

Geophysical Corner

Convolutional Neural Network for Fault Prediction

Interpreter versus machine learning

Machine learning in the geosciences has become increasingly popular in the last few years, but what exactly is machine learning? While it might sound intimidating, machine learning is simply using a variety of algorithms designed to seek out different patterns in the data that the interpreter might not notice easily at first. Machine learning can be divided into two main learning processes: supervised learning, which classifies the data with some user input, and unsupervised learning, which clusters the data.

Overall, machine learning has introduced us to a whole different world that has taken geoscientists by surprise. However, the real question is how much can we trust the machine? How accurate can it be? And the most intriguing question of all – can machine learning replace the interpreter?

We will analyze three machine learning processes to assess the pros and cons of utilization of convolutional neural networks for fault prediction versus interpretations made by the user in a highly complex polygonal fault section of a 3-D seismic reflection dataset.

Convolutional Neural Networks

Our study utilizes convolutional neural networks to predict faults on a time-migrated, 3-D seismic reflection dataset situated on the Rankin Platform, offshore northwest Australia. A convolutional neural network, also known as “CNN,” employs a base dataset trained for a specific task. Then the generated model parameters, or learned features, are transferred to a second CNN target to be trained on a different task (figure 1). CNN application falls in the realm of supervised machine learning because it uses a small training set




Figure 1. Visual representation of a transfer learning process of simple convolutional neural network

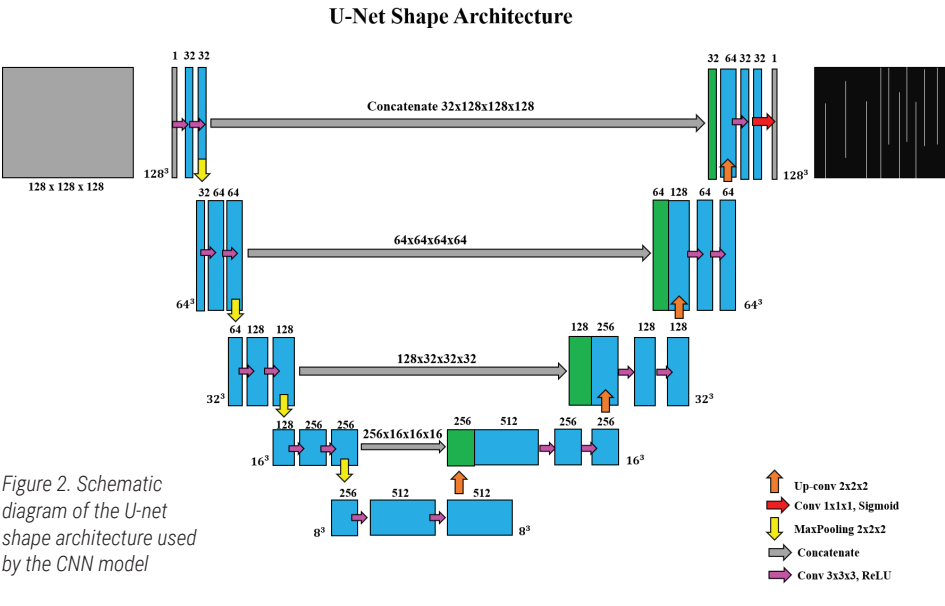


Figure 2. Schematic diagram of the U-net shape architecture used by the CNN model

to find patterns for the data classification. In summary, CNN is used as a feature extractor to classify data or, in our case, to predict faults.

There are four steps used in convolutional neural networks for fault prediction:

- ▶ Constructing the training data samples, which could be the interpretation performed by an interpreter, or a suite of synthetic or modeled data samples with their corresponding fault labels

- ▶ Training the CNN
 - ▶ Validating the results, and
 - ▶ Testing the CNN on the target data volume.
- CNN fault prediction uses an initial network defined by a U-shaped architecture. It is to say, at each step of the training, the amplitude data is convolved with coefficients and later submitted to data reduction. Once the network reaches the bottom of the U, the data gets inverted by going through transposed convolutional

operators, data reduction and activators to predict the location of the faults (figure 2). In our case, we performed the convolutional neural network for fault prediction, fault enhancement and fault skeletonization on polygonal faults that are Paleocene-Eocene age in our seismic reflection data set.

CNN fault prediction uses the original seismic amplitude volume, the CNN configuration and the CNN weights (provided in the software) to generate a fault prediction volume. The generated output is a continuous but lower-resolution volume than a coherence attribute volume. This prompted us to perform fault enhancement, a post-stack attribute that enhances locally planar features, unconformities and other discontinuities, parallel or subparallel to reflector dip within a seismic attribute volume, which generates an output with sharper faults. Subsequently, fault skeletonization was performed, which has the effect of thinning the various faults and other planar features. This algorithm picks on discontinuity features such as faults and unconformities by using dip magnitude and dip azimuth as previously generated inputs (figure 3).

