A Study on the Use of Deep Learning for Detecting Subsurface Structures[1]

Luan Rios Campos[1], Peterson Nogueira Santos[2], Davidson Martins Moreira[3] and Erick Giovani Sperandio Nascimento*[4]

[1,2,3,4]Manufacturing and Technology Integrated Campus – SENAI CIMATEC
Salvador, Bahia, Brazil
[2]National Institute of Science and Technology for Geophysics of Petroleum – UFBA
Salvador, Bahia, Brazil
[1]luan.rios.campos@gmail.com, [2]nogueirapeterson@gmail.com,
[3]davidson.moreira@gmail.com, [4]ericksperandio@gmail.com

** Corresponding Author*

Abstract

Beneath the earth there are many structures, such as different types of rocks and salts. Among them are also hydrocarbons that are a valuable resource for the oil and gas industry. One way of studying sub surfaces is using seismograms, which offers a seismic-wave representation with many valuable information of the area. By studying the patterns within the seismic data one can generate a representation of the subsurface based on some parameters that are able to show each one of underlying structures, such as the velocity that the waves propagated. With the advancement of computer-related technology, such as multi-core processors and GPUs, the processing power of computers have increased and the possibility of working with a much larger amount of data and using new and more powerful computational techniques, such as deep learning, was made possible in a variety of fields. Recently, deep learning methods are being applied to solve many geophysical problems, including the estimation of subsurface structures based on the velocity parameter. This work shows an interdisciplinary approach to estimate velocity models from computer modeling seismograms of non-real sub surfaces using a supervised learning artificial intelligence technique. The results obtained can contribute much to the scientific community as it demonstrates how changes in the seismic data modeling process reflects in the velocity model estimation.

[1] The paper "A Study on the Use of Deep Learning for Detecting Subsurface Structures" was edited by Elsevier Language Editing Services. We would like to express our gratefulness to the reviewers Antônio José da Silva Neto, Reynam Pestana and Alexsandro Cerqueira who took some part of their time to share their comprehensive and detailed thoughts on this works and carefully alert its main issues.

Keywords: *System Sciences and Engineering; Interdisciplinary Research; Deep learning; Geophysics.*

1. Introduction

A variety of structures lie beneath the surface of the earth, including water and different types of rocks and salts. Among them are hydrocarbon fossil fuel resources such as natural gas and petroleum that are extremely valuable for the oil and gas industries because they form our primary source of energy. However, because the subsurface of the earth is vast, deep, and not easily accessible, the identification of such structures is expensive, time-consuming, and sometimes uncertain.

One way of studying the subsurface is by using seismograms, which provide seismic-wave representations with much valuable information on the area. These data are generated by propagating waves in the medium produced by either explosives or thumpers in land acquisitions or by air guns in marine acquisitions.

By studying the patterns within the seismic data, a representation of the subsurface based on some parameters such as the velocity of wave propagation (Al-Yahya, 1989) (Stork, 1992) is obtained. The representation shows each of the underlying structures. This approach is interdisciplinary as geologists and geophysicists rely on computational modeling tools to simulate the physical processes that occur during a seismic acquisition. This facilitates the analysis of seismograms and the development of new approaches to improve the understanding of subsurfaces.

With the advancement of computer technology, such as the development of multi-core processors and graphical processing units (GPUs), the processing power of computers has increased by many folds. This has opened the possibility of working with a much larger amount of data and the use of new and more powerful computational techniques, such as deep learning. Recently, artificial intelligence (Sajeva, 2016) (Datta & Sen, 2016) and deep learning methods have been applied to solve many geophysical problems, including the estimation of subsurface structures based on the velocity parameter (Lewis & Vigh, 2017) (Araya-Polo, Jennings, Adler, & Dahlke, 2018) (Wu, Lin, & Zhou,

2018) (Campos, Nogueira, & Nascimento, Estimating Initial Velocity Models for the FWI Using Deep Learning, 2019) (Campos, Nogueira, & Nascimento, Tuning a Fully Convolutional Network for Velocity Model Estimation, 2019).

This work presents an interdisciplinary approach to discuss the estimation of velocity models based in (Campos, Nogueira, Moreira, & Nascimento, 2019), where we presented a more technical discussion on how different parameters for modeling a wave propagation can affect the training process and the outputs of a deep learning model. An interdisciplinary discussion is fundamental to extend this work and make it available to a broader range of readers.

This paper is organized as follows. Section 2 presents the different contexts of seismic acquisition and modeling, and introduces the concept of the velocity model and why its resolution is important. Section 3 explains the definition of deep learning and how it is performed on data. Section 4 describes the methodology and experiments used in this work. Section 5 discusses the results obtained from our experiments. Section 0 concludes this work.

2. Seismograms and Velocity Models

The exploration of subsurfaces is a highly expensive process for the oil and gas industries. It begins with the study of an area and continues up to the drilling of an extraction well. The first step for studying a subsurface is to acquire seismic information on the area. This is a crucial step for hydrocarbon exploration. It is therefore extremely important to understand how the seismic data is produced. Generally, the acquisition process can either be performed along a line to produce a vertical profile (2D setup) or over the entire area to produce a volumetric data of the subsurface (3D setup).

The seismic acquisition process consists of the following: acoustic energy is usually generated by controlled sources such as explosives or thumpers on land and air guns in the sea. The energy then spreads in the subsurface as a spherical wavefront that propagates along all directions. The energy is reflected and transmitted by the different types of rocks that exist in the region. The reflected signals travel back to the surface and are captured by a set of receivers (geophones in a terrestrial acquisition and hydrophones in a marine acquisition) placed along a predefined geometry. These receivers can record components of the physical velocity vectors from the particle or pressure field (when using hydrophones). Figure 1 illustrates the marine seismic acquisition.

