Contact amustafa9@gatech.edu OR alregib@gatech.edu
https://ghassanalregib.info/

Active Learning with Deep Autoencoders for Seismic Facies Interpretation

Ahmad Mustafa* and Ghassan AlRegib*

ABSTRACT

Machine learning-assisted seismic interpretation tasks require large quantities of labeled data annotated by expert interpreters, which is a costly and time-consuming process. Where existing works to minimize dependence on labeled data assume the data annotation process to already be completed, active learning—a field of machine learning—works by selecting the most important training samples for the interpreter to annotate in real time simultaneously with the training of the interpretation model itself, resulting in high levels of performance with fewer labeled data samples than otherwise possible. Where active learning has been significantly performed for classification tasks with respect to natural images, there exist very little to no works for dense prediction tasks in geophysics like interpretation. We develop a unique and first-of-a-kind active learning framework for seismic facies interpretation using the manifold learning properties of deep autoencoders. By jointly learning representations for supervised and unsupervised tasks and then ranking unlabeled samples by their nearness to the data manifold, we are able to identify the most relevant training samples to be labeled by the interpreter in each training. On the popular F3 dataset, we obtain close to 10 percentage point difference in terms of interpretation accuracy between the proposed method and the baseline with only three fully annotated seismic sections.

*Center for Energy and Geo Processing (CeGP), Omni Lab for Intelligent Visual Engineering and Science (OLIVES), School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA.

INTRODUCTION

Deep learning (DL) has led to major advancements for various computer vision and natural language processing tasks like image classification, object detection and localization, natural scene understanding, speech recognition, and machine translation, to name a few (Szegedy et al., 2013; Krizhevsky et al., 2017; Rahnemoonfar et al., 2021; Deng et al., 2013; Singh et al., 2017). It has also resulted in major breakthroughs in the extremely critical field of medical image analysis, helping in the detection of cellular structures, tissue segmentation, and in diagnosing various diseases (Shen et al., 2017). In exploration geophysics, it has now become a vital component in many seismic interpretation and inversion tasks, as evidenced by the plethora of works in in salt body delineation (Wang et al., 2015; Amin et al., 2017; Shafiq et al., 2017; Di et al., 2018b), fault detection (Di and AlRegib, 2019; Di et al., 2019a,b), facies classification (Alaudah et al., 2019b,c), seismic attribute analysis (Long et al., 2018; Di et al., 2018a; Alfarraj et al., 2018), structural similarity based seismic image retrieval and segmentation (Alaudah et al., 2019a), and seismic inversion (Alfarraj and AlRegib, 2019; Mustafa et al., 2019; Mustafa and AlRegib, 2020).

At the core of this meteoric success of DL-based applications is its ability to learn useful hierarchical representations from raw data. However, to be able to learn meaningful data representations useful for a variety of tasks and datasets requires large quantities of labeled training data. In the absence of this, DL algorithms become highly prone to overfitting. While annotated data is relatively easier to obtain for computer vision related tasks, the costs exorbitantly increase for domains requiring expert assistance with a high level of domain-specific knowledge, for example medical image analysis and seismic data interpretation. For the latter, we have witnessed in recent history works specifically addressing the problem of limited availability of labeled data. In the work by Alaudah et al. (2019a), the authors show how weak labels for various

seismic features in migrated seismic data could be used in a learning-based framework to obtain detailed, pixel wise annotations for features of interest. Similarly, Wu et al. (2019) demonstrated a method to generate synthetic fault data to train a fault detection network on before performing inference for faults in real datasets using the popular machine learning framework of transfer learning. For seismic inversion, Alfarraj and AlRegib (2019) introduced a novel method based on semi-supervised learning to use large quantities of unlabeled seismic trace data in migrated sections to help improve regression performance on predicting elastic properties with limited well data. Mustafa and AlRegib (2020) on the other hand developed a method they dubbed 'joint learning' to train multiple networks for seismic inversion on different datasets to improve generalization performance on the target. Where these methods serve to use available labeled and/or unlabeled data in whatever fashion to maximize performance gains, the popular machine learning framework of active learning aims to capture the most relevant training samples for the oracle (in this case, the interpreter) to label and provide to the machine learning model in real time to increase accuracy with the fewest labeled training data examples.

For most domains, there usually exists a smaller subset of training examples that suffices in creating the representation space for the entire data Katharopoulos and Fleuret (2018). Training the DL model using the identified subset of training instances results in a better generalization performance on unseen test examples than would have been possible otherwise with a similar number of training examples selected arbitrarily. This is in fact what the premise the field of Active Learning is based on (Settles, 2010). The human expert, also called the *Oracle*, labels a small number of training examples in the first cycle. In each subsequent cycle thereafter, the method identifies and has the expert label a small set of training examples in the unlabeled dataset likely to add the most information to the machine learning model, subject to a pre-defined criterion e.g., uncertainty (Settles, 2010).

We propose in this work a novel active learning methodology based on learning reconstruction manifolds with Deep Autoencoders for seismic interpretation. Autoencoders refer to a family of learning models that are trained to reconstruct their inputs. They are designed so that they are only able to reconstruct data sampled from the training distribution, preventing them from regressing to a simple identity mapping. As a useful by-product, they are able to learn the manifold structure of high dimensional data (Martinez-Murcia et al., 2020; Mustafa and AlRegib, 2021). Kwon et al. (2020) utilize such a learned manifold for the task of anomaly detection on image datasets by thresholding the distribution of reconstruction error-based scores on input training examples.

