Double-difference waveform inversion: Feasibility and robustness study with pressure data

Di Yang[1], Mark Meadows[2], Phil Inderwiesen[3], Jorge Landa[3],
Alison Malcolm[4], and Michael Fehler[5]

ABSTRACT

Time-lapse seismic data are widely used to monitor reservoir changes. Qualitative comparisons between baseline and monitor data sets or image volumes provide information about fluid and pressure effects within the reservoir during production. However, to perform real quantitative analysis of such reservoir changes, quantitative estimates of the elastic parameters are required as input parameters to rock-physics-based reservoir models. Full-waveform inversion has been proposed as a potential tool for retrieving subsurface properties, such as P- and S-wave velocities and density by fitting simulated waveforms to seismic data. An extension of this method to time-lapse applications seems straightforward, but, in fact, it requires more tailored processes such as double-difference waveform inversion (DDWI). We used realistic 2D synthetic pressure data examples to compare the performance of DDWI with that of two other inversion schemes: one using the same starting model for both inversions and the other starting the monitor inversion with the final baseline inversion model. The data simulation and inversion were based on acoustic theory. Although P-wave velocity changes were reliably recovered by each inversion method, DDWI was found to deliver the best results when perfectly repeated surveys were used. However, differencing the baseline and monitor data sets, as required by DDWI, could be found to be sensitive to the presence of survey nonrepeatability. To investigate the feasibility of using DDWI in practice, the dependence of DDWI on the quality of the baseline models and its robustness to survey nonrepeatability were studied with numerical tests. Various types of nonrepeatability were considered separately in the synthetic tests, including random noise, acquisition geometry mismatch, source wavelet discrepancy, and overburden velocity changes. A study of the correlation between the levels and types of nonrepeatability and the resulting contamination of the inversion results found that, for pressure data, DDWI was capable of inverting reliably for P-wave velocity changes under realistic survey nonrepeatability conditions.

INTRODUCTION

Knowledge of the changes in reservoir properties resulting from extracting hydrocarbons or injecting fluid is critical for optimizing production. This information can be obtained using time-lapse seismic monitoring technology because the seismic response is sensitive to fluid and pressure effects in the reservoir. Quantitative analysis of time-lapse data uses information, such as seismic amplitude maps or image/time shifts (Arts et al., 2004) to provide a detailed description of reservoir changes based on relationships between the seismic response and changes in pore pressure and fluid saturation (Landrø, 2001; Trani et al., 2011). However, many such analyses are based on a linear approximation of the time-lapse seismic response as a function of reservoir property changes, which

would likely fail if the changes are large, or when nonlinear effects, such as gas saturation changes, are present. Furthermore, map-based methods cannot provide volumetric information, which could better constrain the dynamic reservoir changes in space.

Full-waveform inversion (FWI) aims to estimate the subsurface density and elastic parameters directly from seismic records through a nonlinear data-fitting process (Tarantola, 1984; Virieux and Operto, 2009). Ideally, the extension of FWI to the time lapse is straightforward. Two FWI runs can be conducted for the baseline and monitor data sets, and the difference of the two resulting models should reveal the reservoir property changes. Nevertheless, nonlinear artifacts arising from the nonlinear nature of the inverse problem introduce differences between the inverted models in addition to the real time-lapse changes. To address this problem, Watanabe et al. (2005) apply differential waveform tomography in the frequency domain to crosswell time-lapse data during gas production and show that the results are more accurate for estimating velocity changes in small regions than those obtained using conventional inversion schemes. Onishi et al. (2009) apply a similar strategy to conduct differential traveltime tomography using crosswell surveys. Zheng et al. (2011) propose a double-difference waveform inversion (DDWI) algorithm using time-lapse reflection data in the time domain and demonstrate, with synthetic data, that the method has the potential to produce reliable estimates of reservoir changes. Zheng et al. (2011) compare DDWI with conventional approaches using synthetic ocean-bottom-cable data. Synthetic marine time-lapse streamer data were used for a similar performance comparison in Routh et al. (2012). Asnaashari et al. (2011) conduct a sensitivity study on DDWI results with respect to baseline models. Asnaashari et al. (2012) test the performance of DDWI in the presence of random noise. Maharramov and Biondi (2014) compare DDWI with alternative methods in the frequency domain, including a cross-updating scheme and simultaneous inversions with regularization. The cross-updating scheme inverts the baseline and monitor data in alternate fashion, whereas the simultaneous inversion scheme combines the two data sets in one cost function without data differencing, and it uses a model regularization term to penalize changes outside the reservoir. Raknes et al. (2013) and Raknes and Arntsen (2014) apply DDWI to streamer data that were acquired without time-lapse imaging in mind, and conclude

that the resulting low survey repeatability compromised the DDWI images. A successful real-data application of DDWI is reported by Yang et al. (2013), who use well-repeated ocean-bottom-cable data sets. Given these previous results, it is clear that a comprehensive feasibility study of DDWI under realistic acquisition conditions would help our understanding of real-data applications. In particular, it would be valuable to know in a quantitative sense how practical survey conditions affect the performance and robustness of DDWI, as well as what artifacts might occur in such cases.

In this work, we apply the DDWI methodology of Denli and Huang (2009) to a synthetic data set and investigate the feasibility, advantages, and limitations of DDWI when applied to realistic time-lapse data acquisition scenarios. In the following sections, we first describe the mathematics of DDWI and its implementation. We then compare the performance of DDWI with that of two other inversion schemes using (1) the same initial model for baseline and monitor inversions and (2) the baseline inversion result as a starting model for time-lapse monitor inversion, with both schemes using an acoustic model and noise-free pressure data. The dependence of DDWI on the quality of the baseline model is discussed using several baseline models of increasing accuracy and levels of convergence obtained by applying more iterations of FWI. The success of 4D seismic analysis also relies on the repeatability of the time-lapse surveys. Coprocessed data, in which the nonreservoir-related differences are minimized, give rise to much improved time-lapse results (Rushmere et al., 2010; Campbell et al., 2011). In practice, however, it is very difficult to correct perfectly for the lack of repeatability between surveys (Lumley, 2010). To investigate the effect of survey nonrepeatability on DDWI, several common acquisition mismatches are discussed separately, including contamination with random noise, survey source and receiver positioning errors, source wavelet discrepancies, and seasonal water velocity changes. For simplicity, all our numerical tests are conducted with a 2D acoustic finite-difference model.

METHODOLOGY

Standard FWI can be expressed as a minimization problem with the cost function

$$E_{\text{standard}}(\mathbf{m}) = \frac{1}{2} \|\mathbf{u}(\mathbf{m}) - \mathbf{d}\|^2, \quad (1)$$

where $\mathbf{u}$ is the modeled synthetic waveform sampled at receiver positions, $\mathbf{d}$ is the acquired field data, and $\mathbf{m}$ is the model parameter that is inverted for (see, e.g., Virieux and Operto [2009] for a review of FWI). Many successful real-data applications have also been published (Sears et al., 2010; Prieux et al., 2013; Warner et al., 2013), which establish FWI as a tool for quantitative subsurface property estimation.

To extend FWI to the time-lapse case, the most straightforward strategy is to execute two inversions for the baseline and monitor data sets, respectively. We call this strategy *scheme I* in this paper. As shown in Figure 1, subtraction of the final models should give the property changes between surveys. One could argue that it is wasteful to start the time-lapse monitor inversion in scheme I from an initial model that is independent of the baseline model, and that it would be more reasonable to start from the final model of the baseline inversion. We refer to such a strategy as *scheme II*. As shown in Figure 2, the updated part of the model in the monitor inversion
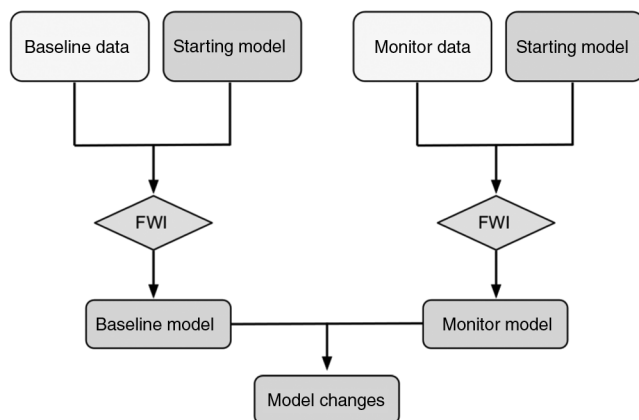

Figure 1. Scheme I: Two independent FWI are conducted for the baseline and monitor data sets, respectively. The model changes are obtained by subtracting the inverted baseline model from the inverted monitor model.

from scheme II should show the correct property changes, provided that the baseline inversion has converged. scheme III, as shown in Figure 3, is DDWI. In DDWI, the monitor inversion also starts from the final baseline inversion model, but the cost function is changed to

$$E_{\mathrm{DDWI}}(\mathbf{m}) = \frac{1}{2} \| (\mathbf{u}_{\mathrm{monitor}} - \mathbf{u}_{\mathrm{baseline}}) - (\mathbf{d}_{\mathrm{monitor}} - \mathbf{d}_{\mathrm{baseline}}) \|^2, \tag{2}$$

where $\mathbf{u}_{\mathrm{baseline}}$ and $\mathbf{u}_{\mathrm{monitor}}$ are simulated waveforms from the baseline inverted model $\mathbf{m}_0$, and the monitor model $\mathbf{m}$, respectively, with $\mathbf{m}$ being iteratively updated. The field data are $\mathbf{d}_{\mathrm{baseline}}$ and $\mathbf{d}_{\mathrm{monitor}}$ from the baseline and monitor surveys, respectively. The name "double difference" comes from the two differences in equation 2, one between baseline and monitor field data sets and one between modeled baseline and monitor data sets. As $E(\mathbf{m})$ is minimized, the property changes $(\mathbf{m} - \mathbf{m}_0)$ corresponding to the data differences are recovered. In practice, before starting the time-lapse inversion, we invert the baseline data set to derive the baseline model $\mathbf{m}_0$. Using $\mathbf{m}_0$, we then generate a synthetic data set $\mathbf{u}_{\mathrm{baseline}}$. To allow the use of standard FWI algorithms, a synthesized monitor data set $\mathbf{d}_{\mathrm{syn}}$ is created by adding the data difference $(\mathbf{d}_{\mathrm{monitor}} - \mathbf{d}_{\mathrm{baseline}})$ to $\mathbf{u}_{\mathrm{baseline}}$. The inverse problem is then reduced to a standard FWI with cost function

$$E(\delta\mathbf{m}) = \frac{1}{2} \| \mathbf{u}_{\mathrm{monitor}}(\mathbf{m}_0 + \delta\mathbf{m}) - \mathbf{d}_{\mathrm{syn}} \|^2, \tag{3}$$

and can be solved by regular FWI solvers. The only extra step compared with schemes I or II is the synthesis of $\mathbf{d}_{\mathrm{syn}}$, which requires a trivial amount of computation in the overall process.

