Facies prediction with Bayesian inference: Application of supervised and semisupervised deep learning

Sagar Singh[1], Ilya Tsvankin[1], and Ehsan Zabihi Naeini[2]

Abstract

Accurate delineation of geologic facies and determination of live fluids from seismic reflection data is crucial for reservoir characterization during petroleum exploration. Facies classification or fluid identification is often done manually by an experienced interpreter, which makes this process subjective, laborious, and time-consuming. Several machine-learning models have been proposed to automate multiclass facies segmentation, but significant practical challenges (e.g., limited scope of labels for training purposes, skewed data distribution, inefficient performance evaluation metrics, etc) still remain. We present supervised and semisupervised Bayesian deep-learning methodologies to improve analysis of seismic facies depending on the scope of the labeled data. The developed networks reliably predict facies distribution using seismic reflection data and estimate the corresponding uncertainty. Therefore, they provide more consistent and meaningful information for seismic interpretation than commonly used deterministic approaches. We apply the proposed deep-learning models to field data from the North Sea to demonstrate the generalized-prediction capabilities of our methodology. In the case of sufficient availability of manually interpreted labels (or facies), the supervised learning model accurately recovers the facies distribution. When the amount of the interpreted labels is limited, we efficiently apply the semisupervised algorithm to avoid overfitting.

Introduction

A reservoir in the geosciences represents a natural storage unit of fluid or gas in the subsurface. The primary focus of hydrocarbon exploration is on reservoirs that contain oil or gas deposits. Characterizing the often complex reservoir structure is a challenging task that requires time-consuming manual interpretation performed by experienced geologists and geophysicists. With a continued increase in the volume of seismic data, the reliance on human interpretation renders this procedure subjective and inefficient (Bond et al., 2007).

Accurate automatic interpretation of lithologic units from seismic reflection data could significantly enhance subsurface characterization workflows. Over the past decade, several machine-learning (ML) models have been proposed for efficient lithologic interpretation. Examples of such ML algorithms include:

1) *k*-means clustering (Napoli et al., 2020);
2) support vector machines (Wrona et al., 2018; Singh et al., 2021);
3) self-organizing maps (Strecker and Uden, 2002; de Matos et al., 2006; Saraswat and Sen, 2012; Zhao et al., 2017);
4) generative topographic mapping (Roy et al., 2014);
5) independent component analysis (Lubo-Robles and Marfurt, 2019);
6) artificial neural networks based on different seismic attributes (West et al., 2002; Saggaf et al., 2003; Singh et al., 2016).

These methodologies operate with different seismic attributes, whose quality is often difficult to evaluate.

The recent progress in computational resources (including graphical processing units [GPUs]) has made it feasible to use deep-learning models for seismic interpretation (Huang et al., 2017; Liu et al., 2017; Guitton, 2018; Shi et al., 2019). In particular, deep-learning networks have the potential to provide highly efficient semantic segmentation because they are capable of extracting essential features in high-dimensional spaces from large-scale data sets.

Although most convolutional networks learn pointwise estimates of their weights, such estimates do not fully encapsulate the uncertainty in the weight values (LaBonte et al., 2020). There are several implications of representing neural network weights by probability

[1]Colorado School of Mines, Center for Wave Phenomena, Golden, Colorado 80401, USA. E-mail: sagarsingh@mines.edu (corresponding author); ilya@mines.edu.

[2]Earth Science Analytics, London KT3 5HF, UK. E-mail: ehsanzabihi@yahoo.com.

Manuscript received by the Editor 10 May 2021; revised manuscript received 20 December 2021; published ahead of production 24 January 2022; published online 3 March 2022. This paper appears in *Interpretation*, Vol. 10, No. 2 (May 2022); p. T279–T290, 14 FIGS., 2 TABLES.
http://dx.doi.org/10.1190/INT-2021-0104.1.

DOI:10.1190/INT-2021-0104.1

distributions. First, introducing randomness into the weights creates a nondeterministic neural network; point estimates are obtained by sampling each weight distribution of the network during every forward-pass computation. This computational procedure is used to determine the network inference. Because of the nondeterministic nature of the network, every output is generally different. Repeated applications of this procedure, called Monte Carlo sampling, can be used for uncertainty analysis (Pham and Fomel, 2020; Zhao and Chen, 2020). Second, such a network is much more difficult to train (especially in 3D), and it is more susceptible to vanishing/exploding gradients of the employed loss function without proper normalization. Because training volumes for nondeterministic networks can be quite large, the batch size is constrained by the available device memory. As a result, that size is often too small for batch normalization and for computing accurate batch statistics. To solve this problem, we use an innovative technique called "group normalization," which adjusts the number of groups (instead of using the entire layer) employed in normalization (Wu and He, 2018).

Our Bayesian architecture employs a general encoder-decoder setup, also known as U-Net, that is widely used for image segmentation (Ronneberger et al., 2015). In this architecture, the encoder half of the network compresses the input into a latent space, whereas the decoder half decompresses the latent representation of the input into a segmentation map (LaBonte et al., 2020).

Although deep learning has become increasingly popular in geophysics, implementation of 2D/3D facies classification is hampered by such issues as insufficient labeled data available for training, imbalanced facies class distribution, and lack of rigorous criteria for performance evaluation. Besides, the facies class distribution in the training set is not always consistent with field data. In this paper, we demonstrate how to address these issues using a facies model derived from North Sea field data.

Because the amount of well data can vary from relatively abundant in older fields to scarce in new prospects, we propose two types of deep-learning workflows. The first is a state-of-the-art attention-gated modified U-Net convolutional model, which relies on a sufficient scope of labeled data. The second is a semisupervised general adversarial network (GAN), which requires only limited labeled data to be incorporated in the training process (Salimans et al., 2016). Both approaches use probabilistic convolutional layers that can measure errors in the weight space. This results in statistically justified uncertainty quantification at the cost of at least doubling the number of trainable parameters and increasing the convergence time along with the sensitivity to hyperparameter optimization. Due to the proper handling of randomness in the convolutional layers, the predicted facies models are more consistent with the geologic structure and reflection data compared to the conventional methods.

The paper starts by discussing supervised and semisupervised deep-learning methodologies and the proposed architecture modifications. Then, we describe the loss functions designed to generate an accurate posterior distribution of network weights and the geologic model employed to test the proposed networks. The methodology is applied to publicly available field data from the North Sea. The consistency of the results from both networks confirms the generalization capabilities of our models.

Methodology

The proposed supervised and semisupervised ML models include two main components: attention layers and uncertainty estimation using Bayesian inference. We present "end-to-end" models designed to operate with raw seismic data and generate comprehensive facies classification results along with the associated uncertainties. Next we describe the architecture of both models and the key steps in implementing supervised and semisupervised learning.

Supervised training model

Supervised learning is designed to learn a function that maps input examples to an output label based on training data-label pairs. It employs labeled training data to establish a nonlinear relationship between the desired output and the input that includes a set of training examples (Mohri et al., 2012).

Deep-learning convolutional neural networks (CNNs) are commonly applied to computer vision (Krizhevsky and Hinton, 2012) and natural language processing. They are fashioned after the neuron connectivity patterns in the human brain, and represent a regularized version of multilayer perceptrons in fully connected networks (Zhu et al., 2020).

