

Ghassan AlRegib, Mohamed Deriche, Zhiling Long,
Haibin Di, Zhen Wang, Yazeed Alaudah,
Muhammad Amir Shafiq, and Motaz Alfarraj

Subsurface Structure Analysis Using Computational Interpretation and Learning

A visual signal processing perspective



©ISTOCKPHOTO.COM/KHADI GANIEV

Understanding Earth's subsurface structures has been and continues to be an essential component of various applications such as environmental monitoring, carbon sequestration, and oil and gas exploration. By viewing the seismic volumes that are generated through the processing of recorded seismic traces, researchers were able to learn from applying advanced image processing and computer vision algorithms to effectively analyze and understand Earth's subsurface structures. In this article, we first summarize the recent advances in this direction that relied heavily on the fields of image processing and computer vision. Second, we discuss the challenges in seismic interpretation and provide insights and some directions to address such challenges using emerging machine-learning algorithms.

Introduction

Seismic interpretation is a critical process in subsurface exploration. The process aims at identifying structures or environments of significant importance in diverse applications. For example, for oil and gas exploration, a successful interpretation can help identify structures (such as faults, salt domes, and horizons, etc.) that are indicators of potential locations of reservoirs. Reliable delineation of natural faults and fractures will help predict crucial geologic deformation events and monitor potential earthquakes. In environmental engineering, capture and storage of carbon dioxide (CO_2) is substantially dependent on accurate subsurface interpretation to find a good geologic trap with minimum leakage probability. When subsurface structures are the primary interest in seismic interpretation, it is often called *structural interpretation*. However, when environments such as patterns of deposition are of concern, it is then called *stratigraphic interpretation*. In this article, we focus on structural interpretation.

Typically, interpreters first assume a geological model based on the analysis of various attributes of the data, in addition to the geological history of the region. They then manually segment the data volume into subvolumes based on the dominant structures contained within each subvolume. Thereafter, interpretation is performed to delineate various subsurface structures of interest. Based on the results, the geological model is

modified, and the process is repeated until the interpreters converge on a geologically plausible and reasonably accurate model. Such process is increasingly becoming more time-consuming and labor intensive with the explosive growth in data. Seismic surveys have grown over the years in both complexity and data sizes. Data sizes from such surveys have increased from hundreds of megabytes of the first three-dimensional (3-D) seismic survey collected in the 1970s to thousands of gigabytes, or even terabytes, today. To tackle this challenging situation, automated or semiautomated interpretation through computational algorithms has been investigated.

For years, image processing theories and algorithms have been employed in seismic interpretation and made essential contributions to the field. Given the richness and fast progress in image processing as well as related areas of computer vision and machine learning, there remains significant room to explore their impact on achieving a higher level of automation in seismic interpretation. Equally important, we believe that seismic interpretation is a challenging problem that continues to have an impact on the advancement of image processing and related fields.

Although a deep understanding of and training in geophysics have been considered prerequisites for developing effective computational interpretation algorithms, the recent progress in machine learning has shined a new light toward the roles that this field can play in domain-specific problems. We believe that, in essence, seismic interpretation fits very well into a generalized

Given the richness and fast progress in image processing as well as related areas of computer vision and machine learning, there remains significant room to explore their impact on achieving a higher level of automation in seismic interpretation.

pipeline for knowledge discovery through imaging applications, in a similar way to that of natural scene analysis and medical imaging analysis. This is illustrated in Figure 1, where all of these processes share a similar pipeline that consists of modules for acquisition, processing, and analysis. For the seismic case, the acquisition is fulfilled by sensing devices such as geophones for land surveys, which collect seismic waves reflected by subsurface structures, caused by a controlled seismic source of energy (e.g., a seismic vibrator). The acquired signals then go through advanced signal- and wave-equation-based processing (normally prestack migration) to yield

data for the stage of analysis, where seismic interpretation takes place. Such a framework aims at transforming a certain physical phenomenon into signals that can be analyzed to understand that phenomenon. The differences lie in the methodologies within each module and the underlying designs that accommodate the physical principles behind the phenomena being observed, which impose certain constraints on the properties of the signals.

The focus of this article is to present the seismic interpretation in the framework of the generalized pipeline for imaging data analysis, where both human-visual-system (HVS)- and learning-based models have the potential to assist in discovering underlying structures that may not have been discovered otherwise. The seismic interpretation pipeline can be redesigned as a labeling problem, modeling after the natural scene labeling problem. Moreover, interpretation algorithms can be designed based




FIGURE 1. An analogy between seismic interpretation, natural scene analysis, and medical imaging analysis. MRI: magnetic resonance imaging. CT: computed tomography.

on human visual attention models, mimicking the process of a human interpreter perceiving and analyzing a large volume of a seismic data set. We hope that this article will make the challenges found in seismic interpretation more accessible to the signal processing community. The start-up tools and data sets provided here will enable the interested readers to get a jump-start in this domain, especially if they do not have a background in geophysics. Furthermore, analyzing the seismic volume from the HVS point of view can unveil excellent opportunities that may not exist otherwise. Finally, solving challenges that are rooted in the unique characteristics of seismic data will help advance the corresponding signal processing and machine-learning theories and algorithms as well.

Subsurface structures and data sets

Before proceeding to the review of common interpretation tasks, in this section, we provide a brief introduction to several subsurface structures of interest, followed by a list of data sets that are commonly used for developing and testing the algorithms. Throughout this article, we will use oil and gas exploration as the example to illustrate subsurface understanding through seismic interpretation for a number of reasons. Compared to other applications, data is more abundantly available with oil and gas exploration. Also, there is rich literature on this subject with various contributions that can provide a benchmark. Furthermore, the depth of the imaged subsurface is deeper in the application of consideration, and that adds to the complexity of the structure to be interpreted. Therefore, the structures and data sets introduced here are typically related to hydrocarbon reservoir identification and characterization, although some of them can also be used

for other applications such as earthquake monitoring and environmental engineering, which is a common practice in the community. Nevertheless, all reviewed and proposed algorithms are applicable in all applications that image the subsurface to locate certain structures.

Common subsurface structures

Subsurface structures are complex because of the massive geologic evolution and deformation over millions of years. A migrated seismic volume, therefore, can contain multitudes of geologic structures such as horizons, unconformities, faults, salt domes, channels, and gas chimneys. Horizons, represented as seismic reflections, are the dominant geologic structures apparent in a seismic volume while the other structures can often be recognized as discontinuities or edges of seismic reflections. We will focus the discussion in this article on the interpretation of faults, salt domes, channels, and gas chimneys. Examples are shown in Figure 2. All of these structures are of great geological implications for hydrocarbon migration and accumulation as discussed next.

Faults

A fault is defined as a lineament or planar surface across which apparent relative displacement occurs in the rocks' layers. The movement of impermeable rocks and sediments along the fault surface creates membranes that hinder the migration of hydrocarbons from source rocks and create structural hydrocarbon traps. Because of the lateral changes in texture across a fault, interpreting such structures can be treated as an edge detection problem that is common in digital image processing. However,

different from edges in natural images, faults do not always display explicit edges, and the visual appearance is typically noisy.

Salt domes

A salt dome is defined as a dome-shaped structure formed by the evaporation of a large mass of salt in sedimentary rocks. Salt domes are impermeable structures that prevent the migration of hydrocarbons and provide entrapment for oil and gas reservoirs. Chaotic reflections are often observed in a salt dome in a form of a distinct texture; thereby interpreting such structures can be treated as a texture segmentation problem in digital image processing. Nevertheless, in salt domes, the boundaries are not explicit due to the underlying physics as well as noise and low-resolution data.

Channels

A channel is defined as a remnant of an inactive river or a stream channel that has been either filled or buried by




FIGURE 2. Examples of four types of subsurface structures essential for hydrocarbon exploration.

younger sediment. Channels, as well as the associated depositional features such as overbanks, lobes, and fans, are important targets of seismic interpretation because river/stream flows often carry sands of high porosity and permeability that are superior traps for hydrocarbon accumulation. Such rock difference gives channels a unique, distinct texture, similar to salt domes; meanwhile, considering its meandering spatial distribution, interpreting such structures can be treated as a curve boundary detection and tracking in digital image processing.

Gas chimneys

A gas chimney is defined as the leakage of gas from a poorly sealed hydrocarbon accumulation in the subsurface. Gas chimneys are often used as a robust hydrocarbon indicator, which implies present or previous vertical migration of hydrocarbons or fluids containing hydrocarbons between different geologic sequences. In a seismic profile, a gas chimney is visible as a vertical zone of poor data quality or push-downs. Therefore, a gas chimney is often featured with a distinct texture. Correspondingly, similar to salt domes, interpreting such structures can be treated as a texture segmentation problem in digital image processing.

Common geophysical data sets

High-quality seismic data are essential for accurate subsurface interpretation with broad applications to both industrial applications (e.g., oil and gas exploration) and environmental studies (e.g., earthquake monitoring and CO₂ storage). Driven by the oil and gas industry, 3-D seismic data have been widely collected around the world in the past few decades, particularly in the areas with high hydrocarbon potential, such as the North Sea, the Middle East, and the Gulf of Mexico. Such data were often confidential during the initial stage of reservoir exploration and production but are then released for public use when the production goal has been achieved. Table 1 lists seven seismic data sets that are commonly used for developing and testing interpretation algorithms. The table includes the geographical location of the surveys, the

dominant subsurface structures, the data size, as well as the spatial dimensions. The first six data sets are available in the typical SEG-Y format, which can be loaded directly using any geophysical exploration software (e.g., OpendTect and Petrel) or popular programming packages such as MATLAB and Python. Large North Sea Data Set of Migrated Aggregated Seismic Structures (LANDMASS) is available in the MATLAB MAT-File format and provides a large number of image patches generated from the F3 block to facilitate structure-based seismic interpretation and machine-learning studies. All links to these data sets and other resources can be found at <https://ghassanalregib.com/datasets-and-resources-for-computational-seismic-interpretation/>.

