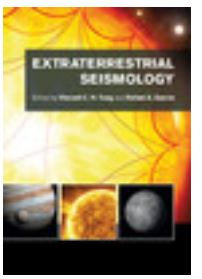


Extraterrestrial Seismology

Edited by Vincent C. H. Tong, Rafael A. García

Book DOI: http://dx.doi.org/10.1017/CBO9781107300668

Online ISBN: 9781107300668

Hardback ISBN: 9781107041721

Chapter

26 - Waveform tomography in geophysics and helioseismology pp. 365-377

Chapter DOI: http://dx.doi.org/10.1017/CBO9781107300668.029

Cambridge University Press

# 26

# Waveform tomography in geophysics and helioseismology

LAURA J. COBDEN, ANDREAS FICHTNER, AND VINCENT C. H. TONG

## 26.1 What is waveform tomography?

Seismic tomography – in which we construct images of a body's interior using seismic waves – is an inverse problem; that is, our goal is to find a model that fits a set of existing data observations. This is much less straightforward than the reverse, forward problem (i.e., generating synthetic data from an existing model) due to the fact that multiple models can fit the same data or, in other words, the solution is non-unique. Furthermore, there may be parts of the model to which the data have no sensitivity, and small errors in the data can propagate into significant errors in the model (e.g., Trampert, 1998). In order to transform data space into model space, a seismic modeling algorithm is required that can generate synthetic data from an initial model, and then update the initial model to minimize the misfit between synthetic and observed data. There is therefore always a trade-off between the computational efficiency of the modeling algorithm and the accuracy or resolution with which it can represent the real seismic structure. Which kind of modeling algorithm is employed in a given situation depends very much on the nature of the structure being imaged, the quality of the data, and the available computational resources.

Traditionally, seismic tomography has used the travel times of wave phases between a seismic source and receiver to infer the sound-speed structure along the path between them. This so-called travel-time tomography, more typically referred to as "time–distance helioseismology" when applied to the Sun, is based upon ray theory, which assumes that waves travel with an infinitely high frequency, in much the same way as light rays propagating through a medium with smoothly varying refractive index, occasionally encountering a sharp interface. Under this approximation the seismic energy propagates along infinitesimally narrow geometric "ray paths." The travel time along a given ray path is then, according to Fermat's principle, simply the line integral of the wave speed along that path:

$$T = \int_l \frac{ds}{c},$$

*Extraterrestrial Seismology*, eds. Vincent C. H. Tong and Rafael A. García. Published by Cambridge University Press.

where $l$ is the ray path, $c$ is the wave speed, $T$ is the travel time, and $ds$ is an increment along the ray path.

Obtaining models of the solar or terrestrial wave-speed structure using travel-time tomography is mathematically simple and computationally efficient. Furthermore, as long as the wavelength of the waves is very much less than their path length, the ray approximation gives a reasonable representation of the region being studied. However, the ray approximation becomes inadequate for complex structures that vary on length scales that are smaller than or comparable with the Fresnel zone of the propagating waves (e.g., Williamson and Worthington, 1993). In reality, seismic waves with a finite frequency do not travel along ray paths between source and receiver, but will sample extensive regions outside the geometric ray path due to scattering. This modifies the structure of the propagating waveform and causes travel times to deviate from their ray-theoretical values.

The idea of waveform tomography is that, during the inversion process, we aim to fit the entire signal waveform recorded on a seismogram, rather than just the travel times of certain peaks within that waveform. Forward modeling of seismic waveforms can proceed by solving the full wave equation, rather than relying on simple approximations such as ray theory to simulate wave propagation. In this case, complex phenomena such as diffractions are accounted for implicitly, rather than imposing explicit conditions to divide the ray into transmitted and reflected components. Modeling entire seismic waveforms requires vastly more computational resources than modeling only the travel times, but it allows us to exploit additional information contained in the amplitudes and/or phases of incoming waves. This extra information may, for example, provide constraints on the density and attenuation structure in a region, which cannot be extracted from travel-time analysis alone. Furthermore, by taking into account the low-frequency components of the seismic signal, the spatial resolution is only limited by the source and receiver distributions. It is thus possible to image structures at sub-wavelength scales, significantly better than the limit of the first Fresnel zone in travel-time tomography (e.g., Pratt *et al*., 1996).

An additional advantage of waveform tomography relates to data processing. In travel-time tomography, travel times are not a directly recorded parameter, but must be picked somehow from the seismic record. This picking can be highly subjective, depending on human judgement, and it may be difficult to extract the timing of a specific peak from triplicated waveforms (i.e., where multiple waveforms arrive at the same receiver within a short and possibly overlapping time interval). Since waveform tomography makes use of the entire wave train recorded at a receiver, these subjective data-picking errors are not introduced, although errors due to noise will still affect the accuracy of the inversion.