Generally, a convolutional network consists of multiple input, hidden, and output layers. The types of hidden layers include convolutional, rectified linear unit (ReLU) (activation function), pooling, fully connected dense, and normalization layers. Deep-learning CNNs outperform the existing methods listed earlier due to their effectiveness and minimal preprocessing requirements compared to other image-processing techniques. The mechanism of hierarchical feature learning enables CNNs to capture the most salient and sensitive feature representations for a range of problems, such as image classification and segmentation. A CNN is different from regular neural networks (e.g., the multilayer perceptron [MLP]) because it uses multiple convolutional layers instead of fully connected layers to filter input data.

CNNs take advantage of the fact that a pixel is more closely related to its neighbors than to more distant locations by connecting each neuron to just a local region of the input volume. Thus, CNNs models are more computationally efficient than fully connected MLPs. Because seismic volumes can be regarded as a particular type of image, we investigate the feasibility and performance

of CNNs in seismic facies segmentation. Typically, facies information is derived from well logs and core analysis. If there is a sufficient number of samples with labels, seismic facies interpretation can be formulated as a supervised classification problem (Feng et al., 2021a).

Models used in image-segmentation problems commonly represent variations of U-Net (Ronneberger et al., 2015). Furthermore, to make segmentation of images more accurate, we employ the attention-gated modified U-Net, a reliable segmentation model that combines the U-Net architecture with an attention mechanism (Li et al., 2020). The main elements of the employed architecture are described next.

Modified U-Net architecture

U-Net has been widely applied to image segmentation problems with some modifications for different segmentation tasks. Figure 1a illustrates our modified U-Net architecture, which employs an encoder-decoder setup. The encoder (left) half of the network compresses the input into a latent space, whereas the decoder (right) half decompresses the latent representation of the input into a segmentation map. The encoder half includes four stages, each with two convolutional and normalization layers followed by a max-pooling layer to reduce the size of the input (LaBonte et al., 2020). Thus, after each layer, the input's depth and height are halved, which reduces the input size by a factor of four.

The first three layers of the encoder part correspond to the three stages included in the decoder part of the network. First, we upsample the output of the previous layer and use convolutional and normalization layers to double the depth and height of the input (Ronneberger et al., 2015). The output from the earlier step is then concatenated with the prepooling output of the corresponding encoder layer; such skip connection facilitates feature-forwarding through the network (Figure 1a). This is followed by application of two more convolutional and normalization layers, the final convolutional layer, and a softmax activation. The resulting volume represents a multiclass segmentation probability map and has the same size as the input.

Attention gate signal

To bridge the results of seismic facies analysis with the input label attributes (features), we employ the attention mechanism strategy (Figure 1b). Attention-gated networks are designed to learn feature-attention probabilities (Oktay et al., 2018). Likewise, we compute attention maps to highlight the regions in the seismic image that contribute to the facies-segmentation process. In addition to the benefits of interpretability, the attention mechanism enforces suppression of the interference from untargeted information (Li et al., 2021).

Bayesian inference

A probabilistic approach to deep learning makes it possible to account for uncertainty, so that lower levels of confidence can be assigned to inaccurate predictions. The sources of uncertainty in the data include measurement errors or noise in the labels ("aleatoric" uncertainty) or in the model, as well as insufficient data availability for effective learning ("epistemic" uncertainty).

Bayesian deep learning allows one to compute epistemic uncertainties by modeling a posterior distribution $P(\mathbf{w}|\mathcal{D})$ over the network weights $\mathbf{w}$ for a given training set of data images and labels ($\mathcal{D}$). In practice, finding the exact posterior is not feasible, but an approximation $q(\mathbf{w}|\theta)$ can be obtained using the so-called variational inference (θ is a variational parameter) by minimizing the

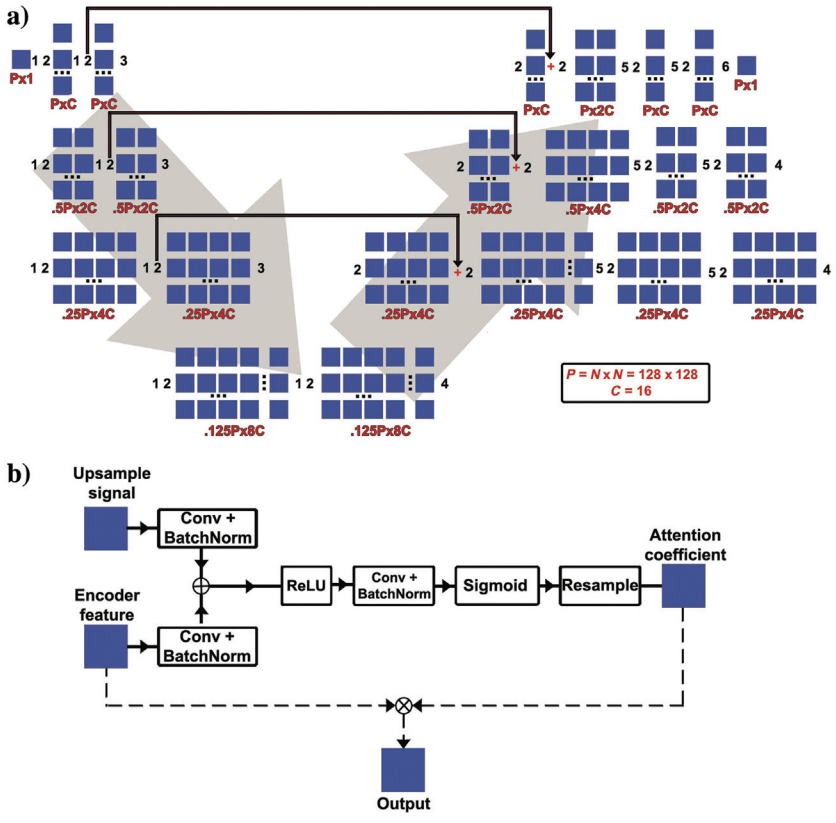

Figure 1. (a) Schematic of the modified U-Net architecture. The measurements are depth, height, and channels. The numbers are — 1) 2D convolution with with 3×3 filter, 1×1 stride, and ReLU activation; 2) group normalization; 3) max pooling of size 2×2; 4) upsampling of size 2×2 followed by Bayesian convolution with 2×2 filters; 5) Bayesian convolution with 3×3 filter, 1×1 stride, and ReLU activation; 6) 2D convolution followed with softmax activation; and **+**: attention gate followed by concatenation. (b) Attention gate signal.

DOI:10.1190/INT-2021-0104.1

Kullback-Leibler (KL) divergence (Graves, 2011). This approach optimizes the weights to prevent overfitting, while simultaneously modeling the weights distribution. The segmentation results are recovered by taking the average over the generated output. The epistemic uncertainty can then be estimated by computing the standard deviation between the samples.

The uncertainty evaluation enhances the interpretability of the model (Figure 1a). Each convolutional Bayesian layer is initialized with a standard normal prior $[P(\mathbf{w}) = \mathcal{N}(0, 1)]$ and employs a flipout estimator (Wen et al., 2018) to approximate the posterior distribution during forward passes. The flipout estimator provides a Monte Carlo approximation of the posterior distribution by integrating over the Bayesian layer's kernel and bias, which significantly lowers the variance (Wen et al., 2018). Bayesian layers are not included in the encoder half of the network to enhance the amount of the information transferred between the original input and the latent space (LaBonte et al., 2020).