SCHEME COMPARISON WITH ACOUSTIC INVERSION

Figure 4 shows the synthetic P-wave velocity and density models used for the numerical study. The dominant geologic structure is the anticline in the center, which lies underneath a sloping water bottom. The layering of the model is very detailed to simulate a realistic sedimentary environment. The synthetic data are generated by finite-difference modeling using 64 sources (250-m spacing) and 680 receivers (25-m spacing), all evenly spaced on the water surface. This mimics an ocean-bottom-cable acquisition after applying reciprocity. The synthetic model grid size is 6.25 m, and the time-step size is 1 ms. We use absorbing boundaries on all sides of the model; therefore, no surface-related multiples are modeled. A standard Ricker wavelet is used as the source time function, and its frequency band is centered at 5 Hz. The low-frequency components are included to enhance the recovery of the low-wavenumber part of the model. An example shot gather from the simulated baseline survey is shown in Figure 5. The direct arrivals are not filtered out; however, they do not contribute to the FWI solution because we use the correct water velocity in the initial model. A time-domain finite-difference solver is used as the wave simulator. The inversion engine is also time domain, and uses a nonlinear conjugate gradient method to obtain the search directions. An L2-norm cost function is adopted, and each inversion is stopped when the same number of iterations is reached each time. The P-wave velocity and density are inverted simultaneously.

For all listed schemes, we first apply FWI to the baseline data to obtain the baseline velocity model. To make the test more realistic, we assume very limited a priori information and use the linearly increasing models shown in Figure 6 for the initial P-wave velocity and density models. The inverted models from FWI are shown in Figure 7. In the inverted velocity model, the anticline is well recovered, and the layers are resolved with a resolution limited by the frequency band of the data. The lower-right and lower-left parts of the model are not as well recovered as the center part because of the limited illumination of the survey. The inverted density model also shows similar structures and appears to recover the true density values successfully. However, the recovery is aided by the correlation between the low-wavenumber components of the density and velocity models used here. There is an inherent ambiguity between
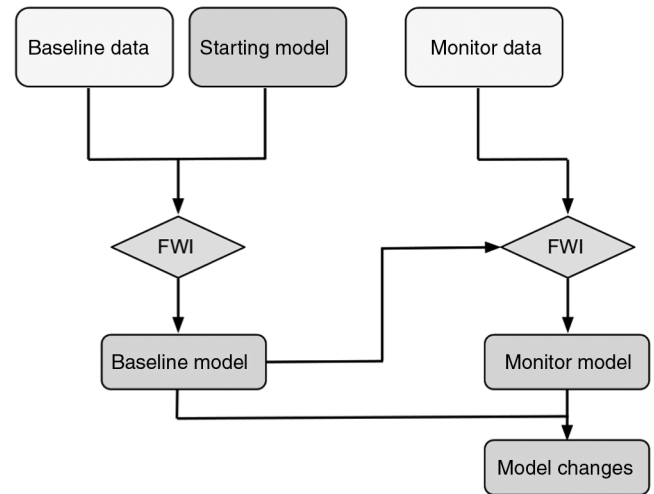

Figure 2. Scheme II: The baseline model is found by FWI with the baseline data set. The monitor inversion starts from the baseline inversion result. The model updates are considered to be model changes between baseline and monitor.
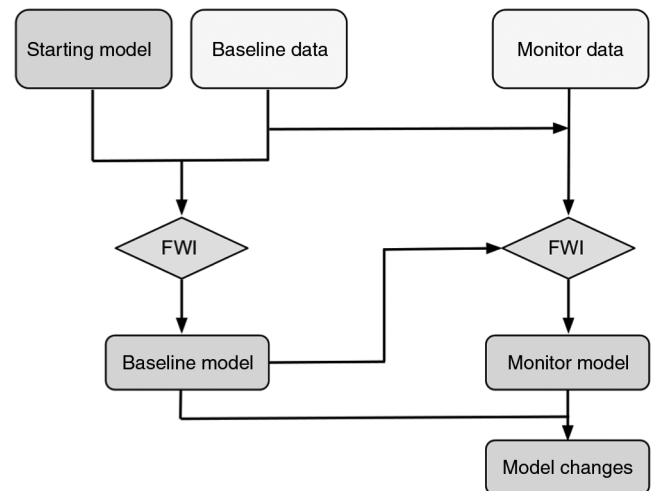

Figure 3. Scheme III: The monitor inversion starts from the baseline inversion result, and inverts the data difference between the baseline and monitor data sets. The model updates are considered to be model changes between baseline and monitor.

velocity and density inversions because of the similar radiation patterns generated by the velocity and density point anomalies (Tarantola, 1986; Virieux and Operto, 2009; Prieux et al., 2013). As a result, the detailed features of our inverted density model are not accurately estimated, even though the gross features are correct.

For constructing the monitor models, we implanted realistic changes in P-wave velocity and density that are typically observed in practice into the baseline models, as shown in Figure 8. The changes are not based on a real reservoir, but they are meant to represent a waterflood scenario with a variety of challenging velocity and density changes. To test FWI's capability of distinguishing such
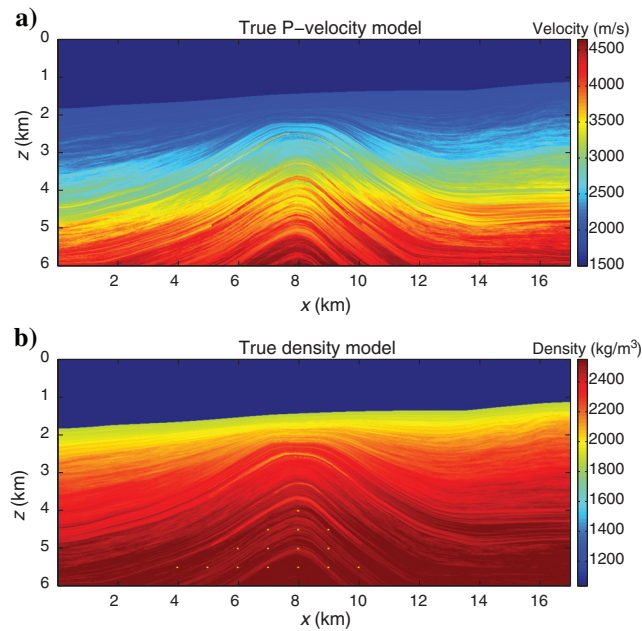

Figure 4. The true baseline (a) P-wave velocity and (b) density models that are used for generating realistic synthetic data for the baseline survey. The small dots at the bottom of the density model are to benchmark inversion resolution and also to demonstrate the ambiguity between P-wave velocity and density.


Figure 5. Shot gather generated by a surface source in the middle of the model.
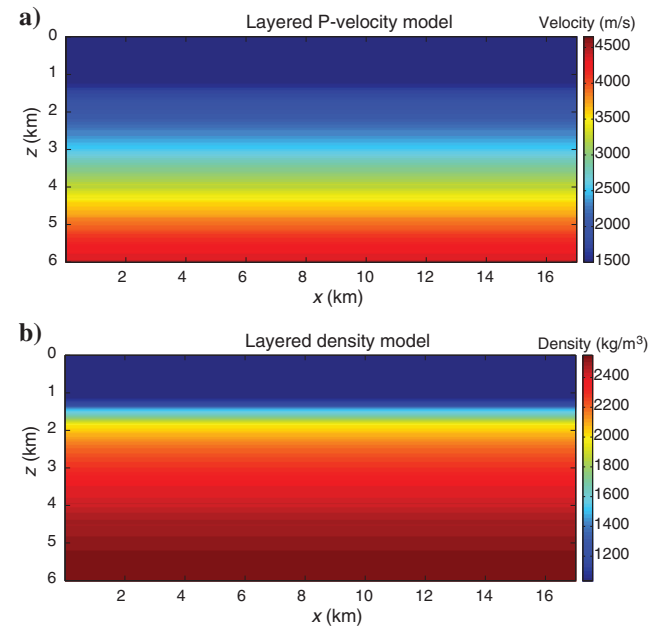