We view the active learning paradigm as the challenge to identify new training samples most dissimilar from the manifold learnt from preexisting training samples. These training samples are also the ones likely to add the *most* information to the learning model about the dataset that it already doesn't have with preexisting training samples. But there is a caveat: we are majorly interested in supervised, discriminative tasks like classification, segmentation etc. In real life, we would not have ground truth labels for these tasks ahead of time. We could however, make decisions about informative training samples based off the reconstruction manifolds learnt via deep autoencoders. To strengthen the link between the learned manifolds for reconstruction and a supervised tasks like segmentation, we present a network architecture that simultaneously learns the representations for the two tasks in a joint learning framework. This way, there is a stronger guarantee that informative training samples identified for the reconstruction task would also apply to the supervised task. We show later that indeed turns out to be the case.

We train an encoder-decoder architecture simultaneously for reconstruction and seismic facies segmentation using the same feature representations within the training phase. In the inference phase, all unlabeled seismic sections/images are scanned and the one with the highest reconstruction error is sampled, labeled, and added to the training dataset for retraining the network for the next cycle. The underlying assumption—based on the shared representation learning framework—is that seismic sections with the highest reconstruction errors are also going to be the ones the network would have performed more poorly on at segmentation. Identifying such training examples would lead to an improved generalization over the whole seismic volume compared to if a similar number of training sections had been sampled randomly or in some other arbitrary fashion. We verify this hypothesis by comparing the proposed work to a baseline data sampling technique in the seismic interpretation domain. Together with this, we also propose a slightly different variant of the method described here that works with total softmax entropy rather than reconstruction error, which we also show to outperform the baseline. This is the first work of this kind in the domain of subsurface imaging, to the best of our knowledge. Our major contributions can be summarized as (1) proposing a novel active learning method based on manifold learning with deep autoencoders, (2) learning informative training samples for segmentation via high reconstruction error sampling of unlabeled sections, made possible by the shared representation learning framework for supervised and unsupervised tasks, and (3) verifying the efficiency of the proposed work via comparison to a baseline sampling approach in the domain of seismic volume interpretation.

METHODOLOGY

Architecture Details

The network architecture used in the study is depicted in Figure 1. It uses a typical encoder-decoder style based configuration where the encoder consists of multiple 2D

convolutional blocks attached in series to extract increasingly higher level features from the input data (i.e., migrated seismic sections). Each encoder block is followed by a maxpooling layer to downsample feature activations and increase the network's receptive field over the input image. Within each block, two convolutional layers are embedded along with rectified non-linearity (ReLU) and batch normalization layers to increase the stability of the network training process. The decoder blocks are constructed to be similar in structure to the encoder blocks except for the first layer that happens to be a transposed convolutional layer with a stride of two. This serves to act as a counterbalance to the maxpooling layers in the encoder, helping increase the spatial resolution of the feature activations by a factor of two at every block. All convolutional layers in the setup described thus far use a kernel size of 3 with padding 1 to preserve the feature spatial resolution. The final feature activations output by the decoder block are made to be of the same size as the the original input. Using independent 1×1 convolutional layers, they are then mapped simultaneously to two different outputs: the segmentation output and the reconstruction of the input seismic image, respectively. The architecture processes a complete seismic section as a grayscale, single-channel tensor to produce its corresponding facies interpretation and reconstruction outputs. The output channel numbers for each network block are indicated in blue in Figure 1.

Active Learning-based Seismic Interpretation Workflow

In this section, we describe two (slightly different) variants of the proposed interpreter-in-the-loop seismic interpretation workflows, each of which is carried out for a certain prespecified number of cycles. As depicted in Figure 2, each cycle consists of four major stages. In the first stage for any given cycle, we train the aforedescribed network architecture on the given training dataset for a certain number of epochs denoted by E_{train}. This is a hyperparameter that we keep constant at 300 for all cycles. The training dataset for the i-th cycle, denoted by $\mathcal{D}_{train}^i = \{(\mathbf{x}_j, \mathbf{y}_j)\}_{j=1}^{|\mathcal{D}_{train}^i|}$ consists of migrated seismic image sections ($\mathbf{x}_j \in \mathbb{R}^{H \times W}$) and their labels ($\mathbf{y}_j \in \mathbb{R}^{H \times W}$) produced by the human interpreter. H and W refer to the height and width of the migrated seismic sections, respectively. Notice that for the very first cycle, this would be a single seismic image-label pair sampled randomly from the volume of unlabeled seismic data, $\mathcal{D}_{oracle} = \{\mathbf{x}_j\}_{j=1}^{|\mathcal{D}_{oracle}|}$ and annotated by the interpreter. $|\mathcal{D}_{train}^i|$ stands for the number of annotated training samples in the labeled dataset while $|\mathcal{D}_{unlabeled}^i|$ denotes those in the unlabeled set for the i-th cycle. The network trains by minimizing the reconstruction and cross-entropy losses on its respective inputs. For each forward pass, the network denoted by $\mathcal{F}$ and parameterized by the set of parameters Θ produces outputs as

$$\hat{\mathbf{y}}_j, \hat{\mathbf{x}}_j = \mathcal{F}_\Theta(\mathbf{x}_j), \tag{1}$$

where $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are the reconstruction and the segmentation outputs, respectively for the j-th seismic section in the training dataset $\mathcal{D}_{train}^i$. The network losses are computed as

