Efficient extraction of seismic reflection with Deep Learning

G. Roncoroni *, E. Forte, L. Bortolussi, M. Pipan

University of Trieste, Department of Mathematics and Geosciences, Italy

ARTICLE INFO

Dataset link: https://github.com/Giacomo-Roncoroni/Efficient_horizons_extraction

Keywords:
Horizon extraction
Deep Learning
Neural Network
Reflection seismic
GPR

ABSTRACT

We propose a procedure for the interpretation of horizons in seismic reflection data based on a Neural Network (NN) approach, which can be at the same time fast, accurate and able to reduce the intrinsic subjectivity of manual or control-points based methods. The training is based on a Long Short Term Memory architecture and is performed on synthetic data obtained from a convolutional model-based scheme, while the extraction step can be applied to any type of field seismic dataset. Synthetic data are contaminated with different types of noise to improve the performance of the NN in a large variety of field conditions. We tested the proposed procedure on 2-D and 3-D synthetic and field seismic datasets. We have successfully applied the procedure also to Ground Penetrating Radar data, verifying its versatility and potential. The proposed algorithm is based on a fully 1-D approach and does not require the input of any interpreter, because the necessary thresholds are automatically estimated. An added benefit is that the prediction has an associated probability, which automatically quantifies the reliability of the results.

1. Introduction

With the increase of 3-D seismic surveys, the quantity of geophysical data used in the interpretation has grown at a phenomenal rate (Dorn, 1998). Therefore, new picking strategies and algorithms have been implemented in order to make data interpretation faster, less subjective and more accurate at the same time.

There is a plethora of different methods used to interpret seismic horizons and their classification is almost impossible since they are intrinsically different in terms of adopted strategy, type of data application (e.g. 2-D or 3-D), expected results (line drawing, horizon extraction, simple geometrical definition of the structures). Typical approaches include manual picking, interpolation, auto-picking, voxel tracking, or surface slicing (Dorn, 1998), often starting on the manual picking of a few control points (a.k.a. seeds) then connected by means of different interpolation algorithms across selected portions of the analyzed data set.

Such an approach is not so trivial especially for 3-D seismic data due to several issues and critical steps including, among the others: erroneous and not coherent control points selection; autopicking of data having low reflection coherence and/or limited signal-to-noise ratio (S/N); intrinsic directivity of data and/or of the selected algorithm; complex geology preventing effective autotracking (Herron, 2014). Multiples, lateral amplitude variations, lateral phase discontinuities related or not to actual geological features, and overall low S/N ratios are further challenges (Hoyes and Cheret, 2011). Most of the existing picking strategies exploit reflection amplitudes and are as a consequence very sensitive to S/N but also to lateral amplitude variations, not always related to acoustic impedance changes. The applied processing flow represents another potential issue because it can introduce both distortions in the horizon's continuity and not real coherencies, which lead to erroneous results.

Recently, different strategies have been suggested in order to extract the horizons directly from the seismic dataset without manual picking (Cubizolle et al., 2015). These new algorithms track surfaces throughout the entire data volume (or within selected portions of it) and are global in nature because they do not need any previous definition of seeds (Stark, 2004; Lomask and Guitton, 2007). In fact, they are able to simultaneously define multiple horizons exploiting the full dimensionality of the data thus offering, at least in favorable conditions faster and less subjective results (Hoyes and Cheret, 2011). Different approaches have been tested: dip-driven algorithms exploit continuity/variations of dips, global optimization methods focus on the entire available dataset, while horizon patching solutions divide the problem into two different steps i.e. the picking and the horizon grouping or definition. They are based on different assumptions and strategies but are generally time consuming due to their computational cost. A comprehensive review with their pros and cons is provided, for instance, in Hoyes and Cheret (2011).

We use a Neural Network (NN) approach to interpret seismic horizons and to reduce the subjectivity of the control points-based procedures, while assuring accurate and robust horizon extraction. This

* Corresponding author.
E-mail address: groncoroni@units.it (G. Roncoroni).