With the increasing availability of cheap random access memory and parallelized computer systems (supercomputers) in the past two decades, the application of waveform tomography has become increasingly feasible. In situations where the resolution of small-scale 3D structures is paramount, and sufficiently high quality data are available, waveform tomography should be the imaging method of choice.

## 26.2 Seismic wave propagation in complex media

The first step in performing waveform tomography is to be able to simulate, accurately, wave propagation through the medium of interest (i.e., the solar or Earth's interior), via numerical solution of the wave equation. For the Earth, the goal is to solve the seismic wave equation, the theoretical derivations of which can be found, for instance, in Aki and Richards (2002). On the Sun, we are ultimately required to solve a highly complex magnetoacoustic wave equation (discussed in Section 26.5).

The earliest fully numerical models of seismic wave propagation through heterogeneous Earth models were mostly based on finite difference methods (e.g., Alterman and Karal, 1968). Finite differencing is a popular method for solving differential equations computationally. It requires construction of a grid. At each node of the grid, values of continuous functions are calculated approximately, by replacing the differentials of the function with discrete difference equations, e.g., $\partial u/\partial t \approx [u(t + \Delta t) - u(t)]/\Delta t$. The advantage of this method is that it can produce accurate solutions and allows modeling of complex media. Finite difference (FD) methods for wave propagation rapidly gained popularity in the 1980s, thanks to both increasing computational resources and the implementation of the staggered grid scheme (Virieux, 1984) originally introduced in fluid dynamics (Harlow and Welch, 1965). In the staggered grid, material properties and dynamic field variables are defined on different sets of grid points, which leads to a substantial reduction of numerical dispersion. Owing to their robustness and simplicity, FD methods continue to be widely used, especially in industrial applications. A comprehensive and modern treatment of FD methods can be found in Moczo *et al*. (2002).

A major drawback of FD methods is the inability to adapt the numerical grid to complex subsurface structure and topography. In exploration geophysics, where most surveying is done in small ($<10^2$ km) regions of relatively flat topography, this is usually not a major concern, but FD methods become increasingly inappropriate at longer length scales on the Earth, when wave propagation effects related to surface topography and undulating seismic discontinuities cannot be ignored. This has led to the development of finite element (FE) methods for the solution of the wave equation. Among the family of FE methods, the spectral element method (SEM) is currently considered to provide the optimal balance between flexibility, accuracy, and computational requirements (e.g., Komatitsch and Vilotte, 1998). In the SEM, the model volume is divided into a number of non-overlapping, hexahedral elements. Within each element, the wavefield and material properties are represented by continuous functions (basis functions), typically Lagrange polynomials of degree 4 to 8. The shape of each element is also defined by low-degree polynomials, and continuity of displacement along the boundaries is enforced. Thus, complex structures and topography are better accounted for than with the discrete FD procedure, and free-surface boundary conditions are implicitly included, unlike in FD modeling where one must explicitly define whether a boundary is reflecting or absorbing wave energy. In contrast to other, more standard, FE methods, the SEM leads to a diagonal mass matrix, which significantly reduces computational costs, since the inversion of a large matrix is not required.

### 26.3 Solution of the full waveform inverse problem

Finite difference methods and the SEM allow us to generate high-fidelity synthetic seismograms for a given Earth or solar model. Since our goal in tomography is to derive a model that best represents a given dataset, the second component to performing waveform tomography is an algorithm that will update an initial-guess model, in order to minimize the misfit between synthetic and observed seismograms. This requires that we (1) define a misfit functional and (2) choose an optimization scheme to locate the global minimum.

Designing suitable misfit functionals remains an active area of research, and good solutions tend to be strongly application-dependent. Many applications are based on the $L_2$ norm: Find the model $m$ that minimizes the $L_2$ distance $\chi$ between the observed seismograms $u_0^r(t)$ and numerically computed (synthetic) seismograms $u^r(t;m)$:

$$\chi(m) = \frac{1}{2}\sum_r \int [u^r(t;m) - u_0^r(t)]^2 dt, \tag{26.1}$$

where $r$ denotes the index of a seismic receiver at position $x^r$. In Equation (26.1), additional data-specific error weights may be introduced. However, the first applications of waveform tomography based on this algorithm (e.g., Bamberger *et al.*, 1982; Gauthier *et al.*, 1986) revealed two problems: (1) The dependence of $\chi$ on the model $m$ can be highly nonlinear, such that gradient-based optimization schemes tend to lock into local minima; and (2) the $L_2$ waveform difference is dominated by large-amplitude waves. Valuable information contained, for instance, in small-amplitude body waves is almost entirely disregarded. Thus, careful choice of a suitable misfit functional is essential for the successful solution of a waveform tomographic inverse problem.

