GEOPHYSICS, VOL. 86, NO. 1 (JANUARY-FEBRUARY 2021); P. A1–A5, 4 FIGS.
10.1190/GEO2020-0042.1

DOI:10.1190/geo2020-0042.1

Interactively tracking seismic geobodies with a deep-learning flood-filling network

Yunzhi Shi[1], Xinming Wu[2], and Sergey Fomel[1]

ABSTRACT

We have designed a deep-learning workflow to interactively track seismic geobodies. The algorithm is based on a flood-filling network, which performs iterative segmentation and moving the field of view (FoV). The proposed network takes the previous mask output, together with the seismic image in a new FoV, as a combined input to predict the mask at this FoV. The movement of the FoV is guided by the flood-filling algorithm to visit and segment the full extent of a geobody. Unlike conventional seismic image segmentation methods, the proposed workflow can not only detect geobodies, but it can also track individual geobody instances.

INTRODUCTION

Traditional seismic interpretation tasks such as fault and salt detection are tedious manual time-consuming processes. While the size of seismic data sets continues to increase, it becomes prohibitively expensive to rely on detailed human work. This motivates many attempts of computational-based methods to automate seismic interpretation.

Deep learning methods, mostly based on convolutional neural networks (CNN), are promising techniques for this subject. Seismic fault analysis is the first topic that is addressed by many deep learning interpretation methods. Previous studies (Araya-Polo et al., 2017; Huang et al., 2017; Guitton, 2018; Guo et al., 2018; Zhao and Mukhopadhyay, 2018) propose various CNN-based patch-wise fault detection algorithms to classify the fault/nonfault attribute on a seismic image voxel given its local image patch. Wu et al. (2018b) additionally predict patch-wise fault plane orientation information. Patch-wise CNN-based methods are prone to expensive computational cost; to mitigate this issue and further improve the quality, some authors (Pham et al., 2018; Shi et al., 2018; Wu et al., 2019) develop methods based on an encoder-decoder architecture to perform salt body, channel, and fault analysis, respectively. These works consider interpretation as image segmentation problems. Zhao (2018) applies the encoder-decoder architecture to facies classification and compares it to traditional patch-wise classification.

However, a critical process leading to the practical interpretation analysis is still missing: These methods generate likelihood images but lack the ability to identify individual geobody instances. This means that interpreters would have to scan through the likelihood images, skeletonize the attributes, and create a geologic model. In the field of neurobiology research, similar issues occur when reconstructing neurons from large electron microscope image data. Such postprocessing to obtain object detections from probability volume is proven to be prohibitively expensive even with an optimized pipeline (Berning et al., 2015). New technologies such as graph cut, cluster analysis, and tracking are necessary to fill this gap in an automated interpretation workflow (Beier et al., 2017; AlRegib et al., 2018). Additionally, the variety of field data demands interaction between the end user and workflow to adjust the algorithm to different situations. Current deep learning-based methods only allow knowledgeable adjustments on the training end, including hyperparameters and training samples, but not on the other end that performs inference on field data.

We propose to adopt the flood-filling network (FFN) algorithm, designed and proposed by Januszewski et al. (2018) for electron microscope neuron reconstruction, as a suitable architecture. Although the encoder-decoder network predicts geobody likelihood, FFN runs similarly but takes geobody likelihood as an additional input channel in a recurrent way. The network iteratively performs prediction in a relatively small field of view (FoV) and takes the output from the previous step as a new input. To interactively track seismic geobodies based on the FFN algorithm starting at a given seed point, the network finds the next possible movement according to the prediction to

Manuscript received by the Editor 4 February 2020; revised manuscript received 28 June 2020; published ahead of production 3 October 2020; published online 14 December 2020.
[1]The University of Texas at Austin, Bureau of Economic Geology, Austin, Texas 78713-8924, USA. E-mail: yzshi08@utexas.edu; sergey.fomel@beg.utexas.edu.
[2]University of Science and Technology of China, School of Earth and Space Sciences, Hefei 230026, China. E-mail: xinmwu@ustc.edu.cn (corresponding author).

