Inversion of vector-acoustic data in a local domain

Colton Kohnke and Paul Sava
Center for Wave Phenomena, Colorado School of Mines

ABSTRACT
Seismic wavefield modeling and inversion requires large computational resources. For applications that involve modeling wavefields through many closely related models (i.e. time-lapse FWI, iterative RTM imaging, etc.) it is of interest to compute the response locally within a region where the model changes instead of computing the response of the entire model. By using the method of multiple point sources, the wavefields inside a local domain can be simulated to mimic the wavefields that propagate in the full domain. The localized wavefields are exact within numerical precision for both forward and adjoint wavefields. FWI gradients for velocity and density model parameters can then be built from the local wavefields to update the model parameters efficiently inside the local domain for a more accurate representation of the subsurface. The wavefields and gradient at each iteration are computed locally at a cost orders of magnitude lower than in the full domain, thus enabling more computationally efficient inversions. Using a synthetic example, we show that inversion inside a local domain is comparable to inversion in the full domain, but is obtainable at a significantly lower cost.

Key words: 4C, computational efficiency, localization, FWI

1 INTRODUCTION

As geophysicists we are constantly pushing our methods and algorithms to give more information about the subsurface using all available information. The physical property models are becoming larger and include more complexity, and the computer effort required to run wavefield simulations through such models require more efficient methods to handle the computations. Meanwhile the exploration targets we are trying to image are becoming more subtle and spatially constrained (i.e. salt boundaries in a specific area, time-lapse changes in a reservoir, unconventional reservoirs, etc.). Even though a region of interest may be small, the models remain large to accommodate wavefield aperture, and therefore the computational time required for a solution remains large. This is inefficient: if the region of interest is small, we should not repeatedly propagate waves in the full model domain, as the majority of this model remains constant while we investigate the region of interest. Rather we should localize the computations to the region of interest such that the wavefields inside this domain match the wavefields that would be obtained by propagation in the full domain.

To further expand upon the problem statement let us consider full waveform inversion (FWI) (Tarantola, 1984; Plessix, 2006; Virieux and Operto, 2009; Symes, 2007; Pratt, 1999) where we may assume to know the background model perfectly, while the perturbation we are trying to image is localized to a small region of interest. If this is the case, the inversion should be constrained to only make model updates within the specified region of interest. To update the model one must construct the gradient of an objective function, for example one that minimizes the difference between the observed and predicted data at the receiver locations. Building the gradient using the adjoint state method (Plessix, 2006) requires the state variable (the wavefield from the sources to the receivers in forward time) and the adjoint variable (the wavefield from the receivers to the sources in reverse time). The gradient is the zero-lag cross-correlation between the state and adjoint variables. Once the model is updated in the direction of the gradient, the state and adjoint variables need to be recomputed in the updated model to find the gradient for the next model update. Since the model is only updating within the small region of interest, much of the wavefield simulation is an expensive repeated computation outside the region of interest. The gradient is sensitive to wavefield changes inside the region of interest, therefore if we could simulate them locally as if they were simulated in the full domain, we could drastically speed-up the inversion process.

The technique described in this paper outlines how to confine wavefields to a local domain and build FWI gradients at a lower computational cost than by using the full model. The gradients are then used for inversion for velocity within the local domain. To accomplish wavefield localization we utilize the method of Multiple Point Sources (MPS) introduced by

Vasmel and Robertsson (2016) (also called Finite-Difference Injection by Robertsson and Chapman (2000) and van Manen et al. (2007)) which reposition the sources and receivers to the local domain and assembles an equivalent experiment. The sources act as boundary conditions that create the wavefield interior to the local domain and cancel the wavefields in its exterior. The MPS method is limited to the same order of scattering as the model used to derive the boundary sources; however, Vasmel (2016), Broggini et al. (2017), and Willemsen (2017) describe how higher order scattering with the full domain can be included by utilizing Green's functions and Immersive Boundary Conditions (IBC). To include the full scattering, the Green's functions from every point on the boundary of the local domain to every other point on the boundary is required, along with the Green's functions between the sources and the local boundary, and between the receivers and the local boundary. These Green's functions only need to be computed once, but they need to be stored, and accessed at every time-step in the simulation. Unfortunately the computation and storage of the required 3D Green's functions is non-trivial, so we choose to refrain from using the higher order scattering at this time.

Applications of localizing wavefields are varied and include any problem where many simulations over similar models are needed, and the models differ only in a localized region. Such problems are fairly common in geophysics. Yang et al. (2012) construct a virtual survey for target-oriented time-lapse FWI by redatuming past a complex overburden. Shabelansky et al. (2013) accomplish a similar goal, but redatum to estimate sensitivity for time-lapse seismic. Willemsen et al. (2016) use a local solver in the frequency domain to invert for salt boundaries using FWI. A similar solver is utilized by Kotsi et al. (2018) for fast sensitivity analysis to time-lapse changes in a local domain. Willemsen and Malcolm (2016) couple an acoustic full domain and elastic local domain and Osorio and Malcolm (2018) localized elastic modeling for amplitude variation with offset (AVO) inversion. Elison et al. (2018) describe how to utilize localization techniques to derive the Green's functions required for Marchenko imaging, and apply boundary conditions on a local survey.

While the IBC method that utilizes Green's function may be more accurate, the MPS method is less costly to implement and offers a useful approximation for localizing wavefields to a local domain. The MPS method is still largely unexplored for inversion purposes, as are the physics of how the wave operator ties in to the adjoint state method to build FWI gradients in the full and local domains. This paper is structured as follows: first we discuss the wave operator, and then define the physics of localization using the MPS method. Then we define the full and local domain FWI gradients, and finally perform local inversions and compared them to their counterparts in the full domain.

2 VECTOR ACOUSTIC OPERATOR

Marine seismic surveys typically use airguns that act as monopole sources, and record the pressure wavefield at streamer locations. However, recent advancements in marine seismic acquisition technology facilitate recording of vector acoustic seismic data, comprised of both pressure and particle velocity wavefields (Robertsson et al., 2008). New marine seismic acquisition systems are capable of generating monopole and dipole type sources (Halliday et al., 2012). Having access to the full wavefield, i.e. both vector and scalar components, gives an opportunity for a more accurate understanding the subsurface. Specifically, using all the components of the wavefield improves reverse-time migration (RTM) by reducing artifacts and improving illumination (Ravasi et al., 2015), and full waveform inversion (FWI) by providing constraints on the subsurface parameters (Fleury and Vasconcelos, 2013; Zhong and Liu, 2019). In order to take advantage of the full pressure and particle velocity wavefield, we need to model the wavefields using the first-order vector-acoustic wave equation. The vector-acoustic wave equation will also be a requirement to confine the wavefields to a local domain using the acoustic representation theorem in de Hoop (1988).