Alternative suggestions for misfit functionals in terrestrial waveform tomography include variants of the $L_1$ norm (e.g., Brossier *et al.*, 2009), as well as instantaneous phase and amplitude misfits (e.g., Bozdag *et al.*, 2011). Fichtner *et al.* (2008, 2013) proposed the use of time- and frequency-dependent phase misfits. In a first step, observed and synthetic seismograms are transformed to the time–frequency domain using, for instance, a moving-window Fourier transform. This results in complex-valued time–frequency representations of observations and synthetics. The $L_2$ distance between the observed phase $\phi_0^r(t,\omega)$ and the synthetic phase $\phi^r(t,\omega;m)$ is then used to define the misfit:

$$\chi(m) = \frac{1}{2}\sum_r \int\int [\phi^r(t,\omega;m) - \phi_0^r(t,\omega)]^2 dt d\omega, \tag{26.2}$$

where again additional weights may be incorporated. In Figure 26.1, we illustrate the measurement of the time–frequency phase misfit for the case of a north–south (N–S) component seismogram recorded at station ARGB in eastern Turkey, at a distance of 9.52° from the earthquake epicentre. Those parts of the recordings where observations and synthetics are not approximately in phase must be eliminated in order to produce meaningful measurements of phase difference. This particular misfit functional does not

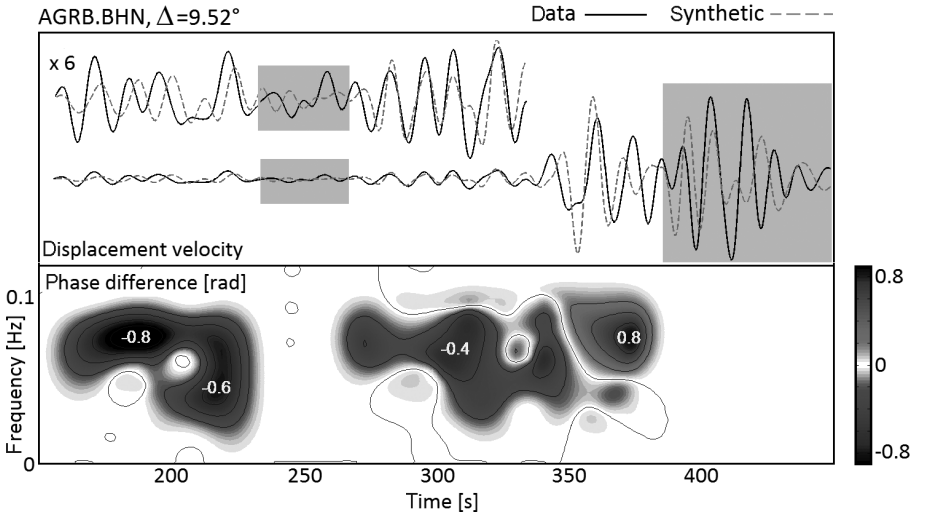


Figure 26.1 Measurement of a time–frequency phase misfit. (Top) Observed (solid) and synthetic (dashed) seismograms for station AGRB in western Turkey (distance to earthquake $\Delta = 9.52^\circ$; see Figure 26.2 for source–receiver geometry). Oscillations prior to 340 s correspond to low-amplitude body waves amplified by a factor of 6 for better visibility. Large-amplitude surface waves appear from 340 to 450 s. Gray shading marks intervals that are excluded from the misfit calculation because waveform differences are too large for phase difference measurements. (Bottom) Phase difference, $\varphi - \varphi_0$, between observations and synthetics as a function of time and frequency. The phase differences between the small-amplitude body waves are as significant as those between large-amplitude surface waves. A black and white version of this figure will appear in some formats. For the color version please refer to the plate section.

require the presence of well-defined body or surface waves; it is comparatively insensitive to the choice of initial model, and it is not biased by the amplitudes of the waveforms.

Owing to the large number of free parameters in a waveform tomographic inversion (thousands to millions), the optimization problem is usually solved using gradient-based techniques, including conjugate gradients or Newton-like methods. Gradient-based optimization algorithms rely on an initial model that is sufficiently close to the global minimum. Since good initial models are often not available, multi-scale strategies (e.g., Virieux and Operto, 2009) have been developed in order to prevent locking into local minima. Multi-scale strategies are based on the observation that misfit functionals are better behaved, i.e., closer to being convex, when low-frequency waves are considered. The iterative optimization therefore starts with the lowest available frequencies, which are used to improve a potentially crude model. After several iterations at low frequencies, the refined model is used as the initial model for the next few iterations at higher frequencies. This cycle is repeated until the model explains the highest-frequency waveforms.

