Targeted nullspace shuttles for full waveform time-lapse seismic monitoring and CO_2 detection thresholds

Kimberly Pike, Kris Innanen and Scott Keating

ABSTRACT

Time-lapse seismic monitoring is a proven technique in hydrocarbon reservoir monitoring and optimization. Time-lapse monitoring can also be extended to full waveform inversion and applied to the measurement, monitoring, and validation of CO_2 sequestration projects. Traditional time-lapse implementations detect changes in the subsurface through differencing baseline and monitor surveys and relies on the repeatability of baseline and monitor survey geometries. The models obtained through full waveform inversion are non-unique, and therefore absolute inferences about time-lapse changes are difficult to make. Through targeted nullspace shuttling, we investigate an approach to find unique baseline and monitor models which minimize the time-lapse difference and preserve the data-fit, though explicitly navigating the nullspace, providing a minimum bound on CO_2 plume uncertainty. Using synthetic examples, this approach demonstrates the ability of nullspace shuttling to detect minimum CO_2 time-lapse changes in sparse monitoring scenarios that are avoided in traditional time-lapse applications.

INTRODUCTION

Time-lapse seismic surveying is a well-established technique that has been used for many decades in the monitoring and optimization of hydrocarbon production (Greaves and Fulp, 1987). More recently, the time-lapse seismic method has been used in the monitoring, measurement, and validation (MMV) of CO_2 sequestration projects to characterize CO_2 plume migration, plume geometry, and plume containment. Previous authors have shown that time-lapse seismic is effective for monitoring injection of CO_2 into saline aquifers (Arts et al., 2004; Brevik et al., 2000; Chadwick et al., 2010).

To recover robust time-lapse changes using seismic monitoring, it is important that the baseline and monitor seismic surveys be repeated in as identical an approach as possible. This requires the replication of identical source and receiver geometries, and the use of the same source type and source parameters, between baseline and monitor surveys. Conventional time-lapse seismic methods also require careful simultaneous processing of baseline and monitor surveys to enable meaningful time-lapse changes to be recovered. Traditional time-lapse implementations, which replicate the baseline survey geometry, take considerable time and economic investment and may be prohibitive for CO_2 monitoring projects. Sparse time-lapse 3D seismic has been piloted at several CO_2 storage sites, including Ketzin, Germany (Ivandic et al., 2012), and Aquistore, Canada (White et al., 2015), with encouraging results. These sparse implementations rely on reflection seismic differencing to characterize temporal changes in the reservoir.

With recent advances in full waveform inversion (FWI), new techniques have emerged enabling recovery of production/injection related time-lapse information from time-lapse seismic datasets, including non-identical recording geometries. CREWES authors Fu and Innanen (2021) present an excellent summary of several time-lapse FWI methods.

Most time-lapse inversion methods are based on either differencing a baseline and monitor inversion result or inverting the difference between baseline and monitor data sets. These approaches can be effective in simple, noise-free models but are challenging in realistic data sets in which non-repeatable effects such as noise and differences in acquisition geometry are present. In addition, the FWI results of the baseline and monitor inversions are non-unique, and dependent on the initial model and numerical optimization scheme used in the inversion. Therefore, the time-lapse difference is also non-unique. The concept of nullspace shuttles in FWI uncertainty analysis was introduced by Keating and Innanen in 2021, and the subsequent application of targeted nullspace shuttling in time-lapse FWI by Keating and Inannen in 2022. Targeted shuttling can be used to explicitly navigate the FWI nullspace, with respect to a user-defined scalar function of model-space location, to find the minimum time-lapse difference due to CO_2 injection.

In this report, we investigate the nullspace shuttling approach for time-lapse FWI applied to CO_2 injection with varying monitor survey geometries. This approach allows for targeted exploration of the baseline and monitor inversion nullspaces, to minimize the time-lapse difference between models while preserving the data-fit of the FWI results. Multiple time-lapse scenarios will be modeled numerically, to assess the utility of the shuttling approach to monitor CO_2 plumes and define detection thresholds with sparse monitor survey geometry. The geological model and acquisition scenarios considered are based on those implemented at Carbon Management Canada's (CMC) Newell County testing facility, where CO_2 has been injected into the Basal Belly River Sandstone (BBRS). In collaboration with the University of Calgary, a permanent VSP receiver array has been deployed and multiple time-lapse field data sets have been acquired. One of the objectives of CMC is to test and demonstrate accurate, cost-effective measurement and verification technologies. This report will explore the application of sparse FWI monitoring using nullspace shuttling as a plausible, cost-effective, CO_2 plume detection method.

THEORY

An FWI result ideally characterizes the true properties of the subsurface, but is highly dependent on the acquisition geometry, initial model, optimization strategy used in the inversion, and is impacted by noise. The time-lapse inversion difference ideally identifies real temporal changes in the subsurface without identifying any changes due to differences in acquisition geometry between baseline and monitor surveys, or noise in the data.

If the inversion problem for a given data set had only one unique solution, the time-lapse problem could be solved by simply taking the difference between baseline and monitor inversion results. However, the model, m, obtained from an FWI inversion corresponds to the lowest value of the objective function, ϕ, which was achieved in the FWI numerical optimization but not the exact, or global minimum. The solution depends on the specific optimization condition used and is inherently distorted due to noise in the data and approximations in the theory. In addition, measurement errors contribute to differences in baseline and monitor data sets, therefore many possible models fit the data equally well. Another model, m^*, with objective function value ϕ^*, could be an equally plausible output of the inversion. Therefore, the FWI problem has a set of solutions, **m,** containing all possible models m^* for which $\phi^* \leq \phi$, that satisfies the inversion conditions equally. We refer to this set of possible models as the nullspace of the inversion problem,

and the objective function-preserving model-space steps as 'nullspace shuttles' (Deal and Nolet, 1996).

