Chapter 2

Finite-frequency tomography using adjoint methods—Methodology and examples using membrane surface waves

Note

This chapter was published as a paper by C. Tape, Q. Liu, and J. Tromp in *Geophysical Journal International* in 2007. Supplemental derivations and tables are included in Appendix A.

Summary

We employ adjoint methods in a series of synthetic seismic tomography experiments to recover surface-wave phase-speed models of southern California. Our approach involves computing the Fréchet derivative for tomographic inversions via the interaction between a forward wavefield, propagating from the source to the receivers, and an "adjoint" wavefield, propagating from the receivers back to the source. The forward wavefield is computed using a 2D spectral-element method (SEM) and a phase-speed model for southern California. A "target" phase-speed model is used to generate the "data" at the receivers. We specify an objective or misfit function that defines a measure of misfit between data and synthetics.

For a given receiver, the remaining differences between data and synthetics are time reversed and used as the source of the adjoint wavefield. For each earthquake, the interaction between the regular and adjoint wavefields is used to construct finite-frequency sensitivity kernels, which we call *event kernels*. An event kernel may be thought of as a weighted sum of phase-specific (e.g., P) *banana-doughnut kernels*, with weights determined by the measurements. The overall sensitivity is simply the sum of event kernels, which defines the *misfit kernel*. The misfit kernel is multiplied by convenient orthonormal basis functions that are embedded in the SEM code, resulting in the gradient of the misfit function, i.e., the Fréchet derivative. A nonlinear conjugate gradient algorithm is used to iteratively improve the model while reducing the misfit function. We illustrate the construction of the gradient and the minimization algorithm, and consider various tomographic experiments, including source inversions, structural inversions, and joint source-structure inversions. Finally, we draw connections between classical Hessian-based tomography and gradient-based adjoint tomography.

2.1   Introduction

Seismic tomography is in a state of transition from ray-based inversions using 1D reference models toward finite-frequency-kernel-based inversions using 3D reference models (*Akçelik et al.*, 2003; *Zhao et al.*, 2005). The transition from ray- to kernel-based inversions has been motivated in part by the pioneering studies of *Marquering et al.* (1999), *Zhao et al.* (2000), and *Dahlen et al.* (2000), which were based on 1D reference models but showed that seismological measurements are sensitive to structure away from the ray path and are affected by wavefront healing. The transition from 1D to 3D reference models has been motivated by computational advances coupled with success in modeling the forward problem of seismic wave propagation in complex media (e.g., *Komatitsch and Vilotte*, 1998; *Komatitsch et al.*, 2002; *Capdeville et al.*, 2003).

The purpose of this paper is to illustrate an approach for "3D–3D" seismic tomography, by which we mean seismic tomography based on a 3D reference model, 3D numerical simulations of the complete seismic wavefield, and finite-frequency sensitivity kernels. The success of 3D–3D tomography depends largely on two factors: (1) the accuracy and efficiency of the technique used to generate 3D synthetic seismograms, and (2) the efficiency of the inver-

sion algorithm. We have implemented numerical methods — the spectral-element method (SEM) — on parallel computers to simulate 3D seismic wave propagation at regional and global scales (e.g., *Komatitsch and Tromp*, 1999; *Komatitsch et al.*, 2004; *Komatitsch and Tromp*, 2002a,b). The inverse problem can be cast as a minimization problem, where the objective or misfit function measures some difference between data and synthetic seismograms computed from a 3D model. Our approach to the inverse problem utilizes adjoint methods (*Tarantola*, 1984; *Talagrand and Courtier*, 1987), which provide the gradient of the misfit function but not its second derivatives, i.e., the Hessian. The efficiency of the inverse algorithm is controlled by the computation of the gradient, which requires only two 3D simulations per earthquake (i.e., the gradient is independent of the number of receivers or the number of measurements), as well as an effectively chosen gradient method.

The framework for 3D–3D tomographic inversions using adjoint methods was developed in exploration geophysics (e.g., *Tarantola*, 1984; *Gauthier et al.*, 1986; *Mora*, 1987; *Pratt et al.*, 1998; *Pratt*, 1999). These studies illustrated the computation of the gradient and the related inversion technique using 2D heterogeneous models and 2D numerical algorithms. Applications of 3D–3D tomographic techniques are presented in *Bijwaard and Spakman* (2000), *Zhao et al.* (2005), *Capdeville et al.* (2005), and *Akçelik et al.* (2003), among others. *Bijwaard and Spakman* (2000) performed 3D ray-tracing though 3D models to iteratively improve a global P-wave model. *Zhao et al.* (2005) used fully numerical methods (finite differencing) to compute traveltime misfit function gradients for 3D models of the greater Los Angeles area. *Capdeville et al.* (2005), using synthetic data, demonstrated a technique of stacking synthetic records that limits the number of forward simulations to one per event (per model iteration); however, the technique requires modification when the data set is incomplete, as is generally the case. *Akçelik et al.* (2003), using synthetic data, illustrated a tomographic inversion using a finite-element method together with an adjoint approach within a conjugate gradient framework. They also addressed multiscale approaches to the inverse problem in an attempt to avoid reaching local minima during the inversion.

This paper is an extension of *Tromp et al.* (2005), which synthesized the work on adjoint methods with studies in finite-frequency tomography (*Marquering et al.*, 1999; *Dahlen et al.*, 2000; *Zhao et al.*, 2000) and time-reversal imaging (*Fink et al.*, 1989; *Fink*, 1992, 1997). In *Tromp et al.* (2005) we illustrated how the computation of a sensitivity kernel for a particular model and a particular type of measurement could be achieved via the interaction of two

wavefields, one constituting the "regular" wavefield traveling from source to receiver, and the other constituting the "adjoint" wavefield traveling from receiver to source, constructed by a suitable time-reversed synthetic seismogram recorded at the receiver. We performed a simple source inversion to illustrate the conjugate gradient algorithm, whereby only the gradient of the misfit function is used to iteratively invert for the source parameters. In this paper, we use the conjugate gradient approach to illustrate wave-speed inversions, source inversions, and joint (source and structure) inversions. In each example, the "observed" seismograms are computed for a "target" model, and the synthetic seismograms are computed from a current model that iteratively improves toward the target model over the course of the inversion. All of the simulations illustrated in this paper were performed on a single Linux PC.

We begin by highlighting the differences between classical and adjoint tomography in the context of a minimization problem. We define *classical tomography* as a Newton inversion scheme that computes model sensitivities for each measurement by constructing the gradient and Hessian of the misfit function (Section 2.3) (e.g., *Woodhouse and Dziewonski*, 1984; *Ritsema et al.*, 1999). In *adjoint tomography* only the gradient is computed, and it is computed via adjoint methods (e.g., *Gauthier et al.*, 1986; *Akçelik et al.*, 2003). In Section 2.5 we illustrate the construction of a *misfit kernel*, which can be thought of as the gradient of the misfit function. In Section 2.6 we show how this gradient is used in the conjugate gradient algorithm to iteratively improve the model. We finish by showing several tomographic experiments, including simultaneous source-structure inversions, as well as a comparison between ray- and kernel-based classical inversions and adjoint tomography.

2.2   General formulation of the inverse problem

Our objective will be to minimize a measure of the misfit between a set of data, for example waveforms or traveltimes, and a complementary set of synthetics. The generated synthetics are based on a model $\mathbf{m}$, for example a set of structural and source parameters, and our aim is to reduce the misfit between the data and the synthetics by making (successive) model corrections $\delta\mathbf{m}$. We define the *misfit function* $F(\mathbf{m})$ to be a measure of misfit between the data and synthetics computed for model $\mathbf{m}$. The function F is alternatively called an "objective" or "cost" function. For example, F could represent least-squares measures of

waveform or traveltime differences.

Let us suppose we have a particular model $\mathbf{m}$, and we wish to obtain an updated model $\mathbf{m} + \delta\mathbf{m}$ that brings us closer to a minimum of the misfit function F (*Nolet*, 1987; *Tarantola*, 2005, Appendix 6.22). We make a quadratic Taylor expansion of $F(\mathbf{m} + \delta\mathbf{m})$:

$$F(\mathbf{m} + \delta\mathbf{m}) \approx F(\mathbf{m}) + \mathbf{g}(\mathbf{m})^T \delta\mathbf{m} + \frac{1}{2}\delta\mathbf{m}^T \mathbf{H}(\mathbf{m})\,\delta\mathbf{m}, \tag{2.1}$$

where the gradient vector $\mathbf{g}(\mathbf{m})$ is defined in terms of the first derivative of the misfit function (also known as the Fréchet derivative) by

$$\mathbf{g}(\mathbf{m}) = \left.\frac{\partial F}{\partial \mathbf{m}}\right|_{\mathbf{m}} , \tag{2.2}$$

and the Hessian matrix $\mathbf{H}(\mathbf{m})$ is defined in terms of the second derivatives of the misfit function by

$$\mathbf{H}(\mathbf{m}) = \left.\frac{\partial^2 F}{\partial \mathbf{m}\partial \mathbf{m}}\right|_{\mathbf{m}} . \tag{2.3}$$

The "$|_{\mathbf{m}}$" dependence is used to emphasize that the preceding variable is evaluated at model $\mathbf{m}$.

The gradient of (2.1) with respect to $\delta\mathbf{m}$ is given by

$$\mathbf{g}(\mathbf{m} + \delta\mathbf{m}) \approx \mathbf{g}(\mathbf{m}) + \mathbf{H}(\mathbf{m})\,\delta\mathbf{m}, \tag{2.4}$$

which can be set equal to zero to obtain the (local) minimum of (2.1):

$$\mathbf{H}(\mathbf{m})\,\delta\mathbf{m} = -\mathbf{g}(\mathbf{m}). \tag{2.5}$$

An updated model $\mathbf{m} + \delta\mathbf{m}$ may be obtained with or without the Hessian $\mathbf{H}$. If the gradient and Hessian (or approximate Hessian) are both available, then the inverse approach is known as a *Newton method*; if only the gradient is available, then it is a *gradient method* (e.g., steepest descent, conjugate gradient). In classical traveltime tomography, one generally has access to both the gradient $\mathbf{g}$ and the Hessian $\mathbf{H}$ of the misfit function, in which case the model update $\delta\mathbf{m}$ may be obtained based on (2.5). For complex, heterogeneous models, computation of the gradient is generally still feasible, but computation of the Hes-

sian is not. In the absence of the Hessian, one can minimize the misfit function using only the gradient (2.2) based on iterative methods.

2.3 Classical tomography

We begin by investigating 2D surface-wave traveltime tomography based on either ray or finite-frequency sensitivity kernels. These classical inversions, which involve access to both the gradient and the Hessian of the misfit function, serve as a reference and standard for subsequent iterative inversions based on only the gradient (Section 2.6). In particular, we will investigate how many iterations of the conjugate-gradient adjoint approach are required to obtain a similar misfit to the data as an inversion based on knowledge of the gradient and Hessian. Of course our ultimate goal is to use the adjoint approach to address inverse problems for fully 3D reference models, when the calculation of the Hessian is generally not feasible, and the experiments in this paper serve as a guide to the implementation and convergence of such iterative inversions.

2.3.1 Theory

The traveltime misfit function may be expressed as

$$F(\mathbf{m}) = \frac{1}{2} \sum_{i=1}^{N} \left[T_i^{\text{obs}} - T_i(\mathbf{m}) \right]^2, \tag{2.6}$$

where T_i^{obs} denotes the observed traveltime for the ith source-receiver combination, $T_i(\mathbf{m})$ the predicted traveltime based on the current model $\mathbf{m}$, and N the number of traveltime measurements. The variation of the misfit function (2.6) is given by

$$\delta F = -\sum_{i=1}^{N} \Delta T_i \, \delta T_i, \tag{2.7}$$

where δT_i is the theoretical traveltime perturbation and

$$\Delta T_i = T_i^{\text{obs}} - T_i(\mathbf{m}) \tag{2.8}$$

denotes the traveltime anomaly. The sign convention for the traveltime anomaly follows that of *Dahlen et al.* (2000) and *Dahlen and Baig* (2002), such that a negative traveltime

indicates a delay in the synthetic arrival relative to the recorded arrival. Throughout this paper, an uppercase delta, Δ, will denote a differential measurement, and a lowercase delta, δ, will denote a mathematical perturbation.

In ray-based tomography, the predicted traveltime anomaly δT_i along the ith ray path may be related to fractional wave-speed perturbations $\delta \ln c = \delta c / c$ based on the relationship

$$\delta T_i = - \int_{\text{ray}_i} c^{-1} \, \delta \ln c \, \mathrm{d}s, \tag{2.9}$$

where $\mathrm{d}s$ denotes a segment of the ith ray.

Taking into account finite-frequency effects, *Marquering et al.* (1999), *Zhao et al.* (2000), and *Dahlen et al.* (2000) demonstrate that the traveltime anomaly may alternatively be related to relative wave-speed perturbations based on a finite-frequency sensitivity kernel $K_i(\mathbf{x})$ for the ith source-receiver combination by

$$\delta T_i = \int_V K_i \, \delta \ln c \, \mathrm{d}^3 \mathbf{x}. \tag{2.10}$$

Marquering et al. (1999) dubbed these finite-frequency kernels "banana-doughnut kernels" on account of their shape in smooth, spherically symmetric Earth models for cross-correlation traveltime measurements. These kernels are also referred to as "sensitivity," "finite-frequency," or "Born" kernels. For our purposes, the key point is that a banana-doughnut kernel does not incorporate the traveltime measurement, whereas the event and misfit kernels discussed in Section 2.5 do incorporate measurements.

Unlike the ray-theoretical expression (2.9), equation (2.10) relates the traveltime anomaly to 3D heterogeneity $\delta \ln c$ throughout the entire Earth model, as seen through the kernel K_i. The relations (2.9) and (2.10) are valid for any model. Frequently the model is chosen to be one-dimensional because this makes the ray and finite-frequency kernel calculations much simpler, but this is not required (*Zhao et al.*, 2005).

Substituting (2.10) into (2.7), we express the variation of the traveltime misfit function for finite-frequency tomography as

$$\delta F = \int_V K \, \delta \ln c \, \mathrm{d}^3 \mathbf{x}, \tag{2.11}$$

where the traveltime *misfit kernel* $K(\mathbf{x})$ is a weighted sum of the kernels $K_i(\mathbf{x})$:

$$K(\mathbf{x}) = -\sum_{i=1}^{N} \Delta T_i \, K_i(\mathbf{x}), \qquad (2.12)$$

such that the weight associated with the kernel for the ith source-receiver combination K_i is the corresponding traveltime anomaly ΔT_i. It is important to note the distinction that misfit kernels $K(\mathbf{x})$ depend on the data, whereas the banana-doughnut kernels $K_i(\mathbf{x})$ are data independent.

To make the tomographic inversions practical, we need to choose a finite set of basis functions in which to expand our model. Let $B_k(\mathbf{x})$, $k = 1, \ldots, M$, denote a set of M basis functions. We expand our fractional phase-speed perturbations, $\delta \ln c(\mathbf{x})$, into these basis functions:

$$\delta \ln c(\mathbf{x}) = \sum_{k=1}^{M} \delta m_k \, B_k(\mathbf{x}), \qquad (2.13)$$

where δm_k, $k = 1, \ldots, M$, represent the perturbed model coefficients, which are determined in terms of the gradient $\mathbf{g}$ and Hessian $\mathbf{H}$ of the misfit function by (2.5).

Next, we determine $\mathbf{g}$ and $\mathbf{H}$ for this classical traveltime tomography problem. Substituting (2.13) into (2.9) and (2.10), respectively, we obtain

$$\delta T_i = \sum_{k=1}^{M} \delta m_k \, G_{ik}, \qquad (2.14)$$

where for ray theory,

$$G_{ik} \equiv \left. \frac{\partial T_i}{\partial m_k} \right|_{\mathbf{m}} = -\int_{\mathrm{ray}_i} c^{-1} B_k \, \mathrm{d}s, \qquad (2.15)$$

whereas for finite-frequency tomography,

$$G_{ik} \equiv \left. \frac{\partial T_i}{\partial m_k} \right|_{\mathbf{m}} = \int_V K_i B_k \, \mathrm{d}^3 \mathbf{x}. \qquad (2.16)$$

We note that in either case G_{ik} will depend on the source-receiver geometry (index i), the choice of basis functions (index k), and the choice of reference model ($\mathbf{m}$).

Using (2.13) we express the variation in the misfit function (2.11) as

$$\delta F = \sum_{k=1}^{M} \int_V KB_k \ \mathrm{d}^3\mathbf{x} \ \delta m_k. \tag{2.17}$$

Upon comparing this result to

$$\delta F = \mathbf{g} \cdot \delta\mathbf{m} = \sum_{k=1}^{M} g_k \ \delta m_k, \tag{2.18}$$

we deduce that the elements of the gradient vector, g_k, are determined by

$$g_k = \frac{\partial F}{\partial m_k} = \int_V KB_k \ \mathrm{d}^3\mathbf{x}, \qquad k = 1, \ldots, M. \tag{2.19}$$

This highlights the simple relationship between the misfit kernel and the gradient of the misfit function. Substituting (2.12) into (2.19), we obtain

$$g_k = -\sum_{i=1}^{N} \int_V K_i B_k \ \mathrm{d}^3\mathbf{x} \ \Delta T_i = -\sum_{i=1}^{N} G_{ik} \ \Delta T_i, \qquad k = 1, \ldots, M, \tag{2.20}$$

which in matrix notation becomes

$$\mathbf{g} = -\mathbf{G}^T\mathbf{d}. \tag{2.21}$$

Here $\mathbf{G}$ is the $N \times M$ design matrix constructed using (2.15) for rays or (2.16) for finite-frequency kernels, a superscript T denotes the transpose, and $\mathbf{d}$ is defined as an N-dimensional data vector of cross-correlation traveltime measurements:

$$\mathbf{d} = (\Delta T_1, \ldots, \Delta T_i, \ldots, \Delta T_N)^T. \tag{2.22}$$

Note that the data vector depends on model $\mathbf{m}$ through the synthetics.

The second derivatives of the misfit function are given by (2.3), and thus the elements of the Hessian $\mathbf{H}$ are given by

$$H_{kk'} = \left.\frac{\partial^2 F}{\partial m_k \partial m_{k'}}\right|_{\mathbf{m}} = \left.\frac{\partial g_k}{\partial m_{k'}}\right|_{\mathbf{m}} = \sum_{i=1}^{N} \left[G_{ik'}G_{ik} + \Delta T_i \left.\frac{\partial^2 T_i}{\partial m_k \partial m_{k'}}\right|_{\mathbf{m}}\right], \tag{2.23}$$

where G_{ik} is defined in (2.16). We introduce an *approximate Hessian* $\tilde{\mathbf{H}}$ by ignoring the second-order terms:

$$\tilde{H}_{kk'} \equiv \sum_{i=1}^{N} G_{ik} G_{ik'}, \qquad k, k' = 1, \ldots, M, \tag{2.24}$$

which in matrix notation is

$$\tilde{\mathbf{H}} \equiv \mathbf{G}^T \mathbf{G}. \tag{2.25}$$

Henceforth, we will refer to $\tilde{\mathbf{H}}$ as the Hessian. This approximation, $\tilde{\mathbf{H}} \approx \mathbf{H}$, characterizes the *Gauss–Newton method* and is exact if the model perturbations are linearly related to the traveltime measurements.

Having established the gradient (2.21) and Hessian (2.25), the model correction $\delta\mathbf{m}$ is determined by (2.5):

$$\mathbf{G}^T \mathbf{G} \, \delta\mathbf{m} = \mathbf{G}^T \mathbf{d}, \tag{2.26}$$

where $\delta\mathbf{m}$ is defined in (2.13), $\mathbf{d}$ is defined in (2.22), and $\mathbf{G}$ is defined according to (2.15) or (2.16).

In general, the Hessian matrix (2.25) is not full rank, which means that its inverse does not exist. To stabilize the inverse problem, one introduces a damping matrix $\mathbf{D}$ typically involving the norm, gradient, or second derivative of the wave-speed perturbations, and a damping parameter γ:

$$\tilde{\mathbf{H}}_\gamma = \mathbf{G}^T \mathbf{G} + \gamma^2 \mathbf{D}. \tag{2.27}$$

The damping parameter γ is chosen in a subjective manner, generally by inspecting a graph that trades off misfit of the solution against complexity of the model. Having stabilized the inverse of the Hessian, the solution to (2.26) may now be expressed as

$$\delta\mathbf{m} = (\mathbf{G}^T \mathbf{G} + \gamma^2 \mathbf{D})^{-1} \mathbf{G}^T \mathbf{d}, \tag{2.28}$$

from which the updated model, $\mathbf{m} + \delta\mathbf{m}$, may be obtained. In Section 2.10 we show how (2.28) is obtained by adding a regularization term to the misfit function. More generally,

for nonlinear inverse problems one uses an iterative Gauss–Newton method to minimize the misfit function. In that case (2.28) is replaced by an iterative expression that relates model $k + 1$ to model k and the initial model (e.g., *Tarantola*, 2005).

2.3.2   Experimental setup

We simulate 2D elastic wave propagation using a spectral-element method (SEM), which combines the flexible spatial parameterization of finite-element methods with the accuracy of pseudospectral methods (e.g., *Komatitsch and Vilotte*, 1998; *Komatitsch and Tromp*, 1999). For simplicity, we consider "membrane waves" (*Tanimoto*, 1990; *Peter et al.*, 2007) traveling in the x-y plane with a vertical (z) component of motion. The elastic wave equation for the vertical component of displacement $s(x, y, t)$ is given by

$$\rho \, \partial_t^2 s = \partial_x(\mu \, \partial_x s) + \partial_y(\mu \, \partial_y s) + f, \tag{2.29}$$

where $\rho(x, y)$ denotes the density distribution and $\mu(x, y)$ the shear modulus. The source $f(x, y, t)$ is given by

$$f(x, y, t) = h(t) \, \delta(x - x_s) \, \delta(y - y_s), \tag{2.30}$$

where $h(t)$ denotes the source time function and (x_s, y_s) the source location. All four membrane edges are absorbing, and attenuation and anisotropy are not incorporated. The relationship between membrane-wave phase speed, c, and rigidity is $\mu = \rho c^2$.

