

Targeted nullspace shuttling in time-lapse FWI

Scott Keating and Kris Innanen

ABSTRACT

Time lapse inversion plays an important role in monitoring applications. Conventional approaches rely primarily on differencing strategies in either data or model space. The results of the model difference based approaches can be strongly influenced by the choice of starting model for baseline and monitor inversions. Different choices can result in different levels of mitigation of non-reproducible survey effects (for instance noise). We propose an approach that substitutes the importance of the initial model choice with explicit navigation of the inversion nullspace. In this strategy, targeted nullspace shuttling is used to find the baseline and monitor models that minimize the difference between models while preserving a desired level of data-fitting. In synthetic examples, this approach demonstrates a significant capacity to mitigate the effects of non-reproducible noise and changing acquisition, and to identify when time-lapse differences fall below the confidence threshold described by nullspace shuttling.

INTRODUCTION

Seismic time-lapse monitoring has significant applications, both in monitoring hydrocarbon reservoirs, and in tracking containment at CO_2 sequestration sites. One potentially powerful approach to seismic time-lapse monitoring is to make use of full waveform inversion (FWI). Many different time-lapse FWI methods have been developed, and a number of them are summarized and tested by in a previous CREWES report: Fu and Innanen (2021). Most time lapse inversion methods are fundamentally based on either 1) differencing a baseline and monitor inversion result, or 2) inverting a difference between baseline and monitor data sets. Even the most naive of these approaches can be effective in simple environments, but coping with more realistic data-sets in which non-reproducible effects, such as noise or differing acquisition geometries, play a role can be very challenging. More complex realizations of these time-lapse inversions often involve multiple inversions, making use of designed differences in the initial model used for each. In this way, these approaches are engaging in nullspace exploration: by starting with a different initial model, they attempt to find different solutions of approximately equal fitness with respect to the FWI objective function. The challenge with this type of approach is that the methods are designed based on the initial model, but their fitness as time-lapse inversion approaches is determined by the resulting points in the nullspace interrogated by an inversion. In general, as the objective function topography is unknown and costly to calculate, it can be very difficult to determine an appropriate set of starting models to achieve a meaningful difference, as the corresponding nullspace locations are unknown at the start of the inversion.

In this report, we investigate a nullspace shuttling approach for time-lapse FWI. This approach allows for targeted exploration of the baseline and monitor inversion nullspaces, which may allow for a time-lapse difference best corresponding to a measure of optimality to be determined.

THEORY

In general, a seismic inversion result depends on both the true properties of the subsurface and the details of acquisition geometry, optimization strategy, noise, etc. An ideal time-lapse inversion result would be one that identifies real time-dependent changes in the subsurface without identifying any spurious changes, such as those that might arise from non-repeatability in the seismic survey. In this way, time-lapse inversion can be thought of as an optimization procedure at the design level: we aim to minimize the differences recovered in the inversion that do not arise from real subsurface changes.

If the inversion problem for a given data-set had only one acceptable solution, the time-lapse problem could be solved simply by differencing inversion results with baseline and monitor data. In practice, noise and measurement errors mean that exact matching of the data is not necessary or desired, and many possible models fit the data acceptably well. In addition, the numerical optimization procedures used in FWI do not typically allow for an exact minimum of the FWI problem to be determined. For these reasons, the FWI problem generally has a set of solutions that all satisfy the inversion conditions equally well. This set of possible models is the nullspace of the inversion problem.

Many time-lapse inversion strategies can be viewed as approaches for finding different points in the nullspace of the monitor and baseline surveys that produce differences more reflective of the true time-lapse change. Figure 1 illustrates, in two dimensions, several common time-lapse FWI approaches. In this figure, the orange ellipse corresponds to the null-space of the monitor inversion, and the green ellipse corresponds to the baseline inversion, the grey dot corresponds to the initial model, and the orange and green dots correspond to the monitor and baseline inversion results. In the parallel strategy (top left), both inversions start from the same initial model. In the sequential and reverse-sequential strategies, one inversion begins from the initial model, while the other begins from the result of the first. In the common model strategy, the mean result of a parallel strategy result is used as a new initial model for a second round of parallel strategy inversion. While the set of strategies illustrated is not comprehensive, Figure 1 illustrates some common features of time-lapse FWI approaches: first, that differences in the initial model for each inversion can change the difference between inversion results, and so the time-lapse estimate, and second, that these changes in the initial model translate into navigations about the nullspace of the inversions. Figure 1 also shows that an attempt to assess the fitness of one time-lapse approach over another must involve a choice of fitness with respect to the null-space differences considered: without knowledge of the true models, the optimal approach must be expressed in terms of optimal nullspace differences.