Loss function and optimizer

We use monotonic KL annealing (Bowman et al., 2016) to improve the convergence of the model because it helps learn the segmentation before applying the KL divergence penalty. We denote the current epoch as E and define the following hyperparameters: the KL starting epoch s, the KL initial weight k_0, and the step value k_1. Then, the KL weight for the current epoch k_E is given by

$$k_E = \begin{cases} k_0 & \text{if } E \le s, \\ \min[1, k_0, +k_1(E - s)] & \text{if } E > s. \end{cases} \tag{1}$$

Variational learning finds the parameters θ of the distribution $q(\mathbf{w}|\theta)$ by minimizing the variational free-energy cost function $\mathcal{F}$, often called the "expected lower bound." That function (Blundell et al., 2015) consists of the sum of the KL divergence and the negative log-likelihood (NLL):

$$\mathcal{F}_i^E(\mathcal{D}_i, \theta) = \frac{k_E}{M} \mathrm{KL}[q(\mathbf{w}|\theta) \| P(\mathbf{w})] - E_{q(\mathbf{w}|\theta)}[\log P(\mathcal{D}_i|\mathbf{w})], \tag{2}$$

where M is the total number of training examples and i is the minibatch. We divide the KL divergence term by M to optimize the minibatch $i \in \{1, 2, \ldots M\}$, as proposed by Graves (2011). This distributes the KL divergence penalty evenly over each minibatch; without this scaling, the KL term dominates equations 1 and 2, causing the model to converge to a posterior with suboptimal accuracy. Equation 2 can be interpreted as a trade-off between satisfying the simplicity prior (the KL term) and fitting the data set (the NLL term).

Semisupervised training model

Semisupervised learning combines a small volume of labeled data with a larger amount of unlabeled data during training procedure. Semisupervised learning model is a combination of supervised learning (trained with labeled data only) and unsupervised learning (trained with no labeled data). Unlabeled data can significantly improve learning accuracy when combined with even a small number of labeled data points. However, acquiring such highly accurate labeled data for training purposes often demands a physical experiment conducted by experienced researchers. Therefore, the high cost of the labeling process typically precludes the preparation of a fully labeled training set; in contrast, preparing unlabeled data is relatively inexpensive. In such situations, semisupervised learning can provide great practical value.

One such semisupervised learning model is GAN, which has shown promise for effective image segmentation using large, unlabeled data by training a generator model (G) via a discriminator (D) (Salimans et al., 2016). For most image-segmentation tasks involving GANs, the discriminator model is also referred to as a classifier. Due to the superior performance of GAN, this procedure can be extended to semisupervised learning, where the discriminator (trained with both supervised and unsupervised data) and a generator model (trained with noise) are trained simultaneously. Unlike traditional neural networks, GANs comprise two competing components, each represented by a neural network: the generator and discriminator. The generator creates samples that come from a targeted, yet unknown, probability distribution (for instance, data points from a Gaussian distribution or channelized reservoir realizations) using random noise in a low-dimensional space as the input. The discriminator is then trained to identify whether the input image is real (obtained from the data) or synthetic (generator output), and then it classifies the real image into its corresponding class. This process allows the discriminator to learn the underlying data features from unlabeled images. Eventually, the generator evaluates the data's underlying probability distribution and starts producing samples with realistic appearance. The result is a supervised classification model that provides accurate generalization of yet unused examples and a generator model that outputs credible image examples from the input data. A general architecture of a semisupervised GAN is illustrated in Figure 2.

We design the discriminator model of the semisupervised GAN to classify a total of $K + 1$ classes, where K denotes class numbers. The discriminator model also includes a new "synthetic" class coming from the generator. As mentioned earlier, the discriminator model of a semisupervised GAN is trained simultaneously in two ways (supervised and unsupervised).

Unsupervised training. Similar to the process in a traditional GAN, unsupervised training estimates whether the input example is real or synthetic. In this mode, the discriminator model learns to extract salient features from a large unlabeled data set.

DOI:10.1190/INT-2021-0104.1

Supervised training. Training with supervised labels allows the discriminator to classify and assign the class label to real examples using the extracted features.

Consider a standard discriminator model that assigns a data point x to one of K possible classes. The output of the classifier is a K-dimensional vector that can be converted into the class probabilities by applying softmax (Liu and Xiang, 2020). The classic GAN (Salimans et al., 2016) is implemented by labeling samples from the GAN generator G with a new "generated" class $K+1$. We use $P_D(y \le K|x)$ and $P_D(K+1|x)$ to determine the probability that x is true or false, respectively. The model D represents both a discriminator and a classifier, which increases the dimension of the output from K to $K+1$ (Liu and Xiang, 2020). Therefore, we can define the discriminator D as

$$P_D(k|x) = \frac{\exp[\boldsymbol{\omega}_k^T \mathbf{f}(x)]}{\Sigma_{i=1}^{K+1} \exp[\boldsymbol{\omega}_i^T \mathbf{f}(x)]}, \tag{3}$$

where $\mathbf{f}(x)$ is a nonlinear vector function, $\boldsymbol{\omega}_k$ is the weight vector for class k and $\{\boldsymbol{\omega}_1^T \mathbf{f}(x), \ldots, \boldsymbol{\omega}_{K+1}^T \mathbf{f}(x)\}$ is the $(K+1)$-dimensional vector. Since a discriminator with $K+1$ outputs is overparameterized, $\boldsymbol{\omega}_{K+1}$ is fixed as a zero vector. We also denote $D := (\omega, f)$ as a discriminator. Similar to the traditional GANs, D and G solve the following two-component minimax problem with the value function J_{GD}:

$$\min_G \max_D J_{GD} = \min_G \max_D (L_D + U_{GD}), \tag{4}$$

where

$$L_D = \mathbb{E}_{(x,y)\sim p(x,y)} \log P_D(y|x, y \le K), \tag{5}$$

$$U_{GD} = \mathbb{E}_{x\sim p(x)} \log P_D(y \le K|x) + \mathbb{E}_{x\sim p_G(x)} \log P_D(K+1|x). \tag{6}$$

Here, L_D is the supervised learning objective for all labeled data and U_{GD} is the unsupervised (i.e., the traditional GAN) objective, in which p is the true data distribution and p_G is the generated distribution (Liu and Xiang, 2020). When G is fixed, the objectives J_{GD} and U_{GD} become J_D and U_D, respectively. After obtaining a satisfactory value of D by optimizing J_{GD}, we use $\operatorname{argmax} P_D(k|x, k \le K)$ to determine the class of the input data x. Similar to other existing semisupervised GAN algorithms, L_D is used as the supervised learning objective by applying only the softmax operator to the former K-dimensional vector of the output of D (Liu and Xiang, 2020).

The result is a classifier model that can achieve high-fidelity results for standard problems after training on a small number of labeled examples (up to a thousand). Additionally, the training process can produce realistic, high-quality images from the generator model.

Stacked discriminator models with shared weights

Although the semisupervised GAN employs distinct unsupervised and supervised models, output layers for one model are reused as the input to the other. Salimans et al. (2016) describe an efficient implementation of this approach, where the supervised model is created with K output classes and a softmax activation function. Then, the unsupervised model uses the output of the supervised model before the softmax activation that yields a normalized sum of the exponential outputs.