Subsurface event detection and tracking

A successful reservoir exploration requires a reliable identification of indicative subsurface structures as introduced in the previous section. Consequently, the majority of existing interpretation algorithms focus on detecting such structures or events in geophysical terms. We have recently proposed an additional task that tracks such structures by creating a semiautomated interpretation workflow. In this section, we review these methods for each of the four key structures, respectively. A depiction of all such algorithms is shown in Figure 3.

Fault detection

In a two-dimensional (2-D) seismic section, faults indicate displacements along fractures, as observed in Figure 2. Because of the geological constraints that are associated with the formation process, faults have two specific features. One is the geological feature, which is the discontinuity along horizons. The other is the geometric feature, i.e., line-like or curved shapes in 2-D seismic sections, which appear as curved surfaces in a 3-D seismic volume. Computational fault detection methods are commonly developed based on these two features. The discontinuity of faults can be characterized by several seismic attributes such as entropy [2], curvature [3], and coherence [4]–[6]. Among them,

Table 1. List of commonly used data sets for interpretation [1].

Data set and Location	Dominant Structures	Size	Domain	Dimensions
F3 block in The Netherlands, North Sea	Salt domes (~1,500 ms); faults (~1,200 ms); gas chimneys (~500 ms)	494 MB	Time	Inline: 651, Crossline: 951, Samples/trace: 463
Stratton field in south Texas, United States	Channels (~845 ms); faults (~2,000 ms)	122 MB	Time	Inline: 100, Crossline: 200, Samples/trace: 151
Teapot Dome in Wyoming, United States	Faults (~5,500 ft)	421 MB	Depth	Inline: 345, Crossline: 188, Samples/trace: 1,601
Great South Basin in New Zealand	Faults (~2,000 ms)	38.1 GB	Time	Inline: 1241, Crossline: 2,780, Samples/trace: 751
Waka Basin in New Zealand	Channels (~1,800 ms)	23.2 GB	Time	Inline: 801, Crossline: 5,756, Samples/trace: 1500
SEAM synthetic data set	Salt domes (~2,000 ms)	4.2 GB	Time and depth	Inline: 1169, Crossline: 1,002, Samples/trace: 851
LANDMASS data set from F3 block	Horizons, chaotic reflections, faults and salt domes	0.95 GB	Time	Horizon patches: 9,385, chaotic patches: 5,140, fault patches: 1,251, salt-dome patches: 1,891




FIGURE 3. Subsurface structures and their corresponding interpretation methods. GLCM: gray-level co-occurrence matrix.

coherence is the most popular one for highlighting faults. Marfurt et al. [4] calculated coherence by comparing the dissimilarity of local regions on the two sides of a fault. Later, Gersztenkorn and Marfurt [5] proposed the eigenstructure-based coherence, which analyzes the eigenstructure of covariance matrices of windowed seismic traces. Recently, a generalized-tensor-based coherence (GTC) attribute [7] has been proposed. The GTC of every voxel in a seismic volume is calculated within a local analysis cube with the size of $I_1 \times I_2 \times I_3$, which can be represented by a third-order tensor, denoted by $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$. Unfolding \mathcal{A} along three modes generates matrix $\mathbf{A}_{(i)} \in \mathbb{R}^{I_i \times \prod_{j \neq i} I_j}$, $i = 1, 2, 3$. Using an eigenstructure analysis on the covariance matrix of unfolding matrix $\mathbf{A}_{(i)}$, the coherence attributes of three modes are obtained as follows:

$$E_c^{(i)} = \lambda_1^{(i)} / \text{Tr}(\mathbf{C}_i) = \lambda_1^{(i)} / \text{Tr}[(\mathbf{A}_{(i)} - \mathbf{1}_i \otimes \boldsymbol{\mu}_i)^T (\mathbf{A}_{(i)} - \mathbf{1}_i \otimes \boldsymbol{\mu}_i)], \quad (1)$$

where $E_c^{(i)}$ is the coherence attribute of mode- i , and $\lambda_1^{(i)}$ represents the largest eigenvalue of covariance matrix \mathbf{C}_i . In addition, $\mathbf{1}_i$ is a matrix of ones with a size of $1 \times \prod_{j \neq i} I_j$, $\boldsymbol{\mu}$ defines the mean of all columns in $\mathbf{A}_{(i)}$, and \otimes represents the Kronecker product. If we take $E_c^{(i)}$, $i = 1, 2, 3$, as the R, G, and B channels of a color image, we arrive at the GTC attribute. In contrast to [5], GTC enhances the details of seismic data and allows a

better fine-tuning and more flexibility to interpreters. In some cases, depending on one seismic attribute does not produce accurate fault delineation. Therefore, enhancement operations such as nonlinear mapping and structure-oriented filtering [8] may be applied to increase the contrast between faults and surrounding structures. In addition, combining various attributes [9] is another common practice in fault interpretation.

In the attribute space, likely fault regions are commonly specified using a hard threshold or selected from local maxima. In the work of Hale [10], faults points with the largest fault likelihoods are selected and connected to construct meshes of quadrilaterals, which are further used to form fault surfaces. However, likely fault regions that may inevitably involve noisy structures are not able to accurately reveal the details of faults. To solve this problem, interpreters utilize the geometric feature of faults. Because of the continuous curved shapes of faults, ant tracking or ant colony optimization algorithm [11] has become a popular method for fault detection. Recently, Wu and Hale [12] proposed to use a simple linked data structure that includes the fault likelihood, dip, and strike to construct complete fault surfaces. In addition, Hough transform, a mapping procedure between attribute volumes and the parameter space, has been widely used to detect lines, circles, planes, and other parametric shapes. Therefore, Wang and AlRegib [13] proposed to delineate faults using line-like features extracted by the Hough transform. Figure 4 illustrates the diagram of the proposed method, in which authors




FIGURE 4. The diagram of the proposed fault detection method, which utilizes the Hough transform to extract fault features and remove false ones using geological constraints.

first highlight likely fault points by applying a hard threshold on the discontinuity map and then extract line-like yellow fault features using the Hough transform. Because of the limitation of the Hough transform, it is inevitable that detected results contain some false features that violate certain geological constraints. On the basis of their appearance, false features are classified into two types: outliers, which are isolated from the others, and neighboring groups, which contain similar features located nearby.

In Figure 4, the red line is fitted from all midpoints of fault features. Therefore, features with larger distances to the fitted fault are treated as outliers and discarded. In contrast, features in a neighboring group will be merged into one fault feature. After false feature removal, the remaining features are connected using blue lines to implement the delineation of faults. In [13], some postprocessing steps are implemented to improve the accuracy of fault delineation. The same concept of Hough transform has also been applied to extract fault planes from discontinuity volumes [14].

Salt-dome detection

To detect salt domes, over the past a few decades, researchers have proposed various seismic attributes and automated workflows based on edge detection, texture, machine learning, graph theory, active contours, and other techniques. Aqrabi et al. [15] proposed a Sobel edge-detection-based approach that delineates salt domes within the Gulf of Mexico data set. Amin and Deriche [16] proposed an approach that highlights small variations in seismic data by detecting edges not only along the x , y , and z directions but also slanted at 45° and -45° . The 2-D Sobel filters along 45° and -45° are shown as follows:

$$S_{45^\circ} = \begin{bmatrix} -2 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix}, \quad S_{-45^\circ} = \begin{bmatrix} 0 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \end{bmatrix}. \quad (2)$$

Berthelot et al. [17] proposed a Bayesian classification approach for detecting salt bodies using a combination of seismic attributes such as dip, similarity, frequency-based attributes, and attributes based on the gray-level co-occurrence matrix (GLCM). A workflow based on seismic attributes and a machine-learning algorithm (i.e., extremely random trees ensemble) that automatically detects salt domes from the Society of Exploration Geophysicists (SEG) Advanced Modeling (SEAM) data set is presented in [18]. Amin and Deriche [19] proposed a supervised codebook-based learning model for salt-dome detection using texture-based attributes. Di et al. [20] proposed an interpreter-assisted approach based on the k-means clustering of multiple seismic attributes that highlights salt-dome boundaries in the F3 block. Shafiq et al. [21] proposed a seismic attribute called *SalSi*, which highlights salient areas of a seismic image by comparing local spectral features based on the 3-D fast Fourier transform. Chopra and Marfurt [22] proposed a seismic disorder attribute to assess randomness and the signal-to-noise ratio (SNR) of data to delineate seismic structures such as faults and salt domes. Wu [23] proposed methods to compute salt likelihoods highlighting salt boundaries, extract oriented salt samples on ridges of salt likelihoods, and construct salt boundaries with the salt samples by solving a screened Poisson surface reconstruction problem.

With most of the aforementioned algorithms, the underlying assumption is that a salt-dome body can be recognized by comparing neighboring planes. Thus, capturing the textures of such planes has been a very effective methodology to delineate salt-dome bodies. To show the effectiveness of texture-based approaches for detection of salt domes, the brief details of a seismic attribute, 3-D gradient of textures (3D-GoT) by Shafiq et al. [24], which describes the texture dissimilarity between neighboring cubes around each voxel in a seismic volume across time (or depth), crossline, and inline directions is presented in this section.




