

Waveform embedding: Automatic horizon picking with unsupervised deep learning

Yunzhi Shi¹, Xinming Wu², and Sergey Fomel¹

ABSTRACT

Picking horizons from seismic images is a fundamental step that could critically impact seismic interpretation quality. We have developed an unsupervised approach, waveform embedding, based on a deep convolutional autoencoder network to learn to transform seismic waveform samples to a latent space in which any waveform can be represented as an embedded vector. The regularizing mechanism of the autoencoder ensures that similar waveform patterns are mapped to embedded vectors with a shorter distance in the latent space. Within a search region, we transform all of the waveform samples to the latent space and compute their corresponding distance to the embedded vector of a control point that is set to the target horizon. We then convert the distance to a horizon probability map that highlights where the horizon is likely to be located. This method can guide horizon picking across lateral discontinuities such as faults, and it is insensitive to noise and lateral distortions. In addition, our unsupervised learning algorithm requires no training labels. We apply our horizon-picking method to multiple 2D/3D examples and obtain results more accurate than the baseline method.

INTRODUCTION

In seismic interpretation, picking and tracking laterally consistent seismic reflectors as seismic horizons is a key step for the purpose of mapping geologic structures, stratigraphy, and reservoir architecture (Wu and Hale, 2013). Seismic horizons represent stratigraphic surfaces of constant geologic time (Vail et al., 1977). Many geologic analyses are performed on horizon surfaces. Therefore, the accuracy of the horizon picking is critical to all of the ensuing interpretation tasks. However, the growth of seismic data volumes significantly

increases the workload and duration of manual horizon extraction. It is necessary to aid human interpreters with automated, yet accurate and reliable, techniques.

Many computer algorithms that automatically extract horizons from seismic images have been proposed. Some methods are based on local slopes that estimate the geometric orientations of seismic reflections. The local slopes can be estimated by structure tensors (Bakker, 2002; Hale, 2009; Wu and Janson, 2017), plane-wave destruction filters (Fomel, 2002), semblance scanning (Marfurt, 2006), 2D log-Gabor filtering (Yu et al., 2013), smooth dynamic image warping (Arias, 2016), and tracking the structure-oriented vector field (Yu et al., 2011, 2012; Di et al., 2018). Fomel (2010) picks horizons by starting from a seed point and recursively following the local slopes. When multiple seed points are available, horizons can be computed by fitting the slopes in the least-squares sense (Lomask et al., 2006; Wu and Hale, 2013, 2015; Zinck et al., 2013; Monniron et al., 2016). However, local slope-based methods often fail in cases of lateral discontinuities (e.g., faults). Luo and Hale (2013), Wu and Hale (2016), and Wu et al. (2016) address the problem by removing the faulting. Xue et al. (2018) incorporate fault slip information in the tracking process. These methods all depend on fault detection and interpretation, which is another challenging task to be automated. On the other hand, some automatic horizon picking methods are based on seismic phase and waveform and, thus, can globally extract horizons by identifying the patterns in the seismic waveforms. Hoyes and Cheret (2011) present a review that summarizes global interpretation methods for 3D horizon mapping. Wu and Fomel (2018) propose a method to accurately extract horizons by computing laterally multigrid slopes by directly correlating seismic traces within multiple laterally coarse grids.

Machine learning shows its potential in seismic interpretation because of its power in solving computer vision problems. Many relevant works demonstrate how to outperform conventional methods by leveraging the power of machine learning. Zhao et al. (2015) review some commonly used machine-learning methods for seismic

Manuscript received by the Editor 16 July 2019; revised manuscript received 21 January 2020; published ahead of production 7 February 2020; published online 08 May 2020.

¹The University of Texas at Austin, Austin, Texas, USA. E-mail: yzshi08@utexas.edu; sergey.fomel@beg.utexas.edu.

²University of Science and Technology of China, School of Earth and Space Sciences, Hefei, China. E-mail: xinmwu@ustc.edu.cn (corresponding author).

© 2020 Society of Exploration Geophysicists. All rights reserved.

facies classification such as principal component analysis, k -means clustering, self-organizing maps, support vector machines, and artificial neural networks. Zhao (2017) shows that computers can learn to perform repetitive tasks in geophysical interpretation and unravel the relationship that underlies the patterns. Figueiredo et al. (2014) propose a clustering method using the growing neural gas algorithm, which organizes the set of voxels according to a global criterion in an unsupervised fashion, and then it classifies seismic waveforms into these clusters and extracts the corresponding horizons. These methods try to learn a transform to map seismic data to a latent space, where it is easier to cluster or separate the data samples for performing classification. However, these methods still have difficulties when dealing with complex real data. The relatively simple transform kernels of these machine-learning algorithms prevent them from learning complex transforms that are required to handle real data. Recently, deep neural networks, especially convolutional neural networks, have been widely applied to seismic interpretation problems because they are suited well for complex visual analysis. Compared with traditional machine-learning methods, deep neural networks can automatically extract useful features during the training process and achieve complex tasks. Deng et al. (2010) design an autoencoder network that learns latent speech embedding that explores the data structure by transforming input data to a latent space and represent each data with an embedded vector. Xie et al. (2016) show that deep embedding clustering can outperform traditional clustering methods. Qian et al. (2018) propose a convolutional autoencoder to perform unsupervised seismic facies analysis from prestack data. Peters et al. (2019) show a seismic horizon tracking method that trains multiresolution neural networks to predict horizons at unknown locations; but this is a supervised method that requires some manual interpretation results.

We propose a new method, waveform embedding, to accurately extract horizons based on a deep convolutional autoencoder network. The network learns to transform seismic waveform samples to a latent space in which all waveforms can be represented as an embedded vector. The additional regularizing mechanism of the autoencoder ensures that similar waveform patterns are mapped to embedded vectors with a shorter distance in the latent space. Starting from one or more control points, we compute an initial horizon estimation using the traditional local slope-based method and locate a search region around this initial horizon to sample waveforms. Then, we train the network with these samples, transform the waveforms to latent space, and compute their embedded distances to the embedded vector corresponding to the control points. We convert these distances to a horizon probability map where shorter distances stand for high similarities and higher probabilities and longer distances stand for low similarities and low probabilities. Finally, we can find a horizon curve (in two dimensions) or surface (in three dimensions) from this probability map by picking through the high values. We show multiple 2D and 3D examples to demonstrate that the proposed method can generate accurate and reliable results across discontinuities and outperform the conventional methods in real data tests.

