Geologically Constrained Full Waveform Inversion With Preconditioning

Antoine Guitton, GeoImaging Solutions Inc., Gboyega Ayeni, Stanford University

Abstract

The waveform inversion problem is inherently ill-posed. Traditionally, regularization terms are used to address this issue. For waveform inversion where the model is expected to have many details reflecting the physical properties of the Earth, regularization and data fitting can work in opposite directions: the former smoothing and the later adding details to the model. In this paper, we constrain the velocity model with a model-space preconditioning scheme based on directional Laplacian filters. This preconditioning strategy preserves the details of the velocity model while smoothing the solution along known geological dips. The Laplacian filters have the property to smooth or kill local planar events according to a local dip field. By construction, these filters can be inverted and used in a preconditioned waveform-inversion scheme to yield geologically meaningful models. We illustrate on synthetic and field data examples how preconditioning with non-stationary directional Laplacian filters outperforms traditional waveform inversion when sparse data are inverted for and when salt is present. Preconditioning could benefit waveform inversion of real data where irregular geometry, coherent noise and lack of low frequencies are present.

Introduction

The goal of waveform inversion is to derive physical properties of the Earth, such as P-wave velocity, S-wave velocity, or density. These properties can be related to the presence of hydrocarbons in the subsurface and their estimation is one of the most important goal in seismic processing. In practice, we try to minimize the function

$$f(\mathbf{m}) = \|\mathbf{u}_{obs} - \mathbf{u}_{mod}\|^{norm} \tag{1}$$

where $\mathbf{m}$ is a vector of physical properties (what we are looking for), $\mathbf{u}_{obs}$ the observed and $\mathbf{u}_{mod}$ the modeled data. Note that f is a 1-D function defined by the choice of a norm. In practice, the ℓ^2 norm is often chosen, but the ℓ^1 norm seems to have more practical needs for its robustness to non-Gaussian noise present in the data (Crase et al., 1990). The minimization of $f(\mathbf{m})$ can be achieved using iterative methods. Tarantola (1984) establishes that the model can be updated as follows:

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \alpha_n . \nabla f(\mathbf{m}_n) \tag{2}$$

where $\nabla f(\mathbf{m})$ is the gradient of $f(\mathbf{m})$ and α_n a step-length that needs to be estimated. It turns out that the computation of the gradient is equivalent to a reverse-time migration (Lailly, 1984).

Unfortunately, although a promising approach, waveform inversion is hampered by many issues. The main one is the presence of local minima in f. To circumvent this problem, the data can be inverted in a multi-scale fashion in the time (Bunks et al., 1995) or frequency domain Sirgue and Pratt (2004). Second, time damping of the input data offer opportunities to focus the inversion on different parts of the data, thus reducing the effects of local minima Brenders et al. (2009).

Traditionally, ill-posed problems can be solved by adding a regularization term to the objective function. Very often, a regularization term that can penalize differences between neighboring points is selected. However, whereas waveform inversion tends to add details to a velocity model, regularization tends to smooth them out, thus working against our primary goal: fitting the data. One way to address these somewhat conflicting goals is to use preconditioning. Here, we show how we can geologically constrain the velocity model by using a non-stationary preconditioning approach. This method requires two ingredients: a dip estimation method and a local dip filtering technique. We use the method of Fomel (2002) for the former and of Hale (2007) for the later.

In this paper we first introduce the waveform inversion approach, with and without preconditioning. We show that preconditioning amounts to a simple change of variable which, in effect, changes the gradient direction. Then, we present our method of local dip filtering, which follows Hale's. Finally, we present synthetic and field data results. These results demonstrate that non-stationary, preconditioned inversion yields geologically plausible models.