FIGURE 5. (a) Computing GoT along the crossline, i.e., x -direction. (b) Salt-dome boundary detected by various delineation algorithms; magenta [15], black [19], blue [24], green: reference salt-dome boundary. (c) A 3-D salt body detected from the F3 block in the North Sea using 3D-GoT [24].

Figure 5(a) illustrates synthetic seismic images in which a green-dashed vertical line separates two textured regions depicted in dotted and striped lines, respectively. To evaluate the GoT in the x -direction (i.e., the crossline), as the center point and its two neighboring cubes move along the blue line, texture dissimilarity function $d(\cdot)$ yields a GoT profile shown at the bottom of Figure 5(a). Theoretically, the highest GoT value corresponds to the highest dissimilarity and is obtained when the center point falls exactly on the texture boundary. Similarly, GoT is also calculated along t - (time or depth) and y - (inline) directions. To improve the delineation efficiency and robustness, GoT employs a multi-scale gradient expressed as follows:

$$\mathbf{G}[t, x, y] = \left(\sum_{i \in \{t, x, y\}} \left(\sum_{n=1}^N \omega_n \cdot d(\mathbf{W}_{i-}^n, \mathbf{W}_{i+}^n) \right)^2 \right)^{\frac{1}{2}}, \quad (3)$$

where \mathbf{W}_{i-}^n and \mathbf{W}_{i+}^n denote the neighboring cubes, n represents the edge length of cubes, and ω_n represents the weight associated with each cube size. To compute dissimilarity between cubes, authors use a perceptual dissimilarity measure based on error magnitude spectrum chaos, which is not only computationally less expensive and performs better than nonperceptual dissimilarity measures, but also highlights texture variations in the most effective manner. The perceptual dissimilarity measure is calculated as follows:

$$d(\cdot) = E \left(\left\| \mathcal{K} \otimes \left(\mathcal{K} \otimes \left| \mathbf{W}_{i-}^n - \mathbf{W}_{i+}^n \right| \right) \right\| \right), \quad i \in \{t, x, y\}, \quad (4)$$

where \otimes represents the tensor product, \mathcal{K} is the Kronecker matrix defined as $\mathcal{K} = D_t \otimes D_x \otimes D_y$, and D_t , D_x , and D_y are

discrete Fourier transform matrices. The output of various salt-dome delineation algorithms on a typical seismic inline and a salt body detected using 3D-GoT [24] from the F3 block in the North Sea are shown in Figure 5(b) and (c), respectively.

Fault and salt-dome tracking

The methods introduced in the previous two sections focus mainly on the detection, or delineation, of faults and salt domes in 2-D sections. To investigate the geological structures of faults and salt domes, interpreters need to repeatedly apply these methods on each section of a seismic volume. However, for a seismic volume with a large size, the repeated detection in every section may impair interpretation efficiency. Because of the slow formation processes of subsurface structures, neighboring sections commonly have strong correlations. In recent years, fault and salt-dome tracking methods have been proposed to utilize correlations between sections to improve interpretation efficiency. Reference [13] borrows the concept of motion vectors in video coding and grouped seismic sections into reference and predicted sections. Faults in the predicted sections can be labeled by detection results in reference sections. Berthelot et al. [25] detected the salt-dome boundary in one time section using the method in [17] and tracked the boundary through adjacent sections by minimizing a defined energy function that maintains boundaries' curvature and smoothness. Similarly, [26] takes advantage of active contour to track salt-dome boundaries in neighboring seismic sections. More recently, Wang et al. [27] have proposed a salt-dome tracking method that extracts the features of salt-dome boundaries in reference sections using tensor-based subspace learning and delineates tracked boundaries by finding points in predicted sections, which are the most similar to reference




FIGURE 6. The diagram of salt-dome tracking, in which features around salt-dome boundaries in reference sections are extracted by tensor-based subspace learning.

ones. Figure 6 illustrates the diagram of feature extraction process using tensor-based subspace learning adopted for salt-dome tracking. Authors manually label the boundaries of salt domes in N_r reference sections, in which $N_c = \lceil N_r/2 \rceil$ is the central one. For each point in the labeled boundary of N_c , its corresponding points in neighboring $(N_r - 1)$ reference boundaries are identified. Since every point corresponds to a patch with the size of $N_p \times N_p$, stacking corresponding patches constructs third-order texture tensors from reference sections, denoted $\{\mathcal{A}_k \in \mathbb{R}^{N_p \times N_p \times N_r}, k = 1, 2, \dots, K_{N_c}\}$, where K_{N_c} represents the number of points in the N_c th reference boundary. Because of strong correlations between neighbors, every point in the N_c th reference boundary corresponds to a tensor group, denoted $\mathcal{G}_k = \{\mathcal{A}_{k-N_s}, \dots, \mathcal{A}_k, \dots, \mathcal{A}_{k+N_s}\}$, which contains $2N_s + 1$ tensors. Using multilinear dimensionality reduction methods such as multilinear principal component analysis and tensor-based linearity projection preserving, one can obtain transformation matrices $\mathbf{U}_k^{(1)}$, $\mathbf{U}_k^{(2)}$, and $\mathbf{U}_k^{(3)}$ and map tensor group \mathcal{G}_k to its subspace as follows:

$$\begin{aligned} \tilde{\mathcal{A}}_m &= \mathcal{A}_m \times_1 \mathbf{U}_k^{(1)} \times_2 \mathbf{U}_k^{(2)} \times_3 \mathbf{U}_k^{(3)}, \\ m &\in \{k - N_s, \dots, k + N_s\}, \end{aligned} \quad (5)$$

where $\times_i, i = 1, 2, 3$ denotes the i -mode product of a tensor by a matrix, and tensor $\tilde{\mathcal{A}}_m$ with fewer dimensions in each mode

contains the features of reference boundaries. These extracted features can help identify tracked boundary points in predicted sections. The local and global comparison of the tracked boundary (green) with the manually labeled one (red) shows the high similarity and accuracy of salt-dome tracking within seismic volumes.

Channel detection

The variety of rocks that are deposited in a channel system change the acquired seismic signals, and, thereby, a channel system can be detected by applying the common edge detection methods on horizontal slices or well-picked horizon surfaces. For example, the coherence attribute [28], [29] was first proposed for identifying the morphology of a channel system, and then a suite of new attributes and algorithms have been developed for enhancing both detection accuracy and noise robustness. However, while successfully highlighting the channel boundaries, such discontinuity analysis fails to depict the rock properties within the channel system, as no such information is considered by the detection operators. To overcome this limitation, the GLCM analysis was introduced from the field of image processing and used for seismic facies analysis [30]. Eichkitz et al. [31] presented a suite of algorithms for generating ten different GLCM attributes from a seismic cube. Di and Gao [32] made it possible to decompose a channel system by integrating the logistic/exponent gray-level transformations with the GLCM attributes. Machine-learning-based

facies analysis has gained popularity, particularly for a deposition system with complexities, and automatically classified all subtle structures, including the meandering channel, overbanks, fans, and lobes. The use of computational classification began soon after the development of seismic attributes in the 1970s with the work by Justice et al. [33]. Barnes and Laughlin [34] reviewed several unsupervised classification techniques, including k-means, fuzzy clustering, and self-organizing maps (SOMs) and emphasized the importance of seismic attributes over the classifiers. Wallet et al. [35] developed the generative topographic mapping for unsupervised waveform classification. Song et al. [36] combined multilinear subspace learning with the SOMs for improved seismic facies analysis in the presence of noise. A comprehensive study of both supervised and unsupervised facies analysis can be found in [37]. With these depositional features in a channel system well differentiated by either seismic attributes or facies analysis, they could then easily be extracted as separate geobodies by seeded tracking. However, reliable differentiation of various depositional features (e.g., overbank, delta, and levee) remains challenging, due to the proximity and overlying distribution in space between each other and, more importantly, their similar reflection patterns in 3-D seismic data. Such a goal can be achieved by analyzing the seismic images at a smaller scale to capture the subtle differences between various features.

Gas-chimney detection

In seismic sections, a gas chimney is visible as vertical zones of poor data quality, chaotic reflections, or push-downs. Therefore, it can be detected using attributes similar to those used for salt-dome detection, such as the coherence [28] and the GoT [24]. The major difference of gas chimneys is the sparse distribution in a seismic volume, and thereby manual identification and interpretation of them is labor intensive, and the conventional tracking tools described in the section “Fault and Salt-Dome Tracking” may also fail in detecting the chimneys that are isolated from each other. However, the machine-learning-based approach offers an efficient solution to such limitation. For example, Heggland et al. [38] combined a set of seismic attributes and the multilayer perception (MLP) to create a chimney cube for a semiautomatic detection, which has been applied to multiple data sets, such as the F3 block [39] and the Taranaki basin in New Zealand [40]. In recent years, researchers also have tried more advanced machine-learning algorithms. For example, Xiong et al. [41] applied adaptive boosting (AdaBoost) to the design of the optimal learning algorithm for identifying gas chimneys, which generated more reliable results than the k-nearest neighbor method. Xu et al. [42] implemented the sparse autoencoder for gas-chimney detection, and the accuracy is greatly improved compared to the traditional MLP algorithm. The sparsity of the spatial distribution of gas chimneys in a seismic data set adds the difficulty of reliable

gas-chimney detection in two ways. First, it limits the amount of training data, and thereby the performance of supervised learning may be affected. Weakly supervised and unsupervised methods would be more applicable. Second, it increases the sensitivity to seismic noise. Reflection pattern-based learning could help improve the noise robustness of gas-chimney detection, such as the chaotic labeling described in the next section given the chaotic reflections in a gas chimney.