$$l_{rec} = \sum_j \|\hat{\mathbf{x}}_j - \mathbf{x}_j\|_2^2, \tag{2}$$

and

$$l_{ce} = \sum_j \text{Cross-entropy}(\hat{\mathbf{y}}_j, \mathbf{y}_j), \tag{3}$$

where l_{rec} and l_{ce} denote reconstruction and cross-entropy losses, respectively. The optimal network parameters, Θ^* are then obtained as the solution to the optimization problem

$$\Theta^* = \arg\min_\Theta \; l_{rec} + l_{ce}. \tag{4}$$

In the second stage of the workflow, the trained network is used to perform inference on the migrated seismic images in the unlabeled dataset $\mathcal{D}_{oracle}$ to obtain a model response profile, $m(\mathbf{x}_j)$, for $\forall \mathbf{x}_j \in \mathcal{D}_{oracle}$. This maybe done in one of two ways. It may either be computed as the l2 reconstruction error

$$m(\mathbf{x}_j) = \|\mathbf{x}_j - \hat{\mathbf{x}}_j\|_2^2 \quad \forall \mathbf{x}_j \in \mathcal{D}_{oracle}, \tag{5}$$

for all unlabeled seismic sections as computed by the network, or it may be calculated as the entropy over softmax probabilities computed over individual pixels and then summed for all pixels in a given section as

$$m(\mathbf{x}_j) = \sum_{p \in \mathbf{x}_j} \sum_K -y_{p_k} \log y_{p_k} \quad \forall \mathbf{x}_j \in \mathcal{D}_{oracle}, \tag{6}$$

where y_{p_k} refers to the network's estimated probability score of pixel p_k belonging to class k (from a total of K classes). The model response output is then sampled to obtain the seismic section corresponding to the highest response value as

$$\mathbf{x}_s = \arg\max_{\mathbf{x} \in \mathcal{D}_{oracle}} m(\mathbf{x}), \tag{7}$$

to be annotated by the interpreter and added to the labeled training set for the next cycle. The whole process is then repeated for a specified number of cycles. In each subsequent cycle, we train the whole network initialized from the solution obtained in the last cycle, as opposed to being retrained from scratch, which as can be seen in Appendix A leads to important conclusions. The complete workflow for both the reconstruction error and entropy-based methods is summarized in Figure 2 as well as in algorithm 1.
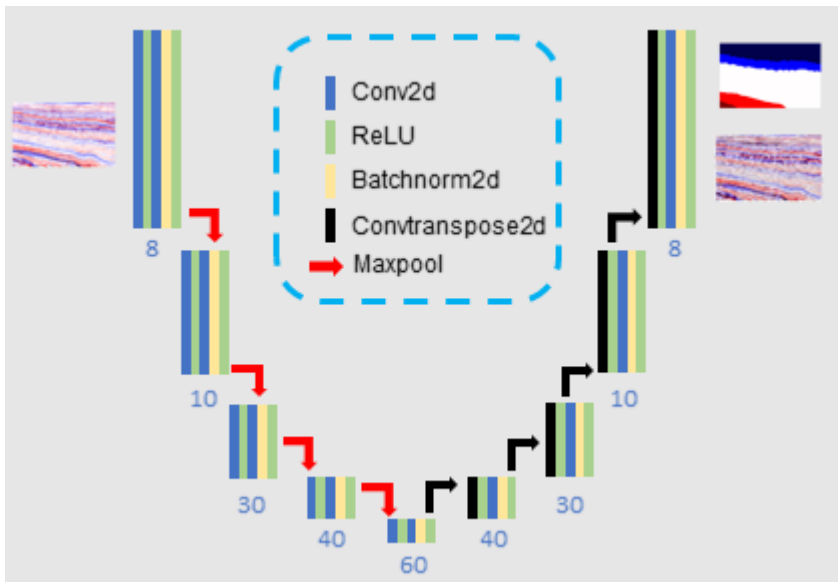
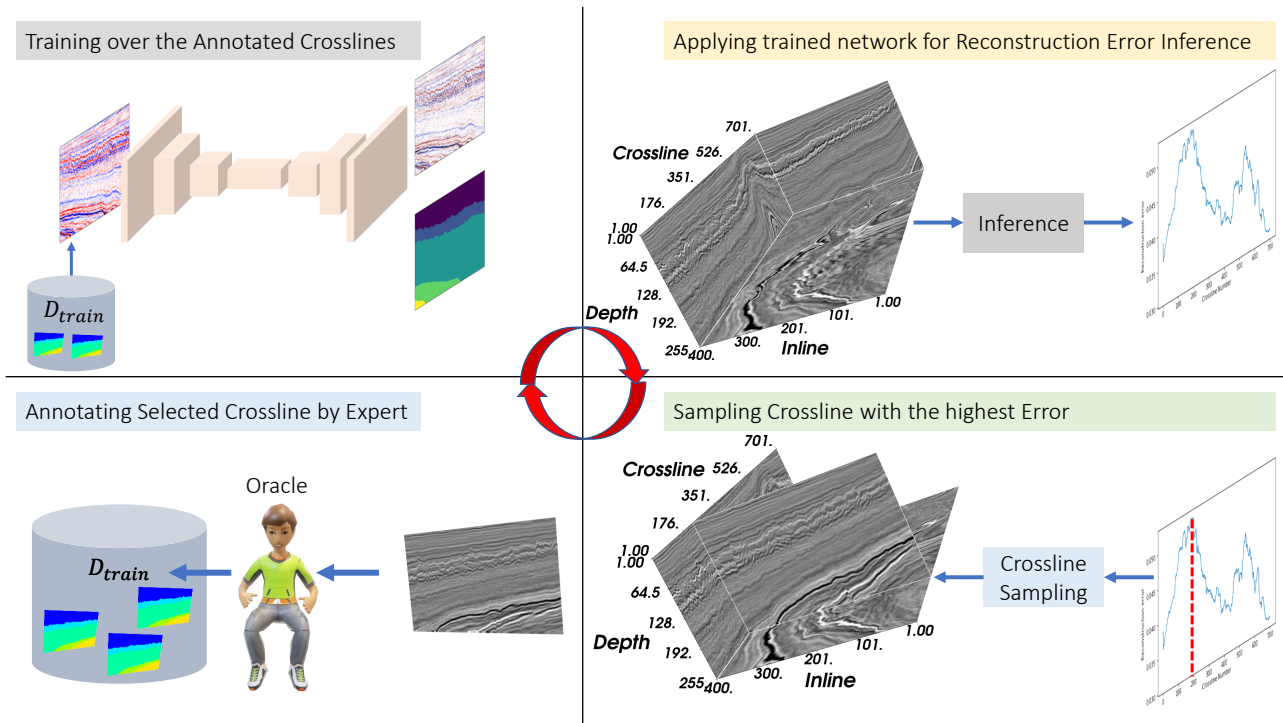Figure 1: Network architecture employed for the proposed method.