Field-data application

The North Sea is rich in hydrocarbon deposits and well studied in the geophysical and geologic literature (Doornenbal et al., 2019). The continental shelf of the North Sea offshore the Netherlands is divided into zones denoted by different letters of the alphabet; smaller areas within those zones are marked with numbers. Here, we use data from one of such areas, a 16×24 km rectangle known as the F3 block. A 3D seismic survey was conducted in 1987 to identify the geologic structures in the F3 block and search for hydrocarbon reservoirs (Figure 3). In addition, a large number of boreholes were drilled inside the F3 block over the years. The data from the 3D seismic survey, along with additional products, were made publicly available by dGB Earth Sciences in 2020.

Using inline 339 from the survey, ConocoPhillips identified nine groups of facies (the labeled data are freely available). Eight facies have distinct patterns in seismic reflection images and one more facies is used to represent the rest of the labeled samples. These facies along with their main lithostratigraphic features are shown in Figure 4 and listed in Table 1.

To demonstrate the performance of the proposed deep neural networks in seismic facies classification, we apply them to the field data from the F3 block.

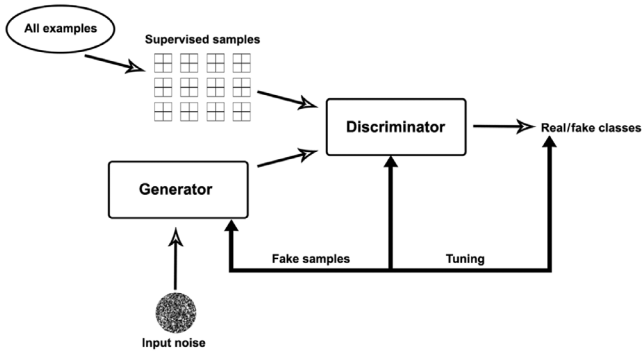

Figure 2. Architecture of the semisupervised GAN includes the generator and the discriminator; both use the network architecture from Figure 1.

We parallelize both developed networks and train them on four NVIDIA Tesla V100 GPUs, each with 12 GB of memory.

Supervised facies prediction

We generate 128,000 training examples by separating the interpreted slice into patches of size 128×128 with the corresponding labels; 30,000 random samples are set aside for validation (Figure 5). Such a large number of diverse labeled samples makes it possible to perform classification using the developed supervised network. To avoid overfitting and extend the network's generalization capabilities, we employ data augmentation that increases the diversity of the training set through application of such random transformations as image rotation, rescaling, shifting, and zooming.

Besides the attention mechanism, Bayesian layers are introduced in the decoder to sample the posterior distribution and obtain meaningful uncertainty maps, which can be directly interpreted as the confidence intervals for the segmentation. Because seismic facies are generally more continuous horizontally than vertically, it may be helpful to employ a different pooling size along the horizontal axis. However, such different pooling sizes do not produce any substantial change in the training for the F3 data.

Figure 6 shows the training history using the Adam optimizer (Kingma and Ba, 2015) with a batch size of 64 and a dynamic learning rate that has the 0.001 starting value. Based on the training history, we stop training at approximately 20 epochs to minimize the possibility of overfitting. The average total training time for supervised learning is around 30 min. The prediction pixel and intersection over union (IOU) (Appendix A) accuracy for the validation set is above 99.9%. In Figure 7, we also show the attention maps for several patches obtained from different locations in Figure 4. The employed accuracy metrics are particularly useful to evaluate the model performance for strongly imbalanced data sets. After successfully training and evaluating the networks, we perform facies prediction for xline 775 (Figure 8). The classification results overlaid on the seismic image (Figure 9) indicate that the predicted target facies distribution is consistent with the seismic data.

Table 1. Facies identified by ConocoPhillips from the seismic data for the F3 block.

Facies	Seismic reflection characteristics
1. Brown	Low coherency
2. Violet	High-amplitude continuous
3. Grass green	Low-amplitude dips
4. Gray	High-amplitude dips
5. Orange	Chaotic
6. Yellow	Low amplitude
7. Magenta	High amplitude
8. Gray	Salt intrusion
9. Turquoise	Everything else

The facies names are taken from Liu et al. (2020).

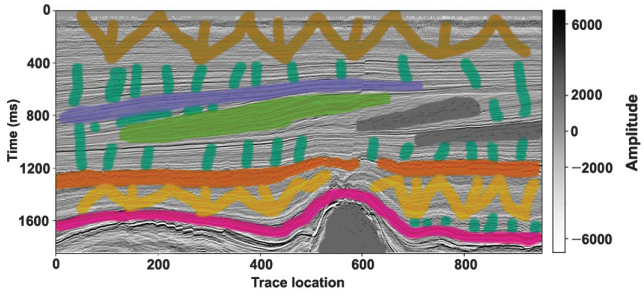

Figure 4. Labeled data set from a 2D inline time-migrated seismic section. Color interpretations are defined in Table 1.

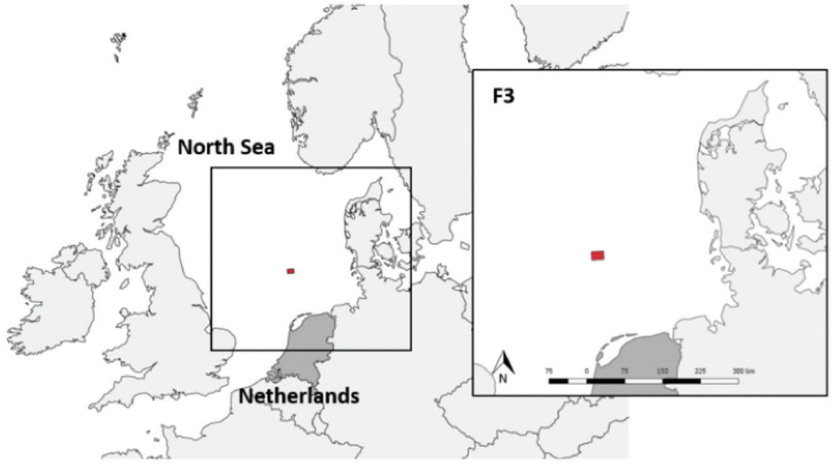

Figure 3. Location of the F3 3D survey in the North Sea, offshore the Netherlands (adapted from Silva et al., 2019).

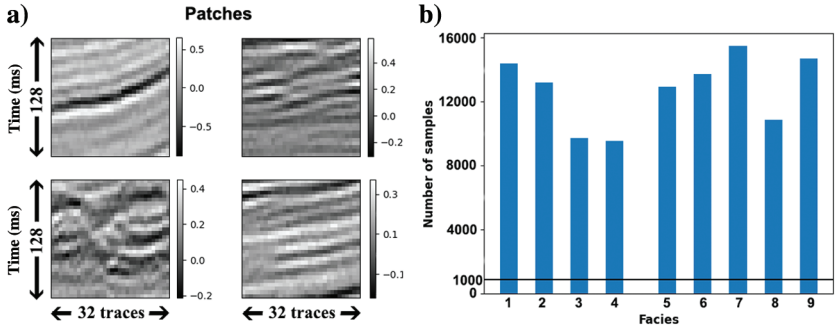

Figure 5. Training data for the ML model. (a) Examples of patches (normalized amplitudes) extracted from the input data and (b) the distribution of labels for each available facies. In total, 1000 samples for each class (marked by the horizontal line on plot [b]) are used in the semisupervised learning model.