Method

In this paper, we use a time domain approach for solving the scalar acoustic wave equation (parametrized in terms of P-wave velocity v_p):

$$\frac{\partial^2 u(\mathbf{x},t)}{\partial t^2} - v_p(\mathbf{x})^2 \nabla^2 u(\mathbf{x},t) = v_p(\mathbf{x})^2 s(\mathbf{x},t). \tag{3}$$

with the source term $s(\mathbf{x},t) = \delta(\mathbf{x} - \mathbf{x}_s) S(t)$ where $S(t)$ is the source function at $\mathbf{x}_s$ and $u(\mathbf{x},t)$ the pressure field. Tarantola (1984) derives the expression of the gradient for the acoustic equation (3) for each component of $\mathbf{m}$ (equal to v_p only in this case).

$$\nabla f(\mathbf{m}_n) = \frac{2}{{\mathbf{v}_p}_n{}^3} \sum_{shots} \sum_t \frac{\partial^2 \overrightarrow{\mathbf{u}_n}}{\partial t^2} \cdot \overleftarrow{\delta \mathbf{u}_n} \tag{4}$$

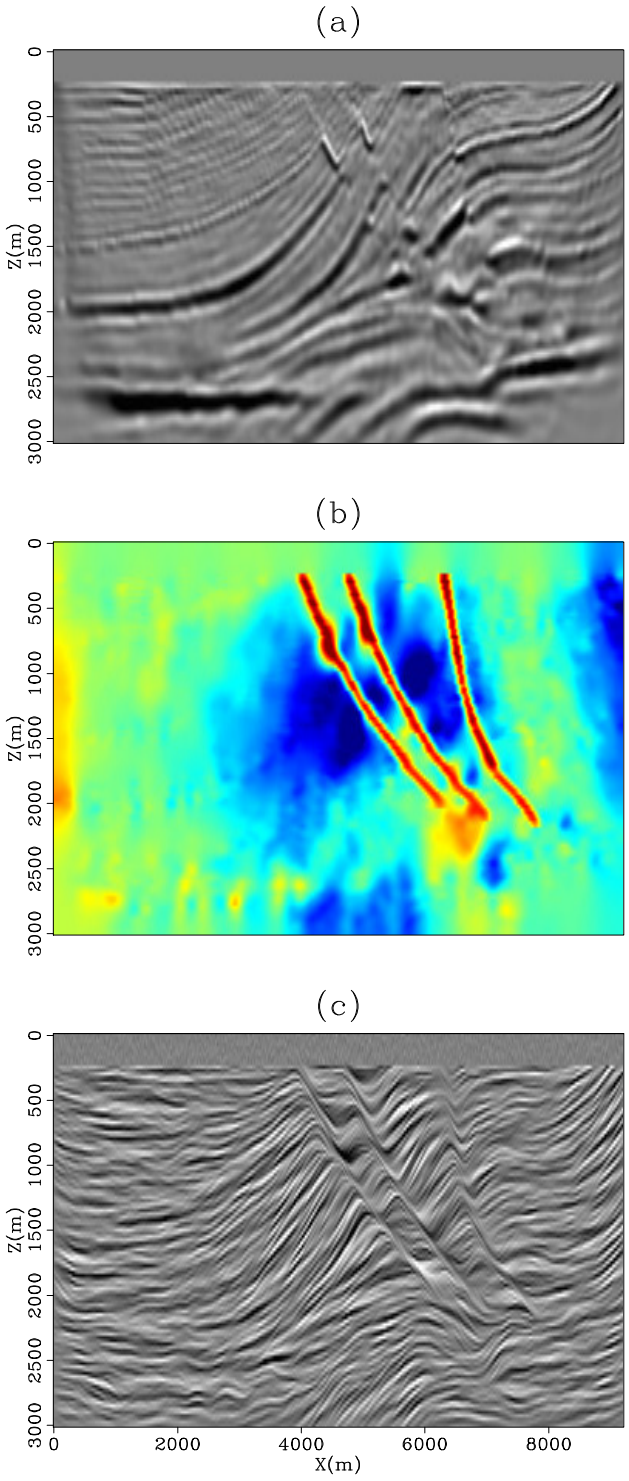

Figure 1: Illustration of the preconditioning operator $\mathbf{S}$. We first migrate the synthetic Marmousi dataset in (a). (b) shows the dip field estimated from the migration result in (a). The faults are picked by hand and smoothed horizontally to preserve discontinuities in the velocity model. Applying $\mathbf{S}$ estimated from (b) to a random field shows in (c) the texture of the migrated section in (a). Note how the water-column is not smoothed and how the fault locations are clearly visible.