Accepted 3 July 2022

can be done by a preliminary training on a large synthetic dataset (100,000 traces with 256 timesteps, i.e. samples each). The Long Short Term Memory (LSTM) architecture fits well the problem because it can handle causality and can be applied to input data of different dimensions. The preliminary subdivision of the input into patches of predefined dimension, commonly required by Convolutive Neural Network, is not necessary in the case of LSTM.

Training the NN with synthetic seismograms offers additional benefits, such as the ability to introduce arbitrarily predefined characteristics that are difficult to find in a single seismic dataset. They can also reproduce different levels of resolution and can be generated at a low computational cost. Last but not least, they overcome the implicit limit of an unknown underground model, as in the case of field datasets, thus providing complete control over the position of the reflectors. We present and critically discuss all the phases of the new proposed flow, namely: the data generation, the training, and the application to both synthetic and field data.

2. Methods

2.1. Generation of the data for the NN training

We train the NN on synthetic data to avoid any link to the field dataset to be interpreted and to have a complete control over the NN performance through the knowledge of the subsurface model that generated the training data. We start from the classical convolutional model of the seismic trace (Yilmaz, 2001) to generate the data:

$$f(t) = w(t) * r(t) + n(t) \tag{1}$$

where *f(t)* is the seismic trace, *w(t)* is the wavelet produced by the seismic source, *r(t)* is the reflectivity function (i.e. the series of reflection coefficients) and *n(t)* is the noise.

We modified the terms in Eq. (1) to generate a variety of synthetic data that may effectively represent that a wide range of conditions commonly encountered in field datasets, in terms of signal and signal-to-noise characteristics.

Number of samples and sampling interval are 256 and 0.002s respectively: we have chosen 256 as the dimension of the data vector representing each seismic trace to reduce the training time and the computational load, without limiting the effectiveness of the training phase.

Reflectivity (*r(t)*) takes as random values in the range $\pm[0.04, 1]$ to avoid reflection coefficient labeled with unrealistic signal-to-noise ratio. Each trace can have a maximum of 7 (seven) reflection coefficients in different and random positions in time ranging from sample 10 to 246.

w(t) is a Ricker wavelet (Wang, 2015), defined in the time domain as:

$$r(\tau) = \left(1 - \frac{1}{2}\omega_p^2\tau^2\right)\exp\left(-\frac{1}{4}\omega_p^2\tau^2\right) \tag{2}$$

where τ is time (in seconds) and ω_p is the most energetic (or dominant) frequency in radiant per second.

The dominant frequency is chosen between [30 Hz and 70 Hz] with a random uniform distribution, in order to cover the common seismic source frequencies.

To better approximate a real seismic acquisition we need to add noise *n(t)*, as seen in Eq. (1).

We added the noise component n(t) in (1) to better reproduce realistic field datasets and introduced different types of noise with different random magnitudes, as shown Fig. 1. Testing of the NN on field data was performed after introduction of the following types of noise, see Fig. 3:

- Noise 1: defined as a pure random noise that is added to the convolved trace. It is *n(t)* defined in Eq. (3)

$$f(t) = w(t) * r(t) + n_1(t) \tag{3}$$

- Noise 2: it represents an "anisotropic noise" and it is a noise added before the convolution with the reflection coefficient series.

$$f(t) = w(t) * (n_2(t) + r(t)) \tag{4}$$

- Noise 3: use of both Noise 1 and Noise 2 in order to simultaneously reproduce the effects that could affect a seismic trace, namely random noise and anisotropic noise.

$$f(t) = w(t) * (n_2(t) + r(t)) + n_1(t) \tag{5}$$

The reference output is defined by a binary indicator (1,0) used to label each sample as reflection/no reflection respectively.