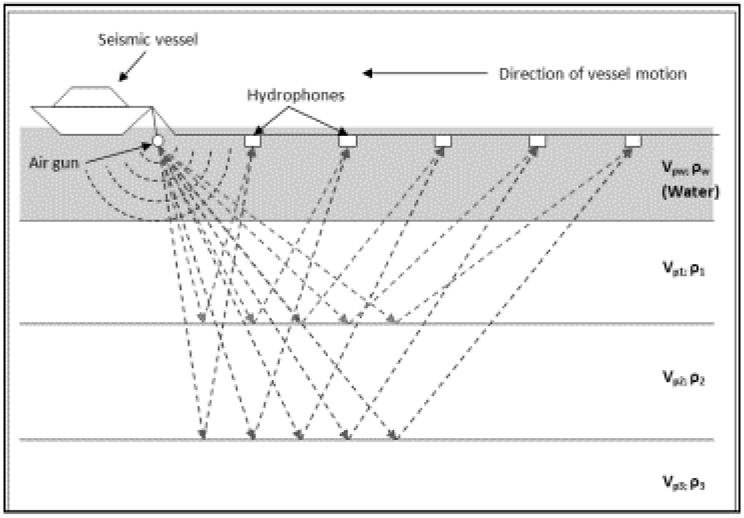

Figure 1: Marine seismic acquisition using a single seismic wave source.
(Committee, 2017)

Seismic modeling attempts to replicate the seismic acquisition process by using algorithms to simulate the sources responsible for the wave propagation and reflection and the wave capture by the receivers. Computers are used to simulate this process for the development, validation, and use of seismic imaging techniques, such as reverse time migration (RTM) (Baysal, Kosloff, & Sherwood, 1983) and full-waveform inversion (FWI) (Virieux & Operto, 2009) (Dos Santos, 2013), to analyze the seismic data in a less expensive and controlled environment. Such methods rely on velocity model representations of the subsurface to operate properly.

The subsurface is represented in a velocity model by the value of the velocity that the wave propagates with in the given region. The velocity values represent the underlying structures that lie in an area because the velocity of the seismic wave varies with the material, i.e., the wave propagates at different speeds in water, salt bodies, sandstones, shales, and other structures.

The higher the resolution of the velocity model, the more realistic is the representation of the subsurface. In other words, the values of the velocity are better defined in higher resolution models and thus, the structure can be better identified during post-processing analysis. Figure 2 shows an example of a smoothed model and a high-resolution velocity model.

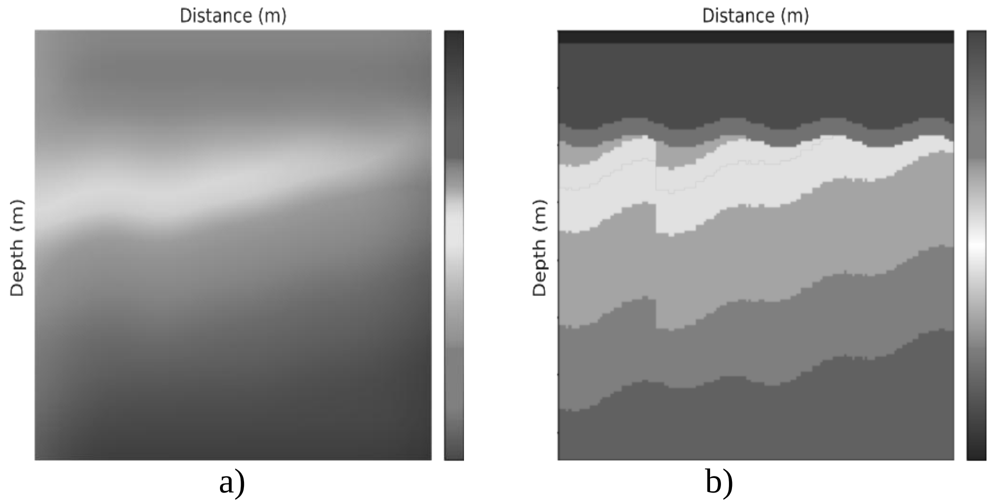

Figure 2: An example of a a) smoothed and b) high resolution velocity models

Velocity models are very important for the analysis of seismograms and decision making, but the production of models with high resolutions is challenging. This has stimulated researchers to experiment with different techniques, including artificial intelligence and deep learning methods.

3. Deep Learning Models

The production of a growing amount of information on a daily basis by all sorts of devices (smartphones, computers, IoT devices, etc.) has increased the need for the development of powerful tools and algorithms that can handle such quantities of data. This, together with the increasing power of computers and their capabilities for processing a wide variety of data, has contributed to the growth in the use of machine learning methods in recent years.

However, the use of machine learning techniques is limited to finding patterns in certain types of data, especially when the data are large, complex, and extremely non-linear. In such cases, for example, in handwritten recognition (LeCun, et al., 1990) and speech or image analysis (LeCun & Bengio, Convolutional networks for images, speech, and time series, 1995), deep learning models can be used to find complex patterns that could not have been found before at performance levels that sometimes even surpass that of human analysis (Krizhevsky, Sutskever, & Hinto, 2012).

Although sometimes confusing, the definitions of artificial intelligence, machine learning and deep learning are different. Artificial intelligence (AI) is defined by the Merriam-Webster dictionary as “a branch of computer science dealing with the simulation of intelligent behavior in computers” (Merriam-Webster, 2020). Machine learning is seen as a subset of AI where the idea is to **train** computer techniques that are able to learn useful representations from the input data (Chollet, 2017). Finally, deep learning is a subset of machine learning that emphasizes the learning capability by adding many layers in succession to extract even more meaningful representations from the data (Chollet, 2017). Figure 3 illustrates a diagram showing how deep learning differs from machine learning and from artificial intelligence approaches.

