Enhanced reverse time migration with neural network-based data-domain preconditioning

Kristian Torres, Mauricio D Sacchi

Abstract

Least-squares reverse time migration (LSRTM) represents a significant advancement over reverse time migration (RTM) by iteratively computing an enhanced broadband reflectivity model. However, this improvement comes at the cost of increased computational complexity. In contrast, single-step methods provide appealing alternatives to iterative LSRTM by effectively balancing image quality and computational efficiency. This study proposes a novel non-iterative approach to enhance RTM images utilizing deep-learning-based preconditioning in the data domain. By training a convolutional neural network (CNN) on input-output pairs of observed and demigrated data, we can compute non-linear filters to precondition the migrated data. The method's effectiveness is demonstrated through tests conducted on the Marmousi data set. Results show that the CNN data preconditioning approach significantly improves amplitude balance and reduces artifacts compared to the original RTM section.

Introduction

RTM is the state-of-the-art technique in seismic imaging, particularly in complex geologies where other migration methods may prove ineffective (Baysal et al., 1983). RTM operates as an adjoint rather than an inverse operator of linearized wave-equation forward modelling (Claerbout, 1984). This distinction gives rise to specific drawbacks, including band-limited effects, irregular amplitudes and low-wavenumber noise arising from backscattering at high contrast boundaries. To address these challenges, LSRTM approximates a generalized inverse operator, thereby improving imaging quality and resolution (Dai et al., 2012).

LSRTM can be posed in the data or image domains (Fletcher et al., 2016). A data-domain implementation attempts to minimize the difference between observed and modelled data such that the inverted reflectivity resembles the recorded seismic data when demigrated. Image-domain LSRTM inverts for a reflectivity model to fit the raw (adjoint) migrated section. However, regardless of the implementation domain, the primary disadvantage of iterative LSRTM stems from the extensive computations required during each iteration. Therefore, implicitly inverting the Hessian can be impractical, especially for 3D data sets, as tens of iterations are typically required to gain substantial improvements.

Since LSRTM hinges on the estimation of the inverse Hessian operator, alternative techniques aim to replicate the Hessian inversion process through cost-effective approximations (e.g. Guitton, 2004; Yu et al., 2006). Single-step methods balance imaging quality and computational efficiency and employ image-domain or data-domain preconditioning to provide appealing solutions. In this context, Guitton (2004) introduces an image-domain method that estimates non-stationary matching filters from migrated and remigrated sections, enabling the approximation of the inverse Hessian in a single iteration. Herrmann et al. (2009) and Wang et al. (2017) extend this method in curvelet-based preconditioning schemes. Similarly, Khalil et al. (2016) propose a data-domain preconditioning technique, which only requires one demigration and two migration operations. Instead of working directly in the image space, this method computes non-stationary filters by matching demigrated and observed data, effectively approximating an inverse to the data resolution matrix. Liu and Peter (2018) and Liu et al. (2019) explore later variations of this idea by applying Wiener and Gabor deconvolution approaches, respectively.

Building upon the data-domain preconditioning method and the recent advent of deep learning algorithms, this study presents a single-step imaging framework that utilizes a convolutional neural network to compute non-linear non-stationary matching filters. Specifically, we employ the symmetrical U-net architecture (Ronneberger et al., 2015) as a data-domain preconditioner, training it with overlapping patches of observed and demigrated data. Following the training process, the filtering procedure can be effortlessly applied to the observations with minimal computational burden compared to seismic modelling and migration. Subsequently, the filtered data is remigrated to return to the image domain, resulting in an RTM outcome that exhibits reduced artifacts. Moreover, by applying preconditioning in the data domain, we remove the requirement for a representative training dataset containing model-space reflectivity labels.

Method

Least-squares reverse time migration

We start by defining the forward modelling equation as

$$\mathbf{d} = \mathbf{Lm}, \qquad (1)$$

where $\mathbf{L}$ represents the Jacobian (Born modelling) operator, and $\mathbf{d}$ and $\mathbf{m}$ are the vectors of linearized data and reflectivity model, respectively.

Assuming an invertible Hessian matrix, $\mathbf{H}$, the closed-form

solution of LSRTM is given by

$$\mathbf{m}_{ls} = [\mathbf{L}^T\mathbf{L}]^{-1}\mathbf{L}^T\mathbf{d}_{\text{obs}} \qquad (2)$$
$$= \mathbf{H}^{-1}\mathbf{m}_{\text{mig}},$$

where $\mathbf{m}_{ls}$ denotes the least-squares solution, and $\mathbf{m}_{\text{mig}}$ represents the RTM image as the result of applying the adjoint of the Jacobian operator, $\mathbf{L}^T$, to the vectorized observed data, $\mathbf{d}_{\text{obs}}$.

