GEOPHYSICS, VOL. 86, NO. 5 (SEPTEMBER-OCTOBER 2021); P. R735–R745, 18 FIGS.
10.1190/GEO2020-0564.1

Deep learning for multidimensional seismic impedance inversion

Xinming Wu[1], Shangsheng Yan[1], Zhengfa Bi[1], Sibo Zhang[2], and Hongjie Si[2]

ABSTRACT

Deep-learning (DL) methods have shown promising performance in predicting acoustic impedance from seismic data that is typically considered as an ill-posed problem for traditional inversion schemes. Most of DL methods are based on a 1D neural network that is straightforward to implement, but they often yield unreasonable lateral discontinuities while predicting a multidimensional impedance model trace by trace. We have developed an improvement over the 1D network by replacing it with a 2D convolutional neural network (CNN) and incorporating the constraints of an initial impedance model. The initial model is fed to the network to provide low-frequency trend control, which is helpful for 1D and 2D CNNs to yield stable impedance predictions. Our 2D CNN architecture is quite simple; however, due to the lack of 2D impedance labels, training it is not straightforward. To prepare a 2D training data set, we first define a random path that passes through multiple well logs. We then follow the path to extract a 2D seismic profile and an initial impedance profile that together form an input to the 2D CNN. The set of well logs (traversed by the path) serves as a partially labeled target. We train the 2D CNN with weak supervision by using an adaptive loss in which the output 2D impedance model is adaptively evaluated at the well logs only in the partially labeled target. Because the training data sets are randomly extracted in all directions in a 3D survey, the trained 2D CNN can predict a consistent 3D impedance model section by section in either the inline or crossline directions. Synthetic and field examples indicate that our 2D CNN is more robust to noise, recovers thin layers better, and yields a laterally more consistent impedance model than a 1D CNN with the same network architecture and the same training logs.

INTRODUCTION

Estimating impedance from seismic data is important for reservoir characterization because the impedance shows high-resolution rock properties that may indicate lithology, porosity, pore fill, and other factors (Latimer et al., 2000). Numerous seismic inversion methods including sparse-spike-based methods (e.g., Oldenburg et al., 1983; Velis, 2008; Zhang and Castagna, 2011; Gholami and Sacchi, 2012; Yuan et al., 2015; Wang et al., 2016; Sui and Ma, 2019), model-based methods (e.g., Smith and Gidlow, 1987; Fu, 2004; Veeken and Da Silva, 2004; Wu, 2017b), and geostatistical or stochastic methods (e.g., Sen and Stoffa, 1991; Mallick, 1995; Ma, 2002; Buland and Omre, 2003; González et al., 2008; Bosch et al., 2010) have been proposed to estimate acoustic or elastic impedance from post- or prestack seismic data.

The basic idea of these methods is to estimate impedance by matching simulated synthetic seismic data with recorded real seismic data. The synthetic data are typically simulated by convolving a seismic wavelet with reflectivities that can be converted from the impedance. In practice, the simple convolutional model, however, may not be suitable to properly describe the recorded real seismic data. This indicates that optimally fitting the synthetic and real seismic data is not necessary to yield an optimal impedance estimation. Potential noise, acquisition errors, and processing errors in the real seismic data pose even more uncertainties to the inversion results. In addition, simulating synthetic seismic data in these methods requires a wavelet which, however, is often unavailable and is hard to estimate in practice. Moreover, assuming that a wavelet is available, more than one impedance solution, when converted to reflectivites and convolved with the wavelet, would match the seismic data. Therefore, in addition to matching the seismic data, most of these methods often carefully introduce various regularization terms, reasonable initial impedance models, and prior knowledge

Manuscript received by the Editor 3 August 2020; revised manuscript received 10 June 2021; published ahead of production 19 July 2021; published online 7 September 2021.
[1]University of Science and Technology of China, School of Earth and Space Sciences, Hefei 230026, China. E-mail: xinmwu@ustc.edu.cn (corresponding author); ssyan33@mail.ustc.edu.cn; zfbi@mail.ustc.edu.cn.
[2]Huawei Technologies Co Ltd, Huawei Cloud EI Product Department, Xi'an 710077, China. E-mail: zhangsibo1@huawei.com; sihongjie@huawei.com.

(e.g., well logs, spatial lithology distribution, and statistic rock properties) as constraints in solving the ill-posed impedance inversion problem. These constraints typically play key roles in these methods to obtain a reasonable impedance solution.

The idea of using machine learning or deep learning (DL) to predict impedance from seismic data is not new. Hampson et al. (2001) show that a simple neural network can work well for impedance prediction by reasonably choosing multiple attributes (derived from well-log and seismic data) as inputs. Recently, more DL methods (e.g., Alfarraj and AlRegib, 2019; Biswas et al., 2019; Das et al., 2019; Mustafa et al., 2019; Puzyrev et al., 2019; Wang et al., 2019; Zheng et al., 2019) have been introduced for impedance estimation when well logs are available for training deep neural networks. DL methods are straightforward to implement without the need of carefully designing or solving complicated optimization systems as in conventional inversion methods. In addition, compared to conventional methods, data-driven DL methods are promising to provide a better estimation of the complex and nonlinear mapping from seismic data to impedance without the need of assuming an approximate forward model to simulate seismic from impedance. However, most of the proposed DL methods are based on 1D neural networks because the well logs, used as training labels, are 1D sequences. The 1D networks may be sensitive to noise in the seismic data and often fail to maintain lateral coherence while predicting a multidimensional impedance model trace by trace. When training logs are limited while the rock properties vary rapidly, such 1D networks especially fail to yield stable or accurate impedance estimations.

We propose to improve DL methods for impedance estimation in two aspects. First, we use an initial impedance model together with seismic data as a mixed input to the DL network, which is similar to using an initial impedance model as a starting model or model-domain regularization in conventional inversion methods. Such an initial impedance model provides a low-frequency trend control to the network to make a stable prediction. Second, we propose a simple but effective 2D convolutional neural network (CNN) to predict more accurate multidimensional impedance models. Preparing 2D training data sets is not as straightforward as in one dimension because a full 2D impedance label is often missing. To prepare a pair of training data sets, we first define a random path that passes multiple well-log positions in an original 3D survey. We then follow this path to extract an input data set (a 2D seismic profile and an initial impedance profile) and a 2D impedance target that is partially labeled by the set of the impedance logs on the path. In this way, we are able to obtain many 2D training data sets defined by numerous random paths that pass through the 3D survey in various directions. With these specially defined 2D training data sets, we train the 2D CNN with weak supervision by using an adaptive loss in which the output 2D impedance profile is adaptively fitted to a partially labeled target. The proposed 2D CNN is robust to noise in the input seismic data and can predict a reasonable 2D impedance model that is structurally consistent with the seismic data. Compared to 1D CNNs, 2D CNN is able to better recover thin layers, which are vertically thin but may laterally extend in a large area.

METHODS

When a seismic data set and well-log data (multiple velocity and density logs) are available in the same survey, we can use them both to train a supervised machine-learning method for predicting an impedance model. We start discussing a simple 1D CNN for predicting an impedance model trace by trace from the seismic data without or with the constraints of an initial impedance model. We further propose a 2D CNN for the impedance prediction and train it with an adaptive loss function. Using the same network architecture as the 1D CNN, the 2D CNN is more robust to noise and can yield more accurate and laterally more consistent impedance models than the corresponding 1D CNN.

Data set

To demonstrate the 1D and 2D CNN methods, we use a subset (Figure 1) of the original 3D synthetic SEG advanced modeling (SEAM) Phase I data set. This data set contains complex salt bodies, and the impedance values appear as sharp variations across the salt boundaries from the surrounding layers into the bodies. In addition, rich lateral and vertical impedance variations are also observed within the surrounding layers. The subset in Figure 1 contains 600 (vertical) ×501 (inline) ×502 (crossline) samples, which cover an area of approximately 15×15 km^2. The impedance model (computed from the density and velocity model) in Figure 1a has been resampled to be consistent with the associated depth-migrated seismic image (Figure 1b).