Figure 2: Proposed active learning setup.

Algorithm 1 Active Learning-based Interpretation Workflow

---
1. Sample a crossline, x_1 from $\mathcal{D}_{oracle} = \{x_j\}_{j=1}^{|\mathcal{D}_{oracle}|}$.
2. Form the initial training dataset, $\mathcal{D}_{train}^0 = \{(x_1, y_1)\}$.
3. Initialize the model, $\mathcal{F}$.
4. For $i = 1, \ldots, N_{cycles}$:
   (a) Train $\mathcal{F}$ for E_{train} epochs on $\mathcal{D}_{train}^{i-1}$
   (b) Obtain model response $m(\mathbf{x})$ for all $x \in \mathcal{D}_{oracle}$.
   (c) Sample the crossline with the highest error,
   $\mathbf{x}_s = \arg\max_{\mathbf{x} \in \mathcal{D}_{oracle}} m(\mathbf{x})$.
   (d) Form training dataset for next cycle as $\mathcal{D}_{train}^i$
   $= \mathcal{D}_{train}^{i-1} \cup \{(x_s, y_s)\}$.
5. Predict facies labels for all $x \in \mathcal{D}_{oracle}$.

EXPERIMENTAL SETUP

Dataset Details

The dataset used to validate our method is the 3D seismic data for the Netherlands offshore F3 block and its corresponding lithostratigraphic interpretations as developed by Alaudah et al. (2019b). The seismic data—originally in time— was depth migrated and used along with borehole data from 26 wells to construct a 3D geological model. Fault planes and horizons were interpreted and six major groups of lithostratigraphic units were identified. Identifying these lithostratigraphic units—or *rock facies*—on the complete migrated seismic volume was then set up as a machine learning-based image segmentation problem to benchmark predictive algorithms. The final migrated seismic volume and its corresponding 3D geological interpretations consisted of 901 crosslines, 601 inlines, and 255 depth samples.

Evaluation Details

We evaluated the proposed method on the dataset using a 3-fold cross-validation procedure. We divided the complete seismic volume into three non-overlapping subsets consisting of 300, 300, and 301 crosslines, respectively. For each subset, the proposed human-in-the-loop interpretation workflow was carried out for a total of five cycles (i.e., $N_{cycles} = |D_{train}| = 5$) keeping the subvolume as the test set, D_{test} and the remaining two subvolumes to be D_{oracle} in algorithm 1. The cross-validation process is illustrated in Figure 3. After every cycle, the trained network would be evaluated for the average segmentation performance over both D_{oracle} and the test set. The segmentation performance for either set was measured in terms of the mean intersection over union (mIOU) for all six classes in the complete set. The process was repeated three times for each of the three possible dataset partitions (for a total of 9 times) and the results averaged to produce the final plots of mIOU versus cycle number for both D_{oracle} and test sets. We obtain two major benefits by evaluating our method this way: firstly, cross-validation helps to avoid selection bias where the the particular partition scheme

may just happen to work by chance for the given method. By testing the method on all possible dataset partitions and averaging the results, we ensure that we do not bias it to any one dataset configuration. Secondly, the method is repeated thrice even for each partition to discount the effect of random initialization on network performance.

We also show the model predictions of rock facies for several randomly selected crossline sections for both D_{oracle} and D_{test} against their ground-truths for one specific configuration of the cross-validation procedure. This was done by ensembling the predictions on the given crosslines by three randomly intialized networks trained using the proposed method. At this point, we should reinforce that while we maintain a separate D_{test} at all times to effectively measure generalization performance for the trained model, this does not mean that D_{oracle} is the training set. Rather, as explained in the methodology section earlier, D_{train} is only a small subset of D_{oracle}. In our case, D_{train} is only allowed to contain a maximum of five crosslines from a total of 600 in D_{oracle}, which is less than 1% of the size of the latter.