The first-order vector-acoustic wave equation is

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \nabla p = \mathbf{f} \tag{1}$$

$$\frac{1}{\rho c^2} \frac{\partial p}{\partial t} + \nabla \cdot \mathbf{v} = q. \tag{2}$$

This equation links particle velocity $\mathbf{v}(\mathbf{x}, t)$ with the pressure $p(\mathbf{x}, t)$ from dipole $\mathbf{f}(\mathbf{x}, t)$ and monopole $q(\mathbf{x}, t)$ sources. The wavefields propagate in a medium with spatially variable density $\rho(\mathbf{x})$ and velocity $c(\mathbf{x})$. Equations 1 and 2 can be written in a compact operator form as

$$\mathbf{L}\mathbf{u} = \mathbf{d}, \tag{3}$$

where

$$\mathbf{L} = \begin{bmatrix} \rho \dfrac{\partial}{\partial t} & \nabla \\ \nabla \cdot & \dfrac{1}{\rho c^2} \dfrac{\partial}{\partial t} \end{bmatrix} \tag{4}$$

$$\mathbf{u} = \begin{bmatrix} \mathbf{v} \\ p \end{bmatrix}, \tag{5}$$

$$\mathbf{d} = \begin{bmatrix} \mathbf{f} \\ q \end{bmatrix}. \tag{6}$$

We solve equation 3 iteratively using finite-differences, by alternating between particle velocity at the current time-step and pressure at the next time-step. The pressure and particle velocity grids are staggered in space and time on a Yee grid (Yee, 1966). The wavefield term $\mathbf{u}$ is the state variable used for building FWI gradients.

The adjoint operator is written as

$$\mathbf{L}^{\dagger}\mathbf{a} = \mathbf{g}, \tag{7}$$

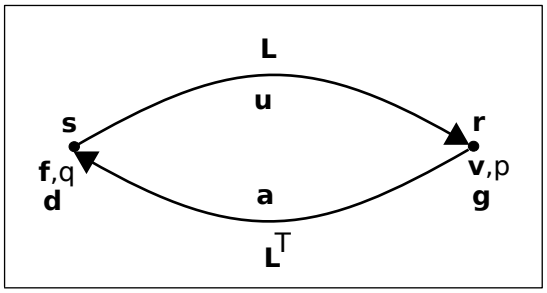


Figure 1. Visual representation of the forward and adjoint operators in the full domain. The forward operator $\mathbf{L}$ creates the state variable $\mathbf{u}$ from the data $\mathbf{d}$ at the source $\mathbf{s}$. The adjoint operator creates the adjoint variable $\mathbf{a}$ from the data difference $\mathbf{g}$ at the receivers $\mathbf{r}$.

where

$$\mathbf{L}^\dagger = \begin{bmatrix} \rho\left(\frac{\partial}{\partial t}\right)^\dagger & -\nabla \\ -\nabla\cdot & \frac{1}{\rho c^2}\left(\frac{\partial}{\partial t}\right)^\dagger \end{bmatrix}, \tag{8}$$

$$\mathbf{a} = \begin{bmatrix} \mathbf{f} \\ q \end{bmatrix} \tag{9}$$

and $\mathbf{g}(\mathbf{x},t)$ is the adjoint source, often the data difference at the receivers in FWI applications, and $\mathbf{a}(\mathbf{x},t)$ is the adjoint variable. The adjoint operator is solved iteratively in reverse time, by alternating between the monopole term (q) at the last time-step and the dipole term ($\mathbf{f}$) at the previous time-step.

Figure 1 gives a visual representation of equations 3 and 7, showing the forward operator $\mathbf{L}$ creating the state variable $\mathbf{u}$ from the data $\mathbf{d}$. The adjoint operator creates the adjoint variable $\mathbf{a}$ from the data difference $\mathbf{g}$. The source locations are indicated by $\mathbf{s}$ and the receiver location by $\mathbf{r}$.

3 WAVEFIELD LOCALIZATION

Our goal is to localize wavefields inside a domain of interest such that the local forward and adjoint wavefields match the wavefields that would propagate within the full domain. To accomplish this, we employ the representation theorem from de Hoop (1995)

$$\begin{aligned} p(\mathbf{x},t) = \oint_{\partial\Omega} [&p(\mathbf{x},t) * \mathbf{G}^{\mathbf{f}}(\mathbf{x},\mathbf{x}',t) \\ &+\mathbf{v}(\mathbf{x},t) * G^{q}(\mathbf{x},\mathbf{x}',t)] \cdot \mathbf{n} dS \text{ for } \mathbf{x} \in \Omega. \end{aligned} \tag{10}$$

In equation 10, $\mathbf{G}^{\mathbf{f}}$ and G^q are the Green's functions in the local medium for dipole and monopole source terms, respectively. The domain of interest (Ω) is bounded by the surface $\partial\Omega$ and has outward pointing normal $\mathbf{n}$. The asterisk ($*$) represents time convolution. Equation 10 states that if we know the pressure and particle velocity on the boundary of a surface, we can find the pressure (and therefore the particle velocity) inside the surface. Convolving the pressure and particle velocity with Green's functions is equivalent to creating monopole and dipole sources on $\partial\Omega$ using the method of Multiple Point Sources (MPS), introduced by Vasmel and Robertsson (2016). We discretize equation 10 and create the boundary source terms

$$\begin{aligned} \mathbf{f}(\mathbf{x},t) &= p(\mathbf{x},t)\mathbf{n} \\ q(\mathbf{x},t) &= \mathbf{v}(\mathbf{x},t)\cdot\mathbf{n} \end{aligned} \tag{11}$$

at every point on $\partial\Omega$. The vector $\mathbf{n}$ in equation 11 is the outward pointing normal of $\partial\Omega$ that has been scaled by the finite-difference grid spacing (Δx, Δy, and Δz) and is written as

$$\mathbf{n} = \begin{bmatrix} \frac{n_x}{\Delta x} \\ \frac{n_y}{\Delta y} \\ \frac{n_z}{\Delta z} \end{bmatrix}. \tag{12}$$

The scaling by the grid spacing comes from discretizing the surface integral in equation 10 and brings the units of the wavefield terms to be consistent with the units of the source terms. Without this scaling, the wavefields inside the local domain have the same shape, but the magnitude is different. The dipole term ($\mathbf{f}$) is computed by converting the scalar pressure on the surface and multiplying by $\mathbf{n}$. The monopole term (q) is found by converting the vector particle velocity with a dot product with $\mathbf{n}$. The sources in equation 11 form the local boundary conditions on $\partial\Omega$ and create the wavefield inside the local domain, while also canceling the wavefield as it leaves the domain.

