GEOPHYSICS, VOL. 89, NO. 2 (MARCH-APRIL 2024); P. R121–R136, 18 FIGS., 3 TABLES.
10.1190/GEO2022-0780.1

Image-domain seismic inversion by deblurring with invertible recurrent inference machines

Haorui Peng[1], Ivan Vasconcelos[1], and Matteo Ravasi[2]

ABSTRACT

In complex geologic settings and in the presence of sparse acquisition systems, seismic migration images manifest as nonstationary blurred versions of the unknown subsurface model. Thus, image-domain deblurring is an important step to produce interpretable and high-resolution models of the subsurface. Most deblurring methods focus on inverting seismic images for their underlying reflectivity by iterative least-squares inversion of a local Hessian approximation; this is obtained by either direct modeling of the so-called point-spread functions (PSFs) or by a migration-demigration process. In this work, we adopt a novel deep-learning (DL) framework, based on invertible recurrent inference machines (i-RIMs), which allows approaching any inverse problem as a supervised learning task informed by the known modeling operator (convolution with PSFs in our case): our algorithm can directly invert migrated images for impedance perturbation models, assisted with the prior information of a smooth velocity model and the modeling operator. Because i-RIMs are constrained by the forward operator, they implicitly learn to shape/regularize output models in a training-data-driven fashion. As such, the resulting deblurred images indicate great robustness to noise in the data and spectral deficiencies (e.g., due to limited acquisition). The key role played by the i-RIM network design and the inclusion of the forward operator in the training process is supported by several synthetic examples. Finally, using field data, we find that i-RIM-based deblurring has great potential in yielding robust, high-quality relative impedance estimates from migrated seismic images. Our approach could be of importance toward future DL-based quantitative reservoir characterization and monitoring.

INTRODUCTION

The ultimate goal of seismic imaging is to retrieve a subsurface model that is as close as possible to the true geologic structures in the subsurface, with amplitudes that could be reliably interpreted in terms of physical parameters, such as impedance. In practice, a seismic migration operator is either used to perform an imaging step in the processing workflow or as an integral part of a more complex inversion approach, such as least-squares migration or waveform inversion. From the point of view of inverse problems, standard migration operators such as those encountered in Kirchhoff prestack depth migration (PSDM) and reverse time migration (RTM) are nonunitary and ill-posed (Claerbout, 1992) because of the band-limited nature of seismic data (Chavent and Plessix, 1999), limited acquisition geometry (Duquet and Marfurt, 1999; Nemeth et al., 1999), and complex geology (Rickett, 2003; Guitton, 2004; Clapp, 2005). As such, migrated images tend to be blurred with unbalanced amplitudes (Gray, 1997). This means that seismic migration by inversion (also known as least-squares migration) is either implicitly or explicitly tied to the problem of deblurring migrated seismic images.

The process of deblurring of seismic images can be formulated as an inverse problem in either the data or image domains (Tarantola, 1987); early methods have focused on the data-domain formulation and were initially developed to invert images based on least-squares methods (Lailly, 1983; Tarantola, 1984) or with asymptotic functionals in a global fashion (Beylkin, 1985; Bleistein, 1987). This type of image-deblurring approach usually requires inverting the

Manuscript received by the Editor 28 December 2022; revised manuscript received 9 October 2023; published ahead of production 8 December 2023; published online 12 February 2024.
[1]Utrecht University, Earth Sciences Department, Utrecht, The Netherlands. E-mail: h.peng@uu.nl (corresponding author); i.vasconcelos@uu.nl.
[2]King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia. E-mail: matteoravasi@gmail.com.

Hessian matrix, which often is, in practice, intractable. Therefore, various attempts have been made to approximate the Hessian or its inverse to alleviate such a computational burden (Lambaré et al., 1992; Chavent and Plessix, 1999; Nemeth et al., 1999; Hu et al., 2001; Rickett, 2003; Guitton, 2004).

Image-domain deblurring can be performed in a target-oriented fashion by either explicitly computing the local Hessian for areas of interest in the migrated images (Valenciano et al., 2006, 2009; Tang, 2009) or by applying fit-for-purpose deblurring filters based on the assumption of local layered media (Yu et al., 2006; Aoki and Schuster, 2009). To further reduce the computation burden, the so-called point-spread functions (PSFs) also can be used as a local Hessian approximation in terms of the response of the imaging system to a point scatterer. They, in fact, represent the local resolution of the migrated image (Lecomte and Gelius, 1998) and capture the dip-dependent illumination effects due to the acquisition geometry and complex geologies (Letki et al., 2015). In this context, seismic images can be deblurred by direct nonstationary deconvolution with their corresponding PSFs (Sjoeberg et al., 2003; Fletcher et al., 2012).

Various approaches exist to compute the PSFs to be used in the deconvolution process. One can conveniently construct them via a demigration-migration process of well-separated point scatterers placed in the smooth velocity background (Fletcher et al., 2012). This approach has the main drawback that crosstalks can arise between neighboring PSFs. Such interference can be avoided by constructing PSFs through analytical approaches (Lecomte, 2008; Toxopeus et al., 2008; Lund et al., 2022), at a reduced computation burden. Deblurring migrated images with PSFs in the image domain has shown increasing potential with applications such as reflectivity estimation from Kirchhoff PSDM inversion (Valenciano et al., 2015; Cavalca et al., 2016; Lund et al., 2022) and RTM (Fletcher et al., 2012; Letki et al., 2015; Fletcher et al., 2016) or impedance inversion from PSDM and RTM (Du et al., 2016, 2021; Wang and Lu, 2016). With such an approach, inversion can be performed in parallel over different local image patches at reduced costs, which also suites 4D applications (Cavalca et al., 2020).

Recent years have seen the rise of deep learning (DL) and its applications in various aspects of geophysics (Yu and Ma, 2021), owing to its ability to assimilate a priori data, while potentially reducing the computation cost compared with traditional algorithms (once training is performed). For the seismic image deblurring problem, one straightforward way to exploit DL is to map low-resolution (prestack or poststack) migrated images to their high-resolution counterparts, using powerful techniques such as deep convolutional neural networks (CNNs) (Wang and Nealon, 2019; Liu et al., 2020a; Lu et al., 2020; Li et al., 2022a; Zhang et al., 2022) or generative adversarial networks (Halpert, 2018; Dutta et al., 2019; Kaur et al., 2020; Zhang and Wang, 2022). Although these approaches have yielded promising results, especially when assisted with the additional prior information of a smooth background velocity model (Kaur et al., 2020), they are purely computer-vision-based direct-mapping methods that are completely separated from the migration process itself.

Alternatively, DL could be leveraged directly within the migration process. For example, Liu et al. (2020b) use a support vector machine as a preconditioner to separate noise from signal in the iterations of least-squares RTM (LSRTM). Motivated by the connection between sparse least-squares migration (SLSM) and CNNs, Liu et al. (2020c) develop the neural network version of SLSM, which aims to find an optimal quasireflectivity image. Similarly, Torres and Sacchi (2022) propose to substitute the projection operators of LSRTM with sets of CNNs and learn an update function at each iteration, while implicitly accounting for the regularization effects.

Attempts also have been made to deblur seismic images with PSFs by DL. Liu et al. (2022) propose using CNNs to map PSFs directly to their inverse in the wavenumber $(K_x - K_z)$ domain, which are then used for the deblurring directly. Although such an inversion was usually considered unstable, their method has yielded reliable results when assisted by a weighting technique, as shown by their benchmark against explicit least-squares inversion by traditional optimization. Different but akin to Liu et al. (2022), we intend to deblur local migrated images with the help of a physics-based DL approach. Our key idea is to directly invert seismic images for their impedance perturbation (i.e., the difference between the true impedance model and its smoothed version) working in overlapping image patches and deconvolving the effect of the PSFs via the so-called invertible recurrent inference machine (i-RIM) network.

Fundamentally, the RIM is a special type of recurrent network architecture designed to solve inverse problems iteratively by explicitly exploiting a known forward operator (Putzky and Welling, 2017). As a more memory-efficient evolution of RIM, i-RIM is modified and combined with specially designed invertible layers to allow (1) on-the-fly computation of the memory-hungry activation functions, thereby relaxing the need for storage of intermediate activations in the forward pass, and (2) more stable training (Putzky and Welling, 2019). Although most machine-learning applications in geophysical inverse problems focus on choosing a suitable learning architecture and a preprocessing pipeline that could replace the knowledge usually embedded in the forward operator, RIM naturally blends the knowledge of the modeling process with the power of neural networks (Vasconcelos et al., 2022). For RIM, the crucial role of the prior information embedded in the forward operator has been shown to benefit the solution of a variety of inverse problems in diverse fields, such as astrophysics (Morningstar et al., 2018, 2019), medical imaging (Lønning et al., 2019; Putzky et al., 2019; Sabidussi et al., 2021), and geophysics (Kuijpers et al., 2020).

To outline the generalization ability of i-RIMs, we deliberately select our training data set to be composed of a single synthetic model (the right part of the SEAM Phase I model [Fehler and Keliher, 2011] referred to as "SEAM right"), that is mostly composed of horizontal fine sedimentary layering and mild lateral variations. During the training process, the i-RIM is initialized with a smooth velocity model and tasked to map the migrated image into its corresponding impedance perturbation, with the guidance of PSFs as the forward operator. Following Putzky et al. (2019), we use i-UNet as the benchmark network to compare to our i-RIM. The i-UNet is a modified UNet architecture (Ronneberger et al., 2015) combined with invertible layers; in this case, i-UNet takes the migrated image and impedance perturbation as input and output, respectively. For testing, we select (1) one synthetic model/data set, which is the left part of the SEAM Phase I model with realistic basin geologic features (referred to as "SEAM left") and (2) the Volve 2D field data set. We refer to Ravasi et al. (2015, 2016) for more details on the preprocessing of the Volve data set.

The rest of this paper is organized as follows. In the "Method" section, we first present the imaging setup, i.e., derive the impedance kernel for Kirchhoff PSDM, which describes the process of

generating migrated images and PSFs. Next, we provide a detailed introduction to the i-RIM network architecture and associated training process. In the "Numerical examples" section, we benchmark i-RIM with the i-UNet network for four cases of increasing complexity: three cases for the SEAM left model (without/with various noise types and with blended shots) and then the Volve field data set.

METHOD

Imaging parameterization: Deblurring for impedance

In this study, we consider a linear least-squares form of PSDM as our imaging algorithm to create seismic images parameterized in terms of impedance kernel, whose mathematical derivation we revisit in this section. Due to our reliance on supervised learning with i-RIM and i-UNet, the choice of parameterization used in imaging and a consistent implementation of the forward operator used to create the training data are key for the successful application of our methodology to field data. Thus, we begin by revisiting the Born approximation parameterized in terms of impedance and velocity. Starting from the linear isotropic acoustic wave equation in terms of velocity c and impedance $Z = \rho c$ (where ρ is the density) (e.g., Clayton and Stolt, 1981), we have

$$\mathcal{L}(p) = \left(\frac{\omega^2}{Zc} + \nabla \cdot \frac{c}{Z}\nabla\right)p = 0, \tag{1}$$

where ω is the angular frequency, $\mathcal{L}$ is the wave equation operator, and p is the pressure field. Next, we define scattering potential V as

$$V = \omega^2\left(\frac{1}{Zc} - \frac{1}{Z_0c_0}\right) + \nabla \cdot \left(\frac{c}{Z} - \frac{c_0}{Z_0}\right)\nabla, \tag{2}$$

where Z_0 and c_0 are the background impedance and velocity models, respectively. When considering constant density ρ, we have $(c/Z) - (c_0/Z_0) = 0$, and therefore equation 2 becomes

$$V = \omega^2\left(\frac{1}{Zc} - \frac{1}{Z_0c_0}\right) = \omega^2\delta\left(\frac{1}{Zc}\right) \tag{3}$$

with

$$\delta\left(\frac{1}{Zc}\right) = \frac{1}{Zc} - \frac{1}{Z_0c_0}, \tag{4}$$

where $\delta(1/Zc)$ denotes the medium perturbation in terms of the inverse of impedance and velocity product. The corresponding scattering-related coefficient α is defined as

$$\alpha = \frac{\delta\left(\frac{1}{Zc}\right)}{\frac{1}{Z_0c_0}} = \frac{c_0^2 - c^2}{c^2}, \tag{5}$$