The application of targeted nullspace shuttles to FWI is described in Keating and Innanen (2021). They propose a scalar function which defines the extent to which an anomalous feature in the FWI result is present. The maximal uncertainty in the feature is considered to be the model step, or nullspace shuttle, which minimizes the function while remaining in the nullspace. The nullspace should include all points that achieve an equivalent or lower objection function values than the inversion output.

In time-lapse FWI we have a nullspace for the baseline, and a separate nullspace for the monitor. Figure 1 illustrates, in two dimensions, a simplistic time-lapse FWI approach. The orange ellipse corresponds to the nullspace of the monitor inversion, and the gray ellipse corresponds to the baseline inversion. The inversion results, denoted by points on the nullspace, are determined by the starting models used for each inversion. Hypothetically, any point on either nullspace is an equally possible FWI model solution. The time-difference obtained from FWI inversion is therefore non-unique. To assess the fitness of a time-lapse difference we must consider the choice of fitness with respect to the true difference, of which we have no knowledge. However, through targeted shuttling, we can directly search the nullspace for an 'optimal' or minimum difference between baseline and monitor, which is a unique solution. This minimum difference will not provide the absolute time-lapse change, but it will provide an estimate of the minimum difference that meets the data-fit obtained through FWI.

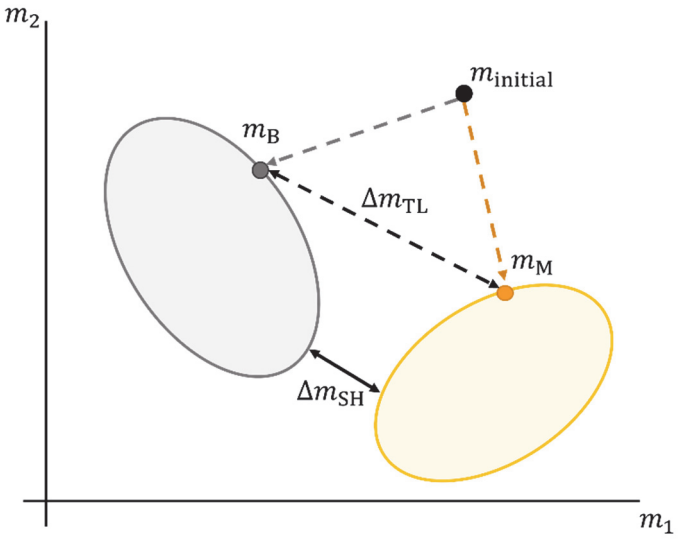

FIG. 1. Conceptual of the nullspace. The baseline FWI model m_{B}, and the monitor model, m_{M}, are a function of the initial model, m_{initial}, and are non-unique; their time-lapse difference, Δm_{TL}, is also non-unique. Through shuttling, we seek, Δm_{SH}, minimum difference between baseline and monitor FWI models.

The method of time-lapse nullspace shuttles is fully described in Keating and Innanen (2022) and summarized as follows. We define a metric ψ, referred to as a hypothesis function. This hypothesis function is small when the difference between baseline and

monitor surveys is also small, and large otherwise. In this context the time-lapse difference can be presented as a shuttling optimization problem in which solve for the minimum difference:

$$\Delta m_{SH} = (m_M + \delta m_M^*) - (m_B + \delta m_B^*) \tag{1}$$

where Δm_{SH} is the optimal time-lapse difference, m_B and m_M are the baseline and monitor FWI results, and δm_B^* and δm_M^* are the optimal shuttled updates given by

$$\begin{aligned}\left(\delta m_{B,}^*, \delta m_M^*\right) = \underset{\delta m_B, \delta m_M}{\operatorname{argmin}} \psi(m_B + \delta m_B, m_M + \delta m_M),\\ \text{subject to } \phi_B(m_B + \delta m_B) < \epsilon_B,\ \phi_M(m_{BM} + \delta m_M) < \epsilon_M,\end{aligned} \tag{2}$$

where $\phi_B$and ϕ_M are the baseline and monitor FWI objective functions, and ϵ_B and ϵ_M are the maximum objective function values, or tolerances, defining the nullspace. We propose a hypothesis function, ψ, which satisfies Equation 2 to find the optimal time-lapse difference, Δm_{SH}, which satisfies the data.

During shuttling we move in towards δm_M^* and δm_B^* by taking maximum steps, while remaining in the nullspace of the baseline and monitor models. Equation 2 is solved iteratively, for baseline and monitor, by solving for the optimal update directions based on local models of ϕ_B and ϕ_M constructed from the gradient (g) and Hessian-vector products ($H\delta m$):

$$\delta m^* = \underset{\delta m^*}{\operatorname{argmin}}\, \psi(m + \alpha(g(m), H(m)\delta m)\delta m), \tag{3}$$

where α is the maximum step expected to stay within the nullspace based on the gradient at the current model and the Hessian-vector product for the current update direction estimate. Once the optimal direction is solved for, we perform a line-search along δm^* to find the true length of the maximum step.

The hypothesis function, ψ, must be defined without knowledge of the true subsurface time-lapse change, yet ideally results in time-lapse changes that are only due to true subsurface changes. The optimal difference should minimize all differences subject to the prior FWI nullspace constraints defined in Equation 2. In the numerical examples presented in this report, we define ψ as the Huber norm of the difference between monitor and baseline models. The Huber norm, which is less sensitive to outliers and strongly convex in the neighbourhood of the minimum (Huber, 1964), allows the L_1 norm to dominate for large time-lapse differences, and the L_2 to dominate for small differences. This approach allows for better convexity and optimization properties as we converge towards the optimal solution (Keating and Innanen, 2022).

The time-lapse formulation described by Equation 2 consists of two shuttling problems, since both baseline and monitor models are updated. By defining ψ to contain both baseline and monitor objective functions the time-lapse problem can be reduced to a single-objective shuttling problem. We can modify the shuttling algorithm described by Keating and Innanen (2021) to be solved sequentially and iteratively. We now have two inner loops - in the first the baseline model is shuttled to minimize the difference from the monitor,

and in the second the monitor is shuttled to match the updated baseline. The outer loop repeats the inner loop optimizations for the specified number of iterations.