Creation of the local sources can be written in operator form as

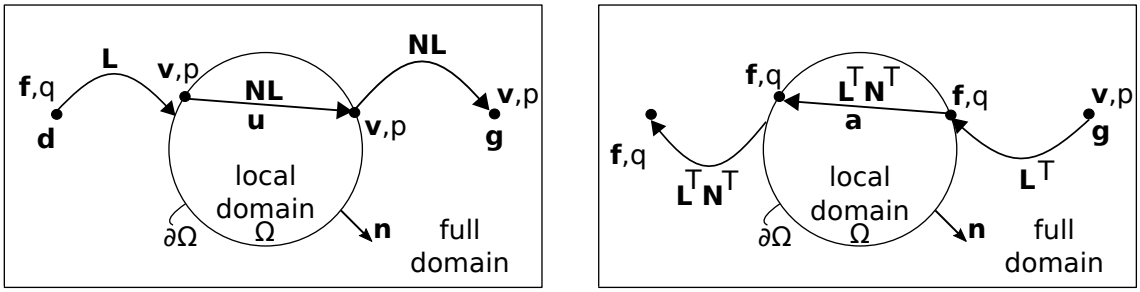
$$\mathbf{d}_L = \mathbf{NWu} \tag{13}$$

where $\mathbf{W}$ is the windowing operator that restricts the wavefield $\mathbf{u}$ to $\partial\Omega$ and

$$\mathbf{N} = \begin{bmatrix} 0 & \mathbf{n} \\ \mathbf{n}\cdot & 0 \end{bmatrix}, \tag{14}$$

is the localization operator. $\mathbf{d}_L$ are the new source terms on $\partial\Omega$ to be used in the forward operator given by equation 3. Using the new sources on $\partial\Omega$, we no longer need to consider the full domain to compute $\mathbf{u}$ inside Ω, as shown in Figure 2(a). The operator $\mathbf{N}$ can also be used to continue the wavefield from the local domain back into the full domain, as shown in Figure 2(a). If we perturb the model inside Ω, the sources derived on the surface still generate $\mathbf{u}$ inside the region, but they do not cancel out the wavefield that leaves the region. The local sources defined in equation 11 do not account for multiple scattering between the full and local domains. In other words, if the model perturbation would cause a wave to scatter outside the local domain and return in the local domain, the return wavefield is not modeled. In order to model the higher order scattering terms we could construct a boundary condition with explicit Green's functions (Vasmel, 2016), but the computation and storage requirement would increase significantly, thus diminishing the gains of localization.

The adjoint scenario is similar, as shown in Figure 2(b),



(a) (b)

Figure 2. A visual representation of the operator chain for a) the forward and b) the adjoint operators. With proper localization, we are able to ignore the operators that describe wave propagation outside the local domain.

with the local adjoint source

$$\mathbf{g}_L = \mathbf{N}^\dagger \mathbf{W}^\dagger \mathbf{a} \tag{15}$$

where

$$\mathbf{N}^\dagger = \begin{bmatrix} 0 & \mathbf{n} \\ \mathbf{n}\cdot & 0 \end{bmatrix}. \tag{16}$$

The adjoint source term on $\partial\Omega$, $\mathbf{g}_L$, produces adjoint wavefields comparable to those produced in the full domain. The adjoint of the localization operator takes the monopole and dipole source terms measured on the boundary of the local domain and converts them to pressure and particle velocity. The adjoint wave propagation operator uses the $\mathbf{g}_L$ source to produce the adjoint wavefield in the local domain, which matches the wavefield in the full domain, but is computed at a much lower cost.

An example illuminating wavefield localization is shown in Figure 3. The sources are indicated by black stars on the left of the model. For this demonstration we only consider the shallowest source point. The receivers are distributed along the vertical black line on the right of the model. The red circle represents the surface of the local domain and is where we measure the wavefield $\mathbf{u}$ to construct the sources for localization.

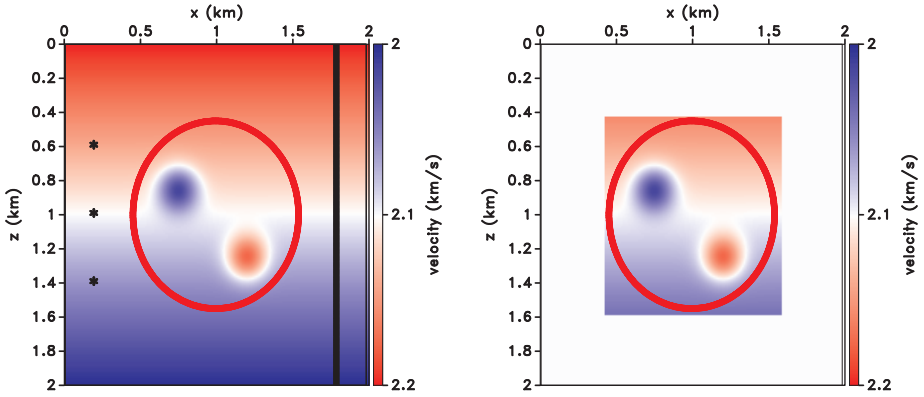
A comparison of the forward pressure wavefields in the full and local domains is shown by Figure 4. The wavefields inside the local domain match those inside the full domain to just under 10% in amplitude and exactly in phase. There are minor artifacts in the wavefields reconstructed near the boundary of the local domain, caused by the finite-difference derivative stencil, but they do not play a significant role in inversion. The sources on the boundary of the local domain cancel out the wavefield that leaves the domain. The adjoint wavefields in the full and local domains are shown in Figure 5, and match eachother to just under 10% in amplitude and exactly in phase.

4 ADJOINT SOURCE

In order to create the FWI adjoint source we need to construct a data difference between observed data, and data simulated with the current model. In the full domain, the data difference is taken at the location of the receivers. Using the local domain we have a choice of where to construct the adjoint source– at the receiver locations, as in the full domain, or at the boundary of the localized domain. If we compute and store the Green's functions between the local domain and the receivers, as in Willemsen et al. (2016), we could use the receiver location to compute the adjoint source. However, we have already stated that the Green's functions are too expensive to compute, store, and apply, so the only feasible place to create the adjoint source is on the boundary of the local domain. To accomplish this, we need a way to relocate the actual receivers of the survey to the boundary of the local domain.

A schematic representation of the full and local domain is shown in Figure 6, where $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are instances of operator $\mathbf{L}$ as in equation 3, and $\mathbf{M}$ is an instance of $\mathbf{N}$ in equation 16. Figure 6(a) shows the full domain scenario where we propagate from the sources to the receivers, calculate the adjoint source at the receivers and then propagate backwards with the adjoint. Figure 6(b) shows a redatuming of the sources on on the local domain as $\mathbf{d}_L$ to localize the state variable $\mathbf{u}$. The adjoint source is still computed at the receivers, as in Willemsen et al. (2016). In this case, we propagate through the local domain using the operator $\mathbf{B}$. Then we use the operators $\mathbf{M}$ and $\mathbf{C}$ to continue the wavefields to the receivers to create the adjoint source, which is propagated backwards using the adjoint operator $\mathbf{C}^\dagger$. Without access to the Green's functions, the operator chain from the local domain to the receivers and back needs to be computed every time the model is updated. This a high computational cost.