From the impedance model, we choose 50 vertical traces as impedance logs displayed with the seismic image in Figure 1b. The positions and lengths of the logs are randomly chosen, where the minimum distance between two logs is larger than 40 samples (1.2 km) and the length of each log is larger than 300 samples (3 km). We use 40 of these extracted logs to train the CNN, whereas 10 logs are used for validation. The true impedance model (Figure 1a) is heavily smoothed by using a 3D isotropic Gaussian filter (half-width $\sigma = 20$ samples) to obtain an initial impedance model (Figure 1c) used as constraints in the CNN methods.

1D CNN

By considering the measured 1D well-log impedance sequences as ground truth, it is straightforward to implement the impedance estimation by using a 1D CNN whose input and output are a seismic trace and the corresponding impedance sequence, respectively.

Network architecture

Figure 2 shows a quite simple 1D CNN that we use for 1D impedance prediction. This CNN consists of a regular convolutional layer, four blocks, and a final output layer. The first regular layer contains 16 1D convolutional filters (each with seven coefficients), which are applied to the input to obtain 16 1D feature vectors. These feature vectors are passed to a rectified linear unit (ReLU) activation and then fed to the next four sequentially linked blocks. Following the four blocks, the final output layer is a convolutional layer with a filter size of one, which yields an impedance sequence. As shown in Figure 2, the backbone of our network is the sequential stack of the four blocks. Each block contains a regular convolutional layer and a ResBlock (He et al., 2016). As shown at the bottom of Figure 2, the ResBlock contains two convolutional layers (each with 16 features or filters) and a skip connection over the two layers. Such ResBlocks are widely used in constructing deep CNN because the skip connections in the ResBlocks can effectively eliminate the problem of the gradient vanishing in the training phase. In addition, in constructing a neural network, we generally do not know the

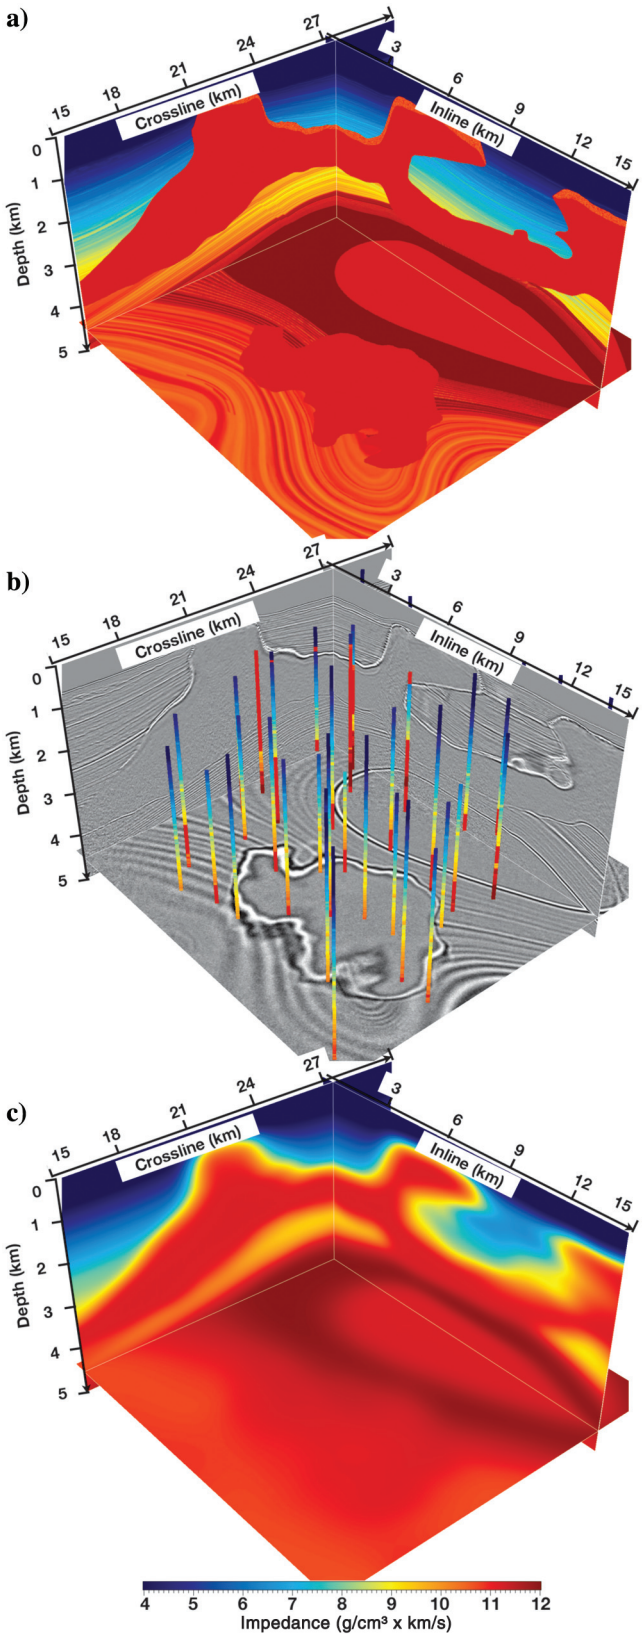

Figure 1. A subset of the synthetic 3D SEAM Phase I data set used to demonstrate 1D and 2D CNNs for impedance prediction.

optimal number of layers, which may depend on the complexity of the data set. Instead of trying to estimate the number of layers, we can construct a deep network with skip connections to bypass the layers that do not add value in overall accuracy. In this way, skip connections make our neural networks dynamic, so that it can automatically tune the optimal number of layers during training.

Training

We train this 1D CNN with two schemes according to the data fed into the network. In the first scheme, we input only a seismic trace to the network; therefore, a training data set pair consists of an input seismic trace and the corresponding target of impedance sequence. We have a total of 40 training logs (Figure 1b) with different lengths (300–600 samples). In training the network, each pair of training data sets is with a fixed length of 300 samples and is randomly extracted from the 40 logs and the corresponding seismic traces.

In the second scheme, the data input to the network includes a seismic trace and a corresponding initial impedance sequence, which is extracted from a heavily smoothed impedance model (Figure 1c). This initial impedance sequence provides a low-frequency impedance trend to constrain the network for predicting a more accurate impedance sequence with more details. In this second scheme, we use the same method as the first scheme to randomly extract training data sets from the same 40 training logs.

In both schemes, estimating an impedance sequence from a seismic trace without or with an initial impedance sequence is a classic regression problem. We therefore train the network by using the following commonly used loss function:

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N}(y[i] - y^p[i])^2, \tag{1}$$

which measures the mean squared error (MSE) between a target impedance log ($y[i]$) and a predicted impedance sequence ($y^p[i]$).

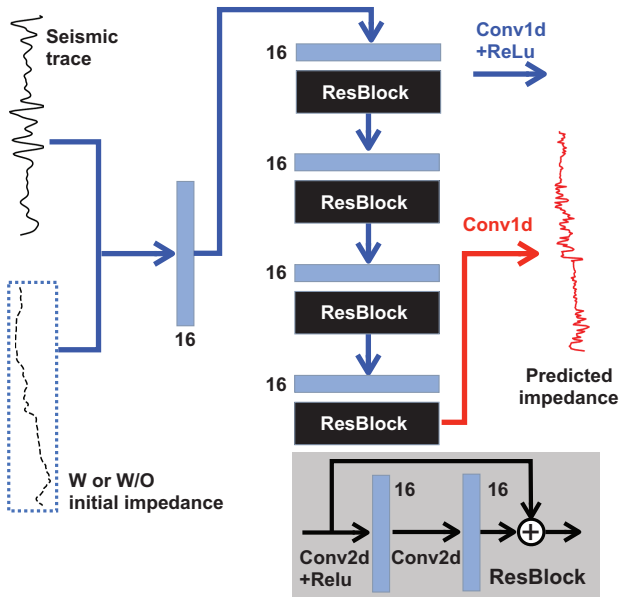

Figure 2. A simple 1D CNN for estimating an impedance sequence from a seismic trace with or without an initial impedance sequence.

