

Curvelet-based migration preconditioning

Felix J. Herrmann¹, Cody R. Brown¹, Yogi A. Erlangga¹, and Peyman P. Moghaddam¹

(September 26, 2008)

Running head:

ABSTRACT

The extreme large size of typical seismic imaging problems has been one of the major stumbling blocks for a successful application of iterative techniques from numerical linear algebra to attain accurate migration amplitudes. These iterative methods are important because they complement theoretically-driven approaches that are hampered by mundane differences to control problems such as finite-acquisition aperture, source-receiver frequency response, and directivity. To solve this problem, we apply the well-know technique of preconditioning that significantly increases the convergence of iterative solvers, making least-squares migration more tangible. First, we discuss different levels of preconditioning that range from corrections for the order of the migration operator to corrections for spherical spreading and position and reflector-dip dependent amplitude errors. While the first two corrections correspond to simple scalings in the Fourier and physical space, the third correction requires an intricate phase-space scaling, which we carry out with curvelets. Aside from providing the appropriate domain for the scaling, curvelets have the additional

¹Seismic Laboratory for Imaging and Modeling, Department of Earth and Ocean Sciences, University of British Columbia, 6339 Stores Road, Vancouver, V6T 1Z4, BC, Canada.

advantage of adding robustness with respect to noise and moderate kinematic errors to our iterative-inversion procedure. We conclude by showing that our combined preconditioner leads to a significant improvement for the convergence of least-squares ‘wave-equation’ migration on a line from the SEG AA’ Salt model.

INTRODUCTION

Over the years, extensive research has been done to reduce the computational costs of (least-squares) seismic imaging. Improvements in this area are particularly important during iterative least-squares migration, where the linear Born scattering operator is inverted with iterative Lanczos methods, such as LSQR (Paige and Saunders, 1982; De Roeck, 2002). Examples of these methods can be found in the literature (see e.g. Nemeth et al., 1999; Chavent and Plessix, 1999; Hu et al., 2001; Kuhl and Sacchi, 2003; Yu et al., 2006).

The most successful methods to reduce the cost of migration are the so-called scaling methods where the action of the compound linearized modeling-migration operator—known as the Hessian or normal operator—is replaced by a diagonal scaling in some domain, see e.g. contributions by Claerbout and Nichols (1994); Rickett (2003); Guitton (2004); Plessix and Mulder (2004), and more recently by Symes (2008) and Herrmann et al. (2008a). These methods vary in degree of sophistication with regard to the estimation of the diagonal through migrated-image to remigrated-image matching. They also differ in the way the scaling is applied—i.e., by 'division' in the physical domain or via sparsity promotion in the curvelet domain, as reported recently by Herrmann et al. (2008a). During all these methods, imaged amplitudes are restored by applying scaling as a post-processing step after migration.

In this paper, we take this line of research a step further by using the above scaling argument to apply the proper preconditioning to the system of equations involved in linearized Born scattering. To illustrate the improvements in migrated images and in the convergence of least-square migration, we consider three levels of preconditioning. First, we correct for the order of the normal operator by introducing a left preconditioning, consisting of a fractional time integration that corresponds to a scaling in the Fourier domain. This first

level of preconditioning follows directly from earlier work on migration-amplitude recovery reported by Herrmann et al. (2008a) and Symes (2008). The next level of preconditioning consists of a canonical diagonal scaling in the physical domain that compensates for spherical spreading of seismic waves. As a final step, we also include a curvelet-domain scaling as part of the right preconditioning. This final step corrects for the remaining amplitude errors that vary spatially as a function of the reflector dip. We conclude by studying the performance of these different levels of preconditioning on the SEG AA' Salt model (O'Brien and Gray, 1996; Aminzadeh et al., 1997), using a reverse-time 'wave-equation' migration code with optimal checkpointing (Symes, 2007).

