.

In order to determine whether the estimated isotropic component is real, or just an artifact of the inversion, we perform the same inversion again but with the constraint of zero trace imposed. Thus we have a seven-dimensional parameter space to search, as we now have only five independent moment tensor components to determine, rather than six. As can be seen in Figure 4.10 the recovered source mechanism provides a better match to the input mechanism than the previous inversion with no constraint on the isotropic component. Again the source depth and time function are well constrained with recovered values of 19.0 km and 1.6 s respectively. A slightly better seismogram fit is achieved (Figure 4.11), and consequently a lower misfit value is obtained (0.179 compared to 0.194). Thus it seems that the estimated isotropic component was an artifact of the inversion, and is not significant.

4.5.4 Variations in velocity structure

We perform a set of synthetic tests to investigate the sensitivity of the inversion to variations in velocity structure. For this set of tests we simulate an event which occurred off the east coast of Honshu, Japan on 14 May 1998 (see Section 5.2 for details of this event). The 'observed' seismograms are simulated for ten stations in the teleseismic distance range, using the Harvard CMT solution for this event with a depth estimate of 19 km. A trapezoidal source time function is used with a rise time of 0.3 s. We investigate the sensitivity of the inversion to variations in the *ak135* velocity model, and the presence of sediments beneath each station. These velocity models are used to calculate the 'observed' seismograms which are then inverted in the usual way, using only *P* wave information.

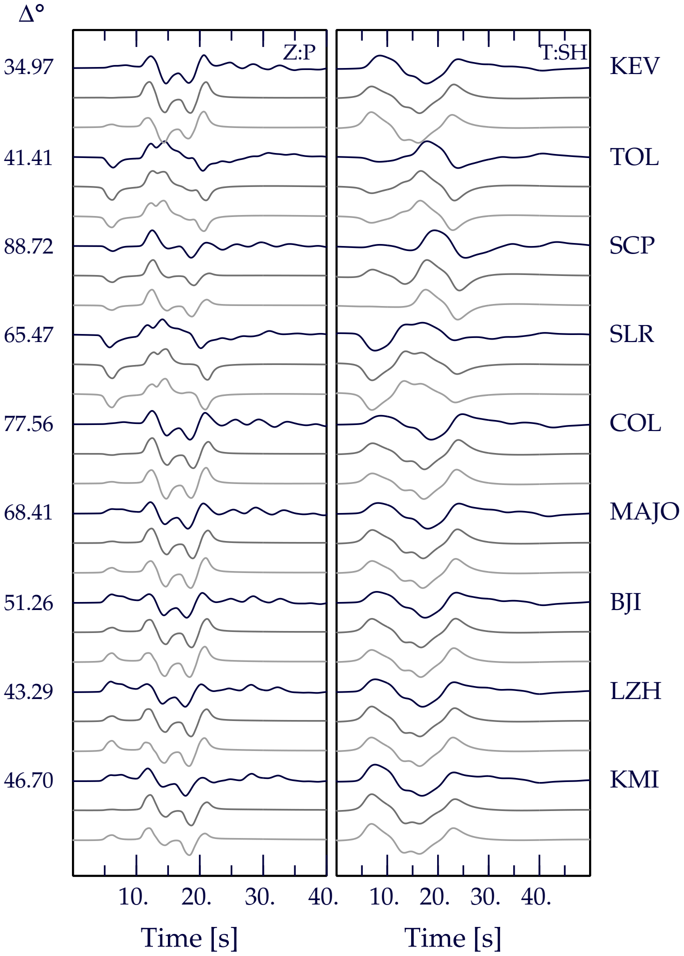

Fig. 4.11. Comparison between the complex 'observed' (black traces) and predicted (grey traces) seismograms for the western Iran event with a perturbed source location, and a moment tensor representation of the source mechanism. The inversion was performed with both an unconstrained (dark grey traces) and constrained (light grey traces) isotropic component. The vertical component of *P* (left), and the transverse component of *S* (right) are shown. The epicentral distances are displayed on the left, and the station names on the right.

We test the sensitivity of the inversion to variations in the *ak135* velocity model by systematically increasing/decreasing the velocities in each layer of the model by 3%. These modified velocity models are used to calculate the 'observed' seismograms using the simple scheme, and are then inverted using only *P* wave information. An inversion is also performed for 'observed' seismograms calculated using the *ak135* model in order to provide a comparison, and the results of this inversion are shown in Figure 4.12. The recovered source depth is 18.9 km which is very close to the true depth of 19 km, and the recovered rise time parameter is 0.3 s. The source mechanism is a good representation of the original mechanism given that we are only