RESULTS AND DISCUSSION

Figure 4(a) depicts the average mIOU computed over the training and evaluation sets (D_{oracle}) on all three folds of the cross-validation procedure after each cycle for three different data sampling strategies. The orange curve corresponds to the interpreter randomly sampling and labeling a crossline section to be added to the set of labeled data the network would train on for the next cycle. This simple yet effective strategy serves as our baseline for the experiments. Adding labeled data in this fashion results in a steady increase in mIOU performance as expected. We also contrast the proposed active learning-based data sampling criteria with the baseline above. This includes sampling the section resulting in the highest softmax entropy (equation 6), and sampling based on distance from learned manifold measured in terms of reconstruction error (equation 5), respectively.Both the proposed methods can be seen to outperform the baseline in terms of mIOU throughout the sampling duration, but more so in the early stages. This is especially important since the interpreter would ideally want to maximize labeling accuracy with the fewest data needed to be annotated. Between the reconstruction and cross-entropy based methods, the former can be observed to perform better. This behavior is reflected across not only D_{oracle}, but also the test section D_{test}, as seen in Figure 4(b). We also demonstrate for one of the cross-validation folds the network facies predictions on select sections from both D_{oracle} and D_{test} after having run the active learning workflow for five cycles, as shown in Figure 5 and Figure 6, respectively. The second column, containing network predictions with the baseline (random) strategy is contrasted with those obtained via reconstruction error-based sampling scheme in the third column. Compared to the ground-truth labels in column one, we observe a marked improvement in performance
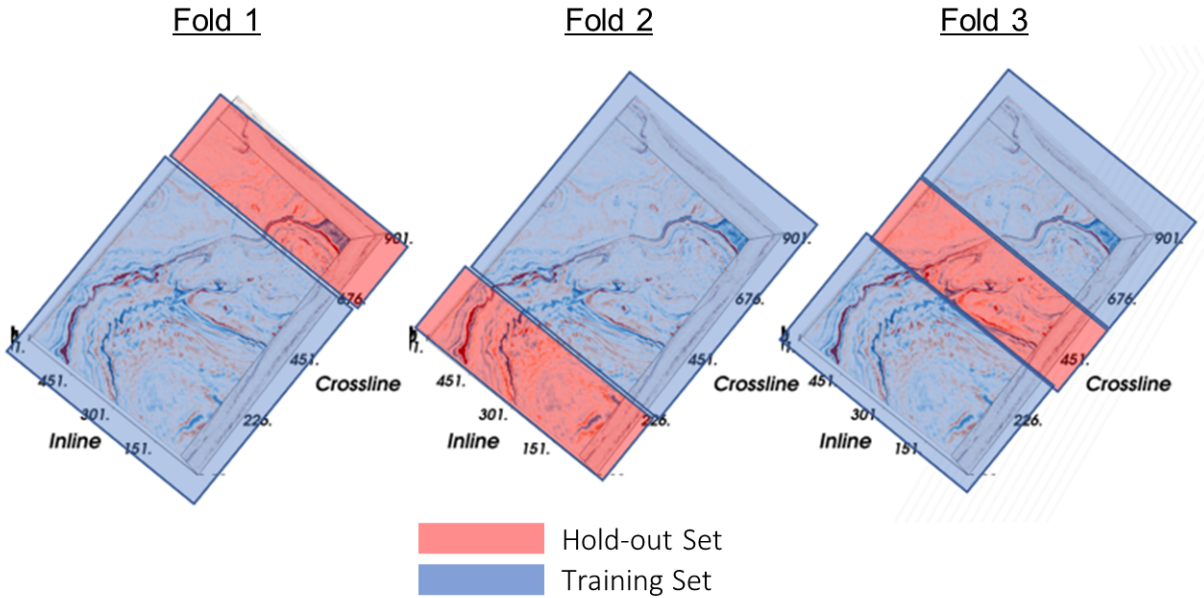
Figure 3: Cross-validation scheme used. The total number of crosslines in the F3 volume are split into three groups as shown, with two at a time used for training and validation (D_{oracle} shaded in blue), and the remaining used as a hold-out test set (D_{test}).

with the latter, especially for the under-represented facies classes. This serves to reinforce our claim about active learning containing the potential to increase performance with fewer labeled data samples.

CONCLUSION

Autoencoders are useful tools to learn data manifolds and characterize individual data samples according to how close (or otherwise) they are from a given data distribution. In this work, we proposed a novel active learning paradigm for the task of seismic facies interpretation on based on the manifold learning property of deep autoencoders. Specifically, the method relies on jointly learning representations for supervised task of facies interpretation and the unsupervised task of seismic reconstruction. This is then used to rank unlabeled data samples by their nearness to the network's learned manifold and samples furthest away are annotated and added to the training dataset. We demonstrate the proposed method is able to achieve a significantly better performance compared to if data were selected randomly for annotaion. We also adapted a commonly used uncertainty metric for active learning in classification tasks (i.e., softmax entropy) for the task of facies interpretation. The proposed method has the potential to significantly reduce the annotation effort for seismic interpreters while ensuring high levels of interpretation accuracy. It need not also be restricted to facies segmentation but can instead be applied to a host of other interpretation tasks involving migrated seismic sections as well.