METHODOLOGY

Initial horizon estimation: Slope-based method

Although local slope-based horizon extraction methods are not capable of tracking horizons correctly across faults, they are relatively efficient to compute an initial horizon estimation. We use structure tensors (Bakker, 2002; Hale, 2009; Wu and Janson, 2017) to estimate

local slopes of the seismic reflections. The 2D structure tensors are 2×2 symmetric positive semidefinite matrices that can be represented as the following eigendecompositions:

$$\mathbf{T} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T, \quad (1)$$

where \mathbf{u} and \mathbf{v} are the normalized eigenvectors corresponding to the eigenvalues λ_u and λ_v , respectively. By assuming $\lambda_u \geq \lambda_v$, the corresponding eigenvectors \mathbf{u} are perpendicular to the seismic reflections, whereas the eigenvectors \mathbf{v} are parallel to the reflections. Equation 1 can be generalized for 3D structure tensors as

$$\mathbf{T} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T + \lambda_w \mathbf{w}\mathbf{w}^T. \quad (2)$$

Without a loss of generality, we can assume $\lambda_u \geq \lambda_v \geq \lambda_w$. Therefore, the normalized eigenvectors \mathbf{u} are perpendicular to the seismic reflections, whereas the eigenvectors \mathbf{v} and \mathbf{w} are parallel to the reflections and perpendicular to each other. The eigenvalues λ_u , λ_v , and λ_w are useful to measure the linearity (2D) or planarity (3D) of the seismic reflections (Hale, 2009; Wu, 2017):

$$c(\mathbf{x}) = \frac{\lambda_u(\mathbf{x}) - \lambda_v(\mathbf{x})}{\lambda_u(\mathbf{x})}, \quad (3)$$

where $0 \leq c(\mathbf{x}) \leq 1$ describes the continuities of the seismic reflections. Assuming that the reflection normal vectors \mathbf{u} always point downward, we can compute the inline and crossline local slopes as

$$p(\mathbf{x}) = -\frac{u_2(\mathbf{x})}{u_1(\mathbf{x})}, \quad \text{and} \quad q(\mathbf{x}) = -\frac{u_3(\mathbf{x})}{u_1(\mathbf{x})}, \quad (4)$$

where u_1 , u_2 , and u_3 are the vertical, inline, and crossline components of the normal vectors \mathbf{u} , respectively, and p and q are the inline and crossline local slopes, respectively.

After obtaining local slopes of the seismic reflections, we can find a slope-based horizon picking result by fitting the slopes in the least-squares sense (Lomask et al., 2006; Wu and Hale, 2013, 2015; Zinck et al., 2013; Monniron et al., 2016):

$$\begin{bmatrix} c(x, z(x)) \frac{\partial z(x)}{\partial x} \\ \mu \frac{\partial^2 z(x)}{\partial x^2} \end{bmatrix} \approx \begin{bmatrix} c(x, z(x)) p(x, z(x)) \\ 0 \end{bmatrix}, \quad (5)$$

where $z(x)$ represents the target horizon. The first equation in equation 5 is a fitting between horizon slopes $\partial z(x)/\partial x$ and local slope estimations $p(x, z(x))$ weighted by e thlinearity attribute $0 \leq c(x, z(x)) \leq 1$ defined in equation 3. The weighting helps to reduce the effects of unstable local slope estimations at discontinuous areas when linearity values are low, and it helps to put more weights on local slope estimations where linearity values are high. The second equation is a regularization term to impose smoothness to the horizon $z(x)$, where μ is a small constant to control smoothness. Equation 5 can be extended to 3D least-squares horizon surface fitting as

$$\begin{bmatrix} c(x, y, z(x, y)) \frac{\partial z(x, y)}{\partial x} \\ c(x, y, z(x, y)) \frac{\partial z(x, y)}{\partial y} \\ \mu \left(\frac{\partial^2 z(x, y)}{\partial x^2} + \frac{\partial^2 z(x, y)}{\partial y^2} \right) \end{bmatrix} \approx \begin{bmatrix} c(x, y, z(x, y)) p(x, y, z(x, y)) \\ c(x, y, z(x, y)) q(x, y, z(x, y)) \\ 0 \end{bmatrix}. \quad (6)$$

To solve those equations, we need several known control points provided as boundary conditions. Suppose we have N_c control points (x_c, z_c) , $c = 1, 2, \dots, N_c$, we can incorporate the hard constraints $z(x_c) = z_c$, $c = 1, 2, \dots, N_c$ to the fitting equations and solve the problem iteratively using a conjugate gradient solver with a constraint preconditioner as discussed by Wu and Hale (2015).

As an example in Figure 1a, we start with a control point (the red cycle) and compute a horizon extraction (the yellow curve) with local slopes least-squares fitting with the method described above. Although the horizon follows seismic reflection closely, it fails to correctly shift at the seismic fault to consistently track the reflection phase. We flatten the seismic image around the horizon extraction as shown in Figure 1b. We can observe that the local slope-based method yields a horizon parallel to the reflections but cannot correctly track the horizon across the fault.

Baseline method: Multigrid correlations

The horizon extraction method that only depends on local slope estimations are often inaccurate in discontinuous areas with noise, chaotic reflections, or faults as shown in Figure 1, other information that can “propagate” across the fault should be used to solve such a problem. Correlating seismic waveforms on the opposite sides of the faults can help to guide the slope-based horizon-tracking algorithm.