Equation 2 suggests that the pseudo-inverse yields an equivalent solution to the normal equations in cases of overdetermined inverse problems. Nevertheless, the seismic inversion system described in equation 1 exhibits both overdetermination, stemming from the abundance of seismic measurements, and underdetermination, arising from the limited receiver coverage of the subsurface and the band-limited nature of the seismic source. The minimum norm solution, $\mathbf{m}_{ls}$, gives one particular solution in the context of underdetermined inverse problems

$$\mathbf{m}_{ls} = \mathbf{L}^T[\mathbf{L}\mathbf{L}^T]^{-1}\mathbf{d}_{\text{obs}}, \qquad (3)$$

where we now assume that the matrix $[\mathbf{L}\mathbf{L}^T]$ is invertible. Thus, an alternative to least-squares imaging can be obtained by defining a data-domain preconditioning operator $\mathbf{P} \approx [\mathbf{L}\mathbf{L}^T]^{-1}$, such that

$$\mathbf{m}_{ls} \approx \mathbf{L}^T\mathbf{P}\mathbf{d}_{\text{obs}}. \qquad (4)$$

A practical way of computing $\mathbf{P}$ is by exploiting the correlation between migration and demigration,

$$\mathbf{d}_{\text{mig}} = \mathbf{L}\mathbf{m}_{\text{mig}} \qquad (5)$$
$$\mathbf{d}_{\text{mig}} = \mathbf{L}\mathbf{L}^T\mathbf{d}_{\text{obs}}$$
$$[\mathbf{L}\mathbf{L}^T]^{-1}\mathbf{d}_{\text{mig}} = \mathbf{d}_{\text{obs}}$$
$$\mathbf{P}\mathbf{d}_{\text{mig}} \approx \mathbf{d}_{\text{obs}},$$

where $\mathbf{d}_{\text{mig}}$ indicates demigrated data. Since $\mathbf{d}_{\text{mig}}$ and $\mathbf{d}_{\text{obs}}$ are known, the elements of $\mathbf{P}$ can be computed, for example, with a non-stationary matching filtering technique minimizing

$$E(\mathbf{P}) = ||\mathbf{P}\mathbf{d}_{\text{mig}} - \mathbf{d}_{\text{obs}}||_2^2 + \lambda R(\mathbf{P}), \qquad (6)$$

where $\mathbf{P}$ is defined as a multidimensional convolutional operator along all spatial axes, and R indicates a suitable regularization term.

CNN-based data-domain preconditioning

We propose to parameterize the preconditioning operator with a CNN, $\mathbf{P}_{\theta_d}$, with θ_d representing the trained CNN weights acting on the data domain. Through this parameterization, we aim to leverage the non-linear representation capabilities of CNNs to enhance deblurring and achieve superior amplitude balancing in RTM images in complex geological scenarios with poor illumination like subsalt regions and steeper reflectors. This choice is motivated by the growing evidence suggesting that CNNs can outperform traditional fixed, linear basis like Fourier and Curvelet transform, which are currently considered state-of-the-art. It also expands on the work of Torres and Sacchi (2023), who propose an analogous model-domain deep-learning-based preconditioning strategy for iterative LSRTM.

Thus, we implement a neural network architecture that focuses on extracting complex features from the high-dimensional data space to learn a non-linear surrogate of the inverse matrix $[\mathbf{L}\mathbf{L}^T]^{-1}$ without the need for intermediate processing steps or diagonal approximations (Liu and Peter, 2018).