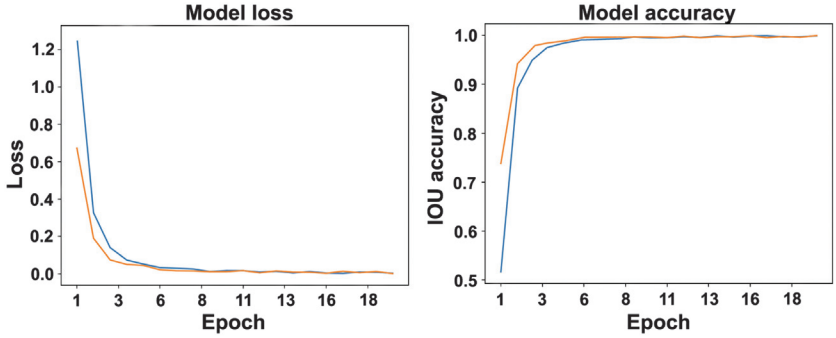

Figure 6. Model training history for the supervised learning. The blue and orange curves represent training and validation, respectively.

To obtain a distribution of the posterior for each grid point, 50 Monte Carlo samples are computed in the rectangular region marked in Figure 9 (so-called epistemic uncertainties; Feng et al., 2021b). We define the confidence intervals for the segmentation as the 33rd and 67th percentiles of the softmax values and the uncertainty as the difference between these percentiles (Figure 10). Note that the uncertainty maps from our Bayesian model provide consistent error estimates across the boundaries of the facies.

Semisupervised facies prediction

Supervised training models perform adequately, if a sufficient volume of borehole data is available. However, supervised learning may fail in areas with just a few drilled wells, where the model becomes overfitted due to a small number of the training data samples. This limits the applicability of supervised algorithms, especially in new exploration areas.

The same model architecture as in the supervised example is used here to demonstrate the feasibility of applying semisupervised GAN to facies classification with a relatively small number of the available labeled data points. We extract 64×64 patches from the interpreted slice and select 1000 training samples for each corresponding label, as opposed to a total of 128,000 samples in the supervised training. The reason for the smaller patch size in semisupervised learning is the increase in the memory requirements caused by incorporating the complex network architecture (Figure 1) for both the generator and discriminator. The generator uses patches of Gaussian realizations as the input to synthesize data with the same size as the training samples ($64 \times 64 \times 1$).

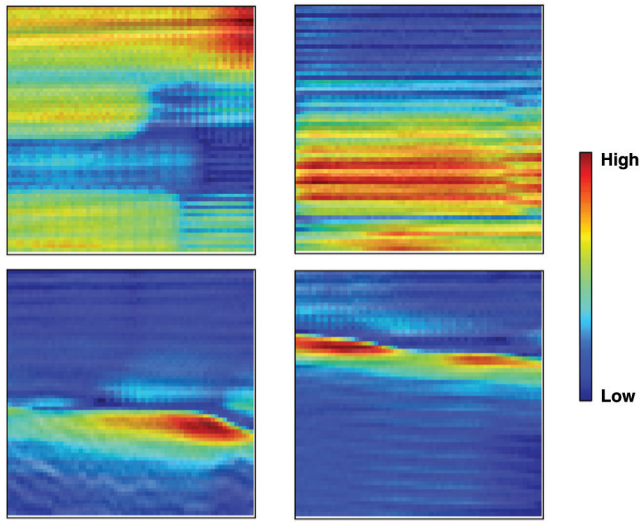

Figure 7. Attention maps of testing patches. The target facies and other untargeted regions show higher energy. Dimensions of the patches are the same as those defined in Figure 5a.

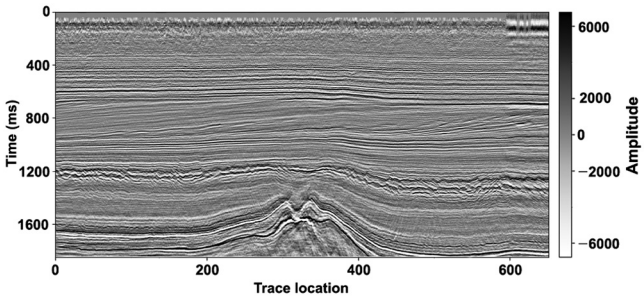

Figure 8. Xline 775 from the F3 block.

DOI:10.1190/INT-2021-0104.1

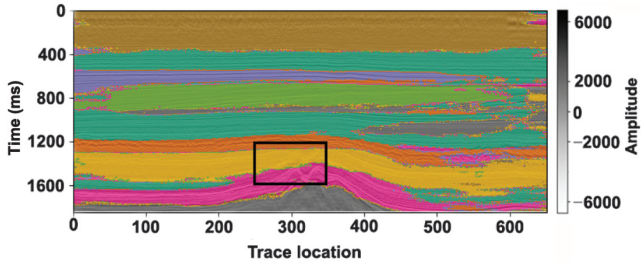

Figure 9. Facies prediction for xline 775 using the supervised training model. The uncertainties are estimated inside the marked rectangular area.

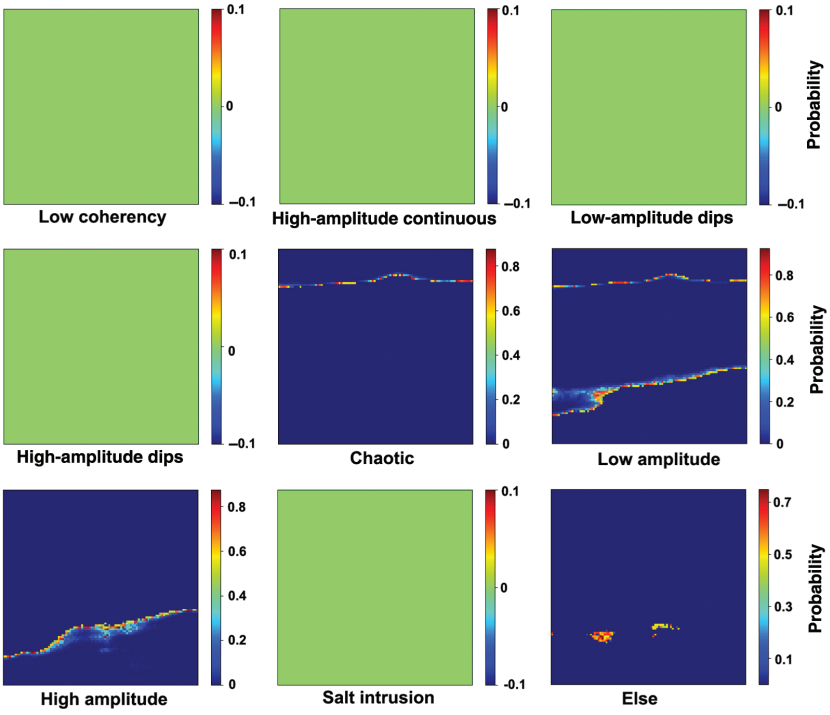

Figure 10. Supervised uncertainty maps of the posterior prediction for the confidence interval between 30% and 70%. Uncertainties are observed mostly along the boundaries between facies.