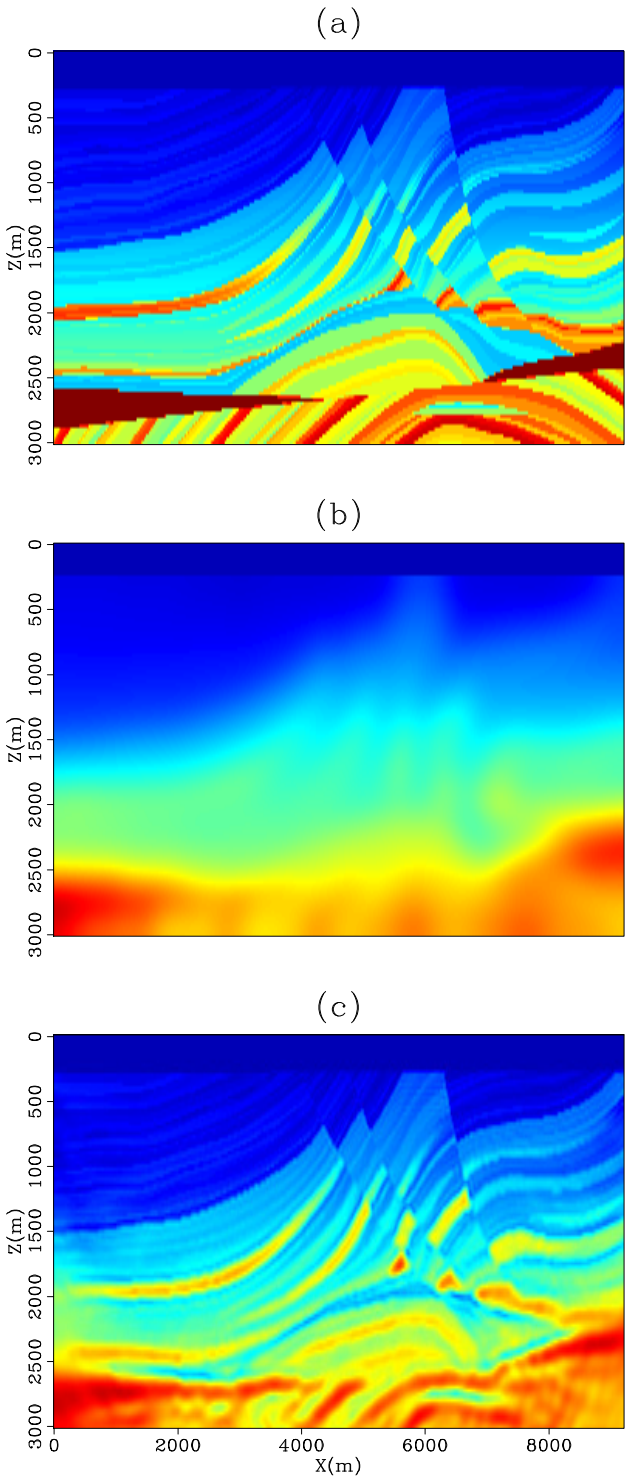

Figure 2: (a) True velocity model used to generate the synthetic dataset. (b) Initial guess obtained by smoothing the true model in (a). (c) Estimated model after waveform inversion. No preconditioning is applied in this case. Four frequency bands were used to bandpass both the source and the data prior to inversion (0-4Hz, 0-8Hz, 0-12Hz, 0-16Hz). The velocity model is recovered very well.

where $\overleftarrow{\delta \mathbf{u}_n}$ is the *backward* propagated residual at iteration n such that $\delta \mathbf{u}_n = \mathbf{u}_{obs} - \mathbf{u}_n$ and $\overrightarrow{\mathbf{u}_n}$ is the *forward* propagated synthetic source. For our iterative method, we opted for the L-BFGS approach of Nocedal (1980). This quasi-Newton approach computes an estimate of the inverse Hessian iteratively by using a user-defined number of solution and gradient vectors. One of the main benefit of this technique is that the Hessian is never explicitely formed, thus involving significant memory savings. With the L-BFGS solver, the model is updated as follows:

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \alpha_n \mathbf{H}_n^{-1} \nabla f(\mathbf{m}_n), \tag{5}$$

where $\mathbf{m}_{n+1}$ is the updated solution, α_n the step length computed by a line search that ensures a sufficient decrease of $f(\mathbf{m})$ and $\mathbf{H}_n \approx \nabla^2 f(\mathbf{m}_n)$ the approximate Hessian. To improve chances of not falling into a local minimum, we follow a multi-scale approach Bunks et al. (1995) where the source and data are bandpassed prior to inversion. We introduce our preconditioning scheme in the following section.