Gradient-based optimizations can be computed with two different approaches: adjoint methods or scattering-integral methods. Adjoint methods (e.g., Fichtner *et al.*, 2006)

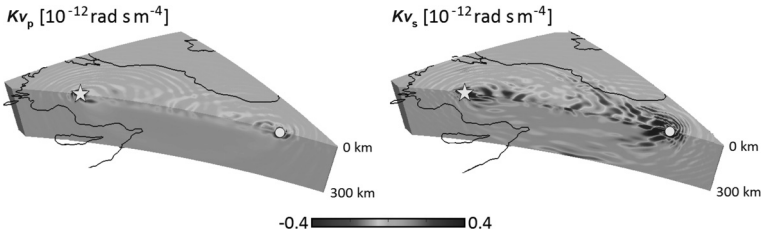


Figure 26.2 Sensitivity kernels $K_{\upsilon p}$ and $K_{\upsilon s}$ for the phase difference measurement in Figure 26.1. The positions of the earthquake source and receiver (station ARGB) are indicated by a star and filled circle respectively. Regions that are brighter or darker than average are those where perturbations in $v_p$ (for the case of $K_{\upsilon p}$) and $v_s$ (for the case of $K_{\upsilon s}$) affect the phase difference measurement. Thus, the kernels indicate where the current model should be modified to reduce the phase misfit. A black and white version of this figure will appear in some formats. For the color version please refer to the plate section.

are closely related to time-reversal techniques in seismic exploration, and to acoustic holography in helioseismology. The difference between observations and synthetics recorded at an array of receivers is time-reversed and re-injected at the receiver positions as sources in a time-reversed or *adjoint* simulation. The resulting wavefield focuses where the synthetic model differs from the true model, and the quality of the focusing depends on the source–receiver coverage. The time correlation of the adjoint wavefield and the regular wavefield yields sensitivity kernels, $K(x)$, which relate infinitesimal variations in the misfit functional, $\delta\chi$, to infinitesimal variations of the model, $\delta m$. The sensitivity kernels for the phase difference measurements shown in Figure 26.1 with respect P-velocity and S-velocity ($K_{\upsilon p}$ and $K_{\upsilon s}$) are shown in Figure 26.2.

The computation of sensitivity kernels via adjoint methods requires one forward and one time-reversed adjoint wavefield simulation per seismic source, irrespective of the number of receivers. It follows that adjoint methods are particularly efficient when the number of sources is small, and the number of receivers is large. Scattering-integral methods (e.g., Chen *et al*., 2007) are most efficient for exactly the opposite scenario: many sources and few receivers. Making use of source–receiver reciprocity (Aki and Richards, 2002), scattering-integral methods compute sensitivity kernels using receiver-side Green functions i.e., wavefields excited by sources located at receiver positions.

## 26.4 Application of waveform tomography in geophysics

The first practical applications of waveform tomography based on solving the wave equation were in exploration geophysics, using controlled (i.e., man-made) seismic sources (e.g., Zhou *et al*., 1995). Tests on borehole seismic data demonstrated that for surveys spanning a few tens of meters, waveform tomography could image structures as small as 2 m, compared with travel-time tomography, which could only resolve down to about 10 m (e.g., Pratt *et al*., 1996). Following this success, waveform tomography was then applied to wide-angle seismic data (survey dimensions: tens to hundreds of kilometers) to image

deep crustal structures (e.g., Pratt *et al*., 1996; Shipp and Singh, 2002). Again, it was shown that waveform tomography could resolve structures up to an order of magnitude smaller than those imaged with travel-time tomography. More recently, it has been possible to apply waveform inversion in non-commercial settings, such as the imaging of buried impact craters (Morgan *et al*., 2011), and on continental length scales ($10^4$ km) (Fichtner *et al*., 2009; Tape *et al*., 2010), with the latter relying on natural sources (i.e., earthquakes). Advances towards global models of the Earth's mantle from SEM-based waveform inversions are also being made (e.g., French *et al*., 2013).

Although waveform tomography is now being used on a wide range of length scales on the Earth, each application is characterized by its own set of challenges and limitations. For example, a key feature in exploration seismics is the use of dense arrays involving large numbers of seismic sources, easily exceeding several thousand. Specific work-arounds are needed to ensure computational efficiency in the wavefield modeling of such large numbers of sources. On the other hand, while fewer sources are typically present in regional or global seismology, a major issue is the heterogeneous distribution of both the earthquake sources (dependent on plate tectonics) and seismic receivers (constrained by financial, geographical, and political factors), leading to large null spaces in the Earth model. Additionally, since the sources are earthquakes, which occur many kilometers below the surface, large uncertainties exist on the source–time function of each event, as well as their spatial coordinates.