Instead, we would like to simulate the situation in Figures 6(c) and 6(d), where we construct $\mathbf{g}_L$ on the local domain and use that as the local adjoint source. To accomplish this we need to find the data on the local domain that would reproduce the observed data at the receivers when propagated using



(a) (b)

Figure 3. The full (left) and local (right) domains. The sources are the black stars on the left of the model and a line of vertical receivers are on the right. The red circle represents the local domain on which synthetic sources are derived for localization.

the operator $\mathbf{C}$. That is, instead of moving the receivers to the local domain using the adjoint, we need to use the inverse of operator $\mathbf{C}$. While the inverses of our wave equation and localization operators do not exist, we can approximate the inverse by solving an iterative least-squares problem. This inversion constructs pressure and particle velocity on the local domain such that when taken through the chain of operators $\mathbf{M}$ and $\mathbf{C}$ minimizes the difference with the observed data at the receivers. By redatuming using the inverse of the wave operator, we are moving our receivers backwards in space and time, undoing the propagation through the full domain until we reach the local domain. The advantage of redatuming through inversion is that we respect the velocity model between the local domain and receivers while undoing propagation through part of the model that does not change. The observed data we obtain on the local region is the same amplitude and phase as the predicted data on the same surface, so they can be subtracted to form the adjoint source $\mathbf{g}_L$. The redatum of the observed data only needs to be computed once, and is a relatively inexpensive operation compared to computing Green's functions, making the data difference at the local domain a more cost effective method.

5 GRADIENT OF THE VECTOR ACOUSTIC WAVE EQUATION

The objective function gradient is obtained by minimizing the augmented functional (Sava, 2014)

$$\mathcal{H} = \mathcal{J} - \mathcal{F}^{\dagger}\mathbf{a}. \tag{17}$$

Here, $\mathcal{J}$ the objective function we are minimizing (i.e. the data difference), $\mathbf{a}(\mathbf{x}, t)$ is the adjoint variable and $\mathcal{F}$ is the physical constraint indicating that the state variable $\mathbf{u}(\mathbf{x}, t)$ is the result of a linear process described by wave equation represented by operator $\mathbf{L}$:

$$\mathcal{F} = \mathbf{L}\mathbf{u} - \mathbf{d} = 0. \tag{18}$$

Using the augmented functional and physical constraints we can derive two gradients: one to update velocity while density is constant, and one to update both velocity and density.

5.1 Constant density

We consider the constant density version of equations 1 and 2, and write the model parameter for inversion as the slowness squared

$$m(\mathbf{x}) = \frac{1}{c^2(\mathbf{x})}. \tag{19}$$

The constrained gradient of the augmented functional in equation 17 is then

$$\frac{\partial \mathcal{H}}{\partial m} = \frac{\partial \mathcal{J}}{\partial m} - \sum_e \left[\left(\frac{\partial \mathcal{F}}{\partial m} \right)^{\dagger} \right] [\mathbf{a}] \tag{20}$$

and implies summation over experiments and time. The derivative of $\mathcal{F}$ with respect to the model parameter can be written

$$\frac{\partial \mathcal{F}}{\partial m} = \frac{\partial}{\partial m}(\mathbf{L}\mathbf{u} - \mathbf{d}) = \begin{bmatrix} 0 & 0 \\ 0 & \dfrac{\partial}{\partial t} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ p \end{bmatrix} = \frac{\partial p}{\partial t}. \tag{21}$$

The gradient is the zero time-lag cross-correlation of the state and adjoint variables:

$$\frac{\partial \mathcal{H}}{\partial m} = \frac{\partial \mathcal{J}}{\partial m} - \sum_e \sum_\tau \delta(\tau) \left(\dot{p} \star \mathbf{a}^q \right). \tag{22}$$

The dot over the pressure indicates the first time derivative. The superscript q of the adjoint variable denotes the scalar component corresponding to a monopole wavefield term.