Wu and Fomel (2018) propose a method that fits a horizon in the least-squares sense with the local slopes and spatially more global reflection correlations: the least-squares horizon with multigrid




Figure 1. (a) We start with a control point (the red cycle) and compute a horizon extraction (the yellow curve) with local slopes least-squares fitting. Although the horizon follows seismic reflection closely, it fails to correctly shift at the seismic fault to consistently track the reflection phase. We sample waveforms around this horizon (the blue mask) to train the unsupervised deep network in the proposed workflow. (b) We flatten the seismic image around the horizon extraction, and we can observe that the local slope-based method yields a horizon parallel to the reflections but cannot correctly track the horizon across the fault.

correlations. The method can correlate reflections that are spatially far away from each other, in laterally coarser grids. The pairwise correlation of waveforms are efficiently computed using dynamic time warping (Hale, 2013). These coarse grid correlations introduce additional constraint equations to equations 5 and 6. These independent constraints can provide multiscale information in tracking the target horizon. In cases in which local slopes cannot correctly follow reflections across discontinuities, the multigrid correlation can provide coarse-grid slopes that help to correct the errors and find the right reflections to continue tracking.

We will use the multigrid correlation method as a baseline conventional horizon extraction method to compare its performance with the proposed waveform embedding deep-learning-based method. Figure 2 shows the result of horizon picking using the multigrid correlation method with the same image and control point as shown in Figure 1. We observe that the multigrid correlation method successfully finds the correct horizon across the fault region because the waveform correlations on the opposite sides of the fault force the horizon to track the correct phase, in this case, at the end of the negative wavelet peak. However, close examination of the flattened image in Figure 2b reveals that the horizon extracted by the multigrid correlation method deviates downward from the correct horizon, from left to right. This is due to the lateral change along the horizon (as shown with the wiggles) misleading the correlation: The reflection wavelet frequency is lower on the right side compared to the left




Figure 2. The result of horizon picking using the multigrid correlation method with the same image and control point as shown in Figure 1. (a) Multigrid correlation method successfully finds a fix across the fault region because the waveform correlations on the opposite sides of the fault force the horizon to track the correct phase. However, close examination of the flattened image in (b) reveals that the horizon extracted by the multigrid correlation method is deviating downward from the correct horizon, from left to right. This is due to the lateral change along the horizon (as shown in the wiggles) misleading the correlation: The reflection wavelet frequency is lower on the right side compared to the left side due to lateral geologic distortion.

side due to the lateral geologic distortion. Although the multigrid correlation method could improve the horizon tracking result across the fault, correlation itself may not be adequate to handle these subtle changes to output a precise and reliable horizon estimation.

Waveform embedding: Deep-learning approach

Compared with conventional methods and traditional machine-learning methods, deep neural networks can automatically extract useful features and during the training process and achieve complex tasks. Deng et al. (2010) design an autoencoder network that learns latent speech embedding and can explore the data structure by transforming input data to a latent space and represent each datum with an embedded vector. We propose to build an autoencoder network and train it with seismic reflection waveform data samples to learn implicit features from these waveforms represented by a latent space.

An autoencoder consists of two parts: the encoder and the decoder, which can be defined as two nonlinear transform operators ϕ and ψ , such that

$$\phi: \mathcal{X} \rightarrow \mathcal{F},$$

$$\psi: \mathcal{F} \rightarrow \mathcal{X},$$

$$\text{where } \phi, \psi = \operatorname{argmin}_{\phi, \psi} \|X - (\psi \circ \phi)X\|^2. \quad (7)$$

The autoencoder network $\mathcal{X} \rightarrow \mathcal{F} \rightarrow \mathcal{X}$ transforms an input waveform data X to the latent space \mathcal{F} , then back to the original data space \mathcal{X} . The training is done by minimizing the reconstruction discrepancy between X and $(\psi \circ \phi)X$. If the network can reconstruct the input X after training with many samples from \mathcal{X} , this means that operator ϕ retains most of the essential information of \mathcal{X} and encodes all possible X into the latent space \mathcal{F} . By incorporating regularization to the \mathcal{F} , the network can learn a useful latent space that separates/clusters the waveforms by their visual features. The vector ϕX in latent space \mathcal{F} is the embedded vector that represents all the features extracted from the waveform X ; therefore, we call this process waveform embedding.

To effectively extract visual features from waveform samples, we build a deep convolutional autoencoder network as shown in

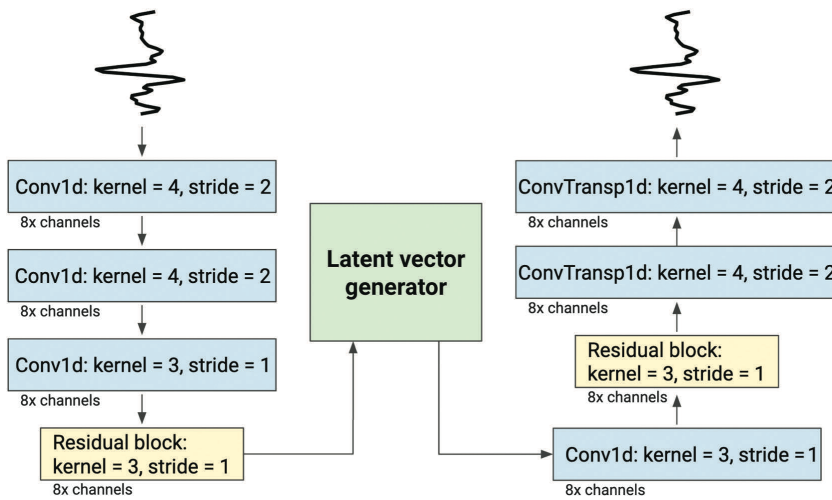


Figure 3. Illustration of the autoencoder network architecture that we implement for the waveform embedding method.