While changing the initial model used in the inversion can allow for a different point in the nullspace of the inversion to be found, the relation between the change in the initial model and that in the nullspace difference is very difficult to predict, effectively requiring the solution of a full inversion problem. A more direct way to interrogate the nullspace in time-lapse applications is through targeted nullspace shuttling of the type described in Keating and Innanen (2021). Through targeted shuttling, we can directly search the nullspace for an optimal difference between baseline and monitor. Supposing that we define a metric ψ that takes a small value for good time-lapse estimates and a large value for bad ones, we

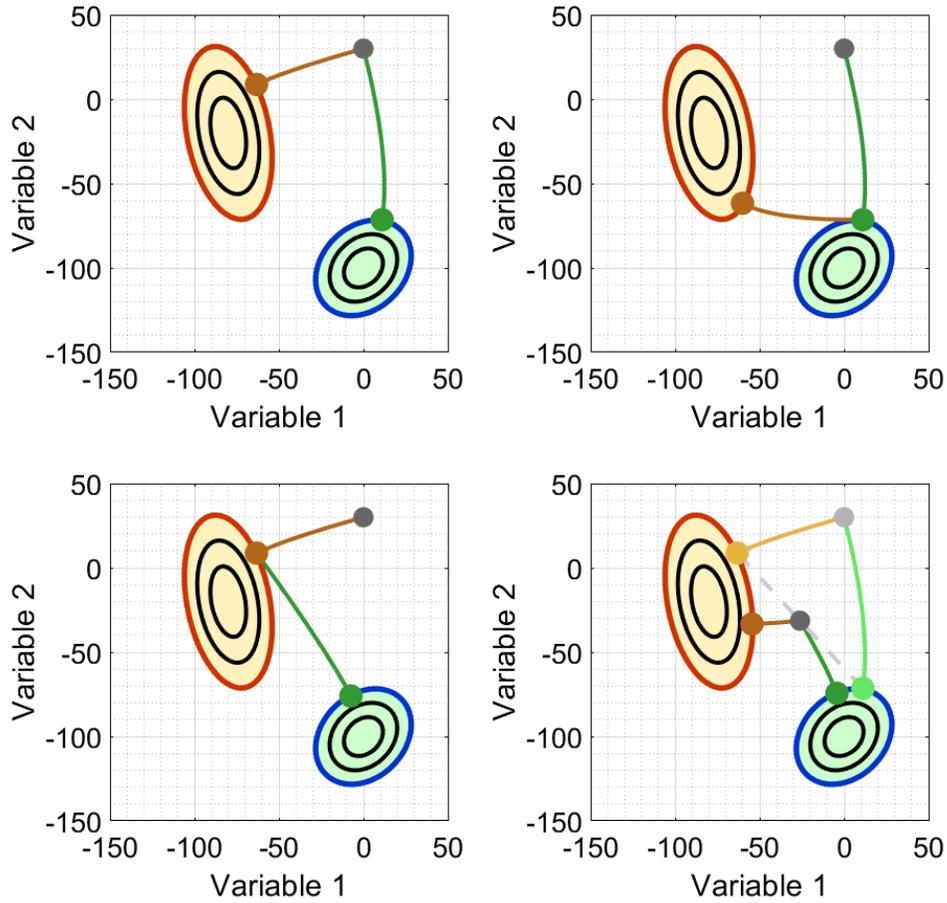


FIG. 1. Two dimensional illustration of some time-lapse inversion strategies, with the baseline inversion nullspace in green and the monitor inversion nullspace in orange. Top left: Parallel strategy: both inversions from the same starting model. Top right: Sequential strategy, with the monitor inversion beginning from the baseline result. Bottom left: Reverse sequential strategy, with the baseline inversion beginning from the monitor result. Bottom right: Common model strategy, an iterated parallel strategy, with the new starting model being the average of the previous inversion results. Note that the time-lapse estimate differs based on the choice of starting model.

can frame a shuttling problem for the optimal difference as follows:

$$\Delta m_{TL} = (m_M + \delta m_M^*) - (m_B + \delta m_B^*), \quad (1)$$

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

$$\begin{aligned} (\delta m_B^*, \delta m_M^*) &= \underset{\delta m_B, \delta m_M}{\operatorname{argmin}} \psi(m_B + \delta m_B, m_M + \delta m_M), \\ &\text{subject to } \phi_B(m_B + \delta m_B) < \epsilon_B, \phi_M(m_M + \delta m_M) < \epsilon_M, \end{aligned} \quad (2)$$

where ϕ_B and ϕ_M are the baseline and monitor FWI objective functions, and ϵ_B and ϵ_M are objective function values defining the nullspace. So, given a shuttling objective ψ we can solve equation 2 to find the best time-lapse difference by the metric ψ that satisfies our data and FWI priors.