Figure 6. The starting (a) P-wave velocity and (b) density models for baseline FWI. The models are obtained by horizontally averaging the true models in Figure 4.
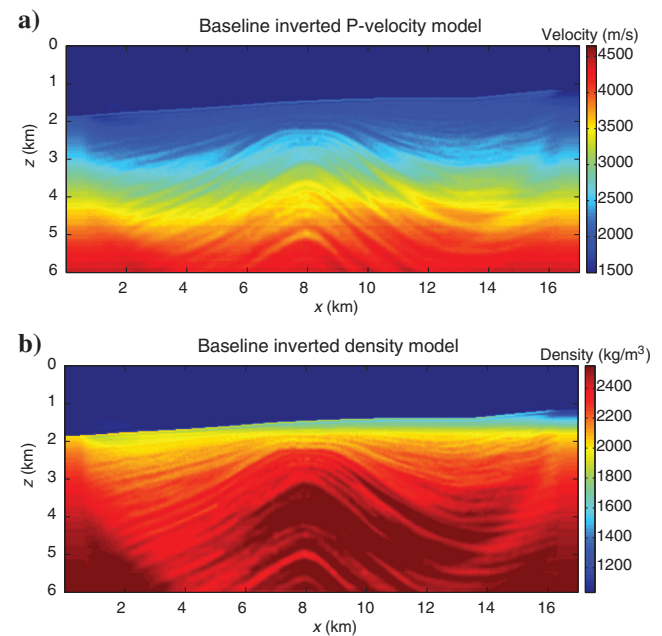

Figure 7. The final baseline (a) P-wave velocity and (b) density models after 60 FWI iterations. The cost function is reduced to less than 5% of its original value.

parameter changes during the inversion, some of the density and P-wave velocity changes have different signs in the three reservoirs. P-wave velocity changes in the shallower reservoirs are smaller than those in the deepest reservoir. By contrast, the relatively shallow density changes are stronger than the deeper changes, and there are no time-lapse density decreases.

The initial P-wave velocity and density models used in the baseline and monitor inversions in scheme I are the same as those shown in Figure 6. The inverted P-wave velocity and density changes from the subtraction between the inverted monitor and baseline models are shown in Figures 9b and 10b, respectively. Major features of the P-wave velocity changes are recovered, including velocity increases and decreases. However, amplitudes of the velocity changes are only partially recovered (approximately 50%). This can be explained by the fact that a typical reservoir thickness for this realistic

model is 80 m, whereas half the seismic wavelength at the reservoir depth is approximately 300 m, as is the theoretical resolution of FWI. As a consequence, the peak velocity values are smoothed out because of limited seismic resolution. In addition to the true changes, other changes that follow the structures in the background are visible, albeit weak in amplitude. The pattern of inverted density changes is basically the same as that of the P-velocity changes. However, the inversion adds some reductions in density, which do not exist in the real density model. It is difficult to differentiate density and P-velocity changes with pressure data because the two parameters generate very similar radiation patterns, even over a large range of scattering angles (Tarantola, 1986; Prieux et al., 2013). P-wave velocity is better constrained by the P-wave kinematic information in the data. Density is estimated from the amplitude information, which is also affected by P-wave velocity. Hence, density changes are not accurately recovered.

In scheme II, we start the time-lapse inversion from the final model that we obtained from the baseline inversion (Figure 7). Figures 9c and 10c show the inverted P-wave velocity and density models, respectively. It is obvious that the background structures, including the seafloor properties, are updated together with the reservoir changes. The resolved changes are also weaker in amplitude compared with those from scheme I.

Scheme III (DDWI) results are shown in Figures 9d and 10d. The results are cleaner than those of the other two schemes, and the amplitude is better recovered as well. However, the density ambiguity is still not resolved. To invert for the density more accurately, a different model parametrization, such as velocity-impedance, would be helpful, as discussed in Tarantola (1986) and Prieux et al. (2013). We will not further address this issue here because it is not directly associated with our analysis of different time-lapse inversion schemes.

The major improvement in scheme III, as compared with schemes I and II, is the removal of the coherent background structures (model residual) that are not related to the actual reservoir changes. These structures correspond to the residuals in the baseline inversion $\mathbf{R}_{baseline} = \mathbf{u}_{baseline} - \mathbf{d}_{baseline}$. Because there is no perfect inversion, $\mathbf{R}_{baseline}$ is never exactly zero. For scheme I, the nonlinearity of the inverse problem leads to different levels of convergence for baseline and monitor inversions at different locations in the model, where $\mathbf{R}_{baseline} \neq \mathbf{R}_{monitor}$. Accordingly, the subtraction between models gives nonzero contributions over the entire model. For scheme II, $\mathbf{R}_{baseline}$ gets injected into the model together
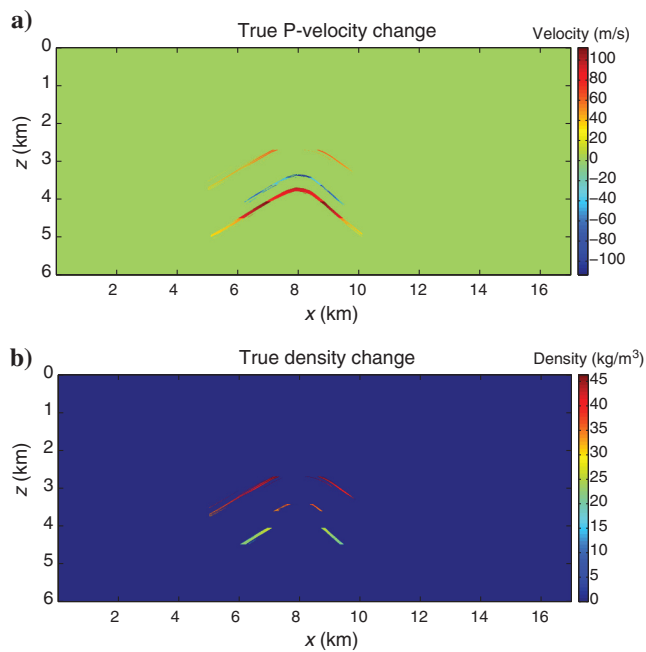

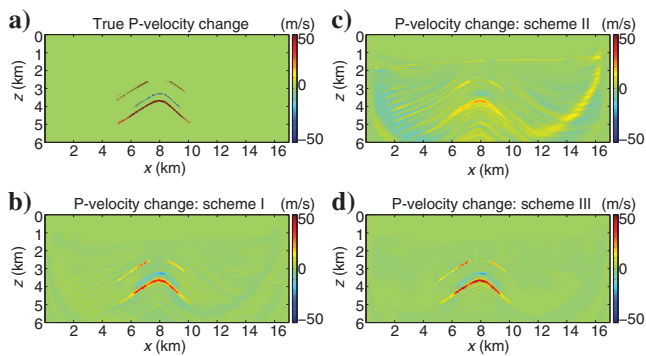Figure 8. The true time-lapse changes in (a) P-wave velocity and (b) density are confined to three reservoirs.

Figure 9. (a) The true time-lapse changes in P-wave velocity, with color scale clipped at ±50 m/s; (b-d) the time-lapse P-wave velocity changes recovered by inversion schemes I, II, and III, respectively.
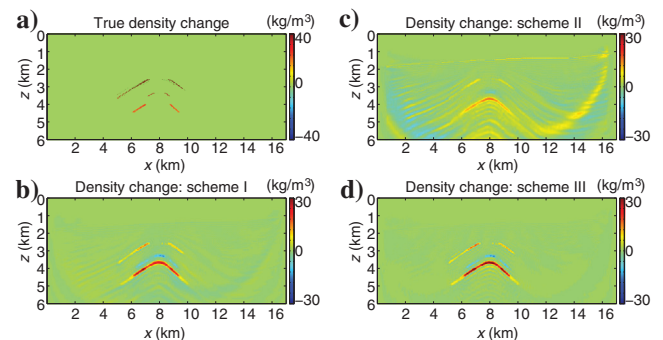

Figure 10. (a) The true time-lapse changes in density, with color scale clipped at ±40 kg/m³; (b-d) the time-lapse density changes recovered by inversion schemes I, II, and III, respectively.

with the real time-lapse signals due to reservoir changes, and it generates model perturbations in regions, where there are no time-lapse changes. As a result, the monitor inversion is trying to further recover the background model, as well as finding the reservoir changes. In scheme III, the common data residuals are subtracted out in the cost function in equation 2. This can be shown by taking the first derivative of equation 2 with respect to $\mathbf{u}_{\text{monitor}}$

$$\frac{\partial E(\mathbf{m})}{\partial \mathbf{u}_{\text{monitor}}} = (\mathbf{u}_{\text{monitor}} - \mathbf{d}_{\text{monitor}}) - (\mathbf{u}_{\text{baseline}} - \mathbf{d}_{\text{baseline}})$$
$$= \mathbf{R}_{\text{monitor}} - \mathbf{R}_{\text{baseline}}. \qquad (4)$$

From this, we see that when $\frac{\partial E(\mathbf{m})}{\partial \mathbf{u}_{\text{monitor}}} = 0$, the cost function $E(\mathbf{m})$ reaches its minimum, where the baseline and monitor data residuals are equal. As a result, the inverted changes are free of background structures because the residuals cancel out perfectly, at least in theory. Similar discussions can be found in Asnaashari et al. (2015).

It is arguable that in scheme II, subtraction of the two models does not provide a fair comparison because the monitor inversion has updated the model for more iterations than the baseline. A more promising approach would be to start from the baseline result and update the baseline and monitor model for the same number of iterations, which, in fact, is equivalent to scheme I with a better starting model. In practice, however, it is difficult to know in advance what starting model is required to achieve a particular level of inversion accuracy. It is clear that the issues with scheme I described above can never be fully eliminated, but they can be mitigated with increasingly better starting models. Similar but more advanced ideas are presented by Maharramov and Biondi (2014) and Yang et al. (2014), both of whom use cross-updating strategies to improve the baseline model and resolve the time-lapse changes at the same time. We do not include such methods within the scope of this study.