Figure 3. The encoder consists of cascading convolutional layers with poolings (downsampling) and nonlinear activations. Because seismic waveforms are one dimensional, we use 1D convolutions. We also introduce a residual convolutional block, introduced by He et al. (2016), to achieve better performance and training efficiency. The encoder ends with a latent vector generator, for which we investigate three different choices: first, directly copying the convolutional output as latent vectors; second, sampling latent vectors from known prior information (usually Gaussian); and third, quantizing the latent space and sampling from the quantizations. These different latent vector generators result in three different types of deep autoencoders: naive autoencoder, variational autoencoder (VAE) (Kingma and Welling, 2013), and vector-quantized VAE (VQ-VAE) (van den Oord and Vinyals, 2017).

On the decoder side, the latent vector will pass through a decoder, which is symmetric to the encoder. Instead of downsampling, the decoder layers upsample by transposed convolutions (Noh et al., 2015; Dumoulin and Visin, 2016) to reconstruct the original waveform. This type of downsample-upsample shape works as a regularization: An information bottleneck at the center of the network forces it to learn essential information in the latent space and discard irrelevant noise features.

Different types of autoencoders

A naive autoencoder passes the output $Z = \phi X$ to the decoder as the input $\bar{X} = \psi Z$. A naive autoencoder is easy to build and train; however, it has no additional regularization on the latent space to better explore the visual features of the waveform (e.g., clustering the samples with similar waveforms).

VAE, proposed by Kingma and Welling (2013), provides a probabilistic manner for describing the observation in latent space. Instead of outputting deterministic embedded vectors, the encoder in VAE outputs a probability distribution for each latent feature. The output of the encoder and decoder can be formulated as

$$Z = \phi X \sim q(Z|X), \quad \text{and} \quad \bar{X} = \psi Z \sim p(X|Z). \quad (8)$$

To control the distribution of the latent vector Z , the following Kullback-Leibler (KL) loss regularizes Z to be independent unit Gaussian distribution $\mathcal{N}(0, I)$:

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(q(Z|X)|p(Z)), \quad (9)$$

where KL divergence D_{KL} is defined as

$$D_{\text{KL}}(P|Q) = \sum_X P(X) \log(P(X)/Q(X)). \quad (10)$$

Van den Oord and Vinyals (2017) propose a variation of the VAE, VQ-VAE, in which the encoder outputs a discrete, rather than continuous, latent space with a data-driven learned prior distribution rather than a static predefined prior. It learns latent transform efficiently by updating a discrete latent prior distribution together with the encoder and decoder while training.

We build these three types of autoencoders and implement them using naive autoencoder

PyTorch library (Paszke et al., 2019). All of these implementations share the same encoder and decoder setups as shown in Figure 3, and they have an 8D latent space to ensure a fair comparison.

Train with waveform data

We train and test the networks with the seismic waveforms shown in Figure 1. We randomly extract waveforms with 32 vertical samples from the blue highlighted region in Figure 1a, and we use these data to train the autoencoder network. Note that although the slope-based method picks an incorrect horizon that fails to handle faults, it provides an approximate trend that the true horizon should lie parallel to it nearby. Unsupervised learning does not require any training label. The slope-based horizon is only used to identify a region from which we sample training data, so it does not matter whether or not it gives an accurate horizon.

During training, we use an independent subset data from the same region to validate the autoencoder performance by measuring the mean-square-error (MSE) of the reconstructed waveform. Figure 4 shows the validation MSE of the three types of autoencoders. The naive implementation and VQ-VAE outperform VAE after convergence. Compared to AE, VAE introduces regularization in the latent space prior; therefore, it would not be as good as AE that does not have any regularization on the latent space. On the other hand, VQ-VAE has a discretized latent distribution with data-driven learned prior rather than static predefined prior; therefore, it can outperform VAE in terms of reconstruction accuracy. Because VQ-VAE achieves the best convergence and enables a regularized latent space prior distribution, we choose this framework for horizon picking.

Convert horizon probability

After training, we can use the encoder part of the network to transform any seismic waveforms to an embedded vector \mathbf{l} . For each pixel \mathbf{x} in the search region (the blue highlighted region in Figure 1) around the initial horizon estimation, we can sample a waveform centered at the pixel and transform it to a latent vector $\mathbf{l}(\mathbf{x})$, including the latent vector \mathbf{l}_c of the control point(s). Because of the distribution of the latent prior, similar waveforms will have embedded vectors with a shorter distance $d = \|\mathbf{l}_1 - \mathbf{l}_2\|_2$. Therefore, we compute a normalized embedded distance corresponding to each pixel \mathbf{x} :

$$d_c(\mathbf{x}) = \frac{\|\mathbf{l}(\mathbf{x}) - \mathbf{l}_c\|_2}{d_0}, \quad (11)$$

where d_0 is a constant normalizing factor that we set as $0.9 \times \max(d)$. Finally, we can convert the embedded distances to a probability map $p(\mathbf{x})$ using radial basis functions:

$$p_c(\mathbf{x}) = e^{-d_c(\mathbf{x})}. \quad (12)$$

Figure 5 shows the probability map, in which brightness represents the horizon probability with regard to the control point. If multiple control points are available ($N_c > 1$), we take the average horizon probability of all the probability maps corresponding to each control point. Using a picking method based on dynamic programming (Sakoe et al., 1990), we can extract a path through the maximum probabilities shown in Figure 6 as the horizon extraction result. Figure 7 summarizes the workflow in a flowchart.

APPLICATION EXAMPLES

By comparing Figure 6 to Figure 2, we can observe that the proposed method correctly tracks the seismic phase along the target horizon without being misled by the lateral distortion. This shows that the waveform embedding method extracts subtle visual features of the waveforms and outperforms waveform correlation. To further demonstrate the superiority of waveform embedding, we apply these methods to more 2D and 3D real data examples.

