Submission code: AYL7X5K4D9

See this and other abstracts on our website: https://home.sbgf.org.br/Pages/resumos.php

Interpreting acoustic drapes on real seismic data using fine-tuned deep learning approaches

Ramon Lopes, Marcos Silva (Petrobras), Matheus Ferreira, Haroldo Júnior, Gleizer Silva (Petrobras), Anderson Vieira, Lucas Gondim Miranda (Cenpes-Petrobras)

Interpreting acoustic drapes on real seismic data using fine-tuned deep learning approaches

Please, do not insert author names in your submission PDF file.

Abstract Summary

The interpretation of seismic sections is a crucial task in the oil and gas industry. Traditionally, this interpretation relies on visual recognition, making it time-consuming and potentially influenced by the varying expertise of observers. While there are machine learning approaches in the literature for identifying facies in seismic data, to the best of our knowledge, research focusing on the identification of acoustic drapes in seismic data is lacking. To fill this gap, we fine-tuned two deep learning architectures using real seismic data. Experiments attest the effectiveness of considered approaches, where accuracy levels exceeded 93%. Therefore, we claim that fine-tuned deep learning models can assist geoscientists in the identification of acoustic drapes in seismic data.

Introduction

The analysis of seismic sections is important for understanding structural and stratigraphic elements in a region of interest, such as in oil exploration areas. Since this is usually a visual task, this analysis is time-consuming and its outcome depends on the subjectivity and experience of the interpreter. In light of this, to mitigate costs, it is possible to train deep learning architectures to assist geoscientists in various seismic interpretation tasks. Next, we present some relevant works in the context of segmentation in seismic data.

Milosavljević (2020) introduced a method based on convolutional networks in the context of salt deposits. In the same direction, Hu et al. (2022) presented a method based on convolutional networks for fault segmentation using 2D seismic sections. In turn, Kaur et al. (2023) investigated convolutional networks and adversarial networks for facies segmentation. Based on our literature review, to the best of our knowledge, there is a lack of works addressing the segmentation of acoustic drapes in Sub Bottom Profiler (SBP) sections in deep-water environments. For example, Figure 1 illustrates a typical SBP section in deep water in the Campos Basin, where we can visualize the layers and the acoustic drape.

The geological drape corresponds to the Quaternary superficial cover, which overlays most of the area, composed of hemipelagic muddy sediments. The use of acoustic geophysical technology may be the first approach adopted for the investigation of marine environments, which are difficult to access and costly to investigate due to their geological nature. Technological advances in acquisition and processing using high-frequency tools have expanded the ability to distinguish and depict the initial layers of the seabed sedimentary cover. As a result, we can highlight a more robust acoustic characterization of the geological drape, as well as the provision of data for submarine projects that allow cost reduction, mitigation, and optimization without compromising environmental safety and being in compliance with legal requirements.

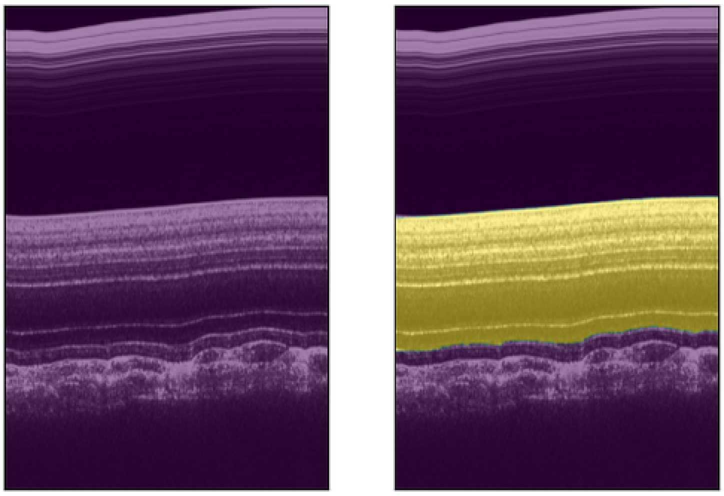

Figure 1: Seismic section and acoustic drape identified by an interpreter.

The understanding of the initial sedimentary layers of the seabed is essential for the use and occupation of these environments since it is a source of information about the local geological history. For example, its comprehension plays a crucial role in the assessment of geological risks, support for environmental licensing and engineering projects related to submarine activities for exploratory purposes in oil fields. Namely, the evaluation of seabed composition and thickness is crucial to determine slope stability and to support projects for the installation of submarine equipment.

In this work, we aim to answer the following question: can deep learning based segmentation models assist geoscientists in the task of interpreting acoustic drapes in SBP sections? To this end, we fine-tuned three semantic segmentation methods using 15 private seismic sections interpreted by geoscientists. In our experiments, the considered models achieved accuracy rates ranging from 93% to 98%. Thus, we claim that fined-tuned deep learning models can assist geoscientists in their interpretation tasks in the context of acoustic drapes.

Deep Learning Architectures

In this section, we discuss the deep learning architectures considered in this work.

U-Net. Ronneberger et al. (2015) proposed an encoder-decoder architecture for semantic segmentation tasks. The encoder - the contraction path - consists of repeated applications of convolution and downsampling operations to extract features from the image. In turn, the decoder - the expansion path - consists of upsampling operations, convolutions, and concatenations to restore the image resolution. Experiments attested the U-Net's robustness, especially in semantic segmentation in the context of medical images, which have low contrast.