Preconditioning

Preconditioning amounts to a change of variable $\mathbf{m} = \mathbf{S}\mathbf{p}$ where $\mathbf{p}$ is a new variable used for the inversion and $\mathbf{S}$ the preconditioning operator. In our case, this operator amounts to a non-stationary deconvolution with local dip-filters. Having introduced the new variable $\mathbf{p}$, the iterative scheme in Equation (5) becomes:

$$\mathbf{p}_{n+1} = \mathbf{p}_n - \alpha_n \tilde{\mathbf{H}}_n^{-1} \nabla \tilde{f}(\mathbf{p}_n), \tag{6}$$

where

$$\nabla \tilde{f}(\mathbf{p}_n) = \mathbf{S}' \nabla f(\mathbf{m}_n) = \mathbf{S}' \nabla f(\mathbf{S}\mathbf{p}_n) \tag{7}$$

and $\mathbf{S}'$ is the transpose of $\mathbf{S}$. Therefore, with preconditioning, we obtain a new gradient direction. In our scheme, we will opt for a preconditioning operator that steers the gradient toward a geologically constrained solution. Note that in Equation (6) the approximate Hessian in equation (5) is blind to this change of variable as it is built from gradient and solution step vectors only. Assuming that we can estimate $\mathbf{S}$ and compute its adjoint and inverse (to accommodate any starting guess $\mathbf{p}_0 = \mathbf{S}^{-1}\mathbf{m}_0$), preconditioning can be easily introduced in any waveform inversion scheme. Once a solution vector $\mathbf{p}_{sol}$ has been found, the final model is obtained by computing

$$\mathbf{m}_{sol} = \mathbf{S}\mathbf{p}_{sol}. \tag{8}$$

Now we present our choice of preconditioning operator $\mathbf{S}$.

Defining the preconditioning operator S

Preconditioning amounts to a change of the gradient direction. For waveform inversion, a gradient that embeds some geological information could help yielding more meaningful velocity models. To this end, we follow the approach of Hale (2007) for the construction of the operator $\mathbf{S}$. Doing so, this operator becomes a non-stationary deconvolution with directional Laplacian filters.

Directional Laplacian filters are built from small wavekill filters $\mathbf{A}$, similar to those of Claerbout (1995). Wavekill filters have the ability to anihilate local planar-events with a given dip. From these filters, new operators $\mathbf{A}'\mathbf{A}$ are formed by autocorrelation. These new operators are then factorized into minimum-phase filters $\tilde{\mathbf{A}}$ such that $\tilde{\mathbf{A}}'\tilde{\mathbf{A}} \approx \mathbf{A}'\mathbf{A}$. Having minimum-phase filters, we can build a stable non-stationary deconvolution operator $\mathbf{S} = \tilde{\mathbf{A}}^{-1}\tilde{\mathbf{A}}'^{-1}$ and its inverse $\mathbf{S}^{-1} = \tilde{\mathbf{A}}'\tilde{\mathbf{A}}$. Because the wavekill filter $\mathbf{A}$ is dip dependent, the operator $\mathbf{S}$ has the ability to smooth along a given direction as well. Therefore, if we can estimate a dip field that contains some desired geological features, the directional Laplacian filters can make the solution of an inverse problem resemble these.

In practice, we estimate and use the directional Laplacian filters as follows: first, we estimate a dip field following the approach of Fomel (2002); then we estimate a bank of directional Laplacian filters for a range of angles; finally we apply the appropriate inverse Laplacian filter on each sample according to the local dip. One added feature of our preconditioning scheme is that the strength of the smoothing can be controlled spatially: different areas with similar dips can be smoothed across different distances. These areas are identified by a weighting function which varies from high values (i.e., little smoothing) to low values (i.e., strong smoothing).