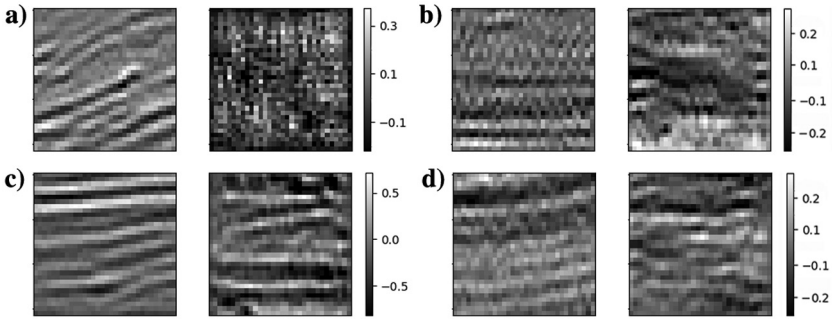

Figure 11. Comparison of the actual labels (left: normalized amplitudes) with the generator output (right: normalized amplitudes) for different epoch numbers: (a) 5, (b) 10, (c) 15, and (d) 20. Dimensions of the patches are the same as those defined in Figure 5a.

We train the generator and discriminator networks on labeled samples extracted randomly from all possible examples and on unlabeled samples evenly extracted from the seismic data along the same inline. In the inference phase, the generator is discarded, and the discriminator is used for the final prediction. The actual and synthetic inputs to the discriminator are illustrated in Figure 11. The training takes about 45 min and is stopped at approximately 20 epochs to avoid overfitting (Figure 12).

The classifier's prediction accuracy on the full data set exceeds 99%. The facies are predicted again for xline 775, and the classification results are overlaid on the seismic data (Figure 13). Compared to the supervised learning results, the semisupervised prediction model is more coherent in the deeper part of the section (for times >1.1 s). The results further improve as the number of the labeled training examples and the patch size are increased (which requires more computing time and memory).

As before, 50 Monte Carlo samples in the marked rectangular region are employed to generate uncertainty maps. The confidence intervals for the segmentation results are set between the 33rd and 67th percentiles of the softmax values (Figure 14).

Discussion

Nondeterministic ML algorithms provide crucial insight into the behavior of predictive models for effective interpretation. The 3D imaging techniques are able to map the interior of objects (e.g., geologic formations), which cannot be studied otherwise. In particular, 3D facies models help interpret the depositional environment and facies distribution using reflection seismic data, which is extremely important in hydrocarbon exploration and development. A critical step in the facies analysis is segmentation, often performed by manually labeling each voxel in a seismic reflection image. However, due to various artifacts and noise in field data, manual segmentation is expensive, irreproducible, and subject to human error.

Deep-learning models, such as CNNs, have revolutionized automated segmentation of 3D images by providing a fast and accurate solution to these problems (LaBonte et al., 2020). For effective application to hydrocarbon exploration, segmentation must include uncertainty

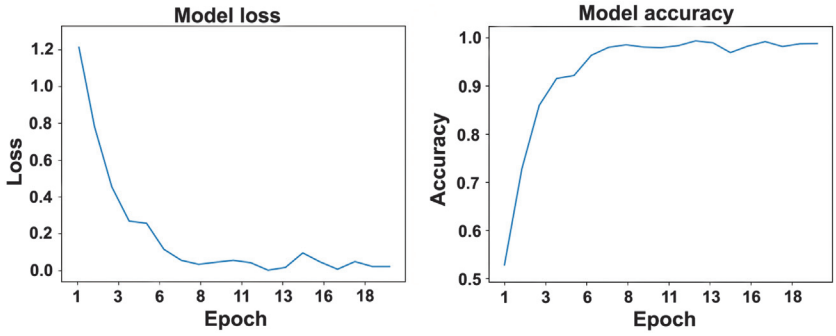

Figure 12. Model training history for test data used in the semisupervised learning.

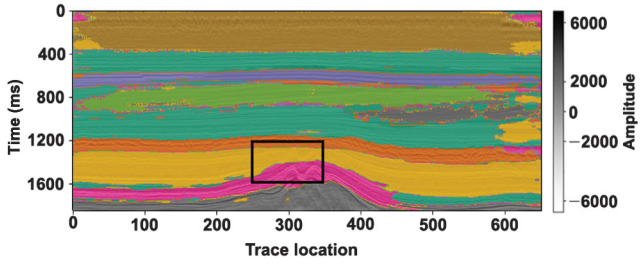

Figure 13. Facies prediction for xline 775 using the semisupervised training model.

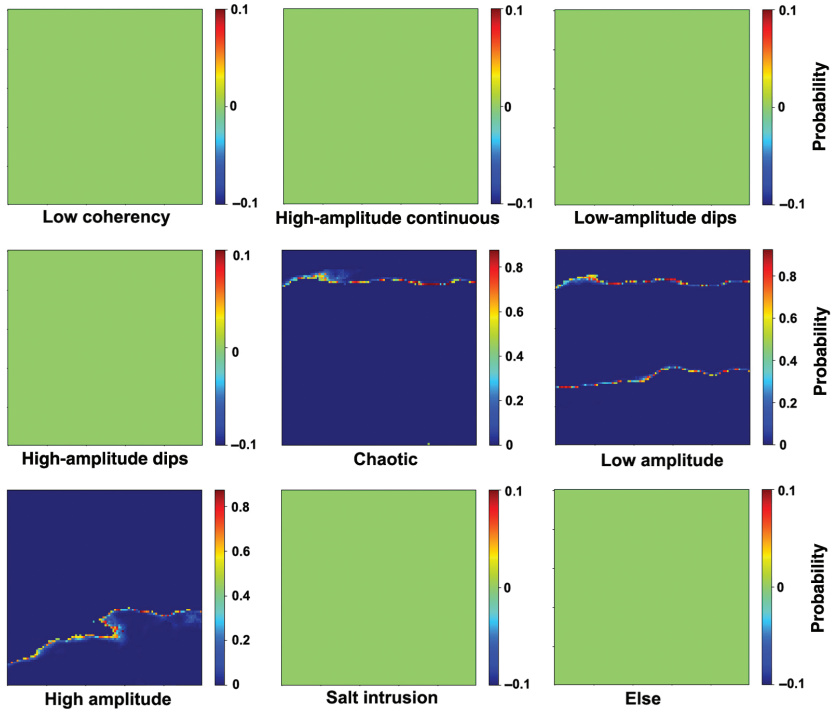

Figure 14. Semisupervised uncertainty maps. According to the supervised uncertainty estimates, the "else" facies has zero probability in that region.

quantification that yields accurate confidence intervals for drilling projects.

The Bayesian inference models proposed here evaluate the uncertainty in the weight space. This results in statistically justified uncertainty quantification at the cost of at least doubling the number of trainable parameters (Gal and Ghahramani, 2016) and increasing the convergence time and sensitivity to hyperparameter optimization (Ovadia et al., 2019). Our Bayesian models effectively predict multiclass segmentation and generate credible uncertainty maps, which can be directly interpreted as the confidence intervals for the segmentation. Although we have leveraged many ideas from previous ML research, the paper presents a new implementation of deep learning for facies classification.