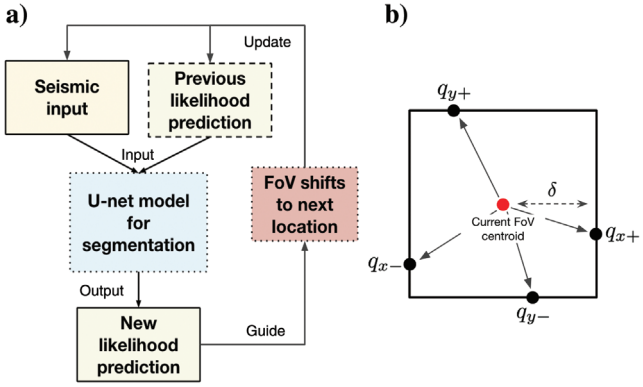

Figure 1. (a) The architecture of the proposed iterative workflow. We use a modified U-net model (Ronneberger et al., 2015) as the segmentation core and form a recurrent style architecture. (b) The scheme of the FoV movement. Starting at the centroid of the FoV, the algorithm searches along the borders with distance δ to the centroid and locates four maximum likelihood points, q_{x-}, q_{x+}, q_{y-}, and q_{y+}, as movement proposals.

track the whole geobody, gradually building a queue to fill the entirety of a geobody. The network predicts each location in this queue until it is exhausted. All geobody instances can be automatically separated and identified in this fashion, with more knowledgeable controls even after training to adjust to various situations such as seed point, tracking length, and connectivity threshold.

MODEL ARCHITECTURE

Encoder-decoder type networks, such as the one proposed by Shi et al. (2018), formulate geologic feature detection as a segmentation problem, which usually takes seismic amplitude or other attributes as input and outputs geobody likelihood values. This type of neural network has a fixed receptive field size, thus limiting the input size for each prediction pass. In seismic images, geologic features such as faults or channels often span a larger area. To cover these areas, it is necessary to either design a network with large receptive field size, which is currently computationally unfeasible (particularly in 3D) or use the sliding window method and then patch all of the window predictions. However, prediction quality degrades near