To illustrate the preconditioning operator $\mathbf{S}$, we show in Figure 1a the migration result of a synthetic dataset based on the 2-D Marmousi model. This result is obtained with Reverse Time Migration (RTM). In real data cases, the dip field could be re-estimated iteratively from a migrated image estimated with the updated velocity field, adding a third outside loop to our waveform inversion algorithm (one for frequency band and one for the muting/masking operator). This possibility is not investigated in this paper.

From the RTM image, we can estimate the local dip field (Figure 1b). This dip field is obtained iteratively with some smoothing using the method of Fomel (2002). We also picked by hand the location of three faults. From these picks, we estimated the dip on the fault and smoothed the local dip horizontally. These three faults were picked to get sharper velocity contrasts. Now, we apply the operator $\mathbf{S}$ to a random field the size of the migrated image in Figure 1a to obtain Figure 1c. We notice that the "texture" of the original migrated image is recovered and that no smoothing is applied in the water layer: for this example, our weigthing function had only two values separated by the water bottom. Finally, we can clearly identify the fault locations. In the next section, we demonstrate that this operator can be used to constrain the solution of waveform inversion.

Examples

We illustrate the geologically-constrained waveform inversion method on a synthetic and field dataset.

Modified Marmousi example

First, we modified the Marmousi 2-D velocity model by adding a 250 m. thick water layer (Figure 2a). We created 184 shots 50 m. apart with a fixed receiver array (369 in total) at the surface using the scalar wave equation (no density). The source is a Ricker-2 wavelet with a fundamental frequency of 8Hz. To illustrate that our inversion works (without preconditioning), we show in Figure 2c the result of waveform inversion with four frequency scales (0-4Hz, 0-8Hz, 0-12Hz, and 0-16Hz) using the starting guess in Figure 2b (obtained by

smoothing the true model in Figure 2a) and using all the shots. There is a very good match between the inverted and true model. Because all the data was used, little would be gained by using preconditioning.

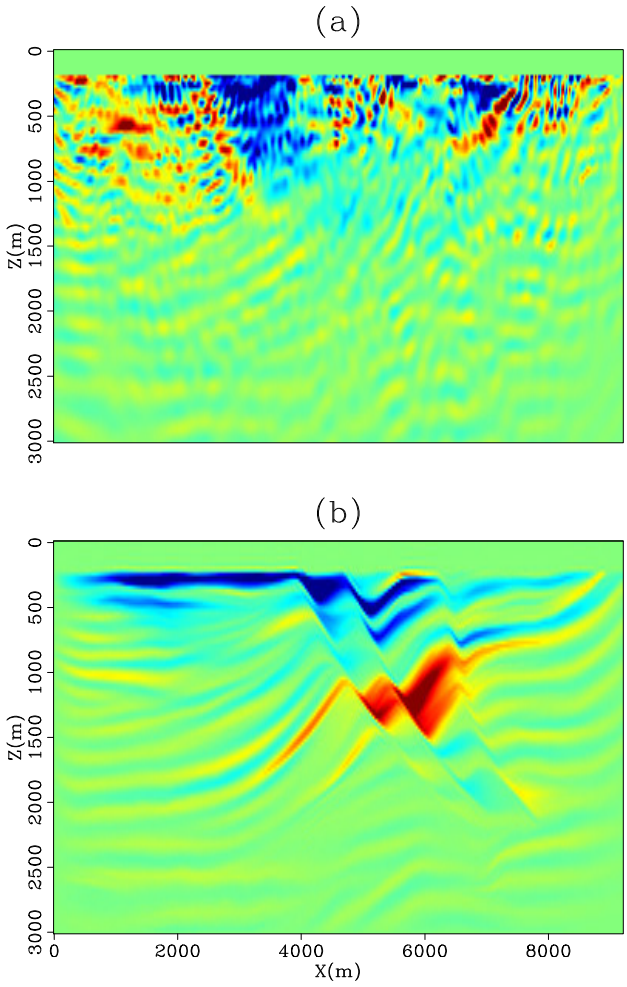

Figure 3: (a) Gradient $\nabla f(\mathbf{m}_n)$ of the unpreconditioned inversion after 4 iterations of the 0-8Hz scale (4 shots, 2.5 Km. apart). (b) Reprojected gradient $\mathbf{S}\nabla\tilde{f}(\mathbf{p}_n)$ of the preconditioned inversion after 4 iterations of the 0-8Hz scale (4 shots, 2.5 Km. apart). With preconditioning, the gradient follows the information captured in the dip field and looks more geologically appealing than in (a).

