Application of Transfer Learning in Fault Detection using U-Net Semantic Segmentation

Shruti, Sumeet Harshad Mavani, Sanjai Kumar Singh, GEOPIC*
Oil and Natural Gas Corporation Limited
Email: shruti@ongc.co.in

Keywords

Convolutional Neural Networks, Transfer Learning, Semantic Segmentation, U-Net, Pre-trained models

Abstract

This research paper aims to provide an overview and application of transfer learning in convolutional neural networks, exploring its underlying principles and applications. Transfer learning involves reusing pre-trained models, which have already been trained on large-scale datasets, and fine-tuning them to new tasks or domains with limited labeled data. In this paper, we have a pre-trained CNN U-Net model, and we have implemented the concept of transfer learning to test the scope of improvement in the performance of our fault detection model using seismic stack data. Various experiments and fine-tuning were carried out, and the best-performing model was used to predict the fault likelihood in 3D space to aid in further fault interpretation.

Introduction

Fault interpretation is of prime importance in the E&P industry since faults play a major role in hydrocarbon migration and entrapment. There has been significant development of tools for computer-aided fault detection and the majority of them are based on the use of seismic attributes. Seismic attributes usually require an expert interpreter to interpret the faults manually. Another way is to train a machine learning (ML) model using a few seismic sections and respective fault labels or annotations. The trained model can then be deployed to generate the fault likelihood volumes to be used for geological fault interpretation.

In traditional machine learning approaches, models are built and trained from the ground up, often requiring substantial labeled data and significant computational resources. However, in most real-world use cases, collecting and annotating large amounts of data is expensive, time-consuming, or even impractical. This limitation is particularly observed in the seismic domain, where creating comprehensive labeled datasets is very challenging due to the high cost of manual annotation. In addition to that, training a deep convolutional model from scratch every time on our continuously expanding datasets is highly computationally expensive and requires a significant amount of time and effort.

Transfer learning has surfaced as a powerful technique in the field of computer vision, specifically in the context of convolutional neural networks. It offers a solution to these challenges by leveraging the knowledge gained from pre-trained models on similar tasks or datasets.

Transfer learning offers several advantages. Firstly, it accelerates the training process by reducing the time and computational resources required for convergence. Secondly, it enables effective learning even when the new dataset is small or lacks to incorporate diversified patterns. By making use of these pre-trained models, transfer learning can mitigate the risk of overfitting and improve the generalization of the model. All in all, it enables effective learning even when the annotated data is limited. We can effectively bootstrap the learning process and benefit from the generalization capabilities of pre-trained models.

In the following sections, we will briefly elucidate the background of the U-Net model and semantic segmentation. Following the literature survey, we will explain the steps involved in the transfer learning process, the structure of the pre-trained model that we have used in our experiment, and the adopted methodologies, shedding light on its performance in comparison to our base model.

Literature Survey

Semantic segmentation is a computer vision task that involves assigning semantic labels to each pixel in an image, enabling the precise delineation and understanding of object boundaries and regions within the image. Unlike image classification, which assigns a single label to the entire image, semantic segmentation aims to provide a pixel-level understanding, where each pixel is labeled with a specific class. The goal of semantic segmentation is to partition an image into coherent regions based on its semantic meaning. This means that pixels belonging to the same category will be assigned the same label, allowing for detailed analysis and interpretation of the image. Figure 1 shows the different image classification and segmentation methods used in computer vision tasks.

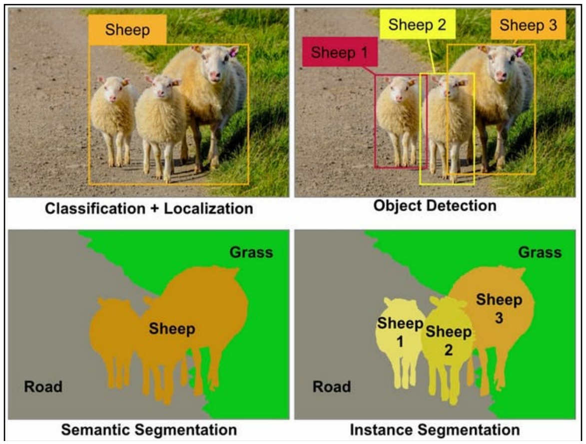

https://nirmalamurali.medium.com/image-classification-vs-semantic-segmentation-vs-instance-segmentation-625c33a08d50