METHODS

Synthetic data

In this report we will consider 2D, acoustic, multi-scale, frequency-domain full-waveform inversion. A synthetic true P-wave velocity baseline model was created using well log data acquired at the CMC facility. Multiple sets of compressional sonic and gamma ray logs were spliced to create a log suite from surface to 525 m depth. The sonic data was converted to P-wave velocity and blocked at major velocity changes, guided by lithological boundaries. A sample of the original and blocked log data, zoomed to the BBRS interval, is shown in Figure 2. The blocked P-wave velocity log was then used to populate a horizontally layered 2D true baseline velocity model comprised of 5 m x 5 m grid cells. This 2D model is the true baseline velocity model for this study. Note that the thickness of the BBRS is approximately 10 m, or 2 grid cells. The BBRS is underlain by shales, and overlain by coals, resulting in a local velocity high within the sand. Log blocking was performed to ensure this velocity signature is preserved in the 2D model.

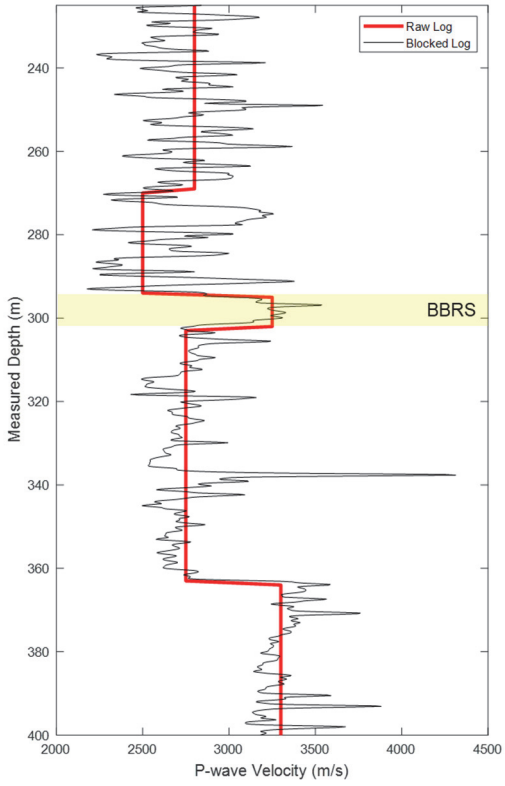

FIG. 2. P-wave velocity log (black) and blocked P-wave velocity log (red), zoomed to show the BBRS and surrounding geological layers.

Macquet et al (2019) examined the behaviour of elastic parameters in the BBRS under simulated CO_2 injection reservoir conditions. Assuming a semi-patchy model and fluid replacement of brine with up to 50% CO_2 saturation, and pore pressure increases of up to 2.7 MPa, they estimated a reduction of P-wave velocity between 20-32%. To simulate a true monitor velocity model the P-wave velocity of the baseline model was reduced by 950 m/s within the BBRS interval. An injection point was assumed to exist in the middle of the model at 500 m offset, and a CO_2 plume extent of 100 m from the injection point. A smooth prior model was generated by applying a gaussian smoother to the true baseline velocity model. The true baseline, monitor, difference, and initial models are shown in Figure 3.

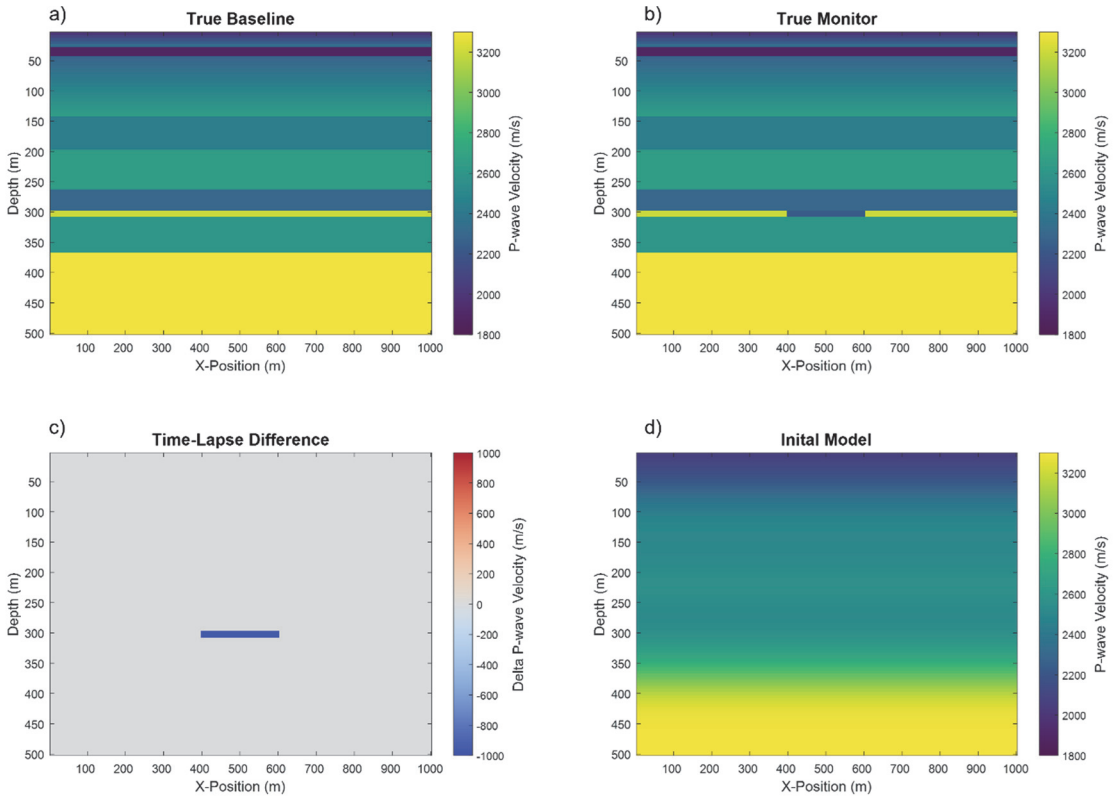

FIG. 3. a) True baseline, b) true monitor, c) time-lapse difference and d) the initial velocity model used for FWI.