such that

$$V = \omega^2\frac{\alpha}{Z_0c_0}. \tag{6}$$

Equations 5 and 6 provide the formalism that informs how a point scatter can be introduced to generate a PSF in terms of impedance-consistent model weighting. It can be shown (e.g., Clayton and Stolt, 1981) that the scattered-field data acquired at the acquisition surface can be expressed as the well-known Born-modeling integral:

$$d_s(\mathbf{x}_r, \mathbf{x}_s, \omega) = \int G_0^r(\mathbf{x}_r, \mathbf{x}, \omega)V(\mathbf{x}, \omega)G_0^s(\mathbf{x}, \mathbf{x}_s, \omega)d\mathbf{x}^2, \tag{7}$$

where $\mathbf{x}, \mathbf{x}_r, \mathbf{x}_s$ indicate the coordinates of the image points, receivers, and sources, respectively, and $G_0^r(\mathbf{x}_r, \mathbf{x}, \omega)$ and $G_0^s(\mathbf{x}, \mathbf{x}_s, \omega)$ are the receiver and source wavefields computed in the background, respectively. From this point forward, we will drop the functional dependencies for simplicity in our derivation and only keep them whenever necessary. Using equation 3, equation 7 becomes

$$d_s = \omega^2\int G_0^r\delta\left(\frac{1}{Zc}\right)G_0^s dx^2. \tag{8}$$

By assuming constant ρ, the impedance kernel for RTM is (Douma et al., 2010)

$$\mathcal{I} = \frac{1}{2}\left[\frac{c}{Z}\nabla^2(pp^\dagger)\right], \tag{9}$$

where $\dagger$ indicates the adjoint operation. Here, we have assumed that the imaging areas are sufficiently far away from the sources and receivers, and therefore have dropped the adjoint-source term. Although equation 9 is initially derived for finite-frequency RTM, it also is valid for the high-frequency-approximation-based PSDM, for which we replace the wave equation-based $p(x, x_s)$ and $p^\dagger(x, x_s)$ with asymptotically derived $G_0^s(\mathbf{x}, \mathbf{x}_s)$ and $\int G_0^r(\mathbf{x}_r, \mathbf{x})d_s^\dagger(\mathbf{x}_r, \mathbf{x}_s)d\mathbf{x}_r^2$, respectively. Together with equation 8, this yields the impedance kernel for PSDM:

$$\mathcal{I} = \omega^2\frac{c}{2Z}\nabla^2\mathbf{I} \tag{10}$$

with

$$\mathbf{I} = G_0^s\int G_0^r\left[\int G_0^{r\dagger}\delta\left(\frac{1}{Zc}\right)G_0^{s\dagger}d\mathbf{x}^2\right]d\mathbf{x}_r^2. \tag{11}$$

It can be seen from equation 10 that the conventional migrated image $\mathbf{I}$ is scaled by the local impedance and velocity, after Laplacian filtering. Moreover, when combined with equation 11, one can observe that the image $\mathcal{I}$ obtained from observations at the surface represents the true medium perturbation $\mathbf{m} = \delta(1/Zc)$ blurred by the Hessian $\mathbf{H}$ as follows:

$$\mathcal{I} = \mathbf{Hm}. \tag{12}$$

Because the Hessian represents multidimensional convolutions with the global, nonstationary imaging-operator PSFs, the target medium perturbation can be inferred by undoing the effects of such PSFs on the migrated image $\mathcal{I}$. Here, we rely on a local PSF approximation to model the effect of the Hessian of the imaging operator. Although the preceding derivation is based on $\mathbf{m} = \delta(1/Zc)$, in this study we choose to define the target model as the impedance perturbation $\mathbf{m} = \delta(Z)$ instead to enable its practical usage for reservoir impedance inversion. Note that this practical modification in the definition of model perturbation only changes the relative scaling of the image,

whereas the overall geologic structure remains the same. This imaging setup choice provides a consistent framework for generating the migrated image and PSFs parameterized in terms of acoustic impedance. The relation in equation 12 is shown in Figure 1, where Figure 1a and 1b shows the local PSF and impedance perturbation, respectively. Figure 1c and 1d shows $\mathcal{I}_H$ and $\mathcal{I}$, respectively, with $\mathcal{I}_H$ being the reconstructed image by the convolution between the PSF in Figure 1a and the impedance perturbation in Figure 1b. Figure 1e shows the difference between Figure 1c and 1d. We observe that $\mathcal{I}$ is an acceptable approximation to $\mathcal{I}_H$, apart from a minor phase shift and the presence of high-frequency noise induced by the Laplacian filter — this supports our use of local PSFs to approximate the effect of the full global Hessian. In this study, we aim at performing deblurring through i-RIM inference, using PSFs as the forward operator. In principle and in practice for deterministic deblurring approaches, high-frequency noise can affect the deblurring process; as such, it can be removed by filtering in a preprocessing step. However, we decide to keep it in the image as we observe that the deblurring process is implicitly regularized by the i-RIM network.

Invertible recurrent inference machines

I-RIMs are a special family of recurrent neural networks, particularly designed with inverse problems in mind. Let us consider the following linear problem with known forward operators:

$$\mathbf{d} = \mathbf{A}\mathbf{m} + \boldsymbol{\epsilon}, \tag{13}$$

where the observed data $\mathbf{d}$ are obtained by applying the forward operator $\mathbf{A}$ to the target model parameter $\mathbf{m}$, in the presence of noise $\boldsymbol{\epsilon}$. As an iterative inverse model, i-RIM takes the general form (Putzky and Welling, 2019)

$$\mathbf{m}_{t+1}, \mathbf{s}_{t+1} = \mathcal{F}(\mathbf{A}, \mathbf{d}, \mathbf{m}_t, \mathbf{s}_t; \boldsymbol{\phi}), \tag{14}$$

where $\mathbf{m}_t$ is the estimated model parameter and $\mathbf{s}_t$ is the auxiliary (memory) variable at iteration t. Here, the parametric function $\mathcal{F}$ is a neural network (NN) with trainable parameter $\boldsymbol{\phi}$. The network is tasked to produce an updated model (alongside an updated memory variable) combining the information of four inputs: whereas the latter two are the current estimates of the model and auxiliary parameters, the former two provide useful information about the data and forward model of the current problem. By unrolling equation 14 (i.e., writing it for each t from 1 to a total number of T), two alternative interpretations can be given to the i-RIM network: first, this can be seen as a recurrent neural network with a single learnable unit with shared parameters across the different steps; second, it can be interpreted as a natural extension of conventional iterative schemes such as the conjugate gradient or the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithms, for which a nonlinear learned transformation is used to update the model parameters at every iteration (instead of a linearly scaled version of the current and past gradients). For these reasons, i-RIMs also are usually referred to as a special case of the wider family of unrolled networks, which are nowadays ubiquitous in any scientific discipline dealing with inverse problems. Finally, the parametric function $\mathcal{F}$ of an i-RIM network is trained in a supervised fashion by minimizing the following loss function:

$$\mathcal{J} = \frac{1}{N}\sum_{i=1}^{N} \mathcal{L}(\mathbf{m}_T^i, \mathbf{m}^i), \tag{15}$$

where N is the number of training samples, $\mathbf{m}_T^i$ is the ith estimated model from the last recurrent step, and $\mathbf{m}^i$ is the corresponding ground truth. Finally, $\mathcal{L}$ is the loss function, which in this work is defined as the structural similarity index (SSIM). It is important to note that the training loss function is designed for the purpose of training/optimization of the NN architecture and is not the same as the metric/objective function that defines the inverse problem and constrains the gradient operator used within the i-RIM to constrain training and inference, which is identified in the following as $\mathbf{D}$.

When combined with an invertible structure h, the update equations of i-RIM for the forward step from t to $t+1$ can be written as

$$\begin{aligned} \mathbf{s}_t' &= \mathbf{s}_t + \nabla\mathbf{D}[\mathbf{d}, \mathbf{A}(\mathbf{m}_t)], \\ \mathbf{m}_{t+1}, \mathbf{s}_{t+1} &= h(\mathbf{m}_t, \mathbf{s}_t'). \end{aligned} \tag{16}$$

Correspondingly, the update equations for the reverse step from $t+1$ to t are

$$\begin{aligned} \mathbf{m}_t, \mathbf{s}_t' &= h^{-1}(\mathbf{m}_{t+1}, \mathbf{s}_{t+1}), \\ \mathbf{s}_t &= \mathbf{s}_t' - \nabla\mathbf{D}[\mathbf{d}, \mathbf{A}(\mathbf{m}_t)], \end{aligned} \tag{17}$$

where h^{-1} is the inverse of h. In this study, the same invertible structure h is shared through all recurrent steps, and therefore we drop the time dependency of h. Note that this choice may not be optimal for

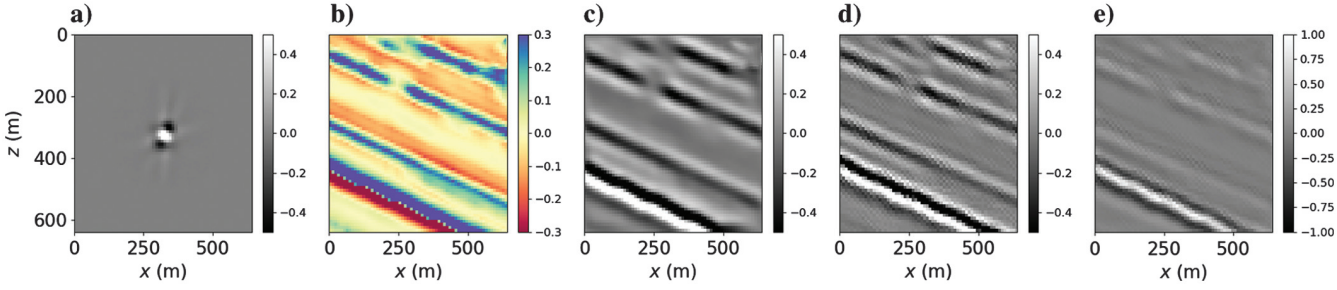

Figure 1. Numerical illustration of equation 12: (a) a PSF representing one column of the Hessian approximation; (b) a patch of the impedance perturbation model $\mathbf{m}$; (c) PSF-reconstructed image patch $\mathcal{I}_H$ by convolving (a) with (b); (d) a patch of the migrated image $\mathcal{I}$ from the PSDM impedance kernel, wherein the high-frequency noises are caused by the Laplacian filter and retained in the input data set for the i-RIM and the i-UNet; and (e) the difference plot between (c and d). The clipping factors used for (a–e) are 50%, 30%, 50%, 50%, 50%, and 100%, respectively.

other applications of i-RIM. Here, $\nabla \mathbf{D}[\mathbf{d}, \mathbf{A}(\mathbf{m}_t)]$ is the gradient of the data misfit term, which — within the i-RIM network architecture — ensures that the estimated $\mathbf{m}_t$ is consistent with the measured data $\mathbf{d}$ under the action of the forward operator $\mathbf{A}$ (Putzky and Welling, 2019). The additional $\mathbf{s}_t'$ is the hidden pairing variables of $\mathbf{s}_t$, introduced such that the NN has an invertible structure. During the training and testing stages, $\mathbf{s}_0$ and $\mathbf{s}_0'$ are initialized with zero, whereas $\mathbf{m}_0$ can be initialized with either the observed data $\mathbf{d}$ or a prior estimate of the target model $\mathbf{m}$. A schematic representation of the forward and reverse steps of the i-RIM recurrent unit is shown in Figure 2a and 2b, respectively. The invertible structure h comprises of L layers in total (from h^1 to h^L). With the property of memory complexity $\mathbf{O}(1)$, h can be layer-wise inverted as

$$h = h^L \cdot h^{L-1} \cdots h^1,$$
$$h^{-1} = (h^L)^{-1} \cdot (h^{L-1})^{-1} \cdots (h^1)^{-1}, \tag{18}$$

where h^L indicates the Lth layer in terms of depth. The forward and reverse computations of one exemplary invertible layer h^L are shown in Figure 3. The unique property of equation 18 allows the backpropagation of the invertible layers, without storing activations in a memory-efficient way.