APPENDIX A

NETWORK REINITIALIZATION VERSUS WARM-STARTING FOR ACTIVE LEARNING

Here, we describe experiments and ablation studies we carried out to investigate the effect on network performance of reinitializing it to random weights versus warm-starting from last cycle's weights in the active learning workflow described earlier. In the work by Bengio et al. (2009), the authors motivate a biology-inspired optimization perspective for neural networks by referring to the way humans and animals learn complicated tasks by first being trained on easier ones. Theorizing that such behavior might also be true of machine learning models, specifically neural networks, they introduce what they term *curriculum learning* to achieve better performance on complex tasks by neural networks. Intuitively, they ground their understanding of the phenomenon by invoking the concept of continuation methods for optimization, where the desired objective is broken down into a series of cost functions $(C_0(\theta), C_1(\theta), \ldots)$ increasing in level of difficulty. In such a situation, optimizing on the easier objective could hypothetically lead to a solution closer to the basin of attraction for an optimal minimum on the more difficult objective as opposed to starting off at a random point in the solution space for the latter. The idea is illustrated with a simple example in Figure A-1. On the surface, it appears to also validate our understanding behind the popularly used framework of transfer learning, where the network performance on a small target dataset is sometimes improved by initializing the network weights
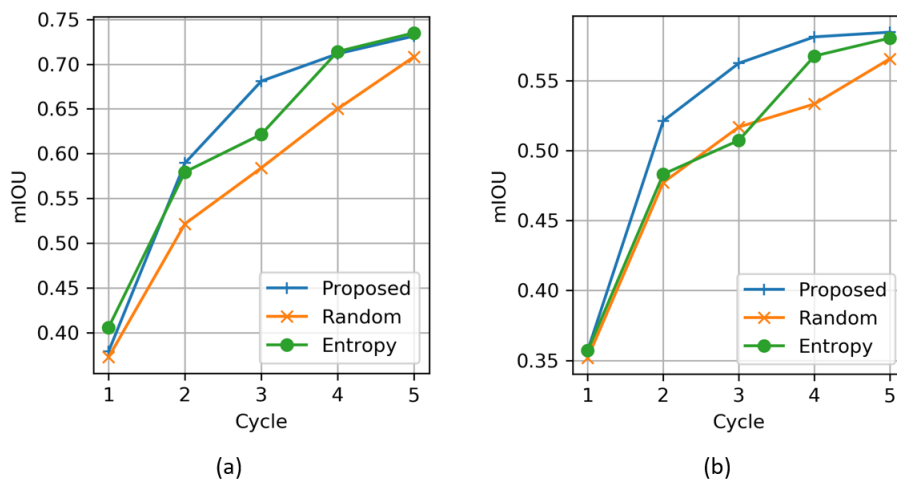
Figure 4: Average mIOUs over five active learning cycles for (a) the training and (b) test volumes respectively.

from the solution to the network optimization on a related dataset. In our active learning experiments, we observed a superior mIOU performance for the network initialized from last cycle's weights compared to one initialized from scratch every cycle *for the same training samples at every cycle.* This is demonstrated in Figure A-2. More specifically, where the blue bars correspond to the network initialized from scratch and trained on all sections obtained until that point, the green bars refer to the network starting from the solution obtained in the last cycle and then trained on the same data examples as the corresponding blue bar. The error curves for the two situations agree with the analysis presented, as shown in Figure A-3. As expected, curriculum learning leads to training error curves starting at lower points from previous cycles.

REFERENCES

Alaudah, Y., M. Alfarraj, and G. AlRegib, 2019a, Structure label prediction using similarity-based retrieval and weakly supervised label mapping: Geophysics, **84**, no. 1, V67–V79.

Alaudah, Y., P. Michałowicz, M. Alfarraj, and G. AlRegib, 2019b, A machine-learning benchmark for facies classification: Interpretation, **7**, no. 3, SE175–SE187.

Alaudah, Y., M. Soliman, and G. AlRegib, 2019c, Facies classification with weak and strong supervision: A comparative study: SEG Technical Program Expanded Abstracts 2019, Society of Exploration Geophysicists, 1868–1872.

Alfarraj, M., Y. Alaudah, Z. Long, and G. AlRegib, 2018, Multiresolution analysis and learning for computational seismic interpretation: The Leading Edge, **37**, 443–450.

Alfarraj, M., and G. AlRegib, 2019, Semisupervised sequence modeling for elastic impedance inversion: Interpretation, **7**, no. 3, SE237–SE249.

Amin, A., M. Deriche, M. A. Shafiq, Z. Wang, and G. AlRegib, 2017, Automated salt-dome detection using an attribute ranking framework with a dictionary-based classifier: Interpretation, **5**, no. 3, SJ61–SJ79.

Bengio, Y., J. Louradour, R. Collobert, and J. Weston, 2009, Curriculum learning: Proceedings of the 26th Annual International Conference on Machine Learning, Association for Computing Machinery, 41–48.

Deng, L., G. Hinton, and B. Kingsbury, 2013, New types of deep neural network learning for speech recognition and related applications: An overview: 2013 IEEE international conference on acoustics, speech and signal processing, IEEE, 8599–8603.

Di, H., M. Alfarraj, and G. AlRegib, 2018a, Three-dimensional curvature analysis of seismic waveforms and its interpretational implications: Geophysical Prospecting.

Di, H., and G. AlRegib, 2019, Semi-automatic fault/fracture interpretation based on seismic geometry analysis: Geophysical Prospecting, **67**, 1379–1391.

Di, H., M. Shafiq, and G. AlRegib, 2018b, Multi-attribute k-means clustering for salt-boundary delineation from three-dimensional seismic data: Geophysical Journal International, **215**, 1999–2007.

Di, H., M. A. Shafiq, Z. Wang, and G. AlRegib, 2019a, Improving seismic fault detection by super-attribute-based classification: Interpretation, **7**, no. 3, SE251–SE267.