The augmented functional $\mathcal{H}$ does not depend explicitly

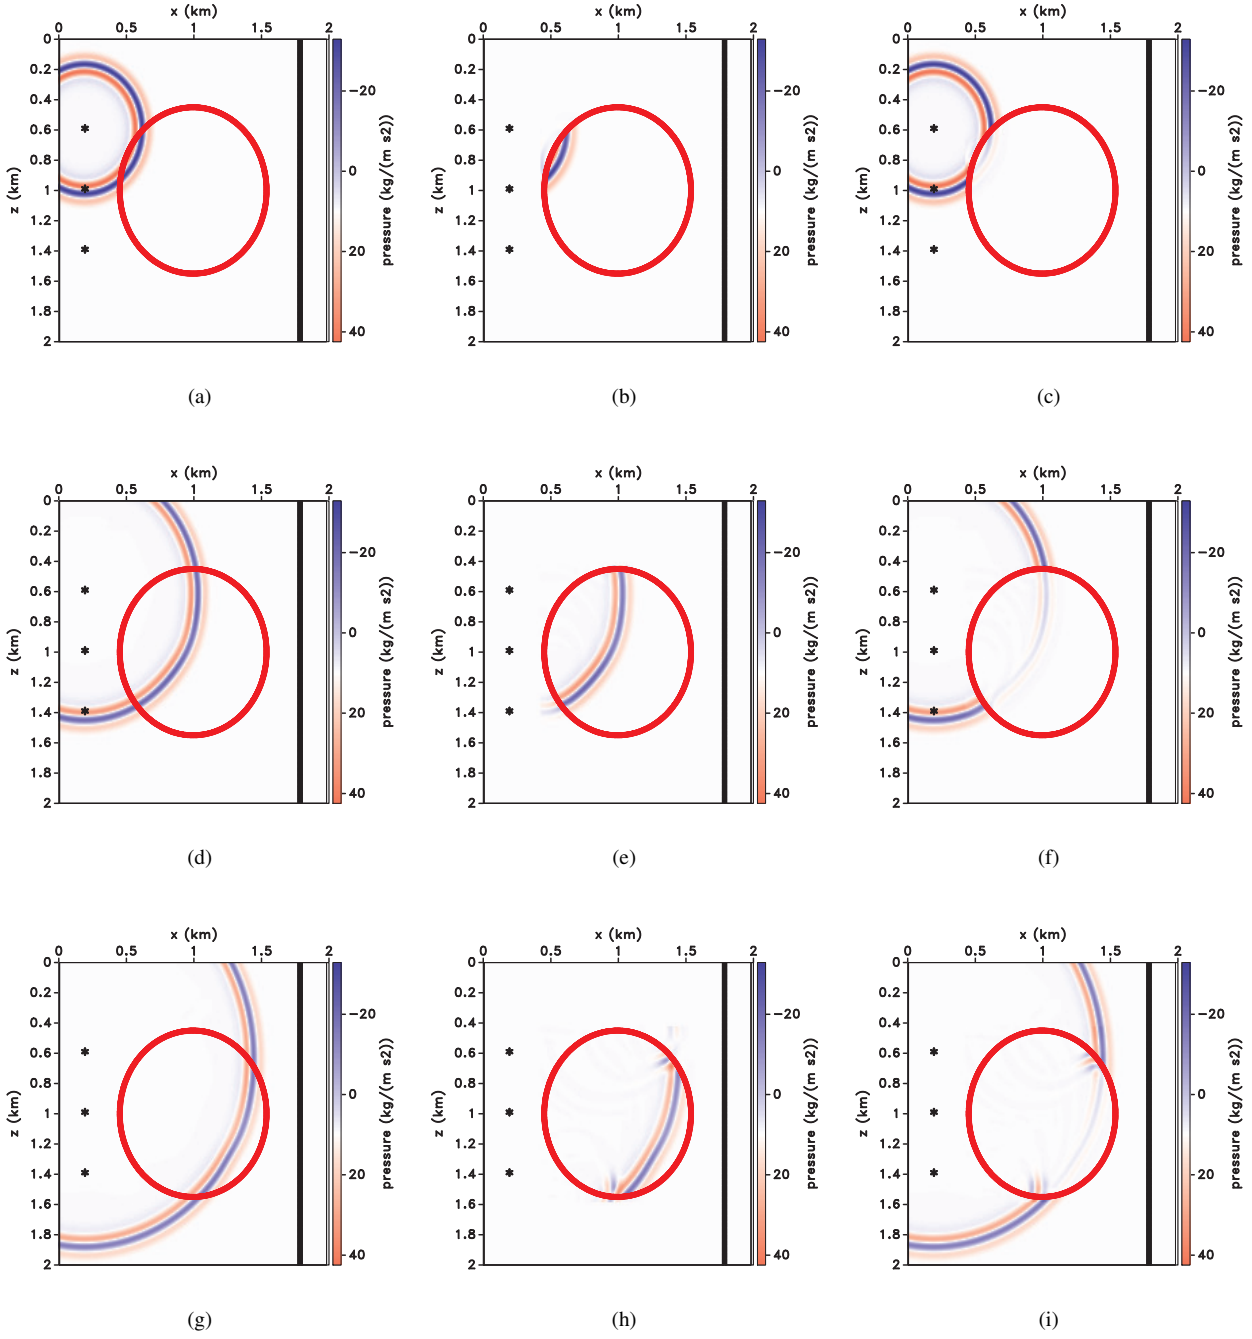


Figure 4. The forward pressure wavefields in the full (left) and local (middle) domains at three timesteps (increasing downward). A difference between the full and local domain is shown by the right column.

on the model parameter m, and therefore $\frac{\partial \mathcal{H}}{\partial m} = 0$, so we write the gradient of the objective function as

$$\frac{\partial \mathcal{J}}{\partial m} = \sum_e \sum_\tau \delta(\tau) \left(\dot{p} \star a^q \right). \tag{23}$$

The gradient in equation 23 can be used by the inversion algorithm to update the model parameter m, after which a new gradient can be computed from the state and adjoint variables. This gradient term using the scalar component of the wavefield is the same as proposed by Akrami et al. (2017), but correctly accounts for the time derivative on the state variable pressure.

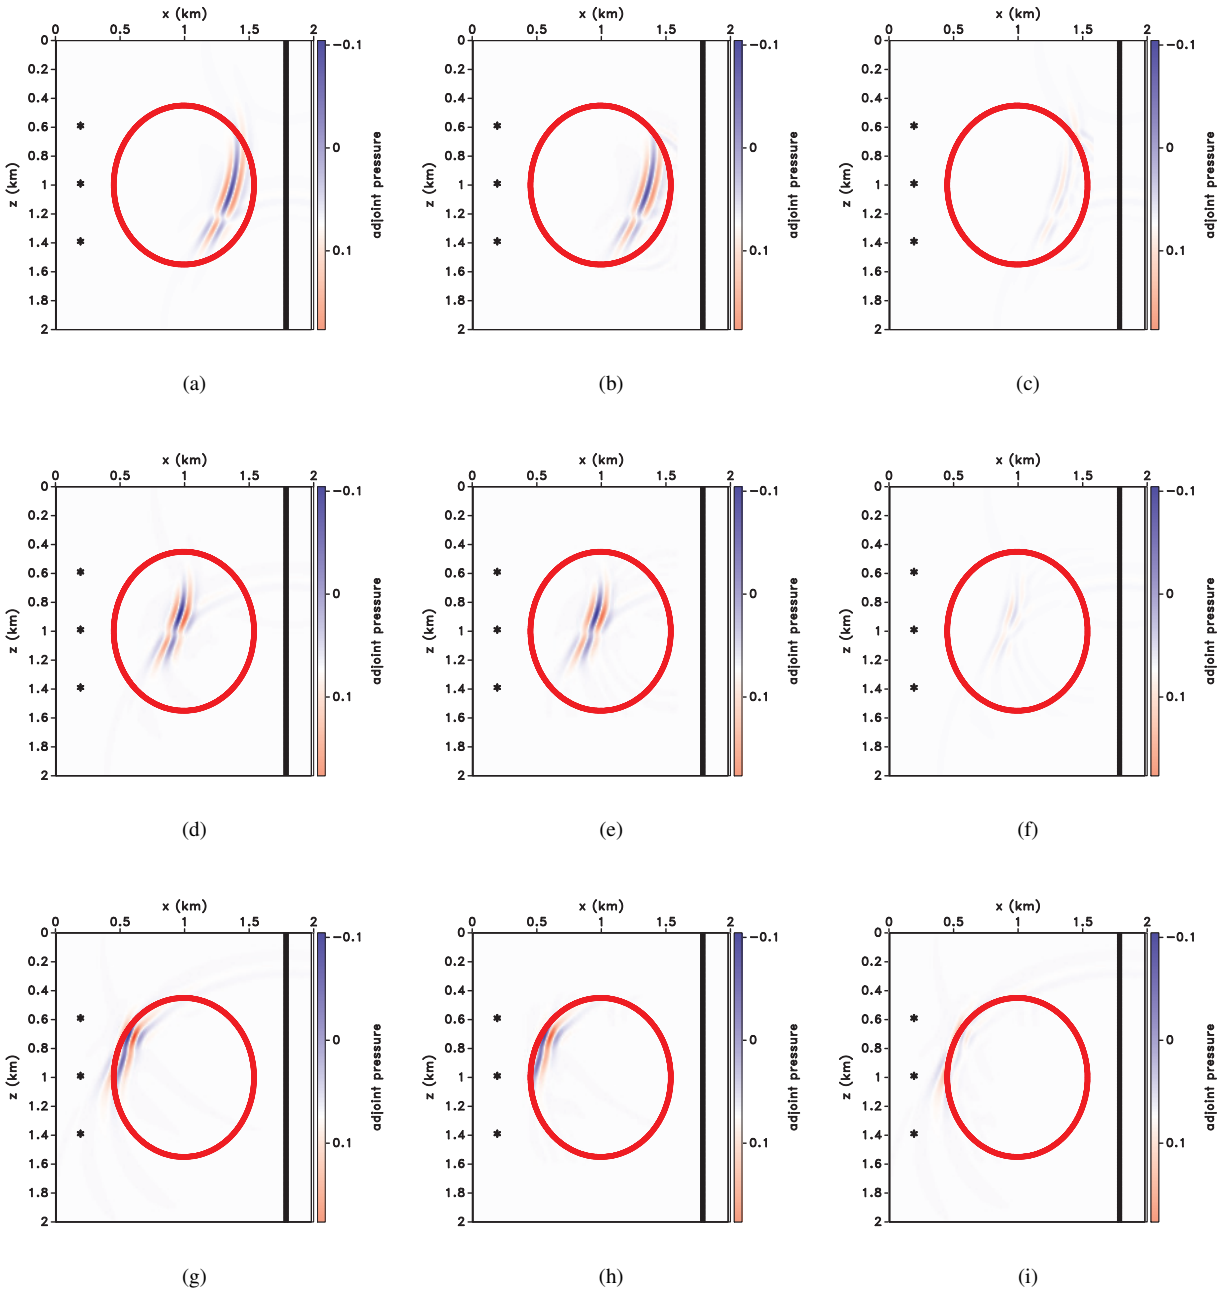
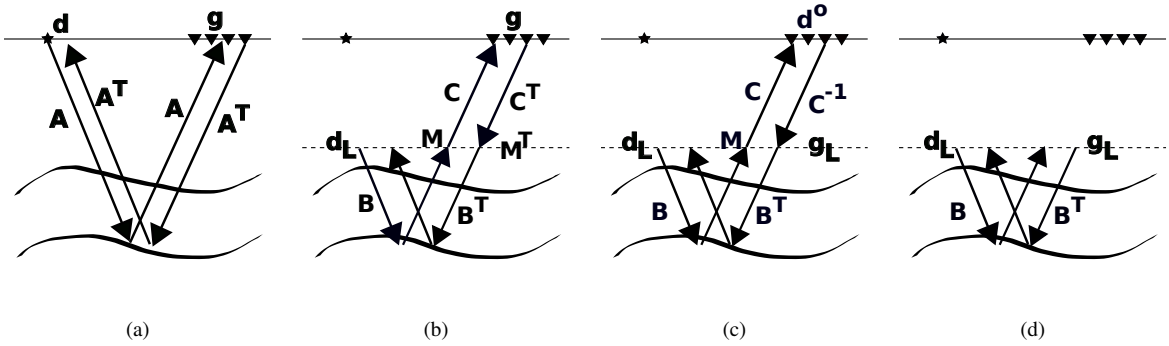


