1

Introduction – modeling, migration, imaging, and inversion

1.1 Seismic data contains information

The fundamental assumption of this book is that seismic data contain decipherable information about Earth properties. We like to think of the Earth's subsurface as composed of regions or layers within which the properties change slowly, separated by boundaries across which properties vary rapidly, perhaps even discontinuously. The regions between boundaries propagate and sometimes distort seismic waves, while each boundary reflects some of the seismic energy. Thus, seismic data tend to contain sequences of discrete "events", each associable with a reflecting boundary, separated by relatively quiet intervals. Buried in the data is information on the location and composition of boundaries, and also on the properties of the regions between them.

Raw seismic data may be difficult to interpret, and seldom reveals the true location or amplitude of reflectors. Events may be obscured by noise and interference. Most locations in the subsurface are illuminated by many sources and receivers, from different directions and distances. Imposed on the data may be a source wavelet long enough to jumble events together and make them difficult to identify individually. Pairs of strong reflectors may produce multiple reflections that mask primary reflections from deeper reflectors. Seismic processing renders the data more interpretable by compressing the source wavelet, identifying and removing multiple reflections, aligning and compositing images of individual events, filtering noise, "migrating" events from apparent to actual locations, and "inverting" for Earth properties.

1.2 Models for propagation and reflection

Seismic migration is based on a simplified picture of seismic reflection data which assumes primary reflections only. Multiple reflections, if they cannot be ignored a priori, are assumed to have been removed from the data. In this picture, a seismic wave travels from a source at or near the Earth's surface to one or more reflectors in the Earth. At each reflector, a reflected wave is generated which travels to receivers placed at or near the Earth's surface.

To migrate seismic data requires physical descriptions of the processes of wave propagation and of reflection. Without a propagation model, one can only wonder how seismic energy got from source to reflector to receiver, and unravelling the process is impossible. Without a reflection model, one cannot interpret reflected energy or relate it to properties of the reflector. The two models provide a descriptive framework from which one can deduce where reflectors are, and an interpretive framework from which one can deduce reflector properties.

In the real Earth, propagation and reflection are intertwined, and treating them as separable entities is based as much on perception as on reality. Surprising as it may seem, this simplified picture usually works out fairly well.

Ideally, descriptions of seismic wave propagation and reflection stem from a common underlying physical model. In practice, however, it may be difficult to unify the two, and separate models are often employed for the two processes. This leads to a propagation model based on a simplified "background" velocity structure, engineered so as to not produce reflections, and a reflection model based on complex, highly localized changes in seismic properties.

This happens in part because of the bandlimited nature of seismic data. The data directly illuminate changes in seismic properties within a range of wavelengths. These changes give rise to what we perceive as reflections. Slower changes do not produce perceptible reflected energy, but do largely determine how seismic waves propagate to and from reflections. The physics of propagation and reflection partially decouple because they inhabit regions of different wavelength. Changes more rapid than the illuminated wavelengths also occur, but they do not directly produce perceptible reflections; rather, they can and do affect the effective seismic properties. They may in fact affect propagation and reflection in different ways, creating an apparent disparity between the reflection and propagation models.

Other factors may also favor separate models for propagation and reflection. The purpose of a seismic model is not to perfectly describe the contents of the real Earth, which is more complex than any model one is likely to devise. The model is meant to allow one to describe, analyze, and interpret seismic processes well enough to derive the information required. To do worse is to fail; to do better is to waste effort. It follows that the physical model that best describes propagation may be different than the one that best describes reflections.

Even so, the most satisfying approach would be to work from a single physical model that describes both reflections and propagation. The inverse scattering approach to seismic processing conforms to this ideal. For the most part, however, the theory and practice of seismic imaging retains two separate models. Seeing value in both approaches, this text is divided into two volumes, the first approaching the inverse seismic problem from the simpler point of view, the second following a more comprehensive approach.

Even within Volume I, there are multiple choices for propagation and reflection models. The best choice depends upon circumstance – upon the geology, upon the data, and upon the processing objectives. If one is interested primarily in the location of reflectors, or if data quality does not support amplitude recovery, then it makes little sense to impose all the trappings of a true-amplitude model on the processing. Generally, we support the principle of minimum complexity: one should use the simplest model that works for the circumstances at hand. While the perfectionist within may cry out for the most realistic model possible, that personality may need to be sedated. It is just as much in error to carry around an unnecessarily realistic Earth model as it is to insist on meaningless precision in calculations. Penurious instincts may need to be suppressed as well: dedicating inadequate resources guarantees failure before one begins.