FWI and modelling parameters

As part of the frequency domain FWI implementation outlined in Keating et al. (2018), synthetic velocity models are used to generate the true frequency domain shot gather data using finite-difference modelling. A 2D grid of 100 x 200 cells (5 m x 5 m grid cell size dimensions) is used for all models. For all true synthetic models, a VSP acquisition geometry was assumed with a receiver array located at x-offset 440 m. The receiver array extends from surface to a depth of 320 m, with receivers spaced at 5 m intervals. All sources are located at the surface, unless otherwise stated, with varying geometries. A minimum-phase wavelet with a dominant frequency of 25 Hz was used for all models. To simulate

challenges with repeatability, source locations were moved between baseline and monitor models, and noise was added to some of the scenarios. The signal-to-noise ratio (SNR) was varied between 5 and 10, where we define SNR as the ratio between the 2-norm of the data over the 2-norm of the noise.

The same prior model was used for all inversions. For each baseline and monitor inversion ten frequency bands of six evenly spaced frequencies were considered, starting with 1-5 Hz and ending with 1-40 Hz. At each band we perform 10 iterations of L-BFGS optimization. In addition, to ensure a good fit for nullspace shuttling, we perform a final refining update of 20 iterations on a band of 16 frequencies from 1 to 40 Hz.

Nullspace shuttling parameters

To estimate the minimum time-lapse changes between baseline and monitor models we apply targeted nullspace shuttling to minimize the Huber norm of the difference between models. The Huber norm objective function can be stated as follows,

$$\psi_i = \sum \begin{cases} \left(m_M^i - m_B^i\right)^2 & \text{for } \left(m_M^i - m_B^i\right)^2 \le \epsilon \\ 2\epsilon\left|m_M^i - m_B^i\right| - \epsilon^2 & \text{for } \left(m_M^i - m_B^i\right)^2 > \epsilon \end{cases} \tag{4}$$

where ϵ is a user-defined threshold. Starting with the baseline, we shuttle each model in turn by performing 15 iterations of L-BFGS optimization to determine the optimal update direction, followed by a line-search in the calculated direction. We shuttle the baseline and monitor each a total of three times.

NUMERICAL EXAMPLES

In this section we present numerical examples of the time-lapse shuttling approach. Several plausible CCS monitoring acquisition geometry scenarios are evaluated to assess the robustness of nullspace shuttling in recovering a CO_2 plume time-lapse feature. For each example, the baseline and monitor FWI results are presented, along with the time-lapse difference before and after shuttling.

In this first example we consider the ideal case where the baseline and monitor model acquisition geometries exactly match and a source is located in each surface cell (i.e., every 5 m). The baseline, monitor, time-lapse difference, and shuttled time-lapse difference are shown in Figure 4. In this figure, and all subsequent figures, the baseline source locations are shown in red, the monitor source locations are shown in cyan, and the receiver locations are shown in white.

Given the high degree of sampling and repeatability, the CO_2 anomaly can be clearly identified in the simple time-lapse difference. The magnitude of the difference is less than the true difference due to band-limitation of frequencies being recovered. The lateral extent of the time-lapse anomaly is close to that of the true CO_2 plume. The minimum CO_2 plume shown in the shuttled difference is nearly identical in lateral extent that of the standard time-lapse difference, however, all non-injection related changes have been removed, providing more definitive plume geometry.