Figure 5. The adjoint pressure wavefields in the full (left) and local (middle) domains at three timesteps (increasing downward). A difference between the full and local domain is shown by the right column.

As an example, consider the velocity in Figure 3 to be the model used to compute the observed data at the receivers. The observed data for the local domain is a perfect redatum from the receivers to the surface of the local domain instead of solving the least squares problem in the preceding section. The background model is the same, but without the Gaussian anomalies, and is used to compute the state and adjoint variables. The gradients for velocity created using equation 23 for the full and local domain is shown by Figure 7. These gradients match in amplitude and in general shape. The high intensity areas of the gradient correspond with the locations of the Gaussian anomalies in the true model. There are also high intensities near the local domain borders, which influence the inversion, but those can be compensated for by incorporating an



(a) (b) (c) (d)

Figure 6. Pictorial representation of the operator chain for the a) full domain b) local domain with difference at the receivers c) the local domain with difference at the surface of the local domain. d) shows what the c) case reduces to after the inversion for $\mathbf{g}_L$ on the surface of the local domain is complete. $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are instances of operator $\mathbf{L}$ as in equation 3, and $\mathbf{M}$ is an instance of $\mathbf{N}$ in equation 16.

additional model objective function term (Sun and Li, 2015), as opposed to only using a data objective function in this example.

5.2 Variable density

Common practice for FWI is to use density as an unconstrained parameter to better match the data after updating the velocity model. This may yield geologically impossible density models which can be a hindrance if the FWI results are used for other applications, such as a starting model in gravity inversion. If we consider the density as an additional model parameter, we can constrain the density update in a way that makes geologic sense. When we consider the density model parameter to be constrained in addition to velocity, we define the model parameters by

$$m_1(\mathbf{x}) = \frac{1}{c(\mathbf{x})^2}, \tag{24}$$

$$m_2(\mathbf{x}) = \rho(\mathbf{x}). \tag{25}$$

Then, similarly to the preceeding case, we write the derivative of $\mathcal{F}$ with respect to each model parameter as

$$\begin{aligned}\frac{\partial \mathcal{F}}{\partial m_1} &= \frac{\partial}{\partial m_1}(\mathbf{Lu}-\mathbf{d}) \\ &= \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{m_2}\frac{\partial}{\partial t} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ p \end{bmatrix} = \frac{1}{m_2}\frac{\partial p}{\partial t}\end{aligned} \tag{26}$$

$$\begin{aligned}\frac{\partial \mathcal{F}}{\partial m_2} &= \frac{\partial}{\partial m_2}(\mathbf{Lu}-\mathbf{d}) \\ &= \begin{bmatrix} \frac{\partial}{\partial t} & 0 \\ 0 & -\frac{m_1}{m_2^2}\frac{\partial}{\partial t} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ p \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathbf{v}}{\partial t} \\ -\frac{m_1}{m_2^2}\frac{\partial p}{\partial t} \end{bmatrix}.\end{aligned} \tag{27}$$

We then find the gradients of the object function with respect to both model parameters as

$$\frac{\partial \mathcal{J}}{\partial m_1} = \frac{1}{m_2}\sum_e \sum_\tau \delta(\tau)\left(\dot{p} \star \mathbf{a}^q\right) \tag{28}$$

$$\frac{\partial \mathcal{J}}{\partial m_2} = \sum_e \sum_\tau \delta(\tau)\left(\sum_i \left(\dot{\mathbf{v}}^i \star \mathbf{a}^i\right) - \frac{m_1}{m_2^2}\left(\dot{p} \star \mathbf{a}^q\right)\right). \tag{29}$$

The derivative with respect to m_1 is a scaled version of what we found in the constant velocity case. The derivative with respect to m_2 has components derived from both vector and scalar components of the state and adjoint variables. The summation over i is for all vector components of the state and adjoint fields. Equations 28 and 29 suggest that the particle velocity wavefield does not influence the velocity model, but all components of the wavefield contribute to the density model. Updating the model parameters is a two-step process. First, we use equation 28 to update the velocity model, then we use equation 29 to update the density model. Then we recompute the state and adjoint variables, as well as the gradient for the next iteration of the inversion.