1.3 Going forward to go back

Modeling, migration, imaging, and inversion of seismic data are dealt with in this Volume. Generally, we assume that other seismic processes, where necessary, have already been performed. We are mostly concerned with the inverse seismic problem, in which Earth properties are inferred from seismic data. However, to go backward successfully, one must be able to go forward, predicting the data for a given set of Earth properties. *Modeling*, as used here, refers to processes for simulating seismic data. The starting point for the simulation may be a map of Earth properties, or, less fundamentally, a map of reflectivity images or an image function.

1.4 Seismic and non-seismic imaging

Migration and inversion have to do with the inverse problem. *Imaging*, narrowly defined, might be considered synonymous with migration. From a broader perspective, virtually all seismic data contains images. If one defines an imaging process as an operation that forms, modifies, or manipulates seismic images, then imaging encompasses modeling, migration, and a host of other processes.

Seismic imaging has relatives in other fields, including radar, sonar, remote sensing, and various forms of medical imaging. All these fields share a common wave-theoretical basis, and have to some extent shared technology and algorithms. However, in some respects the seismic problem is, if not unique, nearly so. Most other imaging problems are very much in the far field, with targets many hundreds or thousands of wavelengths from sources and receivers. In contrast, a 30 hertz seismic wave traveling at 10 000 feet per second (fps) covers only three wavelengths per 1000 feet (300 m). We can usually get away with treating seismic propagation as far field, but not always. Propagation velocities in other disciplines are, with a few exceptions, much less variable than seismic velocities. Seismic data sets tend to be relatively large, and not amenable to processing in real time. These and other differences limit the possibilities for cross-disciplinary technology transfer.

1.5 Motivation for migration

While unmigrated seismic data contain reflector images, they are geometrically distorted; certainly vertically, since images appear in time rather than depth, and generally laterally as well (Claerbout, 1971, 1976; Berkhout, 1982; Stolt & Benson, 1986). *Migration* refers to a process which builds, moves, or "migrates" reflector images as close as possible to the geometric location of the actual reflectors. Figures 1.1 to 1.3 illustrate the process. Figure 1.1 shows a highly simplified two-dimensional dome, with prospective oil reservoirs in the cap and on the flanks.

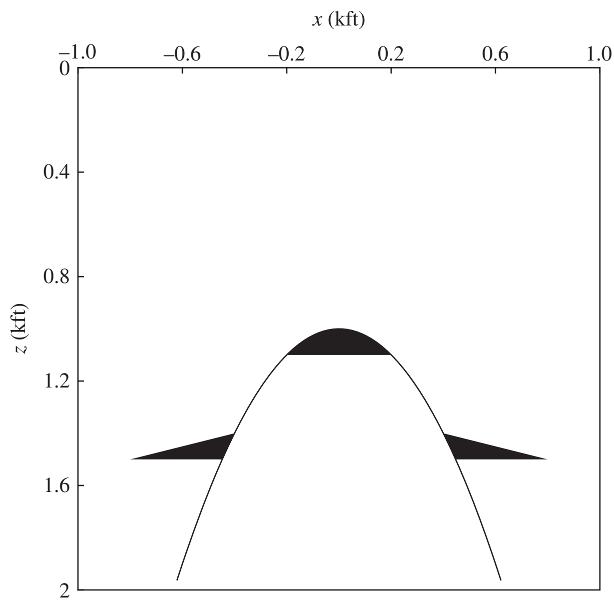

Figure 1.1 Two-dimensional dome model. Seismic velocity is a constant 10 000 fps. The model includes three "reservoirs", one in the dome cap and two on the flanks. 1 kft = 1000 feet.

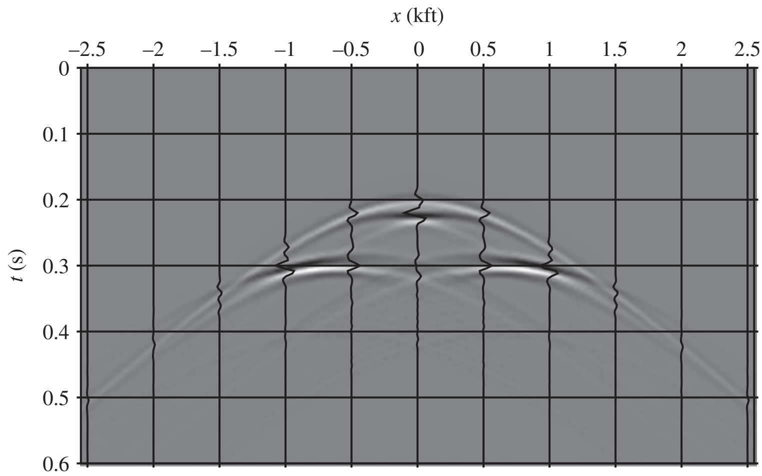

Figure 1.2 Synthesized zero-offset data from dome model. Source–receiver spacing is 20 feet. Time increments are 0.004 s. 1 kft = 1000 feet.