PROBLEM FORMULATION

During seismic imaging, the following system of equations needs to be solved

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{1}$$

where \mathbf{b} is the known data vector, \mathbf{A} the linearized Born scattering operator, and \mathbf{x} the unknown model vector. Contrary to many inverse problems, the matrix \mathbf{A} , albeit extremely large, is reasonably well behaved and a migrated image can be obtained by applying the adjoint of \mathbf{A} to the data vector—i.e., $\tilde{\mathbf{x}} = \mathbf{A}^*\mathbf{b}$ with \mathbf{A}^* the migration operator. The symbol $*$ denotes the adjoint, and $\tilde{\mathbf{x}}$ is the estimate (denoted by the \sim) for the migrated image.

Unfortunately, the output of the above procedure, called migration, produces erroneous results for the amplitudes of the imaged reflectors. To restore these amplitudes, the matrix in Equation 1 can be inverted using the method of least-squares—i.e.,

$$\tilde{\mathbf{x}}_{LS} = (\mathbf{A}^*\mathbf{A})^{-1}\mathbf{A}^*\mathbf{b} = \mathbf{A}^\dagger\mathbf{b}, \tag{2}$$

with $\mathbf{A}^*\mathbf{A}$ the normal or Hessian operator and the symbol \dagger reserved for the pseudo (or generalized, Vogel, 2002) inverse of \mathbf{A} (the system \mathbf{A} is overcomplete and does not permit a straightforward matrix inverse). This expression explicitly solves the following optimization problem:

$$\tilde{\mathbf{x}}_{LS} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2, \quad (3)$$

which corresponds to finding an image vector, \mathbf{x} that after modeling fits the data vector, \mathbf{b} .

Since \mathbf{A} is large, the least-squares inverse needs to be calculated with iterative matrix-free methods such as LSQR (Paige and Saunders, 1982). Even though these iterative methods converge relatively quickly for an accurate background velocity model (see e.g. Nemeth et al., 1999; Chavent and Plessix, 1999; Hu et al., 2001; Kuhl and Sacchi, 2003; Yu et al., 2006), the sheer size of the imaging problem calls for a reduction in the number of iterations—i.e., the number of matrix-vector multiplies. Apparent issues with convergence in practical circumstances, where noise is present and the velocity model is not known exactly, are also a concern for these iterative methods.

In a perfect world, with infinite computational resources, the ideal preconditioning for the algorithm in Equation 1 corresponds to

$$\mathbf{A}\mathbf{M}_R^{-1}\mathbf{u} = \mathbf{b}, \quad \mathbf{x} := \mathbf{M}_R^{-1}\mathbf{u}, \quad (4)$$

with the right-preconditioning matrix, $\mathbf{M}_R := (\mathbf{A}^*\mathbf{A})^{1/2}$, given by the ‘square-root’ of the normal operator. Here, the symbol $:=$ refers to ‘defined as’. In this ideal case, migration recovers the image vector \mathbf{x} , exactly. Unfortunately, in practice (albeit some recent exciting progress has been made by Demanet and Ying, 2008, using discrete symbol calculus for smooth symbols, a development on which we intend to report in the future) the quantity $(\mathbf{A}^*\mathbf{A})^{1/2}$ cannot be computed and we have to resort to appropriate approximations.

In this paper, we propose a combination of left and right preconditioning—i.e., we replace Equation 1 by

$$\mathbf{M}_L^{-1} \mathbf{A} \mathbf{M}_R^{-1} \mathbf{u} = \mathbf{M}_L^{-1} \mathbf{b}, \quad \mathbf{x} := \mathbf{M}_R^{-1} \mathbf{u}, \quad (5)$$

with \mathbf{M}_L^{-1} the left-preconditioning matrix. The migrated and least-squares migrated images are given by $\tilde{\mathbf{x}} = \mathbf{M}_R^{-1} \tilde{\mathbf{u}}$, with $\tilde{\mathbf{u}} = (\mathbf{M}_L^{-1} \mathbf{A} \mathbf{M}_R^{-1})^* \mathbf{M}_L^{-1} \mathbf{b}$, and by $\tilde{\mathbf{x}}_{LS} = \mathbf{M}_R^{-1} \tilde{\mathbf{u}}_{LS}$, with $\tilde{\mathbf{u}}_{LS} = (\mathbf{M}_L^{-1} \mathbf{A} \mathbf{M}_R^{-1})^\dagger \mathbf{M}_L^{-1} \mathbf{b}$, respectively. Our preconditioners are derived from the following three observations: (i) the normal operator is in d dimensions a $(d - 1)$ -order pseudo-differential operator (Ψ DO, see e.g. recent work by Herrmann et al., 2008a; Symes, 2008, and the references therein), (ii) migration amplitudes decay with depth due to spherical spreading of seismic body waves, and (iii) zero-order Ψ DO’s can be approximated by a diagonal scaling in the curvelet domain (see e.g. Herrmann et al., 2008a). These observations allow us to define a series of increasingly more accurate approximations to the ‘square-root’ of the normal operator, leading to better and better preconditioners. Finally, we also argue that using curvelets will add a certain robustness to noise and errors in the velocity model, an observation substantiated by successful applications of this transform in signal separation problems that are hampered by noise and moderate kinematic errors (Wang et al., 2008).