6 SYNTHETIC EXAMPLE

We consider a synthetic model with the background and true velocity models shown in Figure 8. The density model is constant. The sources are on the left side of the model and the receivers are indicated by the vertical line on the right. The local domain is the region of the model between the red and blue vertical lines. The red line represents the surface to which sources are redatumed, from which we generate the state variable. The blue line is the surface to which the receivers are redatumed, from which we create the adjoint variable. In this case, we are assuming a perfect redatum onto the blue surface, as opposed to the least-squares process described above, to eliminate errors from the iterative redatum process.

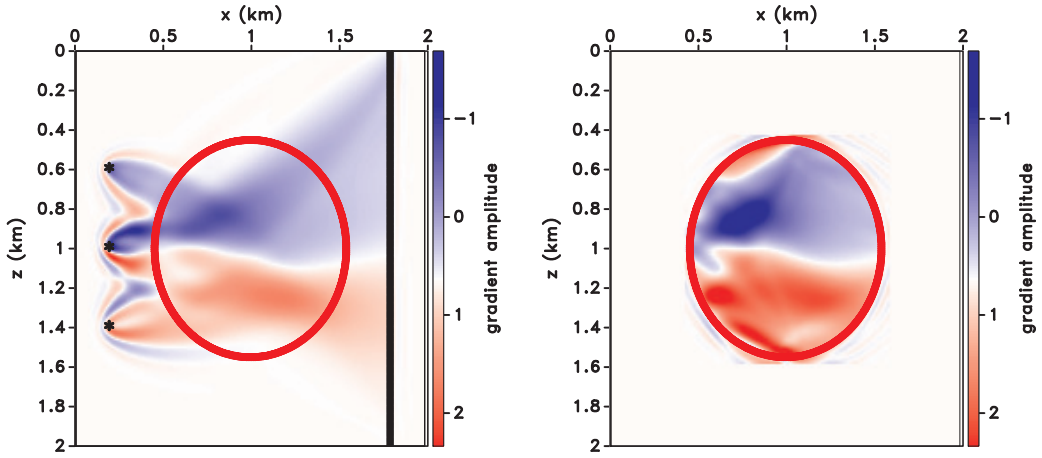


Figure 7. Gradients for the velocity update in the full (a) and local (b) domains. The density model is constant.

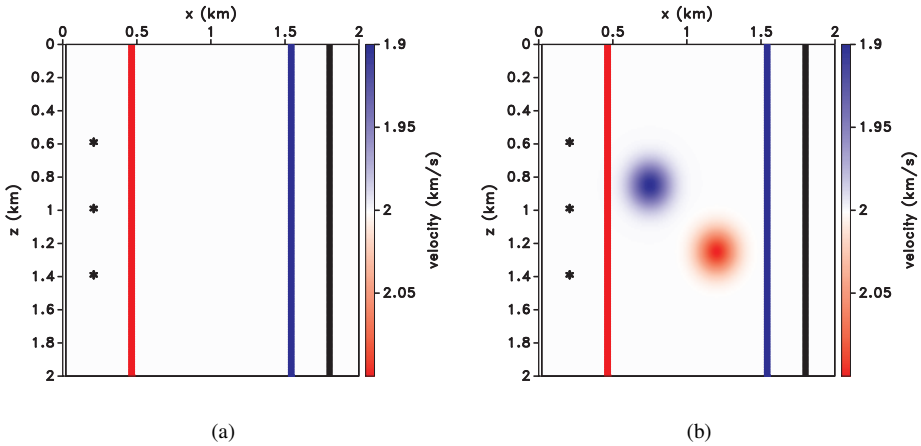


Figure 8. The a) background (starting) velocity model for inversion and b) true velocity model in the full domain. The local domain is the region between the red and blue vertical lines.

We invert for the velocity model that minimizes the data difference at the receivers using a limited memory BFGS algorithm (Nocedal and Wright, 2006). After 10 iterations we obtain the models shown in Figure 9. The high intensity regions of the full domain model correspond with the positive and negative Gaussian anomalies. The local domain inversion also has the Guassian anomalies, but includes artifacts that streak throughout the model. These artifacts could be mitigated by incorporating additional constraints into the model objective functions, such as a model smoothness constraint. The data objective functions as a function of iteration are shown in Figure 10 for the full and local domains. The full domain decreases and starts tapering off near the last iteration. The local domain tapers off to a flat region after about 3 iterations, and obtains a worse data misfit than the full domain, but still recovers the Gaussian anomalies.

7 CONCLUSION

In applications where the model only changes in a local region of interest, it is much more cost efficient to use a local domain for wavefield simulations instead of the full domain. The forward and adjoint wavefields inside the local domain mimic the behavior of the wavefields in the full domain by using the method of multiple point sources, even after model alterations. The local FWI gradient can be constructed using these local wavefields, and be used to update the model. By

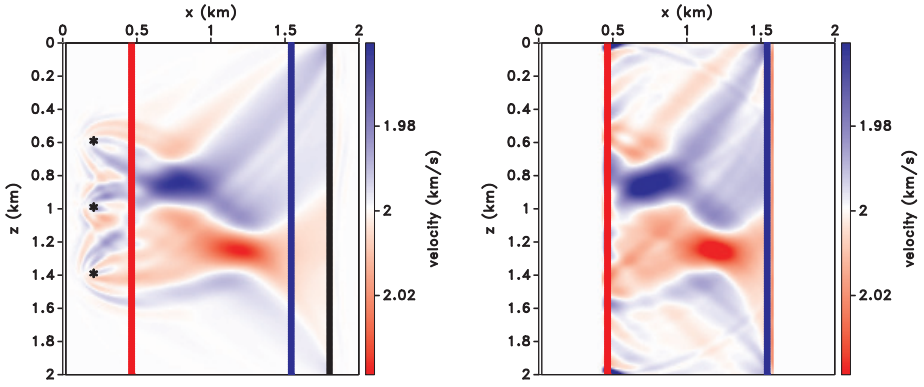


Figure 9. Final velocity model for the full (a) and local (b) domain inversions after 10 iterations of L-BFGS.

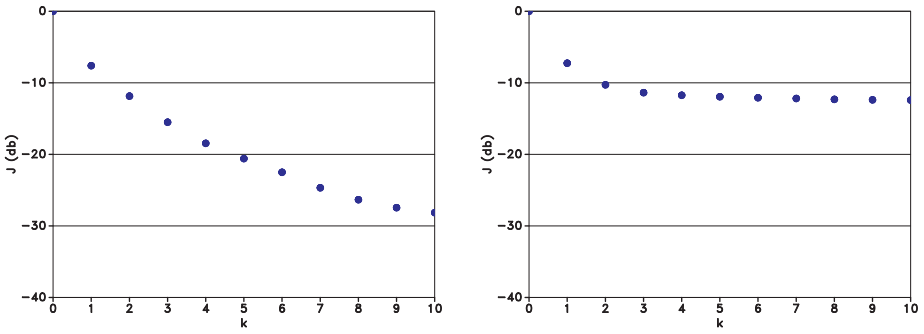


Figure 10. Data-difference objective function for the a) the full domain and b) the local domain inversions for 10 iterations of L-BFGS.

using the vector-acoustic wave equations and the adjoint-state method, we can construct FWI gradients for the velocity and density models, updating both to minimize the data objective function. Traditionally, the density model is used as a fudge factor to match features in the data that the velocity can not explain, but the vector-acoustic wave equation gives a way to incorporate the term into the model update.