We used the following algorithm to generate the input seismic trace and the reference output for the NN training:

for *i in range(number of desired data:)*
n = generate a casual number of horizons
generate n ***reflectivity values***
generate a random ***wavelet***
generate n random ***depth values***
sort depth values
build the ***reflection coefficient series*** *with depth and reflectivity values*
build ***seismic trace*** *by convolving the reflection coefficient series and the wavelet*
Build the ***categorical reference output*** *with the depth values*

This algorithm is very fast and can be iterated to generate large numbers of seismic traces and reference outputs.

2.2. Neural network geometry

We chose the network geometry in Fig. 2 after a grid search to estimate NN hyperparameters, e.g. number of neurons and depth of the NN.

LSTM maintains the causality and the Long-Term memory better fits the physics behind the wave propagation (Hughes et al., 2019).

The Bi-Directional LSTM is able to improve the accuracy of NN classification (Guo et al., 2019): in the present study, it can in particular reduce false positives in the final layer by improving the identification of the correct shape of the wavelet.

The output is driven by a Dense layer with a Softmax activation function that outputs a probability value equal to 1 on the phase with the maximum amplitude of a reflection.

We used the CuDNNLSTM (Hochreiter and Schmidhuber, 1997) implemented in Keras (Chollet et al., 2015), a fast approximation of LSTM on Nvidia CUDA (Chetlur et al., 2014): this means that both the training and the model must be used on a CUDA-compatible GPU.

As optimizer we used AdaMax (Kingma and Ba, 2014), a modified version of Adam with infinity norm, because it performs better on this problem compared to Adam and SGD (Kingma and Ba, 2014).

As loss function we used Categorical Crossentropy (Mannor et al., 2005), defined as:

$$loss = -\sum_{i=0}^{C} t_i log(s_i) \tag{6}$$

Where t_i and s_i are the expected classes and the NN score for each class i in C, respectively summed on the time steps: in this work $C = 2$.

2.3. Training

The training phase was split in two steps (Kavzoglu, 2009): an initial training on a noiseless dataset and a subsequent training on a noisy dataset. This choice was due to the unbalanced output solutions: a direct training on the noisy trace would led to a huge local minimum where the NN outputs only 0s. To avoid this, we initially train the NN on a noiseless dataset. We got good results from such training, as shown in the additional materials: the training phase lasts 2 h and it reaches an accuracy of 0.9995 in test and 0.980 in validation.

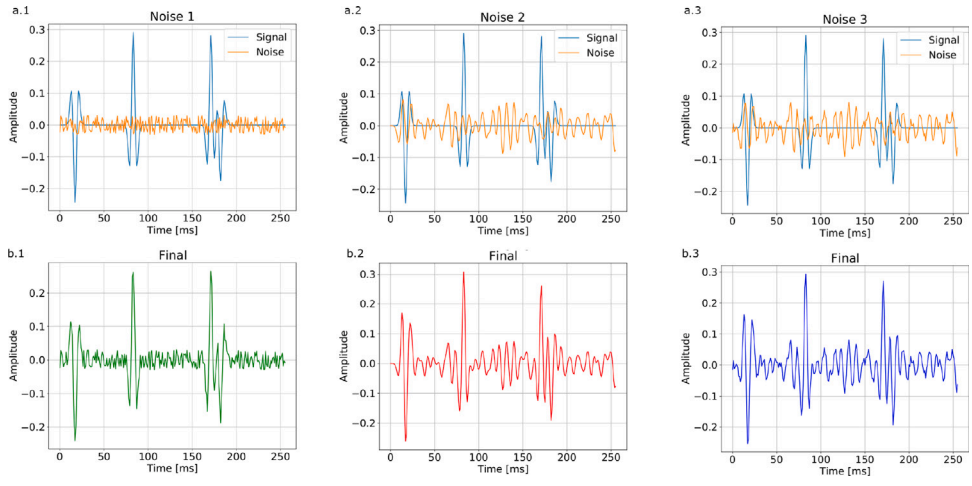

Fig. 1. Three different types of noise (orange) superimposed on a synthetic seismic trace with four reflections (light blue) [a.1, a.2, a.3]; seismic traces resulting from the sum of synthetic trace and noise [b.1, b.2, b.3] (see text for details). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