Figure 1: Various image classification and segmentation methods

U-Net is a popular architecture for semantic segmentation in deep learning. It was proposed by Olaf Ronneberger, Philipp Fischer, and Thomas Brox in 2015, primarily designed for biomedical image segmentation tasks. The U-Net architecture is known for its U-shaped encoder-decoder structure, which allows for the precise localization of objects and high-resolution segmentation maps. Figure 2 shows the architecture of the U-Net model deployed for the fault detection task.

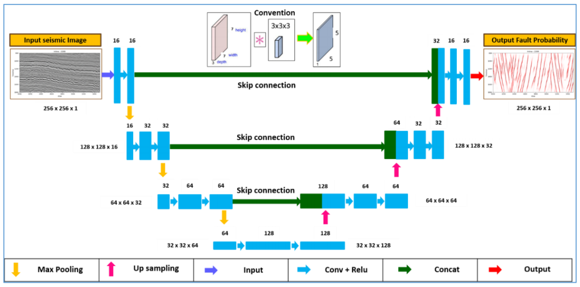

Figure 2: U-Net architecture used for the fault detection task

The U-Net architecture consists of two main parts: the contracting path (encoder) and the expanding path (decoder). The contracting path consists of successive convolutional and pooling layers that shrink the spatial dimensions of the input image, while increasing the number of channels, thereby capturing high-level information. On the other hand, the expanding path consists of upsampling and transposed convolutional layers to gradually restore the spatial resolution, creating a detailed segmentation map. A significant characteristic of U-Net is the skip connections that connect corresponding layers in the contracting and expanding paths, thus combining the local information captured in the contracting path with global information from the expanding path. This allows the model to have a better understanding of fine-grained details, leading to improved segmentation accuracy.

Methodology and Workflow

Transfer learning can be broadly implemented in two ways. The first method is to keep the model architecture intact and unfreeze only certain layers for re-training the model. The second method involves adding new trainable layers to the base model, thereby modifying the architecture of the base model. We have incorporated the following workflow into our transfer learning process:

1. Selection of a base model:

Base models are typically trained on large-scale datasets, having learned to extract general features from the training images. The base model that we have adopted for our experiment is a U-Net architecture with 76 layers consisting of 1 input layer, 1 output layer, 18 convolutional layers, 4 transposed convolutional layers, 18 batch normalization layers, 18 activation layers, 8 max-pooling layers, and 8 dropout layers.

2. **Generation of training data:**

To train the U-Net model, a pair of images were generated corresponding to the feature and label as shown in Figure 3. Each image was gray-scaled, having a size of 256x256 pixels. The labels were manually annotated, and the training dataset of images and labels was created using a self-developed module named "Image Data Generator" to maintain consistency of dimensions and pre-processing applied on input images. In addition to the standard image generation, the images, and labels in the training dataset were also augmented to a rotation of 3 different angles – 11, 45, and 135 degrees to increase the number of training images. Two datasets were generated during the experiment viz., for training the base model and re-training the base model.

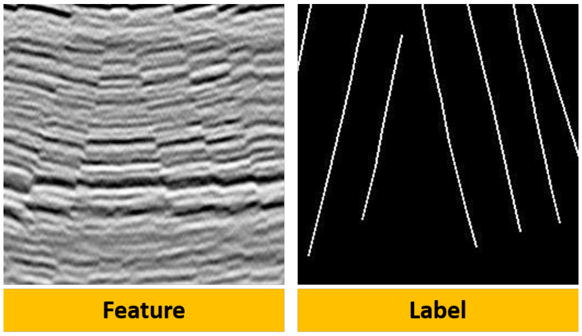

Figure 3: Example of training data (feature and associated label)

The first set of training data consisted of a pair of 10758 images (seismic section and respective annotated fault labels) from 4 different sedimentary basins of India. This dataset was used to train the base model.

The second set of training datasets consisting of a pair of 6000 images was generated from an unseen geographical area. As usual, this dataset was also segregated into the standard 80:20 ratio, with 80 percent of data involved in training, and 20 percent reserved for validation tests.