Sismicanet. This is a private model trained in our company using 50 terabytes of seismic sections. In the backbone, we adopted DINOv2 (Oquab et al., 2024) as the foundational model and UPerNet (Xiao et al., 2018) for fine-tuning. Like the U-Net, the UPerNet also employs an encoder-decoder architecture to perform semantic segmentation. However, it features important structural distinctions compared to the U-Net, which makes it able to simultaneously capture a larger set of attributes of the image to be segmented.

Experiments and Results

In this section, we present the model training details and the results. In this context, we conducted experiments to address the following questions:

Q1 How effective are the pre-trained models for the segmentation task?

Q2 How efficient are the pre-trained models regarding the computational resources expended?

Based on these two questions, one can balance the trade-off between performance and the resources required for training and performing inference in production.

The dataset used consists of SBP sections collected from a remote data acquisition platform using an Autonomous Underwater Vehicle in deep waters of the Campos Basin. Specifically, the SBP seismic lines were collected at a relatively fixed altitude above the seabed using a Chirp source system that emits a pulse with a modulated frequency between 2-16 kHz. First, the dataset was randomly split into train, validation, and test sets, each containing 15, 5, and 7 seismic sections, respectively. The train set was then used to fine-tune all the models, while the validation set was used to select the best hyperparameter configuration for the model. Finally, the test set was used to evaluate the quality of the segmentation. To evaluate the segmentation quality, we employed Intersection Over Union (IoU), also known as Jaccard, which measures the overlap between the region segmented by the network and the true region in the seismic data.

For the sake of reproducibility, we present relevant fine-tuning information. First, each seismic section was split into non-overlapping patches with the original height and a maximum width of 512 to cover the entire image. In the case of the U-Net, we employed Adam as the optimizer with a learning rate of 1×10^{-3}, reduce on plateau as the learning rate scheduler with default configuration, 50 epochs for training, a batch size of 32, the arithmetic mean between binary cross-entropy and dice as the loss function, and finally a set of augmentations from the Albumentations library:[1] Resize to 512×512, HorizontalFlip, GaussianBlur, and RandomBrightnessContrast. Meanwhile, the UPerNet was trained using the mmsegmentation framework[2] for 16,000 iterations, with our proprietary pre-trained weights being loaded into a ViT-Adapter backbone (Chen et al., 2023), with a batch size of 2, AdamW as the optimizer with a learning rate of 6×10^{-5}, binary cross-entropy as the loss function and finally a set of mmsegmentation augmentations: Resize, RandomCrop, RandomFlip, and PhotoMetricDistortion.

Table 1 presents the results obtained on the test set. We report the weights used in each architecture; for example, we utilized respectively Resnet50 and Imagenet weights as the encoder and decoder of the U-Net. To evaluate whether the differences found are statistically significant, we performed a paired two-tailed T-test despite the low number of seismic sections. Thus, we display the symbol • to denote a significant difference at the $p < 0.01$ level compared to the UPerNet architecture with SismicaNet weights.

UPerNet with SismicaNet weights, which was trained in our company using seismic data, shows significantly better results compared to U-Net. In contrast, there is no statistical difference compared to Imagenet, which was originally trained using web-collected images. With respect to the computational burden, the number of parameters in UPerNet is approximately 1,118% higher compared to U-Net, while the former provides a 5% improvement in segmentation quality. Thus, we can conclude that there is a trade-off between computational resources and segmentation quality in light of the incurred costs and available resources by the practitioner.

[1] `https://albumentations.ai/`
[2] `https://github.com/open-mmlab/mmsegmentation`

Table 1: Results on test set. The best result is underlined.

Architecture	Weights	Parameters	IoU (%)
U-Net	Resnet50 / Imagenet	32.569.937	• 93,07
UPerNet	Imagenet	364.169.540	97,75
	SismicaNet		97,82

Conclusion

In this work, we studied the segmentation of acoustic drapes in seismic sections using deep learning architectures. Specifically, in our experiments, we considered the U-Net and UPerNet architectures, the latter with different instantiations regarding the network weights. Based on our experiments, the UPerNet outperformed the U-Net by 5% but at the cost of a substantial increase in the number of network parameters, which turns out to require greater computational resources. Therefore, the results suggest that fine-tuning models for the task at hand is feasible. Furthermore, the U-Net can offer a lower computational burden while still being an effective architecture for those who do not have robust computational resources or financial means to acquire proprietary software. Thus, we believe that this work can motivate future investigations into the fine-tuning of models available in the literature for segmentation in seismic sections to assist geoscientists in interpretation tasks.

References

Chen, Z., Y. Duan, W. Wang, J. He, T. Lu, J. Dai, and Y. Qiao, 2023, Vision transformer adapter for dense predictions.

Hu, G., Z. Hu, J. Liu, F. Cheng, and D. Peng, 2022, Seismic fault interpretation using deep learning-based semantic segmentation method: IEEE Geoscience and Remote Sensing Letters, **19**, 1–5.

Kaur, H., N. Pham, S. Fomel, Z. Geng, L. Decker, B. Gremillion, M. Jervis, R. Abma, and S. Gao, 2023, A deep learning framework for seismic facies classification: Interpretation, **11**, T107–T116.

Milosavljević, A., 2020, Identification of salt deposits on seismic images using deep learning method for semantic segmentation: ISPRS International Journal of Geo-Information, **9**.

Oquab, M., T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, 2024, Dinov2: Learning robust visual features without supervision.

Ronneberger, O., P. Fischer, and T. Brox, 2015, U-net: Convolutional networks for biomedical image segmentation: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Springer International Publishing, 234–241.

Xiao, T., Y. Liu, B. Zhou, Y. Jiang, and J. Sun, 2018, Unified perceptual parsing for scene understanding: Proceedings of the European conference on computer vision (ECCV), 418–434.