Subsurface labeling and classification

In the previous section, we provided a review of seismic interpretation for subsurface event detection and tracking. Although the techniques have significantly reduced the time and effort required for manual interpretation, there is at least one aspect involved that is still done manually, namely, the manual process of extracting subvolumes from a given data volume based on their dominant subsurface structure, so that detection or tracking can be performed on the extracted data. In this section, we discuss a framework we recently developed to address this issue. With this

framework, we attempt to eliminate the aforementioned bottleneck and streamline the interpretation process by building on the recent advances in semantic segmentation and scene labeling.

General framework

Seismic volume labeling is the process of classifying each voxel in a given seismic volume into one of many predefined structures. This process can help classify entire seismic volumes into regions of interest that contain specific subsurface structures. These regions can then be extracted, and various interpretation algorithms can be applied to these regions for more refined results.

While a variety of influential works have been proposed in the area of semantic segmentation [43]–[47], seismic data presents challenges that cannot be immediately solved by the existing methods. First, unlike natural images, where edges between objects are well defined, edges between subsurface structures in seismic data are either not well defined or are characterized by a change in overall texture rather than a sudden change in amplitude. Second, unlike natural images, seismic data is gray scale, and, thus, color features cannot be used to distinguish various structures. Furthermore, there is a severe lack of both labeled seismic data for training and well-established benchmarks for testing various learning-based approaches. This is partly due to the intellectual property concerns in the oil and gas industry. Also, because of the lack of ground truth and the subjective nature of manual interpretation, it is often difficult for different geophysicists to agree on the same interpretation for a given volume.

Naturally, machine-learning techniques are well suited approaches for seismic volume labeling. However, the lack of labeled training data poses a significant challenge. To tackle this

The focus of this article is to present the seismic interpretation in the framework of the generalized pipeline for imaging data analysis, where both human visual-system- and learning-based models have the potential to assist in discovering underlying structures that may not have been discovered otherwise.

problem, we have generated the LANDMASS [48] data set that contains more than 17,000 seismic images extracted from The Netherlands North Sea F3 block [49]. The images are grouped into four classes based on the subsurface structure they contain: horizons, chaotic horizons, faults, and salt domes. Although the data set contains these specific structures, the framework for seismic labeling discussed next can be extended to other seismic structures as well.

In [50], Alaudah and AlRegib proposed using reference exemplars and seismic image retrieval to label seismic volumes in a weakly supervised fashion. First, given a few hand-selected exemplar images, $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_e}]$, that contain subsurface structures belonging to the different classes of interest, an augmented data set, $\tilde{\mathbf{X}}$, is generated in an unsupervised fashion using similarity-based retrieval. This is done to obtain enough data to train a supervised machine-learning model. Then, various features or attributes are extracted from these images to train a classifier. The section “Oversegmentation of Seismic Volumes” describes seismic volume segmentation and its use to enforce the local spatial correlation of the labels and improve the computational efficiency of the approach. Finally, the sections “Labeling of Seismic Volumes” and “Weakly Supervised Pixel-Level Annotation” describe various methods to obtain the final labeling of the seismic volume. The overall framework is outlined in Figure 7.

Building blocks

In the remaining part of this section, we describe the major components of this weakly supervised approach for seismic volume labeling.

Texture attributes for seismic labeling

Seismic images are often well characterized by texture features, or texture attributes, mainly because they are textural in nature. A few classical texture attributes were explored for traditional tasks such as salt-dome detection, but they need to be further examined in the context of seismic labeling. In addition, there are a great number of advanced texture features developed in recent

years for image analysis. They are also potential candidates for the labeling task.

When applied to image processing problems, it is usually desirable for a texture attribute to possess properties such as illumination-, rotation-, and scale-invariance for better robustness.

However, with seismic data, this is not always the case. For example, a vertical slice of a seismic volume (or a seismic section) is characterized by strong directionality, with horizons typically extending in the horizontal direction and faults in the vertical direction. In such cases, being rotation-invariant is no longer a critical requirement for the attributes. Another important difference between a seismic image and a natural texture image is that some subsurface structures (e.g., faults) are of very fine-scale along certain dimensions,

which is not typical with natural textures. Thus, it is important for a texture attribute to be able to capture such fine details.

Recently, some comparative studies were conducted to examine various texture attributes in the context of seismic volume labeling. In one study, the focus was on a group of spatial attributes from the family of local descriptors, including the local binary pattern, a few of its typical variants, and the local radius index [51]. These attributes can represent texture patterns with robustness and computational efficiency. For comparison purposes, the study also included two traditional seismic attributes in the spatial domain, i.e., the GLCM and the semblance. According to the study, the local descriptors and the GLCM are all good attributes for labeling seismic volumes. However, each attribute displayed different characterizing capabilities for different subsurface structures. Thus, they should be selected accordingly if there is a preference for certain structures to be labeled with more reliability.

In a separate study [52], multiresolution attributes in the frequency domain were examined for seismic volume labeling, including the discrete wavelet transform and its nonsubsampling version, Gabor filters, the steerable pyramid, the contourlet transform and its nonsubsampling counterpart, and the curvelet transform. Effective singular values are extracted from each transformed subband and then concatenated into a feature vector.

The ultimate goal of seismic volume labeling is to efficiently and accurately classify entire seismic volumes based on their subsurface structures such as those shown in Figure 8.




FIGURE 7. A block diagram illustrating the weakly supervised seismic volume-labeling approach described in the section “Subsurface Labeling and Classification.”

The major conclusion from the study is that the directional transforms perform better than the nondirectional ones, with the curvelet transform being the best.

Similarity-based retrieval

The attributes described in the previous section are used in the feature extraction step of the workflow depicted in Figure 7. Also, these features can be used to devise similarity measures specialized for seismic images (e.g., [53]–[55]) to be used in the similarity-based retrieval to generate the augmented data, $\tilde{\mathbf{X}}$. Note that texture-based image similarity metrics are different from the generic fidelity and distance measures (also known as *quality metrics*, such as peak SNR and mean square error), in that they capture the content of an image rather than assuming a pixel-to-pixel correspondence.

One particularly useful similarity measure was proposed in [55] for content-based image retrieval, especially for images that are highly textural such as seismic images. The proposed measure uses the singular values of the curvelet transform of an image to form a feature vector. Then, the similarity of any two images is computed as the Czekanowski similarity between their corresponding feature vectors as $\text{SIMILARITY}(\mathbf{x}_1, \mathbf{x}_2) = 1 - (\|\tilde{\mathbf{v}}_1 - \tilde{\mathbf{v}}_2\|_1 / \|\tilde{\mathbf{v}}_1 + \tilde{\mathbf{v}}_2\|_1)$, where $\mathbf{x}_1, \mathbf{x}_2$ are the images and $\mathbf{v}_1, \mathbf{v}_2$ are their corresponding feature vectors. Since the singular values of real images are non-negative by definition, the value of this similarity measure is in the range of $[0, 1]$, where a value of one indicating the two images are identical.

With a proper similarity measure in hand and an exemplar image for each subsurface structure of interest, the training data is formed by retrieving images that are most similar to the exemplar image. The training data set is then used to train a machine-learning model as will be detailed in the section “Labeling of Seismic Volumes.”

To fully explore the potential of deep learning, there are a few open problems that need to be addressed carefully.

Oversegmentation of seismic volumes

An important part of the workflow in Figure 7 is the seismic volume oversegmentation. Segmentation algorithms, like normalized cuts [56], are often used in computational seismic interpretation to extract subsurface structures [57]–[59]. However, oversegmentation here is used as a preprocessing step to enforce local spatial correlation by grouping voxels in the volume that are similar and close to each other. Here, the lack of clearly defined edges between subsurface structures is inconsequential, and instead, each small segment is classified based on its texture content. This step also significantly reduces the computational cost of the labeling, since each segment will be classified once, and the resulting label will then be propagated to all the voxels within that segment.

Graph-based oversegmentation techniques (such as graph cuts [60] and normalized cuts [56]) or gradient ascent-based approaches (such as the watershed algorithm [61] and turbo pixels [62]) are computationally expensive, and therefore not suitable for large seismic volumes. Simpler, and more efficient techniques such as simple linear iterative clustering (SLIC) superpixels [63] are more appropriate for this application. Since SLIC clusters pixels in the CIELAB color space, while seismic images are gray scale, it is possible to obtain a good superpixel oversegmentation by replacing the A and B channels with the x - and y -gradients of the seismic image [50].

Once the oversegmentation is performed, features are extracted from each segment, or from a group of segments centered around a reference segment. These features are later used to classify the voxels within these segments into one of the various subsurface structures that are of interest to the interpreter.

Once the oversegmentation is performed, features are extracted from each segment, or from a group of segments centered around a reference segment. These features are later used to classify the voxels within these segments into one of the various subsurface structures that are of interest to the interpreter.

Labeling of seismic volumes

Once the augmented data $\tilde{\mathbf{X}}$ is created using the methods described in the section “Similarity-Based Retrieval,” texture features can then be extracted from these images to train various machine-learning models for labeling the seismic volume. These models can be support vector machines, convolutional neural networks (CNNs), random forests, or any other machine-learning model suitable for this task. Once these models are trained, we can then extract texture attributes from each segment in the oversegmented seismic volume and classify them. The resulting label is then propagated to all voxels within the segment.