We take southern California as our region of interest (Figure 2.1) in anticipation of eventually improving the present 3D reference wave-speed models (*Hauksson*, 2000; *Magistrale et al.*, 2000; *Süss and Shaw*, 2003). The modeled region is 480 km by 480 km. The numerical simulations are carried out on a planar grid with $N_{\mathrm{glob}} = 25921$ gridpoints. The source time function of the point source (2.30) used in the simulations is a Gaussian of the form

$$h(t) = (-2\alpha^3/\sqrt{\pi}) \, (t - t_s) \, \exp[-\alpha^2(t - t_s)^2], \tag{2.31}$$

where $\alpha = 2\tau_0/\tau$, $\tau_0 = 2.628$ s, $\tau = 20.0$ s is the duration of $h(t)$, and $t_s = 48.0$ s is the origin time (e.g., Figure 2.6a). The duration of each simulation is $T = 240$ s unless

otherwise noted.

The synthetic records are computed using source locations of actual events ($M \geq 4$) recorded in southern California between 1990 and 2005 (Figure 2.1). The initial set of synthetics is computed using a model with homogeneous phase speed c. In general, the synthetics in our experiments are generated from a laterally varying model, while the data are generated from what is designated as the "target" model. Computationally, the model correction is expressed as a fractional perturbation, $\delta c/c = \delta \ln c$, with current phase speed c. In the figures, however, each phase-speed model is plotted as a percent perturbation from the phase-speed value for the initial model. In Section 2.8 we allow for additional perturbations in the source parameters, so that in general the synthetics are computed from a model with perturbed sources and perturbed structure.

2.3.3   2D tomographic example

To illustrate a classical tomographic inversion, we begin by choosing a set of basis functions, $B_k(\mathbf{x})$, in which to expand the fractional wave-speed perturbations $\delta \ln c(\mathbf{x})$ (2.13). We use spherical spline basis functions (*Wang and Dahlen*, 1995; *Wang et al.*, 1998), which are well-suited for regional models where multiscale parameterization is desired because of nonuniform path coverage (e.g., *Boschi et al.*, 2004). (We do not exploit the multiscale aspects here.) An example of a spherical spline basis function is plotted in Figure 2.2b. We choose $M = 286$ spherical spline basis functions to cover the southern California region.

The data are computed using the phase-speed model in Figure 2.1b, and the synthetics are computed for a homogeneous phase-speed model with $c = 3.78$ km/s. We make cross-correlation traveltime measurements between data and synthetics to obtain the data vector $\mathbf{d}$ (2.22). The total number of measurements is $N = N_{\mathrm{events}} \times N_{\mathrm{receivers}} = 25 \times 132 = 3300$.

We illustrate the classical tomographic approach using both rays and banana-doughnut kernels to represent the sensitivities of the measurements to the model parameters. Thus we compute two $N \times M$ design matrices, $\mathbf{G}^{\mathrm{ray}}$ and $\mathbf{G}^{\mathrm{ker}}$, respectively. Figure 2.2a–c shows the computation of a single G_{ik}^{ray} element, and Figure 2.2d–f shows an example for G_{ik}^{ker}. Figure 2.2 illustrates why the choice between kernels or rays may be moot, depending on the resolution of the basis functions. The infinitesimally thin ray path is smeared out by the relatively smooth basis functions. Thus, in our example, $\mathbf{G}^{\mathrm{ray}} \approx \mathbf{G}^{\mathrm{ker}}$, and we will simply use a generic $\mathbf{G}$ to denote either the ray or kernel design matrix.

The (approximate) Hessian matrix, $\tilde{\mathbf{H}} = \mathbf{G}^T\mathbf{G}$, and the gradient vector, $\mathbf{g} = -\mathbf{G}^T\mathbf{d}$, are visualized in Figure 2.3. The Hessian and gradient are determined by the source-receiver geometry and the banana-doughnut kernels (or ray paths), but only the gradient is controlled by the data.

Figure 2.4 shows the model recovery using classical tomography based on a single iteration of the Gauss–Newton method. The recovered model is strongly dependent on the damping parameter γ. When $\gamma \approx 0$, the inverse is unstable and structure is artificially introduced into regions where there is no coverage, i.e., the edges of the domain and the oceans (Figure 2.4a). When $\gamma \to \infty$, the recovered model is simply the initial model (Figure 2.4f), although the spatial pattern is that of the gradient (e.g., compare Figure 2.3c with Figure 2.4g). The reason for this is that for large values of the damping parameter γ the damped Hessian (2.27) is dominated by the damping matrix $\mathbf{D}$, which in our case is the identity matrix $\mathbf{I}$. In this case the solution to the inverse problem given by (2.28) is effectively a scaled version of the gradient $\mathbf{g}$. For the example in Figure 2.4, the L-curve suggests that $\gamma = 10.0$ is a reasonable model selection; this model is shown in Figure 2.4c and Figure 2.20c.

2.4   Computation of the gradient and Hessian

Obtaining the Hessian involves computing banana-doughnut kernels K_i for each source-receiver combination. Thus, the cost of computing the Hessian is the cost of computing all the kernels. For a problem involving N_{events} earthquakes, $N_{\text{receivers}}$ stations, $N_{\text{comp}} = 3$ component seismograms, and N_{picks} measurements per seismogram one would need to calculate $N_{\text{events}} \times N_{\text{receivers}} \times N_{\text{comp}} \times N_{\text{picks}}$ kernels.

In adjoint tomography one computes a misfit kernel K from which only the gradient is obtained. One of the primary benefits of adjoint tomography is that the misfit kernel need not be computed by summing over individual banana-doughnut kernels for each source-receiver pair, as in (2.12). Instead, the measurements, ΔT_i, are incorporated into the adjoint source, which is used to compute the misfit kernel (Section 2.5). This kernel is constructed via the interaction between a forward wavefield and an adjoint wavefield, requiring only two simulations per earthquake (*Tromp et al.*, 2005). So if our inverse problem involves N_{events} earthquakes, obtaining the gradient of the misfit function involves $2N_{\text{events}}$

numerical simulations, i.e., this calculation is *independent of the number of receivers, components, and picks*. The main drawback of adjoint tomography is that the Hessian is not available, which means that iterative techniques must be used to determine the minimum of the objective function.

Thus, a fundamental distinction between classical and adjoint tomography is whether or not individual banana-doughnut kernels are computed. In the context of classical tomography, there are several ways to compute the kernels. For 1D Earth models, they may be calculated cheaply and rapidly, in particular if approximate expressions are used (*Dahlen et al.*, 2000). Using normal modes, *Zhao and Jordan* (2006) computed global finite-frequency kernels for spherically symmetric models. The kernels may be used to construct the design matrix $\mathbf{G}$, which has $N_{\mathrm{events}} \times N_{\mathrm{receivers}} \times N_{\mathrm{comp}} \times N_{\mathrm{picks}} \times M$ elements. The parameterization of the model (2.13) must be carefully considered, since M scales $\mathbf{G}$. Once $\mathbf{G}$ is obtained, the Hessian follows from (2.25).

The computation of the kernels K_i for 3D models may be accomplished in two ways:

1. We may perform an adjoint simulation for every single measurement, which requires a total of $2N_{\mathrm{events}} \times N_{\mathrm{receivers}} \times N_{\mathrm{comp}} \times N_{\mathrm{picks}}$ simulations (two for each measurement). For 3D models the numerical cost is prohibitive.

2. Alternatively, we may invoke source-receiver reciprocity and for every source and receiver calculate and store Green's functions as a function of both space and time. This requires one to perform and store $N_{\mathrm{events}} + 3N_{\mathrm{receivers}}$ simulations: one simulation for each event and one simulation for each receiver component. For realistic 3D simulations the storage requirements are formidable, although for small problems the approach is feasible, as demonstrated by *Zhao et al.* (2005).

Our goal is to improve fully 3D reference models. Therefore, to make the inverse problem tractable, we are forced to consider an approach based on knowledge of the value of the misfit function $F(\mathbf{m})$, its gradient $\mathbf{g}$, but not its Hessian $\tilde{\mathbf{H}}$. Minimization of the misfit function based on this information may be accomplished using a nonlinear conjugate gradient method, as discussed in Section 2.6. But first we demonstrate how we compute the gradient using adjoint methods.

2.5   The gradient: construction of a misfit kernel

In this section we demonstrate how we compute the gradient of the misfit function, $\mathbf{g} = \partial F / \partial \mathbf{m}$, using adjoint methods. The gradient of the misfit function is obtained from (2.19):

$$g_k = \int_\Omega K B_k \, \mathrm{d}^2 \mathbf{x}, \tag{2.32}$$

where for the 2D examples in this paper the integration is over the model surface Ω. Given the misfit kernel, K, and the basis functions, B_k, we can readily compute the gradient of the misfit function. The misfit kernel can also be thought of as a sum of *event kernels*, which we discuss next.

2.5.1   Event kernels

Tromp et al. (2005, Fig. 3) illustrated the construction of a data-independent banana-doughnut kernel based on adjoint methods. In this paper, the kernels we show are *misfit* kernels, whereby the adjoint source is constructed based in part on a set of measurements between data and synthetics.

The construction of misfit kernels based on cross-correlation traveltime measurements is outlined in *Tromp et al.* (2005, Section 4). For membrane waves, motion is restricted to the vertical direction, and the source functions and wavefields are scalar quantities. The source for the adjoint wavefield for a particular event is given by (*Tromp et al.*, 2005, eq. 57)

$$f^\dagger(x, y, t) = -\sum_{r=1}^{N_r} \Delta T_r \frac{1}{M_r} \, w_r(T - t) \, \partial_t s(x_r, y_r, T - t) \, \delta(x - x_r) \, \delta(y - y_r), \tag{2.33}$$

where r is the receiver index, N_r is the number of receivers, ΔT_r is the cross-correlation traveltime measurement over a time window $w_r(t)$, $s(x, y, t)$ is the forward wavefield determined by (2.29), (x_r, y_r) is the location of the receiver, T is the length of the time series, and M_r is a normalization factor. The key point is that the adjoint force comprises time-reversed velocity seismograms, input *at the location of the receivers* and weighted by the traveltime measurement associated with each receiver.

For a given earthquake (event), the interaction between the adjoint wavefield and the

forward wavefield gives rise to the membrane *event kernel*

$$K(x, y) = -2\mu(x, y) \int_0^T \left[\partial_x s^\dagger(x, y, T - t) \partial_x s(x, y, t) + \partial_y s^\dagger(x, y, T - t) \partial_y s(x, y, t) \right] \mathrm{d}t.$$
(2.34)

Note that the misfit between the data and synthetics is incorporated into the adjoint source (2.33), which gives rise to the adjoint wavefield $s^\dagger$. Equation (2.34) is obtained from the expression for an SH β-kernel in *Tromp et al.* (2005), which contains a product of the adjoint and regular deviatoric strain tensors. In the case of the SH (or membrane) waves, there are four nonzero components (two unique) of each deviatoric strain tensor, which leads to (2.34).

Figures 2.5 and 2.6 show the construction of an event kernel for a single source–receiver pair for a cross-correlation traveltime measurement. The source-receiver geometry and forward wavefield are shown in the left column of Figure 2.5. The synthetics are computed for a homogeneous reference model ($c = 3.50$ km/s), and the data are computed for a uniformly perturbed "target" model with $\delta \ln c = 0.1$, i.e., $c(1 + \delta \ln c) = 3.85$ km/s. The cross-correlation traveltime measurement at the receiver is $\Delta T = -9.72$ s, indicating a late arrival of the synthetics with respect to the data. The adjoint source function is constructed by time-reversing the synthetic velocity recorded at the receiver and multiplying by ΔT (Figure 2.6; eq. 2.33).

We now replace the homogeneous target model with the checkerboard target model in Figure 2.7a. Figure 2.8 shows the construction of an event kernel for this target model for multiple receivers, thereby incorporating multiple measurements. Just as in Figure 2.5, the event kernel that forms in Figure 2.8 highlights the regions of the current model that give rise to the (cross-correlation traveltime) discrepancies between the data and synthetics. However, in Figure 2.8 this is more obvious since the model used to generate the data is not simply a homogeneous perturbation but rather a large-scale checker pattern. The event kernel in Figure 2.8 looks qualitatively similar to the phase-speed model in Figure 2.7, except with the opposite sign, which is consistent with (2.11): for the variation of the misfit function to be negative, we invoke a fast, positive (blue) structural perturbation where the kernel is negative (red), and a slow, negative structural perturbation where the kernel is positive.

As shown in (2.33), the amplitude of the adjoint source at a particular receiver, r, is determined in part by the traveltime measurement ΔT_r. Changing the values of ΔT_r changes the weights of the corresponding individual banana-doughnut kernels that comprise the event kernel, something that is explicit in the classical sense (2.12). It is possible to incorporate some measure of weighting at the stage of constructing the adjoint source in order to account for uneven coverage (Figure 2.1), as demonstrated in *Takeuchi and Kobayashi* (2004). Another option is to weight the adjoint sources according to realistic uncertainties associated with each measurement (*Tarantola*, 1984): a measurement with a high uncertainty will have a small amplitude weight, and thus a relatively weak contribution to the event kernel.

2.5.2   Misfit kernels and damping

We define the *misfit kernel* as the sum of the event kernels for a particular model. Thus, the gradient of the misfit function, **g**, is obtained as in (2.32) using the misfit kernel $K(\mathbf{x})$. Figure 2.9 shows the construction of a misfit kernel for 25 events. Note that features of each event kernel are very different, even for the simple checkerboard model in this example (Figure 2.7). Only after summing the event kernels does the pattern (Figure 2.9h) begin to resemble the model used to generate the data (Figure 2.9i).

We apply a smoothing operator to the misfit kernels in order to remove spurious amplitudes in the immediate vicinity of the sources and receivers (Figure 2.10). This is accomplished by convolving (in 2D) the unsmoothed misfit kernel with a Gaussian of the form

$$G(x,y) = \frac{4}{\pi \, \Gamma^2} \, e^{-4\,(x^2+y^2)/\Gamma^2}, \qquad (2.35)$$

where Γ is the full-width of the Gaussian, defined such that at a (polar) distance $r = \Gamma/2$, the Gaussian has amplitude $G(r) = G(0)e^{-1}$; thus Γ is the scalelength of smoothing (Figure 2.10). The choice of Γ is somewhat analogous to the choice of damping parameter γ for the inversion of the Hessian (eq. 2.27), which involves a degree of subjectivity. In the adjoint method, subjectivity may be removed by selecting Γ according to the shortest wavelengths of the waves. It seems sensible to smooth the kernels using scalelengths somewhat less than the wavelengths of the seismic waves resolved in the numerical simulation.

There will exist short-scalelength features and fringes in kernels based on more compli-
cated 2D or 3D models, such as the fringes shown in *Tromp et al.* (2005, Figure 9) for the
P-SV wavefield or in *Zhou et al.* (2004, Figure 13b). The smoothing operation will tend to
remove these sub-resolution features from the kernel. An alternative approach to smoothing
the inversion is to add an explicit damping term to the misfit function (e.g., *Akçelik et al.*,
2002, 2003), as outlined in Section 2.10. This approach leads to an additional term in the
expression for the gradient, which represents the desire to obtain a smooth model. We
prefer to convolve the misfit kernel with a simple Gaussian that represents the resolution
of the simulation, and this is the approach we will take in this paper.

2.5.3   Basis functions

As shown in (2.32), the calculation of the gradient of the misfit function requires a choice of
model parameterization. Which basis functions should one use? In the classical tomographic
example discussed in Section 2.3.3 we used $M = 286$ spherical spline basis functions to
parameterize the model (see Figure 2.2). In adjoint tomography, where the wavefields
and kernels are represented on discretized grids, we can use the basis functions embedded
in the numerical method itself, for example for the SEM we use Lagrange polynomials
(*Komatitsch and Tromp*, 1999). This has the advantage that no restrictions are placed on
the wavelengths of the model, other than that they need to be resolvable by the waves used
in the inversion. This approach increases the number of model parameters dramatically
compared to a classical inversion, but because we do not need to invert a Hessian in the
adjoint approach this is of no consequence.

Any smooth function $f(\mathbf{x})$, where $\mathbf{x} = (x, y)$, that is sufficiently resolved by the SEM
mesh can be expressed in discrete form as

$$f(\mathbf{x}) = \sum_{k=1}^{N_{\text{glob}}} f_k L_k(\mathbf{x}), \tag{2.36}$$

where $k = 1, \ldots, N_{\text{glob}}$ is the index of the N_{glob} global node points, $f_k = f(\mathbf{x}_k)$ is the
functional value at global node $\mathbf{x}_k$, and $L_k(\mathbf{x})$ is a global function defined by

$$L_k(\mathbf{x}) = \begin{cases} l_\alpha(\xi(x, y)) \, l_\beta(\eta(x, y)) & \text{if } \mathbf{x}_k \in \Omega_e \text{ and } k|_{\Omega_e} = (\alpha, \beta), \\ 0 & \text{if } \mathbf{x}_k \notin \Omega_e. \end{cases} \tag{2.37}$$

Here l_α and l_β are Lagrange polynomials of degree α and β, respectively. We use degree 4 polynomials, i.e., 5 Gauss-Lobatto-Legendre points, in the 2D simulations presented in this paper. The invertible mapping from the reference square with points (ξ, η), with $-1 \le \xi \le 1$ and $-1 \le \eta \le 1$, to the deformed quadrilateral spectral-element Ω_e with points (x, y) may be written in the form $\xi = \xi(x, y)$, $\eta = \eta(x, y)$ (e.g., *Komatitsch and Vilotte*, 1998; *Komatitsch and Tromp*, 1999). Note that functions $L_k(\mathbf{x})$ corresponding to global grid points $\mathbf{x}_k$ located on the edges or corners of elements have nonzero contributions from all elements that share the global point. At the kth node,

$$L_k(\mathbf{x}_k) = 1, \tag{2.38}$$

in accordance with (2.36).

The functions $L_k(\mathbf{x})$ are orthogonal but not orthonormal. We may obtain a set of orthonormal basis functions $B_k(\mathbf{x})$ based on the definition

$$B_k(\mathbf{x}) = L_k(\mathbf{x})/A_k, \tag{2.39}$$

where A_k is the square-root-area associated with the kth node:

$$A_k^2 = \int_\Omega L_k^2(\mathbf{x}) \, \mathrm{d}^2\mathbf{x}. \tag{2.40}$$

The B_k are orthonormal in the sense that

$$\int_\Omega B_k(\mathbf{x}) B_{k'}(\mathbf{x}) \mathrm{d}^2\mathbf{x} = \delta_{kk'}, \tag{2.41}$$

and any function can be expanded in terms of these basis functions. For example, we may expand the misfit kernel $K(\mathbf{x})$ in terms of the basis functions $B_k(\mathbf{x})$ as

$$K(\mathbf{x}) = \sum_{k=1}^{N_{\mathrm{glob}}} \tilde{K}_k B_k(\mathbf{x}). \tag{2.42}$$

The expansion coefficients $\tilde{K}_k$ are determined by

$$
\begin{aligned}
\tilde{K}_k &= \int_\Omega K(\mathbf{x}) B_k(\mathbf{x}) \,\mathrm{d}^2\mathbf{x} = \int_\Omega \sum_{k'} K_{k'} L_{k'}(\mathbf{x}) B_k(\mathbf{x}) \,\mathrm{d}^2\mathbf{x} = \sum_{k'} K_{k'} A_{k'} \int_\Omega B_{k'}(\mathbf{x}) B_k(\mathbf{x}) \,\mathrm{d}^2\mathbf{x} \\
&= K_k A_k
\end{aligned} \tag{2.43}
$$

where $K_k = K(\mathbf{x}_k)$ is the value of the misfit kernel at a global grid point, and we have used (2.36) and the orthonormality relation (2.41).

Now let us assume we have computed a misfit kernel $K(\mathbf{x})$. In discrete form, we can write $K(\mathbf{x}_k) = K_k$, since K is defined on the $N_{\mathrm{glob}} = 25921$ global nodes of the SEM mesh. Upon comparing (2.32) with (2.43), we see that, using the basis functions (2.39), the gradient of the misfit function is simply

$$
g_k = K_k A_k. \tag{2.44}
$$

This provides a trivial step from the discretized kernel to the gradient. Using the $M = N_{\mathrm{glob}}$ basis functions in (2.39), the model parameters (2.13) are therefore

$$
\delta m_k = \delta \ln c_k A_k, \tag{2.45}
$$

where $\delta \ln c_k$ is the discrete version of $\delta \ln c(\mathbf{x})$.

2.6 Optimization: iterative improvement of the model

In the previous section we showed how to compute the gradient of the misfit function by summing event kernels (Figure 2.9) and subsequently multiplying by the basis functions of the model (2.32). In this section we illustrate how iterative improvements to the model may be determined based on a nonlinear conjugate gradient algorithm (*Fletcher and Reeves*, 1964). We demonstrated this algorithm for a simple source inversion in *Tromp et al.* (2005, Section 8.1). In Section 2.6.2 we consider a 2D tomographic example.

2.6.1 Conjugate gradient algorithm

The algorithm we use may be summarized as follows: given an initial model $\mathbf{m}^0$, calculate $F(\mathbf{m}^0)$, $\mathbf{g}^0 = \partial F/\partial \mathbf{m}(\mathbf{m}^0)$, and set the initial conjugate gradient search direction equal to minus the initial gradient of the misfit function, $\mathbf{p}^0 = -\mathbf{g}^0$. If $||\mathbf{p}^0|| < \epsilon$, where ϵ is a suitably small number, then $\mathbf{m}^0$ is the model we seek to determine, otherwise:

1. Perform a line search to obtain the scalar ν^k that minimizes the function $\tilde{F}^k(\nu)$ where

$$
\begin{aligned}
\tilde{F}^k(\nu) &= F(\mathbf{m}^k + \nu \mathbf{p}^k) \\
\tilde{g}^k(\nu) &= \frac{\partial \tilde{F}^k}{\partial \nu} = \frac{\partial F}{\partial \mathbf{m}}\left(\mathbf{m}^k + \nu \mathbf{p}^k\right) \cdot \mathbf{p}^k
\end{aligned}
$$

   - Choose a test parameter $\nu_t^k = -2\tilde{F}^k(0)/\tilde{g}^k(0)$, based on quadratic extrapolation.
   - Calculate the test model $\mathbf{m}_t^k = \mathbf{m}^k + \nu_t^k \mathbf{p}^k$.
   - Calculate $F(\mathbf{m}_t^k)$ and, for cubic interpolation, $\mathbf{g}_t^k = \mathbf{g}(\mathbf{m}_t^k)$.
   - Interpolate the function $\tilde{F}^k(\nu)$ by a quadratic or cubic polynomial and obtain the ν^k that gives the (analytical) minimum value of this polynomial.