———, 2019b, Improving seismic fault detection by super-attribute-based classification: Interpretation, **7**, no. 3, SE251–SE267.

Katharopoulos, A., and F. Fleuret, 2018, Not all samples are created equal: Deep learning with importance sampling: Proceedings of the 35th International Conference on Machine Learning, PMLR, 2525–2534.

Krizhevsky, A., I. Sutskever, and G. E. Hinton, 2017, Imagenet classification with deep convolutional neural networks: Commun. ACM, **60**, 84–90.

Kwon, G., M. Prabhushankar, D. Temel, and G. AlRegib, 2020, Backpropagated gradient representations
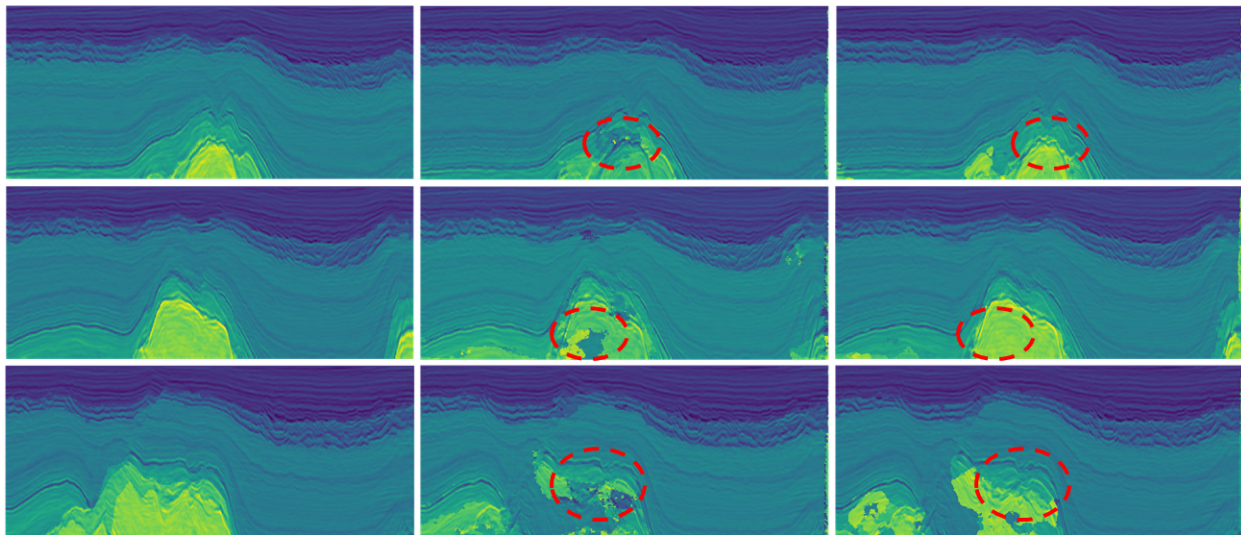
Figure 5: Network predictions after five cycles of active learning for random (second column) and proposed methods (third column) along with ground truth (first column) in the training volume. Each row depicts predictions for a specific section in the volume—500, 550, and 600, respectively.
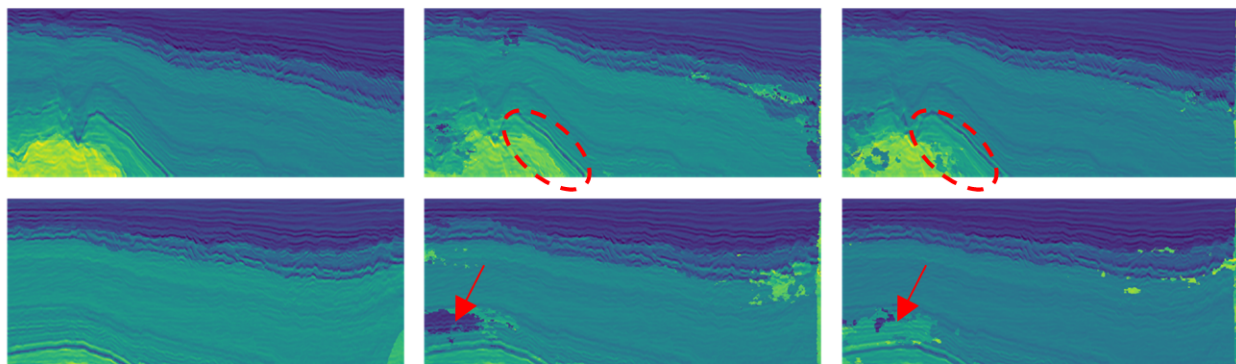

Figure 6: Network predictions after five cycles of active learning for random (second column) and proposed methods (third column) along with ground truth (first column) in the test volume. Each row depicts predictions for a specific section in the volume—50 and 120, respectively.

for anomaly detection: Presented at the Proceedings of the European Conference on Computer Vision (ECCV).

Long, Z., Y. Alaudah, M. A. Qureshi, Y. Hu, Z. Wang, M. Alfarraj, G. AlRegib, A. Amin, M. Deriche, S. Al-Dharrab, and H. Di, 2018, A comparative study of texture attributes for characterizing subsurface structures in seismic volumes: Interpretation, **6**, no. 4, T1055–T1066.

Martinez-Murcia, F. J., A. Ortiz, J.-M. Gorriz, J. Ramirez, and D. Castillo-Barnes, 2020, Studying the manifold structure of alzheimer's disease: A deep learning approach using convolutional autoencoders: IEEE Journal of Biomedical and Health Informatics, **24**, 17–26.