The ultimate goal of seismic volume labeling is to efficiently and accurately classify entire seismic volumes based on their subsurface structures such as those shown in Figure 8. There will naturally be a tradeoff between the accuracy and the efficiency of any algorithm. However, it is better to obtain a rough labeling quickly, rather than an accurate labeling that is time-consuming. Other specialized techniques would later be used for the refined detection and tracking of the subsurface structures.

Weakly supervised pixel-level annotation

In the section “Similarity-Based Retrieval,” we have shown that similarity-based retrieval can be used to extract thousands




FIGURE 8. (a) A weakly supervised labeled seismic section from The Netherlands North Sea F3 block. For reference, the manually labeled seismic section is shown in (b). The chaotic class is in blue, faults are in green, and the salt dome is in red.

of images that have the same subsurface structure to generate our augmented training data $\tilde{\mathbf{X}}$. The retrieved images are then assigned the same image-level label as the query image. While these image-level annotations can be used to train machine-learning models, recent work on weakly supervised labeling such as that proposed in [64] show that these image-level labels can be mapped into pixel-level labels that can be much more useful in training powerful fully supervised deep-learning models such as fully convolutional networks [44]. Figure 9 shows several examples of this mapping for various subsurface structures.

The approach in [64] is based on nonnegative matrix factorization (NMF) [65]. NMF is a commonly used matrix factorization technique that is closely related to many unsupervised machine-learning techniques such as k -means and spectral clustering [66]. NMF decomposes a nonnegative matrix $\tilde{\mathbf{X}} \in \mathbb{R}_+^{N_p \times N_s}$ into the product of two lower-rank matrices $\mathbf{W} \in \mathbb{R}_+^{N_p \times N_f}$, and $\mathbf{H} \in \mathbb{R}_+^{N_f \times N_s}$ such that both \mathbf{W} and \mathbf{H} are nonnegative. In other words,

$$\tilde{\mathbf{X}} \approx \mathbf{W}\mathbf{H} \quad \text{s.t. } \mathbf{W}, \mathbf{H} \geq 0. \quad (6)$$

Here, N_f is the number of components (or the *rank*) of $\tilde{\mathbf{X}}$. In our work, $\tilde{\mathbf{X}}$ represents the augmented data matrix from the section “Similarity-Based Retrieval,” where each column is a single seismic image in vector form. The data matrix $\tilde{\mathbf{X}}$ has N_s such images, each of which is a vector of length N_p . NMF factorizes this data matrix into two nonnegative matrices: a basis matrix \mathbf{W} and a coefficient matrix \mathbf{H} .

The regular NMF problem does not have an analytical solution, and is typically formulated as the following nonconvex optimization problem:

$$\arg \min_{\mathbf{W}, \mathbf{H}} \|\tilde{\mathbf{X}} - \mathbf{W}\mathbf{H}\|_F^2 \quad \text{s.t. } \mathbf{W}\mathbf{H} \geq 0, \quad (7)$$

where, $\|\cdot\|_F$ is the Frobenius norm. Lee et al. [67] showed that NMF can be used to learn a “parts-based” representation, where each feature would represent a localized “part” of the data. In practice, this is rarely achieved using the formulation in (7). To remedy this, the feature matrix \mathbf{W} is initialized using k -means applied on each class in the data matrix $\tilde{\mathbf{X}}$ separately. This initialization simplifies the feature learning and makes \mathbf{W} robust to mislabeled images in $\tilde{\mathbf{X}}$. Then, a sparsity constraint is imposed on these initial features using the following sparsity measure:

$$\rho(\mathbf{w}) = \frac{\sqrt{N_p} - \|\mathbf{w}\|_1 / \|\mathbf{w}\|_2}{\sqrt{N_p} - 1}, \quad (8)$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ are the l_1 and l_2 norms, respectively, and $\rho(\cdot)$ indicates the sparsity of a vector. To enforce this constraint, we follow the algorithm proposed by [68]. Additionally, to make sure that each feature \mathbf{w}_i in the matrix \mathbf{W} represents a single class only, we impose an orthogonality constraint on the coefficients matrix \mathbf{H} . We also add two regularization terms on \mathbf{W} and \mathbf{H} to avoid overfitting. The problem then becomes




FIGURE 9. Results of the weakly supervised pixel-level annotation approach in [64] for various subsurface structures. (a) shows images containing chaotic and fault structures and (b) shows their corresponding pixel-level labels in blue and green respectively. (c) and (d) show various images that contain salt-dome bodies or boundaries, and their corresponding pixel-level labels in red, respectively.

$$\begin{aligned} \arg \min_{\mathbf{W}, \mathbf{H}} & \|\tilde{\mathbf{X}} - \mathbf{W}\mathbf{H}\|_F^2 + \lambda_1 \|\mathbf{W}\|_F^2 + \lambda_2 \|\mathbf{H}\|_F^2 + \gamma_1 \|\mathbf{H}\mathbf{H}^T - \mathbf{B}\|_F^2 \\ \text{s.t. } & \mathbf{W}, \mathbf{H} \geq 0 \text{ and } \rho(\mathbf{w}_i) = \rho_w, \end{aligned} \quad (9)$$

where matrix $\mathbf{B} \in \mathbb{R}^{N_f \times N_f}$ contains random positive real numbers, and λ_1 , λ_2 , γ_1 , and ρ_w are constants. To solve the problem in (9), the following multiplicative update rules for \mathbf{W} , and \mathbf{H} are derived, where

$$\mathbf{W}^{t+1} = \frac{\mathbf{W}^t \odot (\tilde{\mathbf{X}}\mathbf{H}^{tT})_{ij}}{(\mathbf{W}^t\mathbf{H}^t\mathbf{H}^{tT} + \lambda_1\mathbf{W}^t)_{ij}} \quad (10)$$

$$\text{and } \mathbf{H}^{t+1} = \frac{\mathbf{H}^t \odot (\mathbf{W}^{t+1T}\tilde{\mathbf{X}} + \gamma_1(\mathbf{B} + \mathbf{B}^T)\mathbf{H}^t)_{ij}}{(\mathbf{W}^{t+1T}\mathbf{W}^{t+1}\mathbf{H}^t + \gamma_1\mathbf{H}^t\mathbf{H}^{tT}\mathbf{H}^t + \lambda_2\mathbf{H}^t)_{ij}}. \quad (11)$$

Here, \odot represents element-wise multiplication, and the superscript indicates the iteration number. Once \mathbf{W} and \mathbf{H} are initialized, the multiplicative update rules in (10) and (11) are applied successively until both \mathbf{W} and \mathbf{H} converge.

Once \mathbf{W} and \mathbf{H} have converged, each column of \mathbf{H} and \mathbf{h}_n , indicates the features used to construct the n th image. Since every feature in \mathbf{W} should correspond to a single seismic structure, the coefficients of each image can then be mapped into the seismic structures that make up that image. Thus, for every image $n \in [1, N_s]$ we can obtain

$$\mathbf{L}_n = \mathbf{W}^{\text{final}}(\mathbf{Q} \odot (\mathbf{h}_n^{\text{final}} \mathbf{1}_{1 \times N_f})) \quad \forall n \in [1, N_s], \quad (12)$$

where $\mathbf{Q} \in \{0, 1\}^{N_f \times N_f}$ is a cluster membership matrix such that the element $Q_{ij} = 1$ if the feature \mathbf{w}_i belongs to structure j , and $\mathbf{1}_{1 \times N_f}$ is a vector of ones of size $1 \times N_f$. The resulting matrix, $\mathbf{L}_n \in \mathbb{R}_+^{N_f \times N_f}$ shows the likelihood of each seismic structure for each pixel in the image. Then, the pixel-level label for each location i in image n corresponds to the seismic structure given by




FIGURE 10. An illustration of implementing CNNs for salt-body boundary delineation from the poststack seismic amplitude. The detected boundaries are clipped to six vertical sections for quality control. Note the good match between the CNN detection and the poststack seismic images

$$\tilde{\mathbf{y}}_n = \arg \max_j \mathbf{L}_n(:, j). \quad (13)$$

Emerging trends and open problems

From the recent advancement in seismic interpretation research as we have previously reviewed, we observe two important factors that contribute to the success of this endeavor. First, to address the challenges rooted in the ever-increasing data size and complexity, it becomes very critical to leverage the advanced machine-learning techniques, especially those based on deep learning. Second, being a unique type of visual signals, seismic volumes can be interpreted effectively using image analysis algorithms that utilize HVS characteristics and models. In this section, we will briefly discuss the emerging trends and open problems regarding these two aspects.

Deep subsurface learning

Deep learning is one of the most powerful learning techniques available today. As a data-driven approach, it utilizes sophisticated neural networks with deep architectures to uncover complex hidden structures and characteristics directly from a large amount of samples. When applied to subsurface data, deep learning will allow geoscientists to make sense out of the massive data sets with many variables while avoiding the human biases. Naturally, seismic interpretation based on deep learning is emerging as a very promising trend. For example, Waldeland and Solberg [69] applied CNNs to classify salt bodies from 3-D seismic data sets. Huang et al. [70] provided an excellent demonstration of the effectiveness of integrating CNNs and multiattribute analysis from poststack amplitude in detecting faults. An illustration of implementing CNNs for salt-body boundary delineation is given in Figure 10. In this case, a good match is observed between the detected boundaries and the original seismic images. Besides the conventional poststack data used for interpretation, geophysicists are turning their attention to the prestack data. For example, Hami-Eddine et al. [71] investigated a machine-learning approach to optimize the use of both prestack and poststack seismic data. Araya-Polo et al. [72] and Lin et al. [73] proposed using a deep neural network to learn a mapping relationship between the raw seismic data and the subsurface geology so that the labor-intensive processing stage could be avoided.