Checking Results Against a Human Interpreter

Why do we want to use a convolutional neural network for fault prediction?

Our goal is to compare the interpretations made by a geoscientist who knows the background geology of the area of interest and has experience on the interpretation of a seismic reflection data set against the predictions made by

Continued on next page ▶

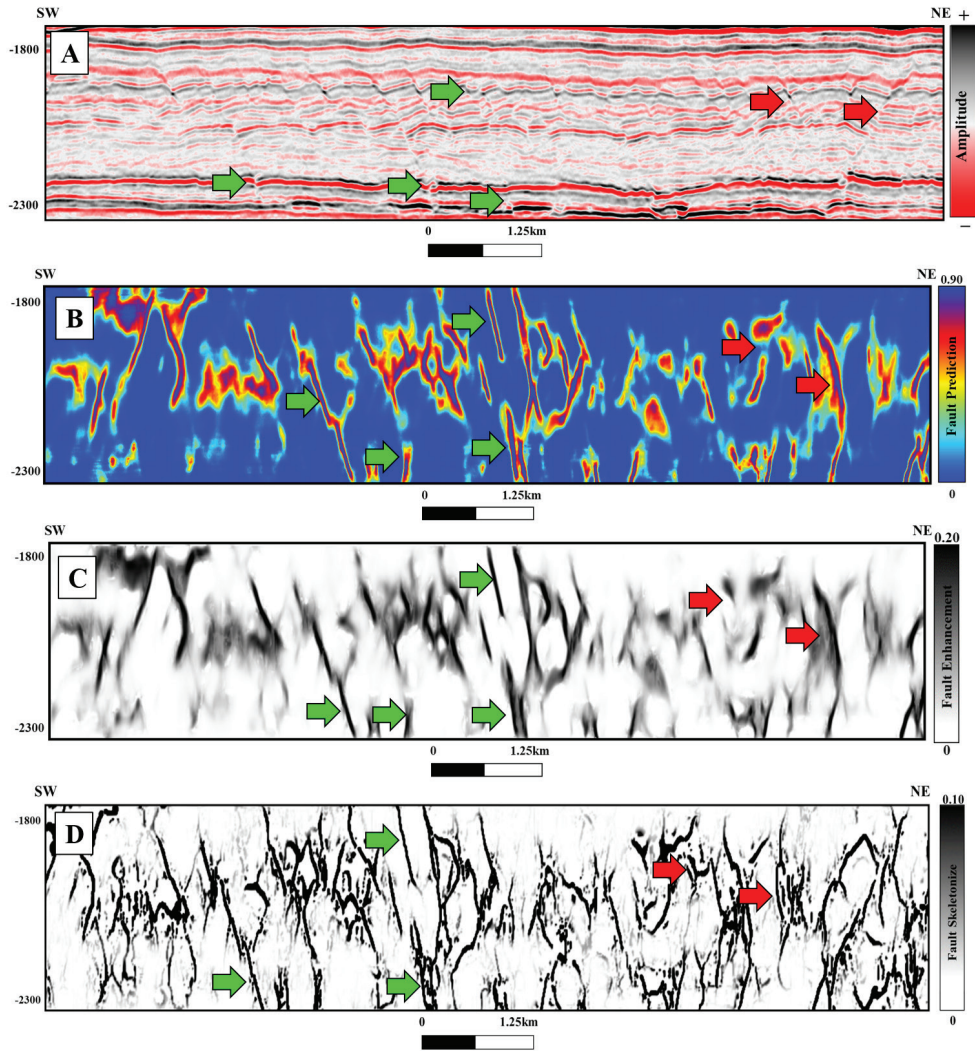


Figure 4. Vertical lines of (A) seismic amplitude, (B) fault prediction, (C) fault enhancement, and (D) fault skeletonize showing good picks, marked by green arrows, and bad picks, marked by red arrows.