We use the Adam method (Kingma and Ba, 2014) to optimize the network parameters by iteratively minimizing the loss function with multiple epochs. In the training process, we use an adaptively decreasing learning rate that starts with 0.001. The training and validation losses, starting from approximately five, smoothly converge to nearly 0.03 and 0.02, respectively, in the first and second schemes after 300 epochs. This indicates that the simple 1D CNN (Figure 2) with 10,001 trainable parameters is good enough to fit the training data sets.

Testing results

With the trained 1D CNN models using the two schemes, we predict the 3D impedance models in Figure 3a and 3b trace by trace from the 3D seismic data (Figure 1b) without or with an initial impedance model (Figure 1c), respectively. With the input of only the seismic data, the 1D CNN yielded a noisy impedance model (Figure 3a) with a lot of unreasonable discontinuities, which is much different from the true impedance model (Figure 3d). Figure 4 shows a comparison of the predicted impedance sequences (the magenta curves) and the ground truth (the black curves) at four validation logs (denoted by the red circles in Figure 5a), which are not included in the training logs. We observe that the impedance sequences (the magenta curves), predicted with only seismic traces, oscillate with depth and are far away from the true impedance sequences (the black curves) as denoted in Figure 4.

Therefore, the trained 1D CNN in the first scheme fails to build a sufficient mapping between the impedance sequences and seismic traces for the whole data set, although it well fits the training data sets because the training loss already converges to 0.03. This failure does not illustrate the failure of the constructed CNN; rather, it indicates that the 40 training logs are not representative

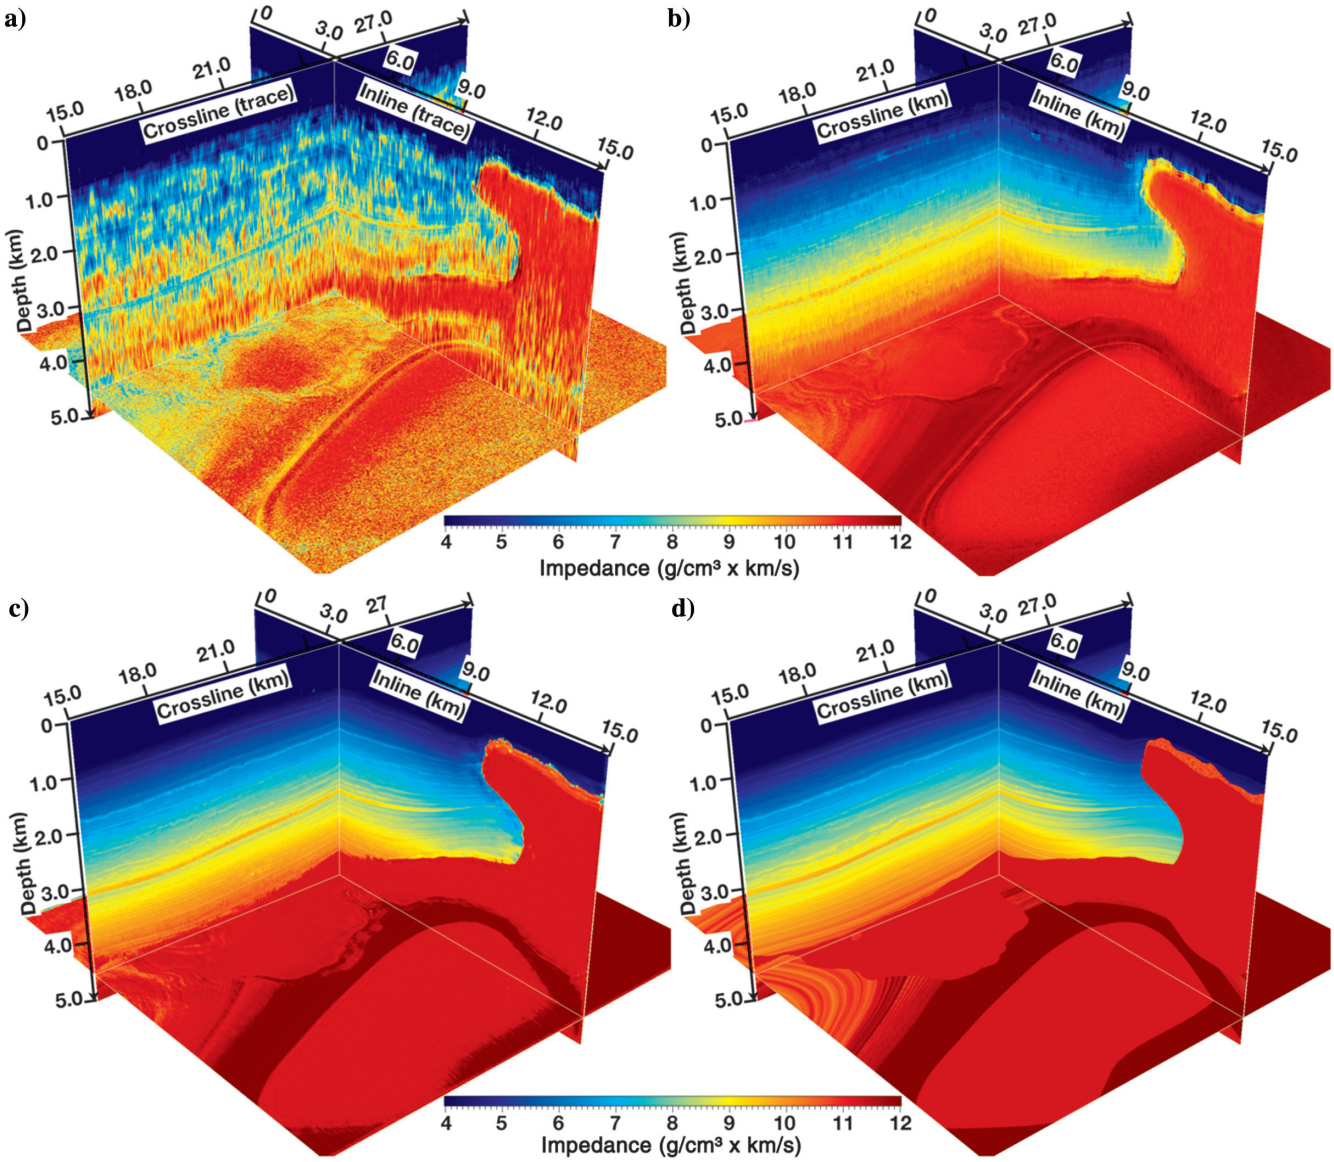

Figure 3. Compared to 3D impedance models predicted by the 1D CNNs (a) without and (b) with an initial impedance model, the one (c) predicted by the 2D CNN matches the ground truth impedance model (d) better.

enough for the CNN to learn the complex mapping between the seismic and impedance for the whole data set. In addition, spatially varying noise in the seismic data may also prevent the trained CNN from making a stable impedance prediction. Moreover, the 1D CNN predicts a 3D impedance model trace by trace, which inherently ignores lateral constraints and thus fails to maintain lateral coherency.

In practice with field data, we could not expect more training logs to improve the performance of the 1D CNN because the number of wells is always limited. Inspired by the idea of using a reasonable initial model as a regularization term or a starting model in conventional inversion methods, we feed an initial impedance sequence (together with a seismic trace) to the 1D CNN to provide a low-frequency trend control in the second scheme. By using initial impedance sequences (extracted from a heavily smoothed impedance model in Figure 1c), the trained 1D CNN can yield an more accurate impedance model (Figure 3b) than the one using only seismic traces (Figure 3a). However, we still observe some noisy features in the predicted impedance model (Figure 3b), some oscillations and mismatches as shown by the blue curves at the validation logs (Figure 4).