BASELINE MODEL DEPENDENCE

With the acoustic synthetic example with perfectly repeated surveys, we observed that DDWI delivers cleaner and better inversion results, at least for the time-lapse P-wave velocity changes. In this section, we investigate how the quality of the baseline models affects the performance of DDWI because the inverted baseline model is a prerequisite for DDWI. We know that DDWI does not resolve
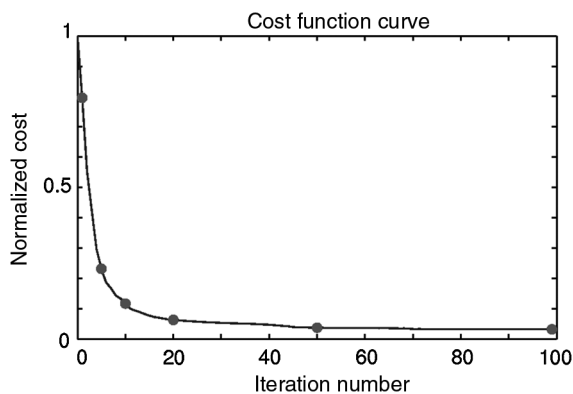
the ambiguity between velocity and density, so we do not present any density results for this test and for all subsequent tests with non-repeatability issues.

Because the baseline model is not further updated in DDWI, the dependence of DDWI on the baseline model accuracy must be investigated to help to decide if the baseline model is reliable enough in practice. As shown in Figure 11, we selected six baseline models corresponding to different convergence levels along the cost function curve for the FWI of the baseline data. With more iterations, the model improves as the predicted data get closer to the recorded data. For each of the baseline models, we generate a synthetic data set $\mathbf{d}_{\text{syn}}$ (equation 3), and run DDWI to invert for the P-wave velocity changes. Figure 12 shows the resulting inverted baseline models. For each of these DDWI applications, the program was stopped when the same number of iterations was reached. The corresponding velocity changes resolved by DDWI starting from each one of the baseline models are shown in Figure 13. Note the different color scale for each subfigure.

It is clear that DDWI gives an improved result with a better baseline model. In Figure 12b, the baseline model has the correct water depth, but is far from the true velocity model. In this case, DDWI fails to invert for the changes correctly, as shown in Figure 13b. Some of the changes have the wrong sign, and the overall recovery in amplitude is poor (by ±15 m/s or more). This is because the baseline model controls the kinematics when the data differences are back-projected. With a poor baseline model, DDWI will project the data differences into the wrong locations, where the signals cannot be correctly stacked. With a better background velocity model, such as those in Figure 12c and 12d, DDWI is able to invert for the velocity changes with the correct sign (Figure 13c and 13d), although the side lobes remain. From Figure 13e–13f, the amplitude recovery improves as the baseline model contains more and more details that better match the reflections and scattering when the data differences are back-propagated.

From these observations, we see that the performance of DDWI certainly depends on the quality of the baseline velocity model. With a good background model such as a smooth migration velocity model without the details, we expect DDWI to be able to invert for

Figure 11. The cost function curve of the baseline inversion. Dots indicate the selected number of iterations: 1, 5, 10, 20, 50, and 99.
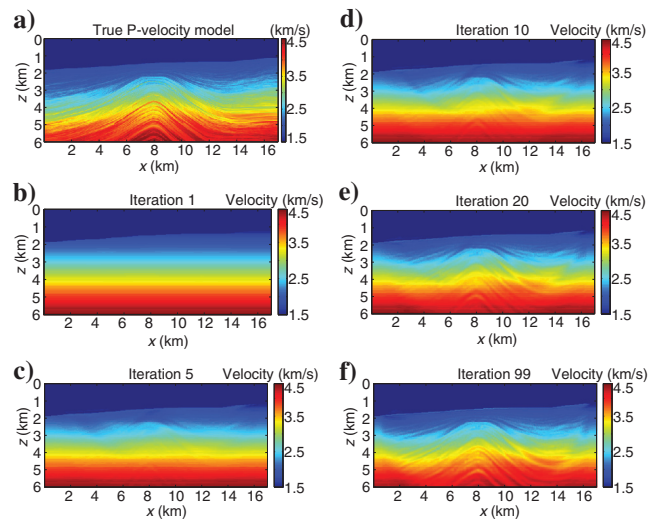

Figure 12. (a) The true baseline P-wave velocity model; (b-f) the baseline P-wave velocity models after 1, 5, 10, 20, and 99 iterations.

at least the approximate locations of the velocity changes. Including more details in the baseline model improves the recovery of the amplitude changes. In practice, when we need to make a judgment on the accuracy of the baseline model is, seismic measures (e.g., gather flatness and data fitting) and prior information (well logs and interpretations) should be used as constraints (Hu et al., 2009; Asnaashari et al., 2013). Within a wide range of convergence levels for baseline inversion (Figure 11), DDWI is robust and capable of delivering reliable results.

SURVEY NONREPEATABILITY

Survey repeatability is a common issue in time-lapse seismic analysis. The successful acquisition of individual surveys does not guarantee quality time-lapse signals. Small deviations between acquisitions can cause significant signal differences between data sets. Because the seismic time-lapse response to reservoir changes is relatively subtle, the true time-lapse signals are easily overwhelmed by data differences caused by survey nonrepeatability. When baseline and monitor surveys are slightly different, schemes I and II will continue to have the same drawbacks as we discussed with perfectly repeated surveys. However, these schemes would not be strongly affected by nonrepeatability because the monitor and baseline data sets are used independently, and subtraction is performed only in the model domain. In contrast, DDWI requires data differencing, which appears to be sensitive to small deviations between surveys. Given the fact that survey discrepancies are inevitable in reality, it is worthwhile to investigate whether DDWI is robust enough to handle realistic nonrepeatability issues before being applied to real data. In this section, we discuss the major causes of nonrepeatable noise, including random noise, source and receiver positioning errors, and source wavelet discrepancies and overburden velocity changes (modeled as static shifts) and their impacts on the performance of DDWI. The

objective here is not to prove that DDWI is better than the other methods, but to demonstrate what to expect with DDWI in different scenarios.

Random noise

Although the sources of random noise vary within and across surveys, they can be effectively characterized by a random distribution of power spectrum and phase. In our study, we investigate the impact of random noise on the baseline inversion and DDWI, so we use a noise model with a uniformly distributed power spectrum and random phase in the frequency-wavenumber domain, and then transform these quantities into the time-offset domain to generate the noise for each shot gather. Therefore, the noise does not change abruptly from trace to trace, but it exhibits a certain spatial correlation. Such an approach also more strongly influences portions of the spectrum in which signals are weak, such as the low-frequency signals that are known to be crucial for obtaining reliable velocities from FWI (Virieux and Operto, 2009). This is illustrated in Figure 14, which shows the power spectra of the combined signal plus noise for the six noise levels we investigated. The same type of noise is added to baseline and monitor data sets. The black dotted curve shows the power spectrum of the clean trace. The colored curves from red to black show the spectra of the noisy traces with different levels of noise contamination. To quantify the noise level, we use the overall energy ratio between the noisy and clean signals of one entire shot gather:

$$r = \frac{\sum_{i=1}^{l} n_i{}^2}{\sum_{i=1}^{l} s_i{}^2} * 100\%, \tag{5}$$

where $n_i{}^2$ is the noise energy of the ith trace, $s_i{}^2$ is the signal energy of the ith trace, and l is the number of traces in a gather. Figure 15a shows a sample trace, and in Figure 15b, the data difference between the noisy monitor and baseline surveys for the same trace is plotted together with the clean data difference trace. Due to the weak reservoir response, the ratio between the noise and the real time-lapse signal is extremely large, even for only 1% noise energy.
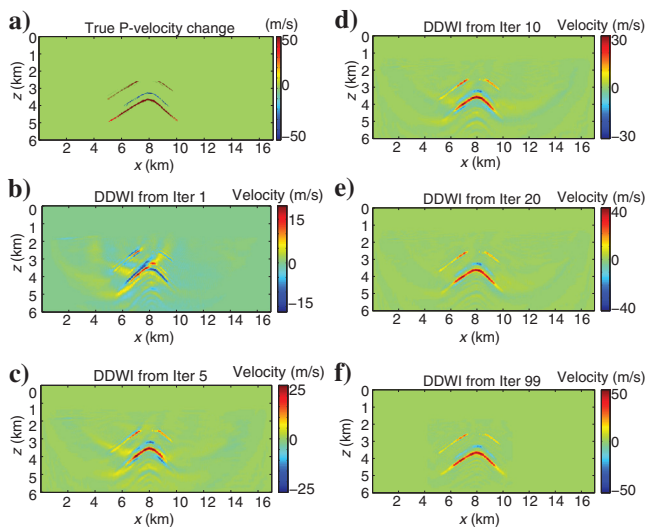

Figure 13. P-wave velocity changes obtained with incorrect baseline velocity models. (a) The true time-lapse changes in P-wave velocity, with the color scale clipped at ±50 m/s and (b-f) the recovered time-lapse P-wave velocity changes by DDWI starting from the baseline models shown in Figure 12b–12f, respectively. Recovery of the velocity changes is clearly improved with better starting baseline models. Note that panel (a) is identical to Figure 10a.
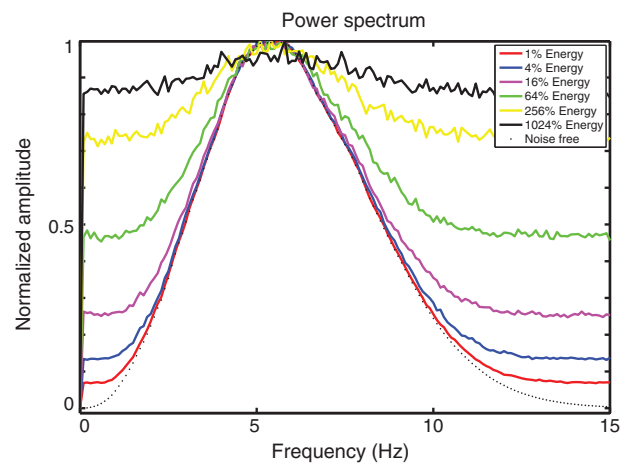

Figure 14. Normalized power spectra of a sample trace with different noise contamination levels. The random noise spectrum obeys a uniform distribution from 0 to 15 Hz. Six noise levels are tested.