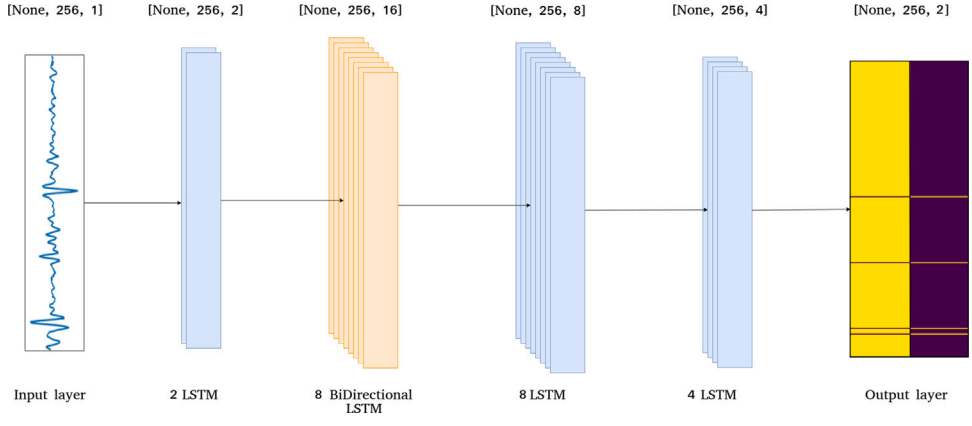

Fig. 2. NN geometry: we set a layer with 2 LSTM (blue, second block from left), 8 BiDirectional LSTM (orange, third block), 2 layers with 8 and 4 neurons each (blue, fourth and fifth blocks) and a final Dense layer with 2 neurons. On the left the input seismic trace and on the right the corresponding reference output. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

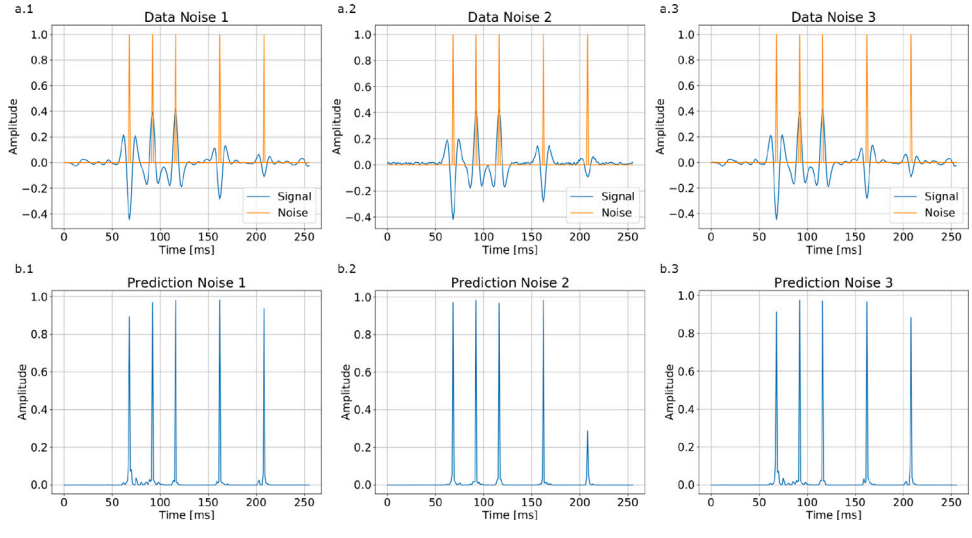

Fig. 3. Noisy traces (blue; see Fig. 1 for the different types of noise) and expected output (orange); (b.1-2-3): results of the NN prediction. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We trained the NN with Adamax optimizer with a learning rate of 0.01, categorical crossentropy as loss function and batch size of 512. (See Figs. 3 and 7.)