The strong scale and source dependence of waveform tomography largely prevents the development of blackbox-like algorithms for waveform tomography that produce excellent Earth models without human interaction. We summarize below key strategies that have been or are being applied on different length scales in the Earth. In Figure 26.3 we show examples of waveform tomography on different length scales.

### *26.4.1 Exploration seismics*

Exploration seismics is mostly oriented towards the discovery and characterization of natural resources, and seismic surveys typically operate on length scales from a few hundred meters to a few tens of kilometers. The biggest problem to overcome in waveform inversion is the need to model hundreds to thousands of sources.

Where significant numbers of sources are present, waveform modeling with FD or FE methods is best handled in the frequency domain (Marfurt, 1984). In particular, only a limited number of frequencies may need to be included to retrieve a satisfactory inversion of the structure (e.g., Pratt *et al*., 1996), and each frequency can be represented by a small discrete system of equations. Solving the wave-propagation problem for a new source then reduces to a simple vector-matrix multiplication, which is relatively fast to compute.

Wave propagation in exploration seismics is usually modeled with an acoustic wave equation, as opposed to the more correct elastic wave equation. This means that only the early-arriving compressional waves can be exploited in the tomography; shear waves,

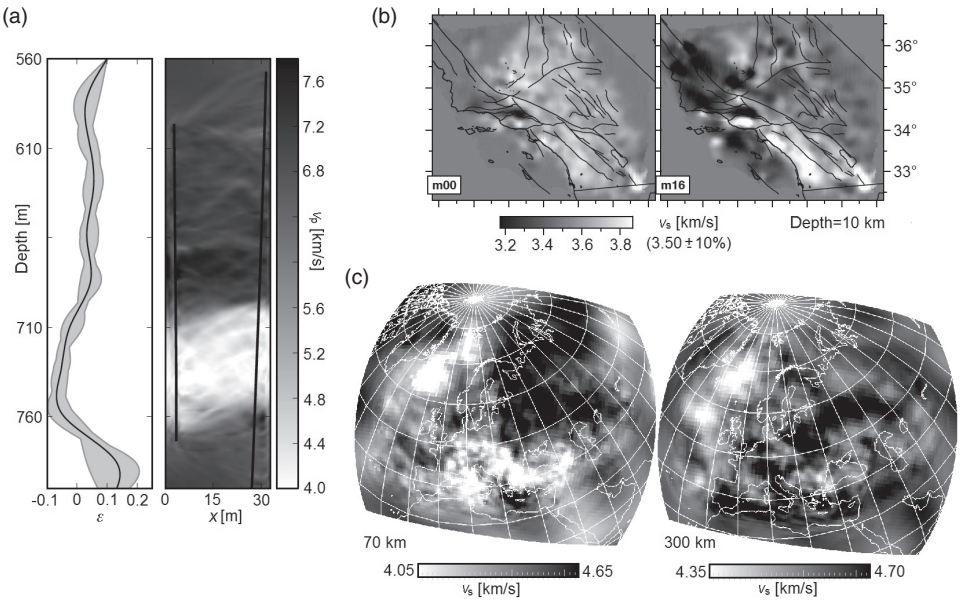


Figure 26.3 Examples of waveform tomographies on various scales in the Earth. (a) Crosshole seismic tomography. The 1D profile of elliptical anisotropy $\varepsilon$ (left panel) was found using simulated annealing to avoid trapping in local minima. The P-velocity model (right panel) was obtained subsequently via 2D frequency-domain waveform tomography. Boreholes are indicated by black lines. (Modified from Afanasiev *et al*., 2014.) (b) Waveform tomography of the Southern California crust (Tape *et al*., 2010) obtained with 3D spectral element simulations in the time domain. Shown are the initial model (left) and the final model (right) obtained after 15 conjugate-gradient iterations. (c) Waveform tomography of Europe and Western Asia, also based on spectral-element time-domain modeling. (Modified after Fichtner *et al*., 2013.)

which do not exist in acoustic media, must be disregarded in the data analysis. Furthermore, simulations are usually limited to 2D or 2.5D, for which wave amplitudes are not correctly modeled. In response to these drawbacks, exploration applications are currently moving towards 3D simulations (e.g., Warner *et al*., 2013), and simulations based on the elastic wave equation (e.g., Guasch *et al*., 2012). These are performed in the time domain, because the explicit solution of the large frequency-domain system of equations is currently not feasible for the 3D elastic case.