NUMERICAL EXAMPLES

Preparation of training and testing data sets

In this section, we present the creation process for the training/testing data sets that will be used to teach the i-RIM and i-UNet networks to deblur migrated seismic images. A single synthetic model, shown in Figure 4, is used for training. In the following, we refer to this model as the SEAM right model. Figure 4a–4c shows the true velocity, smoothed velocity, and impedance perturbation obtained by subtracting the previous two models, respectively. The smoothed model shown in Figure 4b is used to initialize the i-RIM during training. Figure 4d–4f shows the migrated image, PSFs, and the strongly smoothed velocity background, used to generate the PSFs, respectively. The models in Figure 4b and 4f are generated by smoothing the model in Figure 4a with 20 pixels (smoothness factor 0.05) and 1000 pixels (smoothness factor 0.001), respectively. The smoothing process is applied to the inverse of the velocity model, using the Convolve2D operator in PyLops. Here, the source wavelet is a Ricker wavelet with a peak frequency of 20 Hz (see Appendix A). For all the forward modeling in this study, we use the spectral-element-based software Salvus (Afanasiev et al., 2019), without the free surface.

For training, the migrated images, the slightly smoothed velocity models, and the impedance perturbation models are cut regularly into square patches with a fixed size of 640 m. An overlapped random cutting also is applied to augment the training data set to several 20,480 patches in total. The PSFs are generated using the same modeling/migration procedure, by scaling the velocity at each of the point scatterers placed 320 m apart in Figure 4f. The obtained PSFs are interpolated when necessary, then masked with a circular mask and padded to the size of 640 m to reduce the crosstalk effects from the neighboring PSFs. For each benchmark, we perform training using the i-UNet; the same parameters are used wherever possible, in terms of the initial learning rate and the structure of the network. During all the training procedures, the total number of epochs is fixed to 100, and the learning rate is set to decay with a period of 30 epochs and a multiplicative factor of 0.1. All of the input/output patches are normalized to their maximum absolute values, such that i-RIM and i-UNet converge to desired results.

For testing, we design a series of three experiments of increasing complexity for the SEAM left model, namely with and without noise and with blended shots. The SEAM left model has a total of 948 patches. Figure 5a–5c shows the true velocity model, the slightly smoothed velocity model, and the impedance perturbation resulting from the difference between the previous two models, respectively. Figure 5d–5f shows the migrated image, PSFs, and the PSFs-related velocity background, respectively. The models in Figure 5b and 5f are generated by smoothing the model in Figure 5a with 20 pixels (smoothness factor 0.05) and 1000 pixels (smoothness factor 0.001), respectively. To compare the deblurring results for the different cases, we select a large area indicated by the cyan box in Figure 5d and three smaller areas indicated by the blue, red, and purple boxes. Appendix A shows the smoothed velocity model for these small areas (with smoothing of 50 pixels and smoothness

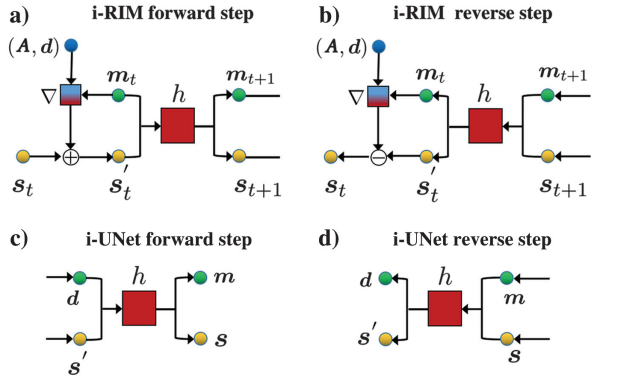

Figure 2. (a and b) The forward and reverse recurrent units of i-RIM, respectively. (c and d) The corresponding units for the i-UNet. All are adapted from Putzky and Welling (2019).

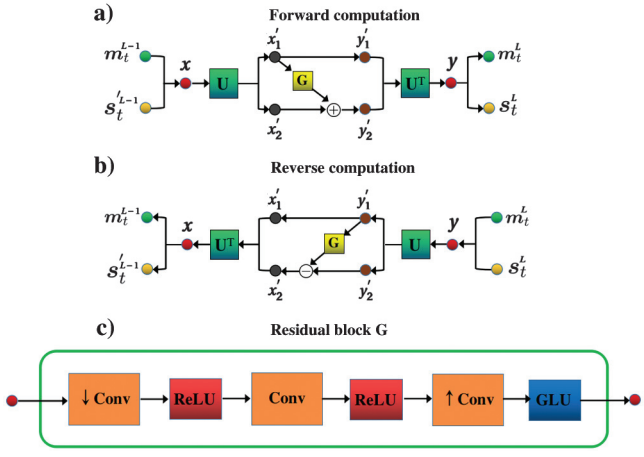

Figure 3. Adapted from Putzky and Welling (2019): (a and b) the forward and reverse computations of one invertible layer h^L, respectively. Here, $\mathbf{U}$ is a specially designed orthogonal 1×1 convolution with the property of $\mathbf{U}^{-1} = \mathbf{U}^T$. (c) The residual block $\mathbf{G}$ for spatial downsampling and upsampling, where GLU represents the gated linear unit. The up and down arrows indicate up and downsampling, respectively.

factor equal to 0.02) used for i-RIM during testing, as well as the unsmoothed velocity model for comparison.

The input patches used for testing are cut by the size of 640 m, with a vertically and horizontally overlapping distance of 320 m.

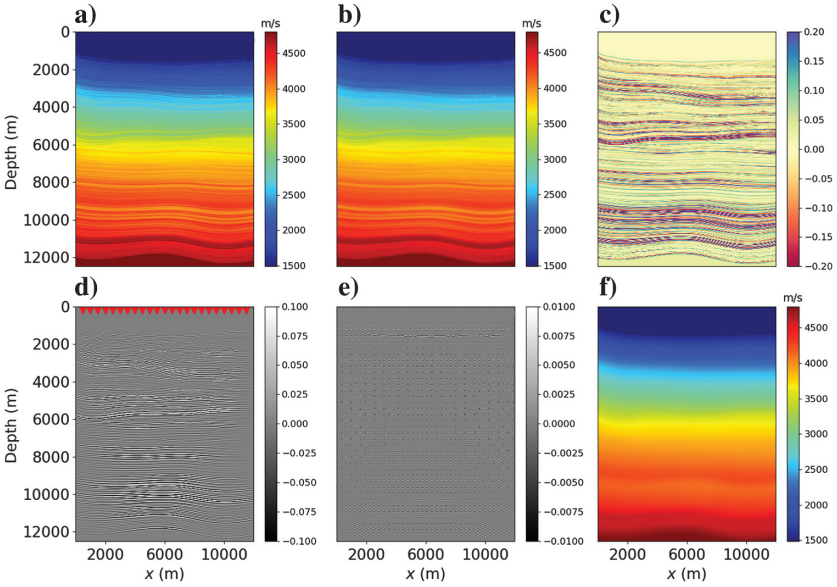

Figure 4. The synthetic model used for training in this study: (a and b) the true and slightly smoothed velocities of the SEAM right model. (c) The difference between (a and b), as the target impedance perturbation. (d–f) The migrated image, PSFs, and the smoothed background velocity used for generating PSFs, respectively. The red triangles in (a) indicate the locations of colocated sources and receivers. The clipping factors used for (a–f) are 100%, 100%, 20%, 20%, 1%, and 100%, respectively.

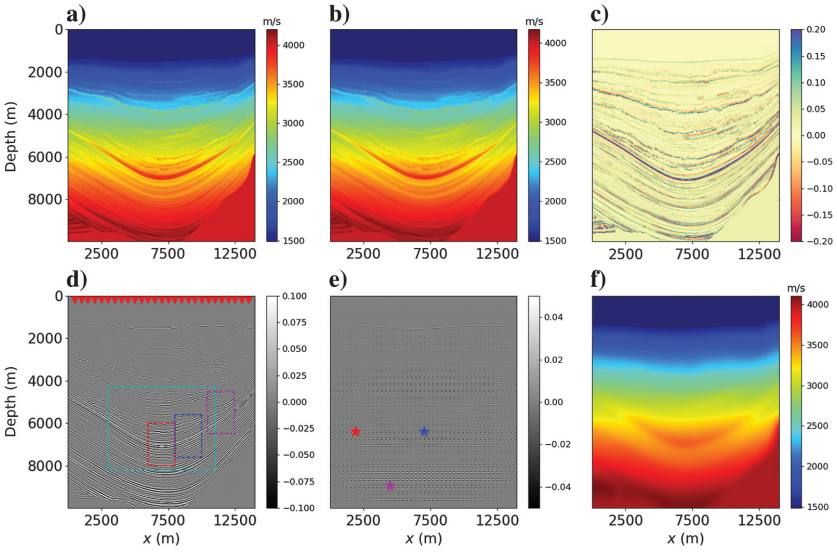

Figure 5. The synthetic SEAM left model used for testing. (a and b) The true and slightly smoothed velocities of the SEAM left model. (c) The difference between (a and b), as the target impedance perturbation. (d–f) The migrated image, PSFs, and the smoothed background velocity used for generating PSFs, respectively. The red triangles in (d) indicate the locations of colocated sources and receivers. The blue, red, and purple boxes in (d) denote the selected areas for the magnified comparison. A larger area indicated by the cyan box is chosen for a more general comparison. The blue, red, and purple stars in (e) are three PSFs selected to illustrate the PSF-deblurring effect. The clipping factors used for (a–f) are 100%, 100%, 20%, 10%, 5%, and 100%, respectively.

Due to the fact that the input data are normalized and also that NNs have an implicit input-dependent scaling factor that cannot be determined in advance, the output impedance perturbation patches are scaled inconsistently when compared with the ground truth. Therefore, two postprocessing steps are proposed to alleviate this issue. First, each output patch of the network is normalized by the maximum absolute value of the corresponding original image patch, before it is normalized to be fed to the network. Second, we use the Patch2D operator in the open-source toolbox PyLops (Ravasi and Vasconcelos, 2020) to assemble the output from i-RIM and i-UNet with a tapered overlapping. This first postprocessing step further reduces the relative scaling mismatch compared to the ground truth impedance perturbation, with an overall scaling factor. The second step yields an image with balanced amplitudes and smooth transitions between patches, albeit with the persistent scaling mismatch.

For the field Volve example, the data set comprises of 100 patches. Figure 6a–6d shows the PSDM-migrated image, the smooth velocity model (provided in the open-source data set) used for initializing the i-RIM network, the velocity background for the computation of the PSFs, and the PSFs generated with the field wavelet, respectively (see Appendix A). The model in Figure 6c is generated by smoothing the model in Figure 6b with 1000 pixels (smoothness factor 0.0001). The entire training and testing setup for the aforementioned scenarios (in total four cases) is summarized in Table 1. For comparison, the training data are always set to be consistent with the testing data, in terms of the presence of noise and/or blending.

Case 1: Testing on SEAM left model

We first show the testing results of image deblurring for the SEAM left model with the Ricker wavelet. For the area indicated by the cyan box in Figure 5d, Figure 7a–7d shows the migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. For the blue box area, Figure 8a–8d shows the magnified plot for the migrated images, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. Figure 8e–8h and 8i–8l corresponds to the areas indicated by the red and purple boxes, separately. In Appendix A, we compare the results of applying i-RIM and i-UNet directly with the PSFs. The training and validation losses of the i-RIM and the i-UNet for this case are shown in Figure 9a and 9b, respectively. We further illustrate the deblurring results by comparing a specific 1D profile in Figure 10a, with its location indicated by the red line shown in Figure 8i.

Case 2: Testing on SEAM left model with noise scenarios

In this second example, we assess the testing results of image deblurring for the SEAM left model with data modeled using a Ricker wavelet and further contaminated by the noise. More precisely, the noise is generated by filtering white Gaussian noise with a Ricker wavelet with a peak frequency of 26 Hz and added to the data directly before migration. Noise also is normalized trace wise such that its maximum amplitude is set to be 20% of the maximum amplitude of each data trace. For the area indicated by the cyan box in Figure 5d, Figure 11a–11d shows the migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. Figure 12a–12d shows the magnified plots for the blue box area, in the same order as in Figure 8a–8d. Figure 12e–12h and 12i–12l corresponds to the areas indicated by the red and purple boxes, separately. The training and validation losses are shown in Figure 9c and 9d, respectively. Similar to case 1, we show the 1D profile comparison in Figure 10b.