Given that the true subsurface time-lapse change is not known, ψ needs to be chosen in a way that does not require access to this information. The ideal time-lapse estimate would contain only true subsurface changes and no differences arising from noise, acquisition, wavelet or other discrepancies between surveys. If the time-lapse changes are well-constrained by the data, then the constraints in equation 2 (which require that we minimize the FWI objective sufficiently well) will act to ensure that any valid nullspace difference includes the time-lapse change. In this case, an optimal difference would be one that minimized differences subject to the constraint of staying in the nullspaces. In other words, if fitting both data sets requires that some minimum time-lapse difference be present, other differences should be minimized (as these will correspond to unwanted differences between surveys). On the other hand, if fitting both data sets does not require a time-lapse difference, this suggests that our data are consistent with no time-lapse change, suggesting either no change or a change smaller than we can confidently detect. With this motivation, we define ψ here as the L_1 norm of the difference between baseline and monitor inversion results.

We choose the L_1 norm here because in the monitoring applications typical for time-lapse analysis, changes are focused in space and (relatively) large, while small changes throughout space are much less plausible. As L_2 promotes the latter over the former, we will focus on the L_1 norm here.

Shuttling

The nullspace shuttling approach we use here is summarized below. It is described in detail in Keating and Innanen (2021). In brief, we attempt to minimize the shuttling objective ψ over possible updates that preserve the FWI objective ϕ :

$$m_\psi = \underset{m}{\operatorname{argmin}} \psi(m) \text{ subject to } \phi(m) < \epsilon, \quad (3)$$

where ψ is the shuttling objective, ϕ is the FWI objective, and ϵ is a tolerance defining the nullspace of the inversion. As a global map of $\phi(m)$ is unavailable in FWI, we iteratively solve equation 3 by solving for optimal updated directions based on local models of ϕ constructed from the gradient (g) and Hessian-vector products ($H\delta m$):

$$\delta m^* = \underset{\delta m}{\operatorname{argmin}} \psi(m + \alpha(g(m), H(m)\delta m)\delta m), \quad (4)$$

where α is the maximum step expected to satisfy $\phi(m + \delta m) < \epsilon$, based on the gradient at the current model and the Hessian-vector product for the current update direction estimate. After solving for the optimal direction, we perform a line-search along δm^* to find the true maximal length of our nullspace step. This process is iterated to approximate a solution to equation 3.

In the time-lapse formulation above (equation 2), there are two shuttling problems as both the baseline and the monitor models are updated. The approach we adopt here to accommodate this complication is to solve these shuttling sub-problems iteratively and sequentially: first the baseline model is shuttled to minimize (in an L_1 sense) the difference from the monitor, then the monitor is shuttled to match the updated baseline, potentially

iterating these steps several times. In this way, the time-lapse problem can be reduced to a set of single-objective shuttling problems of the type discussed in Keating and Innanen (2021).

In our numerical experiments, we define the shuttling objective ψ as the Huber norm of the difference between the baseline and monitor models. This allows for the positive aspects of the L_1 norm (as described above) to dominate for large differences, but at small differences transitions to an L_2 norm with better convexity and optimization properties.

Hessian-vector product calculation

In order to perform the optimization problem in equation 4, Hessian-vector products of the form $H\delta m$ need to be calculated. These products can be efficiently calculated at the same computational cost as a gradient calculation, given that the gradient has already been calculated at the same model-space location (e.g. Métivier et al., 2013), and this approach was used extensively in Keating and Innanen (2021). Hessian-vector product calculation codes can be cumbersome to write, however, and this can be an obstacle to quick implementation of targeted nullspace shuttling approaches. To illustrate that this additional programming work is not critical, we use a finite-difference implementation of our Hessian vector product calculations in this report instead of the analytic calculators used in Keating and Innanen (2021). The finite-difference approximation used is

$$H(m)\delta m \approx \frac{g(m + \gamma\delta m) - g(m)}{\gamma}, \quad (5)$$

where g is the gradient and γ is a small number. While this approach has the same computational cost as the analytic approach, it is easy to program given a gradient calculator and, in our testing, the change does not materially alter the shuttling results.

NUMERICAL EXAMPLES

In this section, we present numerical examples of the time-lapse shuttling approach. In these examples, we consider a 2D, constant-density, acoustic, frequency-domain forward modelling and inversion. The true baseline model, an unscaled model of time-lapse changes, and the initial model used in the inversions, are shown in Figure 2. The true baseline model is used for all of the examples in this section, while the true monitor is generated by scaling the difference term and adding it to the baseline. In this way, we will investigate the time-lapse shuttling approach for a range of different magnitudes of time-lapse change. The models are discretized in 10 m cells, with 1 km depth and 2.75 km horizontal extent. The seven source locations used for the baseline survey are shown by the green stars in Figure 2, while receivers are densely distributed along the top of the model.