PRECONDITIONING

In this section, we introduce different types of preconditioners based on the aforementioned observations. For each preconditioned system, we study the migrated images—i.e., the action of the adjoint of the preconditioned system for scattering. Later, we study the convergence of the iterative solver. The examples are computed for the reflectivity and smooth velocity background model plotted in Figure 1. To test our preconditioner, data with 324 shots is generated using Equation 1. Each shot consists of 176 traces of 6.4s and with a trace interval

of 24m. The maximum offset of the data is 4224m.

Left preconditioning by fractional differentiation

As stated before, the normal operator corresponds under certain conditions—such as the high-frequency limit and smooth background velocity models—to a $(d - 1)$ -order Ψ DO. In 2-D image space, this operator corresponds to the leading-order behavior of a Laplacian—i.e., the action of $(-\Delta)\cdot$ in the physical domain, or to $|\xi|^2\cdot$ with ξ the wave vector in the spatial Fourier domain. In data space, this action corresponds to a multiplication by $|\omega|$ in the temporal Fourier domain (Herrmann et al., 2008a). We compensate for this action by defining the following left preconditioning:

$$\mathbf{M}_L^{-1} := \partial_{|t|}^{-1/2}, \tag{6}$$

where $\partial_{|t|}^{-1/2}\cdot := \mathcal{F}^*|\omega|^{-1/2}\mathcal{F}\cdot$ with \mathcal{F} the Fourier transform and $\mathcal{F}^* = \mathcal{F}^{-1}$ its inverse. We define the level I preconditioner with \mathbf{M}_L^{-1} as above and $\mathbf{M}_R^{-1} = \mathbf{I}$. Comparison of the migrated images before and after left preconditioning (cf. Figures 1(c) and 2(a)) shows that the imprint of the Laplacian on the image is removed—i.e., some of the low-frequency content are restored. However, the migrated image still contains dimming of the amplitudes.

Right preconditioning by scaling in the physical domain

To further correct the amplitudes, we propose to apply a scaling to compensate for the leading-order amplitude error, namely the linear amplitude decay with depth. This decay is linear because the reflected waves travel from the source at the surface down to the reflector, experiencing an amplitude decay proportional to the square-root of the reflector depth (in 2-D), and back up, experiencing another decay—hence the linear dependence. We correct

for this linear amplitude decay by defining the following right-preconditioning matrix:

$$\mathbf{M}_R^{-1} = \mathbf{D}_z := \text{diag}(\mathbf{z})^{\frac{1}{2}}, \quad (7)$$

where $z_i = i\Delta z$, $i = 1 \cdots n_z$, with Δz the vertical sample interval and n_z the number of depth samples. Combined with the left preconditioner \mathbf{M}_L^{-1} , we call this the level II preconditioner. As expected, the results for the migrated image in Figure 2(b) now show a further improvement. However, there remain amplitude variations, e.g., along the major horizontal reflector just above 3500m.