Akin to the previous section, after an initial migration/demigration sequence, the CNN-based preconditioner is trained by minimizing the loss

$$E(\theta) = \frac{1}{N_s}\sum_{i=1}^{N_s}||\mathbf{P}_{\theta_d}(\mathbf{d}_{\text{mig}}^i) - \mathbf{d}_{\text{obs}}^i||_2^2 + \lambda||\theta||_2^2, \qquad (7)$$

over a training dataset of N_s paired samples of observed and demigrated data $\{(\mathbf{d}_{\text{mig}}^i, \mathbf{d}_{\text{obs}}^i)\}_{i=1}^{N_s}$. In this approach, we treat the observed data samples as the labels and the initial demigrated data as the inputs to the network. Finally, we get an enhanced RTM image by first applying the trained network on the observed reflections and then migrating this CNN-preconditioned data,

$$\mathbf{m}_{CNN} = \mathbf{L}^T\mathbf{P}_{\theta_d}(\mathbf{d}_{\text{obs}}) \qquad (8)$$
$$\approx \mathbf{L}^T[\mathbf{L}\mathbf{L}^T]^{-1}\mathbf{d}$$
$$\approx \mathbf{L}^T\mathbf{L}^{T^{-1}}\mathbf{L}^{-1}\mathbf{L}\mathbf{m}$$
$$\approx \mathbf{m}_{ls}.$$

Similar to the previously mentioned data-domain preconditioning methods derived from classic signal processing techniques, the overall cost of our approach amounts to approximately two RTM operations and one demigration. Compared to the cost of migration/demigration of the whole data, the CNN training and inference overhead is negligible.

Neural network architecture and training details

In this study, we chose the U-net architecture for its ability to preserve high-level details. Previously, the U-net has been generally applied as a segmentation precondition in imaging problems (Huang and Huang, 2021), post-stack reflectivity inversion (Torres and Sacchi), and the detection of geological structures such as faults (Wu et al., 2019).

As can be noticed from Figure 1, its architecture consists of an encoding path and a decoding path, where both input and output are images. This setup captures local and global features, enabling the model to learn complex mappings between input and output. The U-Net network incorporates several operations. Each step in the contracting path involves two 3×3 convolutional layers followed by the rectified linear unit (ReLU) activation function. Downscaling is achieved through 2×2 max pooling with a stride of 2. Conversely, the expansive path (right side) utilizes 2×2 upsampling with the same stride, accompanied by two convolutional layers that reduce the number of feature channels. The U-Net also includes skip connections that link the left and right paths to preserve spatial information lost during max pooling. These connections help preserve high-resolution information and facilitate the recovery of fine details in the output.

To train the U-net, we utilize overlapping patches of demigrated and observed data, randomly shuffled before being fed into the network. 10 realizations from the paired

Figure 1: The U-net architecture. The number at the bottom of each convolutional layer indicates the number of channels. **I** refers to the original image size.

training samples are shown in Figure 2a and 2b. When using RTM images for demigration, the energy magnitude of demigrated shot gathers exhibits significant variation compared to the observed data. This magnitude can range from -10^n to 10^n depending on the number of stacking shots and the source energy. We employ a standardization technique to address this data variance and accelerate the convergence of the training loss function. This technique ensures that the training data is transformed to have a mean of zero and a variance of one. Furthermore, we introduce a random horizontal flipping operation to enhance the training process and increase the diversity of training samples.

As the trained U-net will be applied to the observed data in the final step of our workflow (equation 8), we divide the original dataset into 90% training samples and the remaining 10% as validation samples. This division allows for the training of the U-net using most of the data while retaining a separate portion for evaluating the model's performance. Equation 7 was minimized using Adam with 50 epochs, weight decay damping term $\lambda = 10^{-4}$ and a learning rate of 10^{-3}. Figure 3a and 3b show the evolution of the loss and accuracy for both training and validation data concerning the number of epochs, respectively.

Results

We test our approach on the 2D acoustic Marmousi dataset Versteeg (1994). The recorded data are derived from the true velocity model depicted in Figure 4a. The migration velocity presented in Figure 4b is obtained by smoothing the true velocity using a 2D Gaussian smoother. Figure 6c shows the true reflectivity model. In this benchmark scenario, 103 shots are evenly initiated at intervals of 50 m, employing a 30 Hz Ricker wavelet as the source function. The observed data are obtained using a fixed-spreading acquisition setup with 512 receivers at 10 m intervals for each shot-gather. The recording duration covers 3 seconds, with direct waves being eliminated.

We initially present the RTM result for a single shot, migrated at x=1km, as depicted in Figure 5a. Figure 5b illustrates the outcome of migrating the observations after being filtered by the convolutional neural network (CNN). Upon comparing these two images, we observe that the preconditioned result exhibits improved illumination

at deeper reflectors and a noticeable reduction in low-frequency backscattering noise. Subsequently, we showcase the complete RTM section by stacking all individually migrated shot gathers, both without data preconditioning (Figure 6a) and with data preconditioning (Figure 6b) using the CNN. Notably, the amplitude balance is substantially enhanced, and there is a distinct reduction in migration artifacts around the top-right steeper reflectors (highlighted by the top red arrow), as well as a greater focusing power on the deeper reflectors below the bottom high-velocity intrusions (indicated by the bottom red arrow). Additionally, our method effectively mitigates the sparse migration footprints present in the shallower part of the model.