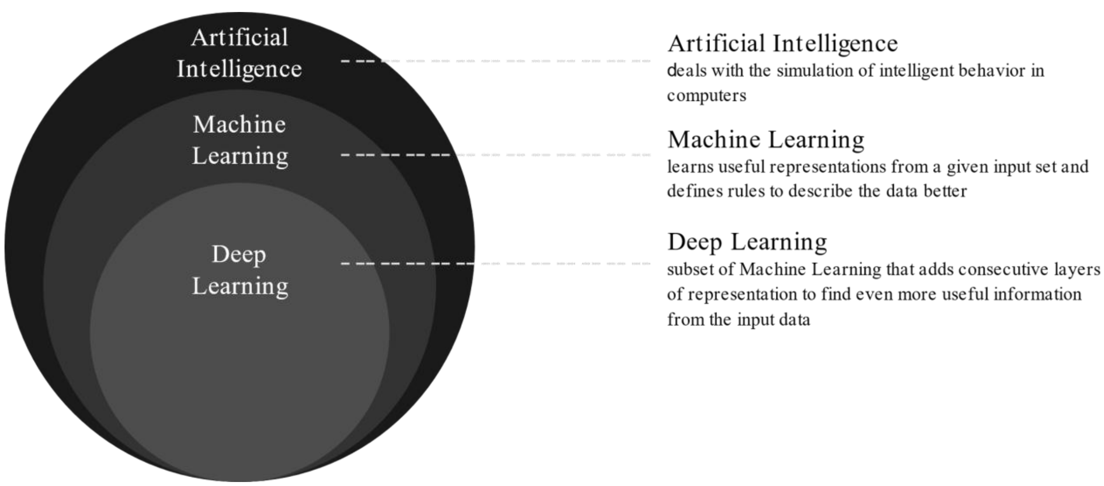

Figure 3: Differences among Artificial Intelligence, Machine Learning and Deep Learning

Deep learning models are usually variations of neural networks (Chollet, 2017). These networks comprise certain types of layers (dense[2], convolutional[3], recurrent[4], etc.) depending on the type of problem and the hyperparameters of the neural network, such as the guiding function for the training process, the

[2] A dense layer has all of its neurons connected to all the features of an input data
[3] A convolutional layer processes a certain region of the input data instead of all of it at once
[4] A recurrent layer has its output fed back to its input

minimization function, and the number of layers. The term “deep” refers specifically to the number of layers these neural networks have.

As the number of layers increases, the capability of the network to handle more complex inputs tends to increase accordingly because of its improved ability to handle the non-linearity in the data and consequently, its ability to detect more patterns to describe the problem in question. However, this may lead to an increase in the computational resource one needs to process such model. Thus, there must be a tradeoff of resources available, complexity of the deep learning model and size of the data to be processed.

Machine and deep learning models can be trained either in a supervised or unsupervised manner. The former (Figure 4) is typically used in regression or classification problems, where each input has a matching known output, whereas in the latter (Figure 5), the model finds patterns without any external guidance and relies only on the input data to form clusters of the most similar samples.

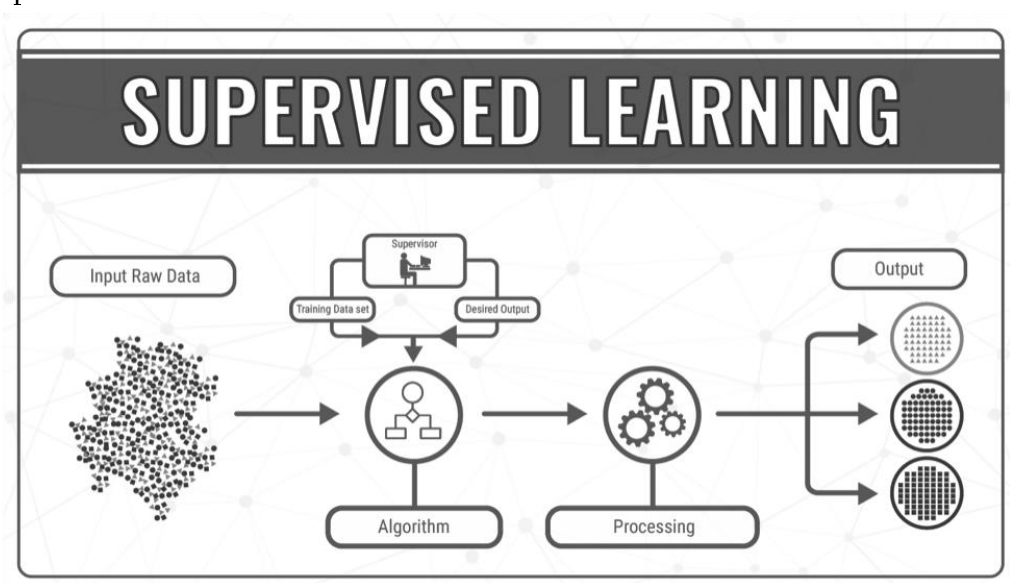

Figure 4: Supervised Learning process
(MC.AI, 2020)