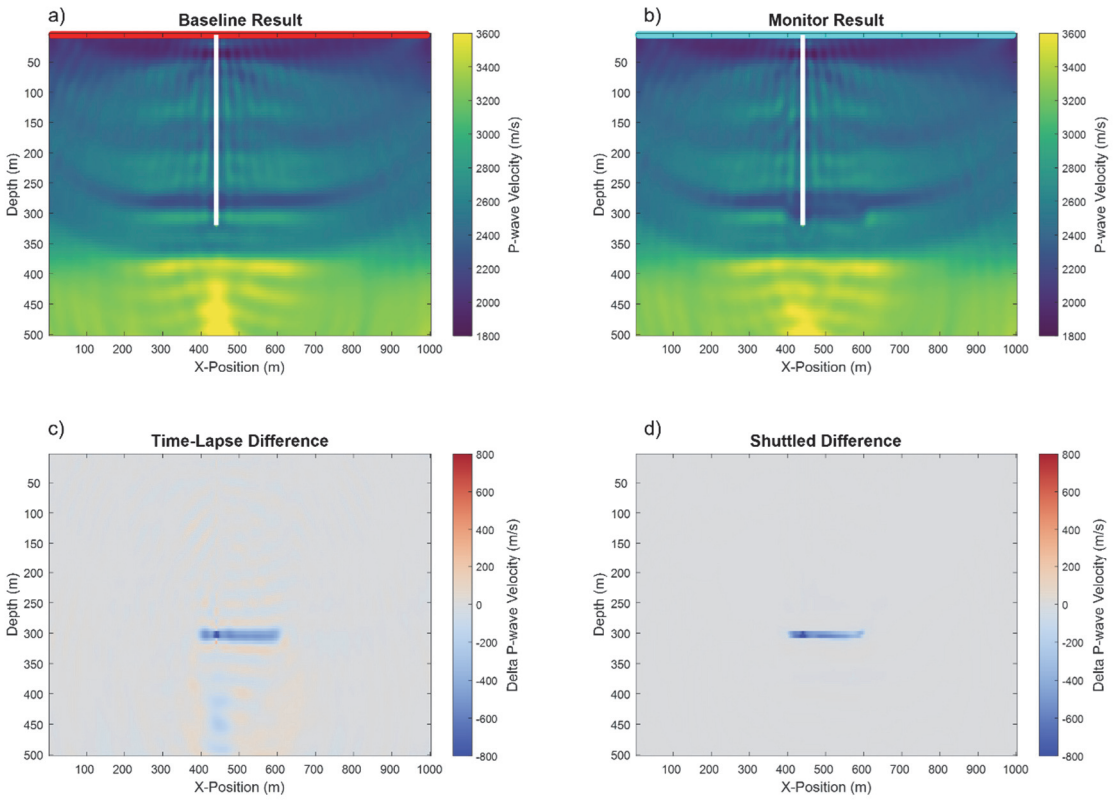

FIG. 4. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of dense, exactly matched acquisition geometry.

In this next example we consider a case where different levels of noise are introduced to the baseline and monitor datasets. As well, the source geometry of the baseline and monitor surveys differ. For this example, and Example 3, the objective is to model a realistic, source-intensive baseline survey and a sparse-source monitor survey to simulate low-cost monitoring scenarios. For the baseline survey, a source interval of 50 m is used; for the monitor survey a source interval of 100 m is used. The FWI outputs, their time-lapse difference and the nullspace shuttled difference are shown in Figure 5. While the CO_2 anomaly can still be identified in the time-lapse difference, additional artifacts, related to noise and acquisition differences are also evident.

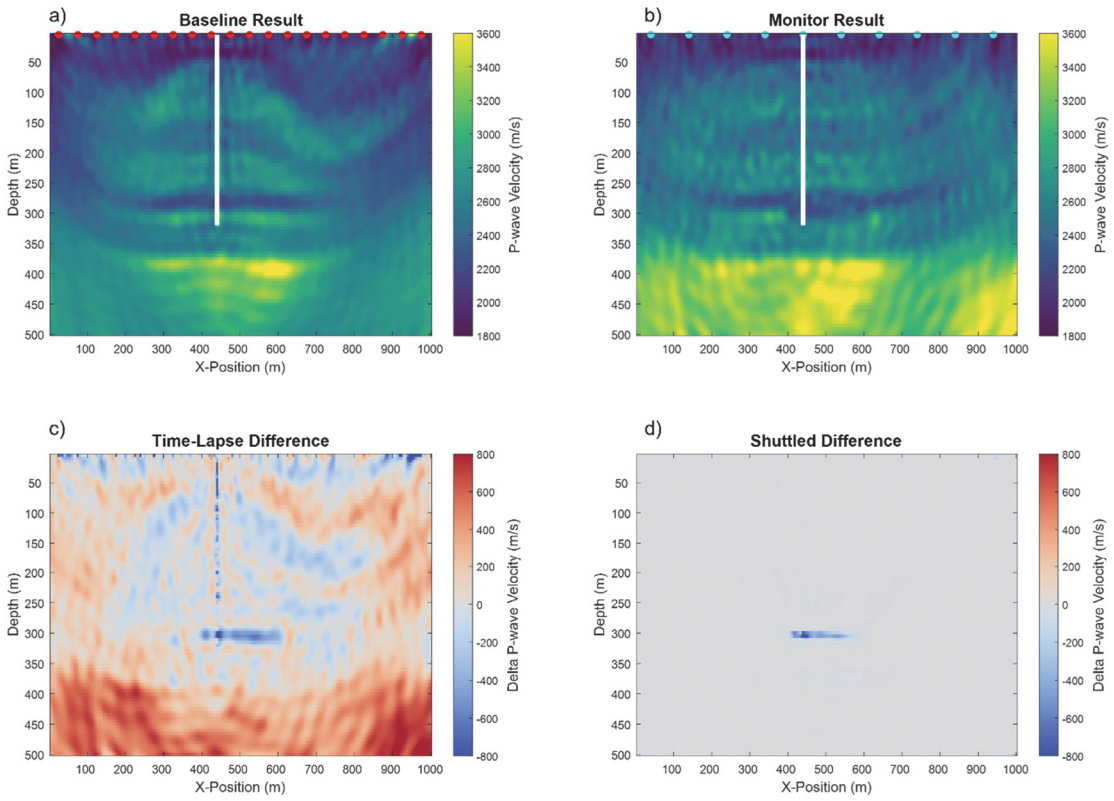

FIG. 5. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of offset unequally spaced source geometry, with added noise.

After nullspace shutting the minimum difference between baseline and monitor surveys only contains energy near the CO_2 plume, all other variations have been removed. As expected, the shuttled anomaly is smaller in extent in comparison with the true plume, since the shuttled difference is the minimum time-lapse difference insisted on by the data, not an absolute difference.

In the next three examples we consider the case of permanent, single source sparse monitor acquisition at various x-positions, with the same dense baseline geometry as implemented in the last example. An SNR of 5 is applied to the monitor cases. First, a single monitor survey source is located at a depth of 10 m and x-position 610 m. The FWI results, time-lapse difference and shuttled difference are show in Figure 6.