Figure 16a–16f shows the baseline FWI P-wave velocity results for each noise level shown in Figure 14. All the inversion results capture the dominant structural features, as well as the fine stratigraphic layers, but they are contaminated by random noise in proportion to the noise level. Although the noise and the primary waves have the same energy level in Figure 16f, FWI still gives a reasonable result due to the strong stacking power of the coherent signals. Figure 17a–17d shows the DDWI results starting from the corresponding baseline FWI results in Figure 16a–16d for each noise level. Only four cases are included because at more than 64% noise energy, the reservoir changes cannot be identified from the image. Below this noise level, DDWI is able to deliver reasonable results, in which reservoir changes are clearly distinguishable even with relatively high noise levels.

We attribute the success of baseline FWI and DDWI in the presence of noise to the coherency in seismic data and their constructive interference during wave propagation through a good velocity model. In FWI, as the data are injected into the model, most of the random noise cancels during propagation, whereas the real signals constructively interfere and produce coherent model updates. This also explains the spotty pattern of the noisy structures in the results in Figures 16 and 17. The random noise we use here is not completely random in space (across traces) because it is generated with a spatial correlation, as is often observed in reality (e.g., the scattering generated by subwavelength inhomogeneities in the subsurface, and acoustic noise from the ocean). Such signals are noise, not signal, and have random properties, but they are not completely random at any scale. Their spatial correlation leads to coherent stacking in the model space to some extent, which will contaminate the final results. In the cases we tested, DDWI is very robust to random noise even when the real signal is not directly observable from the data difference, which can be attributed to the fact that most of the noise energy cancels out due to randomness. The remnant coherent stacking in space generates more noise in the model as the signal-to-noise ratio severely decreases.

Source and receiver positioning error

If the baseline and monitor surveys use the same acquisition geometry, it is straightforward to apply DDWI because the two data vintages can be differenced trace by trace. However, even with advanced GPS guidance during acquisition, the positioning of sources and receivers still contains errors. Here, we have to clarify the difference between inaccurate positioning and nonrepeated positioning. If sources and receivers are placed at different locations between surveys, but the positions are well measured, we can interpolate two surveys to the same grids given enough sampling. However, the positions are not always accurately measured, and interpolation cannot correct this type of error. A small positioning error is expected in a well-repeated monitor survey (Beasley et al., 1997; Yang et al., 2013). Nonetheless, small positioning deviations can generate huge

Figure 16. Baseline models obtained by FWI on noisy data. (a-f) The baseline P-wave velocity models recovered by FWI starting from the same layered model shown in Figure 6a and using progressively higher noise levels from 1% to 1024%, respectively, as depicted in Figure 14. As the noise energy increases, details in the model are more contaminated, but recovery of the dominant structure is very robust to random noise.


Figure 15. (a) A near-offset monitor trace with 1% noise energy. The amplitude of the noise is about the same level as that of the coda waves. (b) Difference between the noise-free monitor and baseline traces (red) and between the noisy monitor and baseline traces (blue). Note that small waveform changes shown in the red trace between approximately 3 and 5 s are obscured by noise in the blue trace.
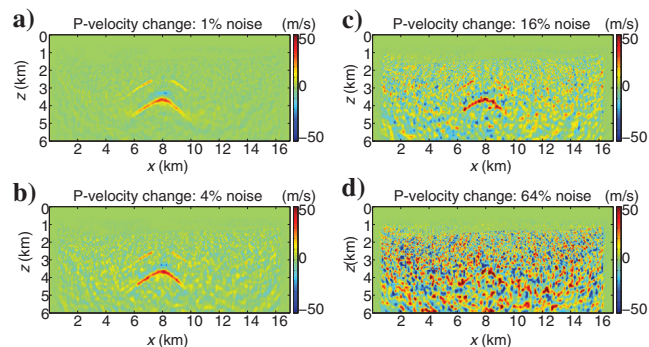

Figure 17. P-wave velocity changes obtained with noisy data. (a-d) The recovered time-lapse P-wave velocity changes obtained from DDWI starting from the baseline models shown in Figure 16a–16d, respectively. Note that in panel (d), it is difficult to identify reservoir changes reliably at a noise energy level of 64%.

data differences over the entire data set; these differences are normally stronger in amplitude than the real 4D signals. In this section, we focus on the impact of this type of error on the performance of DDWI. We assume the baseline survey positioning is known, and use a perturbed survey geometry to generate monitor data. The resulting data differences are input directly into DDWI without correcting for these positioning errors. Due to the inherent limitations of our simplified finite-difference modeling, we are only able to perturb the source and receiver positions by an integral number of grid points (e.g., for a grid spacing of 6.25 m). Two types of perturbations are studied: random perturbations and systematic perturbations. Both types are applied to sources and receivers either simultaneously or individually.

For random source and receiver perturbations, we generate a random sequence of numbers with zero mean that determines if the source or receiver position is perturbed by one grid point to the right or to the left. Figure 18 shows the P-wave velocity changes resolved by DDWI with random perturbations in source positions only, receiver positions only, and combined source and receiver positions for the monitor survey. Despite mild contamination in the background and at the seafloor, the reservoir changes are recovered with an acceptable quality when compared with the clean-data case in Figure 9d. The receiver-only perturbation case (Figure 18c) appears cleaner than that of the source-only perturbation (Figure 18b) because the number of receivers (680) is more than 10 times the number of sources (64); hence, artifacts induced by positioning errors are better canceled out in Figure 18c. Although not attempted here, we expect that including more shots will improve the image in Figure 18b. In Figure 18d, where sources and receivers are perturbed, the artifacts from DDWI show the combined effects of those in Figure 18b and 18c.

For systematic source perturbations, we divide the 64 sources into groups, and perturb each group by the same shift. Figure 19 shows the DDWI results with source positions perturbed in groups, with the 64 evenly spaced sources numbered sequentially from left to right along the top of the model. In Figure 19b, all sources in the monitor survey are shifted one grid point to the right. The recovered velocity changes are slightly weaker in amplitude than those in Figure 18b, but the image is relatively clean. In Figure 19c, sources 1–

32 are perturbed one grid point to the right, and sources 33–64 are perturbed one grid point to the left. In Figure 19d, sources 1–16 and sources 33–48 are shifted one grid point to the right, and sources 17–32 and 49–64 are shifted one grid point to the left. Similar velocity changes are recovered for the last two cases (Figure 19c and Figure 19d), and background artifacts follow the true structures, although they are weaker in amplitude compared with the resolved reservoir changes. The pattern of these artifacts reflects the number of source groups that were perturbed. In particular, the amplitude polarity flip of the artifacts at the seafloor is directly correlated with the position of the shifted sources. As in the random perturbation case, the results for systematically perturbing the receivers are similar to those for source perturbations, although the artifacts are smaller due to more effective stacking.

In practice, source and receiver position errors due to limited GPS accuracy (approximately 1 m) and streamer feathering effects can be larger than what we have tested here (±6.25 m). However, the mismatch between postprocessed baseline and monitor surveys can be reduced to a much lower level by data binning, interpolation, and regularization (Lumley, 2001; Rushmere et al., 2010). In highly repeatable acquisitions, such as ocean-bottom-cable systems, the source positioning mismatch in the raw data can be even smaller than 6.25 m (Beasley et al., 1997; Yang et al., 2013). In practice, errors are likely to arise from a combined effect of systematic and random perturbations. From all the tests above, it is expected that DDWI will be able to deliver good results with mild source and receiver positioning discrepancies. To some extent, the randomness of this error helps to mitigate the artificial patterns seen in the inversion results.