With the dome as initial model, and assuming a constant velocity of 10 000 fps throughout, Figure 1.2 shows the corresponding seismogram for a reflection experiment, in which bandlimited impulsive sources are spaced along the surface at 20 ft intervals, and a response from each source is recorded by a single receiver at the same location as the source. Visible on the seismogram is a dome-like structure much broader than the actual dome. There is some indication of a flat reflector just beneath the top of the dome, though its apparent size and shape do not correlate well with the bottom of the actual reservoir. Reflections from the two reservoirs flanking the dome have transformed into minidomes inside the central dome, leaving one to wonder what and where they really are. An explorer looking at this data would want to know several things: (1) can any potential oil reservoirs be identified; (2) can they be accurately located, and (3) can their size be accurately estimated? Meeting any of these objectives directly from these data would be difficult.

A migration of this data set is shown in Figure 1.3. Like the data it came from, the image is bandlimited, but the size and location of the dome and the reservoirs is accurately recovered. In more realistic situations, where structures are more complex, velocities are variable, and the data are undersampled, the importance of obtaining an accurate image can only increase. In areas of extremely complex geology, reflections may not even be visible in unmigrated data, in which case the only hope for an interpretable image is through a sophisticated imaging algorithm.

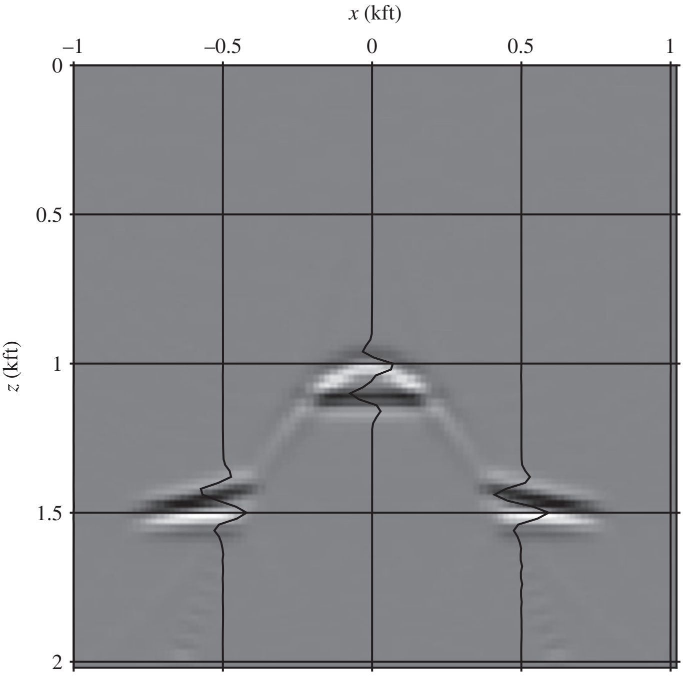

Figure 1.3 Migration of the dome model data. The dome and reservoirs have assumed their geometrically correct positions.

1.6 Time and depth migration

Time migration and depth migration are two types of migration which differ in their ability to image the objective (Schultz & Sherwood, 1980; Stolt & Benson, 1986). Depth migration strives for fidelity both laterally and vertically, whereas time migration leaves the vertical direction in traveltime units. Depth migration requires a detailed model of propagation velocities within the Earth, while time migration needs only an average, or rms (root-mean-squared) velocity structure. Given the right velocity field, depth migration can produce superior images, but time migration is less sensitive to velocity error.