Case 3: Testing on SEAM left model with blended data

Here, we show the testing results of image deblurring for the SEAM left model using blended shots gathers. The same blending scheme, namely three neighboring shots firing simultaneously with random dither time between 0.1 s and 0.5 s, is used to create the training and testing data sets. For the area indicated by the cyan box in Figure 5d, Figure 13a–13d shows the migrated image, ground truth impedance perturbation, and outputs of i-RIM and i-UNet, respectively. Figure 14a–14d shows the magnified plots for the blue box area, in the same order as in Figure 8a–8d. Figure 14e–14h and 14i–14l corresponds to the areas indicated by the red and purple boxes, separately. For this case, the training and validation losses are shown in Figure 9e–9f, respectively, and the 1D profile comparison in Figure 10c. For the blue, red, purple, and cyan box areas, the correlation coefficients between the output from two different networks and the ground truth impedance perturbation are compared in Table 2.

Case 4: Testing on Volve field data

Finally, we present the application of our i-RIM network to the Volve field data set, for the target area indicated by the black box in Figure 6. Figure 15a–15d shows the input migrated image and the outputs of three different i-RIM networks, respectively. These i-RIMs are trained with different choices of the initial learning rates, network depths, and recurrent steps, as shown in Table 3. Figure 15e–15h shows the smoothed velocity used for initializing the i-RIMs and the outputs of three corresponding i-UNets, respectively.

DISCUSSION

Throughout this study, we focus on benchmarking the use of the i-RIM network for seismic image-domain deblurring against more "conventional" DL convolutional architectures such as the i-UNet — with the architectures being used in a supervised learning fashion. RIMs are designed to impose the known forward operator as a constraint within the training process, by building on the structure of recurrent neural networks to mimic the role of gradient-based optimization schemes, routinely applied to linear and nonlinear inverse problems. As such, our choice of focusing on comparing i-RIM versus i-UNet is aimed at understanding the importance of the forward operator in DL-based inference in the context of the image-domain seismic inversion problem. In supervised learning, the choice of training data, architecture, and training-optimization parameters implicitly dictate the priors that are constructed by training. Once a DL network is trained, these implicit priors are set within the final network parameters and weights that have been estimated by training optimization. Thus, when the trained network is applied to a new data set, inference is performed based on the implicit priors built onto the network through training. In our case, we provide the i-RIM and i-UNet networks with the same training data, which we purposefully choose to be "limited" — in terms of subsurface models and acquisition geometries/parameters — with

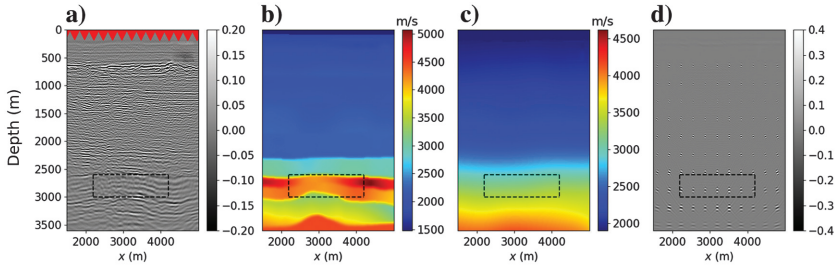

Figure 6. For the Volve data set: (a–d) the PSDM-migrated image, the smooth velocity model (included within the data set) used for initializing the i-RIM, the smooth velocity background used for generating the PSFs, and the PSFs, respectively. The black box indicates the target area. The red triangles indicate the sources and receivers. The clipping factors are 20%, 100%, 100%, and 40%, respectively.

Table 1. Summary for all training and testing scenarios.

			Testing data sets			
			SEAM left model	SEAM left model (noise)	SEAM left model (blend)	Volve field data (noise)
		Wavelet	Ricker wavelet	Ricker wavelet	Ricker wavelet	Field wavelet
Training data set	SEAM right model	Ricker wavelet	Case 1	Case 2	Case 3	
		Field wavelet				Case 4

For the training, adding noises or blending shots also is included, depending on the testing.

the intention of testing the inference capabilities of each network when presented with cases beyond those seen in training. In our original hypothesis, this is where the i-RIM network would in principle present an advantage over the i-UNet network: given the same training data, the i-RIM network is in fact supplied with two additional priors: (1) the generalizing knowledge of the forward operator, i.e., the physics of the inverse problem, and (2) the smoothed velocity model.

Following this thought process, in this work, the i-RIM and the i-UNet networks have been trained using a data set composed of a single, simple geologic model, mostly composed of mildly dipping fine layers. Deblurring of seismic images into their corresponding acoustic impedance perturbation models has been performed on four cases of increasing complexity. The main difference between the i-RIM and i-UNet approaches entails the input of the network; in the former case, we concatenate three different inputs, namely the migrated image, the PSF for the forward operator (i.e., a spatial convolution), and the corresponding smooth velocity model (e.g., a tomography- or full-waveform inversion [FWI]-inferred macromodel). Conversely, the i-UNet only takes a migrated image as the input.

In our experiments, we observe that when applying the trained networks to clean data of the synthetic model SEAM left, i-RIM

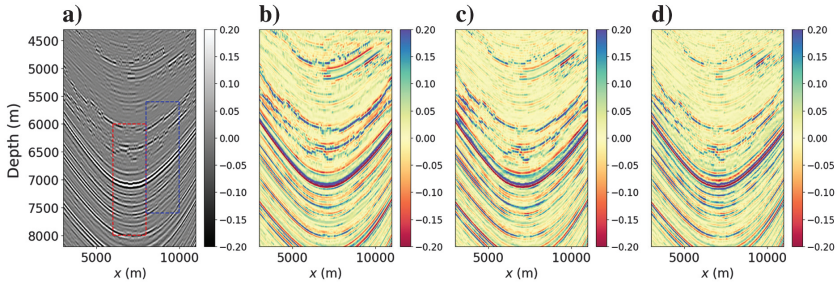

Figure 7. Results of seismic image deblurring for case 1: training on the SEAM right model and testing on the SEAM left model. For the area indicated by the cyan box in Figure 5d: (a–d) the migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. The red and blue boxes correspond to those in Figure 5d. The clipping factor used here is 20%.

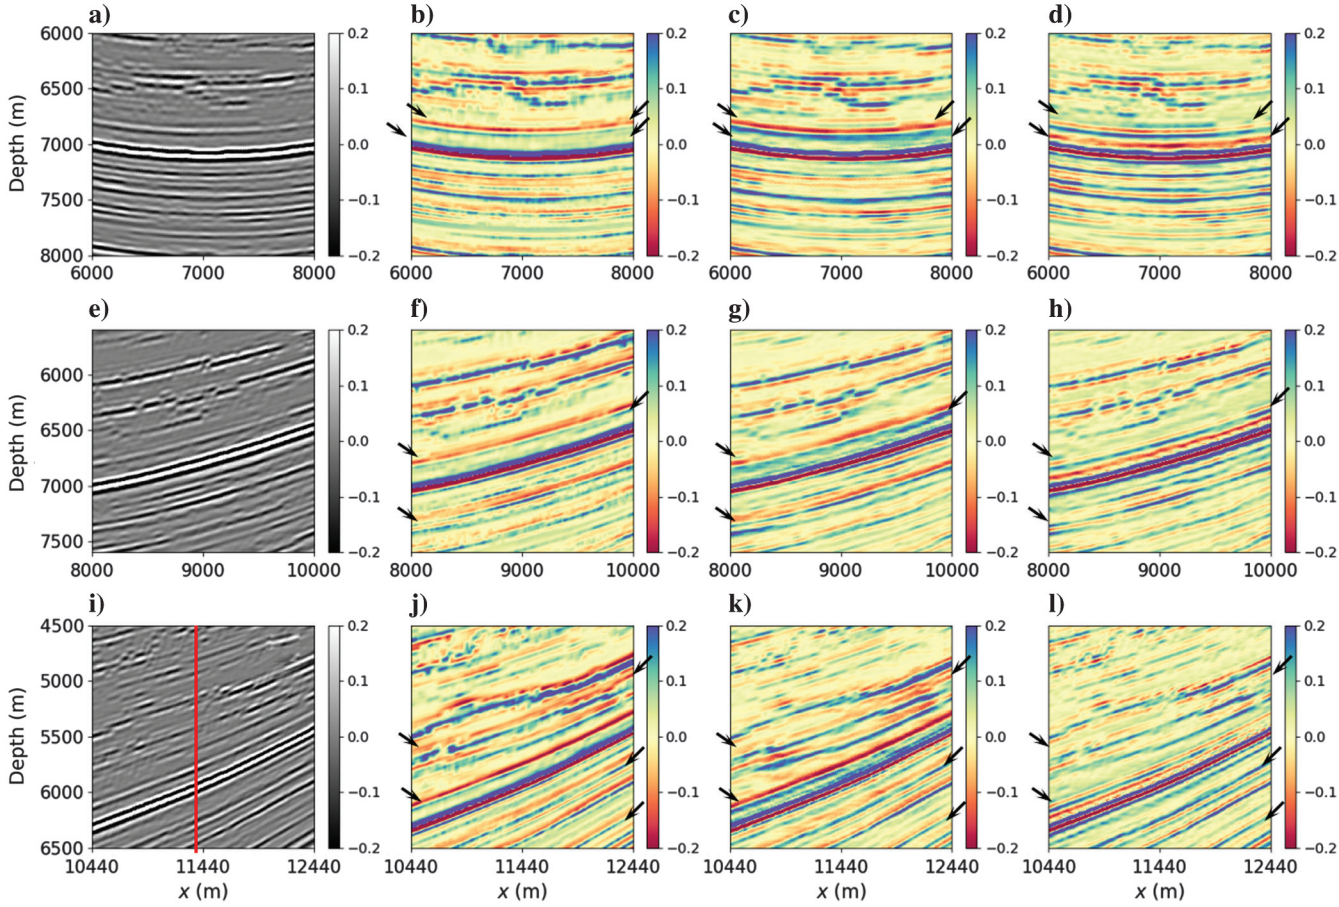

Figure 8. Case 1: magnified plots of seismic image deblurring for the areas indicated by the colored boxes shown in Figure 5d. The migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively, for the areas indicated by (a–d) the blue box, (e–h) the red box, and (i–l) the purple box. Here, we use black arrows to highlight some areas where the differences are clear. The solid red line indicates the location of the 1D profile. The clipping factor used here is 20%.

outperforms i-UNet in terms of amplitude preservation and structure fidelity. Pushing the limits of inference further, in the presence of noise and source-blended data, the i-RIM is consistently more robust; in comparison, the output of i-UNet clearly deteriorates in terms of layer continuity, and structural artifacts start to arise. When comparing these results using the Pearson correlation coefficient as our metric (as it does not depend on the absolute scaling), we can see that the values of i-RIM are consistently higher than those of the i-UNet. For cases 1–3, we observe that the advantage of using i-RIM over the i-UNet increases with (a) increasing layer dips and (b) locations that are closer to the edge of the acquisition aperture (comparing areas indicated by the blue, red, and purple boxes). In view of the quasi-1.5D feature of the single training data set, these two observations demonstrate the superior generalization capability of i-RIM over i-UNet, as shown in Table 2. In particular, observation (b) further implies the importance of providing the knowledge captured into PSFs to the network because this captures the aperture effect of the imaging system.

One thing to notice when examining the training and validation losses is that the networks appear to be overfitting, i.e., the validation loss is constantly noticeably lower than that of the training loss (in which lower is bad given that we use the SSIM metric as our loss function). We believe that this behavior can be partly explained by the difference between the feature of the training and the validation data sets, with the former mostly composed of near-flat thin layers and the latter comprised of much stronger lateral variations and varying layer thickness. Future work will be devoted to further understanding the reasons behind this overfitting phenomenon; for example, we will test if such a gap in the training and validation losses persists when adding additional training data sets that share similar features of the validation data set. In comparison, we also notice that none of the previously published RIM papers shows any training/validation losses (Morningstar et al., 2018, 2019; Lønning et al., 2019; Putzky et al., 2019; Kuijpers et al., 2020; Sabidussi et al., 2021). However, all of our purposely designed synthetic case studies still show that the training and validation losses of i-RIM are

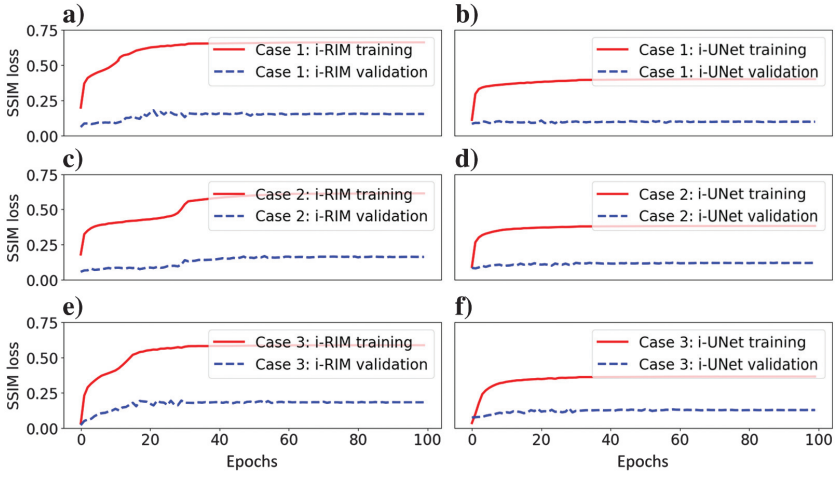

Figure 9. The absolute value of average training (red) and validation (blue) loss per patch of the i-RIM and the i-UNet, respectively, for (a and b) case 1, (c and d) case 2, and (e and f) case 3. Note that for the SSIM metric, higher values indicate better results.