To make a compelling case, we kept only four shots, 2.5 Km. apart. First, we show in Figure 3 a comparison between the gradient without preconditioning $\nabla f(\mathbf{m}_n)$ and the gradient with preconditioning back in the velocity space $\mathbf{S}\nabla\tilde{f}(\mathbf{p}_n)$. Because only 4 shots are present, the unpreconditioned gradient looks noisy and resemble geology in very few locations only. Some authors suggest attenuating the high wavenumbers in the gradient by smoothing it after each iteration (Ravaut et al., 2004), where the size of the smoothing operator in the horizontal and vertical directions is a function of an average wavelength at a given frequency. This bears a resemblance with our proposed scheme but doesn't allow for directional smoothing. On the contrary, thanks to the preconditioning with directional Laplacian filters, the reprojected gradient in Figure 3b shows the geology captured in the dip field of Figure 1b very well.

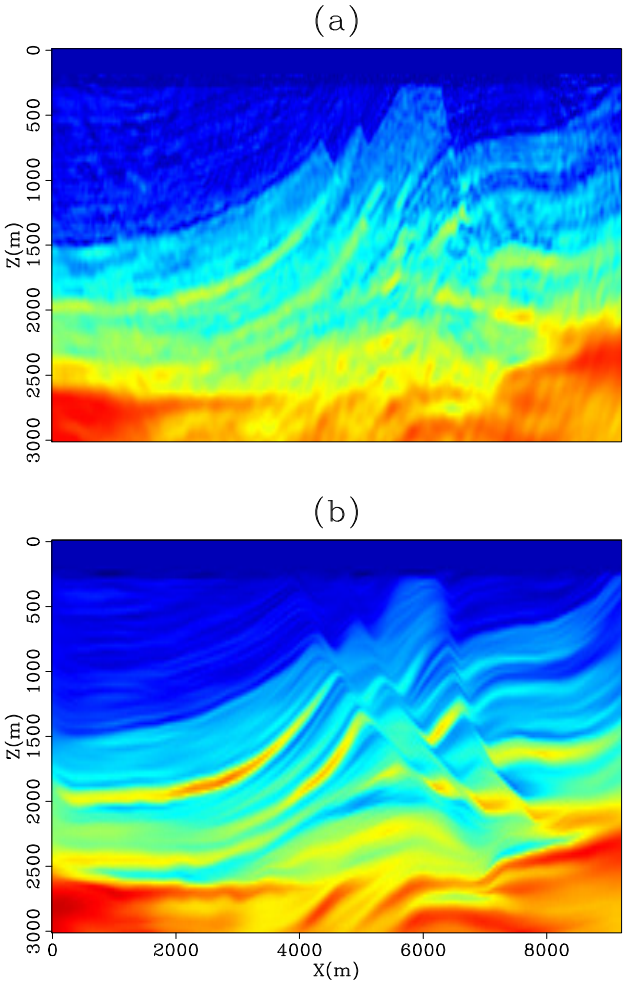

Figure 4: (a) Inversion result for the unpreconditioned inversion. (b) Inversion result for the preconditioned inversion with directional Laplacian filters. The geology at the reservoir level is recovered very well in (b).

Now, we show in Figure 4 the inversion results for the unpreconditioned (Figure 4a) and preconditioned inversion (Figure 4b). Although quite noisy, the unpreconditioned result shows the geology very well: the presence of low frequencies in the data, along with the multi-scale approach, act as regularization operators. This effect will be less pronounced with real data where low frequencies are often missing. The preconditioned inversion result in Figure 4b is much cleaner: the geology at the reservoir level is recovered very well.