3. **Fine-tuning:**

It involves updating the parameters of the pre-trained model to accommodate the new dataset. It is common to freeze the early layers of the pre-trained model during transfer learning, especially when the new dataset is small, to prevent overfitting and preserve the learned low-level features. Fine-tuning the later layers, including the fully connected layers, allows the model to adapt and specialize in the specific tasks.

Observations

We extensively tested numerous transfer learning models and their performances to gain some insights into the impact of freezing certain layers over others. During the above experiments, a combination of different layers was unfrozen, and their impact on the validation dataset was observed. The input and the output layers were frozen and kept the same as that of the base model. All the models generated using transfer learning were compiled using parameters similar to that of the base model. They were first validated on the transfer learning dataset itself and then tested on a completely blind dataset of 60 images.

Figure 4(a) shows a sample image extracted from the KG basin that has been taken for presenting our observations. Figure 4(b) is its corresponding label, which has been marked by a domain expert. Figures 4(c) and 4(d) show the prediction results of the base model and the optimal transfer learning model, respectively.

Figure 5 shows the predictions of various transfer learning models on a completely blind image, taken from a deep marine Andaman area. These models were generated by changing the number of unfrozen layers while re-training the model. Since different numbers of layers were unfrozen in the model used for transfer learning, different outputs were generated in each case, as evident in Figure 5.

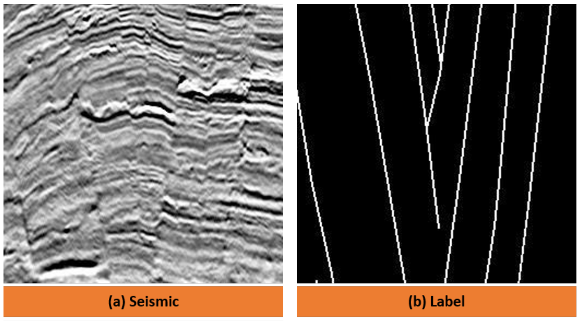

(a) Seismic

(b) Label

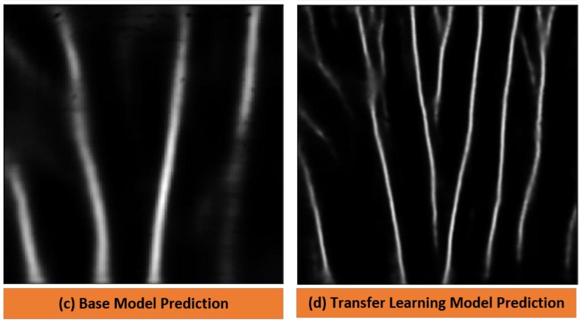

(c) Base Model Prediction

(d) Transfer Learning Model Prediction

Figure 4: (a) sample seismic from the training dataset; (b) its corresponding label; (c) prediction using the base model; (d) prediction using the transfer learning model

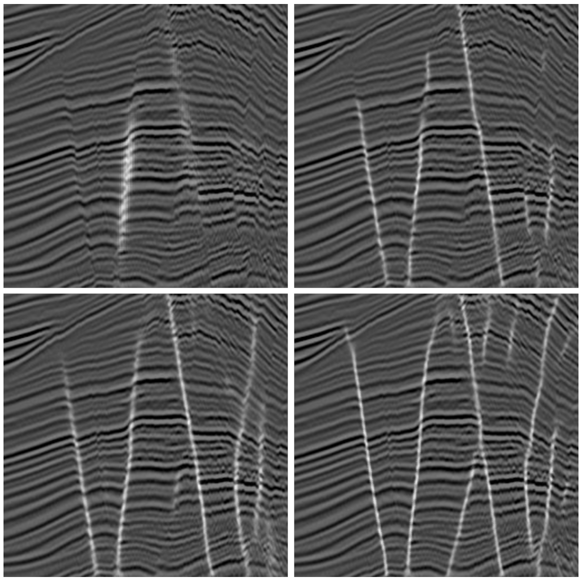

Figure 5: Model performance on a blind image upon unfreezing different layers

Results

In this section, we present the key findings of our experiment, which aimed to investigate the performance of a transfer learning model in our fault detection problem.