In each of our inversions in this section, we invert six frequency bands of six frequencies each. Each band ranges from 1 Hz to a maximum frequency, which begins at 5 Hz for the first band, and rises to 25 Hz for the final band. At each band, we perform 12 iterations of BFGS optimization. With a frequency-domain implementation, the nullspace shuttling we consider here must be performed with respect to the inversion on a specific frequency band, and here we use 12 frequencies from 1 to 25 Hz. To ensure a good data-fit on this band, we

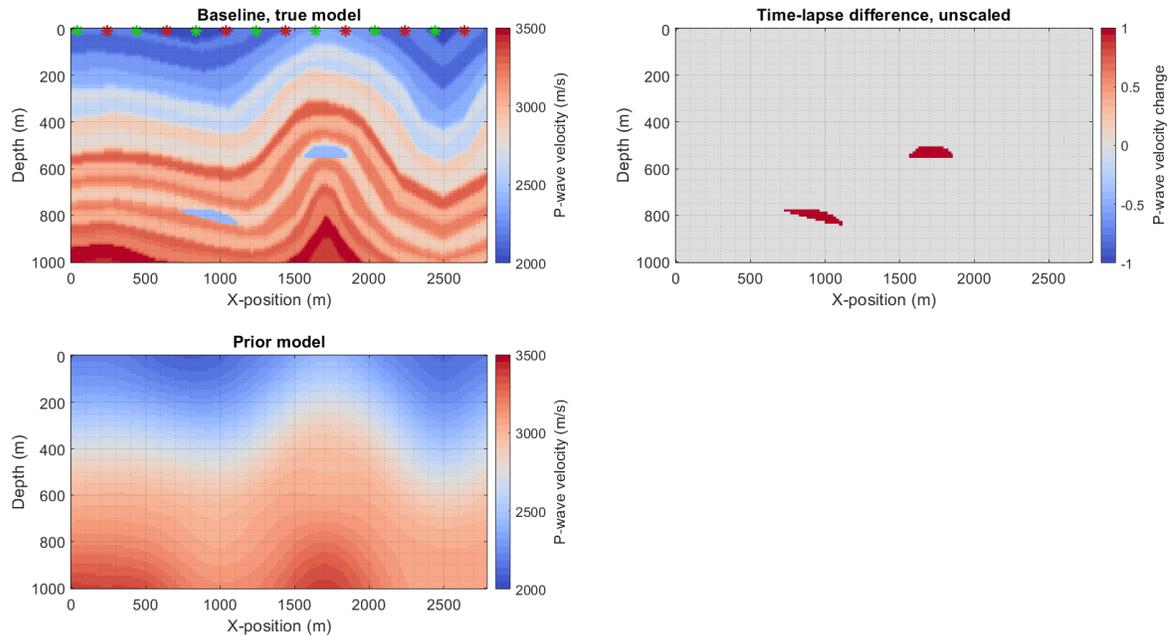


FIG. 2. True baseline (top left), difference between monitor and baseline (top right) and prior model (bottom left) for the synthetic tests. The true monitor model is generated by adding a scaled version of the differences to the true baseline.

perform a final update of 20 iterations on the shuttling frequency band. The monitor and baseline inversions were both initiated from the prior model in Figure 2.

As a first example, we consider the case where the time lapse contrasts are 300 m/s, no noise is present, and the monitor survey geometry exactly reproduces the baseline survey. Figure 3 shows the baseline and monitor inversion results for this case, as well as their difference. Given the high degree of repeatability in this case, a simple differencing provides a good estimate of the time-lapse changes here. A much more challenging case for time-lapse inversion arises when there are significant differences between the baseline and monitor survey. To simulate challenges with repeatability, we repeat the same numerical example, but with the sources in the monitor moved to the location of the red stars in Figure 2, and a small noise term added to both monitor and baseline data-sets (different for each). We choose the noise such that the data SNR is approximately 10, where we define SNR as the ratio between the 2-norm of the data over the 2-norm of the noise. Figure 4 shows the baseline and monitor inversion results for this case. Here, interpretation of the difference between inversion results is much more difficult to interpret. Large changes exist at the locations of the true anomalies, but also exist at many different locations in the model. Without strong prior information, it would be very difficult to identify for which specific regions the time-lapse estimates are accurate and for which they are not. In this case, nullspace shuttling may be appropriate.