1.7 Migration velocity

Even though time migration may be less sensitive to local variations in seismic velocity, neither form of migration can be expected to yield a good image using bad velocities. To illustrate sensitivity to velocity, the next two figures show migrations of the data in Figure 1.2. Figure 1.4 results from migrating with a velocity that is

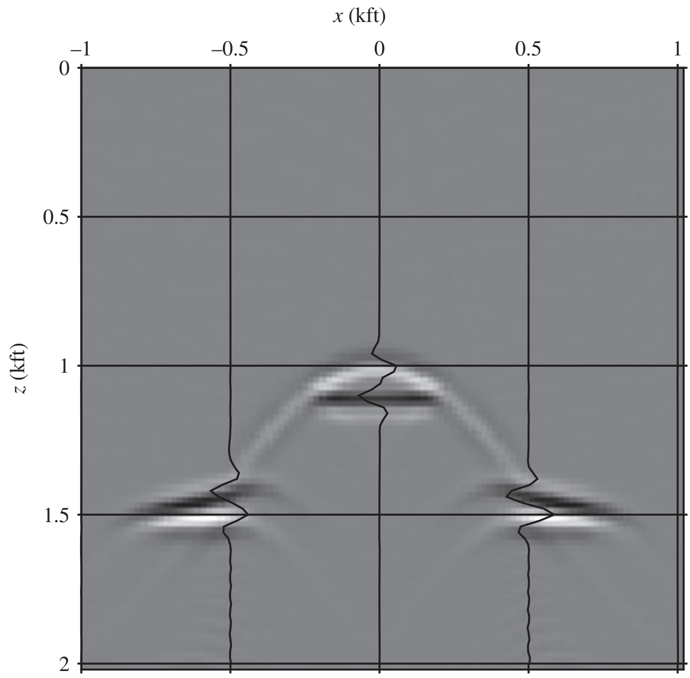

Figure 1.4 Undermigration of the dome model data. Migration velocity is 9500 fps. Neither the dome nor the reservoirs are correctly imaged.

5 percent too low. The dome is too wide, the reservoir terminations are misplaced, and downward-frowning edge diffractions are still very much in evidence.

Figure 1.5 shows a migration with a velocity that is 5 percent too high. The dome is too narrow, the reservoir terminations misplaced, not to mention obscured by upward-smiling diffractions.

These simple examples make it clear that even a small error in migration velocity will result in an inaccurate image. The news, however, is not all bad. Comparing the three migrated images, the correctly migrated one displays a crispness not enjoyed by the other two. It is often possible to spot an undermigrated (overmigrated) image by the presence of downward-pointing (upward-pointing) residual diffractions, and to make velocity adjustments accordingly. Of course, in real life, things are seldom that simple. Determination of correct migration velocities under differing circumstances is given little attention in this text, not because it is unimportant, but because it is a vast, difficult subject, worthy of a volume in itself (see e.g. Liu & Bleistein, 1995; Yilmaz, 2001; Fomel, 2003; Sava *et al.*, 2005).

The migration velocity issue is sometimes viewed as paradoxical, in that on the one hand one cannot perfectly migrate a data set without the correct velocity, while

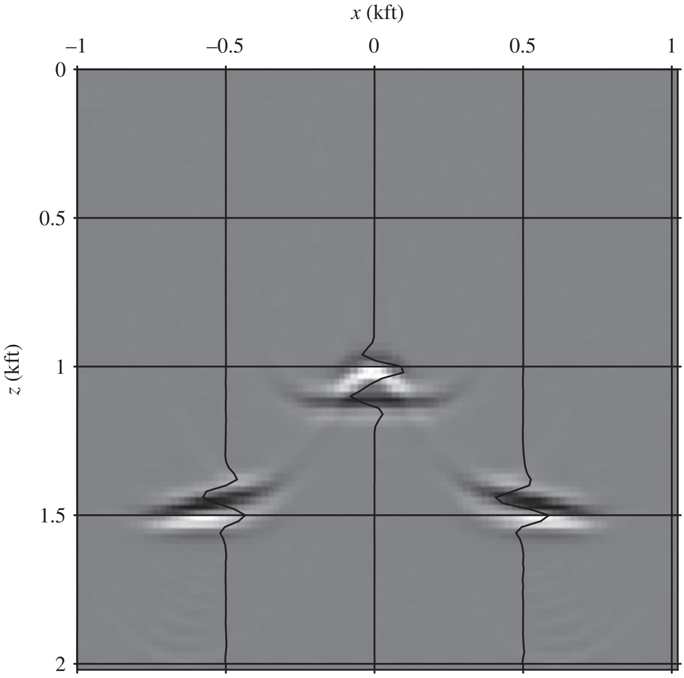

Figure 1.5 Overmigration of the dome model data. Migration velocity is 10 500 fps. Neither the dome nor the reservoirs are correctly imaged.

on the other hand, if the velocity structure were perfectly known, one would already have a structural image of the subsurface. A better view is to acknowledge that the velocity information is imbedded in the data. Extracting this information is a difficult, nonlinear process which is best done concurrently with the imaging. Once this interleaved process is complete, one should have both an accurate velocity map and an accurate seismic image.