On a kilometer scale, the Earth is characterized by strong heterogeneities, which complicate the construction of an initial model sufficiently close to the global minimum of the misfit functional. Today, initial models are mostly constructed using classical travel-time tomography, which constrains the long-wavelength properties of the subsurface. Alternatively, Monte-Carlo-type searches may be used to estimate the broad patterns of an Earth model described by a small number of parameters. This approach was taken, for instance, by Afanasiev *et al*. (2014) (Figure 26.3a).

### *26.4.2 Regional and global seismology*

Finding suitable initial models is a lesser problem on regional ($10^3$–$10^4$ km) to global scales, where radially symmetric reference models (e.g., Dziewonski and Anderson, 1981) explain the travel times of the major seismic wave arrivals to within a few seconds. The challenge lies in the inherent three-dimensionality of the problem and the need to account fully for elastic effects, anisotropy, and visco-elastic dissipation. Consequently, the computational costs of solving the forward problem are orders of magnitude larger than in 2D acoustic simulations that can be justified in certain exploration-scale applications. The seismic sources for waveform tomography on length scales exceeding several kilometers are mostly naturally occurring earthquakes, although the use of ambient seismic noise is an area of active research (e.g., Tromp *et al*., 2010). Owing to the enormous computational requirements, regional- to global-scale full waveform tomographies have only become feasible in the past few years. Tape *et al*. (2010) presented the first full-waveform tomographic model of the Southern California crust, revealing a close correlation between seismic velocity structure and major fault systems (Figure 26.3b). On continental scales, full waveform tomography has established clear relations between seismic anisotropy and mantle flow. (Fichtner *et al*., 2009). An example for Europe and Western Asia is presented in Figure 26.3c.

The highly irregular distribution of both earthquakes and seismic stations over the globe generally results in a heterogeneous resolution. This differs from exploration applications, where more regular source–receiver geometries are available. The quantification of resolution is thus a challenging aspect of regional and global waveform tomography which has so far received little attention. Most resolution analysis is limited to synthetic inversions of chequerboard-like input patterns. Information about the curvature of the misfit functional near its minimum, obtained via second-order adjoints, provides resolution proxies in the vicinity of the optimal model (e.g., Fichtner and Trampert, 2011). However, practicable methods for global uncertainty analysis in large-scale full waveform tomography are not currently available.

## 26.5 Application of waveform tomography in helioseismology

Seismic waves on the Sun are generated in the convection zone, a turbulent layer making up the outer 30% of the Sun. Above the convection zone lies the photosphere, a layer often taken to be the "solar surface." It is believed that seismic waves are excited continuously and stochastically by numerous sources in a narrow layer just beneath ($\sim 75 \pm 50$ km below) the photosphere, as a result of supersonic downflow plumes formed at integranular lanes, which shock the immediate sub-surface layers, channeling fluid kinetic energy into acoustic radiation (e.g., Nordlund *et al*., 2009). The rapidly downwelling material forms thin "intergranular lanes," which surround broad, upwelling regions known as granules (Figure 26.4). The waves are acoustic in nature, and when they reach the photosphere, provided they are below a certain frequency known as the acoustic cut-off frequency of approximately 5.55 mHz, they are prevented from penetrating it due to its sharp density

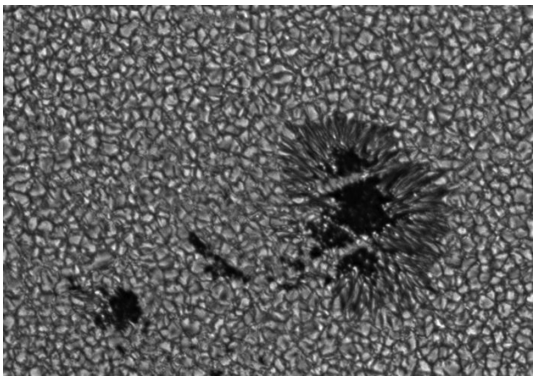

Figure 26.4 Photograph of the photosphere, containing a sunspot (right-hand side). It has a central dark region (umbra) surrounded by radial filaments (penumbra). The smaller polygons are convection cells, known as granules, typically 1–1.5 Mm across. Bright regions are upwelling, dark lines (intergranular lanes) are downwelling. The surface is highly dynamic, with each granule typically lasting about 10 minutes before being replaced by new granules. (Image courtesy of NASA.)

gradient. Each time a wave bounces off the photosphere it causes vertical particle motion, which can be measured via the Doppler shifting of absorption lines in the solar spectrum.