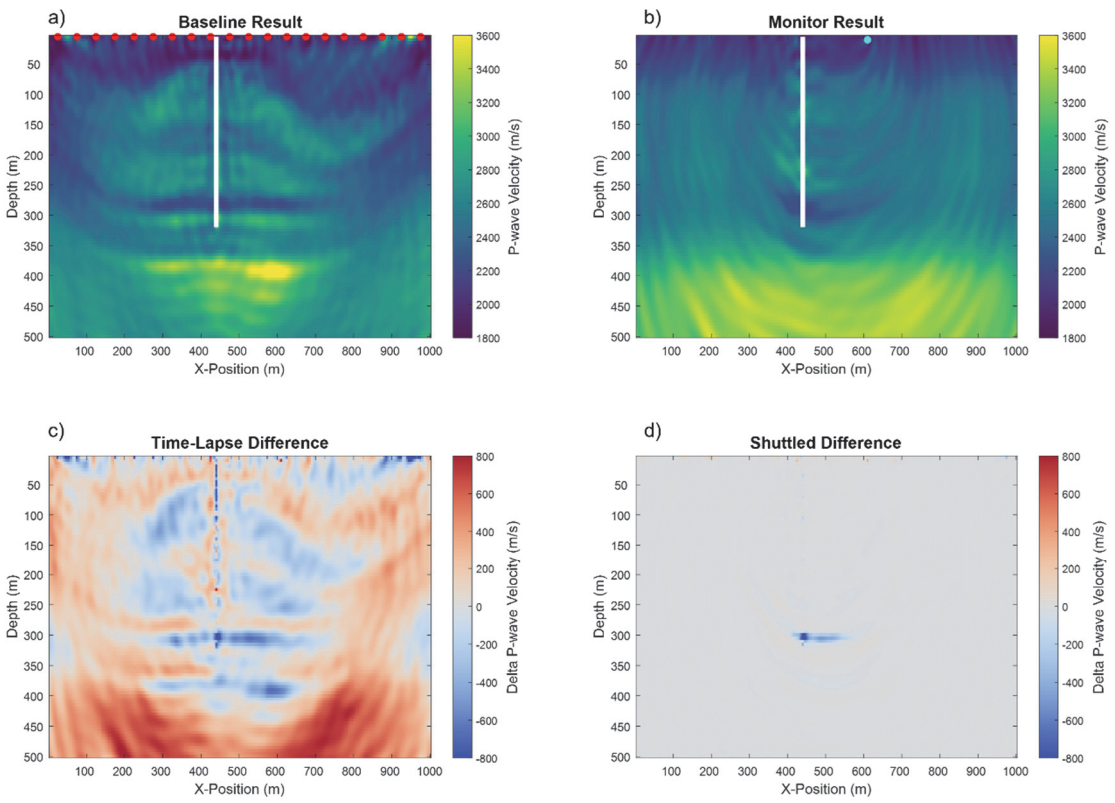

FIG. 6. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of a single source point located at 610 m, with added noise.

The monitor FWI result is uninterpretable; all energy is focused on arc-shaped artifacts. If we examine the velocity values at the receiver array, we note a decrease at the BBRS relative to the baseline survey, but the difference is not convincing. Interestingly, some indication of the velocity contrast at 380 m depth, below the zone of interest evident. The CO_2 plume cannot be clearly identified in time-lapse difference. Horizontal, coherent, low velocity features exist at the zone of interest, but significant artifacts related to the poorly sampled monitor survey are also present. After nullspace shuttling a single, small anomaly remains in the location of the injected CO_2 plume, suggesting the minimum time-lapse difference; all other noise is removed. This is remarkable, given the poor monitor FWI result.

Next, we consider a single source monitor survey, with the source point located farther from the injection and monitor wells at x-position 810 m. The results of this scenario are shown in Figure 7.

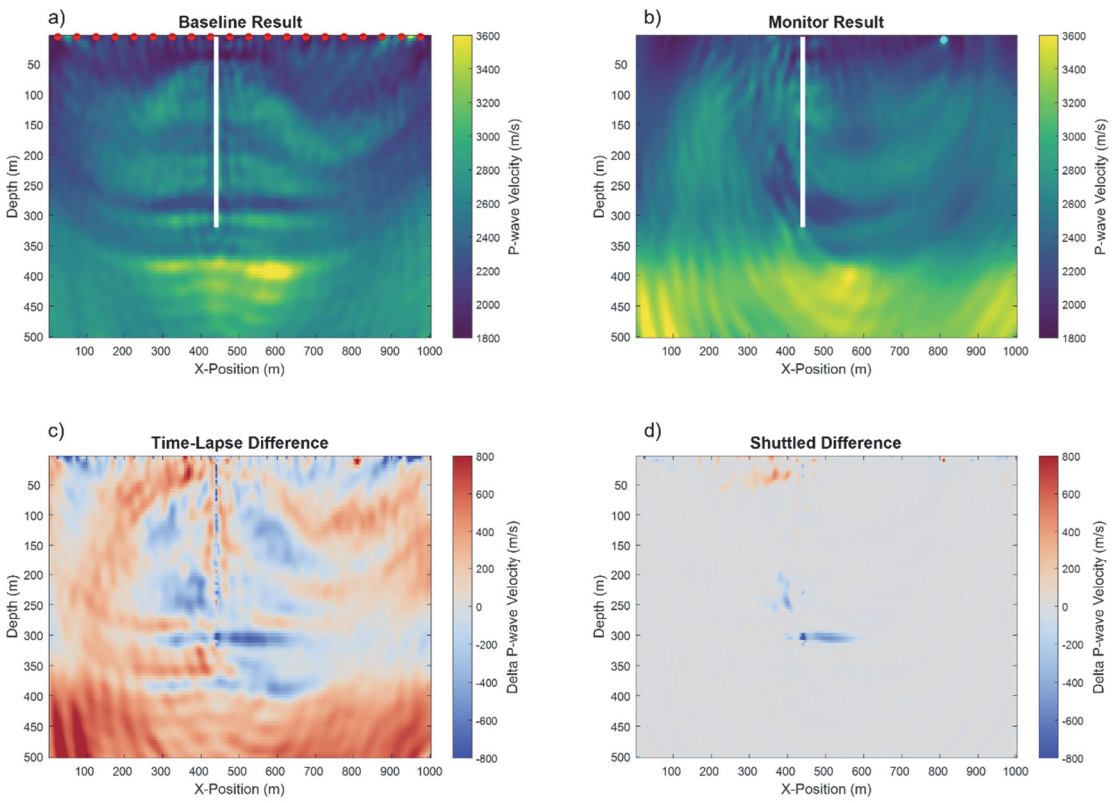

FIG. 7. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of a single source point located at 810 m x-position, with added noise.