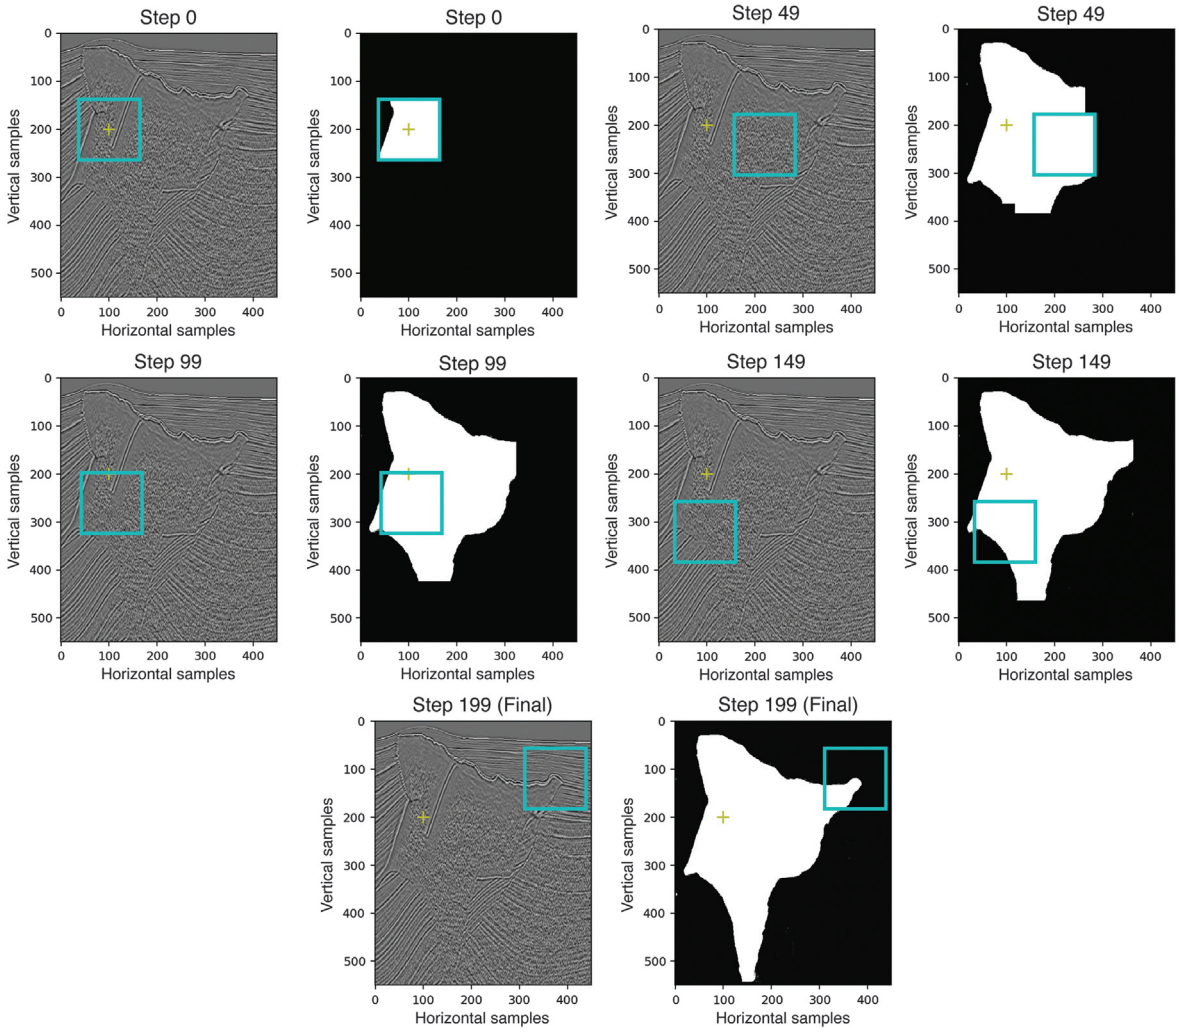

Figure 2. The prediction iteration on a seismic salt body image. Starting from the yellow cross as the seed point, the model performs a segmentation, movement, and likelihood update simultaneously. The cyan box represents the FoV position.

the boundary regions. Therefore, it is natural that we formulate this process as an automatic tracking workflow: The network performs segmentation with the mask as a first channel and images the second channel, where the first channel is replaced by output from the previous pass.

The architecture of the network is shown in Figure 1a. We use a modified U-net (Ronneberger et al., 2015) model as the segmentation core and form a recurrent style architecture. The seismic image within the FoV and a shifted geologic feature likelihood within the same FoV form a two-channel input into the encoder-decoder network. Initially, the FoV is placed at a user-specified seed point inside or near the desired geologic feature. For the first run, the likelihood channel should be empty and disregarded; this run is the same as traditional end-to-end seismic image segmentation. The difference comes when the first run outputs the first likelihood prediction: The proposed method finds the next probable locations within the geobody adjacent to the seed point according to the likelihood map, moves the FoV there, and repeats this process. Each step will update a global likelihood map within the FoV, and all of the steps eventually process the entirety of the geobody.

The movement of the FoV is guided by the flood-filling algorithm according to the following rule: After each run, the algorithm scans through the four boundary lines (or six boundary surfaces in 3D) adjacent to the current FoV centroid with step size δ, as shown in Figure 1b. The scan will find four maximum likelihood points, q_{x-}, q_{x+}, q_{y-}, and q_{y+}, as movement proposals on each boundary line. All of the proposals are checked by these criteria: (1) The likelihood value at the proposed location $L(q)$ is larger than a given threshold T and (2) the distance between the proposed location and all previously visited FoV centroids is larger than step size δ. According to our tests, adjusting T and δ does not impact the precision but only on the computational cost. If a proposed point q_i satisfies the criteria, we add it into a queue $Q = \{q_1, q_2, \ldots\}$. At the next step, Q pops a new FoV centroid and performs a segmentation, likelihood update, and movement proposal in this fashion iteratively until Q is exhausted. Unlike optimization methods that can perform poorly near local minima, the flood-fill algorithm recursively builds a queue on the fly with the aforementioned criteria. The algorithm is guaranteed to end without an explicit terminating condition, instead of entering into an infinite loop.