In summary, due to the spatial variation of rock properties and potential noise in seismic data, a 1D CNN is not able to predict an accurate impedance model from only the seismic traces when the number of training logs is limited. The performance of the same 1D CNN can be significantly improved by using a combined two-channel input of a seismic trace and a reasonable initial impedance sequence, which introduces a low-frequency trend control to stabilize the impedance prediction. However, using an initial impedance model does not essentially eliminate the main limitation of a 1D CNN that makes trace-by-trace predictions, in which the lateral structure features in the seismic data are totally ignored. Taking lateral continuity into account is helpful to laterally estimate more consistent impedance models with better recovered thin layers and

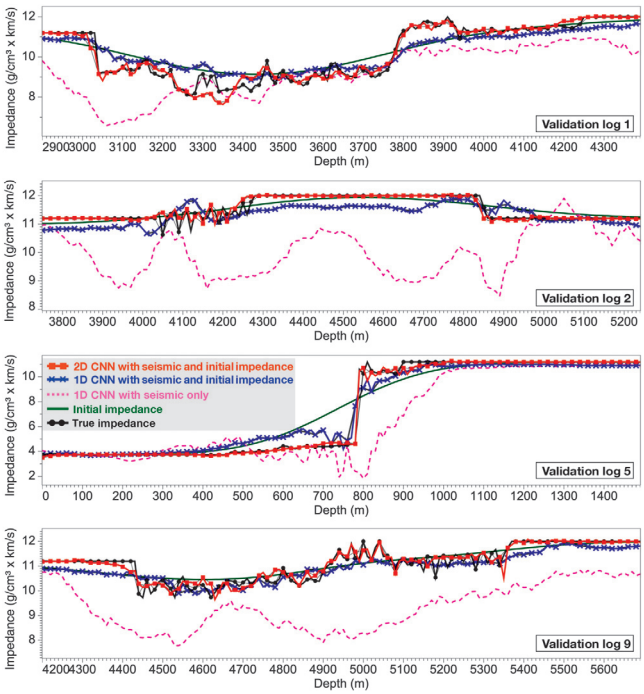

Figure 4. A comparison of the predicted impedance results at validation logs #1, #2, #5, and #9 (denoted by the red circles in Figure 5a), where the 2D CNN method shows the best performance. The impedance values (the red curves) predicted by the 2D CNN match well with the ground truth (the black curves) even at thin layers that rapidly vary with depth.

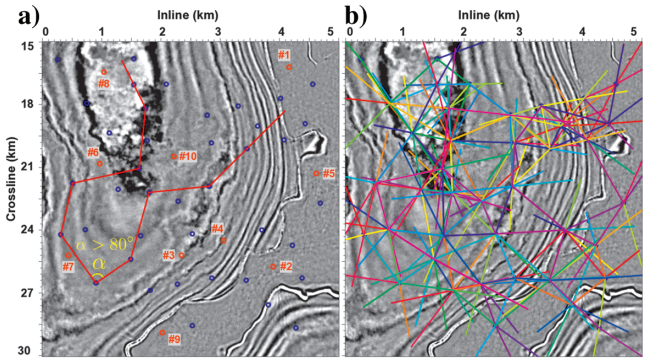

Figure 5. The circles in (a) denote the locations of 40 training logs (blue) and 10 validation logs (red) on a depth slice of the seismic volume in Figure 1a. The red curve in (a) denotes a path that passes through 10 randomly chosen log positions. The colorful curves in (b) represent 100 random paths each of which passes at least through five log positions. There are numerous paths like these in (b) along which we extract the 2D training data sets.

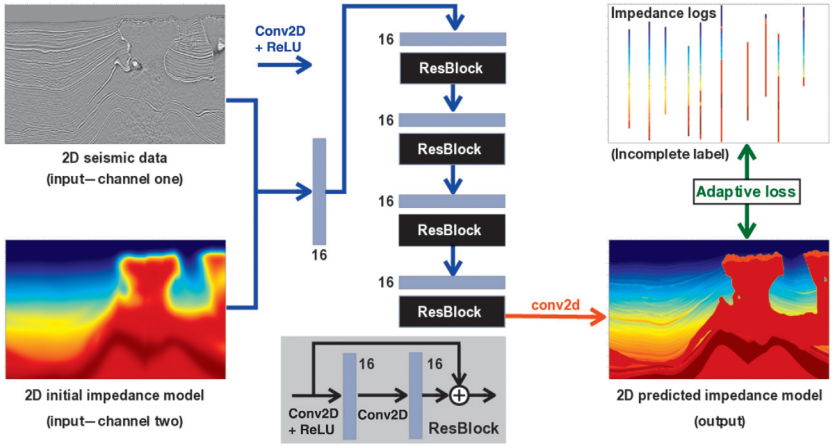

Figure 6. The proposed 2D CNN shares the same architecture of the 1D CNN (Figure 2). The input of the network is a combination of 2D seismic and initial impedance profiles. The output is a predicted 2D full impedance profile, which is fitted to a partially labeled target by using an adaptive loss during the training.

robustness to noise. The thin layers are vertically subtle and can therefore be easily ignored in a 1D trace-by-trace process. However, these thin layers may laterally extend in a large area and can be captured by a multidimensional process that takes into account the constraints of lateral structures in the seismic data.

2D CNN

We propose a 2D CNN to further improve the impedance prediction with the same number of training logs. The architecture that we design for the 2D CNN (Figure 6) is exactly the same as the 1D CNN (Figure 2). We only change the 1D convolutional filters with 3×3 2D filters, which help to extract spatial features in the input seismic data. Accordingly, the input is changed to a combination of 2D seismic initial impedance profiles (the left of Figure 6). The output of the 2D CNN is now a full 2D impedance profile, instead of an impedance sequence in the 1D CNN.

Training data

The main challenge of applying a 2D CNN for impedance prediction is to properly prepare 2D training data sets, in which the measured impedance values (the ground truth) are only 1D sequences at irregularly scattered log positions. To solve this problem, we consider the 2D CNN-based impedance prediction as a weakly or partially supervised learning problem, in which the 2D target impedance profile is partially labeled at only known well-log positions. We then randomly extract 2D training data sets from the original 3D data set and make sure that each extracted 2D data set passes through multiple well logs.

To extract such a 2D training data set, we first define a random path that passes through at least five well logs (the red curve in Figure 5a). We define such a path by sequentially linking a series of randomly chosen well-log positions (the blue circles) with linear segments. There are three simple but helpful considerations to define such a path: (1) In randomly choosing a series of log positions, we need to make sure that any two sequentially adjacent log positions are not the same, (2) the angle between two linked line segments must be greater than 80°, and (3) the path should not start or end with a log position to avoid boundary effects in the CNN. In this way, we are able to randomly pick numerous paths, and we show 100 of them in Figure 5b. These paths link all of the training log positions and cross the whole survey in all directions.

Following each random path, we then extract a pair of 2D training data set (seismic and initial impedance profiles) and a corresponding 2D labeled impedance profile mostly equal to zero except where the path passes through the logs. Figure 7 shows a training data set pair extracted by following the path shown in Figure 5a. We are able to extract numerous pairs of such training data sets that are in two dimensions and can sufficiently capture spatially varying patterns in the 3D data set due to the randomness of the paths.