To obtain a better estimate of the time-lapse changes in the Figure 4 example, we apply targeted nullspace shuttling to minimize the Huber norm of the difference between models. Starting with the baseline, we shuttle each model in turn by performing 15 iterations of L-BFGS optimization on to solve equation 4, followed by a line-search in the calculated

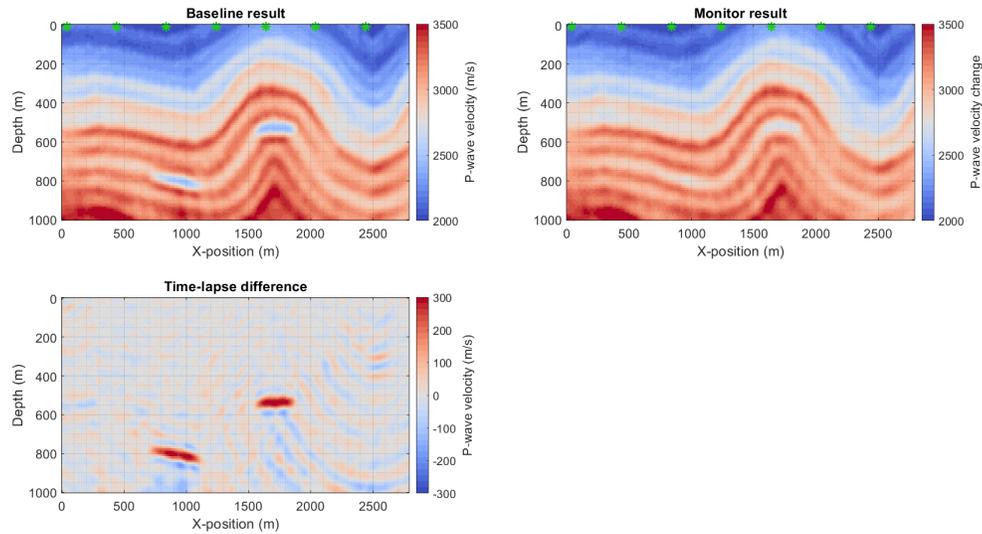


FIG. 3. Baseline inversion (top left), monitor inversion (top right), and their difference (bottom left) for the 300 m/s change case with zero noise and exactly repeated acquisition.

optimal direction. We shuttle the baseline and monitor each a total of three times. The computational cost of the shuttling problem is about two wave-propagations per iteration, similar to the calculation of the gradient in FWI, so our 90 total iterations have a cost comparable to each of the two inversions themselves. The resulting time-lapse differences from this procedure are shown in Figure 5: the left column shows the difference after updating the baseline the first, second and third time, while the right column shows the difference after updating the monitor. The bottom right panel of Figure 5 shows the final result of shuttling. In this case, shuttling is able to largely remove model differences not associated with the true time-lapse changes, and the remaining differences after shuttling provide a much better estimate of the time-lapse changes than the initial difference does.

In Figures 6 to 9 we show the results of repeating the previous example with different values for the magnitude of the true time-lapse change. In each figure the baseline and monitor inversion results are shown on the top row, while the time-lapse differences before and after shuttling are shown in the bottom row. Figures 6, 7, 8 and 9 show the results for 500, 300, 150, and 50 m/s differences, respectively. In the 500 m/s case (Figure 6), the difference without shuttling is largely interpretable, though the shuttled result provides a clearer identification of the true time-lapse changes. In the 300 m/s case (Figure 7, also Figure 5), the difference alone is quite difficult to accurately interpret, but the shuttled result allows for clear identification of the anomalies. In the 150 m/s case (Figure 8), the non-repeatability errors are sufficiently large to make identification of the time-lapse changes from the inversion differences alone almost impossible. After shuttling, however, the shallower anomaly is easy to spot and the deeper anomaly may or may not be interpretable. In the 50 m/s case (Figure 9), the time-lapse differences do not allow for identification of the anomalies, while the shuttled differences have negligible non-zero values throughout. This suggests that the nullspaces of the inversion results overlap each other. In other words, the data in this case are not sufficient to confidently identify a time-lapse change because models exist that satisfy both data sets to within the noise level. In this case, the nullspace

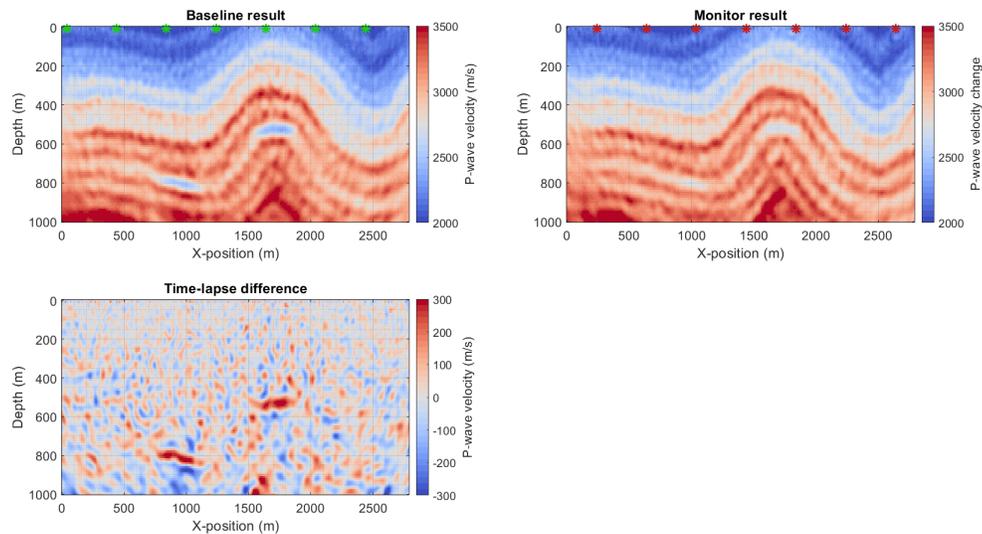


FIG. 4. Baseline inversion (top left), monitor inversion (top right), and their difference (bottom left) for the 300 m/s change case with noise and inconsistent acquisition.