In reality the oscillations recorded at any particular point on the photosphere represent the combined effect of many waves arriving at that point. This produces "noisy" seismograms. In order to extract a coherent signal from the noise, most solar tomography is based on the cross-correlation of pairs of seismic traces to extract the Green function, i.e., the signal that would be obtained if there were a single source beneath the location of one of the data recorders and a receiver at the other.

Full waveform tomography using the approach outlined in Sections 26.1–26.3 has not yet been applied to real solar data. However, waveform tomography has great potential for enhancing the imaging of localized 3D structures near the surface of the Sun, in particular sunspots, and the development of suitable inversion algorithms is already underway (e.g., Hanasoge *et al*., 2011). Sunspots are dark regions of intense magnetism, typically between ten to a few tens of megameters across and, at 4000 K, about 2000 degrees cooler than their surroundings. They have lifetimes lasting anything between a few hours and several months. Although they are well-resolved photographically at the visible surface of the Sun (Figure 26.4), their structure below the photosphere cannot be observed directly and remains somewhat of a mystery. Tomographic imaging has been applied to probe their subsurface structure, and time–distance helioseismology has provided tantalizing first images of the structures and flows underneath sunspots (e.g., Zhao *et al*., 2001). However, Gizon *et al*. (2009) raised concerns over the effect of data filtering procedures on the validity of tomographic models. The presence of magnetic fields causes anisotropy, which needs to be carefully modeled (Hanasoge *et al*., 2012). In addition, the resolution that can be achieved by time–distance methods is limited by the fact that in the shallowest layers of the Sun the

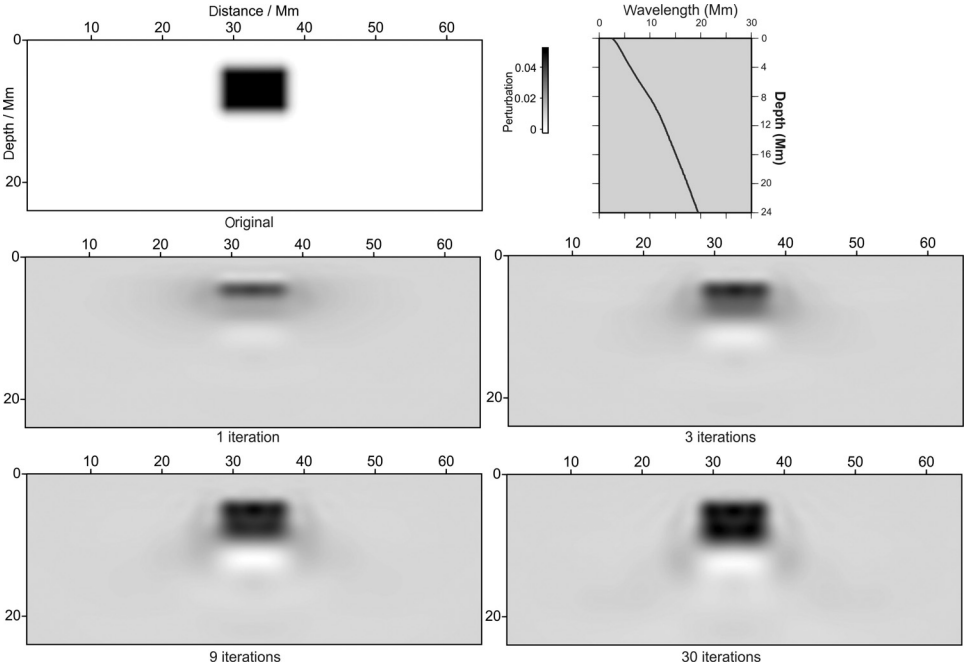


Figure 26.5 Two-dimensional synthetic test showing recovery of a small-sunspot-like structure from full waveform inversion, after Cobden *et al*. (2011). Top row: input seismic structure (left) and seismic wavelength as a function of depth (right) in the Sun's outermost convection zone. Other rows: sound-speed structure recovered from waveform inversion after various iteration numbers of the inversion algorithm. The wavefield is modeled using an acoustic wave equation without convection or magnetic field effects, and assuming a uniform distribution of acoustic sources near the surface. The inversion can retrieve sharp vertical and horizontal boundaries, which are notoriously difficult to obtain in travel-time inversions, as well as accurate imaging of the spatial dimensions and anomaly magnitude of a structure with a size comparable to the signal wavelength.

ratio of the seismic wavelength to path length is high. In particular, structures underneath sunspots have typically been imaged to depths of a little more than 10 Mm, but below this depth, the diameter of sunspots is only two or three times greater than the seismic wavelength at most, and thus their interior structure is potentially beyond the resolution of the ray approximation.