2D applications

The previous examples use a 2D data subset from the Teapot Dome survey. Figure 8 shows the application on multiple horizons with randomly positioned control points. Although the waveforms

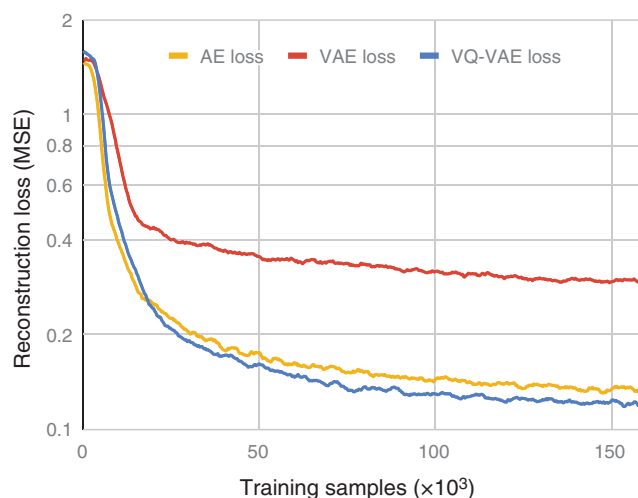


Figure 4. The validation MSE of waveform reconstruction during training of the three types of autoencoders. The naive autoencoder and VQ-VAE outperform VAE after convergence. Compared to AE, VAE introduces regularization in the latent space prior, and it is thus inferior because AE has no regularization on the latent space. On the other hand, VQ-VAE has a discretized latent distribution with a data-driven learned prior rather than statically predefined prior; therefore, it can outperform VAE in terms of reconstruction accuracy. Because VQ-VAE achieves the best convergence and enables a regularized latent space prior distribution, we choose this framework for horizon picking.

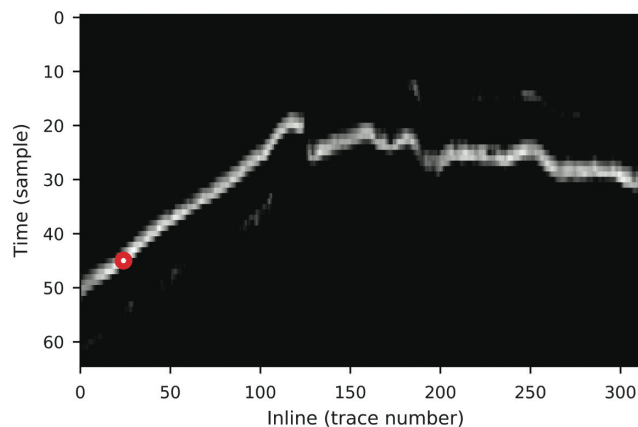


Figure 5. The horizon probability map given in equation 12.

are quite consistent laterally, the structure is complicated by several faults. The slope-only method generates continuous but incorrect horizons in Figure 8a; these horizons will be used as the initial hori-




Figure 6. We extract a path through the maximum probabilities from the probability map shown in Figure 5 and acquire this horizon extraction. Compared with Figure 2, the proposed method correctly tracks the seismic phase along the target horizon without being misled by the lateral distortion. This shows that the waveform embedding method extracts subtle visual features of the waveforms and outperforms waveform correlation.

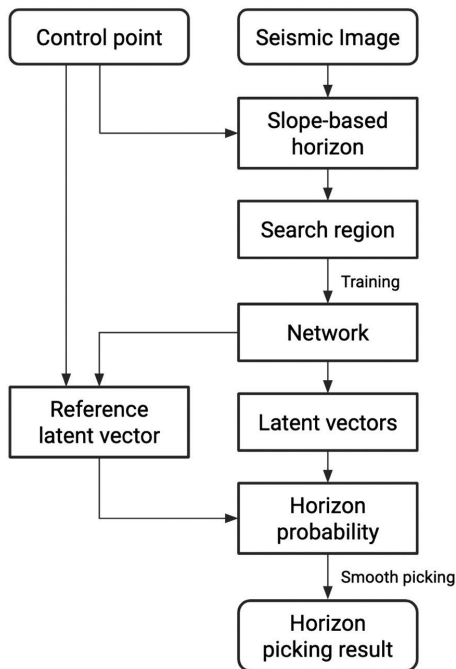


Figure 7. The flowchart of the proposed waveform embedding method.

zon estimation, around which waveform samples are extracted and input to the waveform embedding method for a finer horizon extraction. The multigrid correlation method generates more accurate horizons across the faults in Figure 8b. Even without explicitly detecting and estimating fault attributes as in Luo and Hale (2013), Wu and Hale (2016), Wu et al. (2016), and Xue et al. (2018), this baseline method still guides the horizon to track correct phases by correlating the waveforms in coarse grids. Figure 8c shows the horizon extraction results by the proposed waveform embedding method. Although Figures 2 and 6 already show that the proposed method can generate more accurate and reliable horizons than the




Figure 8. The application on multiple horizons with randomly positioned control points using the 2D data set from the Teapot Dome survey. (a) The slope-only method generates continuous but incorrect horizons. (b) The multigrid correlation method generates more accurate horizons across the faults. (c) The horizon extraction results by the proposed waveform embedding method.




Figure 9. A 3D data set cropped from the Netherlands offshore F3 block seismic data. These data have a group of closely aligned parallel faults and some other faults. The yellow arrows annotate the target horizon that we aim to extract using the aforementioned methods.

baseline method, these results show that the waveform embedding method can also perform equally well as the baseline method in other situations.

3D applications

Figure 9 shows a 3D data set (420 [inline] \times 400 [crossline] \times 100 [time] samples) cropped from the Netherlands offshore F3 block seismic data. These data have a group of closely aligned parallel faults and some other faults. The yellow arrows annotate the target horizon that we aim to extract using the aforementioned methods.

We extract the first horizon result using the slope-only method as shown in Figure 10a, in which the horizon surface is color coded with the amplitude values as shown in Figure 9 (the background seismic images use a grayscale colormap to enhance visualization), and the green cube in the middle represents the position of the control point. Although the horizon follows the trend of the slopes, it fails to correctly track seismic reflections across those faults, and thus results in obvious amplitude variations on the surface (inconsistent colors). Because the control point is located at an amplitude trough, the ideal horizon surface should be colored in red everywhere, and inconsistent colors indicate incorrect jumps across seismic reflections.