Right preconditioning by scaling in the curvelet domain

After applying the left preconditioning, the Hessian can be modeled by a zero-order Ψ DO whose action corresponds to that of a non-stationary dip filter—i.e., we have

$$(\Psi f)(x) \simeq \int_{\xi \in \mathbb{R}^d} e^{j\xi \cdot x} a(x, \xi) \hat{f}(\xi) d\xi, \quad (8)$$

with Ψ the Hessian for the preconditioned modeling operator and $a(x, \xi)$ a space- and spatial-frequency dependent filter known as the symbol. We use the symbol \simeq to indicate a high-frequency approximation. As shown in Herrmann et al. (2008a), the action of the Ψ DO can, after discretization, be approximated by a scaling in the curvelet domain—i.e., we have the following approximate identity

$$\Psi \mathbf{r} \approx \mathbf{C}^* \mathbf{D}_\Psi^2 \mathbf{C} \mathbf{r}, \quad \mathbf{D}_\Psi^2 := \text{diag}(\mathbf{d}^2), \quad (9)$$

accurate for a reference vector \mathbf{r} close enough to the actual image. In this expression, the matrices \mathbf{C} and \mathbf{C}^* stand for the 2D discrete curvelet transform (see e.g. Candès et al., 2006) for which the adjoint equals the pseudoinverse—i.e., we have $\mathbf{C}^T \mathbf{C} = \mathbf{I}$ with \mathbf{I} the identity matrix. The reciprocal of the curvelet-domain scaling coefficients, \mathbf{d}^{-2} , is found by

a remigrated image-to-image matched-filtering procedure that involves the reference vector, typically derived from a conventional migrated image, and the remigrated reference vector. Hence, the cost of calculating the above approximation (cf. Equation 9) is roughly one modeling and one migration (for further details refer to Herrmann et al., 2008a,b). As we will show below, these additional costs are well offset by increased convergence.

By including the above approximation, we define the right preconditioning matrix

$$\mathbf{M}_R^{-1} = \mathbf{D}_z \mathbf{C}^* \mathbf{D}_\Psi^{-1}, \quad (10)$$

which approximately compensates for the remaining amplitude errors. Even though Equation 9 is only an approximation—with an accuracy that improves for increasing wavenumbers—we are in a good position because there is no need for high accuracy when defining preconditioners. We call the combination of Equations 6 and 10, the level III preconditioner. The image obtained with this system is plotted in Figure 2(c) and shows, as expected, further improvement in the amplitudes for the migrated image. Despite this improvement, image artifacts and amplitude errors are still present and can be attributed to the approximation in Equation 9 and to the fact that migration does not correspond to inversion—i.e., $\mathbf{A}^* \mathbf{A} \neq \mathbf{I}$.

CONVERGENCE OF LEAST-SQUARES MIGRATION

Even though the different preconditioning operators defined so far lead to improvements, problems remain in getting high-fidelity images. As reported in the literature, the image quality can be further improved by replacing migration with least-squares migration, as long as the background velocity model is known sufficiently accurately. In that case, the (preconditioned) scattering matrix is inverted, using an iterative method such as LSQR (Paige and Saunders, 1982). After a limited number of iterations, this method inverts

the (preconditioned) system of equations approximately—i.e., the optimization problem of Equation 3 is solved to some problem-dependent accuracy.

The performance of iterative solvers depends on certain properties of the (preconditioned) matrix \mathbf{A} , which include its condition number (ratio of the largest to smallest singular value) and the clustering of the singular values (see e.g. De Roeck, 2002, where these quantities are discussed for a relatively small Kirchoff-based imaging problem). Unfortunately, the size of typical imaging problems precludes actual computation of these properties for the different levels of preconditioning.

Because LSQR minimizes the residual during each iteration (cf. Equation 3), studying its progress as a function of the number of iterations gives us some way to gauge its performance. For this purpose, we introduce the normalized log-based least-squares residual as $\mu_k = 20 \log \|\mathbf{A}\mathbf{u}_k - \mathbf{b}\|_2 / \|\mathbf{b}\|_2$, with \mathbf{u}_k the solution of the (preconditioned) system after k iterations. Following De Roeck (2002), we also track the log-based least-squares model-space residual—i.e., $\nu_k = 20 \log \|\mathbf{A}^*(\mathbf{A}\mathbf{u}_k - \mathbf{b})\|_2 / \|\mathbf{A}^*\mathbf{b}\|_2$. Note that in practice, these latter residuals are never computed because they are not a by-product of LSQR, and require additional matrix-vector multiply per iteration. Contrary to data-space residuals, which possibly contain unmodeled components that may not be in the range of the modeling operator, the model-space residual typically converges to zero.