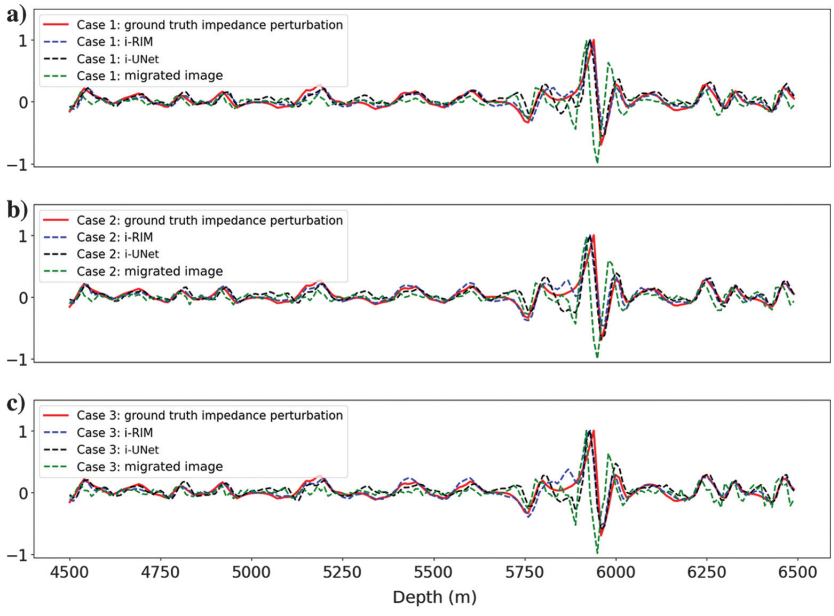

Figure 10. (a–c) The 1D profile plots of the i-RIM and i-UNet for cases 1–3, respectively. The location of the 1D line is indicated by the red line shown in Figures 8i, 12i, and 14i, respectively. Here, the red, blue, black, and green lines correspond to the ground truth impedance perturbation, the i-RIM, i-UNet, and the migrated image, respectively.

consistently higher than those of its counterpart, suggesting its higher data assimilation capacity.

The contrast between the i-RIM and the i-UNet is even more clear when they are applied to the Volve data set that, by virtue of being a field data set, is characterized by subsurface and noise conditions that lie outside those in the training data. In this case, we compare the inference outputs of different i-RIMs and i-UNets trained with three sets of parameters; this allows us to gauge the robustness of each network to hyperparameters given that we do not have access to the ground truth to compare the different outcomes. The results of different i-RIMs appear to be much smoother, cleaner, and consistent, with major image features almost identical, indicating a high degree of robustness to training parameters — which in turn supports the credibility of the i-RIM results. In contrast, the i-UNet output images vary with respect to layer continuity and thickness when different trained networks are used. Although our field data application yields promising results in terms of robustness, bandwidth reconstruction, and structural interpretability, further research is needed toward retrieving quantitatively reliable amplitudes. For example, we believe that by adding attenuation and/or elastic-amplitude effects to the training data set, the quantitative reliability of amplitudes could be further improved. The application of iRIM to the Volve field data set demonstrates that by leveraging the constraint of the forward operator, broader inference capabilities can

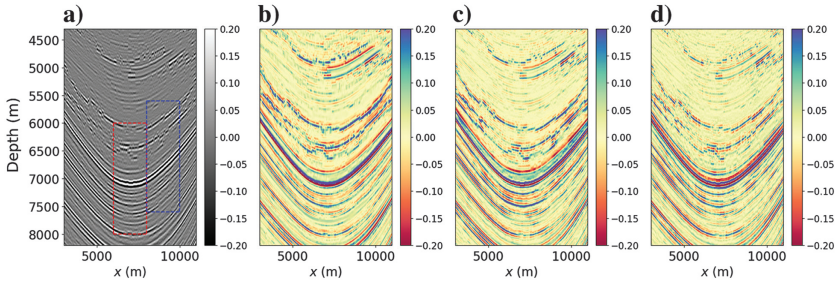

Figure 11. Results of seismic image deblurring for case 2: training on the SEAM right model with noises and testing on the SEAM left model with noises. For the area indicated by the cyan box in Figure 5d: (a–d) the migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. The red and blue boxes correspond to those in Figure 5d. The clipping factor used here is 20%.

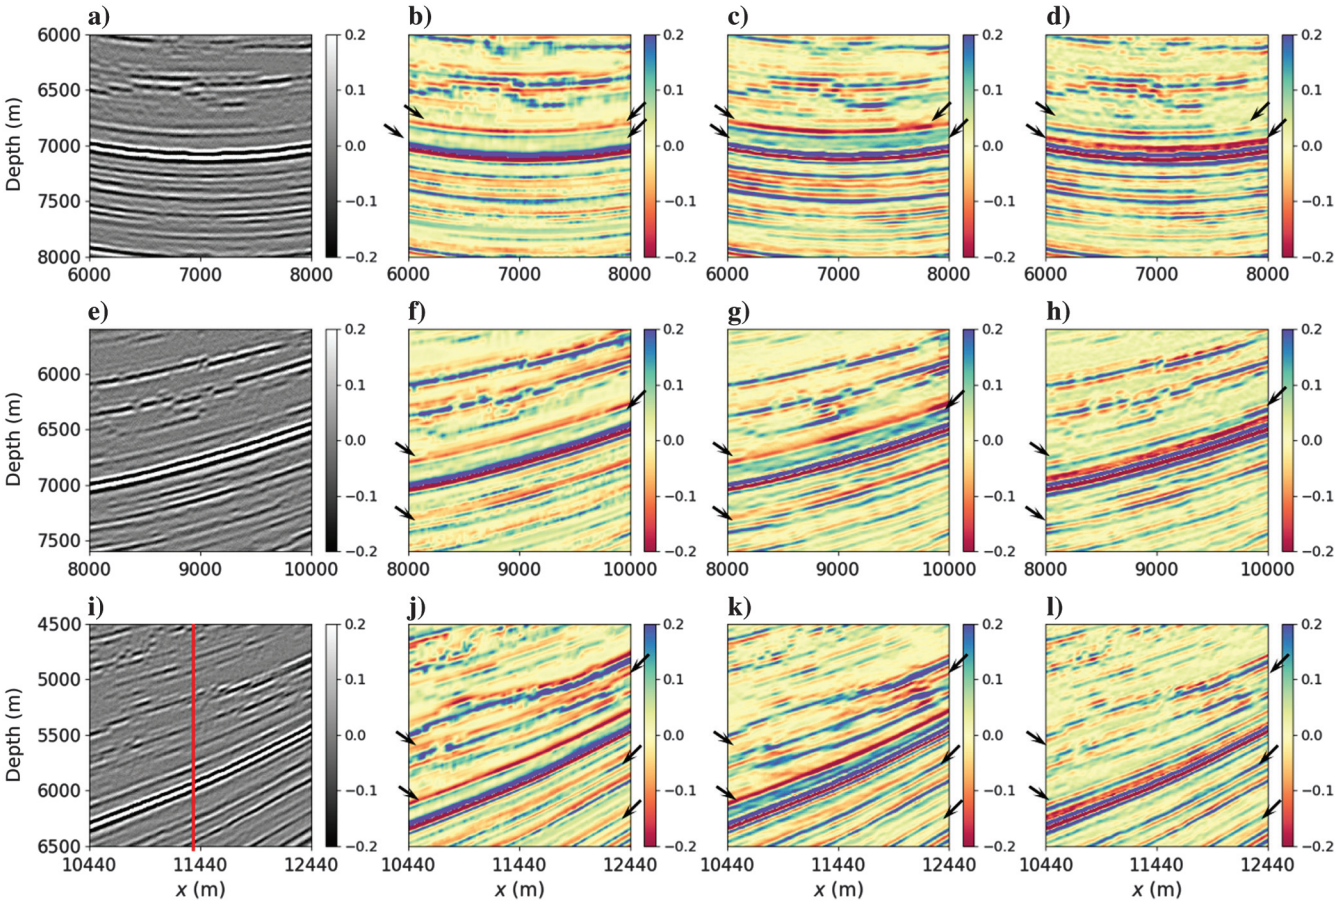

Figure 12. Case 2: magnified plots of seismic image deblurring for the areas indicated by the colored boxes shown in Figure 5d. The migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively, for the areas indicated by (a–d) the blue box, (e–h) the red box, and (i–l) the purple box. Here, we use black arrows to highlight some areas where the differences are clear. The solid red line indicates the location of the 1D profile. The clipping factor used here is 20%.

be reached. This is achieved via limited training information, such as the synthetic SEAM right model, instead of requiring unreasonably complex or large training data, which is usually either inaccessible or intractable.

Here, we have chosen to deblur migrated images parameterized in terms of impedance perturbation, instead of reflectivity as used in previous studies. We stress here that this choice lays the groundwork for quantitative reservoir impedance inversion in future studies. However, in practice, during i-RIM training and inference, the use of the impedance parameterization implies a free parameter regarding the degree of smoothness of the background impedance model — used by the i-RIM for an initial model. In our study, this parameter is selected in training by examining the structural and bandwidth similarity between the reconstructed imaged $\mathcal{I}_H$ and the migrated image $\mathcal{I}$ because we observe that $\mathcal{I}_H$ significantly diverges from $\mathcal{I}$, when the impedance is overly smoothed to define the perturbation. This is because in such cases, the impedance perturbation has strong low-frequency components, which leak into $\mathcal{I}_H$ and are missing in $\mathcal{I}$.

From a more pragmatic perspective, we notice that initializing with a smooth velocity model generally helps to improve the output of i-RIM. In our SEAM left model tests, we use a model that we believe mimics what would have been a reasonable output of FWI for the background model. In the case of the Volve field data set,

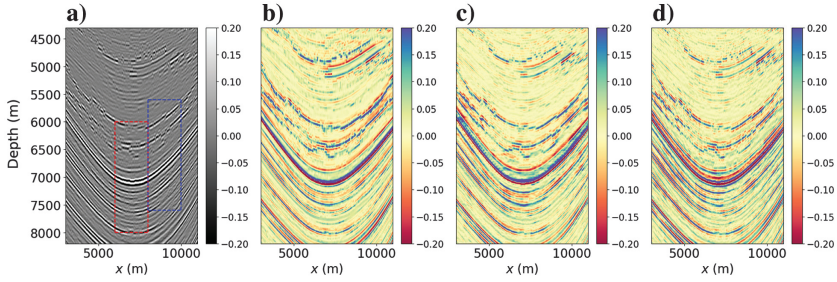

Figure 13. Results of seismic image deblurring for case 3: training on the SEAM right model with blended shots and testing on the SEAM left model with blended shots. For the area indicated by the cyan box in Figure 5d: (a–d) the migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively. The red and blue boxes correspond to those in Figure 5d. The clipping factor used here is 20%.