Similar to the previous scenario, the monitor FWI result is not interpretable. A low velocity zone could be possibly identified at the receiver array, but it is not compelling. The time-lapse difference does indicate a decrease in the zone of interest but is overprinted by noise. The minimum difference after shuttling is similar to that observed in Figure 6d, however residual energy in the near surface persists.

We next consider a single-source monitor survey located at x-position 110 m. This source location is farthest from the CO_2 plume. The results of this scenario are shown in Figure 8.

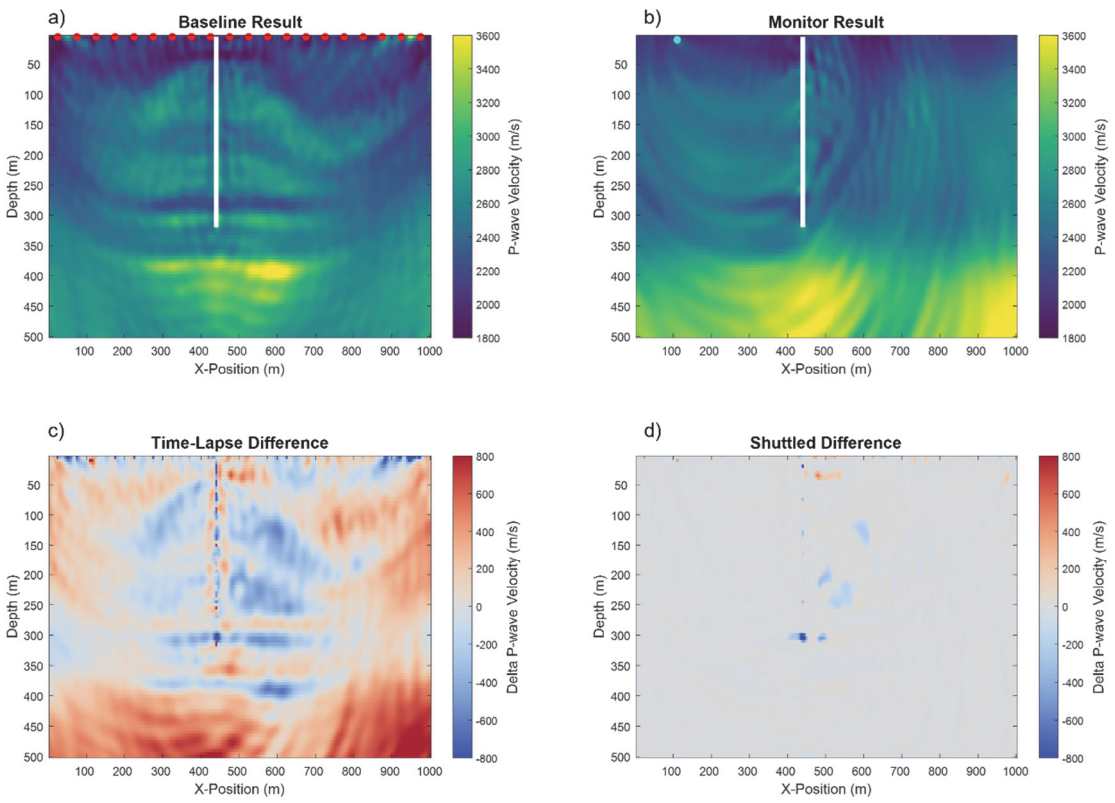

FIG. 8. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of a single source point located at 110 m x-position, with added noise.

In this example, the monitor survey source is located to the left of the receiver array and the CO_2 plume. The time-lapse difference does indicate a decrease in the zone of interest, but as in the previous example, it is overprinted by noise. After nullspace shuttling the CO_2 anomaly cannot be definitively identified. This acquisition geometry is not advisable for monitoring this CO_2 plume.

In the final example, we consider the extreme case of a single source baseline and a single source monitor. The results of this scenario are shown in Figure 9. Such poorly sampled baseline and monitor surveys are not expected to produce favorable results. However, this care does demonstrate the requirement for adequate acquisition effort in either the baseline or monitoring surveys to recover time-lapse changes, with or without nullspace shuttling.