2. Update the model: $\mathbf{m}^{k+1} = \mathbf{m}^k + \nu^k \mathbf{p}^k$, then calculate $\mathbf{g}^{k+1} = \partial F/\partial \mathbf{m}(\mathbf{m}^{k+1})$.

3. Update the conjugate gradient search direction: $\mathbf{p}^{k+1} = -\mathbf{g}^{k+1} + \beta^{k+1}\mathbf{p}^k$, where $\beta^{k+1} = \mathbf{g}^{k+1} \cdot (\mathbf{g}^{k+1} - \mathbf{g}^k)/(\mathbf{g}^k \cdot \mathbf{g}^k)$.

4. If $||\mathbf{p}^{k+1}|| < \epsilon$, then $\mathbf{m}^{k+1}$ is the desired model; otherwise replace k with $k+1$ and restart from 1.

A key decision is the choice of the test parameter, ν_t^k, which determines how far one should go in the direction of the search direction (initially the negative gradient) to obtain the test model. We assume a quadratic form of the misfit function and determine ν_t^k based on this assumption. Computation of $F(\mathbf{m})$ (misfit) and $\mathbf{g}(\mathbf{m})$ (misfit kernel) are expensive in the tomographic problem, and thus we must limit the number of computations as much as possible. Some of these aspects are addressed in Section 2.6.3.

2.6.2   2D tomographic example

Using (2.6), we can define the average traveltime anomaly for a particular model:

$$\overline{\Delta T} = \sqrt{2\, F(\mathbf{m})/N}. \tag{2.46}$$

This gives some physical meaning to the F-values in the plots in this section. Figure 2.11a–i shows one cycle of the conjugate gradient algorithm for the 2D tomographic example. Part (a) shows the phase-speed model used to generate the data (the "target" model), and (b) shows the initial phase-speed model, $\mathbf{m}^0$, used to generate the initial synthetics. The phase speed of the initial model is $c = 3.50$ km/s, the period of the source in the simulations is $\tau = 20$ s, and thus the reference wavelength is approximately $\lambda = c\tau = 70$ km.

Figure 2.11c shows the (smoothed) gradient for this model. The gradient is represented by the slope $\tilde{g}^0(0)$ of a line passing through $[0, \tilde{F}(0)]$ (Figure 2.11d). Quadratic extrapolation with a parabolic minimum at $(\nu_t^0, 0)$ gives the ν-value for a new test model (Figure 2.11d, Section 2.11). Figure 2.11e shows the test model, $\mathbf{m}_t^0$, for which we compute the gradient via the process shown in Figures 2.7–2.10, but now the model is no longer homogeneous. The gradient, shown in Figure 2.11f, is then depicted as the slope of a line passing through $[\nu_t^0, \tilde{F}(\nu_t^0)]$ (Figure 2.11g).

Next, in Figure 2.11g we approximate $\tilde{F}^0(\nu)$ by a cubic polynomial, $P^0(\nu)$, passing through two points, $[0, \tilde{F}^0(0)]$ and $[\nu_t^0, \tilde{F}(\nu_t^0)]$, and having slopes at these points corresponding to the respective gradients. In other words, six values are needed to obtain an analytical minimum of the cubic function: the two models (represented by $\nu = 0$ and ν_t^0), the misfits of these models, and the derivatives at these points (see Section 2.11). The minimum, $[\nu^0, P^0(\nu^0)]$, indicates the expected value of the misfit for the updated model given by $\mathbf{m}^1 = \mathbf{m}^0 - \nu^0 \mathbf{g}^0$, which is shown in Figure 2.11h and represented by the point

$(\nu^0 = 1.2 \times 10^4, 0)$ in Figure 2.11g. Figure 2.11i shows the decrease in the misfit function going from $\mathbf{m}^0$ to $\mathbf{m}^1$. The dashed curve is determined based on the nine iteration points in Figure 2.12.

Figure 2.11 thus constitutes one iteration of the conjugate gradient algorithm. The process is repeated, and the results are shown in Figure 2.12. Each iteration produces a model that looks qualitatively more similar to the target model shown in Figure 2.11a, and generates a lower value of the misfit function (2.6). We draw a best-fitting hyperbola to the $\log_{10}$ values to highlight the convergence.

We next use the seismologically more interesting Rayleigh wave phase-speed model in Figure 2.1. In comparison with Figure 2.11a, this model has variable scalelength and lower amplitude perturbations. The weaker perturbations result in a lower initial misfit, $F(\mathbf{m}^0) = 1182.0$ s^2. Figure 2.13 shows the recovery of an interior portion of the model, where path coverage is good. The basic features in the target phase-speed model (Figure 2.20a) are recovered by the third iteration (Figure 2.13d). The two sets of points in the Figure 2.13f are discussed in the next section. The model obtained after the first iteration, model $\mathbf{m}^1$ shown in Figure 2.13b, looks very similar to the model obtained based on a classical Hessian-based inversion with heavy damping shown in Figure 2.4g. This reflects the fact that in the conjugate gradient approach one is effectively working with an initial approximation to the Hessian that is the identity matrix.

2.6.3   Variations on the conjugate gradient algorithm

Based on the conjugate gradient algorithm outlined in Section 2.6.1, we require $4N_{\text{events}}$ numerical simulations for each iterative improvement of the model: synthetics for $\mathbf{m}^0$, the gradient for $\mathbf{m}^0$, synthetics for test model $\mathbf{m}_t^0$, and the gradient for $\mathbf{m}_t^0$. This information is used to compute the analytical minimum for a cubic polynomial. An alternative approach is to perform $3N_{\text{events}}$ numerical simulations per iteration by neglecting the gradient of the test model and using a quadratic polynomial to compute an analytic minimum (Section 2.11)

A comparison of these two approaches is shown in Figure 2.20f. The initial model for both cases has a misfit of $F(\mathbf{m}^0) = 1182.0$ s^2 for $N = 3300$ seismograms, corresponding to an average traveltime anomaly of $\overline{\Delta T}(\mathbf{m}^0) = 0.85$ s (eq. 2.46). Using $4N_{\text{events}}$ simulations with a cubic polynomial, we obtain a negligible advantage in terms of a better convergence of $F(\mathbf{m})$: for example, $F(\mathbf{m}_{\text{cubic}}^8) = 3.52$ s^2 whereas $F(\mathbf{m}_{\text{quad}}^8) = 3.75$ s^2 (Figure 2.20). To

the eye, the recovered models $\mathbf{m}_{\text{cubic}}^k$ and $\mathbf{m}_{\text{quad}}^k$ are indistinguishable.

An additional part of the conjugate gradient algorithm that can be adjusted is the selection of the test model, which we discuss in Section 2.11. Finally, we note that entrapment into local minima is common in the conjugate gradient method, as addressed in *Akçelik et al.* (2002, 2003). Such local minima may be avoided by using multiscale methods (*Bunks et al.*, 1995). Alternatively, by starting at longer periods, which constrain the long wavelength heterogeneity, and gradually moving to shorter periods, which constrain smaller scale structures, one can also try to avoid local minima.

2.7   Tomographic experiments

The greater the number of events used in the inversion, the better the recovery of the model. Figure 2.14 shows the recovery of the model in Figure 2.15i using 5, 15, and 25 sources, respectively.

Figure 2.15 examines the effect of the smoothing parameter, Γ, on the recovery of three different phase-speed models, each having a scalelength of structural heterogeneity that is proportional to the reference wavelength. Using a smaller Γ we resolve shorter-scalelength structures, whether they are in the target phase-speed model or not. When the scalelength of the smoothing exceeds that of the structure ($\Gamma > \Lambda$), the structure is smoothed out, as expected (Figure 2.15l).

The introduction of random errors into the cross-correlation traveltime measurements, ΔT_i, has essentially no impact on model recovery in our examples. For example, we denote a 50% error in the measurements by $\Delta T_i' = \Delta T_i \left(r + 0.5 \right)$, where $r \in [0, 1]$ is a random number, ΔT_i is the "actual" measurement, and $\Delta T_i'$ is the randomized measurement used in the inversion. In terms of the adjoint method, the introduction of random errors has the effect of changing the amplitude of the various banana-doughnut kernels that comprise the event kernel. Because the coverage in this example is very good, several similar kernels are "stacked" in constructing the event kernel, and thus the random errors effectively cancel.

2.8   Source, structure, and joint inversions

The traveltime differences between data and synthetics may be due to an inaccurate structural model, inaccurate source models, or some combination of both. In this section we

illustrate the simultaneous inversion for structural and source parameters using adjoint methods and the conjugate gradient algorithm. We first describe and illustrate the basic source inversion and then address the joint inversion.

2.8.1   Basic source inversion

A perturbation of the point source (2.30) may be written in the form

$$\delta f(x,y,t) = -\dot{h}(t)\delta t_s\, \delta(x-x_s)\, \delta(y-y_s) + h(t)(\delta x_s\, \partial_{x_s} + \delta y_s\, \partial_{y_s})[\delta(x-x_s)\, \delta(y-y_s)], \quad (2.47)$$

where δt_s denotes a perturbation in the origin time, $(\delta x_s, \delta y_s)$ a perturbation in the source location, and $\dot{h}(t) = \partial h/\partial t = -\partial h/\partial t_s$.

Based on the theory outlined in *Tromp et al.* (2005, Section 8), a change in the traveltime misfit function (2.7) due to a change in the point source is given by

$$\delta F = \int_0^T \int_\Omega \delta f(x,y,t) s^\dagger(x,y,T-t)\, \mathrm{d}x\, \mathrm{d}y\, \mathrm{d}t, \quad (2.48)$$

where $s^\dagger$ denotes the adjoint wavefield, whose sources are time-reversed, measurement-weighted seismograms, injected at the receivers, just as in the case of the structure inversions (2.33). (Here the traveltime measurement is affected by source perturbations only.) Upon substituting (2.47) we obtain

$$\delta F = -\,\delta t_s \int_0^T \dot{h}(t)\, s^\dagger(x_s,y_s,T-t)\, \mathrm{d}t + (\delta x_s\, \partial_{x_s} + \delta y_s\, \partial_{y_s}) \int_0^T h(t)\, s^\dagger(x_s,y_s,T-t)\, \mathrm{d}t. \quad (2.49)$$

We may express (2.49) in terms of the gradient as $\delta F = \mathbf{g} \cdot \delta\mathbf{m}$, where

$$\mathbf{m} = \begin{bmatrix} (x_s^k - x_s^0)/\lambda \\ (y_s^k - y_s^0)/\lambda \\ (t_s^k - t_s^0)/\tau \end{bmatrix}, \quad (2.50)$$

$$\mathbf{g} = \begin{bmatrix} \lambda \int_0^T h(t)\, \partial_{x_s} s^\dagger(x_s,y_s,T-t)\, \mathrm{d}t \\ \lambda \int_0^T h(t)\, \partial_{y_s} s^\dagger(x_s,y_s,T-t)\, \mathrm{d}t \\ -\tau \int_0^T \dot{h}(t)\, s^\dagger(x_s,y_s,T-t)\, \mathrm{d}t \end{bmatrix}. \quad (2.51)$$

Here $\mathbf{m}$ is a three-parameter nondimensionalized model vector describing the source. The source origin time t_s is scaled by the reference period τ, and the source coordinates are

scaled by the reference wavelength $\lambda = c\tau$, where c is the reference phase speed. The gradient vector, $\mathbf{g}$, depends on the model $\mathbf{m}$ through the adjoint wavefield $s^\dagger$: by perturbing the source, the measurement between data and synthetics changes, and thus the adjoint wavefield changes correspondingly.

In the experiments in Section 2.6, the sources for the data and synthetics were identical, whereas the structure was not. We now consider the effects of source perturbations, where the point sources for the initial synthetics are mislocated and initiate at an incorrect time.

Tromp et al. (2005, Fig. 12) demonstrated a two-parameter source inversion based on an adjoint method and the conjugate gradient algorithm. In that example, the two source parameters described the magnitude vector of the point source. In Figure 2.16 we illustrate a three-parameter source inversion for $\delta\mathbf{m} = (\delta x_s, \delta y_s, \delta t_s)$. The structural models for the data and synthetics are identical. We use the adjoint method to compute the gradient (2.51) of the misfit function (2.6). Using the conjugate gradient algorithm, we recover the source by the third iteration.

Finally, we emphasize that all of the equations in this section apply generally for any measurement, for example waveforms or amplitudes. The computed values for the expressions will differ, however, because the adjoint source $f^\dagger$ (and the corresponding adjoint wavefield $s^\dagger$) will vary for each measurement.

2.8.2 Joint inversions

In a joint inversion the sources and structure are initially different from the "target" sources and structure, and we seek to determine both. If we consider the three-parameter source inversion in Figure 2.16, then the model vector for the joint inversion is $\delta\mathbf{m} = [\delta\mathbf{m}_{\mathrm{str}} \, ; \, \delta\mathbf{m}_{\mathrm{src}}]$ with dimension $N_{\mathrm{structure}} + 3N_{\mathrm{event}}$. The misfit function is given by (2.6). We adjust the gradient of the misfit function at each iteration according to a constant, J, computed from the initial gradient:

$$\mathbf{g}^k = \left[J\mathbf{g}_{\mathrm{str}}^k \, ; \, \mathbf{g}_{\mathrm{src}}^k \right], \tag{2.52}$$

denoting a concatenation of the structure gradient $\mathbf{g}_{\mathrm{str}}$ computed via (2.32) and the source gradient $\mathbf{g}_{\mathrm{src}}$ computed via (2.51). The scaling factor J is given by

$$J = \|\mathbf{g}_{\mathrm{src}}^0\|_2 \ / \ \|\mathbf{g}_{\mathrm{str}}^0\|_2, \qquad\qquad\qquad\qquad (2.53)$$

where $\| \cdot \|_2$ denotes the L2-norm of the enclosed vector. The motivation behind (2.52) is that we want the source parameters and structural parameters to have about the same contribution in the gradient in the conjugate gradient algorithm. The exact choice of J, e.g., L1- versus L2-norm, is not important. Note that the factor J is chosen once and for all based on the *initial* structural and source gradients. Also, the gradients $(\mathbf{g}^k, \mathbf{g}_{\mathrm{str}}^k, \mathbf{g}_{\mathrm{src}}^k)$ are all with respect to the misfit function (2.6), evaluated at model $\mathbf{m}^k$.

Figure 2.17 compares a basic source inversion with a joint inversion. In the joint inversion the initial structural model is homogeneous, and the sources are mislocated randomly within 5 km of the target source and have an inaccurate origin time within the range ± 1 s. The two misfit curves in Figure 2.17d show that the joint inversion does almost as well as the basic structure inversion; in fact, it lags the misfit by only one or two iterations. In the final model only the sources on the edges of the grid contain location and timing errors (Figure 2.19c to Figure 2.17f), which is expected since there are few, if any, paths to constrain the structure.

Figure 2.18 shows the recovery of a single source during the joint inversion. It takes approximately 16 iterations to fully recover the source (instead of the 3 iterations in Figure 2.16 for the basic source inversion), although most of the source location is still recovered in the first few iterations. This increase is, of course, due to the gradual improvement of the structural parameters, which initially differ by up to 10% from the target structure.

In an inversion with real data, the initial model is bound to be deficient both in terms of structure and sources. Thus, a joint inversion is a logical approach. Figure 2.19 shows the consequences of neglecting either source or structure in the inversion. In Figure 2.19a–c we invert for structure and assume that the sources are accurate, when in fact they are perturbed as shown in (c). The misfit curve in (a) shows that the conjugate gradient approach appears to be working: the misfit decreases as the structure iterates to updated models. However, it is clear that Figure 2.19b does not represent the true structure, since we know the target model we are trying to recover, as well as its associated misfit curve for the basic structure inversion. This illustrates how (fixed) errors in the source parameters are

mapped into errors in the structural parameters. Figure 2.19d–f shows the opposite scenario: the structure is fixed and assumed to be accurate, and we allow the source parameters to be perturbed to reduce the traveltime misfit. The source parameters adjust themselves from Figure 2.19c (initial) to Figure 2.19f (final) while reducing the misfit.

2.9   Discussion

2.9.1   Three types of sensitivity kernels

We have designated three types of sensitivity kernels: "banana-doughnut kernels," "event kernels," and "misfit kernels." A banana-doughnut kernel (e.g., *Marquering et al.*, 1999) is a phase-specific (e.g., P) kernel for an individual source-receiver combination; for our purposes, the key point is that this kernel does not incorporate the measurement. Alternative names include "finite-frequency," "Born," and "sensitivity" kernel. An event kernel can be thought of as a sum of individual banana-doughnut kernels, such that each kernel in the sum is weighted by its corresponding measurement. Using the adjoint approach, however, the event kernel is not computed by summing banana-doughnut kernels, but rather in one single simulation through the interaction between the forward wavefield and an adjoint wavefield generated by simultaneous fictitious sources for all available arrivals at all available stations (Section 2.5.1). A misfit kernel is simply the sum of event kernels, and may be thought of as a graphical representation of the gradient of the misfit function. In classical tomography, the banana-doughnut kernels are used to compute the gradient and (approximate) Hessian of the misfit function for the Newton approach to the inverse problem. In adjoint tomography, only the misfit kernels are used in the inverse problem.

2.9.2   Classical tomography versus adjoint tomography

In this paper, "classical tomography" refers to Hessian-based inversions, whereby the Hessian is constructed from individual source-receiver paths, either in terms of rays or finite frequency kernels. The Hessian matrix, with a damping parameter γ, can be inverted to obtain structural models. We compute the traveltime anomalies, and thus F, via (2.6), and then compare these values with those obtained from gradient-derived models.

Figure 2.20 shows a comparison among models produced using classical tomography, $\mathbf{m}^{\mathrm{ray}}$ (ray-based inversion) and $\mathbf{m}^{\mathrm{ker}}$ (kernel-based inversion), and the model produced using

adjoint tomography, $\mathbf{m}^{16}$ (16 conjugate gradient iterations). All three models are only subtly different from the target model used to generate the data (Figure 2.20a). The misfit for each approach is summarized in Figure 2.20e. The misfit values for the classical models, $F(\mathbf{m}^{\text{ray}}) = 5.26$ s^2 and $F(\mathbf{m}^{\text{ker}}) = 4.90$ s^2, correspond to average traveltime anomalies of $\overline{\Delta T}(\mathbf{m}^{\text{ray}}) = 0.056$ s and $\overline{\Delta T}(\mathbf{m}^{\text{ker}}) = 0.055$ s (eq. 2.46), indicating that each recovered model explains almost all of the traveltime differences between a homogeneous model and the target model in Figure 2.1b. Two points regarding the two F-values are important: (1) they are essentially the same, which is expected, since the Hessian used in each inversion is very similar (Figure 2.3a); (2) they are met by the conjugate gradient approach by the seventh conjugate gradient iteration. In other words, $F(\mathbf{m}^{\text{ray}}) \approx F(\mathbf{m}^{\text{ker}}) \approx F(\mathbf{m}^7)$; after seven conjugate gradient iterations, we recover a model equivalent to what could be recovered by having the (ray- or kernel-based) Hessian. It is important to note that $\mathbf{m}^{\text{ray}}$, $\mathbf{m}^{\text{ker}}$, and $\mathbf{m}^1$ are based on the homogeneous reference model $\mathbf{m}^0$, but for $k > 1$, the adjoint tomography models $\mathbf{m}^k$ are based on heterogeneous models.

Figure 2.20 might suggest that classical tomography "does pretty well" in comparison with adjoint tomography. This is more or less true for the simple examples in this paper. However, seismic tomography is transitioning from simple 1D reference models to fully 3D reference models. The calculation of a Hessian for 3D reference models is generally not an option, and thus one must resort to iterative, gradient-based algorithms. The results in this paper illustrate that for the problems considered here, such iterative techniques work quite efficiently and converge quickly.

The main advantages of the adjoint tomography approach are fivefold. First, all the complexities that are considered in the forward problem (e.g., *Komatitsch and Tromp*, 2002a,b) can be considered in the inversion. For example, in this paper we have shown finite-frequency sensitivity kernels based on heterogeneous models. But one could also consider fully anisotropic Earth models with 21 elastic parameters for essentially the same numerical cost as an isotropic simulation involving just two parameters. Second, the style of tomography — traveltime, amplitude, waveform — is determined by the choice of the misfit function (*Tromp et al.*, 2005). Given the choice of measurement, one simply determines the associated adjoint source that gives rise to the corresponding kernel. Third, any time segment where the data and synthetics match reasonably well is suitable for a measurement. One does not need to label a particular phase, like P or SS, because the adjoint simulation

will reveal how this particular measurement (or time window) "sees" the Earth model, and the resulting sensitivity kernel will reflect this view. Fourth, the model parameterization is trivial (2.43) and requires only a conservative level of smoothing to remove numerical artifacts in the kernels near the sources and receivers (Section 2.5.2). Furthermore, structure can only be introduced in regions where the kernel (or gradient) is nonzero. This is in contrast to classical tomography, where both the selection of basis functions and the choice of damping involve a certain degree of undesired subjectivity. Finally, the approach scales linearly with the number of earthquakes but is independent of the number of receivers and the number of arrivals that are used in the inversion.

With southern California in mind, say we have $N_{\text{events}} = 150$ earthquakes, $N_{\text{receivers}} = 150$ SCSN stations, $N_{\text{comp}} = 3$ components per seismogram, and $N_{\text{picks}} = 4$ time-windowed measurements per component, for a total of $N_{\text{events}} N_{\text{receivers}} N_{\text{comp}} N_{\text{picks}} = 270,000$ measurements. An adjoint approach would require $2 N_{\text{events}} = 300$ simulations to compute one misfit kernel. A complete 7-iteration conjugate gradient inversion based on cubic interpolation would require $7 \times 3 N_{\text{events}} = 3150$ total simulations. By comparison, a Hessian-based inversion would require individual kernels for the 270,000 measurements, which, for 3D models, is neither computationally feasible nor practical.

2.9.3   Feasibility of 3D–3D tomography

This paper is a step toward "3D–3D tomography," denoting 3D heterogeneity within the reference models and a 3D physical domain for the model, from which we compute finite-frequency sensitivity kernels. (Based on this labeling, the classical tomographic examples in this paper are 0D–2D, whereas the adjoint tomographic examples are 2D–2D.) Presently our SEM codes are set up to compute 3D–3D sensitivity kernels on both regional and global scales (*Liu and Tromp*, 2006, 2008). In this paper we have highlighted some aspects of the inversion process that will be key to limiting the number of wavefield simulations required in the inversion.