We use both these quantities to empirically establish the performance of the different levels of preconditioning that were introduced earlier. The results of this exercise are summarized in Figure 3, which plots the decay of μ_k and ν_k as a function of the number of iterations k . Remember, the cost of each iteration is dominated by a matrix-vector multiply by \mathbf{A} and by \mathbf{A}^* and the overhead of level III preconditioning corresponds to one iteration.

As we move from a single left preconditioner, towards left and right preconditioners, the data residual errors decay faster with a significant improvement obtained by the curvelet-domain scaling. For instance, the residual after 10 iterations for level II (fractional integration and depth weighting) preconditioning is attained by only 5 iterations of the level III preconditioned system (including curvelet-domain scaling), whereas the result after 10 iterations is approximately 2 dB better. Including the computational overhead of the level III preconditioning, there is a net gain of 4 iterations for an improvement of 8 dB in accuracy. Even though the picture for the model space residual is less clear—because of more significant iteration-to-iteration variations, and unexplained behavior for the systems without curvelet-domain preconditioning—a similar trend can be observed for the level III preconditioning. For instance, after 10 iterations, we have an improvement of approximately 4.5 dB.

These improvements in convergence for the preconditioned system also reflect into the least-squares migrated results included in Figure 4. Comparing these images shows a clear improvement for the preconditioned system, plotted in Figure 4(b), over the least-squares result obtained without preconditioning. Moreover, juxtaposing the preconditioned least-squares image with the solution for the preconditioned migrated image after one iteration and depicted in Figure 2(c) shows a significant enhancement of the overall amplitudes and frequency content (cf. Figure 1(a) and 4(b)).

DISCUSSION

Aside from the reported improvements in convergence from the curvelet-domain scaling, we argue that this choice of a transform has the additional benefit of adding robustness to the presence of (coherent) noise and moderate errors in the linearized Born modeling operator. Our arguments in support of this important claim for the practicability of our algorithm are:

- The well-documented approximate invariance of curvelets under the linearized Born scattering operator (Candès and Demanet, 2005; Douma and de Hoop, 2007; Charis and Nguyen, 2008; Anderson et al., 2008). As a consequence, the curvelet-preconditioned system (cf. Equation 10) has columns that are ‘curvelet like’—i.e., the data space is spanned by elements that share the beneficial multiscale and multidirectional properties of curvelets. For instance, small shifts over the support of a curvelet will not adversely affect the corresponding curvelet coefficient.
- The redundancy of curvelet frames (frames are redundant signal representations that generalize the notion of a basis) makes this transform less prone to errors in individual entries in the curvelet vector.
- The redundancy of curvelets spreads (coherent) noise over more coefficients and this property in tandem with their ‘locality’—i.e., only a small subset of the curvelet coefficients contribute to the reconstruction of a particular localized feature in the image space—leads to robustness because only a small fraction of the ‘noise’ will contribute to the reconstruction.

CONCLUSIONS

Because of the size of the seismic imaging problem, preconditioning of least-squares migration is an elusive topic, where traditional approaches from linear algebra have not yet found their way. Lack of direct access to the matrices involved and the cost of evaluating matrix-free implementations of the operators are both to blame. Nonetheless, the first few iterations of the LSQR algorithm towards the solution of least-squares migration are known to make significant progress towards the solution in cases where the background-velocity model is

known with sufficient accuracy. Unfortunately, even for this limited number of iterations, the computational costs are often still prohibitively large for practical problems. Moreover, there may be (coherent) noise, and velocity models may contain errors, which may in certain cases lead to divergence as the number of iterations increases.

The method presented in this paper partly resolves these issues through left and right preconditioning, combined with the use of redundant curvelets. We found that the inclusion of the curvelet-domain scaling as a right preconditioner is important because it restores the amplitudes and leads to faster convergence, at a relatively small computational expense. Aside from this tangible reduction in computational costs of roughly 50%, the use of curvelet frames will add robustness, which in practice translates into less sensitivity to (coherent) noise and errors in the velocity model. This property adds to the practicability of our method. Finally, our approach can also be generalized to exploit well-documented curvelet-domain sparsity on the image.