The cost reduction allowed by using the local domain depends on its relative size relative to the full domain. The main speedup is due to shrinking the number of operations needed to compute the spatial derivatives on a large finite-difference grid. A secondary speedup occurs by lowering the timesteps for modeling, as we do not need to simulate wavefields before the source reaches the domain of interest or after the wavefields leave the domain. These speedups benefit every iteration of the inversion, allowing for more complex inversions to be computed in less time than if they were computed using the full domain.

8 ACKNOWLEDGEMENTS

We thank the sponsors of the Center for Wave Phenomena, whose support made this research possible. The synthetic examples in this paper use the Madagascar open-source software package Fomel et al. (2013) freely available from `http://www.ahay.org`.

REFERENCES

Akrami, S. M., P. Zheglova, and A. Malcolm, 2017, An algorithm for vector data full-waveform inversion: SEG Technical Program Expanded Abstracts 2017, Society of Exploration Geophysicists, 1568–1572.

Broggini, F., M. Vasmel, J. O. A. Robertsson, and D.-J. van Manen, 2017, Immersive boundary conditions: Theory, implementation, and examples: Geophysics, **82**, T97–T110.

de Hoop, A., 1988, Time-domain reciprocity theorem for acoustic wave fields in fluid with relaxation.

de Hoop, A. T., 1995, Handbook of radiation and scattering of waves: Acoustic waves in fluids, elastic waves in solids, electromagnetic waves: Academic Press.

Elison, P., D. van Manen, J. Robertsson, M. Dukalski, and K. de Vos, 2018, Marchenko-based immersive wave simulation: Geophysical Journal International, **215**, 1118–1131.

Fleury, C., and I. Vasconcelos, 2013, Adjoint-state reverse time migration of 4C data: Finite-frequency map migration for marine seismic imaging: Geophysics, **78**, WA159–WA172.

Fomel, S., P. Sava, I. Vlad, Y. Liu, and V. Bashkardin, 2013, Madagascar: open-source software project for multidimensional data analysis and reproducible computational experiments: Journal of Open Research Software, **1**, e8.

Halliday, D., J. O. A. Robertsson, I. Vasconcelos, D.-J. van Manen, R. Laws, K. Özdemir, and H. Grønaas, 2012, Full-wavefield, towed-marine seismic acquisition and applications: SEG Technical Program Expanded Abstracts 2012, Society of Exploration Geophysicists, 1–5.

Kotsi, M., A. Malcolm, and G. Ely, 2018, 4D full-waveform Metropolis Hastings inversion using a local acoustic solver: SEG Technical Program Expanded Abstracts 2018, Society of Exploration Geophysicists, 5323–5327.

Nocedal, J., and S. Wright, 2006, Numerical optimization, *in* Springer.

Osorio, L. E. J., and A. Malcolm, 2018, Amplitude variation with offset (AVO) inversion modeling with a local elastic solver: SEG Technical Program Expanded Abstracts 2018, Society of Exploration Geophysicists, 625–629.

Plessix, R.-E., 2006, A review of the adjoint-state method for computing the gradient of a functional with geophysical applications: Geophysical Journal International, **167**, 495–503.

Pratt, R. G., 1999, Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale model: Geophysics, **64**, 888–901.

Ravasi, M., I. Vasconcelos, A. Curtis, and A. Kritski, 2015, Vector-acoustic reverse time migration of Volve ocean-bottom cable data set without up/down decomposed wavefields: Geophysics, **80**, S137–S150.

Robertsson, J. O., I. Moore, M. Vassallo, K. Özdemir, D.-J. van Manen, and A. Özbek, 2008, On the use of multicomponent streamer recordings for reconstruction of pressure wavefields in the crossline direction: Geophysics, **73**, A45–A49.

Robertsson, J. O. A., and C. H. Chapman, 2000, An efficient method for calculating finite-difference seismograms after model alterations: Geophysics, **65**, 907.

Sava, P., 2014, A comparative review of wavefield tomography methods: CWP Annual Report, **CWP**, 119–144.

Shabelansky, A. H., A. Malcolm, and M. Fehler, 2013, Data-driven estimation of the sensitivity of target-oriented time-lapse seismic imaging to source geometry: Geophysics, **78**, R47–R58.

Sun, J., and Y. Li, 2015, Multidomain petrophysically constrained inversion and geology differentiation using guided fuzzy c -means clustering: Geophysics, **80**, ID1–ID18.

Symes, W. W., 2007, Reverse time migration with optimal checkpointing: Geophysics, **72**, SM213–SM221.

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

van Manen, D.-J., J. O. A. Robertsson, and A. Curtis, 2007, Exact wave field simulation for finite-volume scattering problems: The Journal of the Acoustical Society of America, **122**, EL115–EL121.

Vasmel, M., and J. O. A. Robertsson, 2016, Exact wavefield reconstruction on finite-difference grids with minimal memory requirements: Geophysics, **81**, T303–T309.

Vasmel, M. L., 2016, Immersive boundary conditions for seismic wave propagation: Doctoral thesis, ETH Zurich.

Virieux, J., and S. Operto, 2009, An overview of full-waveform inversion in exploration geophysics: Geophysics, **74**, WCC1–WCC26.

Willemsen, B., 2017, Problems with a localized nature in exploration seismology: Doctoral thesis, Massachusetts Institute of Technology.

Willemsen, B., and A. Malcolm, 2016, Coupling a local elastic solver to a background acoustic model to estimate phase variation: SEG Technical Program Expanded Abstracts 2016, Society of Exploration Geophysicists, 3945–3949.

Willemsen, B., A. Malcolm, and W. Lewis, 2016, A numerically exact local solver applied to salt boundary inversion in seismic full-waveform inversion: Geophysical Journal International, **204**, 1703–1720.

Yang, D., Y. Zheng, M. Fehler, and A. Malcolm, 2012, Target-oriented time-lapse waveform inversion using virtual survey: SEG Technical Program Expanded Abstracts 2012, Society of Exploration Geophysicists, 1–5.

Yee, K. S., 1966, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media: IEEE Transactions on Antennas and Propagation, **14**, 302–307.

Zhong, Y., and Y. Liu, 2019, Source-independent time-domain vector-acoustic FWI: Geophysics, 1–50.