Let us estimate the computational cost of a regional-scale tomographic inversion. As discussed, for 150 earthquakes we require 3150 total simulations for a 7-iteration inversion. Each simulation takes approximately 35 minutes on 72 nodes (144 processors). Thus we can perform 40 runs per day on 72 nodes, and more than 500 runs per day on a 1000-node machine. Therefore, on this kind of hardware the whole inversion can theoretically be

completed in about one week.

To avoid reaching a local minimum in the optimization procedure, we intend to start by using longer-period waveforms, which existing 3D models fit reasonably well, and work our way toward shorter periods. As we improve the model and increase the frequency contents of the waveforms, we expect to not only improve the fit to the current data used to constrain the model, but also to steadily increase the number of picks that is used in the inversion, i.e., more and more parts of the seismograms are expected to be used and matched in the iterative inversion process. Unlike *Akçelik et al.* (2003), our emphasis will be on matching targeted, frequency-dependent body-wave traveltimes and surface-wave phase anomalies, rather than entire waveforms. Waveform tomography is largely controlled by amplitude differences, which are notoriously difficult to fit in seismology. Traveltime or phase, on the other hand, is a robust measure of misfit that has been used for decades to constrain local, regional and global Earth models. From our perspective, the progression from ray-based traveltime tomography to finite-frequency "banana-doughnut" tomography to frequency-dependent adjoint tomography is very natural.

2.10   Appendix A: Regularization

Here we review the fact that stabilizing the Hessian matrix (as in eq. 2.5) via damping is equivalent to adding a regularization term R to the misfit function (2.6):

$$F_R(\mathbf{m}) = F(\mathbf{m}) + R(\mathbf{m}), \tag{2.54}$$

whose gradient is, using (2.20),

$$\frac{\partial F_R}{\partial m_k} = \frac{\partial F}{\partial m_k} + \frac{\partial R}{\partial m_k} = -\sum_{i=1}^{N} G_{ik} \, \Delta T_i + \frac{\partial R}{\partial m_k}. \tag{2.55}$$

There are many options for regularization. For illustrative purposes, we consider regularization according to the wave-speed model itself:

$$R = \tfrac{1}{2}\gamma^2 \int_V (\delta \ln c)^2 \, \mathrm{d}^3\mathbf{x}, \tag{2.56}$$

where γ is the damping parameter. Substituting (2.13), and then differentiating with respect

to the kth model parameter, we obtain

$$
R = \tfrac{1}{2}\gamma^2 \sum_{k=1}^{M} \delta m_k \sum_{k'=1}^{M} \delta m_{k'} D_{kk'}, \tag{2.57}
$$

$$
\frac{\partial R}{\partial m_k} = \gamma^2 \sum_{k'=1}^{M} \delta m_{k'} D_{kk'}, \tag{2.58}
$$

where the $M \times M$ damping matrix $\mathbf{D}$ is given by

$$
D_{kk'} = \int_V B_k \, B_{k'} \, \mathrm{d}^3 \mathbf{x}. \tag{2.59}
$$

If the basis functions are orthonormal, then $\mathbf{D} = \mathbf{I}$, the identity matrix. Substituting (2.58) into (2.55), we obtain

$$
\frac{\partial F_R}{\partial \mathbf{m}} = -\mathbf{G}^T \mathbf{d} + \gamma^2 \mathbf{D} \delta \mathbf{m}, \tag{2.60}
$$

where $\mathbf{D}$ is, for example, (2.59) or (2.63). Substituting this for $\mathbf{g}(\mathbf{m})$ into (2.5), with $\mathbf{H} = \mathbf{G}^T \mathbf{G}$, we obtain

$$
\left(\mathbf{G}^T \mathbf{G} + \gamma^2 \mathbf{D} \right) \delta \mathbf{m} = \mathbf{G}^T \mathbf{d}, \tag{2.61}
$$

which leads to (2.28). Equation (2.61) is known as "Tikhonov" regularization or "ridge regression," and is based on minimizing an L2-norm measure of $\mathbf{D} \, \delta \mathbf{m}$ (e.g., *Hansen*, 1998, Ch. 5). (Typically these two labels refer to the case $\mathbf{D} = \mathbf{I}$.)

Instead, if we regularize using the gradient of the wave-speed model (e.g., *Akçelik et al.*, 2003), we obtain

$$
R = \tfrac{1}{2}\gamma^2 \int_V \boldsymbol{\nabla}(\delta \ln c) \cdot \boldsymbol{\nabla}(\delta \ln c) \, \mathrm{d}^3 \mathbf{x}, \tag{2.62}
$$

then the damping matrix is

$$
D_{kk'} = \int_V (\boldsymbol{\nabla} B_k) \cdot (\boldsymbol{\nabla} B_{k'}) \, \mathrm{d}^3 \mathbf{x}. \tag{2.63}
$$

Regularization according to the *roughness* of the model (e.g., *Zhou et al.*, 2005) leads to

$$R = \tfrac{1}{2}\gamma^2 \int_V \left(\nabla^2 \delta \ln c\right)^2 \mathrm{d}^3\mathbf{x}, \tag{2.64}$$

$$D_{kk'} = \int_V \left(\nabla^2 B_k\right)\left(\nabla^2 B_{k'}\right) \mathrm{d}^3\mathbf{x}. \tag{2.65}$$

Different norms or constraints, as well as combinations of constraints (resulting in multiple damping parameters), may be used in forming $R(\mathbf{m})$. For example, *Akçelik et al.* (2002) advocated the use of L1-based, "total variation" regularization, which avoids smoothing of sharp gradients in material properties. *Akçelik et al.* (2003) applied L1 regularization for the structure gradients and L2 regularization for the source gradients.

2.11 Appendix B: Details of the conjugate gradient algorithm

The computation of the misfit value $F(\mathbf{m})$ and gradient $\mathbf{g}$ takes $2N_{\text{events}}$ simulations. Because each simulation is expensive, it is important to limit the number of simulations in the inverse problem, which we approach using a conjugate gradient algorithm (Section 2.6.1). Two possible areas to aid in this are the selection of the trial step ν_t and the choice of the polynomial (quadratic or cubic) to use in the interpolation. In this section we have omitted the superscript k on quantities to avoid clutter.

2.11.1 Selection of the trial step

The trial step, or test parameter, ν_t, determines how far away from the current model to go in the search direction in order to obtain a test model (and, possibly, test gradient). Given the misfit value, $F(\mathbf{m})$, and the gradient, $\mathbf{g}(\mathbf{m})$, for the current model, the user is faced with determining how far to step in the search direction away from the current model to obtain a test model, for which an additional misfit value will be computed. The gradient vector $\mathbf{g}$ is represented in the conjugate gradient algorithm as a slope, $\tilde{g}(0)$, and the misfit function in the search direction by $\tilde{F}(\nu)$. In the algorithm, we select the test parameter by interpolating $\tilde{F}(\nu)$ using a quadratic polynomial, $Q(\nu)$:

$$Q(\nu) = a\nu^2 + b\nu + c, \tag{2.66}$$

where a, b, and c are determined using the value (r_1) and slope (g_1) for the current model, and a test model location such that the value (r_2) and slope (g_2) of $Q(\nu)$ at ν_t are both zero (see Figure 2.11d). The four values are given by

$$
\begin{aligned}
r_1 &\equiv Q(0) = \tilde{F}(0) = F(\mathbf{m}), \\
g_1 &\equiv Q'(0) = \tilde{g}(0), \\
r_2 &\equiv Q(\nu_t) = 0, \\
g_2 &\equiv Q'(\nu_t) = 0.
\end{aligned}
$$

These equations can be used to determine the coefficients of $Q(\nu)$:

$$
\begin{aligned}
a &= -g_1/(2\nu_t) = g_1^2/(4r_1), \\
b &= g_1, \\
c &= r_1,
\end{aligned}
$$

as well as the test parameter

$$
\nu_t = \frac{-2r_1}{g_1} = \frac{-2F(\mathbf{m})}{\tilde{g}(0)}, \tag{2.67}
$$

which is the value used in the algorithm discussed in Section 2.6.1.

The "test model parabola" $Q(\nu)$ is chosen such that its vertex lies on $F = 0$; however, one could choose the vertex at some $F > 0$ that is determined based on the change in misfit from a previous step. The quadratic extrapolation through $F = 0$ is perhaps too conservative, and computational savings — in the form of better convergence — could be obtained by exploring the choice of the initial step.

2.11.2   Quadratic versus cubic interpolation

As discussed in Section 2.6.3, the tomographer is faced with the choice of using a second- or third-order polynomial, $P(\nu)$, in the interpolation scheme within the conjugate gradient algorithm. Here we outline the formulas required to compute an analytical minimum, $\nu_{\min}$, using each interpolation scheme.

With an order-3 polynomial, one must have four quantities in addition to the test parameter ν_t: the misfit and gradient for the current model, $F(\mathbf{m})$ and $\mathbf{g}$, and the misfit

and gradient for the test model, $F(\mathbf{m}_t)$ and $\mathbf{g}_t$. These values are converted into scalar values for an interpolating polynomial $P(\nu)$:

$$
\begin{aligned}
\nu_1 &= 0, \\
\nu_2 &= \nu_t, \\
p_1 &\equiv P(\nu_1) = F(\mathbf{m}), \\
g_1 &\equiv P'(\nu_1) = \tilde{g}(0), \\
p_2 &\equiv P(\nu_2) = F(\mathbf{m}_t), \\
g_2 &\equiv P'(\nu_2) = \tilde{g}(\nu_t).
\end{aligned}
$$

The cubic polynomial can be written in terms of these quantities as

$$
P(\nu) = a\,(\nu - \nu_1)^3 + b\,(\nu - \nu_1)^2 + c\,(\nu - \nu_1) + d, \tag{2.68}
$$

where

$$
\begin{aligned}
a &= \left[-2(p_2 - p_1) + (g_1 + g_2)(\nu_2 - \nu_1)\right] / (\nu_2 - \nu_1)^3, \\
b &= \left[3(p_2 - p_1) - (2g_1 + g_2)(\nu_2 - \nu_1)\right] / (\nu_2 - \nu_1)^2, \\
c &= g_1, \\
d &= p_1.
\end{aligned}
$$

An analytical minimum of $P(\nu)$ can be obtained when $|c| > 0$:

$$
\nu_{\min} =
\begin{cases}
\nu_1 + \left[-b + (b^2 - 3ac)^{1/2}\right]/(3a) & a \neq 0 \text{ and } b^2 - 3ac > 0 \\
-c/(2b) & a = 0 \text{ and } b \neq 0;\ b^2 - 3ac < 0.
\end{cases}
\tag{2.69}
$$

With an order-2 polynomial, the gradient of the test model — $\mathbf{g}_t$ or g_2 — is not required. In this case, the quadratic polynomial can be written in terms of (2.68) as

$$
P(\nu) = a\,(\nu - \nu_1)^2 + b\,(\nu - \nu_1) + c, \tag{2.70}
$$

where

$$
\begin{aligned}
a &= \left[(p_2 - p_1) - g_1(\nu_2 - \nu_1)\right] / \left(\nu_2^2 - \nu_1^2\right), \\
b &= g_1, \\
c &= p_1 - a\nu_1^2 - b\nu_1.
\end{aligned}
$$

The analytical minimum of $P(\nu)$ is simply

$$
\nu_{\min} = -b/(2a). \tag{2.71}
$$

Based on our experiments, the quadratic interpolation is preferred over the cubic interpola-tion, since it costs $3N_{\text{events}}$ per iteration (versus $4N_{\text{events}}$) and performs only slightly worse (Figure 2.13f).
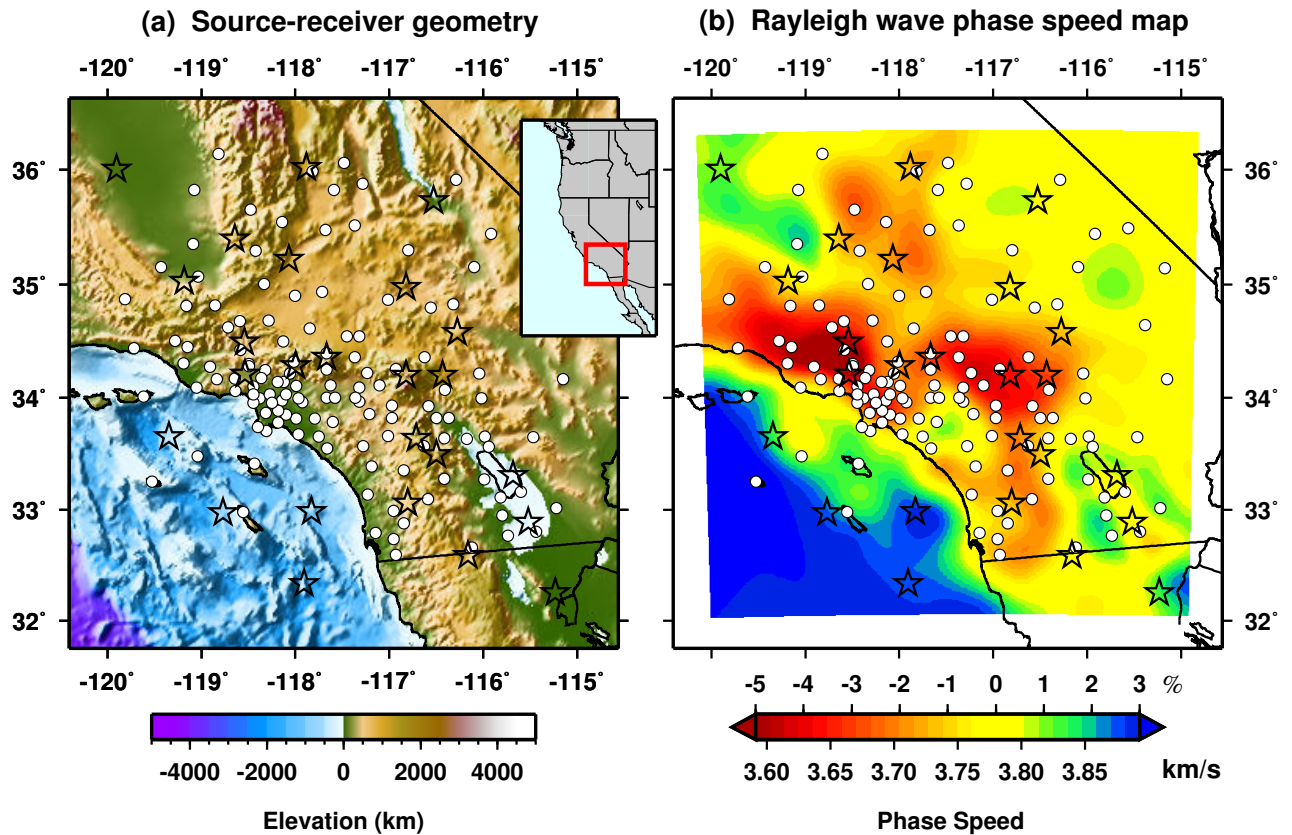
Figure 2.1: Source–receiver geometry for the numerical experiments in this study. The ☆ symbols denote the locations of 25 earthquakes (each has a $M \geq 4.0$ and occurred between 1990 and 2005); the ○ symbols denote the locations of 132 broadband receivers in the Southern California Seismic Network (SCSN). The earthquakes are selected to obtain relatively uniform coverage; all SCSN receivers in the area are included. (a) Topography and bathymetry in the region. (b) Phase speed map for 20 s Rayleigh waves, based on the regional model of *Hauksson* (2000), modified with the Moho map of *Zhu and Kanamori* (2000). This phase speed map is used to generate synthetic data used in some of the inversion experiments.

(a) Ray : i = 126 / 3300 (event 1, rec 126)

(b) Spline : k = 203 / 286

(c) Spline along ray (G_ik = -6.46)

Distance along ray path (10^{05} m)

(d) Kernel : i = 126 / 3300 (event 1, rec 126)

K (φ, θ) (10^{-09} m^{-2} s)

(e) Spline : k = 150 / 286

B (φ, θ)

(f) Kernel * Spline (G_ik = -1.03)

K * B (10^{-09} m^{-2} s)
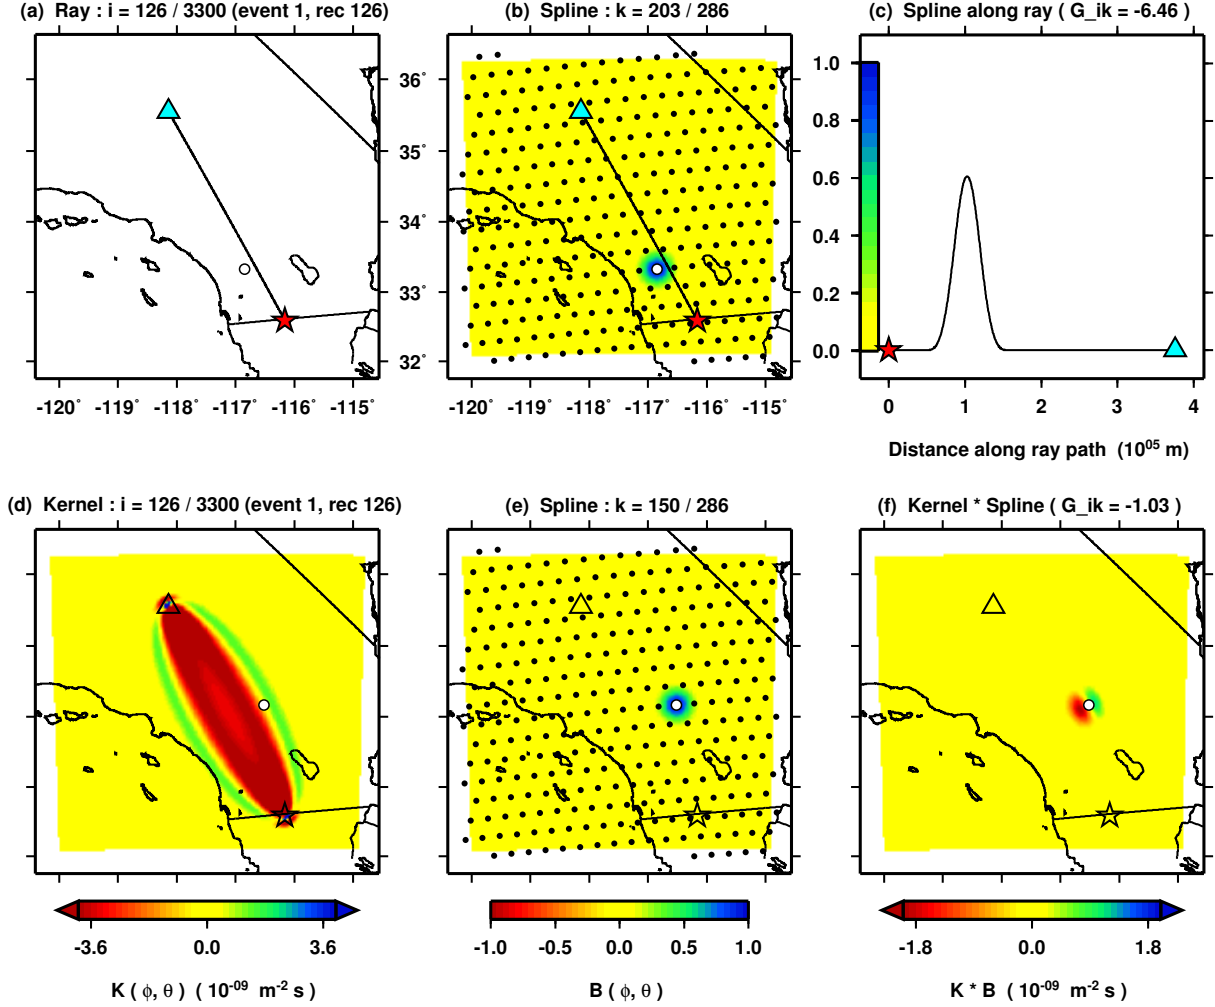
Figure 2.2:   Example computation for an element, G_{ik}, of the design matrix $\mathbf{G}$, using rays (a-c) and finite-frequency kernels (d-f). The row index i is the source-receiver combination, the column index k is the basis function index. The source is denoted by the ☆, the receiver is denoted by the △, and the ∘ shows the center-point of the spherical spline in (b) or (e). (a) Ray path for event number 1 and receiver number 126 (Figure 2.1), corresponding to the $i = 126$ index of the $N = 3300$ ray paths. (b) $B_{203}(\mathbf{x})$, the spherical spline basis function for index $k = 203$. Also shown are the center-points of the $M = 286$ spherical splines. (c) Spline B_{203} evaluated along the ray path. The value of the phase speed for the reference model is constant, so $G_{ik} = (-1/c) \int_{\mathrm{ray}_i} B_k \, ds$ (eq. 2.15). In this example $G_{ik} = -1/(3780 \, \mathrm{m \, s^{-1}}) \, (2.45 \times 10^4 \, \mathrm{m}) = -6.46$ s. (d) Cross-correlation traveltime sensitivity kernel for event number 1 and receiver number 126 (Figure 2.1), corresponding to the $i = 126$ index of the $N = 3300$ kernels. (e) $B_{150}(\mathbf{x})$, the spherical spline basis function for index $k = 150$. Also shown are the center-points of the $M = 286$ spherical splines. (f) The function $K_{126}(\mathbf{x}) \, B_{150}(\mathbf{x})$. The integral of this function gives the value $G_{ik} = \int_{\Omega} K_{126} \, B_{150} \, d^2\mathbf{x} = -1.03$ s. (See Section 2.2.)
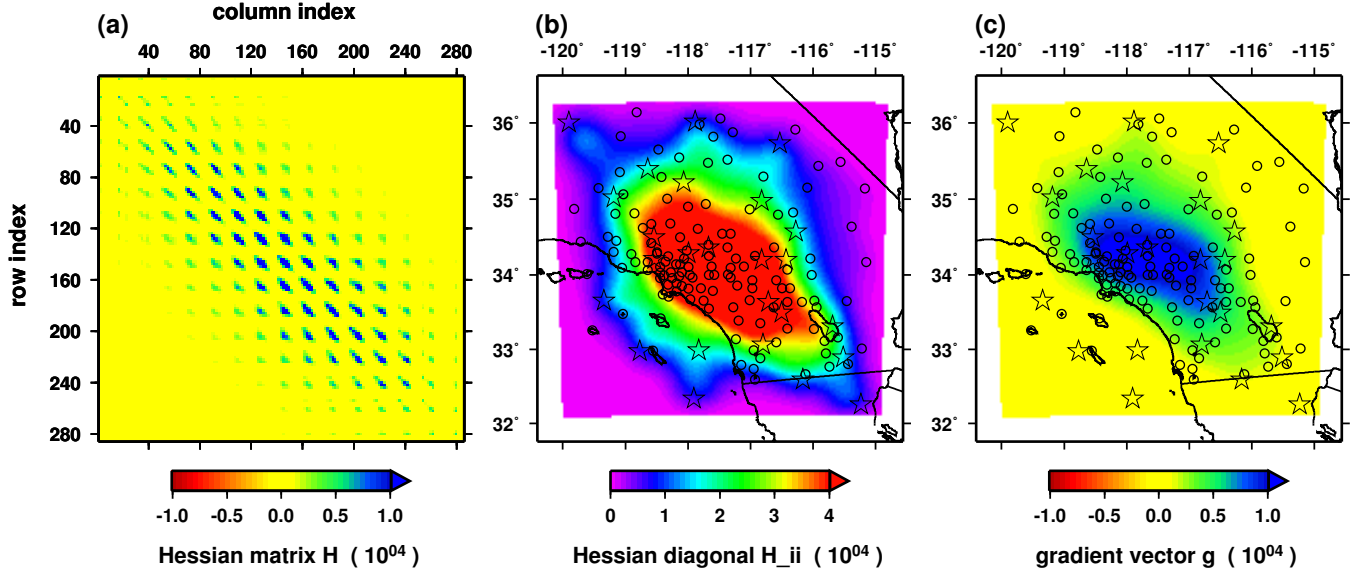
Figure 2.3:   The Hessian matrix and gradient vector for a classical tomography inversion. (a) The Hessian matrix, $\tilde{\mathbf{H}} = \mathbf{G}^T \mathbf{G}$, for the source-receiver geometry shown in Figure 2.1, using finite-frequency kernels based on cross-correlation traveltime measurements. Each element of $\mathbf{G}$ is constructed as shown in Figure 2.2d-f. The Hessian matrix computed using rays, as shown in Figure 2.2a-c, is indistinguishable from the one shown in this figure. In practice, a damping matrix is added to the Hessian to stabilize the inversion (2.27). (b) The diagonal elements of the Hessian matrix, $\tilde{H}_{ii}$, expanded in the spherical spline basis functions to illuminate the spatial pattern (e.g., *Zhou et al.*, 2005, Figure 3). This map is a proxy for spatial coverage and depends on the source-receiver geometry, the basis functions, and the sensitivity kernels. (c) The gradient vector, $\mathbf{g} = -\mathbf{G}^T \mathbf{d}$, expanded in the spherical spline basis functions, whereby $\mathbf{d}$ includes cross-correlation traveltime measurements between data computed for the target phase-speed model in Figure 2.1b and synthetics computed for a homogeneous phase-speed model ($c = 3.78$ km/s). The ☆ symbols denote the sources, and the ○ symbols denote the receivers.