Source wavelet discrepancy

Source wavelets are likely to be different between surveys in real acquisitions. The acquisition conditions (e.g., air-gun types) and initial data processing can introduce discrepancies in source wavelets. These errors are commonly minimized by coprocessing the baseline and monitor data sets. The source wavelet can be shaped by applying matched filtering in the cross-equalization process (Lumley, 2001). However, after all such optimization steps are applied, the resulting wavelets are still likely to have small discrepancies (e.g., a
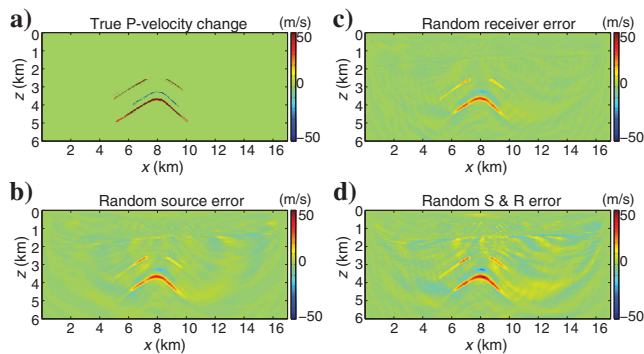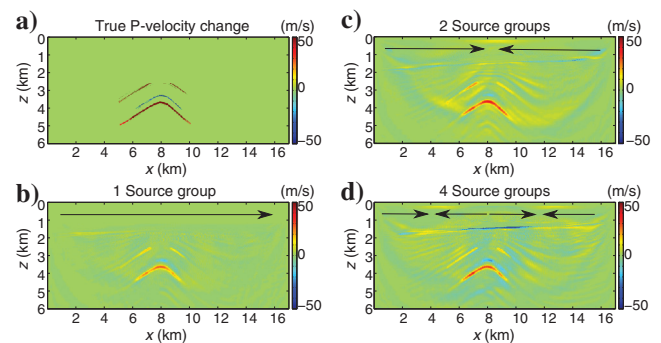

Figure 18. P-wave velocity changes obtained using monitor data with random source and receiver positioning errors. (a) The true time-lapse changes in P-wave velocity, with color scale clipped at ±50 m/s; P-wave velocity changes resolved by DDWI are shown for randomly perturbed (b) source, (c) receiver, and (d) combined source and receiver positions in the monitor survey. Note that panel (a) is identical to Figure 9a.

Figure 19. Effects of systematic shifts in source positions. (a) The true time-lapse changes in P-wave velocity, with the color scale clipped at ±50 m/s. (b-d) P-wave velocity changes resolved by DDWI with the source positions systematically perturbed in the monitor survey in divisions of 1, 2, and 4 groups, respectively. Each group of sources is shifted one grid position (6.25 m) in the same direction either to the right or the left, as shown by the black arrows.

few degrees of phase rotation). In this section, we focus on the impact of phase differences between baseline and monitor source wavelets on the performance of DDWI.

We use a standard zero-phase Ricker wavelet for the baseline survey, and phase-rotated Ricker wavelets for the monitor surveys. Figure 20 shows how the wavelet is distorted gradually with more and more phase rotation. It is apparent that a small discrepancy between source wavelets will cause significant data differences across the entire survey. To simulate the situation, in which we cannot further shape the wavelets by further processing, we obtain the field data difference $\mathbf{d}_{baseline} - \mathbf{d}_{monitor}$ by directly subtracting the baseline and monitor data sets containing the residual phase differences. Synthetic data sets $\mathbf{u}_{baseline}$ and $\mathbf{u}_{monitor}$ are simulated with the same standard zero-phase Ricker wavelet.

Figure 21 shows all the DDWI results with increasing levels of phase rotation in the monitor source wavelet. The inverted P-wave velocity changes are as accurate as those of previous inversions for all the cases tested in terms of location, shape, and amplitude. However, the trend that larger phase rotations give rise to stronger artifacts in the model is also clearly observed. Up to 10°, reservoir changes can be easily distinguished from the incorrectly determined background structures. With larger phase rotations, however, source wavelets in the monitor survey are markedly shifted from the baseline wavelet, and the corresponding data differences are large enough to produce significant model changes that overwhelm the real changes. In practice, a phase difference of less than 10° is generally achievable, in which case DDWI appears to be robust.

It is important to point out that all the inversions in this section are masked (i.e., no water layer was involved). Even for a small phase rotation (e.g., 2°), DDWI cannot converge when the entire model is included in the inversion. This is because the dominant signal, i.e., the major contributor in the L2-norm cost function of equation 2, is the direct arrival. A trivial phase rotation in the source wavelet will generate huge data differences, especially for the dominant phases (e.g., direct waves and water bottom reflections). These data differences are nonphysical, and they cannot be explained by the wave equation without attenuation. For example, a delayed direct arrival indicates a decrease in water velocity between the source and receiver; however, the phase-rotation-

induced delay is frequency-dependent, which means that only a dispersive velocity can explain the traveltime delay. In addition, the data differences are not random enough to cancel each other. As a result, DDWI is not able to find a perturbation in the shallow part of the model (i.e., the water layer) that makes the cost function decrease. When the model is masked, these data differences are not activated in the cost function, allowing DDWI to focus on the reservoir responses. This situation is very similar to the target-oriented inversions presented in Zhang and Huang (2013) and Asnaashari et al. (2015).

Overburden velocity changes

In terms of the effect of model changes on FWI, areas outside the reservoir may be even more important than those inside. In particular, overburden structures may change between surveys. For example, compaction within the reservoir can change the stress field and velocity above and below it (Smith and Tsvankin, 2013). Water velocity also varies seasonally. All such overburden changes will affect the entire data set and cause data differences unrelated to reservoir changes. In this section, we use water velocity changes to represent this type of survey nonrepeatability.

For the synthetic model in this study, the water depth can reach 2000 m. Deepwater production areas like the Gulf of Mexico have water depths of up to 3000 m (Managi et al., 2005). In such water depths, seismic amplitudes and traveltimes can be perturbed significantly even with small variations in water velocity. Factors that influence water velocity include temperature, salinity, and depth. We adopt Medwin's (1975) equation to describe their relationship:

$$v = 1499.2 + 4.6T - 0.055T^2 + 0.00029T^3 + 1.34$$
$$- 0.01T(S - 35) + 0.016D, \qquad (6)$$

Figure 20. The source wavelet with different levels of phase rotation. The blue curve shows the standard Ricker wavelet, which is used as the baseline survey source wavelet. The phases are rotated by the same number of degrees over all frequencies for each monitor wavelet. The five levels of phase rotation are used: 2°, 5°, 10°, 20°, and 30°.
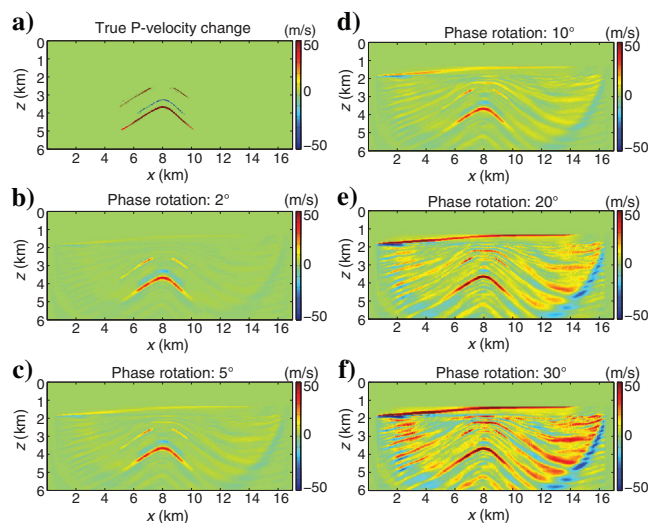

Figure 21. Effects of source wavelet discrepancies between baseline and monitor surveys. (a) The true time-lapse changes in P-wave velocity, with the color scale clipped at ±50 m/s. P-wave velocity changes resolved by DDWI are shown in the remaining plots with the source wavelet in the monitor survey shifted in phase by (b) 2°, (c) 5°, (d) 10°, (e) 20°, and (f) 30°, for all frequencies. Note that panel (a) is identical to Figure 10a.

where T is the temperature (°C), S is the salinity (in parts per thousand, or ppt), D is the water depth (in meter), and v is the velocity (in m/s). For our purpose, it is not necessary to discuss these factors separately. We assume that salinity ($S = 35$ ppt) and sea level stay constant, while temperature changes. Typically, water temperature in the Gulf of Mexico varies at the surface from 30°C in the summer to 15°C in winter, and decreases to 4°C below 1000 m. Water velocity can change by up to 40 m/s at the surface seasonally according to equation 6. From the water surface down to 1000 m, we assume a linear temperature gradient, and compute the water velocity with equation 6. Water velocity is assumed constant ($= 1500$ m/s) less than 1000 m, and acquisition of the baseline survey is assumed to be in winter (with a surface temperature 15°C).

We use 18°C and 30°C for surface water temperatures of two monitor surveys acquired in the spring and the summer, respectively. According to equation 6, the maximum water velocity changes are approximately 9 m/s and 39 m/s, respectively, at the surface. For both cases, we directly difference the monitor and baseline data sets to generate d_{syn}, assuming that no corrections have been made in data processing to account for the water velocity changes. The DDWI results are shown in Figure 22. The reservoir changes are well recovered together with the water velocity changes. As expected, larger water velocity changes generate stronger background noise (Figure 22b) than that for smaller velocity changes (Figure 22a).