Discussion and Conclusion

Identifying optimal preconditioner operators can be a non-trivial task. It requires expertise and extensive experimentation. This work proposes an alternative data-domain preconditioning based on deep-learning filtering. We showed that the proposed approach can potentially deliver enhanced migrated sections with improved focusing capability and amplitude balance around under-illuminated regions. The training stage does not need further pre-processing and requires minimal user interaction. It is also relatively cheap: regarding wave equation operations, the method requires only a single extra migration compared to one LSRTM iteration.

As reported in existing literature (Nichols, 1997), the adjoint operator, which is the final operator employed in our equation 8, cannot generate model components within the null space. Therefore, it should be acknowledged that the proposed data-domain technique might not be particularly successful in aiding the retrieval of a reflectivity model with higher wavenumber content due to the assumptions it is based on. In terms of iterative inversion, future work will focus on combining the proposed data-domain (left) preconditioner with image-space (right) preconditioning to accelerate the convergence of iterative LSRTM and high-resolution reflectivity models.

References

Baysal, E., Kosloff, D. D., and Sherwood, J. W., 1983, Reverse time migration: Geophysics, **48**, no. 11, 1514–1524.
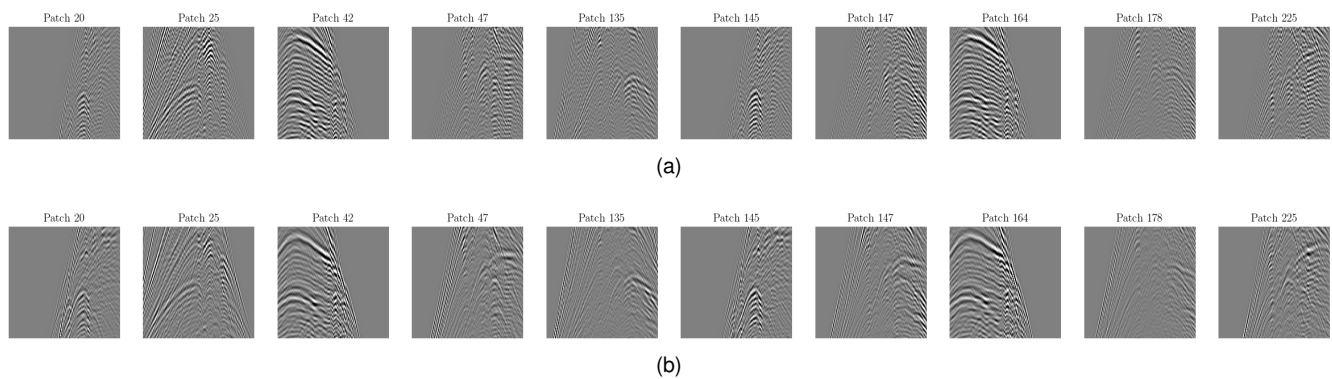
(a)

(b)

Figure 2: (a). 10 random realizations from the patched demigrated shot gathers used as input training data. (b) 10 random realizations from the patched observed shot gathers used as output training labels.

Claerbout, J. F., 1984, Imaging the Earth's Interior:, volume 86.

Dai, W., Fowler, P., and Schuster, G. T., 2012, Multi-source least-squares reverse time migration: Geophysical Prospecting, **60**, no. 4, 681–695.

Fletcher, R. P., Nichols, D., Bloor, R., and Coates, R. T., 2016, Least-squares migration — data domain versus image domain using point spread functions: The Leading Edge, **35**, no. 2, 157–162.

Guitton, A., 2004, Amplitude and kinematic corrections of migrated images for nonunitary imaging operators: Geophysics, **69**, no. 4, 1017–1024.

Herrmann, F. J., Brown, C. R., Erlangga, Y. A., and Moghaddam, P. P., 2009, Curvelet-based migration preconditioning and scaling: Geophysics, **74**, no. 4, A41–A46.

Huang, Y., and Huang, J., 2021, Optimized plane-wave least-squares reverse-time migration via a convolutional network: First International Meeting for Applied Geoscience & Energy, 1350–1354.