Conclusions

Classification of seismic facies is an important interpretation step because it helps visualize and evaluate different geologic settings using reflection data. We present two types of deep neural-network workflows designed for efficient automatic facies analysis. Our experiments suggest that fully supervised convolutional networks produce reliable facies predictions for fields with an abundance of labeled data, whereas semisupervised GANs show promise for new prospects with a limited number of labels. Also, we added an attention scheme to our architecture to generate superior segmentation results. Both proposed models support Bayesian inference, which generates statistically justified uncertainty maps. By estimating the uncertainty in the weight space, the employed Bayesian scheme provides interpretable error quantification for segmentation problems. Both models were successfully tested on seismic field data from the North Sea, which confirms the significant potential of deep-learning methods in quantitative seismic interpretation. Our results demonstrate that deep learning provides a robust framework for combining well information and 3D seismic data in automated interpretation of seismic facies.

Acknowledgments

This work was supported by the Consortium Project on Seismic Inverse Methods for Complex Structures at the Center for Wave Phenomena (CWP). The data used in this study can be accessed at "https://github.com/bolgebrygg/-MalenoV."

Data and materials availability

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

Appendix A

Evaluation metrics

To objectively evaluate the performance of our models on field data, we use the evaluation metrics that are described in the computer-vision literature (Alaudah et al., 2019). If the set of pixels belonging to class i is denoted by G_i, and the set of pixels in class i by F_i, the correctly classified pixels form the set $G_i \cap F_i$.

The first evaluation metric is the pixel accuracy (PA) that represents the percentage of the correctly classified pixels over all classes:

$$\mathrm{PA} = \frac{\sum_i |F_i \cap G_i|}{\sum_i |G_i|}, \tag{A-1}$$

where $|\cdot|$ denotes the number of elements in a set.

The second metric is the IOU, which is the number of elements of the intersection of G_i and F_i normalized by the number in elements of their union set,

$$\mathrm{IOU}_i = \frac{|F_i \cap G_i|}{|F_i \cup G_i|}. \tag{A-2}$$

This metric quantifies the overlap between the two sets and should be equal to unity if and only if all pixels were correctly classified. Averaging the IOU over all classes yields the mean intersection over union (mean IOU),

$$\text{mean IOU}_i = \frac{1}{n_c}\sum_i \mathrm{IOU}_i. \tag{A-3}$$

References

Alaudah, Y., P. Michalowicz, M. Alfarraj, and G. AlRegib, 2019, A machine-learning benchmark for facies classification: Interpretation, **7**, no. 3, SE175–SE187, doi: 10.1190/INT-2018-0249.1.

Blundell, C., J. Cornebise, K. Kavukcuoglu, and D. Wierstra, 2015, Weight uncertainty in neural networks: International Conference on Machine Learning, **37**, 1613–1622.

Bond, C. E., A. D. Gibbs, Z. K. Shipton, and S. Jones, 2007, What do you think this is? "Conceptual uncertainty" in geoscience interpretation: GSA Today, **17**, 4–10, doi: 10.1130/GSAT01711A.1.

Bowman, S. R., L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio, 2016, Generating sentences from a continuous space: Conference on Computational Natural Language Learning, 10–21.

de Matos, M. C., P. L. Osorio, and P. R. Johann, 2006, Unsupervised seismic facies analysis using wavelet transform and self-organizing maps: Geophysics, **72**, no. 1, P9–P21, doi: 10.1190/1.2392789.

Doornenbal, J. C., H. Kombrink, R. Bouroullec, R. A. F. Dalman, G. De Bruin, C. R. Geel, A. J. P. Houben, B. Jaarsma, J. Juez-Larr´e, M. Kortekaas, H. F. Mijnlieff, S. Nelskamp, T. C. Pharaoh, J. H. Ten Veen, M. Ter Borgh, K. Van Ojik, R. M. C. H. Verreussel, J. M. Verweij, and G. -J. Vis, 2019, New insights on subsurface energy resources in the Southern North Sea Basin area: Geological Society, Special Publications, 494, doi: 10.1144/SP494-2018-178.

Feng, R., N. Balling, D. Grana, J. S. Dramsch, and T. M. Hansen, 2021a, Bayesian convolutional neural networks for seismic facies classification: IEEE Transactions on Geoscience and Remote Sensing, **59**, 8933–8940, doi: 10.1109/TGRS.2020.3049012.

Feng, R., D. Grana, and N. Balling, 2021b, Uncertainty quantification in fault detection using convolutional neural networks: Geophysics, **86**, no. 3, M41–M48, doi: 10.1190/geo2020-0424.1.

Gal, Y., and Z. Ghahramani, 2016, Dropout as a bayesian approximation: Representing model uncertainty in deep learning: International Conference on Machine Learning, **48**, 1050–1059.

Graves, A., 2011, Practical variational inference for neural networks: Conference on Advances in Neural Information Processing Systems, 2348–2356.

Guitton, A., 2018, 3D convolutional neural networks for fault interpretation: 80th Annual International Conference and Exhibition, EAGE, Extended Abstracts, 1–5, doi: 10.3997/2214-4609.201800732.

Huang, L., X. Dong, and T. E. Clee, 2017, A scalable deep learning platform for identifying geologic features from seismic attributes: The Leading Edge, **36**, 249–256, doi: 10.1190/tle36030249.1.

Kingma, D. P., and J. Ba, 2015, Adam: A method for stochastic optimization: International Conference on Learning Representations.

Krizhevsky, I. S., and G. Hinton, 2012, ImageNet classification with deep convolutional neural networks: Communications of the ACM, **60**, 84–90, doi: 10.1145/3065386.

LaBonte, T., C. Martinez, and S. A. Roberts, 2020, We know where we don't know: 3D Bayesian CNNs for credible geometric uncertainty: arXiv:1910.10793.

Li, C., Y. Tan, W. Chen, X. Luo, Y. Gao, X. Jia, and Z. Wang, 2020, Attention UNet++: A nested attention-aware U-Net for liver CT image segmentation: IEEE International Conference on Image Processing, 345–349.

Li, F., H. Zhou, Z. Wang, and X. Wu, 2021, ADDCNN: An attention-based deep dilated convolutional neural network for seismic facies analysis with interpretable spatial-spectral maps: IEEE Transactions on Geoscience and Remote Sensing, **59**, 1733–1744, doi: 10.1109/TGRS.2020.2999365.

Liu, L., R. Lu, J. Li, and W. Yang, 2017, Seismic lithofacies computation method based on deep learning: International Geophysical Conference, SEG, Global Meeting Abstracts, 649–652, doi: 10.1190/IGC2018-430.

Liu, M., M. Jervis, W. Li, and P. Nivlet, 2020, Seismic facies classification using supervised convolutional neural networks and semisupervised generative adversarial networks: Geophysics, **85**, no. 4, O47–O58, doi: 10.1190/geo2019-0627.1.

Liu, X., and X. Xiang, 2020, How does GAN-based semi-supervised learning work? arXiv:2007.05692.

Lubo-Robles, D., and K. J. Marfurt, 2019, Independent component analysis for reservoir geomorphology and unsupervised seismic facies classification in the Taranaki Basin, New Zealand: Interpretation, **7**, no. 3, SE19–SE42, doi: 10.1190/INT-2018-0109.1.