In response to this limitation, several finite-frequency approximations, such as the Fresnel zone, Born, and Rytov approximations, which allow for a single scattering event between the source and receiver, have been applied to solar acoustic data, and such approximations have allowed solar subsurface structure to be imaged to a greater depth and higher accuracy than with the ray approximation (e.g., Couvidat *et al*., 2004). Exploratory numerical simulations of solar wave propagation in 2D, based on the acoustic wave equation, have indicated that the application of full waveform tomography to solar data could potentially allow an even finer spatial resolution (Tong *et al*., 2003a,b; Cobden *et al*., 2011, Figure 26.5), as

well as the inference of additional parameters such as density and attenuation (Tong *et al*., 2003c; Tong, 2005). Nonetheless, the application of waveform tomography to helioseismic data presents several major challenges that are not encountered in terrestrial seismology.

In terrestrial seismology, seismic waveforms are modified by the elastic properties of the effectively static media through which they are transmitted. The transmission speed of the waves (km/s) is much shorter than the time scales over which the material is convecting (mm–cm/year), and tomography gives us a "snapshot" image of structures that are evolving on time scales of millions of years. However, on the Sun, the wave speeds (km/s) are comparable to the time scales of some convective motions, such that the propagation velocities of acoustic waves between two surface points are dependent not only on the temperature and density of the intermediate material, but also on the velocity with which it is convecting. Wave speeds are further modified by the spatially and temporally variable magnetic fields that are characteristic of plasmas such as the Sun, and seismic perturbations inside sunspots (e.g., Braun and Birch, 2008) may be produced by both thermal and magnetic effects. Inferring the magnetic field strength using measurements of propagation times of seismic waves is particularly challenging because heliomagnetism generates seismic anisotropy between field-parallel and field-perpendicular wave paths, and thus, wave paths that intersect at $90^\circ$ are required to study it, which likely have encountered different inhomogeneities along their mostly non-intersecting paths. At the same time, numerical modeling tells us that magnetic effects are non-trivial and sensitivity kernels that do not realistically account for them will lead to the simulation of incorrect wave travel times (Crouch *et al*., 2011; Braun *et al*., 2012). Thus, the need to include magnetic and convection effects makes complete numerical modeling of wave propagation much more challenging, both theoretically and computationally. Inclusion of magnetic effects in numerical models is a very recent development (Hanasoge *et al*., 2011; Crouch *et al*., 2011), but convective mass transfers have not yet been realistically incorporated.

A second problem in helioseismology is the stochastic and continuous nature of the acoustic sources, leading to the need to cross-correlate pairs of traces to extract a coherent signal. In order to converge at the desired Green function, and to enhance the signal-to-noise ratio, a large amount of data must be averaged. On the Sun, the rapid temporal evolution (hours to weeks) of the target structures prohibits extensive temporal averaging. To compensate, data are typically averaged over a patch or annulus of neighboring pixels. This spatial data-averaging would likely cancel out any improvements in the resolution of images that could be obtained in waveform tomography. Additionally, waveform amplitude analysis is hampered by the fact that the amplitudes of cross-correlated waveforms are influenced by the distribution of the acoustic sources (Cobden *et al*., 2010), which in turn is unknown. A promising solution to these problems may lie in the adjoint method of Tromp *et al*. (2010) for computing sensitivity kernels for cross-correlation observations. In this case the misfit functional can be any feature of the cross-correlated signal (travel times, waveforms, etc.). Applying this method to the Sun, a small number of "master pixels" are selected to be cross-correlated with any number of other pixels; computation of sensitivity kernels requires only three numerical simulations per master pixel. This method relies

on temporal averaging of cross-correlation time series rather than spatial averaging, thus avoiding the loss of resolution associated with spatial averaging, and it can accommodate non-uniform source distributions. It remains to be tested how successfully it can be applied to time-limited solar structures.

## 26.6 Future directions

In geophysics, waveform tomography is a relatively young, but fully established, method for obtaining high-resolution images of the Earth's interior. Active research is targeted at applying the method to increasingly larger areas and greater depths in the Earth, at shorter wavelengths with higher spatial resolution, as well as elucidation of physical properties beyond isotropic wavespeeds, such as density and attenuation. On the Sun, the application of full waveform tomography to real data remains a challenging but realistic prospect within our grasp. It will require the combination of a code that can realistically simulate both wave generation by stochastic sources and wave propagation through a fully magnetohydrodynamic medium, with an efficient inversion algorithm based on fitting cross-correlated time series. These three components have all been developed separately and future endeavors to combine them may allow us an unparalleled view into the internal structure of sunspots.