In fact, background noise will exist even if the exact water layer velocity model is used in DDWI. We can write the data difference as

$$\delta \mathbf{d} = \mathbf{d}_{\text{monitor}} - \mathbf{d}_{\text{baseline}}$$
$$= \mathbf{G}(\mathbf{m})\delta \mathbf{m}_{\text{water}} + \mathbf{G}(\mathbf{m})\delta \mathbf{m}_{\text{reservoir}} + \ldots, \quad (7)$$

where $\mathbf{G}(\mathbf{m})$ is the sensitivity matrix of the wavefield with respect to a model perturbation at the vicinity of the true model $\mathbf{m}$; $\delta \mathbf{m}_{\text{water}}$ and $\delta \mathbf{m}_{\text{reservoir}}$ are model perturbations in the water layer and reservoir,

respectively. Neglecting higher order terms, the major contributor to the data differences is first-order scattering caused by water velocity and reservoir changes. If we managed to obtain the exact water velocity and reservoir changes, the cost function residual would be

$$(\mathbf{u}_{\text{monitor}} - \mathbf{u}_{\text{baseline}}) - (\mathbf{d}_{\text{monitor}} - \mathbf{d}_{\text{baseline}})$$
$$= (\mathbf{G}(\mathbf{m}_0) - \mathbf{G}(\mathbf{m}))\delta \mathbf{m}_{\text{water}}$$
$$+ (\mathbf{G}(\mathbf{m}_0) - \mathbf{G}(\mathbf{m}))\delta \mathbf{m}_{\text{reservoir}} + \ldots, \quad (8)$$

where $\mathbf{G}(\mathbf{m}_0)$ is the sensitivity matrix based on the inverted baseline model $\mathbf{m}_0$. Because we cannot obtain the exact baseline velocity model $\mathbf{m}$, the data residual does not go to zero even with correct water velocities and reservoir properties. Instead, the data difference in the cost function tries to use perturbations in addition to the correct time-lapse changes to minimize the misfit, generating spurious model updates outside the reservoir that are not from time-lapse effects. From equation 8, we would expect this type of background noise in DDWI even without water velocity changes. However, if the baseline inversion is successful, $\mathbf{G}(\mathbf{m}_0) - \mathbf{G}(\mathbf{m})$ should be much weaker than $\mathbf{G}(\mathbf{m})$ because $\mathbf{G}(\mathbf{m}_0)$ approximates $\mathbf{G}(\mathbf{m})$ after the inversion. As a consequence, the second-order scattering caused by the reservoir changes and imperfect baseline model $((\mathbf{G}(\mathbf{m}_0) - \mathbf{G}(\mathbf{m}))\delta \mathbf{m}_{\text{reservoir}})$, is significantly weaker than that caused by the water velocity changes $(\mathbf{G}(\mathbf{m}_0) - \mathbf{G}(\mathbf{m}))\delta \mathbf{m}_{\text{water}})$, and the first-order scattering from the reservoir $(\mathbf{G}(\mathbf{m})\delta \mathbf{m}_{\text{reservoir}})$. Therefore, the second-order reservoir scattering is not strong enough to contaminate the result. One example of this weak second-order scattering is the case without water velocity changes as presented in Figure 9d. After all, this type of artifact is due to the incorrect back-propagation of data differences in the inaccurate baseline model. The level of the model inaccuracy determines the level of artifacts.

Although the results of the two cases presented here are of good quality and interpretable, DDWI is not able to overcome water velocity differences by itself because, once the data difference is taken, the inversion does not differentiate between the sources of these signal changes. DDWI could be improved if we processed the time-lapse data set carefully with a calibrated water velocity before taking the data difference. After this, DDWI would function as if there were no water velocity changes.

DISCUSSION

As we observed from the mathematical derivations and the synthetic tests, the advantage of DDWI over the other two time-lapse inversion schemes discussed here is that the common data residuals are subtracted out and do not generate background velocity updates that are unrelated to reservoir changes. It is important for interpreters to make decisions based on clean and meaningful images, in which the reservoir information is not contaminated by background noise. However, in practice, data subtraction is intuitively dangerous whenever at least some of the differences between data sets do not originate from the reservoir response. When these nonreservoir signals are included in the cost function, it is reasonable to expect that DDWI will produce artifacts in the inverted images by attempting to fit such data. What we observe, however, from our synthetic study does not obey this intuition. Neither strong random noise nor mild survey nonrepeatability severely harms the performance of DDWI. It is worth clarifying that the mechanism of this
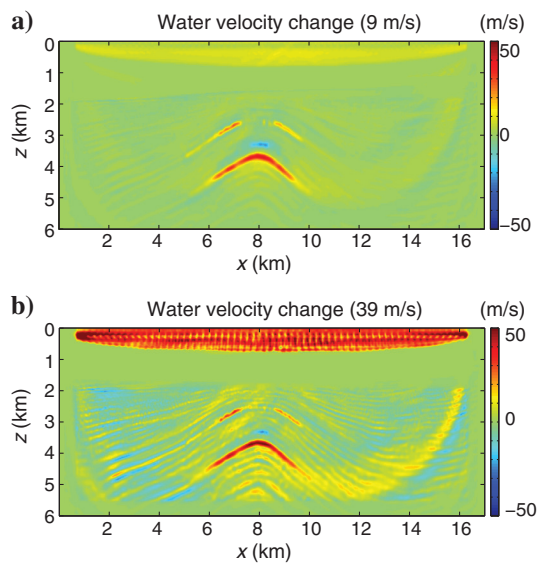

Figure 22. P-wave velocity changes resolved by DDWI with the water velocity in the monitor survey increased by a maximum of (a) 9 m/s and (b) 39 m/s at the surface.

robustness is not only due to DDWI, but also to the merits of FWI itself. Unlike linear imaging methods (e.g., reverse time migration), FWI or DDWI does not directly map all of the data into the model domain. Instead, these methods look for a model perturbation that can explain the data via the wave equation. It is not difficult to understand that random noise has little effect on DDWI because most of the noise energy is stacked out during back-propagation. Data differences caused by survey nonrepeatability are strong and coherent, and, therefore, are not fully stacked out during back-propagation. However, these differences do not produce strong velocity updates because they cannot easily be generated by wave-equation-based velocity perturbations. Although reservoir changes are generally well resolved by DDWI, data differences arising from nonrepeatability effects have relatively small contributions to the final results. Obviously, when the nonrepeatability becomes severe, some of the data differences will lead to spurious model perturbations. The major information we would like to deliver therefore, to achieve a successful time-lapse waveform inversion, baseline and monitor data sets still need to be carefully coprocessed to mitigate nonrepeatability effects before applying DDWI. If any combination of noise effects we have tested in this study applies to the same data set, the aggregated effect will deteriorate the performance of DDWI.

Our use of the term realistic in this study is not intended to address acquisition design, but rather focuses on the nonrepeatability of time-lapse surveys. Although we believe that our 2D examples are sufficient for testing the performance of DDWI, their acquisition geometries differ significantly from those commonly used in practice. For example, 3D time-lapse ocean-bottom-cable (OBC)/ocean-bottom-node (OBN) acquisitions suffer only from the location errors of shots, which are often dense enough for interpolation. The 3D time-lapse streamer data tend to have more serious issues with nonrepeatability because sources and receivers can have errors in their inline and crossline directions. In addition, time-lapse streamer data can have different offset ranges, which would limit the capability of FWI. However, as long as the same offset range is used for baseline and monitor inversions, DDWI can still adequately resolve the time-lapse changes.

We limited our study to acoustic data and P-wave velocity inversion because our goal is to demonstrate the effect of survey nonrepeatability. The real earth is elastic, not acoustic, and often anisotropic and attenuating. In an acoustic earth with local P-velocity or density changes, acoustic DDWI would be effective in identifying fluid property changes that give rise to the changes in seismic properties. When additional physics is added, more sophisticated inversion methodologies would be needed, depending on how strong the time-lapse effects are. S-velocity and anisotropy might change when reservoir compaction occurs, and attenuation might change over time in a producing gas field. Such changes would be reflected as data differences in DDWI; if the inversion engine is acoustic, false changes in P-velocity would likely be created to explain these data differences. The DDWI framework, however, is not limited to acoustic models, and it is straightforward to include more physics in the inversion if needed. The same principle also applies to schemes I and II. The fundamental conclusions from our tests should still hold for more advanced physics.

In this work, we have not included any tests of nonrepeatability for schemes I and II. However, it is not difficult to gain insight into such tests based on our understanding of scheme III results. Because baseline and monitor data sets are used separately for schemes I and

II, they will suffer from contaminated data in the same way as would the inversion of a single vintage of data. Contaminated data would lead to parameter estimation errors in baseline and monitor models, which would introduce spurious model differences. Such differences, however, would be better constrained than those from DDWI. One obvious example would be for sources that are shifted significantly between surveys (e.g., several hundred meters), in which case DDWI will likely fail because the cost function in equation 2 becomes inappropriate. Schemes I and II might still work properly in this case because both vintages cover identical areas, and subtraction only happens between final models after inversion takes place (although the previously discussed issue of background model errors still exists). The advantages of DDWI certainly diminish as the level of nonrepeatable noise increases, although the degradation is gradual. In practice, it is safer to try all three schemes and compare them side-by-side if the level of survey nonrepeatability is uncertain. The results presented here can be used as a reference to help to identify artifacts and choose the method that yields the most reliable result.

What we mean by "realistic" in this study is not about the total acquisition, but we only focused on the nonrepeatability of time-lapse surveys. Although we think our 2D examples are sufficient for the purpose of testing the performance of DDWI, it is quite far away from a commonly used acquisition in practice. The 3D OBC/OBN acquisitions would only have location errors with shots, which are often dense enough for interpolation. The 3D streamer data tend to have more issues because source and receiver might have errors in inline and crossline directions. In addition, time-lapse streamer data could have different offset ranges that would limit the capability of FWI. However, as long as the same offsets are chosen for the baseline and monitor, DDWI can still resolve the time-lapse changes.

We limited our study to acoustic data and P-wave velocity inversion because we only want to demonstrate the effect of survey nonrepeatabilities. The real earth is always elastic, and often anisotropic and attenuating. If all the other physics stay the same, and only the P-wave velocity or density is locally changed, an acoustic DDWI would be effective such as in reservoir fluid substitution scenarios. When more physics are changed over time, more sophisticated inversions would be needed depending on how strong these changes are S-wave velocity and anisotropy might change when reservoir compaction occurs. Attenuation effect might change when a gas field is in production. These changes would be reflected in the data differences in DDWI. If the inversion is still acoustic, false changes in P-wave velocity would likely be created to explain these data differences. However, the DDWI frame work is not limited to acoustic and it is straightforward to include more physics in the inversion. The same principle also applies to schemes I and II. The conclusions derived from our tests still hold for more advanced physics.