In this section, we demonstrate the workflow on salt body interpretation. The seismic image is cropped from the SEAM Phase 1 synthetic data set (Fehler and Keliher, 2011). We selected eight crossline 2D slices as training data. The training labels are manual annotations generated by an optimal path picking method (Wu et al., 2018a).

We set the FoV size for 2D salt body interpretation to 127×127. In the network, we downsample the inputs three times with 2× max pooling every two convolutional layers (six layers in total) in the encoder section and symmetrically upsample in the decoder section. Because the input dimension is downsampled 8× in total before upsampling, the FoV is zero-padded to 128×128, and it is cropped to the original size at the final output.

To train the model to move the FoV effectively, we adopt a "dynamic subregion" scheme to constrain the movements of the FoV during training. A limited number of subregions are randomly sampled from each training image so that all subregions are cen-

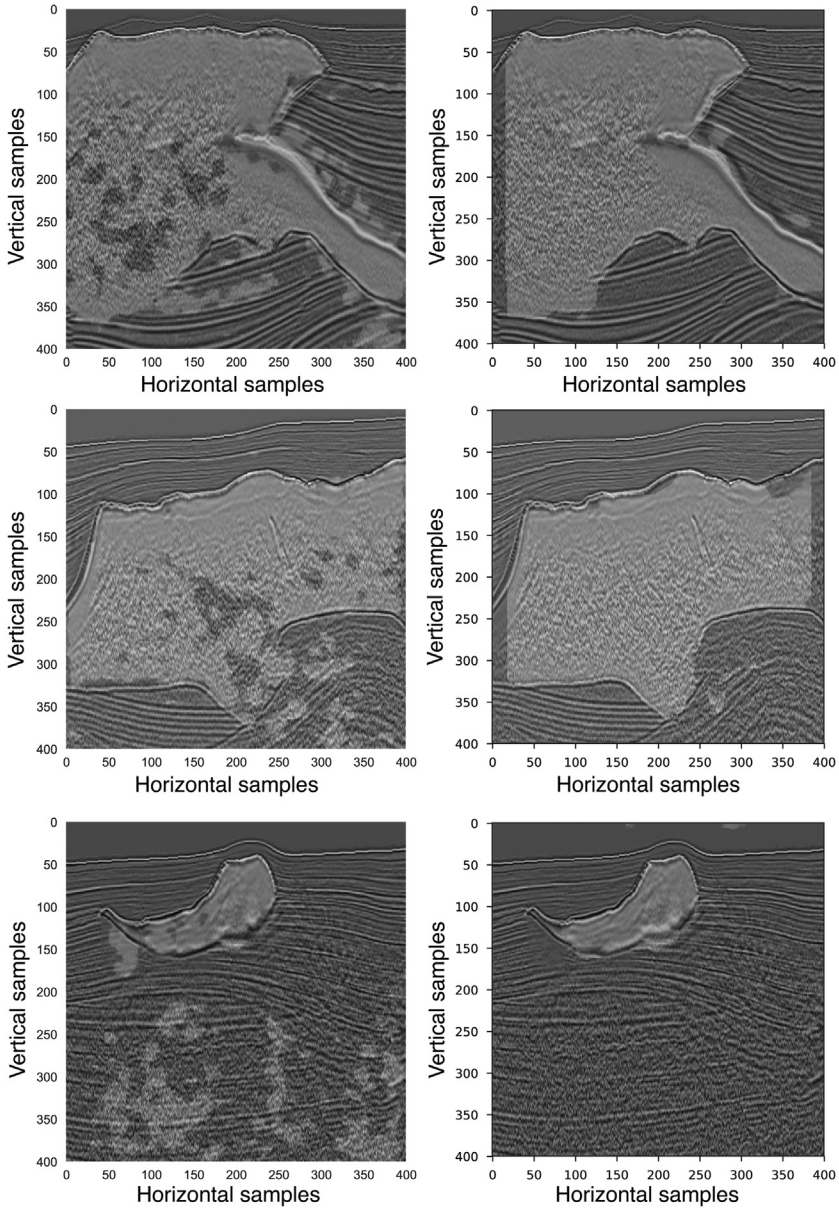

Figure 3. Comparison of the predictions on several seismic images using the previous work by Shi et al. (2018) (the left column) and using the proposed method (the right column). The prediction salt likelihood is overlaid on the seismic image. Notice that a "blind zone" exists on the boundaries of the right column images. This is because we limit the FoVs to not move outside the image, leaving a blind zone on the boundary with size δ.