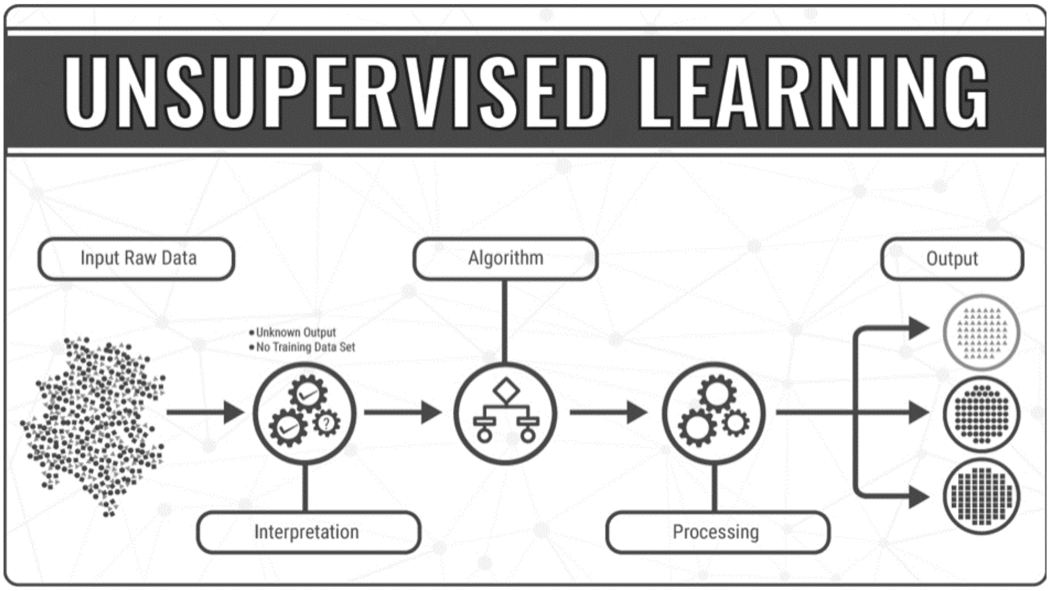

Figure 5: Unsupervised Learning process
(MC.AI, 2020)

4. Methodology and Experiments

We propose to treat the velocity model estimation problem as a regression problem and solve it using a supervised deep learning approach. The input is the seismic data and the output, i.e., the target, is a velocity model that represents the subsurface on which the seismic acquisition/modeling is performed. Therefore, our task can indeed be seen as a regression problem whose goal is to find a function that converts seismograms into velocity models.

To form our dataset, we first created 1,020 bi-dimensional synthetic velocity models of marine regions 3000 m in depth and 3000 m in length with the geophysically significant structures separated into layers, folded layers, rocks sliding between the layers (fault areas) and vertical variation in the velocity value[5]. An example of velocity model can be seen in Figure 6. The process of creating a set of velocity models is described with more details in Figure 7.

After the models were created, we applied seismic modeling to each of the models. The seismic modeling has many parameters that can be adjusted before starting the process. These parameters include the positions and numbers of

[5] The velocity within each layer is constant but varies between the different layers

sources and receivers, the dimensions of the subsurface, the timestep the wave travels before being recorded by the receivers and the peak frequency. In this work, both the sources and receivers were placed along the same geometrical axis during the modeling. We chose this configuration for the sources and receivers because this is a likely scenario during actual data acquisition when exploring unknown subsurfaces as illustrated in Figure 1.

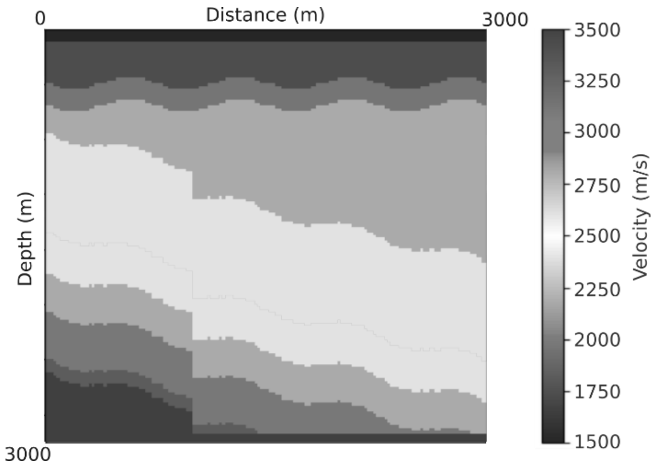

Figure 6: A synthetic velocity model containing 10 layers, folding, dipping and fault areas

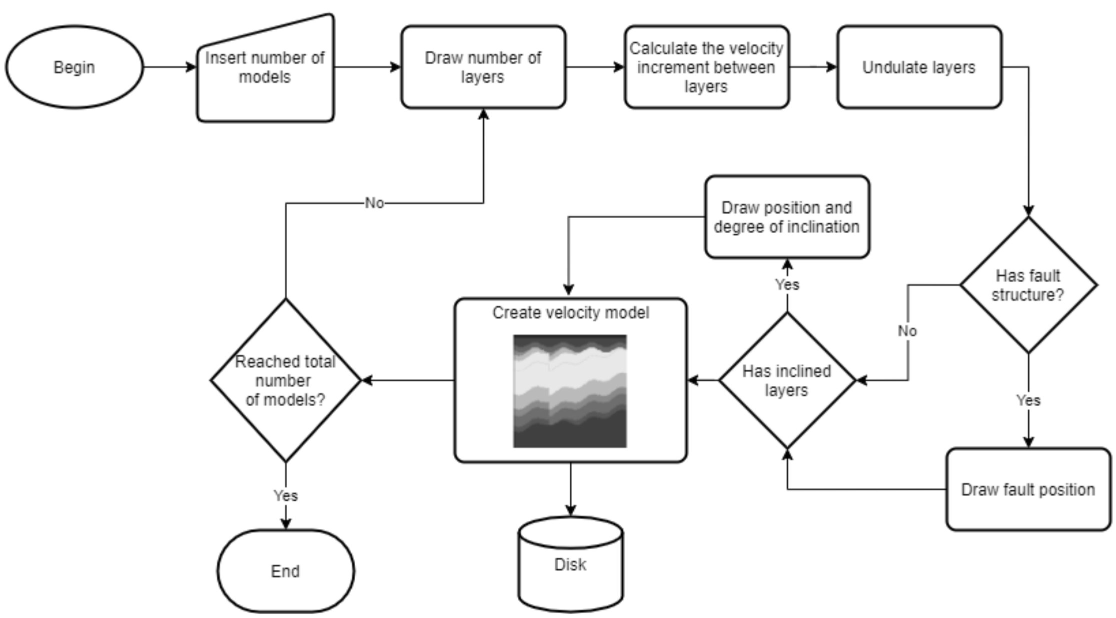

Figure 7: Flowchart illustrating the process for creating synthetic velocity models

Our experiments follow two strategies: in one, we vary the number of sources[6] and in the other, we vary the peak frequency[7] for the seismic acquisition. After generating the seismic data for a given velocity model, we create the pair of data that will be used in the deep learning model and after this is done for every velocity model create, we compose the dataset for that seismic modeling configuration as described in Figure 8.