We use the multigrid correlation method to extract the second horizon result as shown in Figure 10b. It is easy to notice the improvement from the slope-only method: The amplitude is now much more consistent across the horizon surface. We can also observe those parallel faults in a staircase shape from the horizon surface.

Finally, Figure 10c shows the horizon extracted by the proposed waveform embedding method. We can see a further improvement in terms of amplitude consistency compared with the conventional method. Except for a few transition regions with large fault throws and noisy waveforms, the horizon surface is colored in red almost everywhere as expected from the control point amplitude. Figure 11 shows the corresponding 2D slice view.

Figure 12 presents another more complex 3D data set (600 [inline] \times 600 [crossline] \times 400 [time] samples) that contains larger fault throws and more significant lateral distortions, which is a challenging task for any automated horizon extraction method.

Figure 13 shows the horizon extraction results using the aforementioned methods from a single control point located on the left side of the figure. Figure 14 shows the corresponding 2D slice view. Due to the large fault slip nearby, the slope-only method mis-

tracks the horizon to a shallower layer on the right side of the figure and also has very inconsistent amplitude values. The multigrid correlation method, despite being better than the slope-only method in terms of consistent amplitudes, still tracks the incorrect horizon.




Figure 11. A 2D slice view of the horizon results shown in Figure 10. The slice is selected along the direction perpendicular to most of the faults (left to right in the 3D view). The slope-only method (yellow), multigrid correlation method (lime), and proposed waveform embedding method (cyan).




Figure 12. Another more complex 3D data set that contains larger fault throws and more significant lateral distortions, which could be a challenging task for any automated horizon extraction methods.




Figure 10. Comparison of the horizon extraction results. The green cube in the middle represents the position of the control point. The horizon surface is color coded with the amplitude values to visualize the amplitude consistency. The (a) slope-only method, (b) multigrid correlation method, and (c) proposed waveform embedding method.

The proposed method can extract a horizon that is tracking the correct trend of the horizon by going deeper gradually from left to right. However, due to lateral waveform distortions, the horizon quality deteriorates from left to right with more and more inconsistent layers.

Figure 15 shows similar horizon extraction results in Figure 13, but with a different control point located in the middle of the figure. Figure 16 shows the corresponding 2D slice view. The slope-only

method generates a rather flat horizon surface. The multigrid correlation method also fails to deal with the large fault throws and yields a flat horizon surface that tracks incorrect horizons. The proposed method, with this new control point, performs better than Figure 13c on the central and right part of the horizon surface and also tracks the incorrect horizon to the left.

Figure 17 shows the horizon extraction results with both control points available to these methods. Figure 18 shows the correspond-




Figure 13. The horizon extraction results using the aforementioned methods from a single control point located on the left side of the figure. The green cube in the middle represents the position of the control point. The horizon surface is color coded with the amplitude values to visualize amplitude consistency. (a) Slope-only method, (b) multigrid correlation method, and (c) the proposed waveform embedding method.




Figure 14. The 2D slice view of the horizon results shown in Figure 13. The slice is selected along the direction perpendicular to most of the faults (left to right in the 3D view). The slope-only method (yellow), multigrid correlation method (lime), and proposed waveform embedding method (cyan).




Figure 16. The 2D slice view of the horizon results shown in Figure 15. The slice is selected along the direction perpendicular to most of the faults (left to right in the 3D view). The slope-only method (yellow), multigrid correlation method (lime), and proposed waveform embedding method (cyan).




Figure 15. Similar horizon extraction results in Figure 13, but with a different control point located in the middle of the figure. The green cube in the middle represents the position of the control point. The horizon surface is color coded with the amplitude values to visualize amplitude consistency. (a) The slope-only method, (b) multigrid correlation method, and (c) proposed waveform embedding method.




Figure 17. The horizon extraction results with both control points available to these methods. The green cube in the middle represents the position of the control point. The horizon surface is color coded with the amplitude values to visualize amplitude consistency. (a) The slope-only method, (b) multigrid correlation method, and (c) proposed waveform embedding method.




Figure 18. The 2D slice view of the horizon results shown in Figure 17. The slice is selected along the direction perpendicular to most of the faults (left to right in the 3D view). The slope-only method (yellow), multigrid correlation method (lime), and proposed waveform embedding method (cyan).

ing 2D slice view. Although the local slope-based method now captures a correct fault trend, we can still observe inconsistent amplitude values from the surface colors. The multigrid correlation method improves significantly and shows a very consistent amplitude surface with blue colors almost everywhere. However, the proposed method also improves significantly and outperforms the conventional method in terms of consistent amplitude, accurate horizon tracking, and much sharper display of all the faults. This demonstrates that waveform embedding horizon picking method can handle large fault throws and significant lateral distortions in real data.

For all the tests involving deep-learning training/inference, we use a GTX 1080 graphics processing unit (GPU) card to perform the computations. In the last field example, we set sampling height to 60 samples and we randomly sample 32-length waveform traces from this $600 \times 600 \times 60$ region. The inference requires passing all waveform traces within this search region through the encoder of the trained network. In total, $600 \times 600 \times (60 - (32 - 1)) = 10,440,000$ waveform samples are input to the network. Including the following dynamic warping process, the total machine time of this workflow is 21–23 min. Meanwhile, the multigrid correlation method would take 10–15 min depending on the tracking parameters. We can see that, with the acceleration of GPU, the proposed method spends the same order of computational time compared to the baseline method.

DISCUSSION

Due to the complexity of geology structures in field surveys, there are still some limitations of the proposed method to practical applications. First, we assume a limited fault slip; therefore, we could sample from a relatively narrow zone around the initial slope-based horizon estimation. However, the larger fault slip will require a wider sampling region, thus increasing the amount of computations. Besides, the quality of the regularized picking using dynamic warping could suffer from instability when encountering large jump across the fault. Moreover, the proposed method could not handle unconformities well because they could either terminate a horizon or significantly change the local waveforms on the same horizon, with both posing challenges to the proposed method based on measuring waveform similarity.