Khalil, A., Hoeber, H., Roberts, G., and Perrone, F., 2016, An alternative to least-squares imaging using data-domain matching filters: 2016 SEG International Exposition and Annual Meeting.

Liu, Q., and Peter, D., 2018, One-step data-domain least-squares reverse time migrationlsrtm: Geophysics, **83**, no. 4, R361–R368.

Liu, Q., Lu, Y., Sun, H., and Zhang, H., 2019, Single-step data-domain least-squares reverse-time migration using gabor deconvolution for subsalt imaging: IEEE Geoscience and Remote Sensing Letters, **17**, no. 1, 13–16.

Nichols, D., 1997, A simple example of a null space and how to modify it: Stanford Exploration Project Report, **82**, 185–192.

Ronneberger, O., Fischer, P., and Brox, T., 2015, U-net: Convolutional networks for biomedical image segmentation: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, 234–241.

Torres, K., and Sacchi, M. D., Deep decomposition learning for reflectivity inversion: Geophysical Prospecting.

Torres, K., and Sacchi, M., 2023, A deep-learning inverse hessian preconditioning for iterative least-squares migration: 84th EAGE Annual Conference & Exhibition, 1–5.

Versteeg, R., 1994, The marmousi experience: Velocity model determination on a synthetic complex data set: The Leading Edge, **13**, no. 9, 927–936.

Wang, M., Huang, S., and Wang, P., 2017, Improved iterative least-squares migration using curvelet-domain hessian filters: 2017 SEG International Exposition and Annual Meeting.

Wu, X., Liang, L., Shi, Y., and Fomel, S., 2019, Faultseg3d: Using synthetic data sets to train an end-to-end convolutional neural network for 3d seismic fault segmentation: Geophysics, **84**, no. 3, IM35–IM45.

Yu, J., Hu, J., Schuster, G. T., and Estill, R., 2006, Prestack migration deconvolution: Geophysics, **71**, no. 2, S53–S62.

Acknowledgments

(a)

(a)

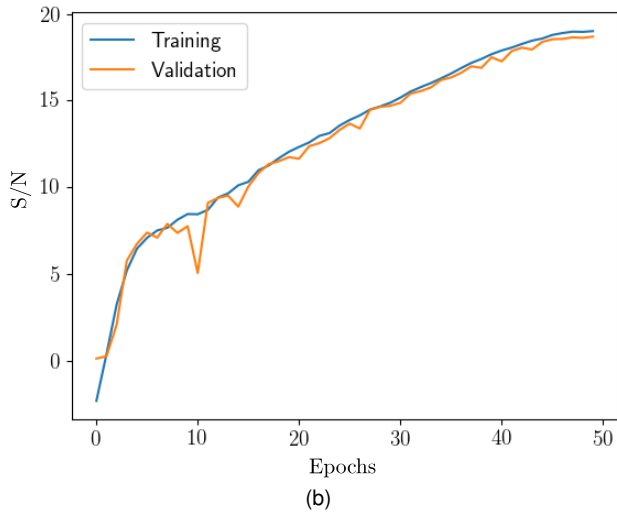Figure 4: (a). Marmousi velocity model. (b) Migration velocity model.

(b)

Figure 3: (a) The normalized training and validation loss functions versus the number of epochs. A similar pattern between both curves indicates that our method does not suffer from the overfitting problem. (b) Evolution of the signal-to-noise ratio (dB) $= 10 \times \log_{10} ||\mathbf{d}_{\text{mig}}||_2^2 / ||\mathbf{d}_{\text{mig}} - \mathbf{d}_{\text{obs}}||_2^2$, versus the number of iterations to measure reconstruction accuracy.


(a)

(b)