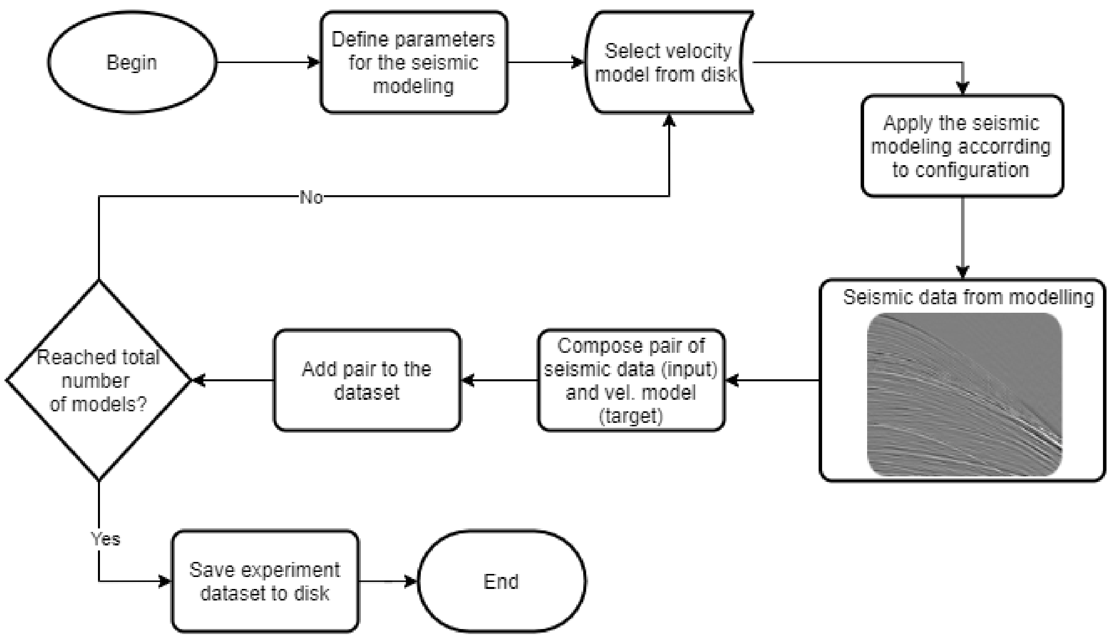

Figure 8: Flowchart that demonstrates how the seismic data is created and how a dataset for an experiment is composed

We use a fully convolutional network (Long, Shelhamer, & Darrel, 2015), specifically the U-Net architecture (Ronneberger, Fischer, & Brox, 2015), to perform the estimation of the velocity model. This network consists of two parts. The first part is called the encoder[8] and the second part is called the decoder[9]. Identifying both the important aspects of the seismic data and where they are located is essential for the velocity model estimation problem as the

[6] As the sources are responsible for propagating the waves, we expect to change the information redundancy contained in the seismic data in the first strategy because fewer sources generate less redundant information and more sources generate more redundant information.
[7] We expect to vary the amount of detail in the seismic data because the lower the frequency of the acquisition/modeling, the less details and complexity the seismic data has.
[8] The encoder reduces the size of the input and simplifies the input for the detection of usable information in the data.
[9] The decoder resizes the output of the encoder to the size of the output to determine the location of useful information.

seismic data contain location-sensitive information on the substructure, i.e., some parts of the seismogram represent certain regions of the subsurface better.

For each seismic modeling configuration, we trained a different instance of the deep learning model but preserved its architecture, i.e., the training of a model with a set of seismograms neither contributed to nor interfered with the training of the other models. In order words, the configuration of the models remained the same throughout every experiment, but each model was trained with a set of seismogram results from its own unique seismic modeling configuration separately from the other models, as shown in Figure 9.

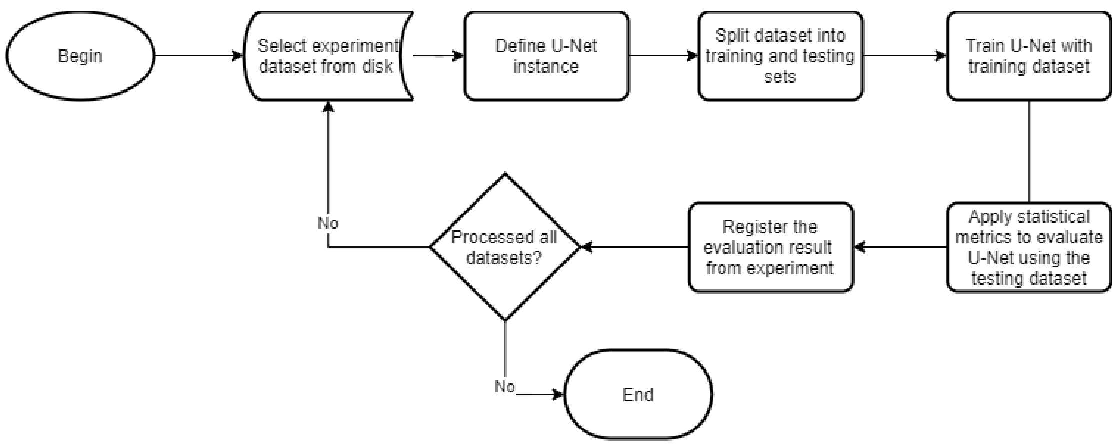

Figure 9: Flowchart that demonstrates the process for training and evaluating a U-Net instance for all experiments available

The evaluation of the estimations from the deep learning model relative to their ground-truth target counterparts was performed in two different ways. We first analyzed the estimation results qualitatively by comparing the plot of the ground-truth model with the plots of the estimated velocity models; these models were generated by the deep learning models trained with seismograms from different seismic modeling configurations as input. We then performed a quantitative analysis by calculating some statistical metrics to assess the quality of the regression methods where we now considered the entire data set instead of a single sample.