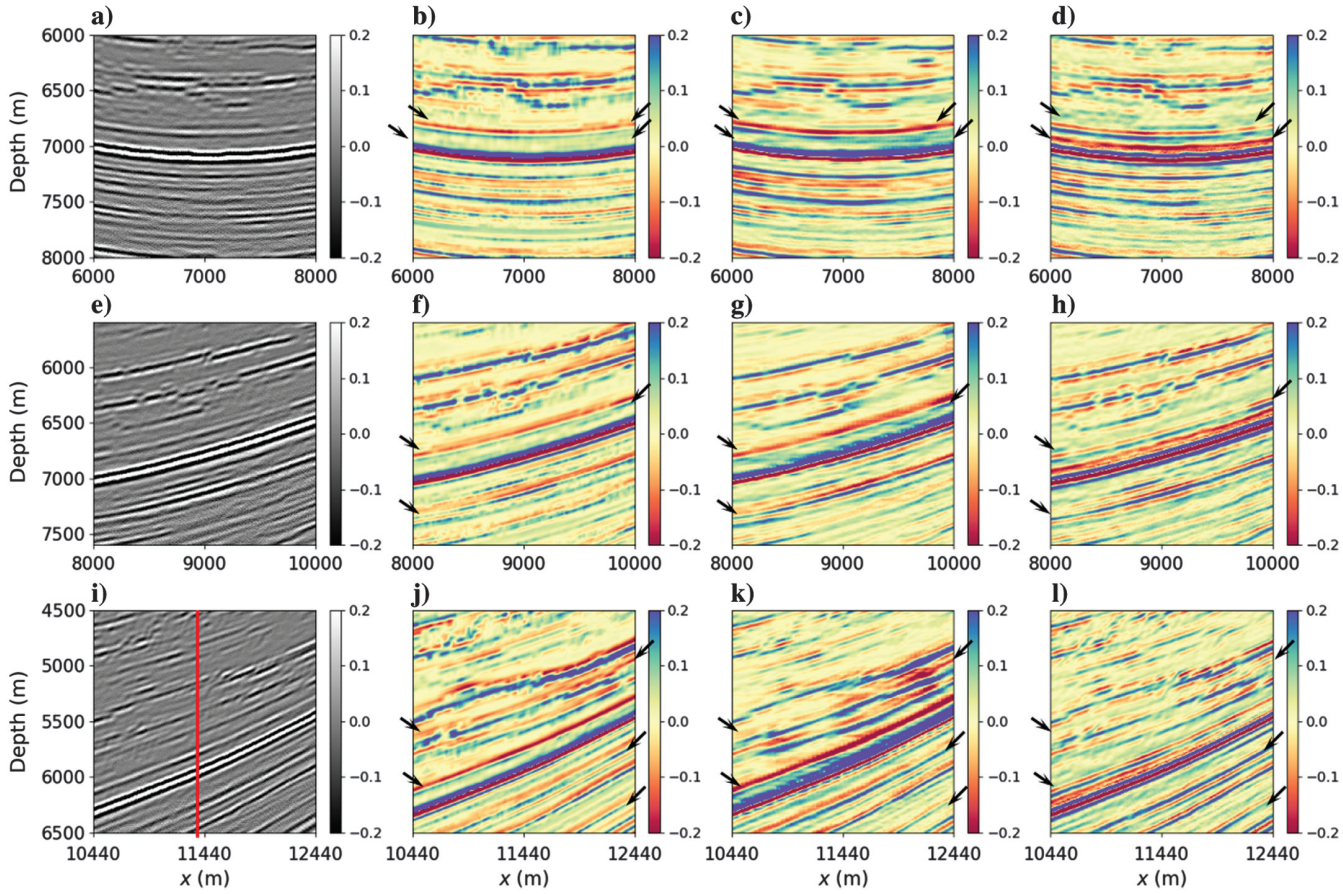

Figure 14. Case 3: magnified plots of seismic image deblurring for the areas indicated by the colored boxes shown in Figure 5d. The migrated image, the ground truth impedance perturbation, and the outputs of i-RIM and i-UNet, respectively, for the areas indicated by (a–d) the blue box, (e–h) the red box, and (i–l) the purple box. Here, we use black arrows to highlight some areas where the differences are clear. The solid red line indicates the location of the 1D profile. The clipping factor used is 20%.

DOI:10.1190/geo2022-0780.1

much-smoothed velocity models are used, resulting from tomographic velocity analysis of the original data. In our experience, this careful consideration in supplying an initial background model is only necessary because of the choice of the impedance perturbation parameterization. According to our previous studies (Vasconcelos et al., 2022), deblurring a migrated image parameterized as reflectivity instead only requires initializing the i-RIM with the input migrated image itself. We believe that this could be a result of the characteristic image-feature differences between impedance perturbation and reflectivity models because the former contains high-frequency and low-frequency components (i.e., outside the seismic bandwidth), whereas the latter mainly contains the high-frequency part (i.e., within the seismic bandwidth). Initializing the i-RIM with a smooth velocity model compensates for the missing low-frequency component, thus enhancing the results for the mapping of the impedance perturbation. When opting for reflectivity models (Vasconcelos et al., 2022), we observe a clear phase shift between the reconstructed PSF and the migrated image patches. However, this issue can be dealt with by including extra filtering operators into the i-RIM forward operator.

In this study, we have conducted seismic image deblurring in two dimensions. However, we foresee that our approach should in principle extend to three dimensions as well, owing to the fact that the convolutional relation between PSFs and local images holds as a straightforward extension of the two dimensions. This convolutional relation has been previously used in, e.g., 3D seismic image deblurring (Cavalca et al., 2020; Liu and Fu, 2021) and 3D microscopy deblurring (Bruce and Butte, 2013) — these being examples of deterministic PSF deblurring by optimization. In the context of DL-based 3D deblurring, Li et al. (2022b) propose a PSF-based DL approach that is specifically designed for 3D microscopy image deconvolution — but overall 3D DL-based deblurring, and in particular the use of RIMs in 3D image-processing, remains the subject of future research. For our seismic imaging problem, on the one hand, we can take advantage of fast 3D PSDM algorithms, or that PSFs can be constructed analytically in many cases, to further speed up the process of producing PSFs for training. However, with increasing structural complexity included in training data, 3D migration costs could become an issue in training; this will perhaps be another argument in favor of i-RIMs over architectures with no forward-operator constraints, whereas the RIM may be important to extend the benefits of limited training data in the context of 3D imaging. On the other hand, the memory cost of the networks

Table 2. Comparison of correlation coefficients for the results of i-RIM and i-UNet of cases 1–3.

Target area	Networks	Case 1	Case 2	Case 3
Blue box	i-RIM	0.831	0.800	0.759
	i-UNet	0.760	0.760	0.753
Red box	i-RIM	0.866	0.829	0.799
	i-UNet	0.801	0.789	0.781
Purple box	i-RIM	0.863	0.825	0.815
	i-UNet	0.775	0.758	0.761
Cyan box	i-RIM	0.829	0.835	0.814
	i-UNet	0.773	0.774	0.767

Areas indicated by the blue, red, purple, and cyan boxes are shown in Figure 5d.

Table 3. Summary of i-RIM and i-UNet parameters for all training and testing scenarios.

	Case 1	Case 2	Case 3	Case 4
Initial learning rate (e^{-4})	2	2	2	5, 2, 2
Network depth	5	4	5	4, 5, 6
i-RIM recurrent steps	5	5	5	3, 5, 6
Memory cost (Mb)				
i-RIM	1.9	1	1.9	1, 1.9, 2.7
i-UNet	1.9	1	1.9	1, 1.9, 2.7
Inference time cost (ms)				
i-RIM	13.4	11.0	13.4	7.2, 13.4, 18.7
i-UNet	2.5	2.0	2.5	2.0, 2.5, 3.0

Here, i-RIM and i-UNet share the same parameters, in terms of the initial learning rate and NN depth, with recurrent steps only applied to i-RIM. In our tests, i-RIM and i-UNet have the same memory costs because the weights are shared throughout the recurrent steps for the former. The memory cost is estimated after one epoch of forward and backward process, without including the input/output patches. The inference time cost per image patch is estimated by repeatedly running the forward operation 100 times for 100 input patches. The three parameters for case 4 correspond to Figure 15b–15h, for i-RIMs and i-UNets, respectively.

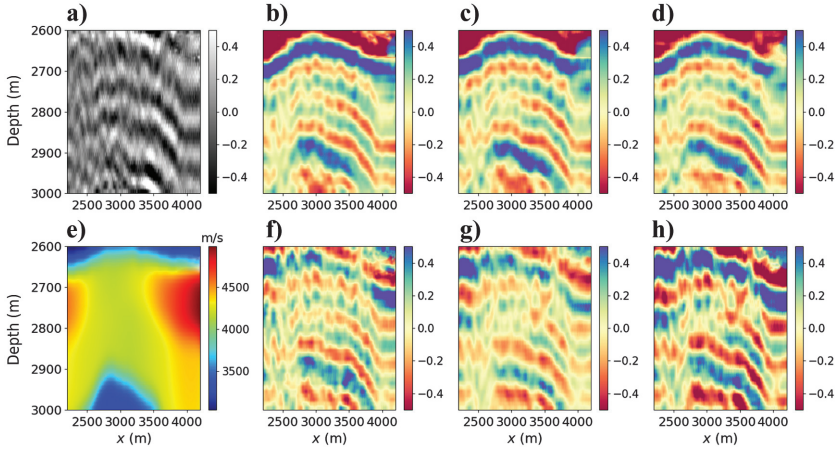

Figure 15. Results of seismic image deblurring for case 4: training on the SEAM right model with the field wavelet and noises and testing on the target area of the Volve field data. (a–d) The migrated image and the outputs from three different i-RIMs (with parameters summarized in Table 2, respectively. (e–h) The smooth velocity model used for initializing i-RIMs and outputs from three corresponding i-UNets. The clipping factor used for (e) is 100% and for the rest 50%, respectively. The position of the target area is indicated by the black box shown in Figure 6.

themselves will likely not increase drastically, thanks to the use of invertible layer structures. Furthermore, in three dimensions, we expect the role of the forward operator to possibly be more important because now the feature space is expanded with more parameters, and PSFs can still capture the prior information of the additional dimension of the imaging system. In addition, eventual i-RIM-based inference for 3D deblurring is likely to significantly outperform 3D deconvolution by optimization in terms of computational efficiency, thus potentially enabling fast, wide-bandwidth image deblurring of large 3D seismic volumes.

Although our study performs a benchmark of the i-RIM network against the i-UNet network (in which no modeling operator is involved in the forward pass) for the reasons mentioned previously, it is important to consider the implications of DL-based deblurring for practical image-domain processing. First, even in two dimensions, it is well known that deterministic image-domain deblurring using PSFs through conventional, optimization-based least-squares inversion typically requires (1) case-dependent expert-user inputs such as preconditioning operators, (2) additional free parameters for regularization, which typically require user-based tuning, and (3) a large number of iterations to converge. In comparison, the data-driven i-RIM used in this study, once properly trained, is able to (1) relieve the need for human intervention in image-domain deblurring, (2) attain implicit regularization and preconditioning by learning, and (3) perform deblurring-by-inference within only a few seconds for all image patches. Thinking forward toward 3D applications, in which the computational cost of least-squares inversion escalates rapidly and the need for optimization-parameter tuning becomes a bigger hurdle, our i-RIM could potentially provide a significant advantage in the future for 3D image-domain impedance inversion. Of course, further research is necessary to bridge the study presented in this paper to a full-fledged 3D image application and benchmark against the optimization-based least-squares inversion approach (Du et al., 2016, 2021).

CONCLUSION

Making explicit use of the relation between PSFs and local migrated images, we apply the i-RIM network to the problem of deblurring of seismic images into impedance perturbation models; the proposed approach uses 2D convolutions with precomputed PSFs as the forward operator and is assisted with the prior information of a smooth velocity model. With a limited training data set, we show that i-RIM is able to consistently outperform its i-UNet counterpart (representing a well-established convolutional NN architecture) with respect to noise regularization and structure consistency. This is attributed to the information of the forward operator acting as a prior for the i-RIM, helping the trained network perform well when used for inference on situations beyond those represented in the training data set. We support these observations with results from a series of tests on different synthetic models and acquisition scenarios and a North Sea field data set.

Our work adds new elements to previous studies of DL applications to seismic image deblurring in that: (1) we formulate deblurring as an inverse problem with a known forward operator constraint, informed by depth-domain PSFs and the additional prior information on background velocity, and (2) we map images directly into impedance perturbation models, instead of reflectivity, aiming to provide a DL-based tool for quantitative reservoir characterization and monitoring. Based on our results and other related studies, we see the i-RIM as a powerful tool for addressing inverse problems that are widely present in geophysics, given known and computationally feasible forward operators — as long as these can be treated as local problems to take advantage of current RIM architectures.