tered within the valid geobody. In the first epoch, the subregions have the same size as the FoV, so the network only learns how to segment the input image without actually moving the FoV. After that, the subregions grow in size until hitting the 256×256 limit on each new epoch. This way, the network gradually learns to move the FoV based on previous segmentation results further and further away from the starting point. We train the model for five epochs. Because intersection-over-union is sensitive to the classifying threshold, we use the area under the curve (AUC) (Hanley and McNeil, 1982) to measure the precision-recall relationship with different classifying threshold values. The training took fewer than 5 hours with a GTX 1080-Ti GPU. The AUC reaches 99.23% on the validation data set that is held out from the training.

SALT BODY EXAMPLE

Figure 2 demonstrates the iteration process on a seismic salt body image after training. Starting with a seed point indicated by the yellow cross (representing the user input), the model performs a segmentation, movement, and likelihood update simultaneously at each iteration. Within 201 steps, the salt body is fully explored and the iteration is terminated. This process took 13 s to complete with a GTX 1080-Ti GPU. Note that the number of iterations can be decreased with a larger FoV moving step size.

Figure 3 shows more examples that compare the results of the proposed method and those from the previous work by Shi et al. (2018). Although the model is trained with images selected from crossline sections (the proposed method and the work to be compared), some of these test examples use images selected from inline sections. We can notice a significant improvement in the prediction quality with the FFN method, especially where the image is noisy near salt base regions. The reflectivity also looks more consistent and more intact with the proposed method, without many false negatives inside the geobody.

FAULT-PICKING EXAMPLE

We generated 2000 synthetic 2D fault image samples for training and another 2000 images for testing via the method described by Wu et al. (2019). These images are all 128×128 in size and contain various types of geologic transformations including deposition, folding, shearing, and faulting. We set the FoV size in this case as 79×79. The training lasts for 10 epochs, reaches an average of 97.42% AUC on the test data, and takes 9 hours to complete with a GTX 1080-Ti GPU.

Figure 4 shows examples of fault tracking, in which the yellow cross stands for the starting point, the red dots stand for all the FoV centers that the FoV has visited during the iteration, and the fault likelihood map is overlaid on the background image. Figure 4a, 4b, and 4c shows