To fully explore the potential of deep learning, there are a few open problems that need to be addressed carefully. First, the lack of labeled data is a serious obstacle. Unlike natural image classification problems, public domain data sets with large sets of labeled samples are rare for seismic interpretation. This severely limits the application of supervised learning techniques. Alternatively, weakly supervised or unsupervised learning will be more realistic choices in this case. Other techniques such as generative adversarial networks (GANs) and active learning, which will be discussed shortly, can also provide an alternative approach to supervised learning.

Second, more advanced network architectures need to be explored for interpretation. Current works mainly use CNNs as the core deep-learning structure. To account for the strong

correlations along different dimensions in a seismic volume, networks such as the long short-term memory [74] and the neural Turing machines [75] can be incorporated.

Third, as a data-driven approach, deep learning relies heavily on the provided data. Thus, a concern with deep learning is about falsely labeled data. Such samples will be misleading in the training of the network. To tackle this problem, various regularization techniques such as dropout can be adopted. If applicable, weakly supervised or unsupervised learning should be adopted.

Finally, seismic imaging and processing that precede the interpretation will inevitably introduce errors or uncertainties. How do these errors propagate through the learning networks? How do they affect the learning performance? Moreover, in general, how do we quantify the uncertainty associated with these learning approaches? These are all issues of practical significance for seismic interpretation, and should be carefully addressed.

GANs

Among various deep-learning techniques, generative models are a set of models that have been quite popular recently. Generative models are systems that can learn and generalize the probability distribution of data from training samples. The outcome of such systems can be either the explicit probability distribution or samples from the learned models. Some models can be designed to do both tasks. One specific type of generative model that has gained popularity recently is the GAN. GANs are used to generate samples that come from the same distribution as some given training samples. The task is achieved by designing two separate systems that compete against each other. More specifically, GANs consist of 1) a generator that can transform random noise into an image that looks like the training samples, and 2) a discriminator that takes an image as input and estimates the probability that this image is from the same distribution as that of the training samples. After proper training, the generator will be able to produce images that resemble images from the training samples.

GANs have been employed in various applications such as image synthesis, image interpolation and inpainting, style transfer, and next-frame prediction [76]–[79]. More details about GANs and their useful applications can be found in [80]. However, in the context of seismic interpretation, we highlight three applications of GANs that are particularly useful. The first application is to use GANs to train a model with unlabeled or partially labeled data (e.g., [81]). A second application is model-based seismic interpolation using GANs to interpolate seismic traces or generate new seismic sections at superresolution. A third application is to utilize GANs for style transfer of seismic data sets, which are very diverse in terms of acquisition resolution, reflector strength, and size of structures. Due to the diversity of such data sets, an algorithm trained on one seismic data set will not necessarily work on other data sets. Using GANs, it is possible to transfer the

Saliency detection is one of the promising directions that can not only assist interpreters in the interpretation process but also help develop better attention models for seismic interpretation.

style of one data set along with its labels to another data set that can be used to improve the robustness of the algorithm on a different data set.

Figure 11 shows sample seismic images generated by a deep convolutional GAN (DCGAN) [78]. The samples are arranged such that each column contains images of the same class. The network was trained on the entire LANDMASS data set [48].

Active learning

In a typical interpretation process, an interpreter works in cycles of interpretation-examination to improve the results gradually. It is desirable if computational approaches can mimic this interactive and iterative process. Among all varieties of machine-learning techniques, active learning emerges as a perfect fit to fulfill this purpose. The inherent interactivity will involve a user to assist in making intermediate decisions, guiding the learning process to follow the correct path. More importantly, active learning will also help alleviate the problem of insufficient manually labeled data, as it adds samples to the set of labeled data during each iteration. A typical active learning process starts with a limited set of manually labeled data samples. Initial models are trained using the limited labeled data and then used to select a subset from a much larger set of unlabeled samples. The chosen samples are provided as queries to an interactive annotation procedure, where they are labeled by a human annotator (or a group of annotators to reduce the chance of a biased decision). Carefully designed computer algorithms may also be employed to perform the labeling in place of human annotators. Finally, the labeled queries are added to the existing labeled data set, with which a new iteration of learning can be performed. For active learning, the query selection is the key to the whole process and the most active area of current research.




FIGURE 11. Sample seismic images generated using DCGAN. Each column in the figure contains generated images that represent one of the four classes of LANDMASS data set, i.e., (a) horizons, (b) faults, (c) chaotic horizons, and (d) salt domes.

Attention models inspired by the HVS

In addition to machine learning, we believe the other important factor that will impact seismic interpretation research is the HVS modeling. For visual analysis, HVS reduces the amount of the sensory data information, also known as the *task-free visual search*, by focusing on the perceptually salient segments of visual data that conveys the most useful information about the scene. For example, as an interpreter traverses through a seismic volume or a video, the HVS that is sensitive to structures and variations in amplitude and surrounding environment registers the relative motion of seismic structures across all three dimensions, not all of which are characterized by strong seismic reflections or texture contrast. Therefore, computational algorithms inspired by HVS can be used to develop attention models that can effectively automate the structural interpretation of seismic volumes. Among many, one such attention model is defined as a *saliency* model, which attempts to predict the interesting areas in images and videos, typically called salient regions, by relying on low-level visual cues. In seismic interpretation, visual saliency is important to predict the attention of human interpreters and highlight areas of interest in seismic sections, which cannot only help in the automation of the interpretation process but also develop attention models that can be applied to new data sets. Recently, a seismic attribute called *SalSi* was proposed in [21] and [82], which highlights salient areas within seismic volumes to assist interpreters in structural interpretation of salt domes and faults. Furthermore, by incorporating a priori information into saliency detection, attention models can be used to mimic the behavior of interpreters looking at seismic sections. Therefore, saliency detection is one of the promising directions that can not only assist interpreters in the interpretation process but also help develop better attention models for seismic interpretation.

It is a challenging task to train a machine to understand and reveal various structures of Earth's subsurface. Yet, once achieved, it will be a truly rewarding accomplishment that has a large impact on the economy and society. Through its fast advancement in the past decade, seismic interpretation has already demonstrated its high potential in fulfilling this mission. With a broader and deeper collaboration between the geophysical and the signal processing communities, we expect to witness more breakthroughs in the near future in seismic-interpretation-enabled subsurface understanding.

Acknowledgments

We would like to acknowledge the support of the Center for Energy and Geo Processing at the Georgia Institute of Technology, and King Fahd University of Petroleum and Minerals (Project GTEC 1401-1402).

Authors

Ghassan AlRegib (alregib@gatech.edu) received his B.S degree in electrical engineering in 1997 from King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia, and his Ph.D.

We expect to witness more breakthroughs in the near future in seismic interpretation-enabled subsurface understanding.

degree in electrical and computer engineering in 2003 from the Georgia Institute of Technology (Georgia Tech), where he is a professor in the School of Electrical and Computer Engineering. In 2012, he was named the director of Georgia Tech's Center for Energy and Geo Processing. He has authored/coauthored more than 200 articles in international journals and conference proceedings and has been issued several U.S. patents. He is a member of the IEEE Image, Video, and Multidimensional Signal Processing Workshop and the IEEE International Workshop on Multimedia Signal Processing Technical Committees (2015–2020). He received several awards including the ECE Outstanding Junior Faculty Member Award and the Steven A. Denning Faculty Award for Global Engagement. He has been an active consultant to global corporations and educational organizations as well as a witness expert on several patent infringement cases. He is the technical program chair of the 2020 IEEE International Conference on Image Processing. He is a Senior Member of the IEEE.

Mohamed Deriche (mderiche@kfupm.edu.sa) received his B.Sc. degree in electrical engineering from the National Polytechnic School, Algeria, in 1986, and his Ph.D. degree in signal processing from the University of Minnesota in 1994. He worked at Queensland University of Technology, Australia, before joining King Fahd University of Petroleum and Minerals (KFUPM) in Dhahran, Saudi Arabia, where he leads the signal processing group. He has published more than 200 papers in multimedia signal and image processing. He delivered numerous invited talks and chaired several conferences including GlobalSIP-MPSP, IEEE Gulf (GCC), Image Processing Tools and Applications, and TENCON (a Region 10 conference). He has supervised more than 35 M.Sc. and Ph.D. students and is the recipient of the IEEE Third Millennium Medal. He also received the Shauman Best Researcher Award, and both the Excellence in Research and Excellence in Teaching Awards at KFUPM.

Zhiling Long (zhiling.long@ece.gatech.edu) received his Ph.D. degree in electrical engineering from Mississippi State University. He is currently a postdoctoral fellow in the Center for Energy and Geo Processing at the Georgia Institute of Technology. His research interests include signal and image processing, machine learning, and human perception models, with applications to seismic interpretation and computer vision.

Haibin Di (haibin.di@ece.gatech.edu) received his B.E. degree in geophysics from Ocean University of China, in 2008 and his Ph.D. degree in geology from West Virginia University in 2016. In summer 2016, he joined the Center for Energy and Geo Processing at the Georgia Institute of Technology as a postdoctoral fellow. His primary research interests include machine-learning-assisted three-dimensional seismic interpretation, reservoir characterization, and seismic attribute analysis.