shuttling does identify the anomaly, but it does identify that any anomaly present is below our detection threshold.

DISCUSSION

It is likely that each of the shuttled time-lapse estimates corresponds to the result that would have been obtained by inverting the baseline and monitor data sets and differencing the results, given the appropriate choice of starting models for each inversion. The appropriate choice of starting models, however, is not easy to determine, and the computational cost of finding the appropriate starting models is likely dramatically more costly than the approach we adopt here.

While the shuttling results we obtain here incur a cost of approximately one additional FWI problem, this cost can likely be mitigated with intelligent optimization strategies. Figure 5 shows that most of the model noise is eliminated after a few iterations, and only the larger residuals persist afterwards. This is likely a consequence of the Huber norm objective we use in ψ , which prioritizes reducing all residuals above a certain size equally. With better optimization strategies (e.g. a momentum term) or a better choice of objective function, similar results could likely be achieved with significantly less computation.

CONCLUSIONS

Time-lapse FWI has significant potential for monitoring applications, but robustness to non-repeatability in survey conditions is a significant limiting factor. The minimum time-lapse change between two acceptable inversion results should be a relatively robust measure of the time-lapse changes that the data confidently support. Finding this difference by changing the initial models used in the respective inversions could be extremely challenging. A more direct approach is to use targeted nullspace shuttling to find the op-