Figure 5: (a). Single-shot RTM. (b) Single-shot RTM with CNN-based preconditioning.
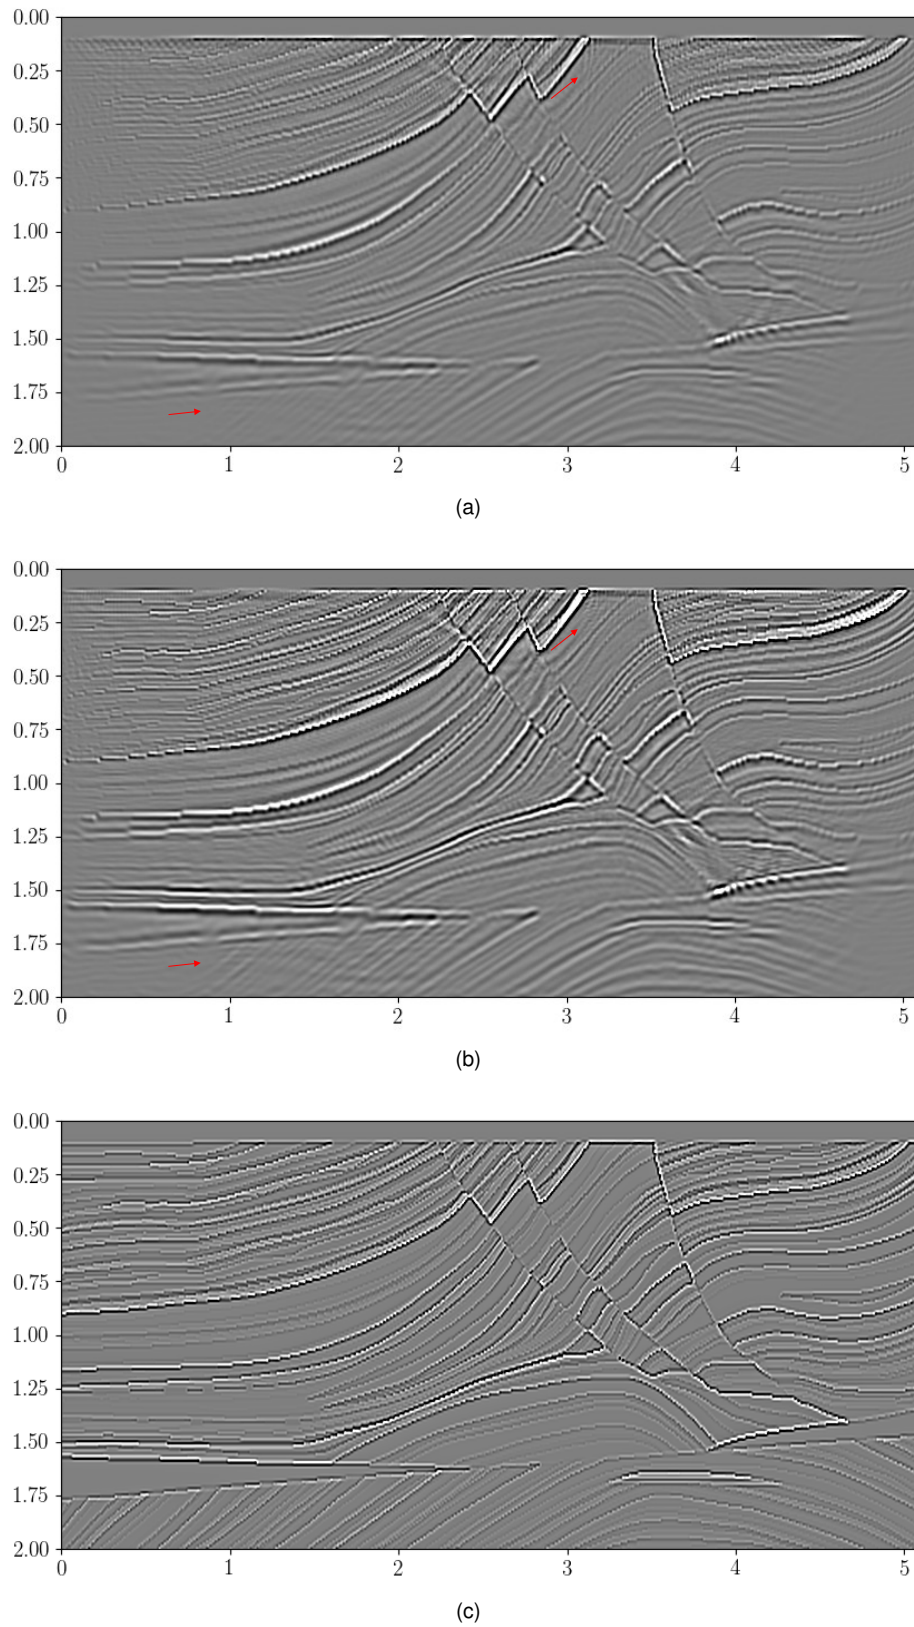
(a)

(b)

(c)

Figure 6: (a). Stacked RTM section. (b) Stacked RTM section with CNN-based preconditioning. (c) True reflectivity.