Mohri, M., A. Rostamizadeh, and A. Talwalkar, 2012, Foundations of machine learning: The MIT Press.

Napoli, O. O., V. Martins do Rosario, J. P. Navarro, P. M. C. Silva, and E. Borin, 2020, Accelerating multi-attribute unsupervised seismic facies analysis with RAPIDS: arXiv, abs/2007.15152.

Oktay, O., J. Schlemper, L. L. Folgoc, M. J. Lee, M. Heinrich, K. Misawa, K. Mori, S. G. McDonagh, N. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, 2018, Attention U-Net: Learning where to look for the pancreas: Conference on Medical Imaging with Deep Learning.

Ovadia, Y., E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. V. Dillon, B. Lakshminarayanan, and J. Snoek, 2019, Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift: Conference on Neural Information Processing Systems.

Pham, N., and S. Fomel, 2020, Uncertainty estimation using Bayesian convolutional neural network for automatic channel detection: 90th Annual International Meeting, SEG, Expanded Abstracts, 3462–3466, doi: 10.1190/segam2020-3427239.1.

Ronneberger, O., P. Fischer, and T. Brox, 2015, U-Net: Convolutional networks for biomedical image segmentation: International Conference on Medical Image Computing and Computer-assisted Intervention, Springer, 234–241.

Roy, A., A. S. Romero-Peláez, T. J. Kwiatkowski, and K. J. Marfurt, 2014, Generative topographic mapping for seismic facies estimation of a carbonate wash, Veracruz Basin, southern Mexico: Interpretation, **2**, no. 1, SA31–SA47, doi: 10.1190/INT-2013-0077.1.

Saggaf, M. M., M. N. Toksöz, and M. I. Marhoon, 2003, Seismic facies classification and identification by competitive neural networks: Geophysics, **68**, 1984–1999, doi: 10.1190/1.1635052.

Salimans, T., I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, 2016, Improved techniques for training GANs: 30th Annual Conference on Neural Information Processing Systems, 2234–2242.

Saraswat, P., and M. K. Sen, 2012, Artificial immune-based self-organizing maps for seismic facies analysis: Geophysics, **77**, no. 4, O45–O53, doi: 10.1190/geo2011-0203.1.

Shi, Y., X. Wu, and S. Fomel, 2019, SaltSeg: Automatic 3D salt segmentation using a deep convolutional neural network: Interpretation, **7**, no. 3, SE113–SE122, doi: 10.1190/INT-2018-0235.1.

Silva, R., L. Baroni, R. Ferreira, D. Civitarese, D. Szwarcman, and E. V. Brazil, 2019, Netherlands dataset: A new public dataset for machine learning in seismic interpretation: arXiv, abs/1904.00770.

Singh, S., A. I. Kanli, and S. Sevgen, 2016, A general approach for porosity estimation using artificial neural network method: A case study from Kansas gas field: Studia Geophysica et Geodaetica, **60**, 130–140, doi: 10.1007/s11200-015-0820-2.

Singh, S., I. Tsvankin, and E. Zabihi Naeini, 2021, Elastic FWI for orthorhombic media with lithologic constraints applied via machine learning: Geophysics, **86**, no. 4, R589–R602, doi: 10.1190/geo2020-0512.1.

Strecker, U., and R. Uden, 2002, Data mining of 3D post-stack seismic attribute volumes using Kohonen self-organizing maps: The Leading Edge, **21**, 1032–1037, doi: 10.1190/1.1518442.

Wen, Y., P. Vicol, J. Ba, D. Train, and R. Grosse, 2018, Flipout: Efficient pseudo-independent weight perturbations on mini-batches: 6th International Conference on Learning Representations.

West, B. P., S. R. May, J. E. Eastwood, and C. Rossen, 2002, Interactive seismic facies classification using textural attributes and neural networks: The Leading Edge, **21**, 1042–1049, doi: 10.1190/1.1518444.

Wrona, T., I. Pan, R. L. Gawthorpe, and H. Fossen, 2018, Seismic facies analysis using machine learning: Geophysics, **83**, no. 5, O83–O95, doi: 10.1190/geo2017-0595.1.

Wu, Y., and K. He, 2018, Group normalization: European Conference on Computer Vision.

Zhao, T., and X. Chen, 2020, Enrich the interpretation of seismic image segmentation by estimating epistemic uncertainty: 90th Annual International Meeting, SEG, Expanded Abstracts, 1444–1448, doi: 10.1190/segam2020-3424987.1.

Zhao, T., F. Li, and K. J. Marfurt, 2017, Constraining self-organizing map facies analysis with stratigraphy: An approach to increase the credibility in automatic seismic facies classification: Interpretation, **5**, no. 2, T163–T171, doi: 10.1190/INT-2016-0132.1.

Zhu, R., X. Tu, and J. X. Huang, 2020, Deep learning on information retrieval and its applications, *in* H. Das, C. Pradhan, and N. Dey, eds., Deep learning for data analytics: Academic Press, 125–153.

Sagar Singh received an Integrated Master of Technology degree (2016) in geophysical technology from the Indian Institute of Technology Roorkee. After graduation, he joined an IT firm where he worked on location platforms applying computer vision and machine learning for improving digital maps. Before joining the Center for Wave Phenomena in the fall of 2017, as a Ph.D. candidate, he worked in the fields of high-dimensional imaging and inversion, and high-performance computing at NGRI (Hyderabad), KAUST (Saudi Arabia), and SNU (South Korea). He is currently working as a data scientist at NVIDIA. His research interests include imaging, inversion, deep learning, and uncertainty analysis.

Ilya Tsvankin received an M.S. (1978) and a Ph.D. (1982) in geophysics from Moscow State University in Russia. From 1978 to 1989, he worked at the Institute of Physics of the Earth in Moscow and was deputy head of laboratory "Geophysics of Anisotropic Media." After moving to the United States in 1990, he became a consultant at the Amoco Production Research Center in Tulsa. Since 1992, he has been on the faculty of Colorado School of Mines, where currently he is professor of geophysics and co-leader of the Center for Wave Phenomena. For his pioneering work on seismic anisotropy, in 1996 he received the Virgil Kauffman Gold Medal Award from the SEG. Among his other recognitions are the SEG Best Paper in Geophysics Award (2009), election to Fellowship of the Institute of Physics (2011), SEG Honorary Membership Award (2015), and SEG Outstanding Educator Award (2020). The third edition of his monograph "Seismic signatures and analysis of reflection data in anisotropic media" was published by SEG in 2012. He also coauthored (with Vladimir Grechka) the book "Seismology of azimuthally anisotropic media and seismic fracture characterization" (SEG, 2011). Since 2002, he has been teaching a short course on anisotropy as part of the SEG Continuing Education Program. His memberships include SEG, EAGE, AGU, and Sigma Xi. His research has focused on seismic modeling, inversion, and processing for anisotropic media, fracture characterization, time-lapse seismic, and nonlinear elasticity.

Ehsan Zabihi Naeini received a B.S. in physics, and an M.S. and a Ph.D. in geophysics. He has more than 15 years extensive software and technology development experience in geoscience software industry and has published numerous papers. His research interests include machine learning, data science, signal processing, physics-driven subsurface analysis, and inversion.

DOI:10.1190/INT-2021-0104.1