After the training on the noiseless dataset, we retrained this model 3 times on 3 different datasets with the noise levels and the noise types described in the previous section. The main objective of this training phase is to evaluate the performances of the trained NN on synthetic and field datasets to infer the optimum solution to generate synthetic data for the application of NN on field seismic data.

2.4. Prediction

A 1-D approach can be adversely affected by a high level of random noise, although the prediction method is quite robust.

We use the ensemble learning technique to reduce the prediction uncertainty: this technique uses multiple learning algorithms to obtain better predictive inferences (Mendes-Moreira et al., 2009).

In detail, we tested two solutions, namely prediction with different NNs trained on a dataset with the same characteristics and prediction with the same NN on a single trace and on its inverted version in time.

The two approaches produced similar results and we thus decided to use a single NN to reduce the required training effort.

This methodology generates 2 different predictions, and we combined them with the geometric mean, as it gives better results than the arithmetic mean.

This can be explained by the nature of the prediction: we predict values of probability [0-1], so if NN1 predicts 0.5 (a mid value due to a noisy area) and NN2 0.001 (a very low value), the arithmetic mean would produce a value of $\simeq 0.25$, while the geometric approach would penalize more this prediction ($\simeq 0.02$). An example is shown in Fig. 4. We can see how two wrong peaks at 25 ms and at 780 ms, in the first prediction, are muted by the geometric mean. We further normalized the maximum value of the first reflection to one to avoid exceedingly large values that can affect the NN performance.

3. Results

3.1. Test on 2-D synthetic data

The first test of the NN performances on 2-D datasets generated with the same workflow as the training dataset produced almost perfect results, as expected. So we moved onto a more complicated example using the Marmousi model (Martin, 2004). Fig. 5 shows part of the data and the prediction, which is given as a probability set that associates a probability value to each point in the section: the value indicates the probability of the point to be a reflector, i.e. to belong to a reflecting surface. Therefore a threshold above which the point is labeled as a reflector has to be set. We directly estimate the optimum threshold by evaluating the number of points classified as reflectors vs the threshold (Fig. 6). We perform this task by using the algorithm described in Satopaa et al. (2011). The threshold is set at the sharp inflection point clearly visible in the resulting curve, thus limiting the subjectivity of the choice (Fig. 6).

3.2. Test on 2-D field data

In order to test the proposed methodology we use at first a 2-D marine seismic profile of the WS10 exploration project, obtained in autumn 2010 in the west Mediterranean Sea by the Istituto Nazionale di Oceanografia e Geofisica Sperimentale (OGS), which also performed the data processing (Geletti et al., 2014). The selected portion of the seismic profile images a rifted margin of the eastern Sardo-Provençal Basin characterized by a faulted salt dome and by a portion of an almost undisturbed sedimentary sequence (Fig. 8). For such reason, the analyzed data represent an interesting and complex test for the proposed procedure. We focused on this portion also because it is exactly the same used in Forte et al. (2016) (Fig. 5a therein) to test an automated picking and phase assessment approach based on phase seismic attributes. The NN is able to properly extract all the main horizons, both where they are sub-horizontal (i.e. in the shallow part) and where they exhibit a significant dip (i.e. along the flanks of the salt dome). As desired, horizons interrupt at the fault location (f labels in Fig. 8),while correctly no horizons are detected within the salt dome (sd labels in Fig. 8). The latter result is definitely not easily achievable with traditional picking methods based on 1-D, 2-D or 3-D approaches, because some lateral coherent events can be detected even if they are not actually related to real reflectors (see results in Fig. 5a in Forte et al. (2016)).

NN is also able to extract the high amplitude reflector below the salt dome, while no significant features are detected in the deepest part of the section below the salt dome where reflectors are not continuous also because of some coherent noise due to over-migration effects. We point out that the above described prediction procedure can be applied to automatically extract the reflectors from any kind of dataset after the training phase is concluded on a set of randomly generated data containing various types and levels of noise.