The message we want to deliver in this paper is not significantly dependent on 2D versus 3D or acoustic versus elastic. The focus here is only on how DDWI is going to perform when the data are not perfectly repeated because all the researches before this assumed identical surveys, which is impossible. We agree that 3D simulations will make the tests closer to reality, but we also think the 2D tests are sufficient enough for the purpose of this research.

What we have not included in this work is to test all the nonrepeatability issues on schemes I and II. However, with all the under-

standings, we obtained from the existing tests, it is not hard to gain some insights for them. Because data sets are used separately, schemes I and II will suffer from the contaminations for the same reason as when surveys are perfectly repeated. In addition, deviated surveys would likely lead to deviated parameter estimations, which introduce spurious model differences. However, these deviations are better constrained than those in DDWI. One obvious example would be when the sources between the two surveys are shifted significantly (e.g., several hundred meters), DDWI is definitely going to fail because the data set subtraction is not valid. Schemes I and II might still work properly because two surveys cover similar areas, and the subtraction only happens between final models, although the old issue of deviated estimations of the background model still exists. The advantage of DDWI will certainly degrade as the nonrepeatabilities get stronger, and the degradation is gradual. In practice, it is safer to try all the schemes and compare them side by side when the level of survey nonrepeatability is uncertain. What we presented here can be used as a reference to help identify the artifacts and choose a more reliable result from these methods.

CONCLUSION

In summary, our synthetic examples show that DDWI gives better results than conventional inversion schemes by suppressing background model updates when surveys are well repeated. The investigation of nonrepeatable noise shows that within a practical range of data quality (e.g., a few degrees of phase rotation, a few meters of positioning error, etc.), DDWI is robust enough to give a reliable estimate of time-lapse P-wave velocity changes within the reservoir.

ACKNOWLEDGMENTS

REFERENCES

Arts, R., O. Eiken, A. Chadwick, P. Zweigel, L. van der Meer, and B. Zinszner, 2004, Monitoring of CO$_2$ injected at Sleipner using time-lapse seismic data: Energy, **29**, 1383–1392, doi: 10.1016/j.energy.2004.03.072.

Asnaashari, A., R. Brossier, S. Garambois, F. Audebert, P. Thore, and J. Virieux, 2011, Sensitivity analysis of time-lapse images obtained by differential waveform inversion with respect to reference model: 81st Annual International Meeting, SEG, Expanded Abstracts, 2482–2486.

Asnaashari, A., R. Brossier, S. Garambois, F. Audebert, P. Thore, and J. Virieux, 2012, Time-lapse imaging using regularized FWI: A robustness study: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: 10.1190/segam2012-0699.1.

Asnaashari, A., R. Brossier, S. Garambois, F. Audebert, P. Thore, and J. Virieux, 2013, Regularized seismic full waveform inversion with prior model information: Geophysics, **78**, no. 2, R25–R36, doi: 10.1190/geo2012-0104.1.

Asnaashari, A., R. Brossier, S. Garambois, F. Audebert, P. Thore, and J. Virieux, 2015, Time-lapse seismic imaging using regularized full-waveform inversion with a prior model: Which strategy?: Geophysical Prospecting, **63**, 78–98, doi: 10.1111/1365-2478.12176.

Beasley, C., R. Chambers, R. Workman, K. Craft, and L. Meister, 1997, Repeatability of 3-D ocean bottom cable seismic surveys: The Leading Edge, **16**, 1281–1286, doi: 10.1190/1.1437782.

Campbell, S., C. Lacombe, R. Brooymans, H. Hoeber, and S. White, 2011, Foinaven: 4D processing comes up trumps: The Leading Edge, **30**, 1034–1040, doi: 10.1190/1.3640527.

Denli, H., and L. Huang, 2009, Double-difference elastic waveform tomography in the time domain: 79th Annual International Meeting, SEG, Expanded Abstracts, 2302–2306.

Hu, W., A. Abubakar, and T. M. Habashy, 2009, Simultaneous multifrequency inversion of full-waveform seismic data: Geophysics, **74**, no. 2, R1–R14, doi: 10.1190/1.3073002.

Landrø, M., 2001, Discrimination between pressure and fluid saturation changes from time-lapse seismic data: Geophysics, **66**, 836–844, doi: 10.1190/1.1444973.

Lumley, D., 2010, 4D seismic monitoring of sequestration: The Leading Edge, **29**, 150–155, doi: 10.1190/1.3304817.

Lumley, D. E., 2001, Time-lapse seismic reservoir monitoring: Geophysics, **66**, 50–53, doi: 10.1190/1.1444921.

Maharramov, M., and B. Biondi, 2014, Joint full-waveform inversion of time-lapse seismic data sets: 84th Annual International Meeting, SEG, Expanded Abstracts, 954–959.

Managi, S., J. J. Opaluch, D. Jin, and T. A. Grigalunas, 2005, Technological change and petroleum exploration in the Gulf of Mexico: Energy Policy, **33**, 619–632, doi: 10.1016/j.enpol.2003.09.007.

Medwin, H., 1975, Speed of sound in water: A simple equation for realistic parameters: Journal of the Acoustical Society of America, **58**, 1318–1319, doi: 10.1121/1.380790.

Onishi, K., T. Ueyama, T. Matsuoka, D. Nobuoka, H. Saito, H. Azuma, and Z. Xue, 2009, Application of crosswell seismic tomography using difference analysis with data normalization to monitor CO$_2$ flooding in an aquifer: International Journal of Greenhouse Gas Control, **3**, 311–321, doi: 10.1016/j.ijggc.2008.08.003.

Prieux, V., R. Brossier, S. Operto, and J. Virieux, 2013, Multiparameter full waveform inversion of multicomponent ocean-bottom-cable data from the Valhall field. Part 1: Imaging compressional wave speed, density, and attenuation: Geophysical Journal International, **194**, 1640–1664, doi: 10.1093/gji/ggt177.

Raknes, E., and B. Arntsen, 2014, Time-lapse full-waveform inversion of limited-offset seismic data using a local migration regularization: Geophysics, **79**, no. 3, WA117–WA128, doi: 10.1190/geo2013-0369.1.

Raknes, E. B., W. Weibull, and B. Arntsen, 2013, Time-lapse full waveform inversion: Synthetic and real data examples: 83rd Annual International Meeting, SEG, Expanded Abstracts, 944–948.

Routh, P., G. Palacharla, I. Chikichev, and S. Lazaratos, 2012, Full wavefield inversion of time-lapse data for improved imaging and reservoir characterization: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: 10.1190/segam2012-1043.1.

Rushmere, E., M. Dyce, S. Campbell, and A. Hill, 2010, 4D acquisition and processing: A North Sea case study on the relative contributions to improved 4D repeatability: Geological Society of London Petroleum Geology Conference Series, **7**, 517–522, doi: 10.1144/0070517.

Sears, T., P. Barton, and S. Singh, 2010, Elastic full waveform inversion of multicomponent ocean-bottom cable seismic data: Application to Alba field, U.K. North Sea: Geophysics, **75**, no. 6, R109–R119, doi: 10.1190/1.3484097.

Smith, S., and I. Tsvankin, 2013, Sensitivity of compaction-induced multicomponent seismic time shifts to variations in reservoir properties: Geophysics, **78**, no. 5, T151–T163, doi: 10.1190/geo2012-0361.1.

Tarantola, A., 1984, Inversion of seismic reflection data in the acoustic approximation: Geophysics, **49**, 1259–1266, doi: 10.1190/1.1441754.

Tarantola, A., 1986, A strategy for nonlinear elastic inversion of seismic reflection data: Geophysics, **51**, 1893–1903, doi: 10.1190/1.1442046.

Trani, M., R. Arts, O. Leeuwenburgh, and J. Brouwer, 2011, Estimation of changes in saturation and pressure from 4D seismic AVO and time-shift analysis: Geophysics, **76**, no. 2, C1–C17, doi: 10.1190/1.3549756.

Virieux, J., and S. Operto, 2009, An overview of full-waveform inversion in exploration geophysics: Geophysics, **74**, no. 6, WCC1–WCC26, doi: 10.1190/1.3238367.

Warner, M., A. Ratcliffe, T. Nangoo, J. Morgan, A. Umpleby, N. Shah, V. Vinje, I. Štekl, L. Guasch, C. Win, G. Conroy, and A. Bertrand, 2013, Anisotropic 3D full-waveform inversion: Geophysics, **78**, no. 2, R59–R80, doi: 10.1190/geo2012-0338.1.

Watanabe, T., S. Shimizu, E. Asakawa, and T. Matsuoka, 2005, Differential waveform tomography for time-lapse crosswell seismic data with application to gas hydrate production monitoring: 75th Annual International Meeting, SEG, Expanded Abstracts, 2323–2326.

Yang, D., M. Fehler, A. Malcolm, F. Liu, and S. Morton, 2013, Double-difference waveform inversion of 4D ocean bottom cable data: Application to Valhall, North Sea: 83rd Annual International Meeting, SEG, Expanded Abstracts, 4966–4970.

Yang, D., A. Malcolm, and M. Fehler, 2014, Time-lapse full waveform inversion and uncertainty analysis with different survey geometries: Presented at 76th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.20141120.

Zhang, Z., and L. Huang, 2013, Double-difference elastic-waveform inversion with prior information for time-lapse monitoring: Geophysics, **78**, no. 6, R259–R273, doi: 10.1190/geo2012-0527.1.

Zheng, Y., P. Barton, and S. Singh, 2011, Strategies for elastic full waveform inversion of time-lapse ocean bottom cable (OBC) seismic data: 81st Annual International Meeting, SEG, Expanded Abstracts, 4195–4200.