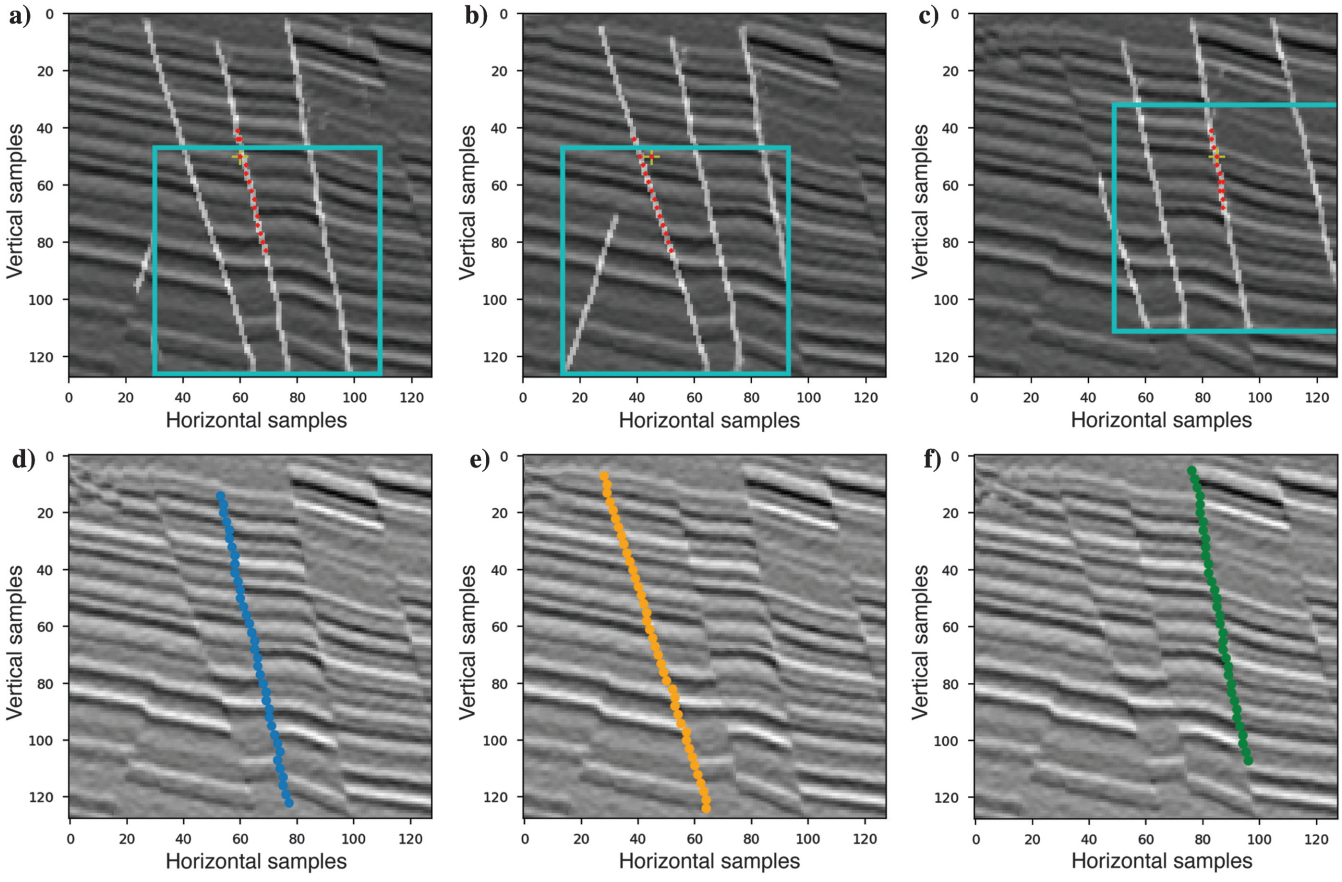

Figure 4. Demonstration of separate instance picking using the proposed method. In this example, fault tracking starts from the yellow cross as the seed point. The red dots stand for all of the FoV centers that the FoV has visited during the iteration. The three faults in the image are picked in (a-c) and are highlighted in different colors in (d-f), respectively.

DOI:10.1190/geo2020-0042.1

examples of picking only one fault instance with other faults in the vicinity. In Figure 4b, even though the seed point is slightly off from the correct fault position, flood filling can automatically correct to align with the fault. Figure 4d, 4e, and 4f shows all three fault instances generated in Figure 4a–4c, respectively. This demonstrates the ability to separate geobody instances with the same attribute.

DISCUSSION

Transferring this method to 3D could be more potential. In addition to the salt body and fault interpretation, the proposed method can be suitable for 3D channel tracking (Pham et al., 2018). However, special care should be taken when extending this method to field data. For example, real images will be noisy, salt bodies may not be perfectly migrated during interpretation, and the frequency range of the training data should match the field data.

However, the proposed method may not perform well trying to use physical information to separate overlapping instances, for example, crossing faults with different strike/dip. This is because each FoV searches for the proposal for the next movement in all directions equally. Wu et al. (2018b) show that deep learning can predict fault dip and strike angle. Because fault dip and strike angle are important parameters that describe fault geometry, integrating them can help guide the flood-filling search pattern to solve the crossing-fault issue.

CONCLUSION

We propose a recurrent style geobody tracking workflow based on an FFN algorithm. The workflow provides two major improvements over previous methods: The tracking algorithm allows, for instance, separation during segmentation, and the atomic design allows for more interaction on the user side to control the model application on various data sets.

We tested the model on 2D salt body synthetic examples and fault-picking synthetic examples. The salt body synthetic example shows that the iterative pattern in the model architecture improves the segmentation significantly compared to the encoder-decoder network architecture. The fault-picking synthetic examples demonstrate that the proposed workflow can separate multiple instances with the same classification attribute.