In order to verify this capability and evaluate the performances, we tested the NN on totally different data, i.e. a Ground Penetrating Radar (GPR) dataset collected in a Glacier in the Eastern Alps (for further details about this data please refer to Colucci et al. (2014)). Although GPR and reflection seismic techniques differ radically concerning sources characteristics and physical parameters, they are both based on the same physical wave theory (Ursin, 1983). Therefore, processing and analysis techniques used for seismic data may be adapted, at least from the theoretical point of view, also to GPR datasets.

The test data have 802 samples per trace with a 0.454 ns sampling interval, while the spatial sampling interval (0.15 m) is not relevant in a 1-D procedure.

In the example provided in Fig. 8 the NN is able to extract all the relevant reflectors, which are related, from top to bottom to completely different glaciological and geomorphological units, namely: snow and firn (sf), firn (f), debris (d) ice (i) and bedrock (b). The procedure performs quite well on continuous reflectors (like for instance sf-f contact) but also on highly scattered levels (like d). The extraction of the bedrock is much more difficult mainly because it is not continuous, it is characterized by diffuse scattering and it has a lower signal-to-noise ratio than the shallower units.

3.3. Test on 3-D synthetic data

We also tested the approach on 3-D synthetic datasets to deepen the analysis of the proposed method.

The datasets used are taken from Wu et al. (2019, 2020) and were originally generated to train a Convolutive Neural Network for fault recognition (Wu et al., 2019), and for paleokarst detection (Wu et al., 2020). We chose these datasets to test our methodology because they are a good approximation to field data.

In order to test the flexibility of our methodology we directly applied our method to the datasets without preliminary processing or additional training to check its flexibility and robustness. The dataset taken from Wu et al. (2019) has been generated with a convolutive approach to simulate a faulted environment (Wu et al., 2019).

We tested the NN on a $128 \times 128 \times 128$ samples seismic volume obtained by random selection from this dataset, The results are quite good, despite the huge number of close reflections; as shown in Fig. 9 we can see the continuity of the horizons in the 3 dimensions and also easily identify the main faults. The second dataset was also generated with a convolutive approach to simulate a paleokarst environment (Wu et al., 2020). From this dataset, we extracted a $256 \times 256 \times 256$ seismic volume. The prediction is fast and quite accurate (see Fig. 10) and we can identify the collapsing structures as high noise area, the 3-D continuity of the horizons and their interruptions in the karstified areas.