Gulf of Mexico example

This sections shows the results of our proposed scheme applied to a 2-D Gulf of Mexico dataset. This dataset, which has been extensively used in the past to benchmark multiple attenuation techniques (special edition of The Leading Edge, January 1999), comes from the Mississippi Canyon. For our inversion, we decided to keep the multiples in the data. This clearly increases the non-linearity of the problem with the advantage of not including processing artifacts during the inversion (Sirgue et al.,

2009).

The starting velocity model was derived in two steps. First, we applied a sediment and salt flood procedure to define a rough model. Second, we performed two iterations of reflection tomography to refine it. The migration result is then used to estimate a dip field that will later serve the estimation of the directional smoothing operators. Now, we show in Figures 5b and 5a the results of FWI with and without preconditioning, respectively. To increase resolution, we ran the last three iterations of the geologically constrained FWI without preconditioning. Clearly, the preconditioned result of Figure 5b is better behaved above salt than the unpreconditioned result in Figure 5a. Without preconditioning, artifacts due to sharp contrasts between salt and sediment velocities, and the discretization of the top of salt interface, dominate the image.

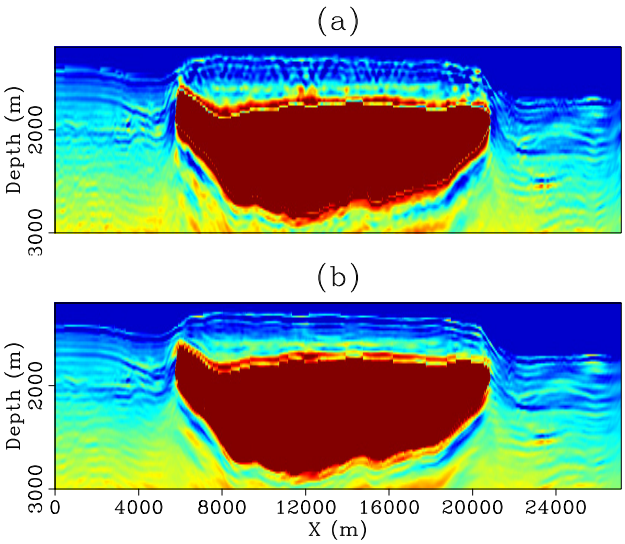

Figure 5: FWI results without (a) and with (b) preconditioning. Artifacts above salt are attenuated with preconditioning.

To better assess the quality of the inversion, Figure 6 displays a cube of the recorded data (6a) and of the final residual for the last scale (0-12 Hz) for the preconditioned inversion only (6b). We do not display the residual cube without preconditioning since it looks very similar to the one in Figure 6b. In Figure 6, the front panel is a common-shot section, the side panel a common-offset section, and the top panel a constant-time section. Comparing Figures 6a and 6b, we notice that the inversion is quite successful at matching the data outside the salt boundaries (before X=5,000 m and after X=21,000 m): the common-shot section shows that many events are succesfully inverted for. However, the residual remains strong where velocity contrasts are important (e.g., water-sediment and salt-water transitions.) This observation is consistent with previous work from Barnes and Charara (2009) and illustrate the limitations of the acoustic approximation with field data where sharp velocity contrasts are present.

Conclusions

Preconditioning waveform inversion with non-stationary directional Laplacian filters can yield geologically meaningful velocity models. It can help decrease artifacts due to acquisition geometry or inconsistencies in the data (not shown here). We anticipate that preconditioning can play a bigger role with real data where low frequencies are often lacking, where data are noisy and where the acquisition geometry is irregular.

Acknowledgements

The authors thank GeoImaging Solutions Inc., Repsol Sinopec Brasil SA and Geo Imaging Solucoes Tecnologicas em Geociencias Ltda for permission to publish these results. They also thank the sponsors of the Stanford Exploration Project for their support and WesternGeco for providing the field data example.