Upon unfreezing the right numbers of convolutional, activation, and batch normalization layers, we obtained the optimal predictions of our transfer learning model. 96.443% of accuracy was achieved on validation data, which was an improvement of ~5% from the base model. Finally, we applied the best-performing transfer learning model on a completely blind seismic section and compared the results of its performance with the base model. Figure 6(a) shows the faults identified by the base model, while Figure 6(b) shows the faults identified by our transfer learning model.

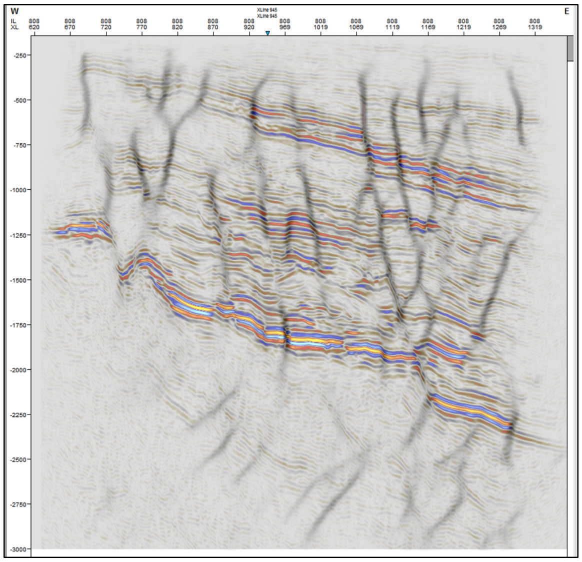

Figure 6(a): Fault detection using the base model

It is evident from Figure 6(b) that the transfer learning model has performed extremely well compared to the base model.

Conclusion

In conclusion, our experiment provides compelling evidence regarding the effectiveness of transfer learning in semantic segmentation for geological fault identification. The results demonstrate improved clarity and crispness of faults in comparison to the results of the base model. These findings highlight the potential of transfer learning technique in identifying discontinuities using a deep learning model, without having to retrain the model from scratch.

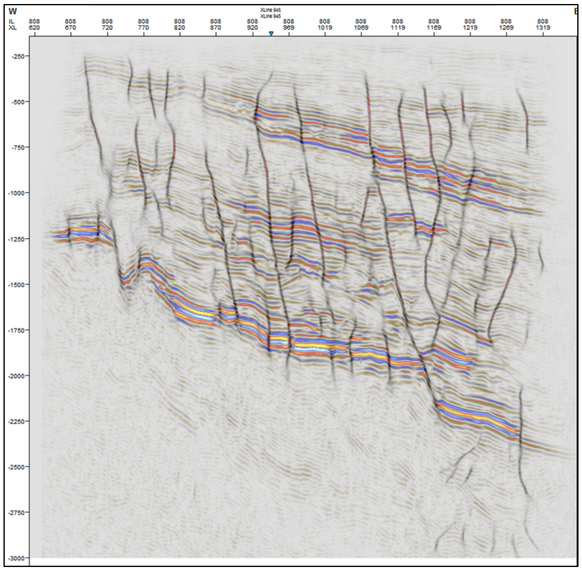

Figure 6(b): Fault detection using the transfer learning model

Further, the transfer learning model has a high potential to perform even better in other geographical areas as well, given that it has been trained on a small labeled training dataset from the same area beforehand.

References

Xinming Wu, Yunzhi Shi, Sergey Fomel, and Luming Liang (2018) Convolutional neural networks for fault interpretation in seismic images

Axelle Pochet, Pedro Henrique Bandeira Diniz, Hélio Lopes, Marcelo Gattass (2018), Seismic Fault Detection Using Convolutional Neural Networks trained on Synthetic Post stacked Amplitude Maps

Ronneberger et al., 2015, U-Net: Convolutional Networks for Biomedical Image Segmentation, MICCAI

Acknowledgments

The authors are thankful to the management of Oil & Natural Gas Corporation Limited (ONGC) especially GEOPIC for providing the opportunity and dataset to carry out this project. The authors would also express their gratitude to all the staff posted at GEOPIC for their support extended during the various stages of the project in the form of technical, software, hardware, and human resources.

The views expressed in this paper are solely those of the authors and need not necessarily be that of ONGC.