We used five statistical metrics: the mean squared error (MSE), the mean absolute error (MAE), the coefficient of correlation (Pearson r), the coefficient of regression (R^2), and the factor of two (fac2). The results should ideally show a value as close to 0 as possible for the first two metrics and as close to 1 as possible for the remaining metrics. The MSE was used as the function to be minimized by the deep learning model during training in order to guide the deep learning model to its optimal configuration.

5. Results

Our discussion starts with the graphical analysis of the results. We chose for this analysis a ground-truth velocity model with 8 folding layers and dipping layers, as depicted in Figure 10a. This velocity model also has a fault area indicated in the figure by a yellow ellipsis. The remaining velocity models in Figure 10 are the results of estimations performed by the U-Net using seismic data generated by different configurations of the seismic modeling described above as input.

We can infer from this analysis that the estimated velocity models c, f, and g in Figure 10 have the best deep learning results. They have well positioned layers, clearly visible folding and dipping layers, and fair estimates of the fault structure in the model. Furthermore, the velocity values are consistent with the ground-truth reference.

However, we cannot generalize that these models performed well for all the samples by analyzing only a single sample. The performance of the models for this sample might only be a fortuitous coincidence, and therefore, further analysis with more samples is needed. It is however impractical to manually check every estimate from all the experiments to evaluate the models. Hence, proceeding with a quantitative analysis is desirable as we can quickly extend the performance evaluation to all the samples in the dataset by applying reliable methods to determine the quality of the regression achieved.

Table *1* displays the results obtained from the U-Net in all the experiments and the time it took to completely train the deep learning model. It can be seen that the computational time taken to train the deep learning model was more sensitive to variations in the number of sources than variations in the peak frequency. This is because the size of the seismic dataset became larger as we increased the number of sources during the modeling/acquisition, whilst changing the peak frequency changed the values obtained during the modeling/acquisition, but not the size of the data.