Adaptive loss

In our 2D CNN method shown in Figure 6, we expect the output to be a full 2D impedance profile although the label profile is incomplete with known values at log positions only. To deal with the problem of the mismatch between the output and the label data

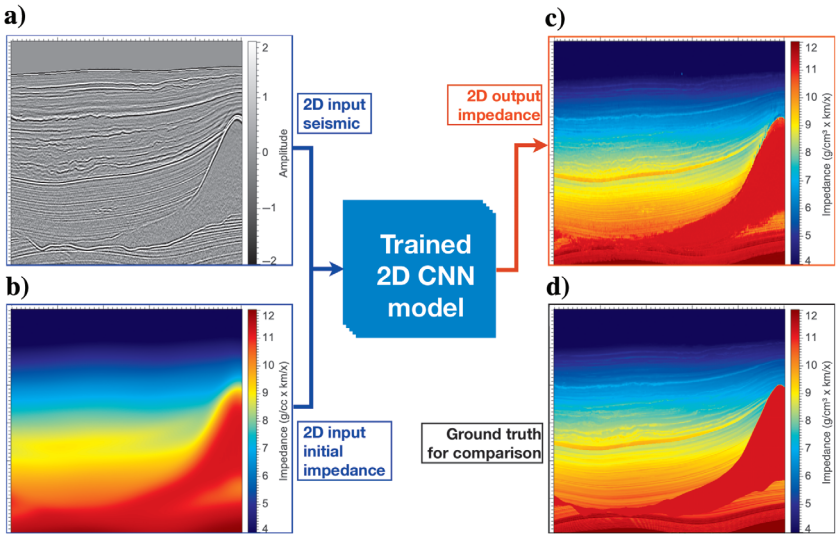

Figure 8. The inference step of the 2D CNN method: With an input pair of (a) seismic and (b) initial impedance profiles, our trained two dimensions directly predicts in (c) a full 2D impedance profile, which is consistent with the ground truth in (d).

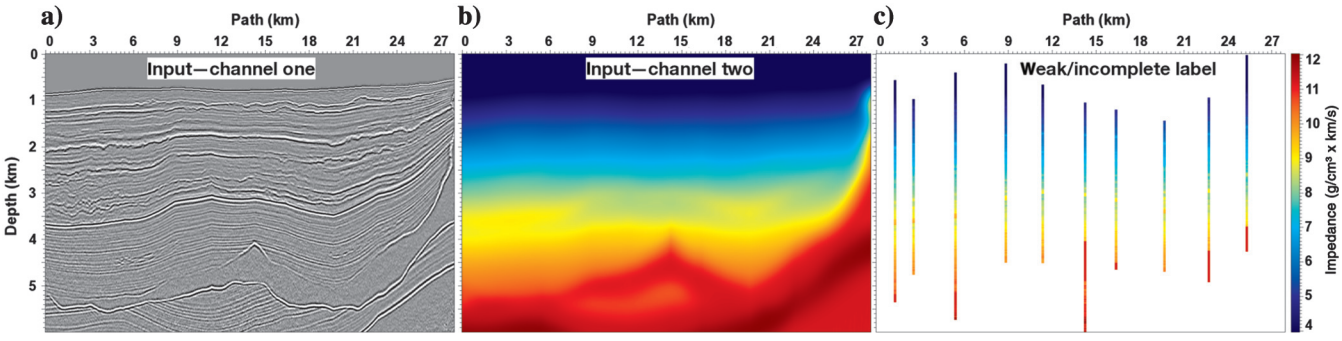

Figure 7. (a) A 2D seismic amplitude profile, (b) a 2D smoothed impedance profile, and (c) 10 impedance logs are extracted on the red path shown in Figure 5a. The 2D seismic and smoothed impedance profiles form a two-channel input to our 2D CNN, whereas the set of 10 impedance logs is the corresponding target (a weak or partial label).

in training the CNN, we define an adaptive MSE loss function as follows:

$$\mathcal{L}=\frac{1}{\sum_{i=1}^{N}\sum_{j=1}^{M}w[i,j]}\sum_{i=1}^{N}\sum_{j=1}^{M}w[i,j]\cdot(y[i,j]-y^{p}[i,j])^{2}, \quad (2)$$

where $y[i,j]$ represents a 2D target impedance profile (label data) and $y^{p}[i,j]$ denotes a predicted impedance profile. The term $w[i,j]$ represents a 2D binary mask, which is equal to one at the well-log positions and is zero elsewhere. In training the 2D CNN with this loss, we adaptively fit the 2D output with the incompletely labeled target impedance profile at only the positions with well-log values. However, the 2D convolutional filters in the 2D CNN still continuously process the input 2D data and extract 2D feature maps to predict a full 2D impedance profile with spatially consistent values. In the inference step shown in Figure 8, our trained 2D CNN model directly predicts a full 2D impedance profile (Figure 8c) from an input pair of seismic (Figure 8a) and initial impedance (Figure 8b) profiles. The predicted 2D impedance profile (Figure 8c) is highly consistent with the ground truth in Figure 8d.

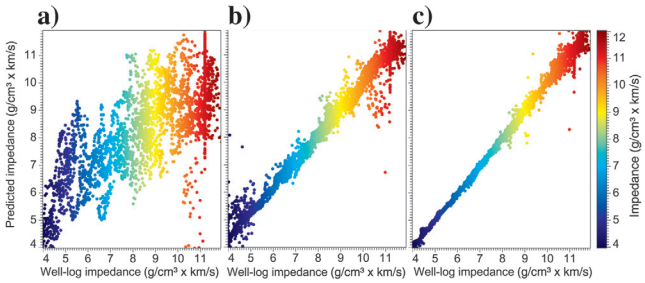

Figure 9. Crossplots at the 10 validation logs: the ground truth impedance against the predicted impedance by using (a) the 1D CNN with seismic data only, (b) the 1D CNN with the seismic and initial impedance, and (c) the 2D CNN with the seismic and initial impedance.

Methods / Validation data	MSEs by 1D CNN with seismic only	MSEs by 1D CNN with seismic and initial impedance	MSEs by 2D CNN with seismic and initial impedance
Validation log 1	1.3940	0.1888	**0.0265**
Validation log 2	2.9415	0.3183	**0.0241**
Validation log 3	0.8722	0.0833	**0.0124**
Validation log 4	1.5199	0.0717	**0.0123**
Validation log 5	2.1104	0.1963	**0.0188**
Validation log 6	2.2405	0.0632	**0.0167**
Validation log 7	1.9171	0.0686	**0.0135**
Validation log 8	2.4492	0.0707	**0.0167**
Validation log 9	2.1172	0.2584	**0.0211**
Validation log 10	1.2863	0.0922	**0.0073**
The whole volume	1.7694	0.2122	**0.0928**

Figure 10. The MSEs of the impedance results are predicted at the 10 validation logs and for the whole volume.

Testing results and accuracy metric

Because the size of the input to our 2D CNN is not fixed, we can directly apply the trained 2D CNN to the original 3D data set to predict a 3D impedance model section by section in the inline or crossline directions. Figure 3c shows the predicted 3D impedance model, which displays laterally more consistent layered impedance features than the one computed by the 1D CNN (Figure 3b). Compared to the impedance models (Figure 3a and 3b) estimated by the 1D CNNs, the one predicted with the 2D CNN (Figure 3c) is less affected by the noise in the seismic data and better recovers thin layers. Figure 4 shows more details of the predicted impedance values at four validation logs. We observe that the impedance values (the red curves) predicted by the 2D CNN match well with the ground truth (the black curves) even at the thin layers in which the impedance values vary rapidly with depth. Figure 9 shows the crossplots between the validation well-log

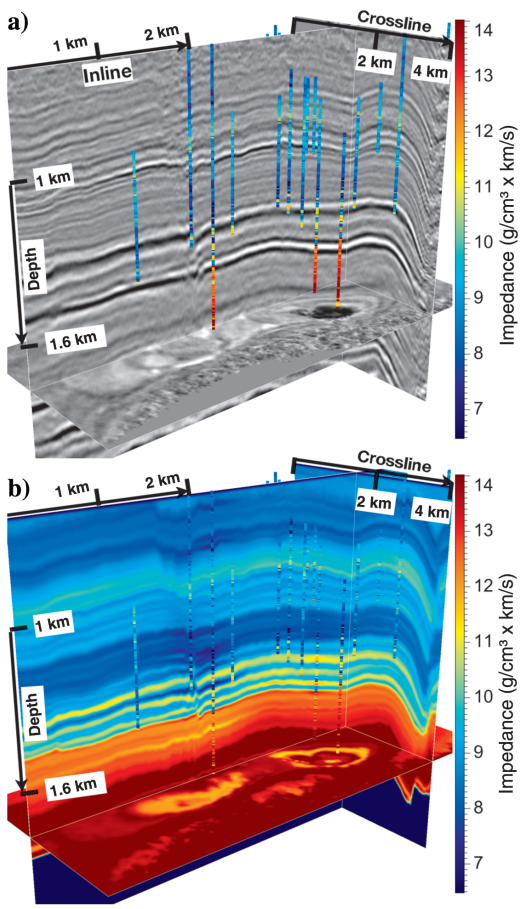

Figure 11. (a) The 3D field seismic volume and 27 impedance logs are extracted from the freely available Teapot Dome data set. (b) An initial impedance model is interpolated from the impedance logs while conforming with the seismic structures. Some smoothing has been applied to the interpolated impedance model.