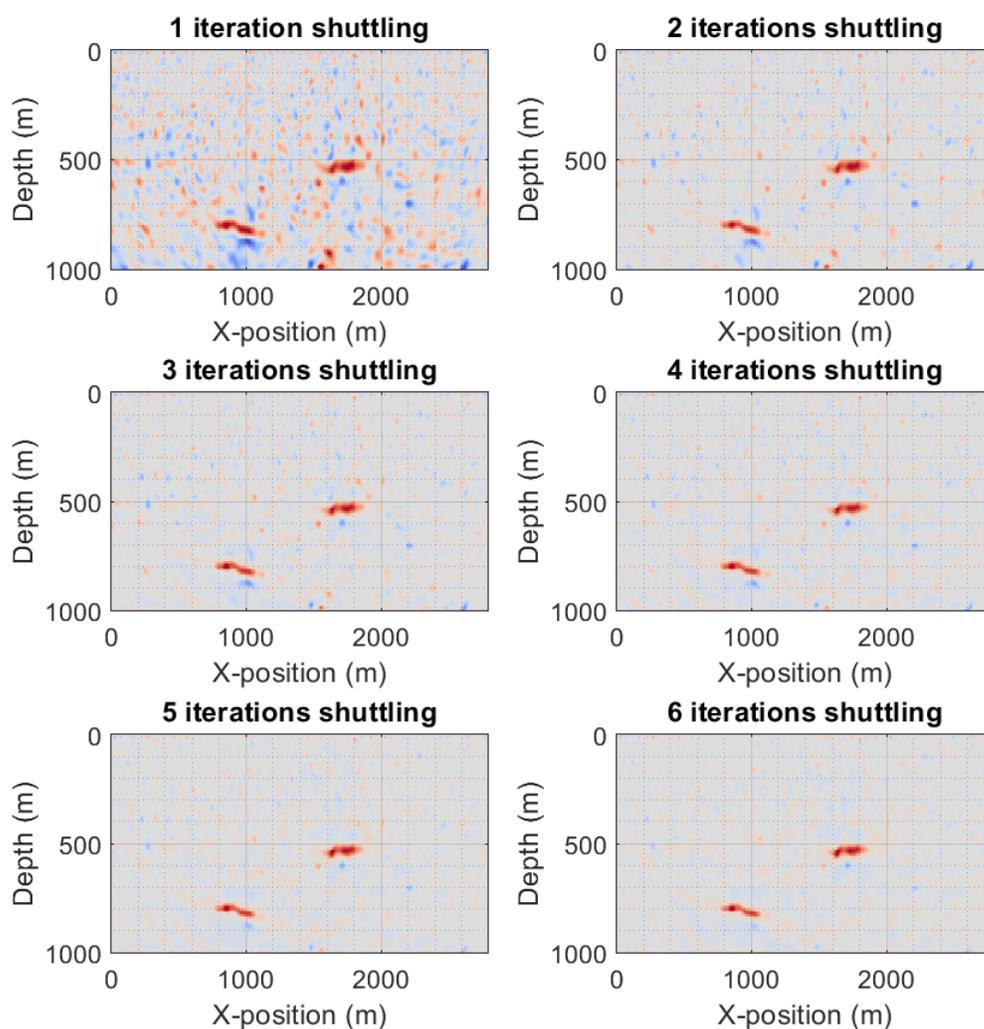


FIG. 5. Time-lapse estimate from Figure 4 after iterating targeted nullspace shuttling. Odd-number iterates update the baseline, even iterates update the model. These plots are on the same color-scale as Figure 4.

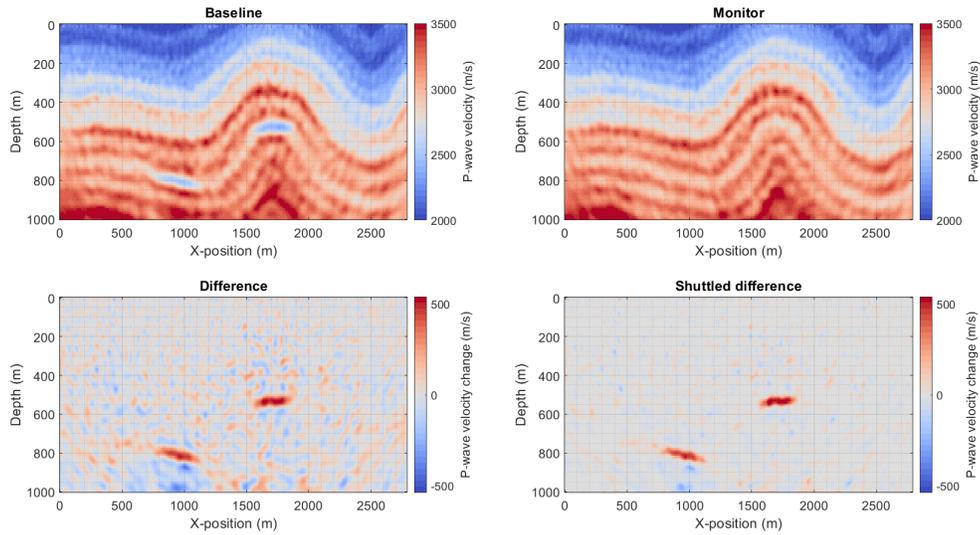


FIG. 6. Baseline inversion (top left), monitor inversion (top right), their difference (bottom left), and the shuttled difference (bottom right) for the 500 m/s change case with noise and inconsistent acquisition.

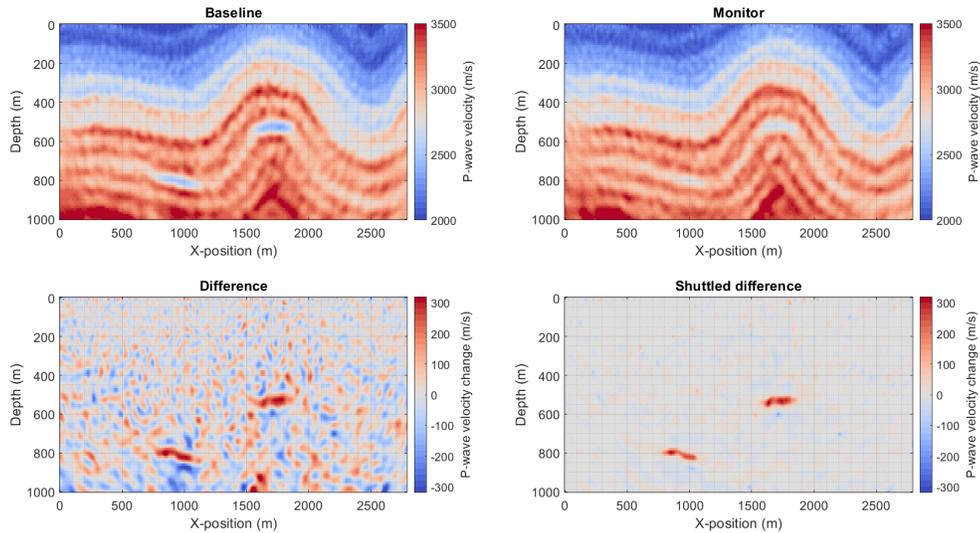


FIG. 7. Baseline inversion (top left), monitor inversion (top right), their difference (bottom left), and the shuttled difference (bottom right) for the 300 m/s change case with noise and inconsistent acquisition.