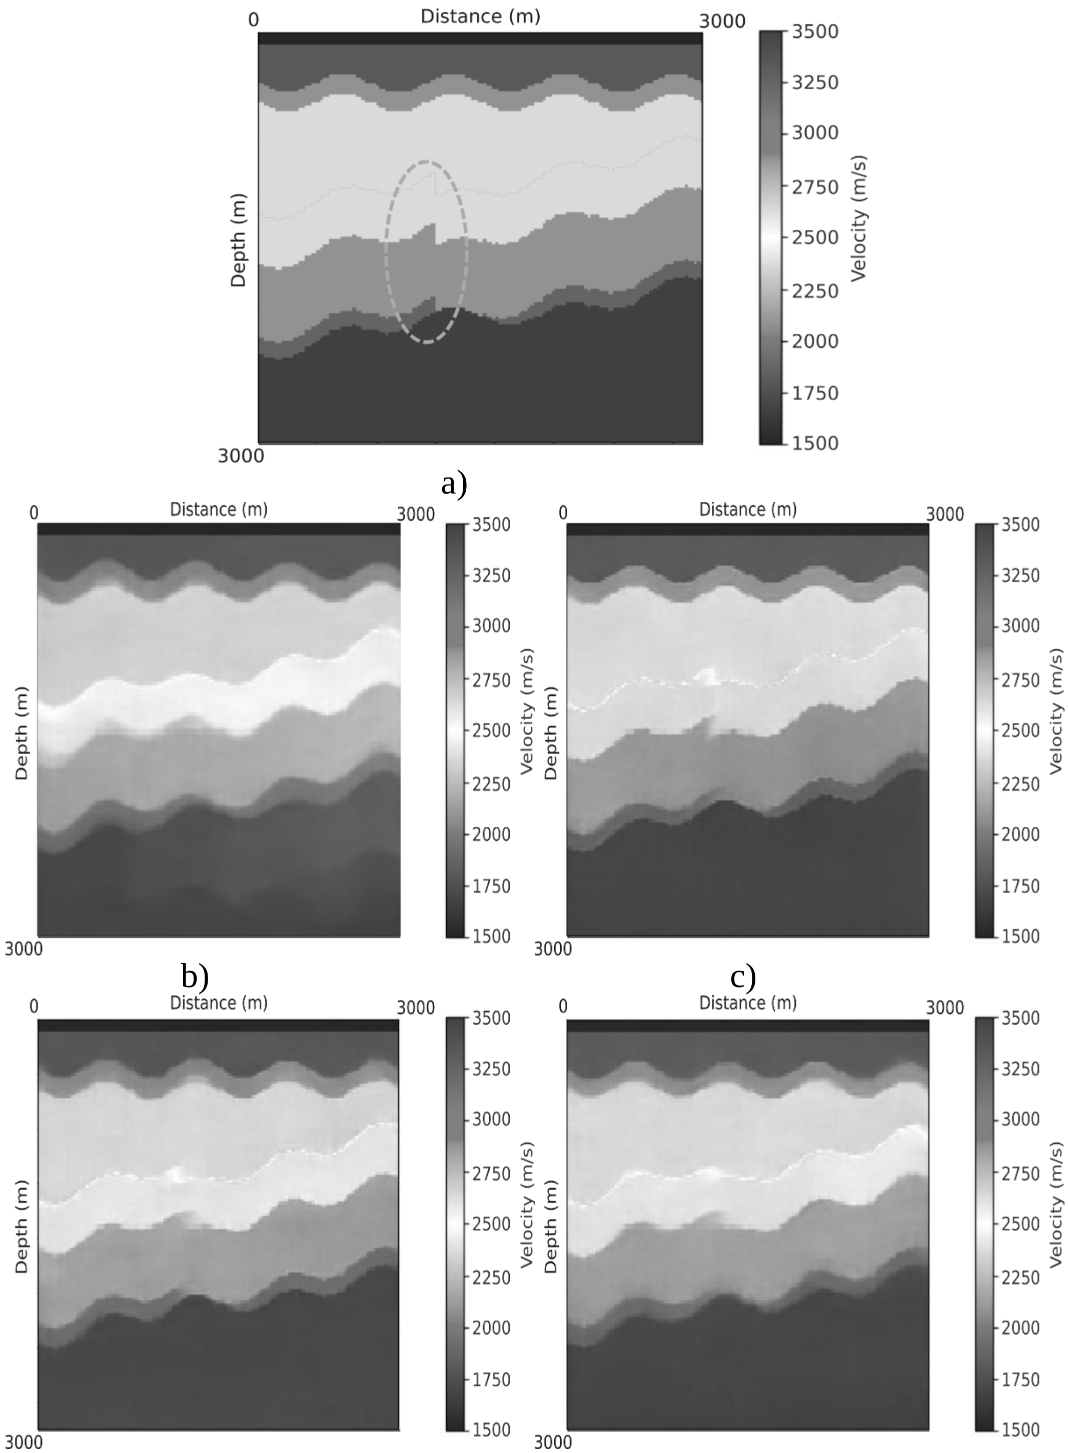

a)

b) c)

d) e)

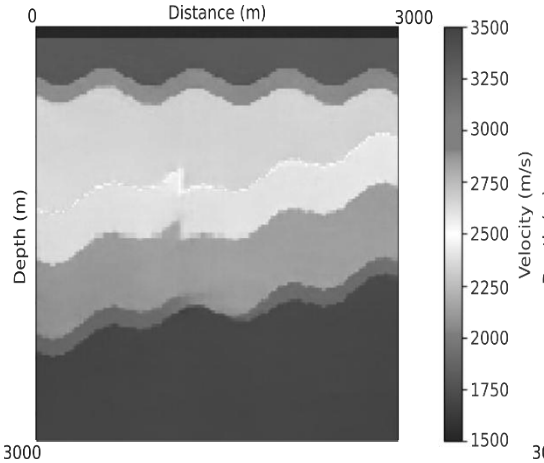

f)

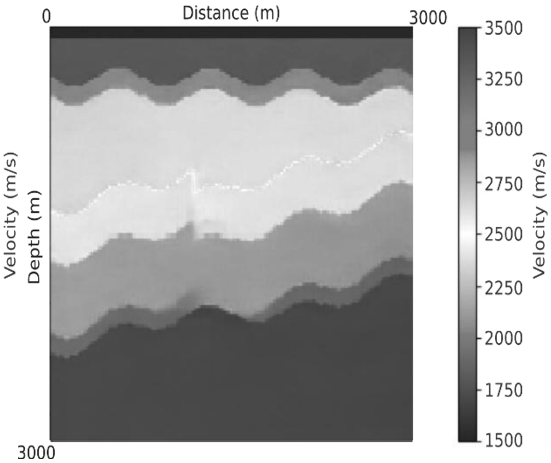

g)

Figure 10: a) Ground-truth velocity model and graphical results obtained with the experiments using b) 1 source, c) 10 sources, d) 25 sources, and e) 50 sources all modeled with a peak frequency of 4 Hz, and 25 sources with a peak frequency of f) 8 Hz and g) 16 Hz.

Table 1. Evaluation statistical metrics and the time for training (in hours) the U-Net in each of the experiments

	Peak Frequency (Hz)	**Time (h)**	**MSE**	**MAE**	**R^2**	**r**	**fac2**
1 source	4	7.19	14172	75.39	0.966	0.983	0.999
10 sources	4	7.43	7313	45.41	0.978	0.989	1.0
25 sources	4	8.10	6837	44.19	0.980	0.990	1.0
25 sources	8	8.14	6126	46.79	0.982	0.991	0.999
25 sources	16	8.09	7578	54.69	0.980	0.990	1.0
50 sources	4	9.07	7207	49.72	0.980	0.990	0.999

Apart from analyzing the computational time, it can also be seen from Table 1 that the experiments have broadly similar statistical metrics. Although they are close to one another, these metrics contain much information on how each model performed. For example, having only one source yields a high MSE, indicating that the U-Net did not progress much during training and hence detected fewer useful patterns in the input data that could be translated into the velocity model compared to the models trained with the other configurations.

Following this line of thought, it is reasonable to assume that having more sources will lead to better results. We can indeed see that the MSE decreased by almost half when 10 sources were modeled instead of 1. The other statistical metrics also improved significantly. This pattern was repeated when the sources were increased from 10 to 25 at a constant peak frequency of 4 Hz. However, all the metrics deteriorated when the number of sources was increased from 25 to 50. We believe that this deterioration is due to the large size of the velocity model. Because the subsurface and consequently, the velocity model, are relatively small, increasing the number of sources excessively might increase the input data complexity and introduce data artifacts that distort the U-Net learning process and lead to degraded performance.

An interesting result was obtained when the peak frequency of the seismic modeling was changed. After increasing the peak frequency from 4 to 8 Hz, all the statistical metrics except the MAE improved even further and we achieved the best results in this work. However, when the peak frequency was doubled again, the metrics became worse than the experiment with 50 sources. Since the peak frequency changes the level of details in the seismic data and consequently, the numerical values in the data, we believe that changing the values and increasing the details excessively could cause the U-Net to misidentify certain patterns in the seismic data and associate them erroneously with the velocity model.