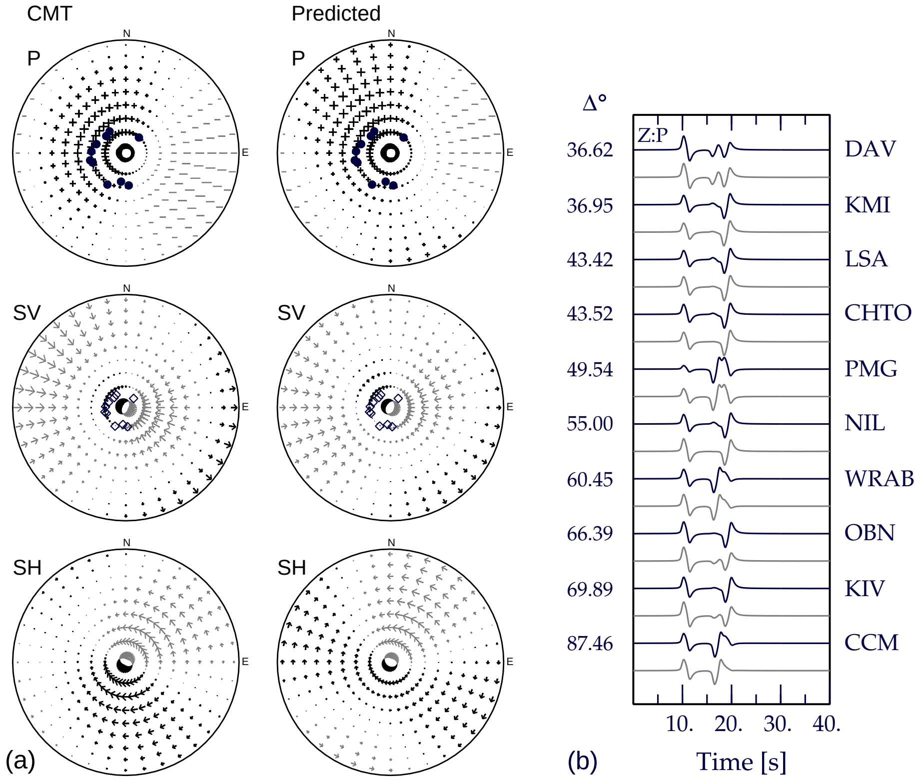

Fig. 4.12. Results of a *P* only waveform inversion for the Honshu event, simulated using the *ak135* model. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the CMT solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the ‘observed’ (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

using *P* information, and the seismogram match is good. Figure 4.13 displays the results of an inversion where the ‘observed’ seismograms are calculated using the *ak135* model with a 3% decrease in the velocity of each layer. The recovered source depth of 19.7 km is slightly deeper than the true depth, as would be expected, but it is not significantly so. The recovered rise time parameter is 0.3 s, and the recovered source mechanism is close to the original mechanism. The synthetic seismograms match the ‘observed’ seismograms very well. Figure 4.14 displays the results of an inversion where the ‘observed’ seismograms are calculated using the *ak135* model with a 3% increase in the velocity of each layer. Again the recovered mechanism is a good representation of the original mechanism. As would be expected, a slightly shallow source depth of 18.1 km is obtained, but it is within 1 km of the true depth. A rise time parameter of 0.3 s is recovered, and the seismogram fit is good, though there are noticeable differences at some of the stations. These tests provide an extreme case, with a systematic increase/decrease in the velocity of every layer. In the real earth we