(a) Model for γ = 0.01

(b) Model for γ = 1.00

(c) Model for γ = 10.00

(d) Model for γ = 31.62

(e) Model for γ = 100.00

(f) Model for γ = 1000.00

(g) Model for γ = 1000.00

(h) L-curve for the models

(i) Target phase speed model (data)
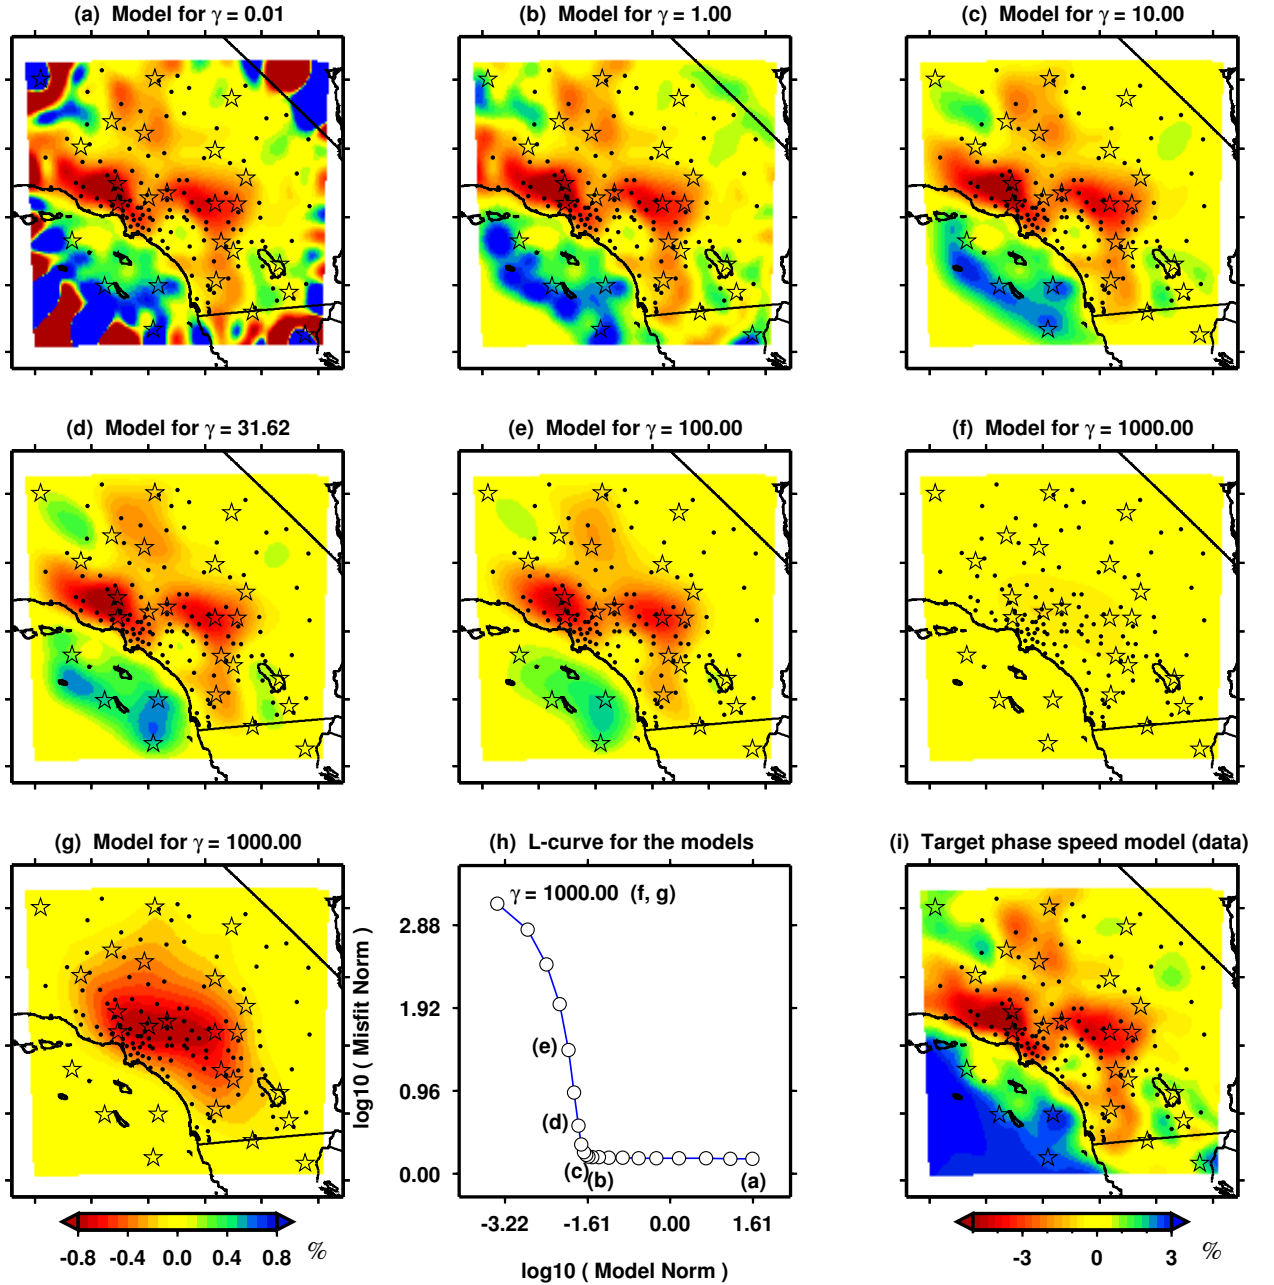
Figure 2.4:   Model recovery and damping in classical tomography, illustrated for an inversion using 3300 banana-doughnut kernels. Each model is computed via $\mathbf{m} = -\tilde{\mathbf{H}}_\gamma^{-1}\mathbf{g}$, where $\tilde{\mathbf{H}}_\gamma = \mathbf{G}^T\mathbf{G} + \gamma^2\mathbf{I}$ is the Hessian matrix with damping parameter γ, and $\mathbf{g} = -\mathbf{G}^T\mathbf{d}$ is the gradient vector (Figure 2.3c). The undamped Hessian matrix, $\tilde{\mathbf{H}}_0$, is shown in Figure 2.3a-b. (a)–(f) Recovered phase-speed models for various values of γ. The color scale for each model is shown below (i). (g) Same as (f), but with a more saturated color scale to show its resemblance to the gradient (Figure 2.3c). (h) *L*-curve illustrating the trade-off between misfit norm and model norm, that is, $\|\mathbf{Gm} - \mathbf{d}\|_2$ versus $\|\mathbf{m}\|_2$. Note that this measure of misfit is not the same as $\mathbf{d}^T\mathbf{d}$, the misfit function in (2.6). The γ values for the model-points are spaced by uniform $\log_{10}$ increments. (i) Target phase-speed model used to generate the data (Figure 2.1b). The ☆ symbols denote the sources, and the ● symbols denote the receivers. See Section 2.3.3.
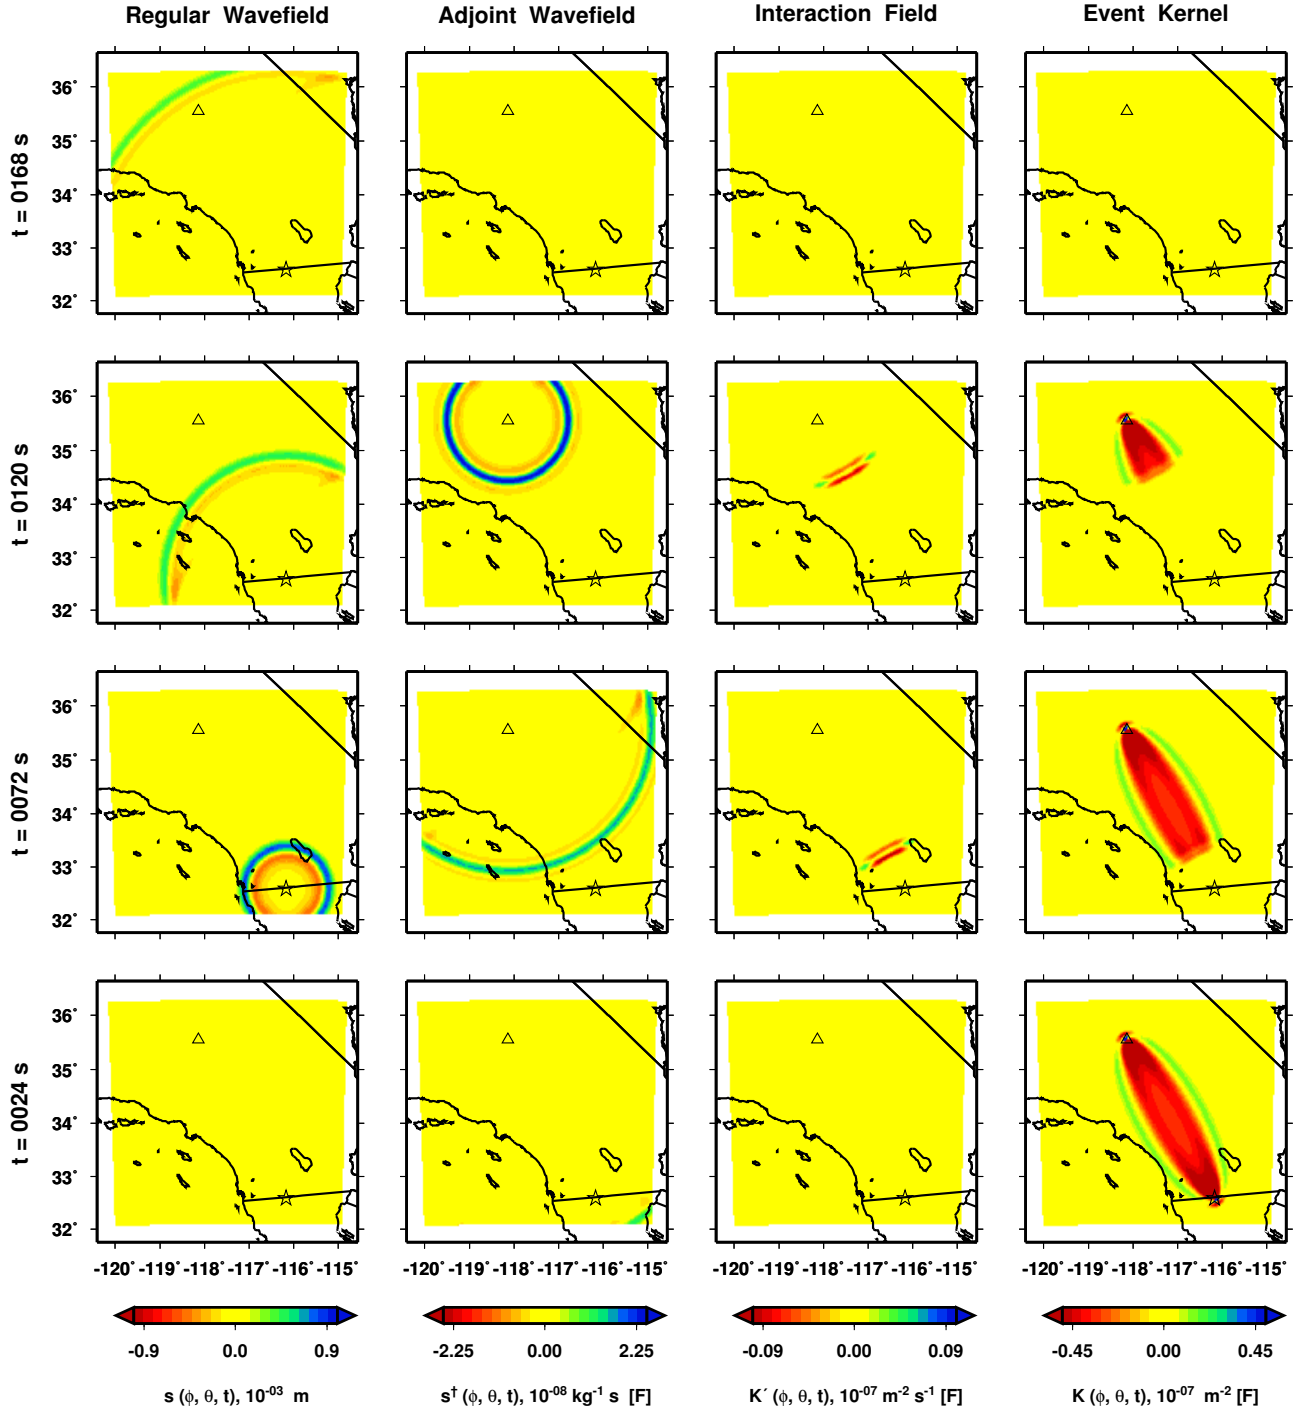
Figure 2.5: Sequence of interactions between the regular and adjoint wavefields during the construction of a traveltime cross-correlation event kernel $K(\mathbf{x})$. The ☆ symbol denotes the source, and the △ symbol denotes the receiver. Each row represents the time-step indicated on the left. In this case, with only a single receiver and a uniform model perturbation, the event kernel resembles a banana-doughnut kernel $K_i(\mathbf{x})$. The event kernel is constructed via the interaction between the forward wavefield (first column) and the adjoint wavefield (second column). The interaction field (third column) is the instantaneous product of the two wavefields, which is integrated to form the event kernel (fourth column). The event kernel shows the region of the current model that gives rise to the discrepancy between the data and the synthetics. The regular source function and adjoint source function are shown in Figure 2.6.
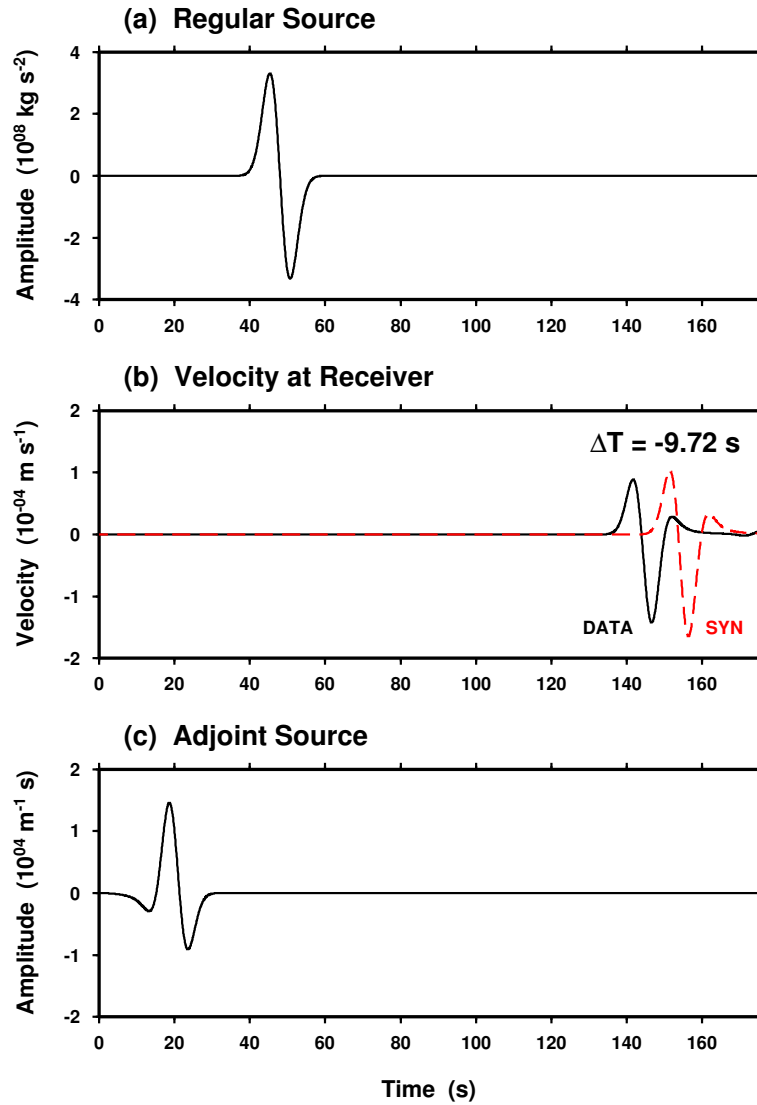
Figure 2.6:   Construction of an adjoint source function used in calculating the membrane surface-wave event kernel in Figure 2.5. The traveltime sign convention is shown in (2.8), such that $\Delta T < 0$ represents a delay of the synthetics with respect to the data. The duration of the simulation is $T = 175$ s.
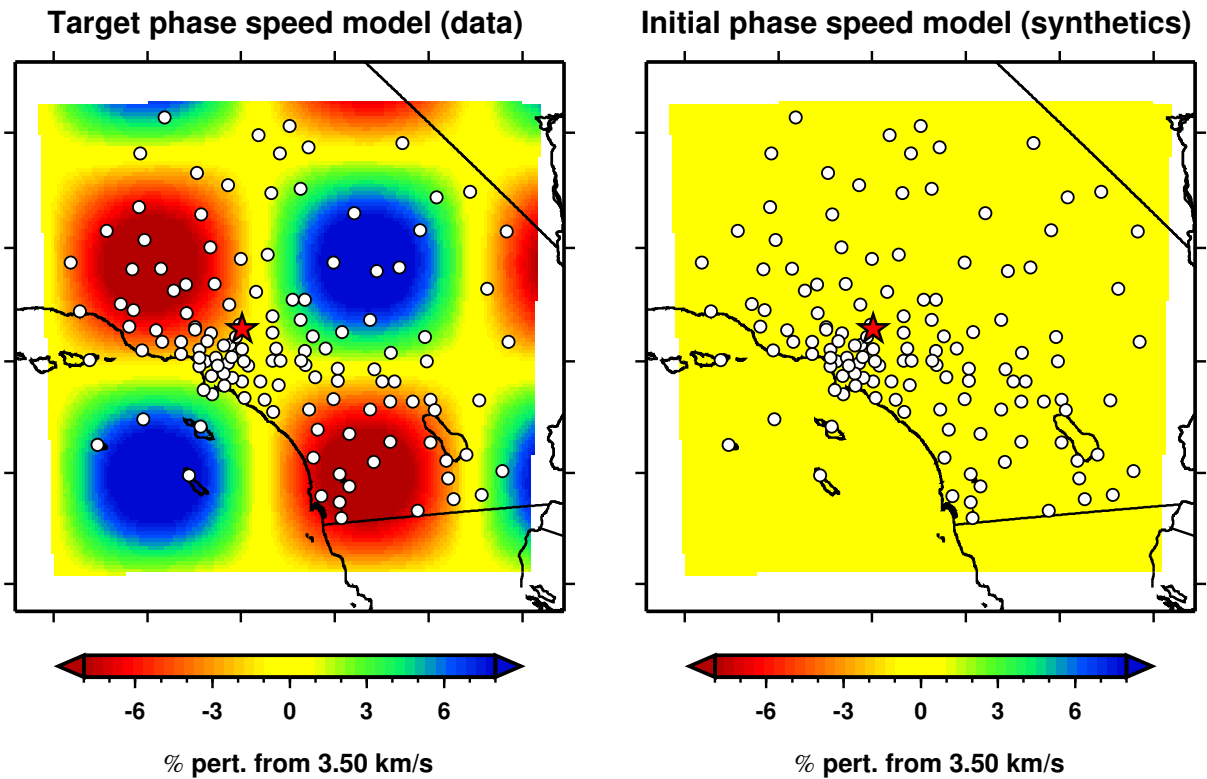
Figure 2.7: Experimental setup for the event kernel shown in Figure 2.8. The data are computed using the target phase-speed model, and the synthetics are computed using the initial model. The minimum and maximum percent perturbations in the target model are ±10%. The red star is the event location, and the circles denote the 132 receivers. For plotting purposes, the gridpoints are converted to longitude-latitude points, which results in the nonrectangular appearance of the boundary of the grid.
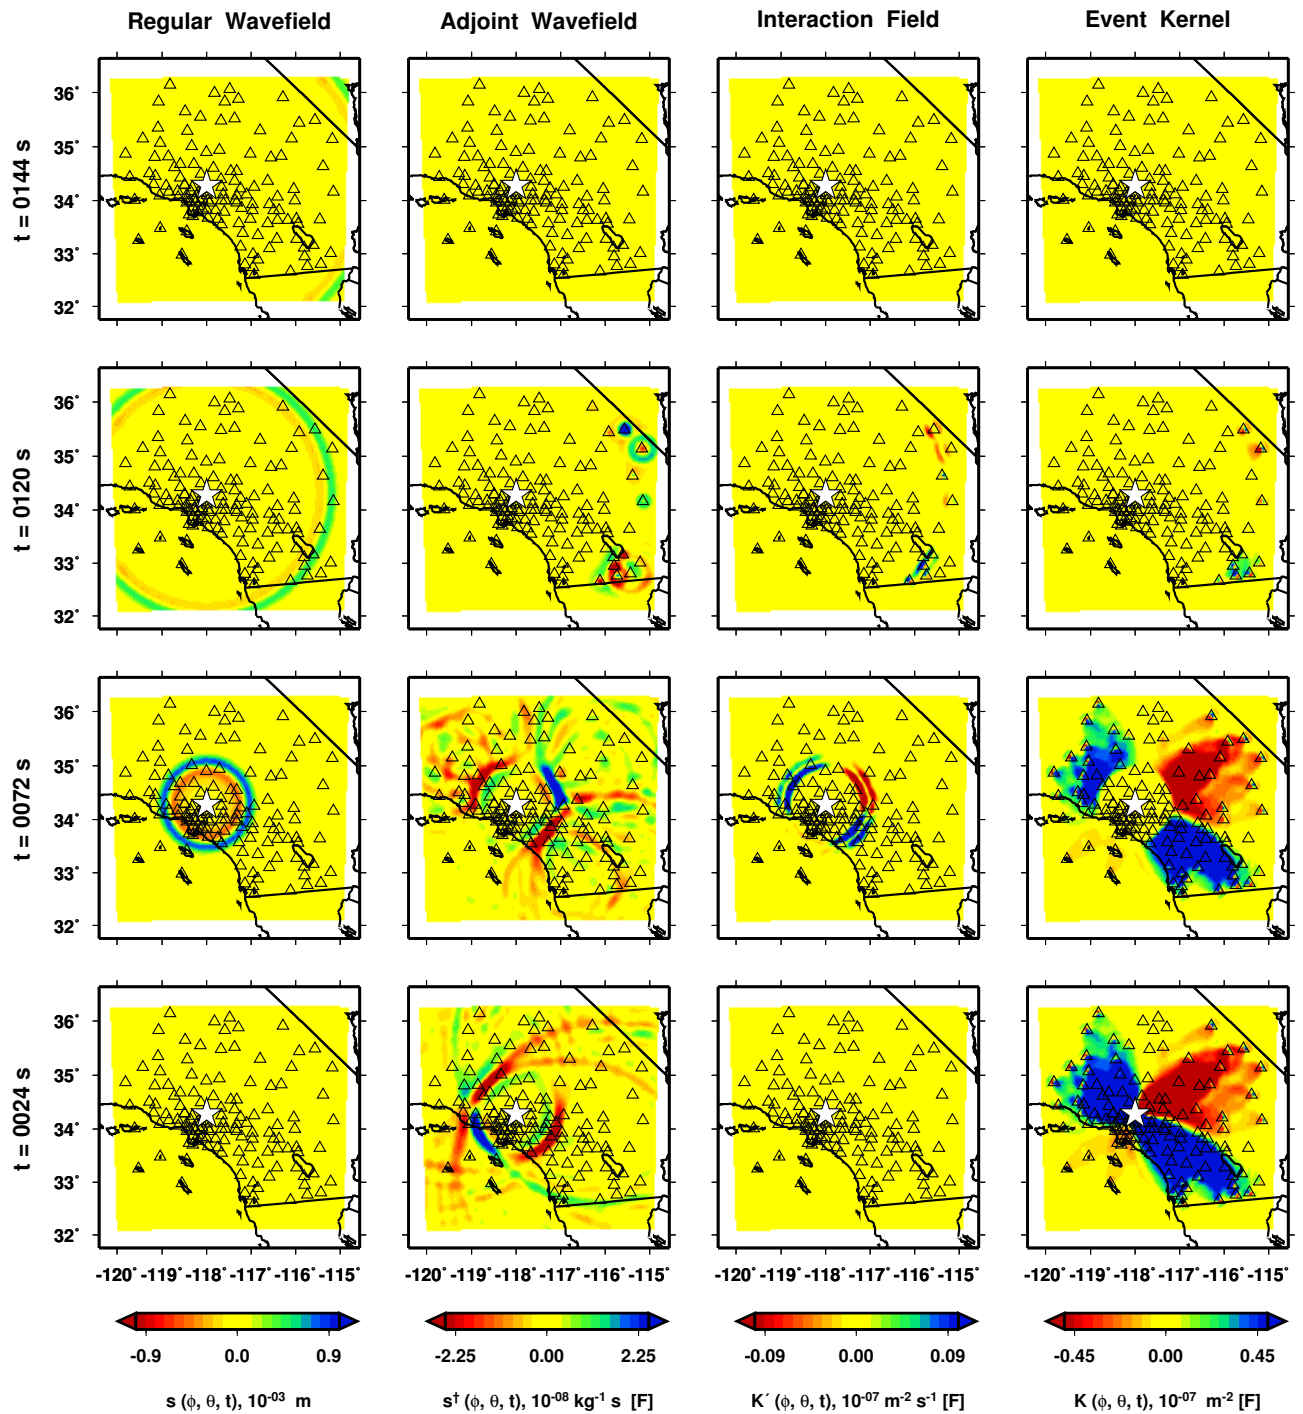
Figure 2.8:   Formation of an event kernel for multiple receivers. The phase-speed models used to generate the data and synthetics are shown in Figure 2.7. See Figure 2.5 for details. In comparison with Figure 2.5, here the event is in a different location, there are 132 receivers instead of one, and the data are generated from a checkerboard model, not a uniformly perturbed model.
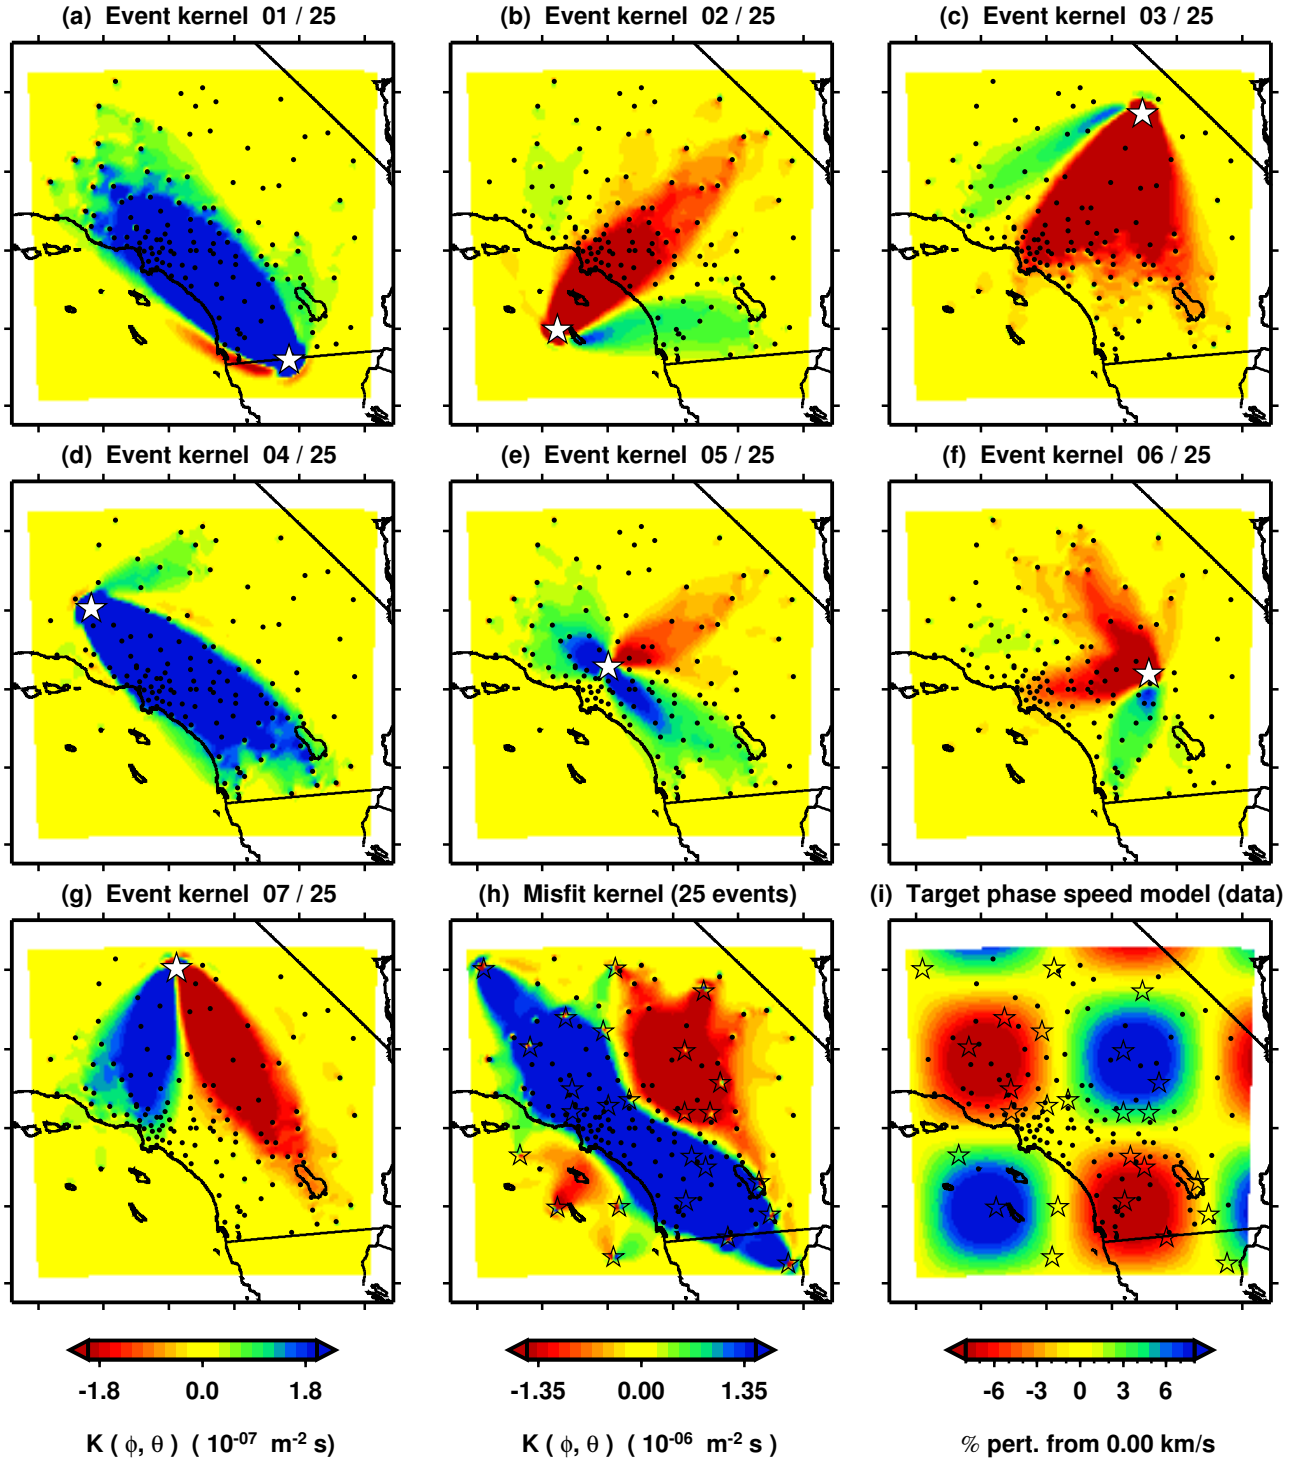
Figure 2.9: Construction of a misfit kernel. (a)–(g) Individual event kernels, each constructed via the method shown in Figure 2.8 (which shows Event 5). The color scale for each event kernel is shown beneath (g). (h) The misfit kernel is simply the sum of the 25 event kernels. (i) The source-receiver geometry and target phase-speed model. There are a total of $N = 25 \times 132 = 3300$ measurements that are used in constructing the misfit kernel. (See Section 2.5.)
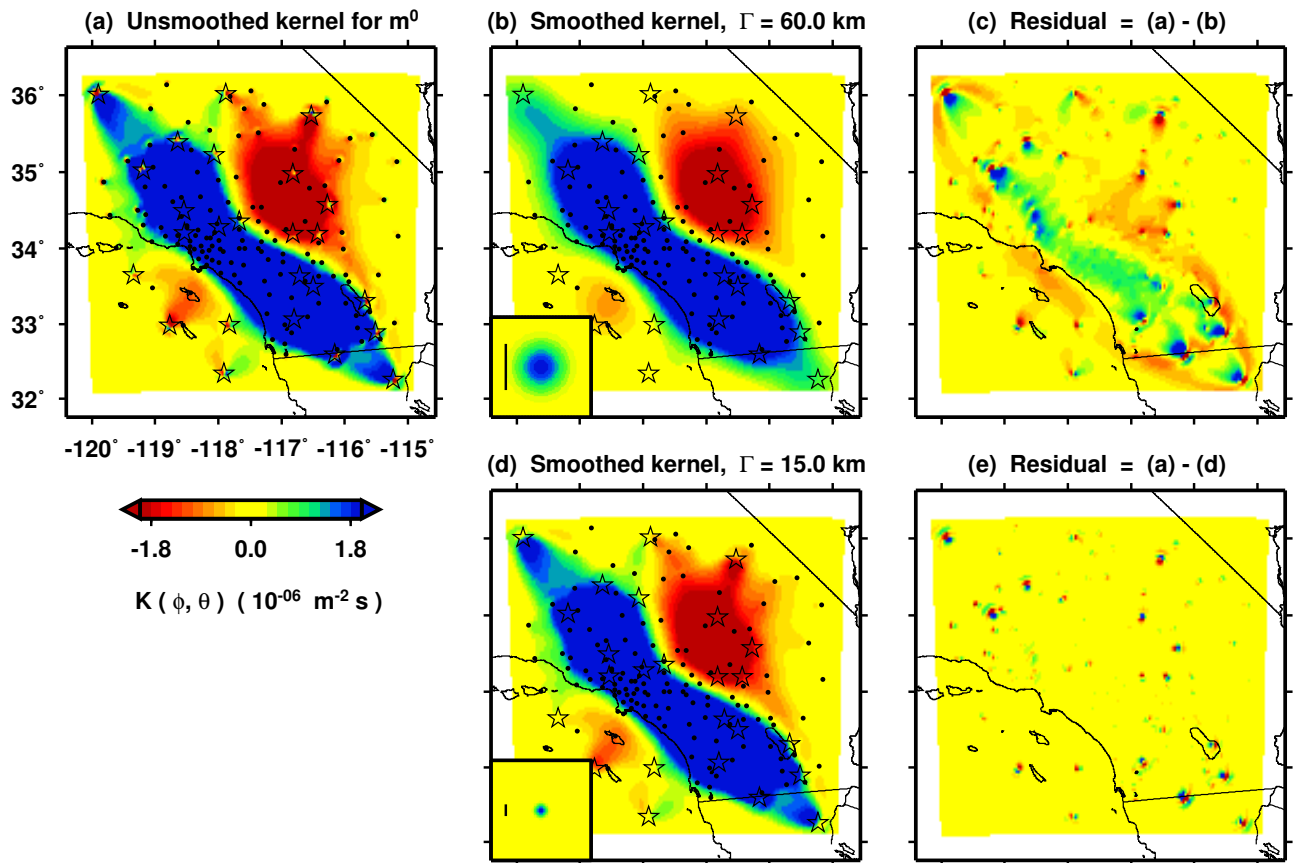
Figure 2.10:   Smoothing the misfit kernel. (a) Unsmoothed misfit kernel (Fig. 2.9h). (b)–(c) Smoothed misfit kernel, with residual, obtained via convolution of a Gaussian function (bottom left inset) with (a). The parameter $\Gamma = 60$ km controls the width of the Gaussian and, thus, the degree of smoothing; its value is plotted as a line next to the Gaussian. (d)–(e) Same as (b)–(c), but for less smoothing ($\Gamma = 15$ km). Note that the source and receiver labels are not plotted in the residual plots. (See Section 2.5.2.)