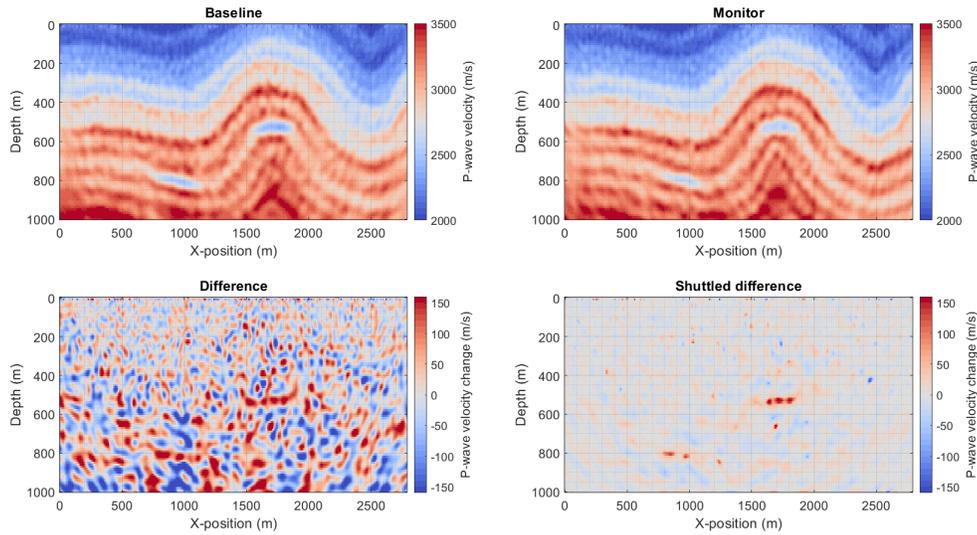


FIG. 8. Baseline inversion (top left), monitor inversion (top right), their difference (bottom left), and the shuttled difference (bottom right) for the 150 m/s change case with noise and inconsistent acquisition.

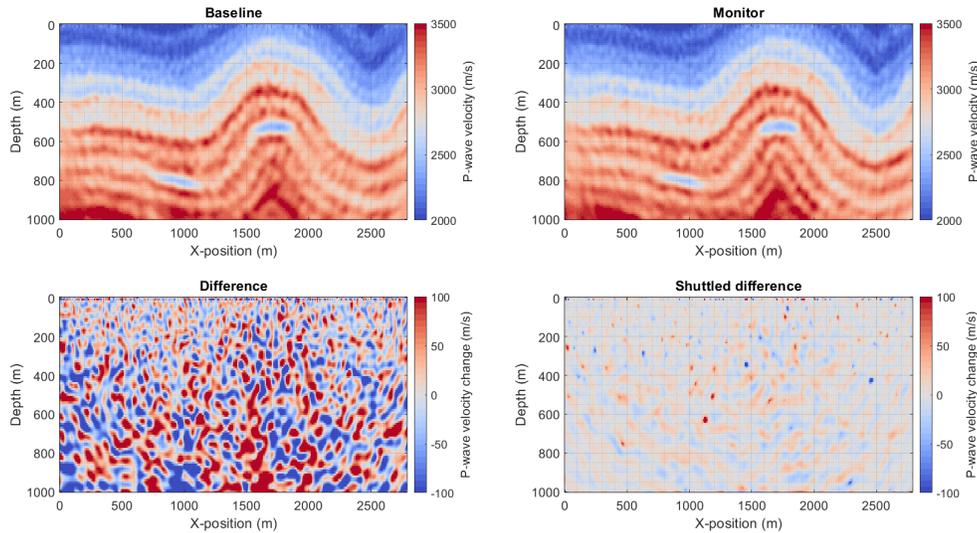


FIG. 9. Baseline inversion (top left), monitor inversion (top right), their difference (bottom left), and the shuttled difference (bottom right) for the 50 m/s change case with noise and inconsistent acquisition.

timal difference. We apply a targeted nullspace shuttling time-lapse inversion approach to noisy synthetic data with inconsistent acquisition geometries. In several cases, the null-space shuttling approach is able to identify time-lapse changes much more accurately than a naive parallel inversion strategy. In a case with very small contrasts, the shuttling inversion approach suggest either no change or a change below the detection threshold.

ACKNOWLEDGEMENTS

The authors thank the sponsors of CREWES for continued support. This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through the grant CRDPJ 543578-19. Scott Keating was also supported by the Canada First Research Excellence Fund, through the Global Research Initiative at the University of Calgary. We gratefully acknowledge Xin Fu for productive discussions on time-lapse inversion.

REFERENCES

- Fu, X., and Innanen, K. A., 2021, Stepsize sharing in time-lapse full-waveform inversion: CREWES Annual Report, **33**.
- 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.
- Métivier, L., Brossier, R., Virieux, J., and Operto, S., 2013, Full waveform inversion and the truncated Newton method: *SIAM Journal on Scientific Computing*, **35**, No. 2, no. 2, B401–B437.