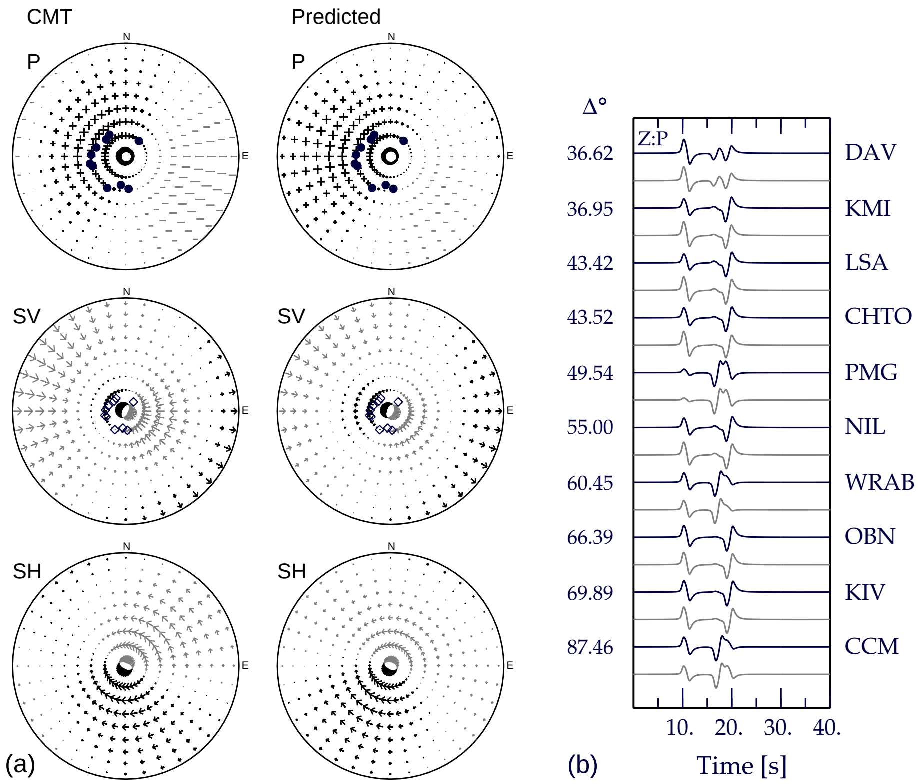

Fig. 4.13. Results of a *P* only waveform inversion for the Honshu event, simulated using the *ak135* model with a 3% decrease in the velocity of each layer. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the CMT solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

would not expect such systematic velocity fluctuations, and therefore any differences would be effectively averaged out along the entire ray path. Even so, these extreme tests have shown that both the source depth and mechanism are relatively insensitive to the velocity model used.

We test the sensitivity of the inversion to the presence of shallow sediments by including a low velocity sedimentary layer beneath each station. The 'observed' seismograms are calculated using the full teleseismic representation which allows for the effects of reverberations in the shallow source and receiver structures. These 'observed' seismograms are then inverted in the usual way, using only *P* wave information. We have performed tests with varying sediment thicknesses (1 km, 2 km and 5 km) and in each case the recovered source depth is within 1 km of the input depth, and the recovered source mechanism is a good representation of the original mechanism. The seismogram fit is pretty good in general, though the presence of

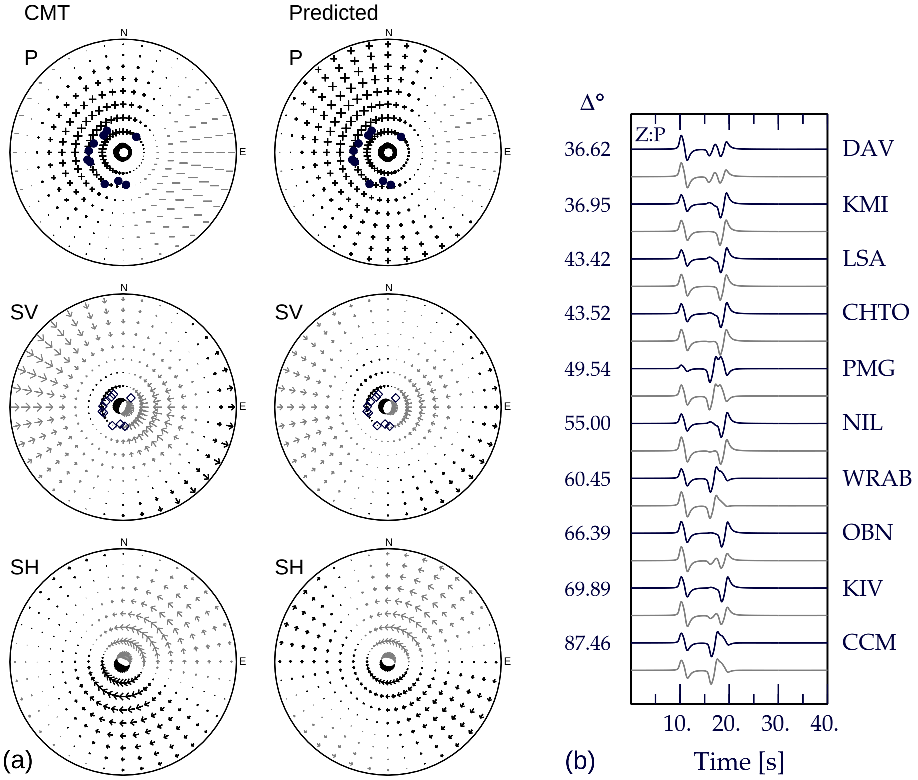

Fig. 4.14. Results of a *P* only waveform inversion for the Honshu event, simulated using the *ak135* model with a 3% increase in the velocity of each layer. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the CMT solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the ‘observed’ (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

sediments means that some of the arrivals at some of the stations are not fit very well, particularly for a thickness of 5 km. We display the inversion results for a 2 km thick sedimentary layer in Figure 4.15. It is clear that the presence of the sedimentary layer changes the pulse shape (see for example Figure 4.12 for comparison), and produces quite large reverberations in the ‘observed’ seismograms. Even so, the inversion works well, with a recovered source depth of 19.1 km. The recovered source mechanism fits the observed mechanism reasonably well. Thus the inversion is robust, even in the presence of a surface sedimentary layer beneath each receiver. This is an extreme test case, as in reality you would not expect sediments beneath every station. The more likely scenario is that surface sediments would be present under just a few of the stations.