(a) Target model (data)

(b) Initial model m⁰ (synthetics)

(c) Gradient (smoothed) for m⁰

(d) Estimate test model m⁰ᵗ

(e) Test model m⁰ᵗ

(f) Gradient (smoothed) for m⁰ᵗ

(g) Cubic interpolation to get m¹

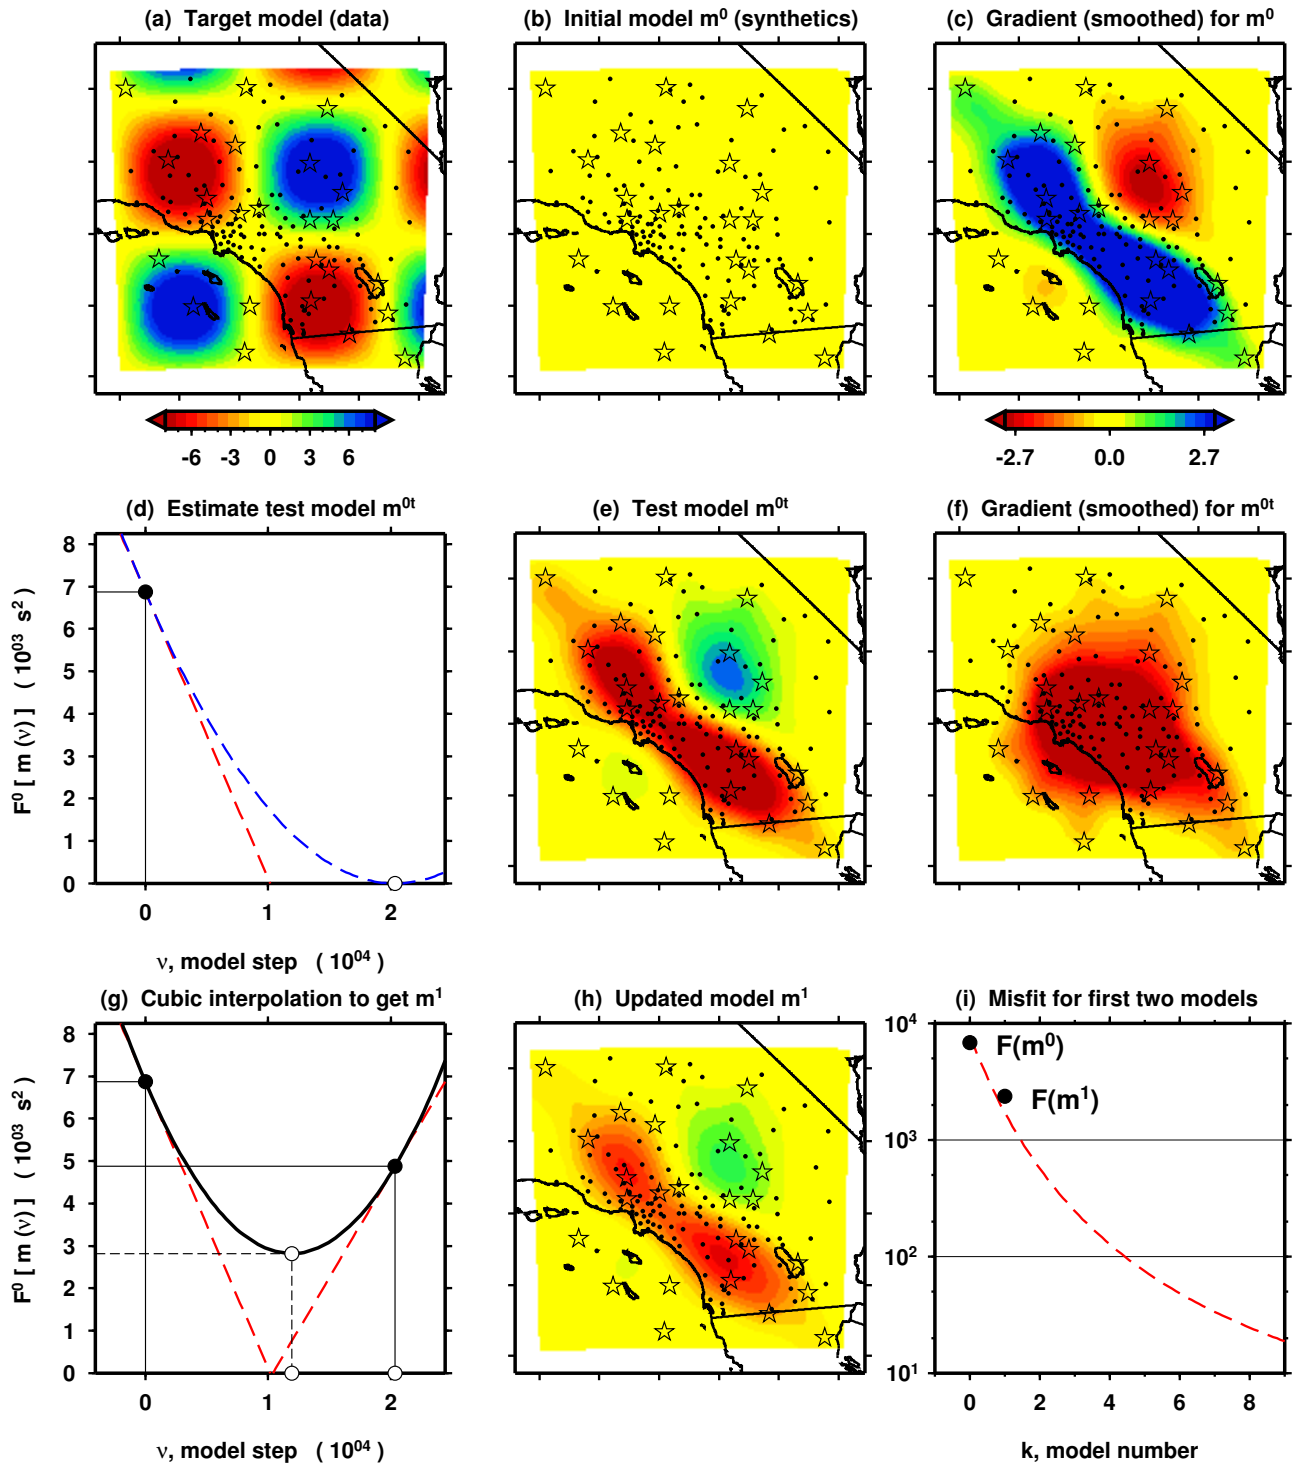**(h) Updated model m¹**

(i) Misfit for first two models

Figure 2.11:   (Caption on following page)