ACKNOWLEDGMENTS

The authors would like to thank C. C. Stolk and W. W. Symes for their input on migration-amplitude recovery and for the use of the reverse-time migration code. We also would like to thank Tamas Nemeth for insightful discussions and the authors of CurveLab (www.curvelet.org). The examples were prepared with Madagascar (rsf.sf.net), supplemented by SLIMpy operator overloading, developed by S. Ross-Ross. This work was in part financially supported by the NSERC Discovery (22R81254) and CRD Grants DNOISE (334810-05) of F.J.H. and was in part carried out within the SINBAD project with support, secured through ITF, from BG Group, BP, Chevron, ExxonMobil and Shell.

REFERENCES

- Aminzadeh, F., J. Brac, and T. Kunz, 1997, 3-D Salt and Overthrust Model. SEG/EAGE 3-D Modeling Series, No. 1: Society of Exploration Geophysicists, Tulsa.
- Anderson, F., M. V. de Hoop, H. Smith, and G. Uhlmann, 2008, A multi-scale approach to hyperbolic evolution equations with limited smoothness: *Communications in Partial Differential Equations*, **33**, 988–1017.
- Candès, E. J. and L. Demanet, 2005, The curvelet representation of wave propagators is optimally sparse: *Comm. Pure Appl. Math*, **58**, 1472–1528.
- Candès, E. J., L. Demanet, D. L. Donoho, and L. Ying, 2006, Fast discrete curvelet transforms: *SIAM Multiscale Modeling and Simulation*, **5**, 861–899.
- Chauris, H. and T. Nguyen, 2008, Seismic demigration/migration in the curvelet domain: *Geophysics*, **73**, S35–S46.
- Chavent, G. and R. E. Plessix, 1999, An optimal true-amplitude least-squares prestack depth-migration operator: *Geophysics*, **64**, 508–515.
- Claerbout, J. and D. Nichols, 1994, Spectral preconditioning: Technical Report SEP-82, Stanford Exploration Project.
- De Roeck, Y., 2002, Sparse linear algebra and geophysical migration: A review of direct and iterative methods: *Numerical Algorithms*, **29**, 283 – 322.
- Demanet, L. and L. Ying, 2008, Discrete symbol calculus. (Submitted for publication. Available at <http://math.stanford.edu/~laurent/papers/DSC.pdf>).
- Douma, H. and M. V. de Hoop, 2007, Leading-order seismic imaging using curvelets: *Geophysics*, **72**, S231–S248.
- Guitton, A., 2004, Amplitude and kinematic corrections of migrated images for nonunitary imaging operators: *Geophysics*, **69**, 1017–1024.

- Herrmann, F. J., P. P. Moghaddam, and C. C. Stolk, 2008a, Sparsity- and continuity-promoting seismic imaging with curvelet frames: *Journal of Applied and Computational Harmonic Analysis*, **24**, 150–173.
- Herrmann, F. J., D. Wang, and D. J. E. Verschuur, 2008b, Adaptive curvelet-domain primary-multiple separation: *Geophysics*, **73**, A17–A21.
- Hu, J., G. T. Schuster, and P. A. Valasek, 2001, Poststack migration deconvolution: *Geophysics*, **66**, 939–952.
- Kuhl, H. and M. D. Sacchi, 2003, Least-squares wave-equation migration for AVP/AVA inversion: *Geophysics*, **68**, 262–273.
- Nemeth, T., C. Wu, and G. T. Schuster, 1999, Least-squares migration of incomplete reflection data: **64**, 208–221.
- O’Brien, M. and S. Gray, 1996, Can we image beneath salt?: *The Leading Edge*, **15**, 17–22.
- Paige, C. C. and M. A. Saunders, 1982, LSQR: An algorithm for sparse linear equations and sparse least squares: *ACM TOMS*, **8**, 43–71.
- Plessix, R. and W. Mulder, 2004, Frequency-domain finite difference amplitude-preserving migration: *Geophysical Journal International*, **157**, 975–987.
- Rickett, J. E., 2003, Illumination-based normalization for wave-equation depth migration: *Geophysics*, **68**.
- Symes, W. W., 2007, Reverse time migration with optimal checkpointing: *Geophysics*, **72**, SM213–SM221.
- , 2008, Approximate linearized inversion by optimal scaling of prestack depth migration: *Geophysics*, **73**, R23–R35.
- Vogel, C., 2002, *Computational Methods for Inverse Problems*: SIAM.
- Wang, D., R. Saab, O. Yilmaz, and F. J. Herrmann, 2008, Bayesian wavefield separation by

transform-domain sparsity promotion: *Geophysics*, **73**, A33–A38.