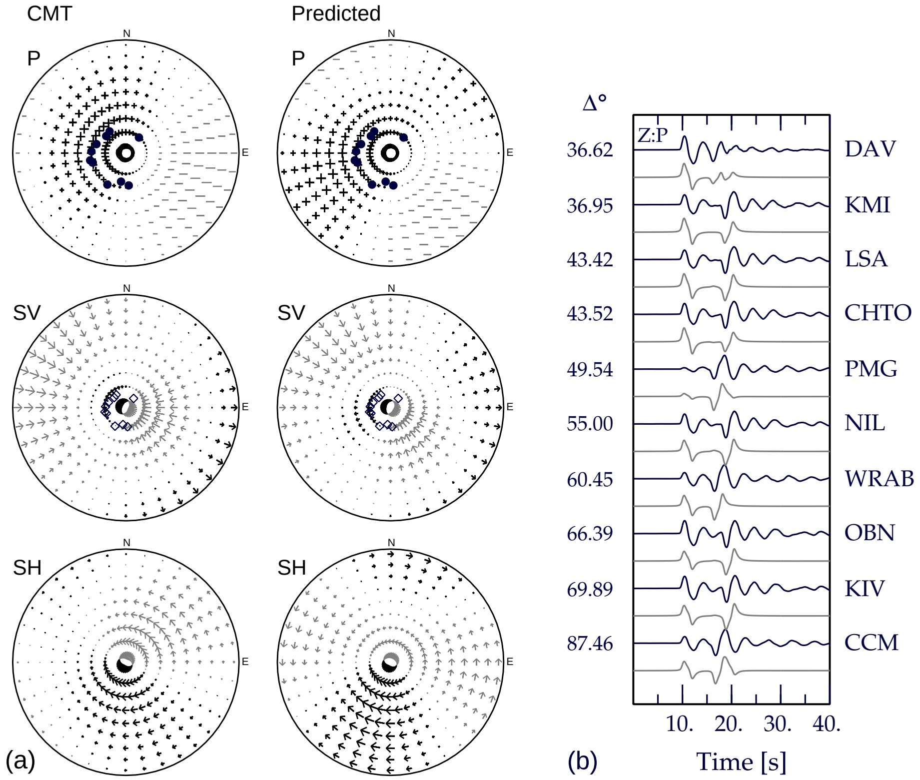

Fig. 4.15. Results of a *P* only waveform inversion for the Honshu event, simulated using the *ak135* model with a 2 km thick sedimentary layer beneath each receiver. (a) Source radiation patterns for *P*, *SV* and *SH* waves using the CMT solution (left), and the predicted source mechanism obtained from the inversion (right). (b) Comparison between the 'observed' (black traces) and predicted (grey traces) seismograms for the vertical component of *P*. The epicentral distances are displayed on the left, and the station names on the right.

4.6 Discussion

The synthetic inversion tests we have undertaken demonstrate that our inversion scheme using the neighbourhood algorithm is able to achieve good matches to both simple noise-free seismograms, and more accurate noise-free seismograms. We have deliberately limited the number of stations used in the synthetic tests, and so need to take advantage of both *P* and *S* data to exploit differences in the source radiation patterns. Even a limited amount of *S* information can help to resolve ambiguities present in the *P* data alone. In particular, the inclusion of just a few *SV* observations can provide significant constraints on the isotropic component. Both of the source representations work well with the neighbourhood algorithm. We have found that the use of a full moment tensor allows for a more general source mechanism than the use of a double-couple plus isotropic component representation, and therefore gives better results.

In order to test the robustness of the inversion scheme in the presence of noise, we have performed a number of synthetic tests with a perturbed source location. In all our tests the source time function and source depth have been well recovered, but limited sampling of the source radiation pattern with just a few stations means that it is possible to obtain source mechanisms which give a good representation of the seismograms but which differ significantly from the true model. Often a large isotropic component estimate is obtained, which results in a poor correspondence between the input and predicted source mechanisms. In this case it is desirable to also perform an inversion constrained to just a double-couple to see whether the isotropic estimate is real or just an artifact of the inversion. In particular, for all the synthetic earthquake tests we have performed, the misfit is lower when the isotropic component is restricted indicating that the data are compatible with a double-couple model.

Synthetic tests have also been performed to investigate how sensitive the inversion is to variations in velocity structure. We have used extreme test cases, with a systematic increase/decrease of the velocity in every layer of the *ak135* model, and with low velocity sediments present beneath every station. These examples have shown that both the source depth and mechanism are relatively insensitive to the velocity model used, even in these extreme cases.