1.8 Full-wave and asymptotic migration

Migration methods can be distinguished by whether they employ a full wave equation or an asymptotic or far-field approximation to it. Asymptotic methods assume that wavefronts can be treated as locally planar, with a well-defined direction and local wavenumber. The distinction between asymptotic and full-wave-equation methods is more equivocal than one might think, in that many ostensibly full-wave-equation algorithms have allowed some asymptotic assumptions to slip in, noticed or unnoticed, along the way. Nevertheless, the distinction is worth making. For an individual primary wavefront, the asymptotic approximation is for the most part pretty well satisfied. A complete asymptotic description of a wave in complex media, however, may be difficult or cumbersome. Multiple reflections, largely untreated in Volume I of this book, are more likely to violate the asymptotic assumption. On the plus side, asymptotic algorithms are simple and powerful, and are capable of extracting useful information even under conditions where the assumptions behind them are not completely met. Dealing with missing or limited data may be simplified with asymptotic approximations. It is easier to think (and interpret) in terms of localized events than of extended wavefields, even if a wavefield description is more complete and accurate. The asymptotic approximation certainly has a place in seismic imaging, hence is introduced without hesitation in Volume I.

1.9 Seismic migration and inversion

Seismic migration is an inverse process, and one would not be wrong to refer to seismic migration as seismic inversion. Typically, however, the term *inversion* is reserved for either (1) a process in which a reflector image (hopefully migrated) is used to predict quantitative changes in physical properties; or (2) a process which combines migration with property prediction. Sometimes the term *migration–inversion* is used to mean either the latter process or a migration algorithm in which the imaged reflection amplitudes are physically meaningful. Use of this term suggests that ordinary migration is indifferent to amplitude, which is not exactly the case. All migration algorithms produce amplitude as well as phase, and the amplitude is often interpretable, provided one understands what the algorithm has done. However, one must be aware that many migration algorithms distort amplitudes by making approximations or by using physically inadequate descriptions of wave propagation or reflection.

1.10 True amplitude migration

In a similar vein, one also finds the terms *true-amplitude migration* and *amplitude-preserving migration* (see, e.g., Black *et al.*, 1993; Schleicher *et al.*, 2007). Again, these terms are sometimes used interchangeably. Purists will claim that there is no such thing as true amplitude, let alone true-amplitude migration. Those who use the term may be thinking of “true” not in the sense of being absolutely correct, but in the sense of faithfully preserving the amplitude information. Standards may depend on need and level of ambition. Ideally, one would like an amplitude twice as large to correspond to twice the reflectivity. Practically, one might have to settle for this condition to be satisfied within some range of depths, or dips, or incident angles.

Sometimes, requirements are more rigorous: the amplitudes of multiple reflections depend on the absolute magnitudes of the constituent primaries.

1.11 Linear and nonlinear processes

A useful way to look at modeling and migration is as forward and inverse mappings between a model space and a data space. For convenience, we often treat these mappings as linear operations, though in reality, they are not. Multiple reflections are obviously nonlinearly related to the Earth model. Multiples aside, the amplitudes of primary reflections are also nonlinear. Consider, for example, a scalar plane wave traveling at velocity c_0. An abrupt change in velocity to c_1 will produce a reflection of strength $R = (c_1 - c_0) / (c_1 + c_0)$, nonlinear in the velocity change except in the limit as c_1 approaches c_0. One can argue that the data amplitude is proportional to reflection strength. This leads one to devise linear forward and inverse mappings between data and reflectivity, leaving the nonlinearities to be dealt with by subsequent processing and interpretation.

Subsurface complexity, especially if not fully known a priori, can also introduce nonlinearities. Depth migration requires that the velocity structure between a reflector and the sources and receivers be accurately specified. If the velocities are imperfectly specified, the resulting images may be distorted, spurious, or missing. Recovery, where possible, is a nonlinear operation.

In Volume I, we generally treat the forward and inverse mappings as linear, in the knowledge that Volume II will address the full, nonlinear problem.

At heart, seismic migration is a simple concept. Imaging algorithms, though diverse, share this underlying simplicity. Seismic data, however, are complex, and computational resources are finite. Expectations continue to rise, and as imaging algorithms become more ambitious they also become more complicated.