Yu, J., J. Hu, G. T. Schuster, and R. Estill, 2006, Prestack migration deconvolution: *Geophysics*, **71**, S53–S62.

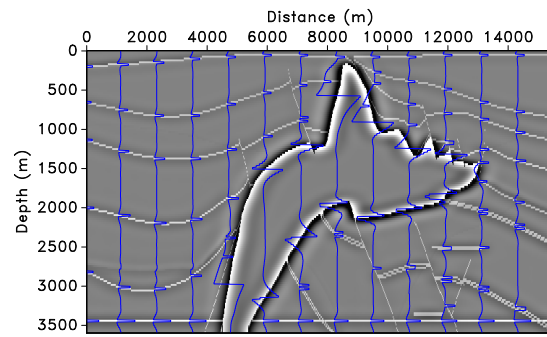
LIST OF FIGURES

1 The SEG/EAGE AA' salt model. **(a)** Reflectivity defined by the high-pass filtered velocity model. **(b)** Smoothed velocity model. **(c)** The migrated image according to Eq. 1. This image suffers from deteriorated amplitudes, especially under the high-velocity salt and for steep reflectors and faults.

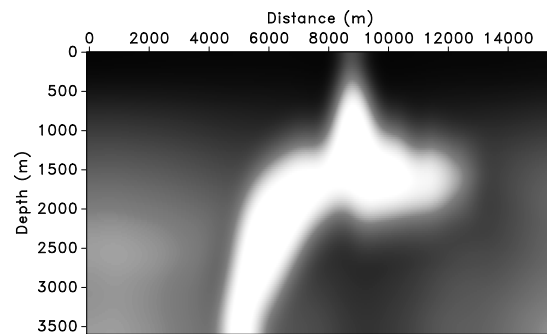
2 Migrated images for different levels of preconditioning. **(a)** Result for left preconditioning (level I, cf. Equation 6). **(b)** Result for left-right (including depth-correction) preconditioning (level II, cf. Equation 7). **(c)** The same but now including curvelet-domain scaling (level III, cf. Equation 10).

3 Residual decays for different levels of preconditioning. The dotted blue lines corresponds to least-squares migration without preconditioning, the dash-dotted lines to level I preconditioning, the dashed black lines to level II preconditioning, and the red solid lines to level III preconditioning. **(a)** Plot for the decay of the data-space normalized residues μ_k as a function of the number of LSQR iterations. **(b)** The same but now for the model-space normalized residuals ν_k .

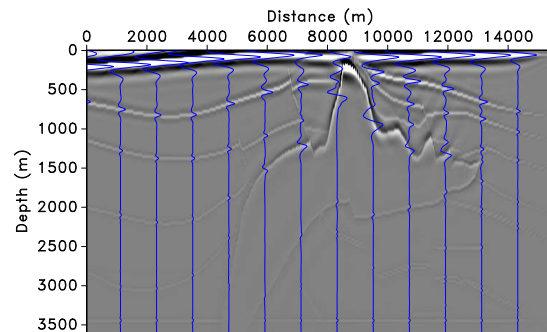
4 Least-squares migration without and with preconditioning. **(a)** Least-squares migrated image. **(b)** Least-squares image with level III preconditioning. Notice the improvement in the recovered reflectivity.