To solve the aforementioned issues, we consider the following techniques as future works to improve this method. Introducing uncertainty estimation of the network prediction can help to improve the reliability and quality control capability of the results (Gal and Ghahramani, 2016; Leibig et al., 2017; Feng et al., 2018), thus allowing some postprocessing to refine the result. Meanwhile, dilated convolution (Yu and Koltun, 2015; Strubell et al., 2017) can significantly increase the receptive size of a network with the same number of layers compared to regular convolution, thus allowing the larger input waveform length. Adding more context in the input extends the capability of this method to complex fault or unconformity structures.

CONCLUSION

We propose waveform embedding, an automatic horizon picking method based on the deep-learning autoencoder. We implement a VQ-VAE network to reconstruct seismic waveform samples, while learning to transform the waveforms to embedded vectors in a latent space. The regularization of the network ensures that the latent space assigns similar waveforms to closer embedded vectors in the latent space, therefore allowing us to use this latent space to effectively measure the similarity between waveforms. Given one or more control points, we first estimate an initial horizon using an efficient least-squares local-slope fitting method, and then extract waveforms in the nearby region of this initial horizon estimation. We train the network with these waveforms and compute the embedded distances to the control point, and then we convert

the embedded distances to horizon probabilities. Finally, a horizon curve (2D) or surface (3D) can be picked from the probability map.

We use 2D/3D real data application examples to demonstrate the superiority of the proposed method over conventional methods. Because deep learning automatically extracts important features into the latent space during training, the proposed method can discover subtle visual features in the waveforms. Even with large fault throws and lateral distortions that challenge conventional methods, the horizons extracted by the proposed method still track correct seismic reflections with consistent phases.

ACKNOWLEDGMENTS

This research is financially supported by the sponsors of the Texas Consortium for Computational Seismology. We thank NVIDIA GPU Grant Program for donating a Titan Xp GPU for research computation.

DATA AND MATERIALS AVAILABILITY

Data associated with this research are confidential and cannot be released.