ACKNOWLEDGMENTS

The authors are grateful to R. Abma and D. Abma for their review and comments. We appreciate the financial support from the sponsors of the Texas Consortium for Computational Seismology. We thank the Texas Advanced Computing Center and the NVIDIA GPU Grant Program for providing the computational resources.

DATA AND MATERIALS AVAILABILITY

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

REFERENCES

AlRegib, G., M. Deriche, Z. Long, H. Di, Z. Wang, Y. Alaudah, M. Shafiq, and M. Alfarraj, 2018, Subsurface structure analysis using computational interpretation and learning: A visual signal processing perspective: arXiv:1812.08756.

Araya-Polo, M., T. Dahlke, C. Frogner, C. Zhang, T. Poggio, and D. Hohl, 2017, Automated fault detection without seismic processing: The Leading Edge, **36**, 208–214, doi: 10.1190/tle36030208.1.

Beier, T., C. Pape, N. Rahaman, T. Prange, S. Berg, D. D. Bock, A. Cardona, G. W. Knott, S. M. Plaza, L. K. Scheffer, and U. Koethe, 2017, Multicut brings automated neurite segmentation closer to human performance: Nature Methods, **14**, 101–102, doi: 10.1038/nmeth.4151.

Berning, M., K. M. Boergens, and M. Helmstaedter, 2015, SegEM: Efficient image analysis for high-resolution connectomics: Neuron, **87**, 1193–1206, doi: 10.1016/j.neuron.2015.09.003.

Fehler, M. C., and P. J. Keliher, 2011, SEAM phase I: Challenges of subsalt imaging in tertiary basins, with emphasis on deepwater Gulf of Mexico: SEG.

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.

Guo, B., L. Li, and Y. Luo, 2018, A new method for automatic seismic fault detection using convolutional neural network: 88th Annual International Meeting, SEG, Expanded Abstracts, 1951–1955, doi: 10.1190/segam2018-2995894.1.

Hanley, J. A., and B. J. McNeil, 1982, The meaning and use of the area under a receiver operating characteristic (ROC) curve: Radiology, **143**, 29–36, doi: 10.1148/radiology.143.1.7063747.

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.

Januszewski, M., J. Kornfeld, P. H. Li, A. Pope, T. Blakely, L. Lindsey, J. Maitin-Shepard, M. Tyka, W. Denk, and V. Jain, 2018, High-precision automated reconstruction of neurons with flood-filling networks: Nature Methods, **15**, 605–610, doi: 10.1038/s41592-018-0049-4.

Pham, N., S. Fomel, and D. Dunlap, 2018, Automatic channel detection using deep learning: 88th Annual International Meeting, SEG, Expanded Abstracts, 2026–2030, doi: 10.1190/int-2018-0202.1.

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.

Shi, Y., X. Wu, and S. Fomel, 2018, Automatic salt-body classification using a deep convolutional neural network: 88th Annual International Meeting, SEG, Expanded Abstracts, 1971–1975, doi: 10.1190/segam2018-2997304.1.

Wu, X., S. Fomel, and M. Hudec, 2018a, Fast salt boundary interpretation with optimal path picking: Geophysics, **83**, no. 3, O45–O53, doi: 10.1190/geo2017-0481.1.

Wu, X., L. Liang, Y. Shi, and S. Fomel, 2019, FaultSeg3D: using synthetic datasets to train an end-to-end convolutional neural network for 3D seismic fault segmentation: Geophysics, **84**, no. 3, IM35–IM45, doi: 10.1190/geo2018-0646.1.

Wu, X., Y. Shi, S. Fomel, and L. Liang, 2018b, Convolutional neural networks for fault interpretation in seismic images: 88th Annual International Meeting, SEG, Expanded Abstracts, 1946–1950, doi: 10.1190/segam2018-2995341.1.

Zhao, T., 2018, Seismic facies classification using different deep convolutional neural networks: 88th Annual International Meeting, SEG, Expanded Abstracts, 2046–2050, doi: 10.1190/segam2018-2997085.1.

Zhao, T., and P. Mukhopadhyay, 2018, A fault detection workflow using deep learning and image processing: 88th Annual International Meeting, SEG, Expanded Abstracts, 1966–1970, doi: 10.1190/segam2018-2997005.1.

Biographies and photographs of the authors are not available.