Zhen Wang (zwang313@gatech.edu) received his B.Sc. and M.S. degrees in information and communication engineering

from Shanghai Jiao Tong University, China. He is currently a Ph.D. student at the Georgia Institute of Technology. His research interests include digital signal and image processing, machine/deep learning, subspace learning, seismic interpretation, and geoprocessing.

Yazeed Alaudah (alaudah@gatech.edu) received his B.Sc. degree in electrical engineering from King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia, and his M.S. degree in electrical and computer engineering from the Georgia Institute of Technology, where he is currently pursuing his Ph.D. degree, working under the supervision of Prof. Ghassan AlRegib in the Center for Energy and Geo Processing. His research interests include digital signal and image processing, machine learning, deep learning, and seismic interpretation.

Muhammad Amir Shafiq (amirshafiq@gmail.com) received his B.Sc. degree in electrical engineering and his M.S. degree in electronics engineering from the University of Engineering and Technology, Taxila, and the Ghulam Ishaq Khan Institute of Engineering Sciences and Technology, Topi, Pakistan, in 2006 and 2008, respectively. His work at National Engineering and Scientific Commission was related to signal processing, embedded systems and real-time implementations. Currently, he is a Fulbright Ph.D. scholar at the Georgia Institute of Technology working under the supervision of Prof. Ghassan AlRegib in the Center for Energy and Geo Processing. His research interests include digital signal and image processing, machine learning, and seismic interpretation.

Motaz Alfarraj (motaz@gatech.edu) received his B.Sc. degree in electrical engineering from King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia, and his M.S. degree in electrical and computer engineering from the Georgia Institute of Technology, where he is currently a Ph.D. student working under the supervision of Prof. Ghassan AlRegib in the Center for Energy and Geo Processing. His research interests include digital signal and image processing, machine learning, deep learning, and seismic interpretation.

References

[1] (2017). Data sets and resources. [Online]. Available: <https://ghassanalregib.com/datasets-and-resources-for-computational-seismic-interpretation/>

[2] I. Cohen, N. Coult, and A. A. Vassiliou, "Detection and extraction of fault surfaces in 3D seismic data," *Geophysics*, vol. 71, no. 4, pp. P21–P27, 2006.

[3] A. Roberts, "Curvature attributes and their application to 3D interpreted horizons," *First Break*, vol. 19, no. 2, pp. 85–100, 2001.

[4] K. J. Marfurt, R. L. Kirrin, S. L. Farmer, and M. S. Bahorich, "3-D seismic attributes using a semblance-based coherency algorithm," *Geophysics*, vol. 63, no. 4, pp. 1150–1165, 1998.

[5] A. Gersztenkorn and K. J. Marfurt, "Eigenstructure-based coherence computations as an aid to 3-D structural and stratigraphic mapping," *Geophysics*, vol. 64, no. 5, pp. 1468–1479, 1999.

[6] P. Bakker, "Image structure analysis for seismic interpretation," Ph.D. dissertation, Delft Univ. Technol., Delft, The Netherlands, 2002.

[7] Y. Alaudah and G. AlRegib. (2016). A generalized tensor-based coherence attribute, in *Proc. 78th European Association Geoscientists and Engineers Annu. Conf. and Exhibition*. [Online]. Available: <http://earthdoc.eage.org/publication/publicationdetail/s/?publication=85573>

[8] G. C. Fehmers and C. F. Höcker, "Fast structural interpretation with structure-oriented filtering," *Geophysics*, vol. 68, no. 4, pp. 1286–1293, 2003.

[9] H. Di, M. Shafiq, and G. AlRegib. (2017). Seismic fault detection based on multi-attribute support vector machine analysis, pp. 2039–2044. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/segam2017-17793533.1>

[10] D. Hale, "Methods to compute fault images, extract fault surfaces, and estimate fault throws from 3D seismic images," *Geophysics*, vol. 78, no. 2, pp. O33–O43, 2013.

[11] S. I. Pedersen, T. Randen, L. Sonneland, and Ø. Steen, "Automatic fault extraction using artificial ants," in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, 2002, pp. 512–515.

[12] X. Wu and D. Hale, "3D seismic image processing for faults," *Geophysics*, vol. 81, no. 2, pp. IM1–IM11, 2016.

[13] Z. Wang and G. AlRegib, "Interactive fault extraction in 3-D seismic data using the Hough transform and tracking vectors," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 99–109, Mar. 2017.

[14] Z. Wang and G. AlRegib, "Automatic fault surface detection by using 3D Hough transform," in *Proc. Expanded Abstracts Society Exploration Geophysicists 84th Annu. Meeting*, 2014, pp. 1439–1444.

[15] A. A. Aqrawi, T. H. Boe, and S. Barros, "Detecting salt domes using a dip guided 3D Sobel seismic attribute," in *Proc. Expanded Abstracts Society Exploration Geophysicists 81st Annu. Meeting*, 2011, pp. 1014–1018.

[16] A. Amin and M. Deriche, "A new approach for salt dome detection using a 3D multidirectional edge detector," *Appl. Geophys.*, vol. 12, no. 3, pp. 334–342, 2015.

[17] A. Berthelot, A. H. Solberg, and L. J. Gelius, "Texture attributes for detection of salt," *J. Appl. Geophys.*, vol. 88, pp. 52–69, Jan. 2013.

[18] P. Guillen, G. Larrazabal, G. González, D. Bumber, and R. Vilalta, "Supervised learning to detect salt body," in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, 2015, ch. 351, pp. 1826–1829.

[19] A. Amin and M. Deriche, "Salt-dome detection using a codebook-based learning model," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 11, pp. 1–5, 2016.

[20] H. Di, M. A. Shafiq, and G. AlRegib, "Multi-attribute k-means cluster analysis for salt dome boundary detection," in *Proc. Extended Abstract 79th European Association Geoscientists and Engineers Conf. and Exhibition*, Paris, France, 2017. doi: 10.3997/2214-4609.201700915.

[21] M. A. Shafiq, T. Alshawi, Z. Long, and G. AlRegib, "Salsi: A new seismic attribute for salt dome detection," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Mar. 2016, pp. 1876–1880.

[22] S. Chopra and K. J. Marfurt, "Understanding the seismic disorder attribute and its applications," *Leading Edge*, vol. 35, no. 8, pp. 695–702, 2016.

[23] X. Wu, "Methods to compute salt likelihoods and extract salt boundaries from 3D seismic images," *Geophysics*, vol. 81, no. 6, pp. IM119–IM126, 2016.

[24] M. Shafiq, Z. Wang, G. AlRegib, A. Amin, and M. Deriche, "A texture-based interpretation workflow with application to delineating salt domes," *Interpretation*, vol. 5, no. 3, pp. SJ1–SJ19, 2017.

[25] A. Berthelot, A. H. Solberg, E. Morisbak, and L.-J. Gelius, "3D segmentation of salt using texture attributes," in *Proc. Expanded Abstracts Society Exploration Geophysicists 82nd Annu. Meeting*, 2012, pp. 1–5.

[26] M. A. Shafiq and G. AlRegib, "Interpreter-assisted tracking of subsurface structures within migrated seismic volumes using active contour," in *Proc. 78th European Association Geoscientists and Engineers Conf. and Exhibition*, May 2016, pp. 1–5.

[27] Z. Wang, T. Hegazy, Z. Long, and G. AlRegib, "Noise-robust detection and tracking of salt domes in postmigrated volumes using texture, tensors, and subspace learning," *Geophysics*, vol. 80, no. 6, pp. WD101–WD116, 2015.

[28] M. Bahorich and S. Farmer, "3-D seismic discontinuity for faults and stratigraphic features: The coherence cube," *Leading Edge*, vol. 14, no. 10, pp. 1053–1058, 1995.

[29] X. Wu, "Directional structure-tensor-based coherence to detect seismic faults and channels," *Geophysics*, vol. 82, no. 2, pp. A13–A17, 2017.

[30] D. Gao, "Latest developments in seismic texture analysis for subsurface structure, facies, and reservoir characterization: A review," *Geophysics*, vol. 76, no. 2, pp. W1–W13, 2011.

[31] C. G. Eichkitz, J. Amtmann, and M. G. Schreilechner, "Calculation of grey level co-occurrence matrix-based seismic attributes in three dimensions," *Computers Geosci.*, vol. 60, pp. 176–183, Oct. 2013.

[32] H. Di and D. Gao, "Nonlinear gray-level co-occurrence matrix texture analysis for improved seismic facies interpretation," *Interpretation*, vol. 5, no. 3, pp. SJ31–SJ40, 2017.

[33] J. Justice, D. Hawkins, and G. Wong, "Multidimensional attribute analysis and pattern recognition for seismic interpretation," *Pattern Recognit.*, vol. 18, no. 6, pp. 391–399, 1985.

[34] A. E. Barnes and K. J. Laughlin, "Investigation of methods for unsupervised classification of seismic data," in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, 2002, pp. 2221–2224.

[35] B. C. Wallet, M. C. de Matos, J. T. Kwiatkowski, and Y. Suarez, "Latent space modeling of seismic data: An overview," *Leading Edge*, vol. 28, no. 12, pp. 1454–1459, 2009.