REFERENCES

- Arias, E., 2016, Estimating seismic reflection slopes: Master's thesis, Colorado School of Mines.
- Bakker, P., 2002, Image structure analysis for seismic interpretation: Delft University of Technology.
- Deng, L., M. L. Seltzer, D. Yu, A. Acero, A. Mohamed, and G. Hinton, 2010, Binary coding of speech spectrograms using a deep auto-encoder: Presented at the 11th Annual Conference of the International Speech Communication Association.
- Di, H., D. Gao, and G. AlRegib, 2018, 3D dip vector-guided auto-tracking for weak seismic reflections: A new tool for shale reservoir visualization and interpretation: *Interpretation*, **6**, no. 4, SN47–SN56, doi: [10.1190/INT-2018-0053.1](https://doi.org/10.1190/INT-2018-0053.1).
- Dumoulin, V., and F. Visin, 2016, A guide to convolution arithmetic for deep learning: arXiv preprint, arXiv:1603.07285.
- Feng, D., L. Rosenbaum, and K. Dietmayer, 2018, Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3D vehicle detection: 21st IEEE International Conference on Intelligent Transportation Systems (ITSC), 3266–3273.
- Figueiredo, A., P. Silva, M. Gattass, F. Silva, and R. L. Miliúdi, 2014, A seismic facies analysis approach to map 3D seismic horizons: 84th Annual International Meeting, SEG, Expanded Abstracts, 1501–1505, doi: [10.1190/segam2014-1382.1](https://doi.org/10.1190/segam2014-1382.1).
- Fomel, S., 2002, Applications of plane-wave destruction filters: *Geophysics*, **67**, 1946–1960, doi: [10.1190/1.1527095](https://doi.org/10.1190/1.1527095).
- Fomel, S., 2010, Predictive painting of 3D seismic volumes: *Geophysics*, **75**, no. 4, A25–A30, doi: [10.1190/1.3453847](https://doi.org/10.1190/1.3453847).
- Gal, Y., and Z. Ghahramani, 2016, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning: International Conference on Machine Learning (ICML), 1050–1059.
- Hale, D., 2009, Structure-oriented smoothing and semblance: CWP Report, **635**, 261–270.
- Hale, D., 2013, Dynamic warping of seismic images: *Geophysics*, **78**, no. 2, S105–S115, doi: [10.1190/geo2012-0327.1](https://doi.org/10.1190/geo2012-0327.1).
- 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 (CVPR), 770–778.
- Hoyes, J., and T. Cheret, 2011, A review of global interpretation methods for automated 3D horizon picking: *The Leading Edge*, **30**, 38–47, doi: [10.1190/1.3535431](https://doi.org/10.1190/1.3535431).
- Kingma, D. P., and M. Welling, 2013, Auto-encoding variational Bayes: arXiv, abs/1312.6114.
- Leibig, C., V. Allken, M. S. Ayhan, P. Berens, and S. Wahl, 2017, Leveraging uncertainty information from deep neural networks for disease detection: *Scientific Reports*, **7**, 17816, doi: [10.1038/s41598-017-17876-z](https://doi.org/10.1038/s41598-017-17876-z).
- Lomask, J., A. Guitton, S. Fomel, J. Claerbout, and A. A. Valenciano, 2006, Flattening without picking: *Geophysics*, **71**, no. 4, P13–P20, doi: [10.1190/1.2210848](https://doi.org/10.1190/1.2210848).
- Luo, S., and D. Hale, 2013, Unfaulting and unfolding 3D seismic images: *Geophysics*, **78**, no. 4, O45–O56, doi: [10.1190/geo2012-0350.1](https://doi.org/10.1190/geo2012-0350.1).
- Marfurt, K. J., 2006, Robust estimates of 3D reflector dip and azimuth: *Geophysics*, **71**, no. 4, P29–P40, doi: [10.1190/1.2213049](https://doi.org/10.1190/1.2213049).
- Monnirou, M., S. Frambati, S. Quillón, Y. Berthoumieu, and M. Donias, 2016, Seismic horizon and pseudo-geological time cube extraction based on a Riemmanian geodesic search: IEEE 12th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), 1–5.
- Noh, H., S. Hong, and B. Han, 2015, Learning deconvolution network for semantic segmentation: Proceedings of the IEEE International Conference on Computer Vision (ICCV), 1520–1528.
- Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, and L. Antiga, 2019, PyTorch: An imperative style, high-performance deep learning library: *Advances in Neural Information Processing Systems (NeurIPS)*, 8024–8035.
- Peters, B., J. Granek, and E. Haber, 2019, Multi-resolution neural networks for tracking seismic horizons from few training images: *Interpretation*, **7**, no. 3, SE201–SE213, doi: [10.1190/INT-2018-0225.1](https://doi.org/10.1190/INT-2018-0225.1).
- Qian, F., M. Yin, X.-Y. Liu, Y.-J. Wang, C. Lu, and G.-M. Hu, 2018, Unsupervised seismic facies analysis via deep convolutional auto-encoders: *Geophysics*, **83**, no. 3, A39–A43, doi: [10.1190/geo2017-0524.1](https://doi.org/10.1190/geo2017-0524.1).
- Sakoe, H., S. Chiba, A. Waibel, and K. Lee, 1990, Dynamic programming algorithm optimization for spoken word recognition: *Readings in Speech Recognition*, **159**, 224.
- Strubell, E., P. Verga, D. Belanger, and A. McCallum, 2017, Fast and accurate entity recognition with iterated dilated convolutions: arXiv preprint, arXiv:1702.02098.
- Vail, P. R., R. M. Mitchum, and S. Thompson, 1977, Seismic stratigraphy and global changes of sea level — Part 5: Chronostratigraphic significance of seismic reflections, in C. W. Payton, ed., *Seismic stratigraphy — application to hydrocarbon exploration: AAPG Memoir*, **26**, pp. 99–116.
- van den Oord, A., and O. Vinyals, 2017, Neural discrete representation learning: *Advances in Neural Information Processing Systems (NeurIPS)*, 6306–6315.
- Wu, X., 2017, Directional structure-tensor-based coherence to detect seismic faults and channels: *Geophysics*, **82**, no. 2, A13–A17, doi: [10.1190/geo2016-0473.1](https://doi.org/10.1190/geo2016-0473.1).
- Wu, X., and S. Fomel, 2018, Least-squares horizons with local slopes and multigrid correlations: *Geophysics*, **83**, no. 4, IM29–IM40, doi: [10.1190/geo2017-0830.1](https://doi.org/10.1190/geo2017-0830.1).
- Wu, X., and D. Hale, 2013, Extracting horizons and sequence boundaries from 3D seismic images: 83rd Annual International Meeting, SEG, Expanded Abstracts, 1440–1445, doi: [10.1190/segam2013-0296.1](https://doi.org/10.1190/segam2013-0296.1).
- Wu, X., and D. Hale, 2015, Horizon volumes with interpreted constraints: *Geophysics*, **80**, no. 2, IM21–IM33, doi: [10.1190/geo2014-0212.1](https://doi.org/10.1190/geo2014-0212.1).
- Wu, X., and D. Hale, 2016, Automatically interpreting all faults, unconformities, and horizons from 3D seismic images: *Interpretation*, **4**, no. 2, T227–T237, doi: [10.1190/INT-2015-0160.1](https://doi.org/10.1190/INT-2015-0160.1).
- Wu, X., and X. Janson, 2017, Directional structure tensors in estimating seismic structural and stratigraphic orientations: *Geophysical Journal International*, **210**, 534–548, doi: [10.1093/gji/ggx194](https://doi.org/10.1093/gji/ggx194).
- Wu, X., S. Luo, and D. Hale, 2016, Moving faults while unfaulting 3D seismic images: *Geophysics*, **81**, no. 2, IM25–IM33, doi: [10.1190/geo2015-0381.1](https://doi.org/10.1190/geo2015-0381.1).
- Xie, J., R. Girshick, and A. Farhadi, 2016, Unsupervised deep embedding for clustering analysis: International Conference on Machine Learning (ICML), 478–487.
- Xue, Z., X. Wu, and S. Fomel, 2018, Predictive painting across faults: *Interpretation*, **6**, no. 2, T449–T455, doi: [10.1190/INT-2017-0171.1](https://doi.org/10.1190/INT-2017-0171.1).
- Yu, F., and V. Koltun, 2015, Multi-scale context aggregation by dilated convolutions: arXiv preprint, arXiv:1511.07122.
- Yu, Y., C. Kelley, and I. Mardanava, 2011, Automatic horizon picking in 3D seismic data using optical filters and minimum spanning tree (patent pending): 81st Annual International Meeting, SEG, Expanded Abstracts, 965–969, doi: [10.1190/1.3628233](https://doi.org/10.1190/1.3628233).
- Yu, Y., C. L. Kelley, and I. M. Mardanava, 2012, Seismic horizon autopicking using orientation vector field: U.S. Patent 8,265,876.
- Yu, Y., C. Kelley, and I. Mardanava, 2013, Volumetric seismic dip and azimuth estimation with 2D log-Gabor filter array: 83rd Annual International Meeting, SEG, Expanded Abstracts, 1357–1362, doi: [10.1190/segam2013-0046.1](https://doi.org/10.1190/segam2013-0046.1).
- Zhao, T., 2017, Machine assisted quantitative seismic interpretation: Ph.D. thesis, University of Oklahoma.
- Zhao, T., V. Jayaram, A. Roy, and K. J. Marfurt, 2015, A comparison of classification techniques for seismic facies recognition: *Interpretation*, **3**, no. 4, SAE29–SAE58, doi: [10.1190/INT-2015-0044.1](https://doi.org/10.1190/INT-2015-0044.1).
- Zinck, G., M. Donias, J. Daniel, S. Guillon, and O. Lavialle, 2013, Fast seismic horizon reconstruction based on local dip transformation: *Journal of Applied Geophysics*, **96**, 11–18, doi: [10.1016/j.jappgeo.2013.06.010](https://doi.org/10.1016/j.jappgeo.2013.06.010).