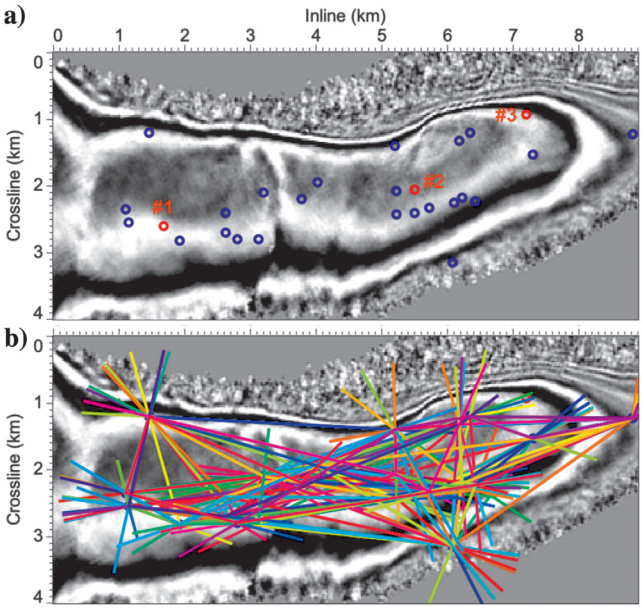

Figure 12. The circles in (a) denote the locations of 27 logs overlain on a depth slice of the seismic volume (Figure 11a), where the blue and red circles represent the logs used for training and validation, respectively. The colorful curves in (b) represent random paths each of which passes through at least five training logs. We can define numerous random paths like these and extract a lot of 2D training data sets along the paths.

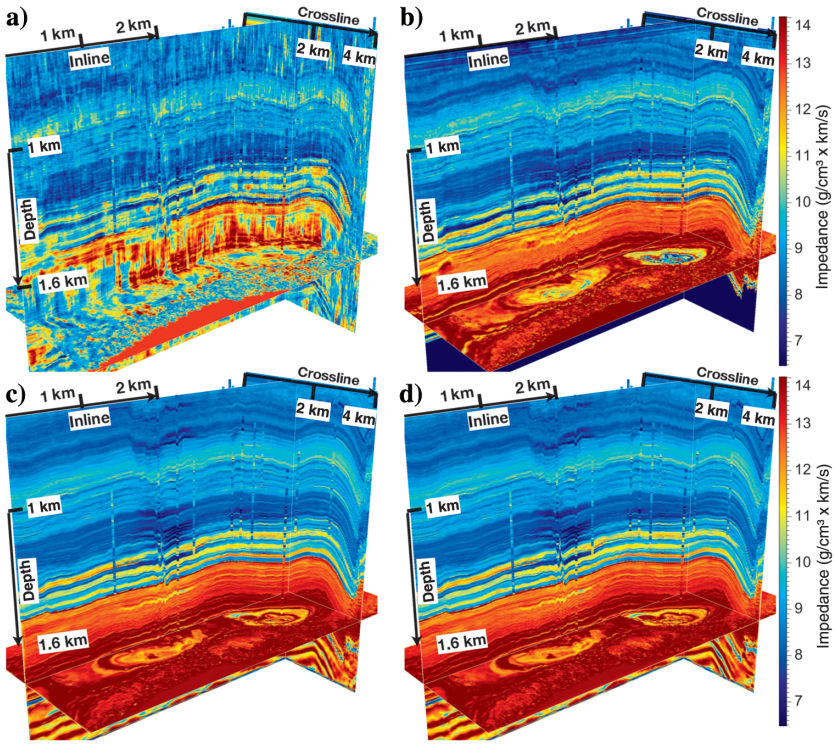

Figure 13. Comparison of the predicted 3D impedance models by using (a) the 1D CNN with only seismic data, (b) the 1D CNN with seismic and interpolated impedance, and (c and d) 2D CNN with seismic and interpolated impedance. With the trained 2D CNN, we are able to predict consistent 3D impedance models section by section in either the (c) crossline or (d) inline directions.

impedance values and the ones predicted at the validation logs. In Figure 9c, we observe that the impedance values predicted by the 2D CNN show the highest correlation with the well-log impedance values.

Figure 10 shows a quantitative comparison of the MSE of the predicted impedance results using the 1D and 2D CNNs. Although trained using the same 40 training logs, the 2D CNN result shows significantly smaller MSE than the results by the 1D CNNs at the 10 validation logs and for the whole 3D impedance volume. Note that the validation logs of #1, #2, #5, and #9 show relatively larger errors for all three methods, and we display the predictions at these logs in Figure 4.

APPLICATION

To further illustrate the effectiveness of the methods, we apply them to the freely available Teapot Dome data set (Anderson, 2009). Figure 11a shows a 3D field seismic volume and 27 impedance logs that are extracted from the Teapot Dome data set. Although hundreds of wells are available with the seismic data, we choose only 27 logs shown in Figure 11a after discarding the wells whose density or velocity logs are missing. These 27 logs have been tied to the seismic data that have already been converted to depth. From these 27 impedance logs, we compute a 3D impedance model (Figure 11b) using a seismic-structure-guided interpolation method (e.g., Hale, 2010; Naeini and Hale, 2015; Karimi et al., 2017; Wu, 2017a). The interpolated impedance model is generally consistent with the measured impedance-log values and laterally matches the seismic structures. It is used as an initial impedance model that, together with the seismic data, are fed into the 1D and 2D CNNs. We expect the initial impedance model to provide a low-frequency control and the CNNs to fill in details and correct interpolation errors in this initial model based on the seismic amplitudes.

Figure 12a shows the spatial positions of the 27 logs, which are overlain on a depth slice of the seismic volume. We use 24 logs (the blue circles) for training, whereas the remaining three logs (the red circles) are used for validation. The 1D training data sets are randomly chosen at the positions of these 24 logs, and the vertical length of each training data set is set to be 72 samples long, which is the length of the shortest log. The 2D training data sets are extracted by following random paths that pass through at least five training logs. Figure 12b shows 100 random paths, which traverse the survey in various directions.

Figure 13a and 13b shows 3D impedance models predicted trace by trace by using the 1D CNN without and with the interpolated initial impedance model, respectively. It is obvious that using an initial impedance model is helpful for the 1D CNN to predict a much more stable result. By applying the trained 2D CNN to the 3D data set section by section in the inline and crossline directions, we obtain two consistent 3D impedance models as shown in Figure 13c

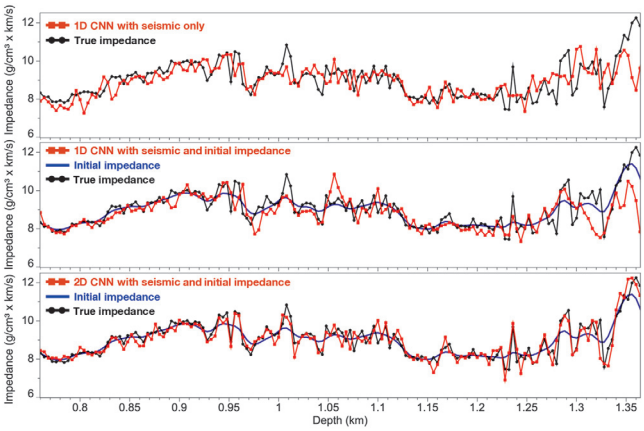

Figure 14. Predicted impedance results at the #1 validation log (highlighted by the red circle in Figure 12a).