(a)




(b)

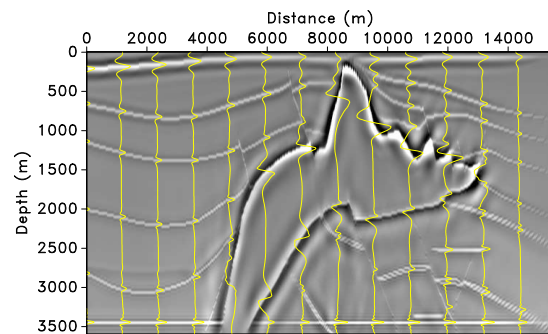


(c)


Figure 1: The SEG/EAGE AA' salt model. **(a)** Reflectivity defined by the high-pass filtered velocity model. **(b)** Smoothed velocity model. **(c)** The migrated image according to Eq. 1. This image suffers from deteriorated amplitudes, especially under the high-velocity salt and for steep reflectors and faults.



(a)




(b)




(c)

Figure 2: Migrated images for different levels of preconditioning. **(a)** Result for left preconditioning (level I, cf. Equation 6). **(b)** Result for left-right (including depth-correction) preconditioning (level II, cf. Equation 7). **(c)** The same but now including curvelet-domain scaling (level III, cf. Equation 10).

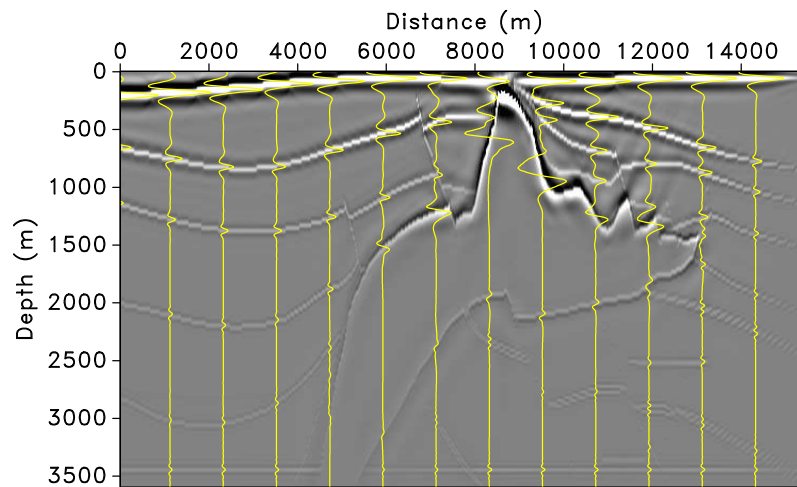


(a)

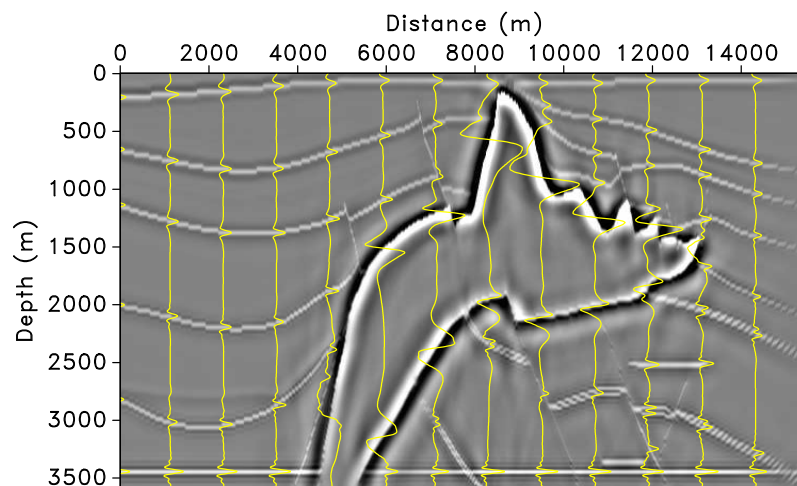


(b)

Figure 3: Residual decays for different levels of preconditioning. The dotted blue lines corresponds to least-squares migration without preconditioning, the dash-dotted lines to level I preconditioning, the dashed black lines to level II preconditioning, and the red solid lines to level III preconditioning. (a) Plot for the decay of the data-space normalized residues μ_k as a function of the number of LSQR iterations. (b) The same but now for the model-space normalized residuals ν_k .



(a)



(b)

Figure 4: Least-squares migration without and with preconditioning. **(a)** Least-squares migrated image. **(b)** Least-squares image with level III preconditioning. Notice the improvement in the recovered reflectivity.

Herrmann *et. al.* –