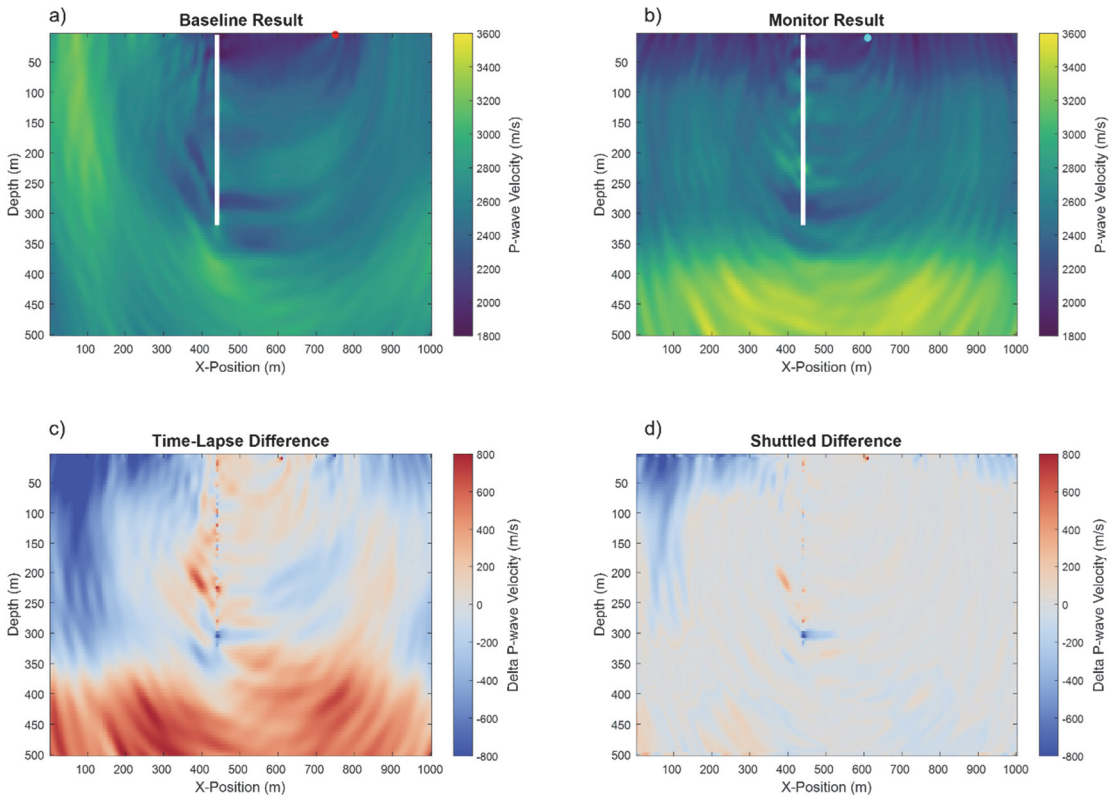

FIG. 9. a) Baseline inversion, b) monitor inversion, c) time-lapse difference and d) shuttled difference for the CO_2 injection case of a single baseline source located at 710 m and a single monitor source located at 610 m x-position, with added noise.

DISCUSSION

The application of nullspace shuttling to time-lapse FWI based on these synthetic models continues to show promise. The objective of nullspace shuttling in this report was to provide the minimum time-lapse difference between baseline and monitor, while remaining in the baseline and monitor nullspaces. The hypothesis function must be carefully defined, and results must be interpreted in the context of the hypothesis statement. In most cases nullspace shuttling was able to effectively isolate a minimum time-lapse difference due to CO_2 injection, and successfully removed noise unrelated to subsurface changes. This was especially remarkable in the case of sparse monitoring. A potential reason why shuttling is working so well is the linking of the baseline and monitor FWI results through the hypothesis function. This will require further investigation.

While comparing results from different scenarios is tempting, the shuttling optimization may not be at the same point in convergence, since number of iterations define the termination of the shuttling operation. Implementing a criterion-based cutoff such as momentum may be advantageous.

CONCLUSIONS

Time-lapse FWI has significant potential for CO_2 monitoring applications. Nullspace shuttling, in conjunction with an understanding of the expected elastic response to CO_2 injection, may reduce the costly requirement of repeatability in baseline and monitor survey geometry. Nullspace shuttling may be useful in evaluating acquisition geometry scenarios for baseline and monitoring surveys. An appropriately posed hypothesis function may allow for the detection of time-lapse changes, even when a traditional time-lapse difference does not.

ACKNOWLEDGEMENTS

The sponsors of CREWES are gratefully thanked for their continued support. This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through grant CRDPJ 543578-19. This work was done in collaboration with the Carbon Management Canada.

REFERENCES

Arts, R., Eiken, O., Chadwick, A., Zwegel, P., van der Meer, B., and Zinszner, B., 2004, Monitoring of CO2 injection at Sleipner using time-lapse seismic data: Energy, 29, 1383–1392.

Brevik, I., Eiken, O., Arts, R.J., Lindeberg, E., and Causse, E., 2000, Expectations and results from seismic monitoring of CO2 injection into a marine aquifer: 62nd EAGE Meeting, Glasgow, Paper B-21.

Chadwick, A., Williams, G., Delepine, N., Clochard, V., Labat, K., Sturton, S., Buddensiek, M.L., Dillen, M., Nickel, M., Lima, A.L., Arts, R., Neele, F., and Rossi, G., 2010, Quantitative analysis of time-lapse seismic monitoring data at the Sleipner CO2 storage operation: Leading Edge **29**, 170–177.

Deal M.M., and Nolet, G., 1996, Nullspace shuttles: Geophysical Journal International, **124**, No.2, 372-380.

Fu, X., and Innanen., K.A., 2021, Stepsize sharing in time-lapse full-waveform inversion: CREWES Annual Report, **33**.

Greaves, R.J., nad Fulp, T.J., 1987, Three-dimensional seismic monitoring of an enhanced oil recovery process: Geophysics, **52**, 1175-1187.

Huber, P.J., 1964, Robust Estimation of a Location Parameter: Ann. of Mathematical Statistics.: **35**, 73-101.

Keating., S.D., Li, J., and Innanen, K.A., 2018, Viscoelastic FWI: solving for q_p, q_s, v_p. v_s and density: CREWES Annual Report, **30**.

Keating., S.D., and Innanen, K.A., 2021, Null-space shuttles for targeted uncertainty analysis in full-waveform inversion: Geophysics, **86**, No. 1., R63-R76.

Keating., S.D., and Innanen, K.A., 2022, Targeted nullspace shuttling in time-lapse FWI: CREWES Annual Report, **34**.

Ivandic, M., Yang, C., Luth, S., Cosma, C., and Juhlin, C., 2012, Time-Lapse Analysis of Sparse 3D Seismic Data from the CO2 Storage Pilot Site at Ketzin, Germany: Journal of Applied Geophysics, **84**, 14-28.

Macquet, M., Lawton, D., Saeefdar, A., and Osadetz, K., 2019, A feasibility study for detection thresholds of CO2 at shallow depths at the CaMI field research station, Newell County, Alberta Canada: Petroleum Geoscience, **25**, 509-518.

White, D. J., Roach, L. A. N., and Roberts, B., 2015, Time-lapse seismic performance of a sparse permanent array; experience from the Aquistore CO2 storage site: Geophysics, **80**, WA35–WA48.