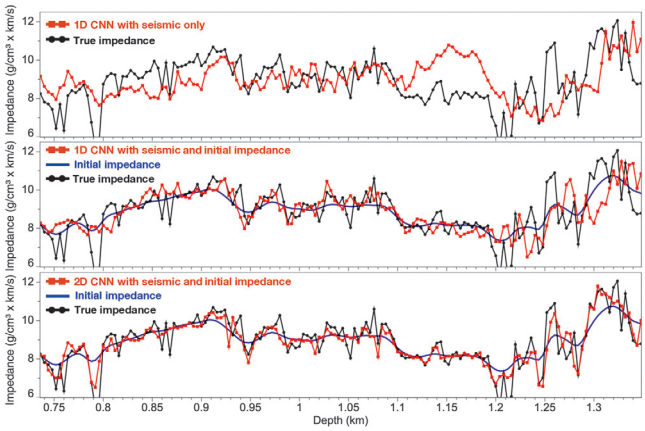

Figure 15. Predicted impedance results at the #2 validation log (highlighted by the red circle in Figure 12a).

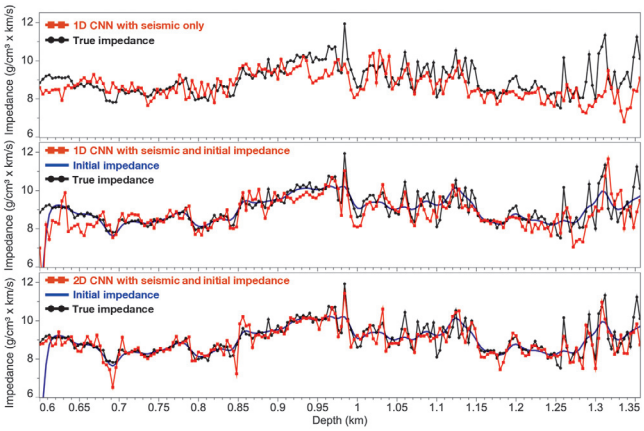

Figure 16. Predicted impedance results at the #3 validation log (highlighted by the red circle in Figure 12a).

Methods / Validation logs	MSEs by 1D CNN with seismic only	MSEs by 1D CNN with seismic and initial impedance	MSEs by 2D CNN with seismic and initial impedance
Validation log 1	0.6051	0.5500	**0.1689**
Validation log 2	0.6997	0.5599	**0.2411**
Validation log 3	1.7195	1.0255	**0.3854**

Figure 18. MSEs of the predicted impedance results at the three validation logs.

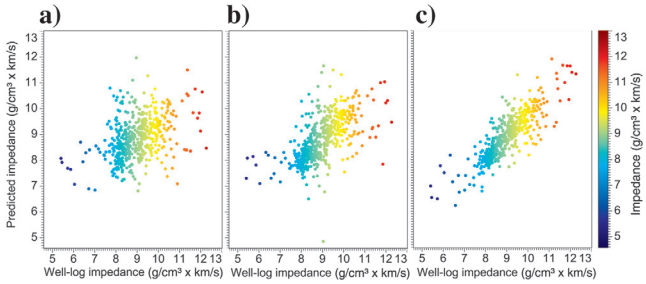

Figure 17. Crossplots at the three validation logs: the well-log impedance against the predicted impedance by using the (a) 1D CNN with only seismic data, (b) the 1D CNN with seismic and initial impedance, and (c) the 2D CNN with seismic and initial impedance.

and 13d. This indicates that the 2D training data sets extracted along the random paths are able to capture the spatially varying structures in the 3D seismic data. Compared to the 1D CNN result in Figure 13b using an initial impedance model, the impedance models (Figure 13c and 13d) by the 2D CNN show laterally more consistent features and clearer thin layers.

Figures 14, 15, and 16 show a comparison of the predicted impedance sequences at the three validations logs. Compared to the results using the 1D CNN without (the magenta curves) and with (the blue curves) the initial impedance model (the green curves), the impedance traces predicted by the 2D CNN (the red curves) visually match better the measured impedance logs (the black curves). This observation is supported by the crossplots (Figure 17) between the validation well-log impedance values versus the predicted ones. We observe that the impedance values predicted by the 2D CNN show the highest correlation with the well-log impedance values in Figure 17c. Figure 18 displays a quantitative comparison of the MSE of the predicted impedances results at the three validation logs. We observe that the 2D CNN predictions also show the highest accuracy (or smallest MSE) at each of the three validation logs.

CONCLUSION

When seismic data and impedance logs are available, a supervised DL method provides a simple and promising way to predict an impedance model from seismic amplitude data without the need for carefully designing and solving large linear or nonlinear systems as is often the case in conventional impedance inversion methods. By considering 1D impedance logs as labels and the corresponding seismic trace as inputs, it is straightforward to train a 1D CNN to

predict an impedance model. However, the performance of such a 1D CNN is limited to the number of available logs. Estimating an impedance model from a seismic data set with noise and rapidly varying structures or rock properties is especially challenging to a 1D CNN trained with a limited number of sparsely measured logs. In these cases, the trained CNN is not sufficient to present a complex and spatially varying relationship between the seismic data and the impedance.

It is necessary to introduce reasonable constraints to improve the simple 1D CNN method, which is similar to solving ill-posed problems in conventional inversion schemes by introducing a good starting model or regularization constraints. One feasible way is to input an initial impedance sequence (together with a seismic trace) to the 1D CNN, which significantly improves the stability and accuracy of the impedance prediction as demonstrated in the synthetic and field examples. A further way to improve CNN-based impedance prediction is to use a 2D CNN whose 2D convolutional filters are helpful for incorporating lateral structure constraints to predict a laterally more consistent impedance model. However, training a 2D CNN is not straightforward because we do not have a known 2D impedance profile to use as a label. We design an adaptive loss to train the 2D CNN by weak supervision from an incomplete impedance label that contains multiple measured impedance logs. By using the same network architecture and training logs as the 1D CNN, the 2D CNN is less sensitive to noise and is able to predict a laterally more consistent impedance model. In addition, the 2D CNN can extract lateral structure features in the seismic data to more accurately recover thinner layers, which are vertically subtle but may laterally extend in a large range.

The architecture of the CNNs used in this paper is quite simple, but it is sufficient for the impedance prediction problem. We tried more sophisticated architectures that show similar performance, but they require greater computational costs. In our experience, more training logs, more prior knowledge constraints, and more proper use of these available data are more significant factors than designing a more complicated network to improve CNN-based impedance prediction.

Some limitations remain in our method. One arises from an assumption that the patterns of impedance logs at the training wells are sufficiently like those in the prediction areas. This assumption indicates that our 1D and 2D CNN models, trained on one seismic survey, may not be generalized well on (or successfully applied to) a different survey. This means that for each survey, we may need to train a new CNN model by using the new training data sets of seismic traces and impedance logs. For the two examples in this paper, we have trained different 1D and 2D CNN models. Although we need to retrain a model for a different survey, it is still acceptable because the training costs are computationally affordable. In our examples, training 1D and 2D CNNs takes approximately minutes and hours, respectively, by using one graphics processing unit. Generalization over different surveys remains a common challenge for most CNN models in different geophysical problems including impedance inversion. We have tried to use more diverse training data sets (real data sets from multiple surveys and a lot of synthetic ones based on physical models) to train a more generalized model. However, we found that the generalization of the trained CNN models for impedance (or other rock properties) prediction is not as good as for seismic structural interpretation tasks (e.g., fault, horizon, and geobody detection). It would be risky to directly apply a trained CNN model to predict impedance from seismic on various surveys. Transfer learning might be a potential way to improve the generalization of a trained model on a new survey, in which the model parameters are adjusted by fitting the new training data sets in the new survey.