Mustafa, A., M. Alfarraj, and G. AlRegib, 2019, Estimation of acoustic impedance from seismic data using temporal convolutional network: SEG Technical Program Expanded Abstracts 2019, Society of Exploration Geophysicists, 2554–2558.

Mustafa, A., and G. AlRegib, 2020, Joint learning for seismic inversion: An acoustic impedance estimation case study: SEG Technical Program Expanded Abstracts 2020, Society of Exploration Geophysicists, 1686–1690.

——, 2021, Man-recon: Manifold learning for reconstruction with deep autoencoder for smart seismic interpretation: 2021 IEEE International Conference on Image Processing (ICIP), 2953–2957.

Rahnemoonfar, M., T. Chowdhury, A. Sarkar, D. Varshney, M. Yari, and R. R. Murphy, 2021, Floodnet: A high resolution aerial imagery dataset for post flood scene understanding: IEEE Access, **9**, 89644–89654.

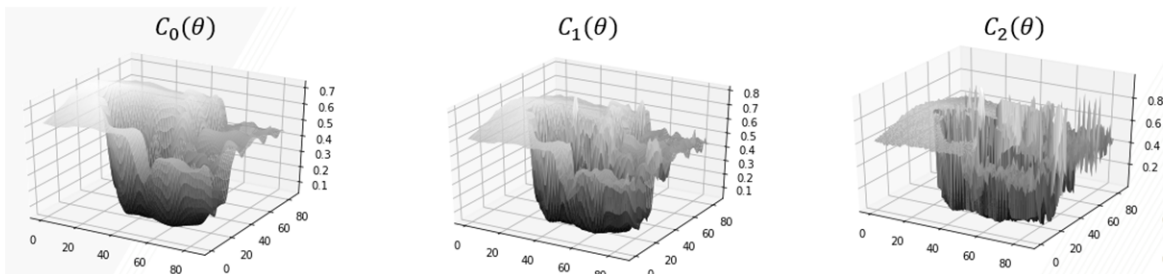Settles, B., 2010, Active learning literature survey: Tech-

Figure A-1: Curriculum learning motivated as an example of continuation method-based optimization. The desired objective is broken into sub objectives $(C_0(\theta), C_1(\theta), ...)$ increasing in level of difficulty. Optimizing on easier objectives makes it easier to find a good minimum on more difficult ones.
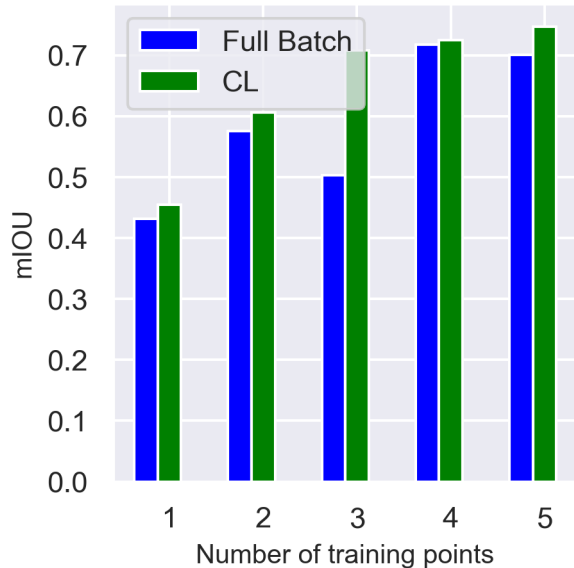

Figure A-2: Full batch (FB) training of the network is carried out for the training samples returned by the network for each cycle (blue bars) and contrasted with the curriculum learning-based proposed method (CL) for the same training samples (green bars). The latter is shown to outperform the former in terms if mIOU at every cycle.
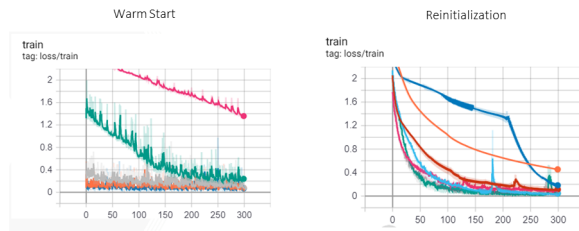

Figure A-3: Network training curves during a five-cycle active learning workflow for curriculum learning-based warm start strategy (left) versus random initialization (right). Curriculum learning leads to training error curves starting at lower points from previous cycles.

nical report.

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

Shen, D., G. Wu, and H.-I. Suk, 2017, Deep learning in medical image analysis: Annual review of biomedical engineering, **19**, 221–248.

Singh, S. P., A. Kumar, H. Darbari, L. Singh, A. Rastogi, and S. Jain, 2017, Machine translation using deep learning: An overview: 2017 international conference on computer, communications and electronics (comptelix), IEEE, 162–167.

Szegedy, C., A. Toshev, and D. Erhan, 2013, Deep neural networks for object detection: Presented at the Advances in Neural Information Processing Systems, Curran Associates, Inc.

Wang, Z., T. Hegazy, Z. Long, and G. AlRegib, 2015, Noise-robust detection and tracking of salt domes in postmigrated volumes using texture, tensors, and subspace learning: Geophysics, **80**, WD101–WD116.

Wu, X., L. Liang, Y. Shi, and S. Fomel, 2019, Faultseg3d: Using synthetic data sets to train an end-to-end convolutional neural network for 3d seismic fault segmentation: Geophysics, **84**, IM35–IM45.