Figure 2.11 (*previous page*): The conjugate gradient algorithm applied to a 2D tomographic example. An extended explanation can be found in Section 2.6.2. The algorithm is repeated to obtain the models in Figure 2.12. (a) Target phase-speed model used to generate the data. (b) Phase-speed model used to generate the initial synthetics. The period of the source is $\tau \approx 20$ s, the reference wave speed is $c = 3.50$ km/s, and thus the reference wavelength is $\lambda \approx 70$ km. (c) Misfit kernel — corresponding to the gradient of the misfit function — constructed as illustrated in Figures 2.8–2.10, with smoothing parameter $\Gamma = 30$ km. This kernel highlights the regions of model (b) that need to be improved to reduce the misfit between data and synthetics. (d) Representation of the misfit of the initial model (b) and the initial gradient (c) in the conjugate gradient algorithm. The misfit is denoted by the •, and the gradient is denoted by the red dashed line. The white circle indicates the test model obtained by quadratic extrapolation of the gradient through $F = 0$. (e) Test model $\mathbf{m}_t^0$ corresponding to the white circle in (d). (f) Gradient associated with the test model in (e). (g) Cubic interpolation of two misfit values, $F(\mathbf{m}^0)$ and $F(\mathbf{m}_t^0)$, and two gradients, shown in (c) and (f). The analytical minimum provides ν^0, the distance away from $\mathbf{m}^0$ (b) in the direction of (c) that is taken to obtain the first updated model, $\mathbf{m}^1$. (h) First updated model, $\mathbf{m}^1$, corresponding to the white circles in (g) with $\nu^0 = 1.24 \times 10^4$. (i) Misfit values for the first two models. The red dashed curve is taken from Figure 2.12f.
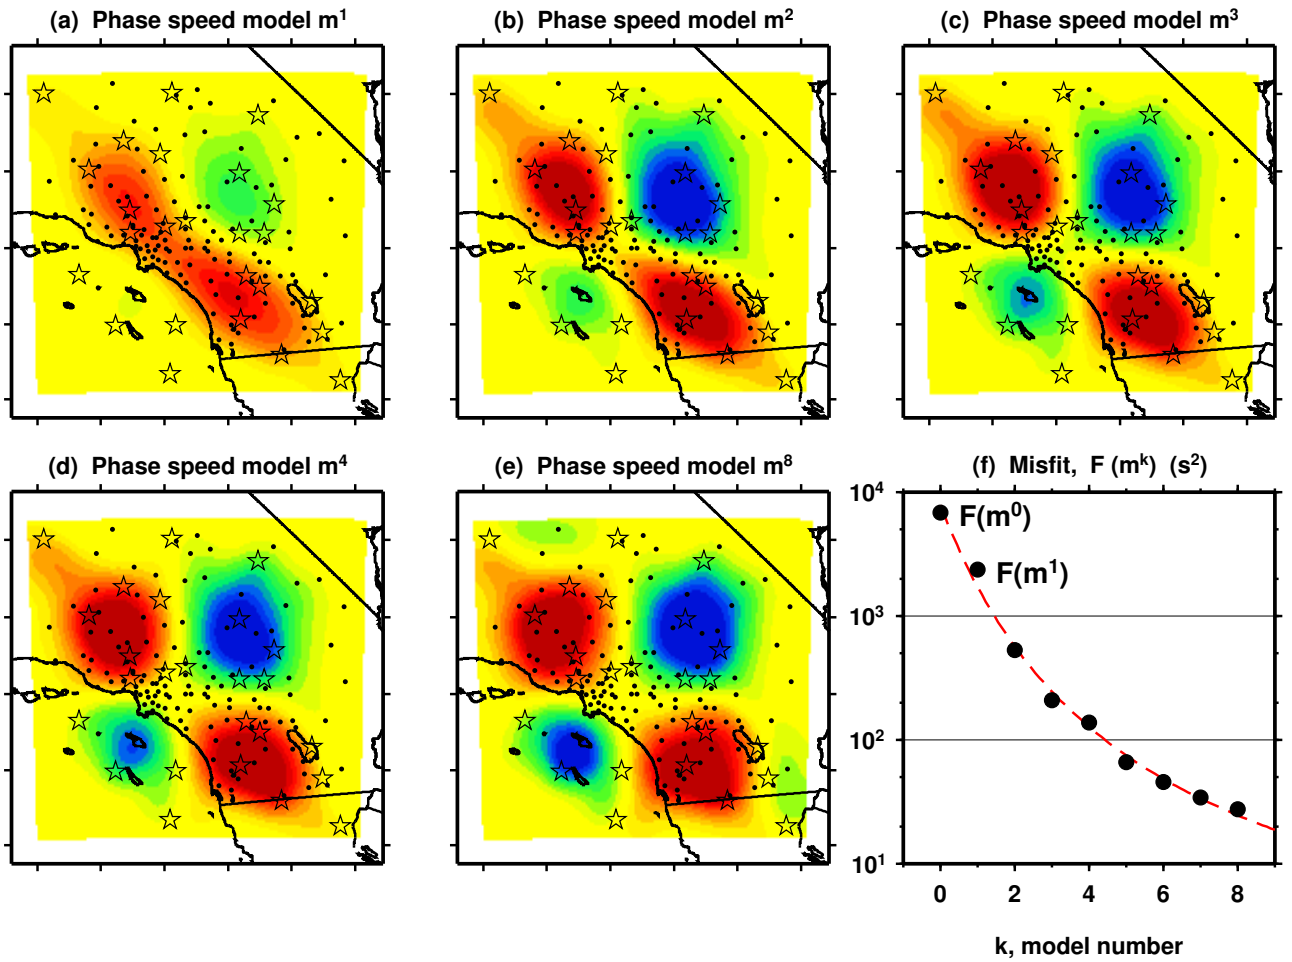
Figure 2.12: Iterative improvement of the reference phase-speed model using the conjugate gradient algorithm illustrated in Figure 2.11. An extended explanation can be found in Section 2.6.2. The first iteration in Figure 2.11 produces $\mathbf{m}^1$ (a), which becomes the current model, from which we obtain $\mathbf{m}^2$ (b), and so on. The red dashed hyperbolic curve in (f) is drawn to accentuate the reduction in misfit.
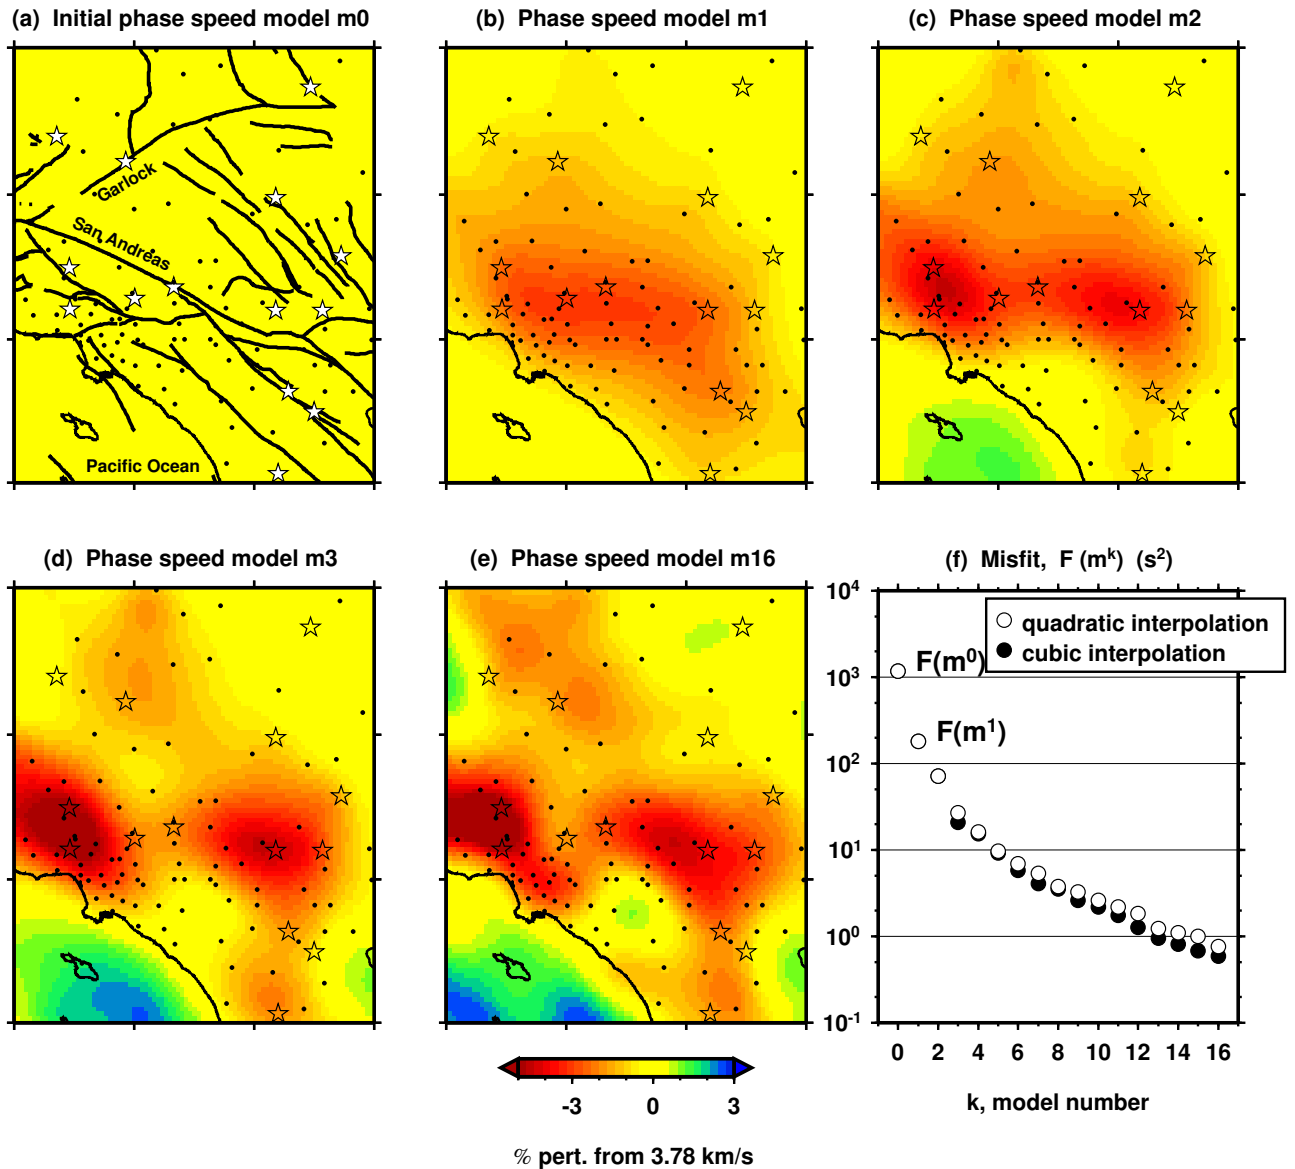
(a) Initial phase speed model m0   **(b) Phase speed model m1**   **(c) Phase speed model m2**

(d) Phase speed model m3   **(e) Phase speed model m16**   **(f) Misfit, F (m^k) (s^2)**

% **pert. from 3.78 km/s**

Figure 2.13:   Adjoint tomography recovery of a Rayleigh wave phase-speed model (Figure 2.20a). Here we show an interior portion of the southern California region with sufficient path coverage. The color scale for each model is shown below (e). (a) Initial phase-speed model $\mathbf{m}^0$. Faults of *Jennings* (1994) are drawn only for scale. (b)–(e) Iterations $\mathbf{m}^1$, $\mathbf{m}^2$, $\mathbf{m}^3$, and $\mathbf{m}^{16}$. (f) Reduction in the misfit function (2.6) using cubic interpolation (•) versus quadratic interpolation (∘) in the conjugate gradient algorithm (Section 2.6.3).
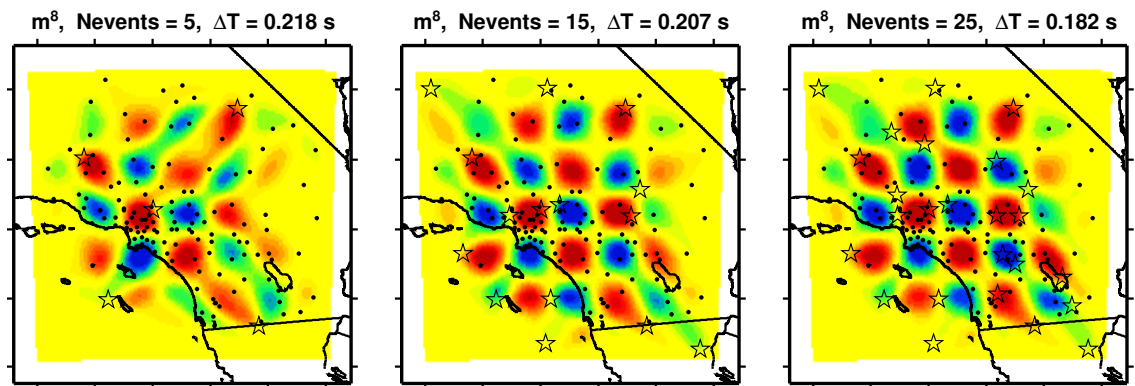
Figure 2.14:   Effect of the number of events on the recovery of the phase-speed models. Data are generated from the phase-speed model in Figure 2.15i. The average traveltime anomaly, $\overline{\Delta T}$, is computed from the misfit function value, $F(\mathbf{m}^8)$, using (2.46). As expected, $\overline{\Delta T}$ decreases as we increase the number of events.
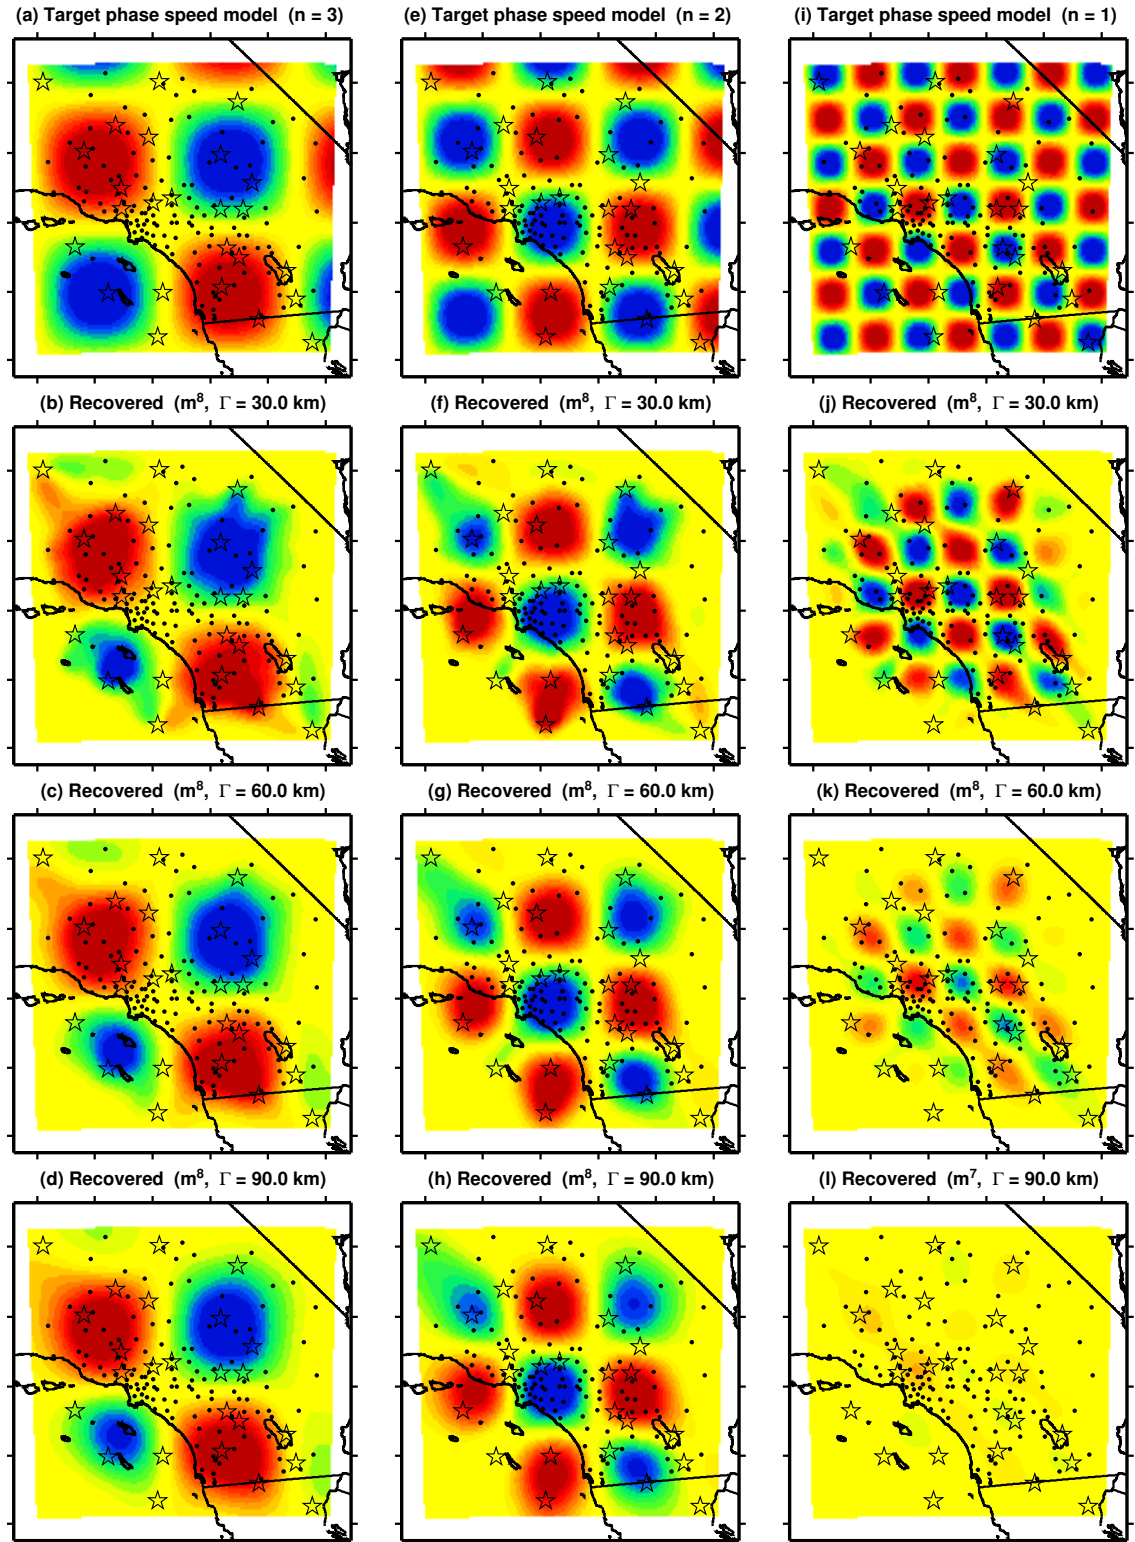
(a) Target phase speed model (n = 3) **(e) Target phase speed model (n = 2)** **(i) Target phase speed model (n = 1)**

(b) Recovered (m^8, Γ = 30.0 km) **(f) Recovered (m^8, Γ = 30.0 km)** **(j) Recovered (m^8, Γ = 30.0 km)**

(c) Recovered (m^8, Γ = 60.0 km) **(g) Recovered (m^8, Γ = 60.0 km)** **(k) Recovered (m^8, Γ = 60.0 km)**

(d) Recovered (m^8, Γ = 90.0 km) **(h) Recovered (m^8, Γ = 90.0 km)** **(l) Recovered (m^7, Γ = 90.0 km)**