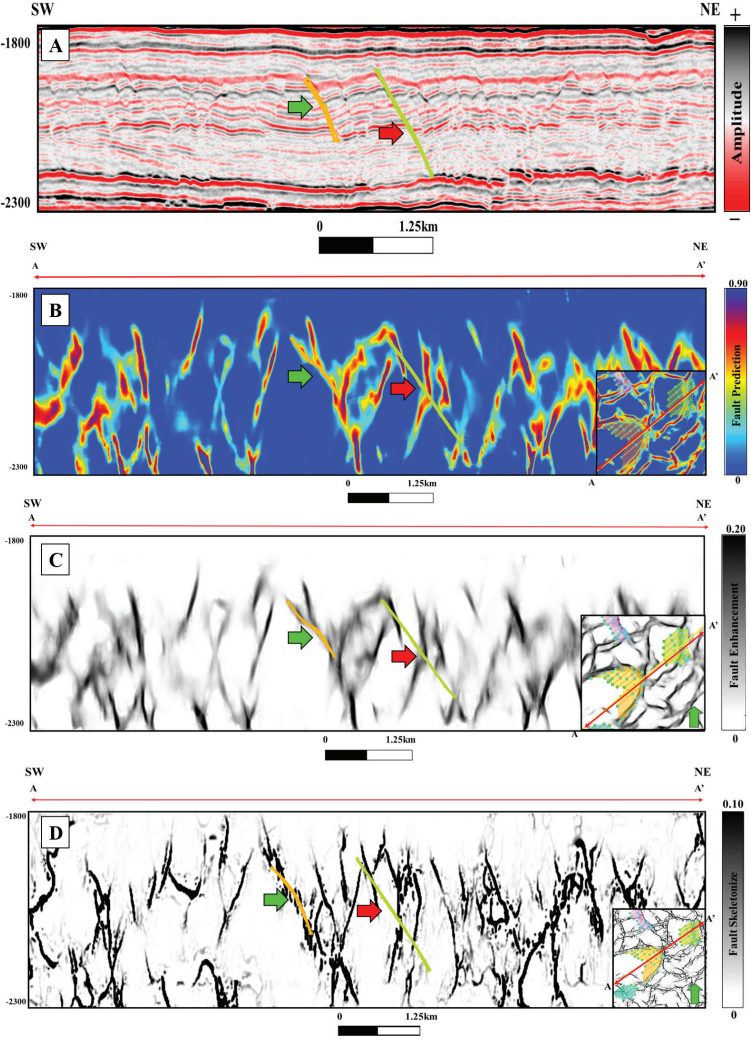


Figure 5. Vertical cross-sections illustrating colored faults manually picked by the interpreter versus fault prediction generated by the algorithms with arrows indicating good picks, green arrows, and bad picks, red arrows. (A) Seismic amplitude, (B) fault prediction, (C) fault enhancement, (D) skeletonize.

◀ Continued from previous page

the machine and draw comparisons and conclusions. We applied CNN for fault prediction on a section of our seismic volume characterized by the presence of polygonal faults that are Paleocene-Eocene age. Polygonal faults are non-tectonic extensional normal faults with low displacements (throw more than 100 meters). They are usually developed at the initial stages of consolidations of sediments and form a polygonal arrangement in plan-view, hence their name. These faults are most of the time linked to the dewatering and compaction of fine-grained sediments, and they play a vital role in fluid migration because they have the potential to reduce or enhance the sealing capacity of successions.

Figure 4 shows a comparison of a vertical seismic segment with the equivalent segments from fault prediction, fault enhancement and fault skeletonization processes. Figure 4A displays a simple seismic amplitude cross-section of the polygonal fault area with green arrows indicating good picks and red arrows indicating bad picks. If we observe the exact figure, we notice that the other algorithms look slightly different (figure 4). However, the most important feature that we need to notice from figure 4 is where and how each algorithm picks the faults compared to the seismic amplitude volume. Some areas, marked by the green arrows, have good picks, meaning that the same faults seen on the seismic amplitude volume were picked by the machine on the other three equivalent algorithms successfully.

Nonetheless, in other areas, the machine could not pick some discernible faults that are present on the seismic amplitude volume (figure 4B, C, D), leaving us with questions about the accuracy of