Another limitation arises from the training data sets. First, our 1D and 2D CNN methods require multiple well logs for constructing training data sets. The applicability of the CNN methods will be limited when the number of available well logs is limited. Second, accurate seismic-well ties are required in constructing the training data sets for the 1D and 2D CNNs. The upper bound of the prediction accuracy of a trained CNN model depends on the number of training logs and the accuracy of the seismic well ties. In the future, to reduce both limitations of the CNN methods, we may determine a way to incorporate physical constraints into CNNs to build data- and physics-driven models.

ACKNOWLEDGMENTS

The field seismic and well-log data are provided by the Rocky Mountain Oilfield Test Center. This research was supported by the National Science Foundation of China under grant no. 42050104.

DATA AND MATERIALS AVAILABILITY

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

REFERENCES

Alfarraj, M., and G. AlRegib, 2019, Semi-supervised learning for acoustic impedance inversion: 89th Annual International Meeting, SEG, Expanded Abstracts, 2298–2302, doi: 10.1190/segam2019-3215902.1.

Anderson, T., 2009, History of geologic investigations and oil operations at teapot dome: Presented at the Annual Convention, AAPG.

Biswas, R., M. K. Sen, V. Das, and T. Mukerji, 2019, Prestack and poststack inversion using a physics-guided convolutional neural network: Interpretation, **7**, no. 3, SE161–SE174, doi: 10.1190/INT-2018-0236.1.

Bosch, M., T. Mukerji, and E. F. Gonzalez, 2010, Seismic inversion for reservoir properties combining statistical rock physics and geostatistics: A review: Geophysics, **75**, no. 5, 75A165–75A176, doi: 10.1190/1.3478209.

Buland, A., and H. Omre, 2003, Bayesian linearized AVO inversion: Geophysics, **68**, 185–198, doi: 10.1190/1.1543206.

Das, V., A. Pollack, U. Wollner, and T. Mukerji, 2019, Convolutional neural network for seismic impedance inversion: Geophysics, **84**, no. 6, R869–R880, doi: 10.1190/geo2018-0838.1.

Fu, L.-Y., 2004, Joint inversion of seismic data for acoustic impedance: Geophysics, **69**, 994–1004, doi: 10.1190/1.1778242.

Gholami, A., and M. D. Sacchi, 2012, A fast and automatic sparse deconvolution in the presence of outliers: IEEE Transactions on Geoscience and Remote Sensing, **50**, 4105–4116, doi: 10.1109/TGRS.2012.2189777.

González, E. F., T. Mukerji, and G. Mavko, 2008, Seismic inversion combining rock physics and multiple-point geostatistics: Geophysics, **73**, no. 1, R11–R21, doi: 10.1190/1.2803748.

Hale, D., 2010, Image-guided 3D interpolation of borehole data: CWP Report 656.

Hampson, D. P., J. S. Schuelke, and J. A. Quirein, 2001, Use of multiattribute transforms to predict log properties from seismic data: Geophysics, **66**, 220–236, doi: 10.1190/1.1444899.

He, K., X. Zhang, S. Ren, and J. Sun, 2016, Deep residual learning for image recognition: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.

Karimi, P., S. Fomel, and R. Zhang, 2017, Creating detailed subsurface models using predictive image-guided well-log interpolation: Interpretation, **5**, no. 3, T279–T285, doi: 10.1190/INT-2016-0051.1.

Kingma, D. P., and J. Ba, 2014, Adam: A method for stochastic optimization: CoRR, abs/1412.6980.

Latimer, R. B., R. Davidson, and P. Van Riel, 2000, An interpreter's guide to understanding and working with seismic-derived acoustic impedance data: The Leading Edge, **19**, 242–256, doi: 10.1190/1.1438580.

Ma, X.-Q., 2002, Simultaneous inversion of prestack seismic data for rock properties using simulated annealing: Geophysics, **67**, 1877–1885, doi: 10.1190/1.1527087.

Mallick, S., 1995, Model-based inversion of amplitude-variations-with-offset data using a genetic algorithm: Geophysics, **60**, 939–954, doi: 10.1190/1.1443860.

Mustafa, A., M. Alfarraj, and G. AlRegib, 2019, Estimation of acoustic impedance from seismic data using temporal convolutional network: 89th Annual International Meeting, SEG, Expanded Abstracts, 2554–2558, doi: 10.1190/segam2019-3216840.1.

Naeini, E. Z., and D. Hale, 2015, Image- and horizon-guided interpolation: Geophysics, **80**, no. 3, V47–V56, doi: 10.1190/geo2014-0279.1.

Oldenburg, D., T. Scheuer, and S. Levy, 1983, Recovery of the acoustic impedance from reflection seismograms: Geophysics, **48**, 1318–1337, doi: 10.1190/1.1441413.

Puzyrev, V., A. Egorov, A. Pirogova, C. Elders, and C. Otto, 2019, Deep neural networks for 1D impedance inversion: Presented at the Annual Convention, ASEG.

Sen, M. K., and P. L. Stoffa, 1991, Nonlinear one-dimensional seismic waveform inversion using simulated annealing: Geophysics, **56**, 1624–1638, doi: 10.1190/1.1442973.

Smith, G., and P. Gidlow, 1987, Weighted stacking for rock property estimation and detection of gas: Geophysical Prospecting, **35**, 993–1014, doi: 10.1111/j.1365-2478.1987.tb00856.x.

Sui, Y., and J. Ma, 2019, A nonstationary sparse spike deconvolution with anelastic attenuation: Geophysics, **84**, no. 2, R221–R234, doi: 10.1190/geo2017-0846.1.

Veeken, P., and M. Da Silva, 2004, Seismic inversion methods and some of their constraints: First Break, **22**, 47–70, doi: 10.3997/1365-2397.2004011.

Velis, D. R., 2008, Stochastic sparse-spike deconvolution: Geophysics, **73**, no. 1, R1–R9, doi: 10.1190/1.2790584.

Wang, K., L. Bandura, D. Bevc, S. Cheng, J. DiSiena, A. Halpert, K. Osypov, B. Power, and E. Xu, 2019, End-to-end deep neural network for seismic inversion: 89th Annual International Meeting, SEG, Expanded Abstracts, 2418–2422, doi: 10.1190/segam2019-3216464.1.

Wang, L., Q. Zhao, J. Gao, Z. Xu, M. Fehler, and X. Jiang, 2016, Seismic sparse-spike deconvolution via toeplitz-sparse matrix factorization: Geophysics, **81**, no. 3, V169–V182, doi: 10.1190/geo2015-0151.1.

Wu, X., 2017a, Building 3D subsurface models conforming to seismic structural and stratigraphic features: Geophysics, **82**, no. 3, IM21–IM30, doi: 10.1190/geo2016-0255.1.

Wu, X., 2017b, Structure-, stratigraphy-, and fault-guided regularization in geophysical inversion: Geophysical Journal International, **210**, 184–195, doi: 10.1093/gji/ggx150.

Yuan, S., S. Wang, C. Luo, and Y. He, 2015, Simultaneous multitrace impedance inversion with transform-domain sparsity promotion: Geophysics, **80**, no. 2, R71–R80, doi: 10.1190/geo2014-0065.1.

Zhang, R., and J. Castagna, 2011, Seismic sparse-layer reflectivity inversion using basis pursuit decomposition: Geophysics, **76**, no. 6, R147–R158, doi: 10.1190/geo2011-0103.1.

Zheng, Y., Q. Zhang, A. Yusifov, and Y. Shi, 2019, Applications of supervised deep learning for seismic interpretation and inversion: The Leading Edge, **38**, 526–533, doi: 10.1190/tle38070526.1.

Biographies and photographs of the authors are not available.