Figure 2.15: Effect of the degree of smoothing and scalelength of heterogeneity on the recovery of the phase-speed models. The factor n is given by $\Lambda = n\lambda$, where Λ is the scalelength of heterogeneity and $\lambda = c\tau = 70$ km is the reference wavelength. The smoothing parameter, Γ, is constant for each row. (See Figure 2.10 and Section 2.7.)
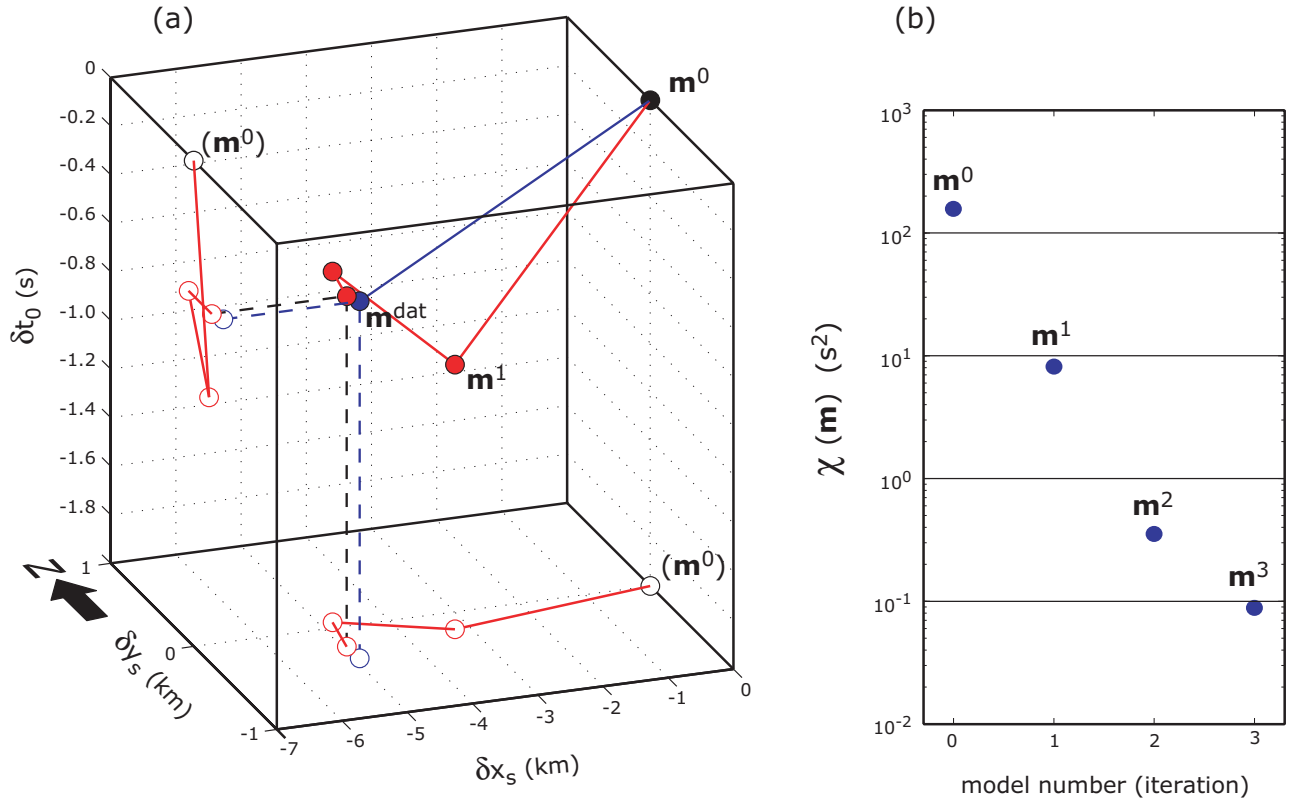
Figure 2.16:   Basic source inversion: source recovery using an unperturbed (fixed) wave-speed structure. The model vector, $\mathbf{m} = (\delta x_s, \delta y_s, \delta t_s)$, represents the source with respect to the initial model, where (x_s, y_s) is the location and t_s is the origin time. The data are generated using the target source model $\mathbf{m}^{\mathrm{dat}}$. The initial source model for the synthetics, $\mathbf{m}^0$, initiates 0.53 s late with respect to the data and is mislocated by 4.93 km at an azimuth of N85.8°E with respect to the data. The initial source parameters are randomly chosen from a mislocation "patch" with a 5 km radius and a timing error range of ± 1 s. (a) Iterative improvement of the source model toward the target source model. White circles show the projections of the model points onto horizontal and vertical planes, respectively; these are shown to aid in the perspective. (b) Reduction in the traveltime cross-correlation misfit (2.6) for the source models shown in (a). (See also Figure 2.18.)
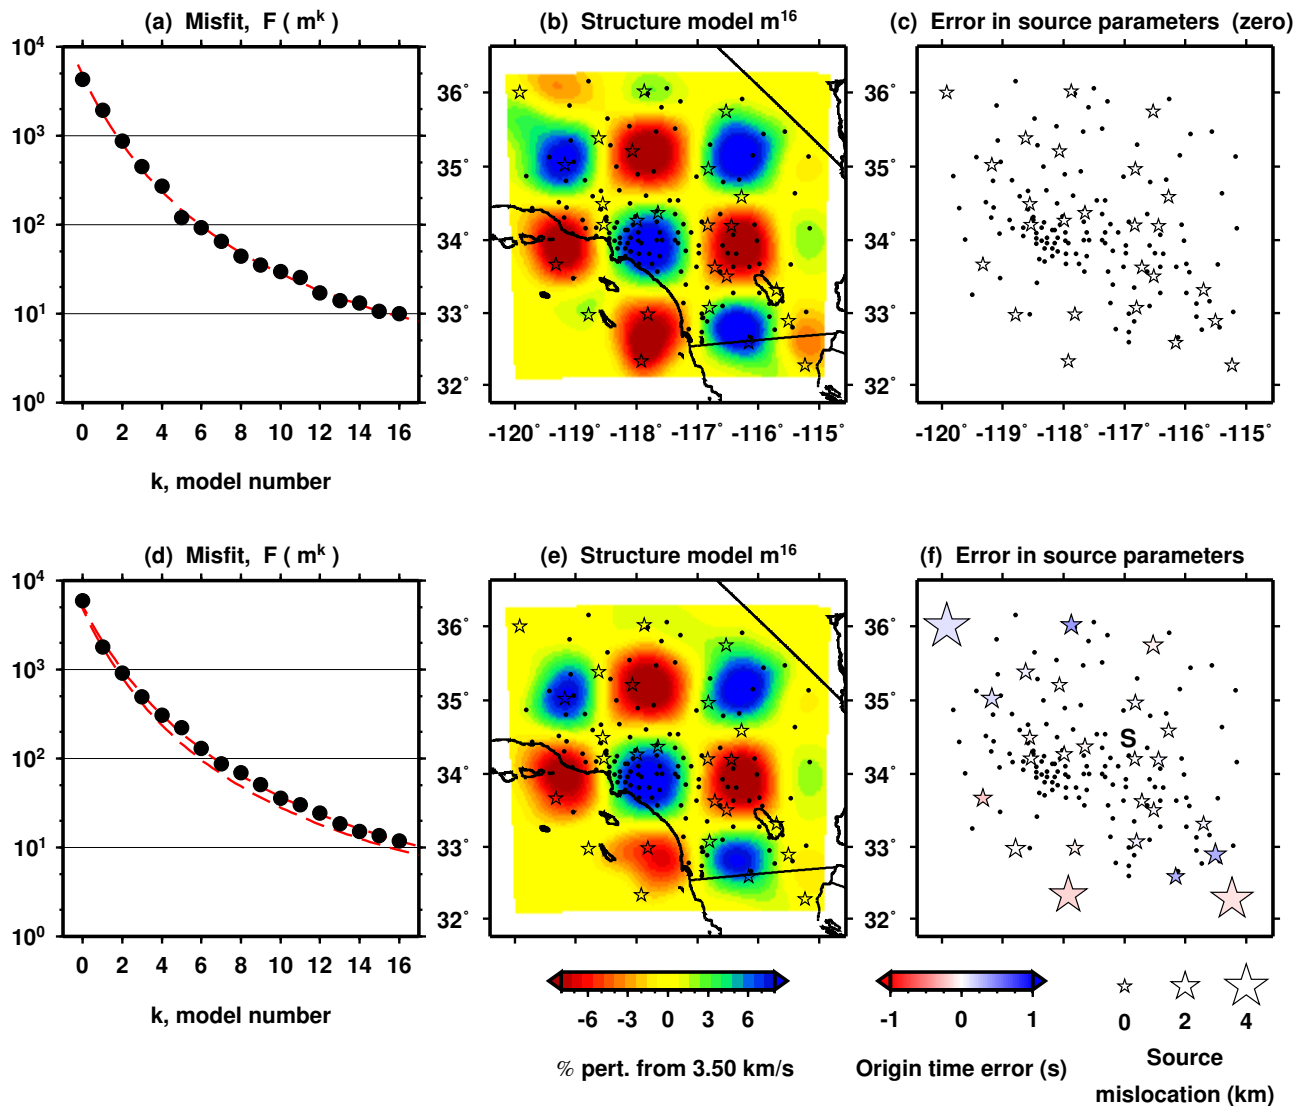
Figure 2.17: Joint inversion for source and structural parameters. The initial structural model is homogeneous. The traveltime cross-correlation misfit function values in (a) and (d) are computed from (2.6). The data are generated using Figure 2.15e. (a) Reduction in misfit for a basic structure inversion, i.e., one in which the structure of the initial model differs from that of the data, but the sources are always identical to those that generated the data. (b) Recovered model $\mathbf{m}^{16}$. Color scale is shown in (e). (c) Error in recovered source parameters. In the basic source inversion, the sources for the data and synthetics are identical and hence there is no error. Key is shown in (f). (d) Reduction in misfit for a joint inversion. The lower dashed curve is the basic structure inversion in (a). (e) Recovered model $\mathbf{m}^{16}$. Subtle differences from (b) can be seen near the edges. (f) Error in recovered source parameters. The initial error in the source parameters is shown in Figure 2.19c. Sources near the edges have the largest remaining error. The recovery of the source parameters for the event labeled S is shown in Figure 2.18.
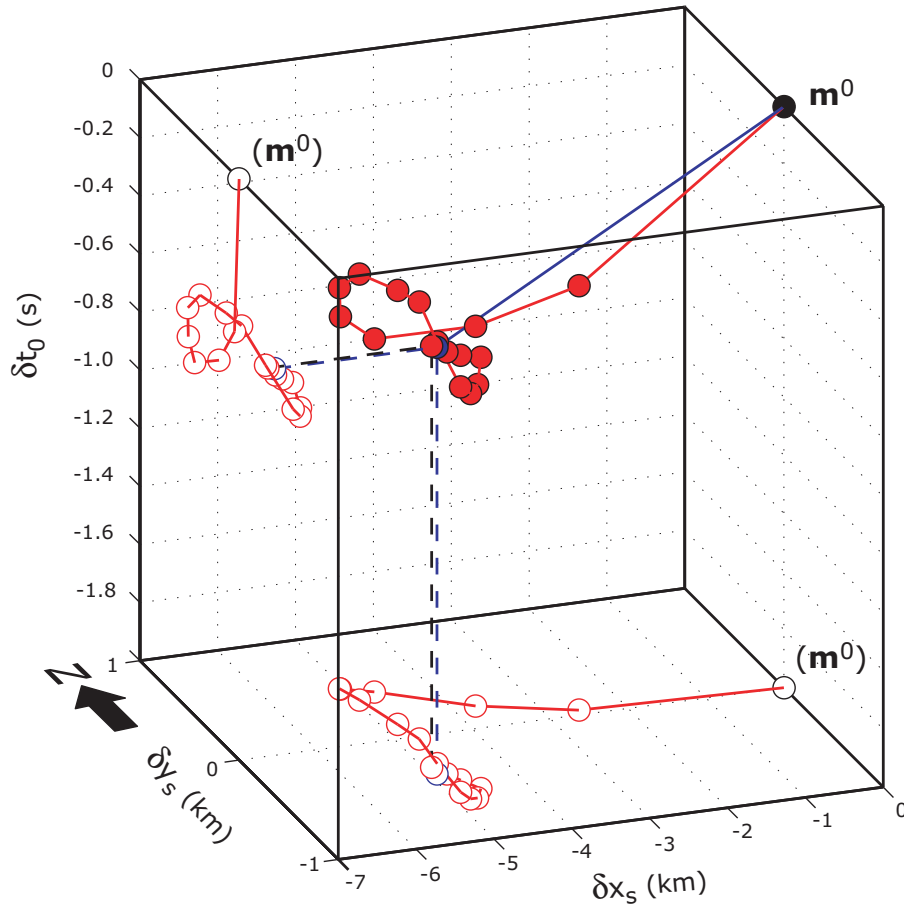
Figure 2.18: Source recovery of a particular event during the joint inversion shown in Figure 2.17d-f. The source location is denoted by the S in Figure 2.17f. By the sixteenth iteration, the source is nearly identical to the source used to generate the synthetics. The recovered structure is shown in Figure 2.17e. Compare with Figure 2.16a, which is the same source perturbation, $(4.93\,\mathrm{km}, -0.53\,\mathrm{s})$, but for a basic source inversion.
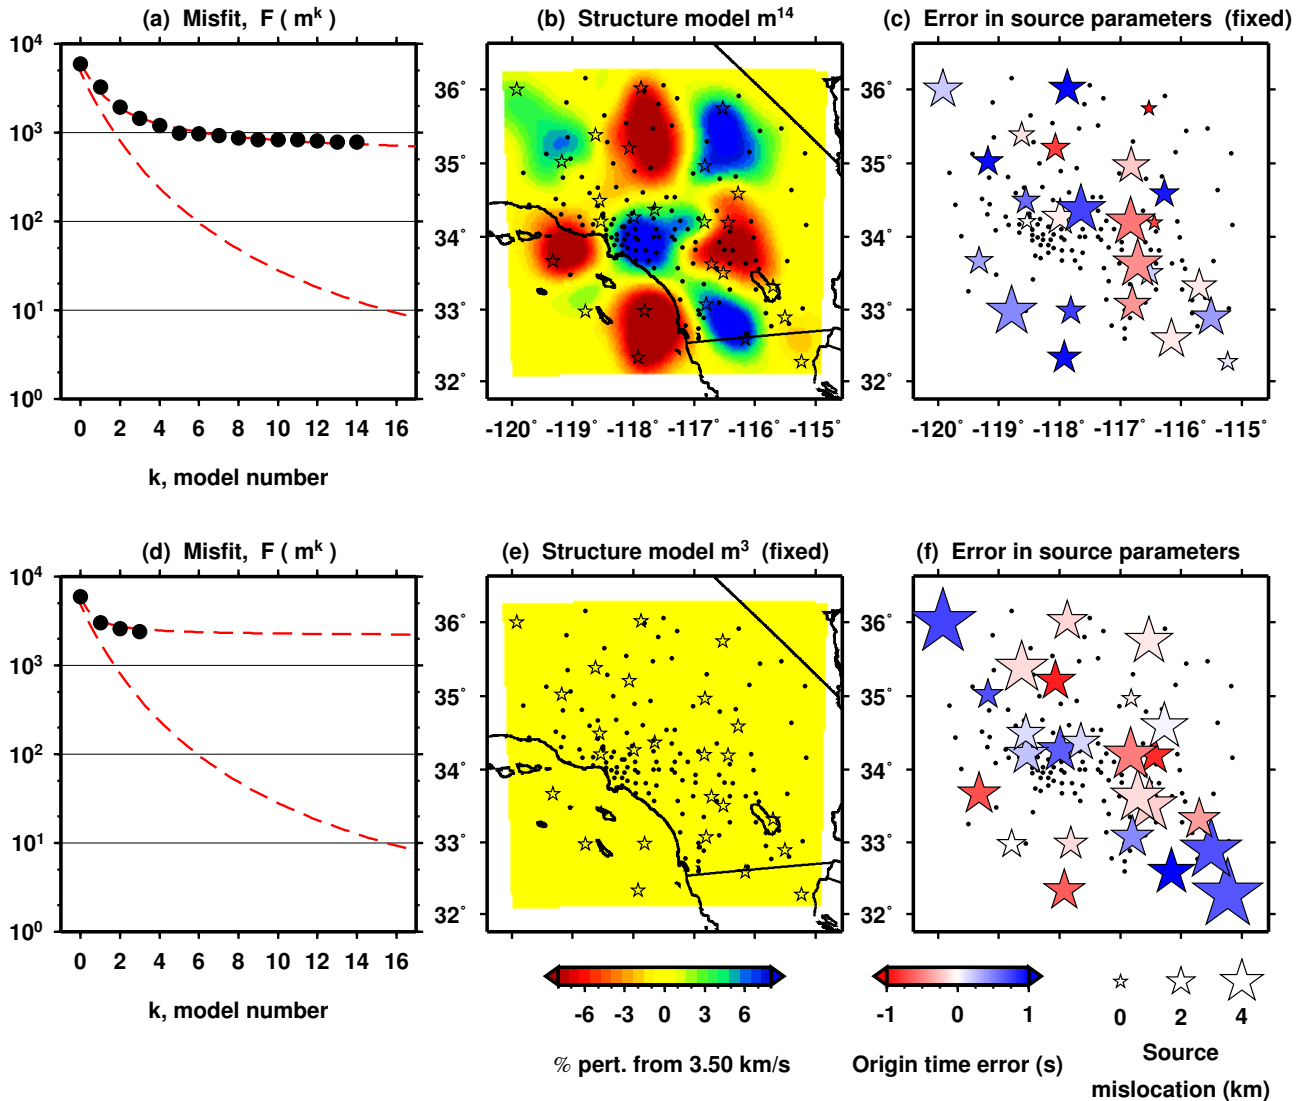
Figure 2.19:   Mapping source errors onto structure and vice versa. The initial source and structural model parameters are different from the target source and structural parameters. The traveltime cross-correlation misfit function values in (a) and (d) are computed from (2.6); the number of values is <17, because the stopping criterion for the conjugate gradient algorithm was reached. The data are generated using Figure 2.15e. (a) Reduction in misfit for a structure inversion, whereby the source errors are fixed. The lower dashed curve is the basic structure inversion in Figure 2.17a. (b) Recovered model $\mathbf{m}^{14}$. Color scale is shown in (e). Note the discrepancy with Figure 2.17b. (c) Error in source parameters used in the inversion. Key is shown in (f). (d) Reduction in misfit for a source inversion, whereby the structure errors are fixed. The lower dashed curve is the basic structure inversion in Figure 2.17a. (e) Structure model used in the inversion. (f) Error in recovered source parameters. The initial error in the source parameters is shown in (c).
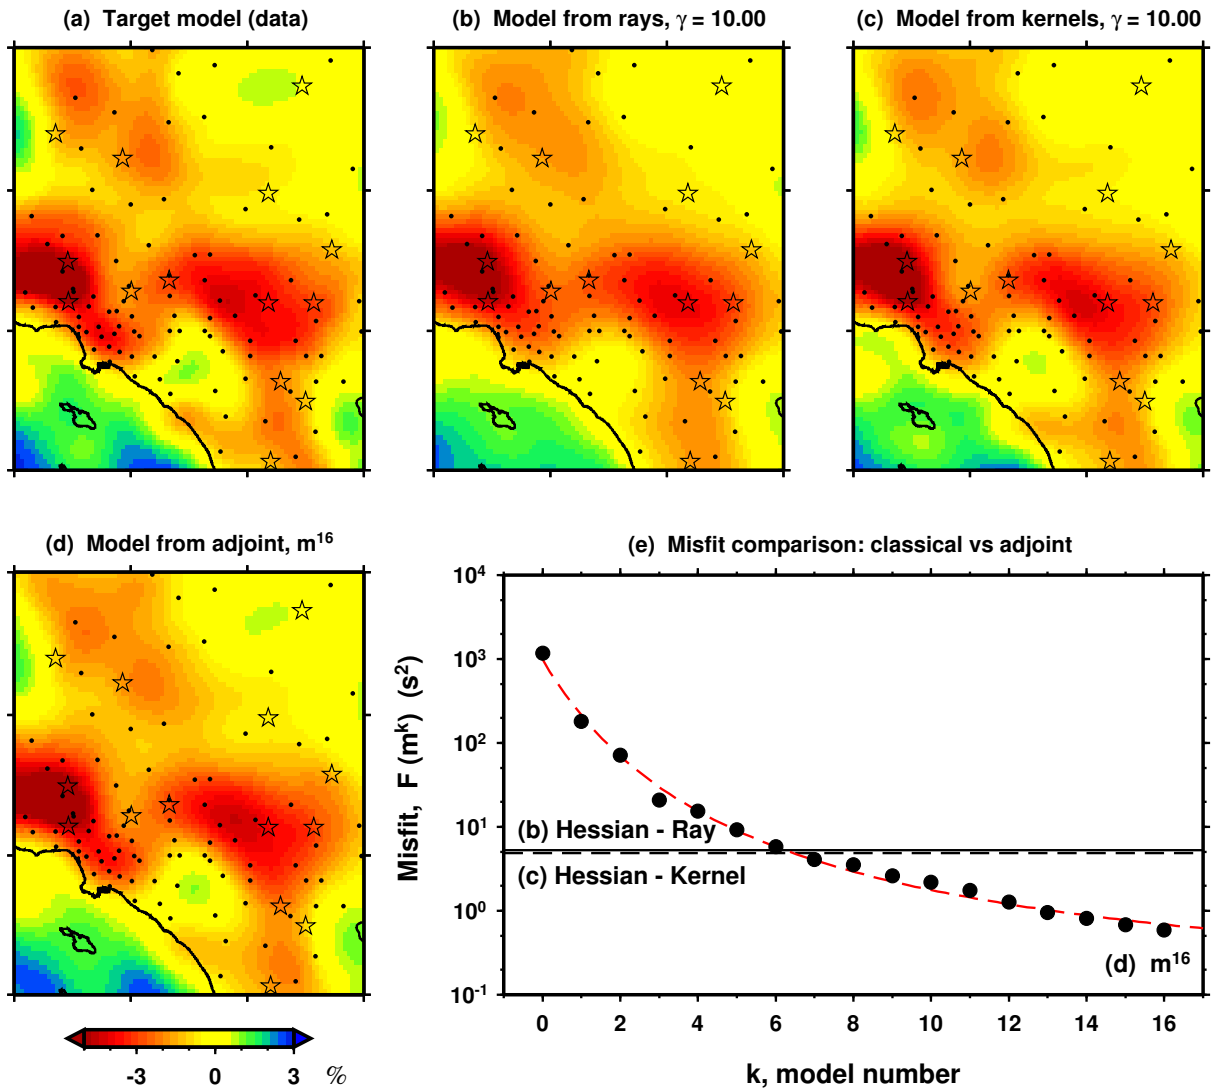
Figure 2.20:   Comparison of recovered phase-speed models for classical and adjoint to-
mography. (a) Target model used to generate the data; expanded version is shown in
Figure 2.1b. (b) Recovered model $\mathbf{m}^{\mathrm{ray}}$ using classical, ray-based inversion. The damping
parameter γ is defined in (2.27). (c) Recovered model $\mathbf{m}^{\mathrm{ker}}$ using classical, kernel-based in-
version (Figure 2.4c). (d) Recovered model $\mathbf{m}^{16}$ using the adjoint method and a conjugate
gradient algorithm (Figure 2.13e). (e) Misfit comparison for the three approaches (eq. 2.6).
The horizontal lines denote the misfit computed for the ray- and kernel-based models shown
in (b) and (c). (See Section 2.9.2 for details.)