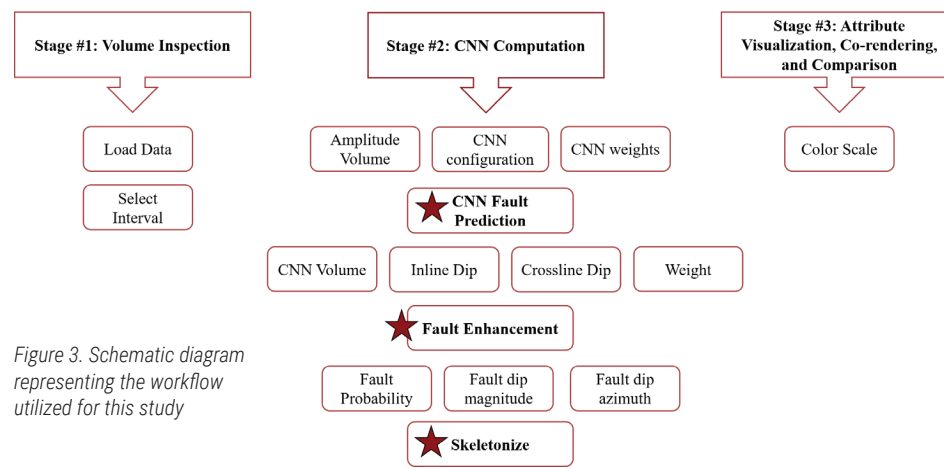


Figure 3. Schematic diagram representing the workflow utilized for this study

the algorithms when picking these highly complex networks of faults. Moreover, looking back at figure 4A, the seismic amplitude volume, we notice that there are various faults with similar dips spread

throughout the vertical line that the other algorithms do not seem to catch at all (figure 4 B, C). The only algorithm closer to representing what we observe on the amplitude volume is CNN skeletonize (figure 4D).



Christ Ramos recently received her bachelor's degree in geology with academic distinctions from the University of Oklahoma. She has worked as an undergraduate research assistant since her junior year under the guidance of Heather Bedle for the Subsurface Detective Agency and the Attribute Assisted Seismic Processing and Interpretation consortiums. Currently, she accepted a fully sponsored offer from the geosciences master's program at the University of Massachusetts Amherst. She will be conducting research under Michele Cooke on the evolution of fault systems by directly documenting fault growth with physical experiments and the use of numerical models to investigate fault propagation within simple and complex systems in southern California. She is an active member of AAPG, The Geological Society of America, and the Society of Exploration Geophysicists.



Heather Bedle received a bachelor's in physics from Wake Forest University and then worked as a systems engineer in the defense industry. She later received a master's and a doctorate from Northwestern University. After nine years of working with Chevron, she Instructed at the University of Houston for two years and is now an assistant professor at the University of Oklahoma. Her primary research interests merge seismic interpretation, rock physics, and attribute analysis. She is an active member of AAPG and the Society of Exploration Geophysicists.

If we evaluate figure 5, we have a vertical seismic amplitude cross-section with the equivalent segments of the other algorithms against the location of faults interpreted by the user. The user interpreted the faults only by using the amplitude volume and indicating the discontinued reflectors and changes in the dip as faults (figure 5A). Figure 5B shows a CNN fault prediction vertical line against two manually picked faults. As we notice, in figure 5B, the southwestern fault is relatively close to the location and dip that the machine predicted. That is to say, it is a good pick. However, the northeastern fault is a bad pick because the manually picked fault does not have a close location and dip as the fault interpreted by CNN fault prediction. The algorithm shows that the wrongly predicted fault location has a complete opposite dip compared to the user-interpreted fault on the seismic amplitude volume, or it is not even close.

We have CNN fault enhancement volume on the next cross-section against the same previous manually picked faults (figure 5C). As shown in figure 5C, the northeastern fault once more does not correspond to the location and dip of the fault interpreted by the machine. Hence, it is a terrible pick. The last volume generated by using a convolutional neural network is fault skeletonize (figure 5D). Fault skeletonization uses the CNN fault prediction and fault enhancement volume plus fault dip magnitude and dip azimuth to generate a sharper volume that picks on every discontinuous structure and unconformities (figure 3).

Nonetheless, the algorithm fails to characterize the same northeastern fault that the interpreter manually picked. Moreover, if we look closely at the seismic amplitude volume, there are many identifiable faults that an interpreter will pick due to the discontinuous reflectors

and changes in dip (figure 5A). However, upon observing the equivalent vertical section for all the other algorithms. We notice that none of the algorithms pick on these faults (figure 5B, C, D). Is it because they are too small for the machine to see them? Or because the algorithm only picks on the most prominent faults? Or is it due to the difference in the training data used in comparison to these types of faults that we want to pick?

An Excellent Tool, With Limitations

Machine learning is an excellent tool that saves time, speeds up lengthy processes and facilitates interpretations, but like every other suitable tool, it comes with its limitations. That is why it should be a common practice always to scroll, perform quality control throughout the data, and make sure the algorithm worked correctly and represents reasonable results that match your intuition or interpretation before using it for final purposes. While we believe that machine learning is a potential tool that can be used for different purposes, it will not completely replace the in-context geologic knowledge and interpretations made by experts in the area of interest. ☒

(Editors Note: The Geophysical Corner is a regular column in the EXPLORER, edited by Satinder Chopra, Founder and President of SamiGeo, Calgary, Canada, and a past AAPG-SEG Joint Distinguished Lecturer.)