ACKNOWLEDGMENTS

The authors thank all the sponsors of the Utrecht Consortium for the Subsurface Imaging (UCSI) for funding this project. We also thank the software Salvus for spectral-element-based modeling. In addition, we appreciate the open-source code scikit-fmm for computing traveltime efficiently. We thank the open-source software PyLops for the implementation of the Patch2D operator. We also thank all the reviewers who helped to improve this paper. Finally, we thank P. Putzky for sharing the i-RIM and i-UNet codes online.

DATA AND MATERIALS AVAILABILITY

Data associated with this research are available and can be obtained by contacting the corresponding author.

APPENDIX A

WAVELET COMPARISON, SMOOTHED VELOCITY MODEL INPUT FOR I-RIM, AND PSF DEBLURRING COMPARISON

The Ricker wavelet and the field wavelet used in this study are given in the time domain and the frequency domain in Figure A-1a and A-1b, respectively. Sections of the unsmoothed velocity model and their smoothed counterparts that are used to initialize the i-RIM are shown in Figure A-2a–A-2f, respectively.

Here, we compare the results of applying the i-RIM and the i-UNet to deblurring PSFs directly. Figure A-3a–A-3c shows the PSF, and outputs of i-RIM and i-UNet in the space-time (x-t) domain, respectively, corresponding to the red star in Figure 5e, with Figure A-3d–A-3f showing the same results in the $K_x - K_z$ domain. Similarly, Figure A-3g–A-3l and A-3m–A-3r corresponds to the blue and purple stars in Figure 5e, respectively. Note that neither the i-RIM nor the i-UNet is provided with point-scatterer images

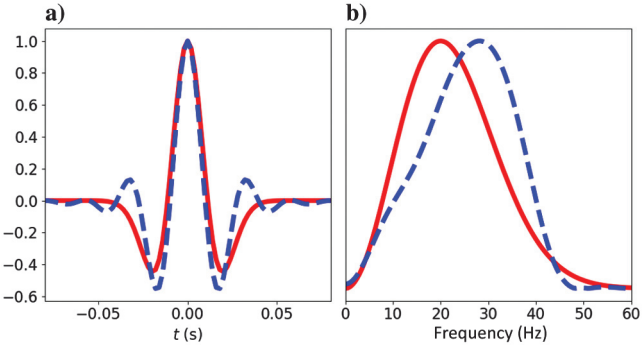

Figure A-1. The two wavelets used in this study: for cases 1–3, the Ricker wavelet with a peak frequency of 20 Hz, indicated by the solid red line, and for case 4, the field wavelet extracted from the Volve field data set, indicated by the dashed blue line. (a and b) Plots in the time and frequency domains, respectively.

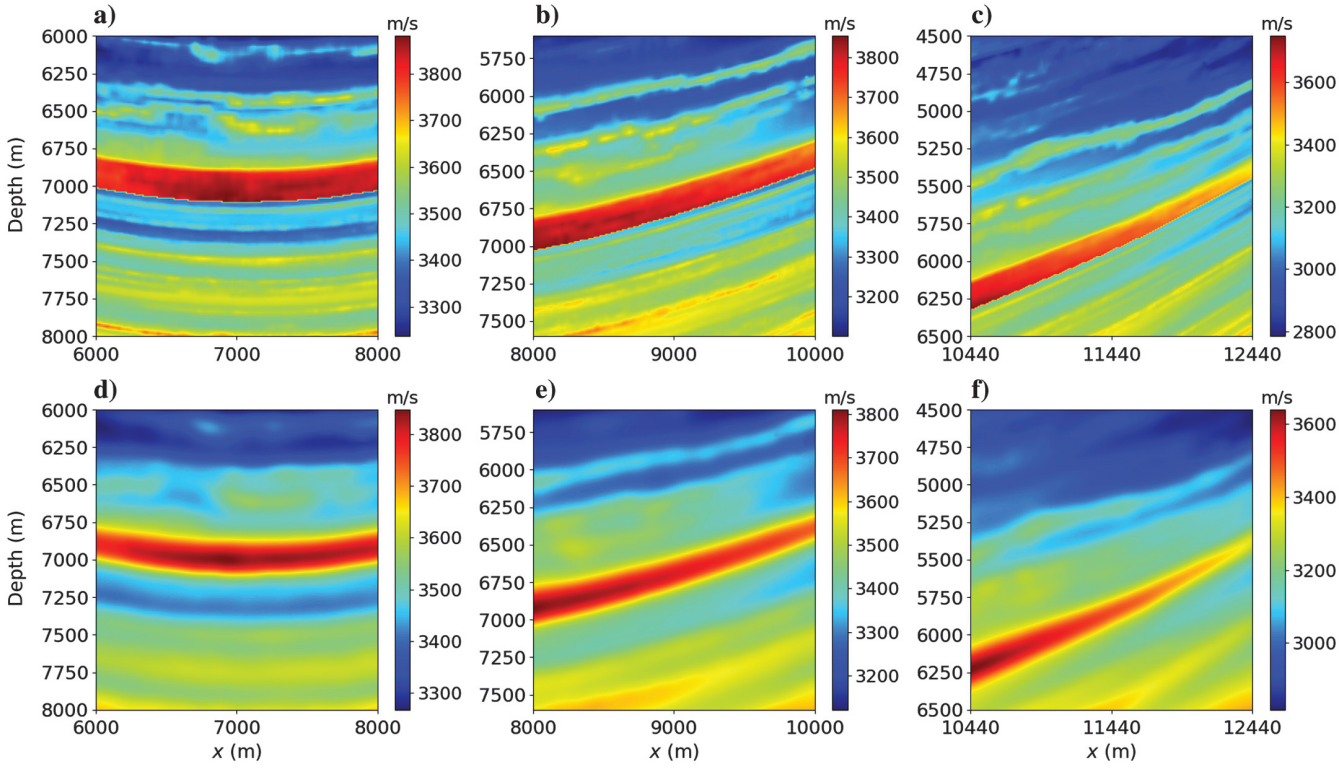

Figure A-2. (a–c) The unsmoothed velocity of sections corresponding to the areas indicated by the red, blue, and purple boxes in Figure 5d, respectively, and (d–f) the corresponding smoothed sections with 50 pixels and a smoothness factor of 0.02, respectively, which resemble typical velocity models built by FWI. We use (d–f) to initialize the i-RIM.

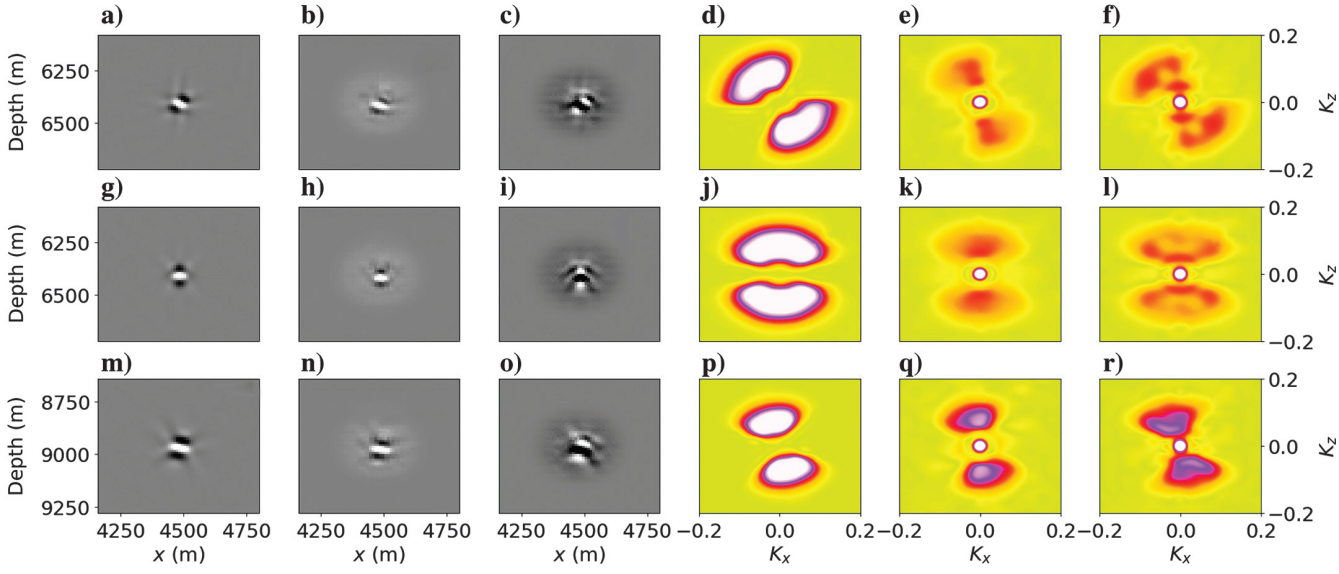

Figure A-3. Case 1: for the position indicated by the red star in Figure 5e, the PSF and the outputs of i-RIM and i-UNet in, respectively, (a–c) the *x*-*t* domain and (d–f) the $K_x - K_z$ domain. For the position indicated by the blue star in Figure 5e, the PSF and the outputs of i-RIM and i-UNet in, respectively, (g–i) the *x*-*t* domain and (j–l) the $K_x - K_z$ domain. For the position indicated by the purple star in Figure 5e, the PSF and the outputs of i-RIM and i-UNet in, respectively, (n–o) the *x*-*t* domain and (p–r) the $K_x - K_z$ domain. The PSFs are applied with a smooth circular mask, which also is applied to the outputs of i-RIM and i-UNet. The clipping factors used are 50% for the *x*-*t* domain and 60% for the $K_x - K_z$ domain.

during training; therefore, testing PSF deblurring on point scatterers is outside the range of the original training data. This is because point scatterers contain all dips, whereas the training images are dip limited. In addition, the i-RIM is initialized with the PSF itself, instead of a smooth background velocity used for image deblurring shown previously. This initialization is different from that used during the training stage.

REFERENCES

Afanasiev, M., C. Boehm, M. van Driel, L. Krischer, M. Rietmann, D. A. May, M. G. Knepley, and A. Fichtner, 2019, Modular and flexible spectral-element waveform modelling in two and three dimensions: Geophysical Journal International, **216**, 1675–1692, doi: 10.1093/gji/ggy469.

Aoki, N., and G. T. Schuster, 2009, Fast least-squares migration with a deblurring filter: Geophysics, **74**, no. 6, WCA83–WCA93, doi: 10.1190/1.3155162.

Beylkin, G., 1985, Imaging of discontinuities in the inverse scattering problem by inversion of a causal generalized Radon transform: Journal of Mathematical Physics, **26**, 99–108, doi: 10.1063/1.526755.

Bleistein, N., 1987, On the imaging of reflectors in the earth: Geophysics, **52**, 931–942, doi: 10.1190/1.1442363.

Bruce, M. A., and M. J. Butte, 2013, Real-time GPU-based 3D deconvolution: Optics Express, **21**, 4766–4773, doi: 10.1364/OE.21.004766.

Cavalca, M., R. Fletcher, and P. Caprioli, 2016, Least-squares Kirchhoff depth migration in the image domain: 78th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.201601374.

Cavalca, M., R. P. Fletcher, M. Shadrina, C. Leone, and L. Leon, 2020, Depth-domain inversion for enhanced quantitative interpretation: First EAGE Conference on Seismic Inversion, doi: 10.3997/2214-4609.202037038.

Chavent, G., and R. E. Plessix, 1999, An optimal true-amplitude least-squares prestack depth-migration operator: Geophysics, **64**, 508–515, doi: 10.1190/1.1444557.

Claerbout, J. F., 1992, Earth soundings analysis, processing versus inversion: Blackwell Scientific Publications.

Clapp, M. L., 2005, Imaging under salt: Illumination compensation by regularized inversion: Ph.D. thesis, Stanford University.

Clayton, R. W., and R. H. Stolt, 1981, A Born-WKBJ inversion method for acoustic reflection data: Geophysics, **46**, 1559–1567, doi: 10.1190/1.1441162.

Douma, H., D. Yingst, I. Vasconcelos, and J. Tromp, 2010, On the connection between artifact filtering in reverse-time migration and adjoint tomography: Geophysics, **75**, no. 6, S219–S223, doi: 10.1190/1.3505124.

Du, X., R. Fletcher, and M. Cavalca, 2016, Prestack depth-domain inversion after reverse time migration: 78th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.201601429.