The results obtained with these experiments are extremely interesting. They show that a deep learning model can be applied to even a simple seismogram with only a single source to estimate a velocity model with a fairly good resolution. The results also show that different configurations of the seismic modeling produce different outputs in the deep learning model and that increasing the number of sources or the peak frequency excessively does not necessarily lead to better estimated velocity models.

6. Conclusions

We demonstrated in this work how a deep learning model can be used to estimate the velocity models of subsurfaces from seismic data. The seismic data were generated from seismic modeling with different configurations from a common set of synthetic ground-truth velocity models. We could thus evaluate how the different configurations influenced the performance of our U-Net deep learning model and consequently, the resolution of the estimated velocity models.

We conclude that the U-Net can estimate velocity models from seismograms acquired from even just a few sources because the size and complexity of the subsurface is not high. However, we believe that as the size grows and more complex structures such as salt bodies are considered, more sources would generate more patterns for the deep learning model to detect and thus generate better estimates. Moreover, because the U-Net was shown to be insensitive to high peak frequencies, i.e., large amounts of detail in the seismic data, modeling with middle-range frequencies produces the best estimations.

The results are promising as the U-Net can estimate velocity models with considerably high resolution from seismograms produced by seismic modeling with different sets of parameters. This indicates that the method can give an initial estimation even if the seismogram is produced from a limited configuration. Future research will focus on experiments dimensionality reduction applied on the seismic data and different deep learning architectures.

Acknowledgments

The study is fully supported by the SENAI Cimatec University Center, in Bahia, Brazil. We want to thank the SENAI Cimatec, its Reference Center on Artificial Intelligence and its Supercomputer Center for Industrial Innovation (CS2i).

References

Al-Yahya, K. (1989). Velocity analysis by iterative profile migration. *Geophysics*, 718-729.

Araya-Polo, M., Jennings, J., Adler, A., & Dahlke, T. (2018). Deep-learning tomography. *The Leading Edge*, 58-66.

Baysal, E., Kosloff, D. D., & Sherwood, J. W. (1983). Reverse time migration. *Geophysics*, 1514-1524.

Campos, L. R., Nogueira, P., & Nascimento, E. (2019). Estimating Initial Velocity Models for the FWI Using Deep Learning. *Proceedings of the 16th International Congress of the Brazilian Geophysical Society.* Brazilian Geophysical Society.

Campos, L. R., Nogueira, P., & Nascimento, E. (2019). Tuning a Fully Convolutional Network for Velocity Model Estimation. *Offshore Technology Conference Brasil.* Offshore Technology Conference.

Campos, L. R., Nogueira, P., Moreira, D., & Nascimento, E. (2019). An Empirical Analysis of the Influence of Seismic Data Modeling for Estimating Velocity Models with Fully Convolutional Networks. *Journal of Systemics, Cybernetics and Informatics*, 26-32.

Chollet, F. (2017). *Deep Learning with Python.* Manning Publications CO.

Committee, I. M. (2017). Source Separation for Simultaneous Seismic Data Acquisition. *IEEE Dataport.* http://dx.doi.org/10.21227/H2TP46.

Datawider. (s.d.). *How Deep Learning is Different from Machine Learning*. Fonte: Datawider Website: https://datawider.com/how-deep-learning-is-different-from-machine-learning/

Datta, D., & Sen, M. K. (2016). Estimating a starting model for full-waveform inversion using a global optimization method. *Geophysics*, R211-R223.

Dos Santos, A. W. (2013). *Waveform inversion applied to the analysis of seismic velocities using a multi-scale approach.* Salvador: Universidade Federal da Bahia.

Krizhevsky, A., Sutskever, I., & Hinto, G. E. (2012). Imagenet classification with deep convolutional neural network. *Advances in neural information processing systems*, (pp. 1097-1105).

LeCun, Y., & Bengio, Y. (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*.

LeCun, Y., Boser, B. E., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. E., & Jackel, L. D. (1990). Handwritten digit recognition with a back-propagation network. *Advances in neural information processing systems*, (pp. 396-404).

Lewis, W., & Vigh, D. (2017). Deep learning prior models from seismic images for full-waveform inversion. *SEG Technical Program Expanded Abstracts 2017* (pp. 1512-1517). Society of Exploration Geophysicists.

Long, J., Shelhamer, E., & Darrel, T. (2015). Fully convolutional networks for semantic segmentation. *Proceedings of the IEEE conference on computer vision and pattern recognition*, (pp. 3431-3440).

MC.AI. (26 de February de 2020). *Understanding supervised, unsupervised, and reinforcement learning*. Fonte: MC.AI Website: https://mc.ai/understanding-supervised-unsupervised-and-reinforcement-learning/

Merriam-Webster. (27 de October de 2020). *Dictionary: Artificial Intelligence*. Fonte: Merriam Webster Dictionary: https://www.merriam-webster.com/dictionary/artificial%20intelligence

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical image computing and computer-assisted intervention* (pp. 234-241). Springer.

Sajeva, A. (2016). Estimation of acoustic macro models using a genetic full-waveform inversion: Applications to the Marmousi model Genetic FWI for acoustic macro models. *Geophysics*, R173-R184.

Stork, C. (1992). Reflection tomography in the postmigrated domain. *Geophysics*, 680-692.

Virieux, J., & Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. *Geophysics*, WCC1-WCC26.

Wang, W., Yang, F., & Ma, J. (2018). Velocity model building with a modified fully convolutional network. *SEG Technical Program Expanded Abstracts 2018* (pp. 2086-2090). Society of Exploration Geophysicists.

Wu, Y., Lin, Y., & Zhou, Z. (2018). InversionNet: Accurate and efficient seismic waveform inversion with convolutional neural networks. *SEG Technical Program Expanded Abstracts 2018* (pp. 2096-2100). Society of Exploration Geophysicists.