- [36] C. Song, Z. Liu, Y. Wang, X. Li, and G. Hu, "Multi-waveform classification for seismic facies analysis," *Computers Geosci.*, vol. 101, pp. 1–9, Apr. 2017.
- [37] T. Zhao, V. Jayaram, A. Roy, and K. J. Marfurt, "A comparison of classification techniques for seismic facies recognition," *Interpretation*, vol. 3, no. 4, pp. SAE29–SAE58, 2015.
- [38] P. Meldahl, R. Heggland, B. Bril, and P. de Groot, "The chimney cube, an example of semi-automated detection of seismic objects by directive attributes and neural networks: Part I; methodology," in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, 1999, pp. 931–934.
- [39] H. Hashemi, D. Tax, R. Duin, A. Javaherian, and P. De Groot, "Gas chimney detection based on improving the performance of combined multilayer perceptron and support vector classifier," *Nonlinear Processes Geophys.*, vol. 15, no. 6, pp. 863, 2008.
- [40] D. Singh, P. C. Kumar, and K. Sain, "Interpretation of gas chimney from seismic data using artificial neural network: A study from MAARI 3D prospect in the Taranaki basin, New Zealand," *J. Natural Gas Sci. Eng.*, vol. 36, pp. 339–357, Dec. 2016.
- [41] W. Xiong, Z. Wan, X. Bai, H. Xing, H. Zuo, K. Zhu, and S. Yang, "Adaboost-based multi-attribute classification technology and its application," in *Proc. 76th European Association Geoscientists Engineers Conf. and Exhibition*, 2014. doi: 10.3997/2214-4609.201412000.
- [42] P. Xu, W. Lu, and B. Wang, "Multi-attribute classification based on sparse autoencoder: A gas chimney detection example," in *Proc. 79th European Association Geoscientists and Engineers Conf. and Exhibition*, 2017. doi: 10.3997/2214-4609.201700920.
- [43] H. Zhu, F. Meng, J. Cai, and S. Lu. (2016). Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation, *J. Vis. Commun. Image Represent.*, vol. 34, pp. 12–27. [Online]. Available: <http://dx.doi.org/10.1016/j.jvcir.2015.10.012>
- [44] E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional networks for semantic segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 640–651, Apr. 2017.
- [45] A. Bansal, X. Chen, B. Russell, A. Gupta, and D. Ramanan. (2016). PixelNet: Towards a general pixel-level architecture. [Online]. Available: <http://arxiv.org/abs/1609.06694>
- [46] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," arXiv Preprint 1606.00915, June 2016.
- [47] C. Farabet, C. Couprie, L. Najman, and Y. LeCun, "Learning hierarchical features for scene labeling," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1915–1929, Aug. 2013.
- [48] (2015). LANDMASS: Large north-sea dataset of migrated aggregated seismic structures, CeGP. [Online]. Available: <http://cegp.ece.gatech.edu/codedata/LANDMASS/default.htm>
- [49] (1987). dGB Earth Sciences, The Netherlands Offshore, The North Sea, F3 Block – Complete. [Online]. Available: <https://opendtect.org/osr/pmwiki.php/Main/Netherlands/OffshoreF3BlockComplete4GB>
- [50] Y. Alaudah and G. AlRegib, "Weakly-supervised labeling of seismic volumes using reference exemplars," in *Proc. IEEE Int. Conf. Image Processing*, Sept. 2016, pp. 4373–4377.
- [51] Z. Long, et al., "A comparative study of texture attributes for characterizing sub-surface structures in migrated seismic volumes," CeGP Tech. Rep., 2017. [Online]. Available: https://ghassanalregibdotcom.files.wordpress.com/2018/01/long2017_survey.pdf
- [52] M. Alfarraj, Y. Alaudah, Z. Long, G. AlRegib, and M. Deriche, "A survey of multiresolution texture attributes for computational seismic interpretation," CeGP Tech. Rep., 2017. [Online]. Available: https://ghassanalregibdotcom.files.wordpress.com/2018/01/alfarraj2017_survey.pdf
- [53] J. Zujovic, T. N. Pappas, and D. L. Neuhoff, "Structural texture similarity metrics for image analysis and retrieval," *IEEE Trans. Image Process.*, vol. 22, no. 7, pp. 2545–2558, 2013.
- [54] Y. Alaudah and G. AlRegib, "A curvelet-based distance measure for seismic images," in *Proc. IEEE Int. Conf. Image Processing*, 2015, pp. 4200–4204.
- [55] M. Alfarraj, Y. Alaudah, and G. AlRegib, "Content-adaptive non-parametric texture similarity measure," in *Proc. IEEE 18th Int. Workshop Multimedia Signal Processing*, 2016, pp. 1–6.
- [56] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, Aug. 2000.
- [57] A. D. Halpert, R. G. Clapp, and B. Biondi, "Salt delineation via interpreter-guided 3D seismic image segmentation," *Interpretation*, vol. 2, no. 2, pp. T79–T88, 2014.
- [58] D. Hale and J. Emanuel. (2005). Seismic interpretation using global image segmentation, pp. 2410–2413. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/1.1817860>
- [59] J. Lomask, R. G. Clapp, and B. Biondi, "Application of image segmentation to tracking 3D salt boundaries," *Geophysics*, vol. 72, no. 4, pp. P47–P56, 2007.
- [60] R. Zabih and V. Kolmogorov, "Spatially coherent clustering using graph cuts," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 27 June–2 July 2004, Washington, D.C., 2004, pp. 437–444.
- [61] L. Vincent and P. Soille, "Watersheds in digital spaces: An efficient algorithm based on immersion simulations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 6, pp. 583–598, June 1991.
- [62] A. Levinstein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi, "Turbopixels: Fast superpixels using geometric flows," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 12, pp. 2290–2297, Dec. 2009.
- [63] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Ssstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, Nov. 2012.
- [64] Y. Alaudah and G. AlRegib. (2017). A weakly supervised approach to seismic structure labeling, in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, pp. 2158–2163. [Online]. Available: <http://library.seg.org/doi/abs/10.1190/segam2017-17793533.1>
- [65] D. Lee and H. Seung. (2001). Algorithms for non-negative matrix factorization, *Adv. Neural Inform. Process. Syst.*, no. 1, pp. 556–562. [Online]. Available: <http://papers.nips.cc/paper/1861-algorithms-for-non-negative-matrix-factorization>
- [66] C. Ding, X. He, and H. D. Simon. (2005). On the equivalence of nonnegative matrix factorization and spectral clustering, in *Proc. 5th Society Indian Automobile Manufactures Int. Conf. Data Mining*, no. 4, pp. 606–610. [Online]. Available: <http://www.siam.org/proceedings/datamining/2005/dm05.php>
- [67] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [68] P. O. Hoyer, "Non-negative matrix factorization with sparseness constraints," *J. Mach. Learn. Res.*, vol. 5, pp. 1457–1469, Nov. 2004.
- [69] A. Waldeland and A. Solberg, "Salt classification using deep learning," in *Proc. 79th European Association Geoscientists and Engineers Conf. and Exhibition*, 2017. doi: 10.3997/2214-4609.201700918.
- [70] L. Huang, X. Dong, and T. E. Clee, "A scalable deep learning platform for identifying geologic features from seismic attributes," *Leading Edge*, vol. 36, no. 3, pp. 249–256, 2017.
- [71] K. Hami-Eddine, B. de Ribert, P. Durand, and P. Gascue, "A growing machine learning approach to optimize use of prestack and poststack seismic data," in *Proc. Society Exploration Geophysicists Tech. Program Expanded Abstracts*, 2017, pp. 2149–2152.
- [72] M. Araya-Polo, T. Dahlke, C. Frogner, C. Zhang, T. Poggio, and D. Hohl, "Automated fault detection without seismic processing," *Leading Edge*, vol. 36, no. 3, pp. 208–214, 2017.
- [73] Y. Lin, G. Guthrie, D. Coblenz, S. Wang, and J. Thiagarajan. (2017). Towards real-time geologic feature detection from seismic measurements using a randomized machine-learning algorithm, pp. 2143–2148. [Online]. Available: <https://library.seg.org/doi/abs/10.1190/segam2017-17775517.1>
- [74] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A search space odyssey," in *Proc. IEEE Trans. Neural Networks and Learning Systems*, 2016, pp. 2222–2232.
- [75] A. Graves, G. Wayne, and I. Danihelka, "Neural Turing machines," arXiv Preprint, arXiv:1410.5401, 2014.
- [76] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," arXiv Preprint, arXiv:1703.10593, 2017.
- [77] C. Ledig et al. "Photo-realistic single image super-resolution using a generative adversarial network," arXiv Preprint, arXiv:1609.04802, 2016.
- [78] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," arXiv Preprint, arXiv:1511.06434, 2015.
- [79] W. Lotter, G. Kreiman, and D. Cox, "Unsupervised learning of visual structure using predictive generative networks," arXiv Preprint, arXiv:1511.06380, 2015.
- [80] I. Goodfellow, "NIPS 2016 tutorial: Generative adversarial networks," arXiv Preprint, arXiv:1701.00160, 2016.
- [81] J. T. Springenberg, "Unsupervised and semi-supervised learning with categorical generative adversarial networks," arXiv Preprint, arXiv:1511.06390, 2015.
- [82] M. Shafiq, T. Alshawi, Z. Long, and G. AlRegib, "The role of visual saliency in the automation of seismic interpretation," *Geophys. Prospect.*, 2017. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1111/1365-2478.12570/full>. doi: 10.1111/1365-2478.12570.