Du, X., R. Fletcher, and M. Cavalca, 2021, Prestack least-squares reverse-time migration for elastic property estimation: First International Meeting for Applied Geoscience & Energy, SEG, Expanded Abstracts, 196–200, doi: 10.1190/segam2021-3580936.1.

Duquet, B., and K. J. Marfurt, 1999, Filtering coherent noise during prestack depth migration: Geophysics, **64**, 1054–1066, doi: 10.1190/1.1444613.

Dutta, P., B. Power, A. Halpert, C. Ezequiel, A. Subramanian, C. Chatterjee, S. Hari, K. Prindle, V. Vaddina, A. Leach, R. Domala, L. Bandura, and M. Mascaro, 2019, 3D conditional generative adversarial networks to enable large-scale seismic image enhancement: Second Workshop on Machine Learning and the Physical Sciences.

Fehler, M., and P. J. Keliher, 2011, SEAM phase 1: Challenges of subsalt imaging in tertiary basins, with emphasis on deepwater Gulf of Mexico: SEG.

Fletcher, R. P., S. Archer, D. Nichols, and W. Mao, 2012, Inversion after depth imaging: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: 10.1190/segam2012-0427.1.

Fletcher, R. P., D. Nichols, R. Bloor, and R. T. Coates, 2016, Least-squares migration — Data domain versus image domain using point spread functions: The Leading Edge, **35**, 157–162, doi: 10.1190/tle35020157.1.

Gray, S. H., 1997, True-amplitude seismic migration: A comparison of three approaches: Geophysics, **62**, 929–936, doi: 10.1190/1.1444200.

Guitton, A., 2004, Amplitude and kinematic corrections of migrated images for nonunitary imaging operators: Geophysics, **69**, 1017–1024, doi: 10.1190/1.1778244.

Halpert, A. D., 2018, Deep learning-enabled seismic image enhancement: 88th Annual International Meeting, SEG, Expanded Abstracts, 2081–2085, doi: 10.1190/segam2018-2996943.1.

Hu, J., G. T. Schuster, and P. A. Valasek, 2001, Poststack migration deconvolution: Geophysics, **66**, 939–952, doi: 10.1190/1.1444984.

Kaur, H., N. Pham, and S. Fomel, 2020, Improving the resolution of migrated images by approximating the inverse Hessian using deep learning: Geophysics, **85**, no. 4, WA173–WA183, doi: 10.1190/geo2019-0315.1.

Kuijpers, D., I. Vasconcelos, and P. Putzky, 2020, Reconstructing missing seismic data through deep learning with recurrent inference machines: 82nd Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.202011046.

Lailly, P., 1983, The seismic inverse problem as a sequence of before stack migration: Conference on Inverse Scattering, Theory and Applications, SIAM, 206–220.

Lambaré, G., J. Virieux, R. Madariaga, and S. Jin, 1992, Iterative asymptotic inversion in the acoustic approximation: Geophysics, **57**, 1138–1154, doi: 10.1190/1.1443328.

Lecomte, I., 2008, Resolution and illumination analyses in PSDM: A ray-based approach: The Leading Edge, **27**, 650–663, doi: 10.1190/1.2919584.

Lecomte, I., and L. Gelius, 1998, Have a look at the resolution of prestack depth migration for any model, survey and wavefields: 68th Annual International Meeting, SEG, Expanded Abstracts, 1112–1115, doi: 10.1190/1.1820082.

Letki, L., K. Darke, and Y. A. Borges, 2015, A comparison between time domain and depth domain inversion to acoustic impedance: 85th Annual International Meeting, SEG, Expanded Abstracts, 3376–3380, doi: 10.1190/segam2015-5897069.1.

Li, J., X. Wu, and Z. Hu, 2022a, Deep learning for simultaneous seismic image superresolution and denoising: IEEE Transactions on Geoscience and Remote Sensing, **60**, 1–11, doi: 10.1109/TGRS.2021.3057857.

Li, Y., Y. Su, M. Guo, X. Han, J. Liu, H. D. Vishwasrao, X. Li, R. Christensen, T. Sengupta, M. W. Moyle, and I. Rey-Suarez, 2022b, Incorporating the image formation process into deep learning improves network performance: Nature Methods, **19**, 1427–1437, doi: 10.1038/s41592-022-01652-7.

Liu, C., and H. Fu, 2021, Enhance 3D seismic images resolution by deconvolving point spread function: First International Meeting for Applied Geoscience & Energy, SEG, Expanded Abstracts, 2759–2763, doi: 10.1190/segam2021-3594599.1.

Liu, C., M. Sun, N. Dai, W. Wu, Y. Wei, M. Guo, and H. Fu, 2022, Deep learning-based point-spread function deconvolution for migration image deblurring: Geophysics, **87**, no. 4, S249–S265, doi: 10.1190/geo2020-0904.1.

Liu, W., Q. Cheng, L. Liu, Y. Wang, and J. Zhang, 2020a, Accelerating high-resolution seismic imaging by using deep learning: Applied Sciences, **10**, 2502, doi: 10.3390/app10072502.

Liu, X., Y. Chen, and L. Huang, 2020b, Least-squares reverse-time migration with a machine-learning-based denoising preconditioner: 90th Annual International Meeting, SEG, Expanded Abstracts, 2993–2997, doi: 10.1190/segam2020-3427177.1.

Liu, Z., Y. Chen, and G. Schuster, 2020c, Deep convolutional neural network and sparse least-squares migration: Geophysics, **85**, no. 4, WA241–WA253, doi: 10.1190/geo2019-0412.1.

Lønning, K., P. Putzky, J.-J. Sonke, L. Reneman, M. W. A. Caan, and M. Welling, 2019, Recurrent inference machines for reconstructing heterogeneous MRI data: Medical Image Analysis, **53**, 64–78, doi: 10.1016/j.media.2019.01.005.

Lu, Y., H. Sun, X. Wang, Q. Liu, and H. Zhang, 2020, Improving the image quality of elastic reverse-time migration in the dip-angle domain using deep learning: Geophysics, **85**, no. 5, S269–S283, doi: 10.1190/geo2019-0250.1.

Lund, H., C. Purcell, and L. Casasanta, 2022, Analytical point spread functions for image domain least-squares Kirchhoff inversion: 84th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.202210110.

Morningstar, W. R., Y. D. Hezaveh, L. P. Levasseur, R. D. Blandford, P. J. Marshall, P. Putzky, and R. H. Wechsler, 2018, Analyzing interferometric observations of strong gravitational lenses with recurrent and convolutional neural networks: arXiv preprint, doi: 10.48550/arXiv.1808.00011.

Morningstar, W. R., L. P. Levasseur, Y. D. Hezaveh, R. Blandford, P. Marshall, P. Putzky, T. D. Rueter, R. Wechsler, and M. Welling, 2019, Data-driven reconstruction of gravitationally lensed galaxies using recurrent inference machines: The Astrophysical Journal, **883**, 14, doi: 10.3847/1538-4357/ab35d7.

Nemeth, T., G. T. Schuster, and P. A. Valasek, 1999, Least-squares migration of incomplete reflection data: Geophysics, **64**, 208–221, doi: 10.1190/1.1444517.

Putzky, P., D. Karkalousos, J. Teuwen, N. Miriakov, B. Bakker, M. Caan, and M. Welling, 2019, i-RIM applied to the fastMRI challenge: arXiv preprint, doi: 10.48550/arXiv.1910.08952.

Putzky, P., and M. Welling, 2017, Recurrent inference machines for solving inverse problems: arXiv preprint, doi: 10.48550/arXiv.1706.04008.

Putzky, P., and M. Welling, 2019, Invert to learn to invert: Advances in Neural Information Processing Systems.

Ravasi, M., and I. Vasconcelos, 2020, PyLops — A linear-operator Python library for scalable algebra and optimization: SoftwareX, **11**, 100361, doi: 10.1016/j.softx.2019.100361.

Ravasi, M., I. Vasconcelos, A. Curtis, and A. Kritski, 2015, Multi-dimensional free-surface multiple elimination and source deblending of Volve OBC data: 77th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.201413355.
Ravasi, M., I. Vasconcelos, A. Kritski, A. Curtis, C. A. D. C. Filho, and G. A. Meles, 2016, Target-oriented Marchenko imaging of a North Sea field: Geophysical Journal International, **205**, 99–104, doi: 10.1093/gji/ggv528.
Rickett, J. E., 2003, Illumination-based normalization for wave-equation depth migration: Geophysics, **68**, 1371–1379, doi: 10.1190/1.1598130.
Ronneberger, O., P. Fischer, and T. Brox, 2015, U-Net: Convolutional networks for biomedical image segmentation: 18th International Conference on Medical Image Computing and Computer-Assisted Intervention.
Sabidussi, E., S. Klein, M. Caan, S. Bazrafkan, A. den Dekker, J. Sijbers, W. Niessen, and D. Poot, 2021, Recurrent inference machines as inverse problem solvers for MR relaxometry: Medical Image Analysis, **74**, 102220, doi: 10.1016/j.media.2021.102220.
Sjoeberg, T. A., L.-J. Gelius, and I. Lecomte, 2003, 2-D deconvolution of seismic image blur: 73rd Annual International Meeting, SEG, Expanded Abstracts, 1055–1058, doi: 10.1190/1.1817453.
Tang, Y., 2009, Target-oriented wave-equation least-squares migration/inversion with phase-encoded Hessian: Geophysics, **74**, no. 6, WCA95–WCA107, doi: 10.1190/1.3204768.
Tarantola, A., 1984, Linearized inversion of seismic reflection data: Geophysical Prospecting, **32**, 998–1015, doi: 10.1111/j.1365-2478.1984.tb00751.x.
Tarantola, A., 1987, Inverse problem theory: Methods for data fitting and model parameter estimation: Elsevier Science Publication Company, Inc.
Torres, K., and M. Sacchi, 2022, Least-squares reverse time migration via deep learning based updating operators: Geophysics, **87**, no. 6, S315–S333, doi: 10.1190/geo2021-0491.1.
Toxopeus, G., J. Thorbecke, K. Wapenaar, S. Petersen, E. Slob, and J. Fokkema, 2008, Simulating migrated and inverted seismic data by filtering a geologic model: Geophysics, **73**, no. 2, T1–T10, doi: 10.1190/1.2827875.
Valenciano, A., S. Lu, N. Chemingui, and J. Yang, 2015, High resolution imaging by wave equation reflectivity inversion: 77th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.201412903.
Valenciano, A. A., B. Biondi, and A. Guitton, 2006, Target-oriented wave-equation inversion: Geophysics, **71**, no. 4, A35–A38, doi: 10.1190/1.2213359.
Valenciano, A. A., B. L. Biondi, and R. G. Clapp, 2009, Imaging by target-oriented wave-equation inversion: Geophysics, **74**, no. 6, WCA109–WCA120, doi: 10.1190/1.3250267.
Vasconcelos, I., H. Peng, M. Ravasi, and D. Kuijpers, 2022, Deep learning in seismic inverse problems with recurrent inference machines: 84th Annual International Conference and Exhibition, EAGE, Extended Abstracts, doi: 10.3997/2214-4609.202211146.
Wang, E., and J. Nealon, 2019, Applying machine learning to 3D seismic image denoising and enhancement: Interpretation, **7**, no. 3, SE131–SE139, doi: 10.1190/INT-2018-0224.1.
Wang, Y., and W. Lu, 2016, Discontinuity enhancement based on time-variant seismic image deblurring: Journal of Applied Geophysics, **135**, 155–162, doi: 10.1016/j.jappgeo.2016.10.008.
Yu, J., J. Hu, G. T. Schuster, and R. Estill, 2006, Prestack migration deconvolution: Geophysics, **71**, no. 2, S53–S62, doi: 10.1190/1.2187783.
Yu, S., and J. Ma, 2021, Deep learning for geophysics: Current and future trends: Review of Geophysics, **59**, e2021RG000742, doi: 10.1029/2021RG000742.
Zhang, H., and W. Wang, 2022, Imaging domain seismic denoising based on conditional generative adversarial networks (CGANs): Energies, **15**, 6569, doi: 10.3390/en15186569.
Zhang, W., J. Gao, Y. Cheng, Z. Li, X. Jiang, and J. Zhu, 2022, Deep-learning for accelerating prestack correlative least-squares reverse time migration: Journal of Applied Geophysics, **200**, 104645, doi: 10.1016/j.jappgeo.2022.104645.

Biographies and photographs of the authors are not available.