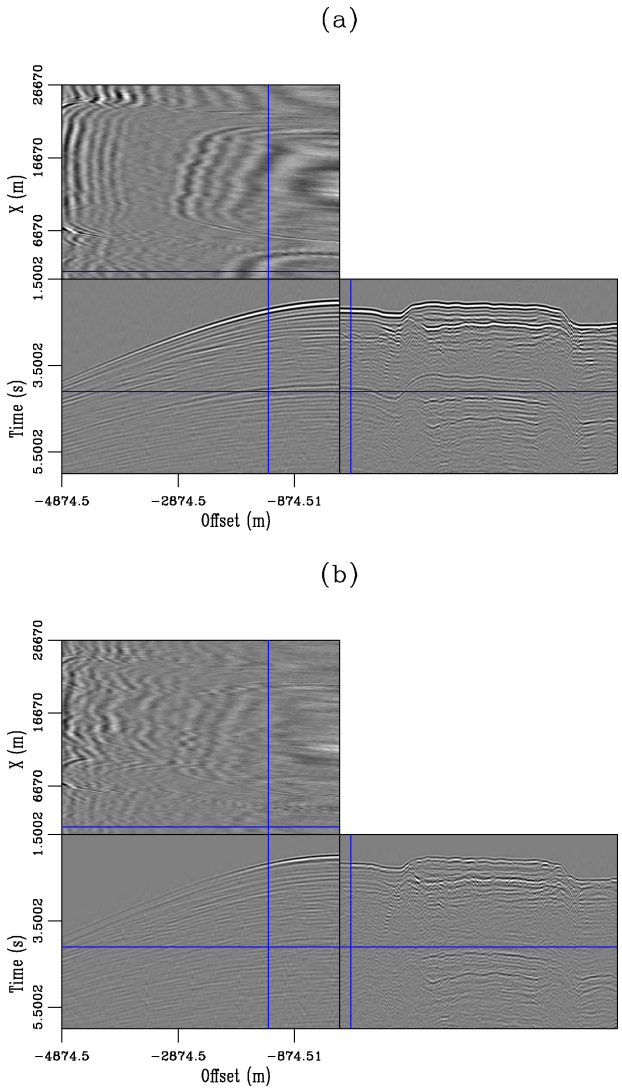

Figure 6: Slices through (a) the input data and (b) data residual volumes after FWI with preconditioning.

References

Barnes, C., and M. Charara, 2009, The domain of applicability of acoustic full-waveform inversion for marine seismic data: Geophysics, **74**, WCC91–WCC103.

Brenders, A., R. Pratt, and S. Charles, 2009, Waveform tomography of 2-d seismic data in the canadian foothills - data preconditioning by exponential time-damping: EAGE expanded abstract, U041.

Bunks, C., F. Saleck, S. Zaleski, and G. Chavent, 1995, Multiscale seismic waveform inversion: Geophysics, **60**, 1457–1473.

Claerbout, J., 1995, Earth Soundings Analysis: Processing Versus Inversion: Blackwell Scientific Publications.

Crase, E., A. Pica, M. Noble, J. McDonald, and A. Tarantola, 1990, Robust elastic nonlinear waveform inversion: Application to real data: Geophysics, **55**, 527–538.

Fomel, S., 2002, Applications of plane-wave destruction filters: Geophysics, **67**, 1946–1960.

Hale, D., 2007, Local dip filtering with directional laplacians: CWP report, 91–102.

Lailly, P., 1984, The seismic inverse problem as a sequence of before stack migrations: Conference on Inverse Scattering.

Nocedal, J., 1980, Updating quasi-newton matrices with limited storage: Mathematics of Computation, **95**, 339–353.

Ravaut, C., S. Operto, L. Improta, J. Virieux, A. Herrero, and P. dell'Aversana, 2004, Multi-scale imaging of complex structures from multi-fold wide-aperture seismic data by frequency-domain full-wavefield inversions: Application to a thrust belt: Geophysical Journal International, **159**, 1032–1056.

Sirgue, L., O. Barkved, J. V. Gestel, O. Askim, and J. Kommedal, 2009, 3d waveform inversion on valhall wide-azimuth obc: EAGE expanded abstract, U038.

Sirgue, L., and R. Pratt, 2004, Efficient waveform inversion and imaging: A strategy for selecting temporal frequencies: Geophysics, **69**, 231–248.

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