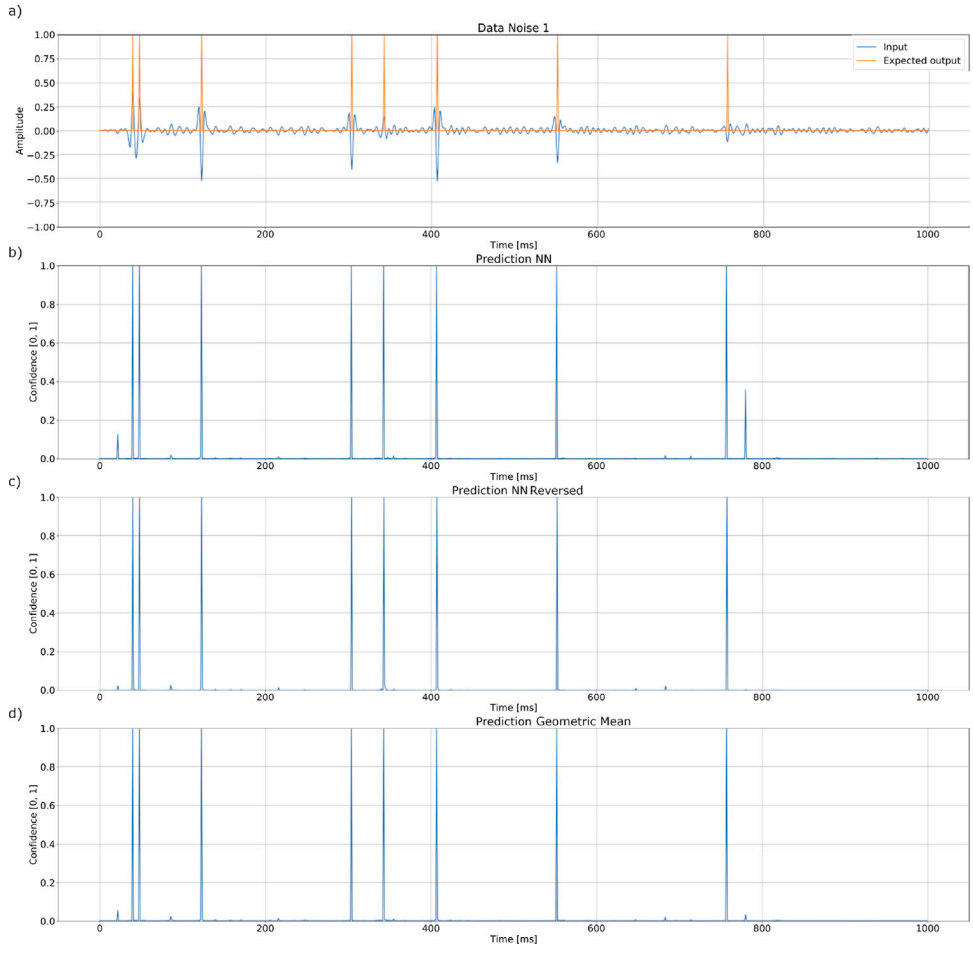

Fig. 4. (a): Input data and expected output, (b): results of the NN prediction, (c): results of the prediction performed on the input trace reversed in time, (d): geometric mean of the two predictions.

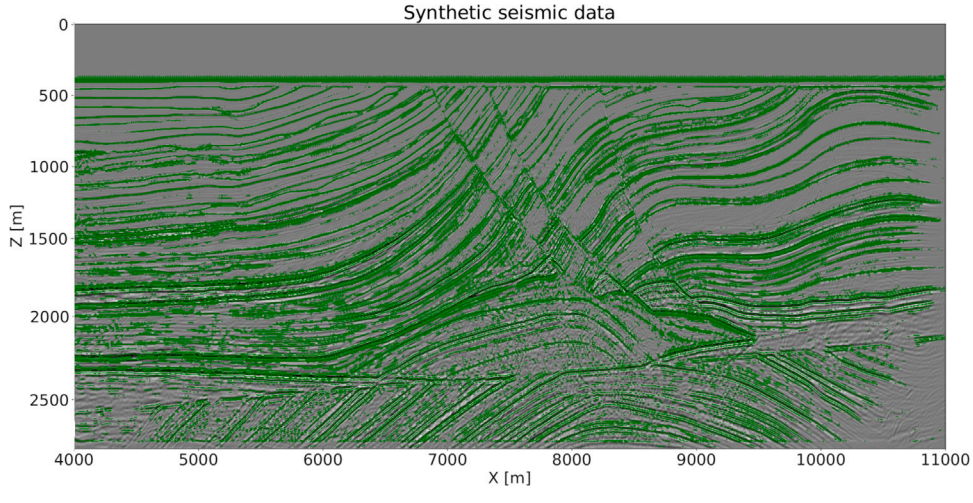

Fig. 5. Results of the horizon extraction (green) performed by the NN on the Marmousi Dataset, (green), superimposed on the input data. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4. Discussion

The proposed methodology substantially reduces the computational costs of the generation process compared to other approaches. The use of Finite Differences methods for the forward modeling typically requires huge amounts of memory and computational time for the generation of datasets like the ones here considered. We therefore choose the convolutive approach, which has the added advantage of better handling different types of superimposed noise (see Section 2.1).

Another key factor is using fully synthetic datasets for the training phase: avoiding field data at this stage prevents possible bias due to manual selection of horizons and implicit limitations in data characteristics due to the fact that field data can only include a limited range of possible geological situations. Furthermore, the more classic approach,

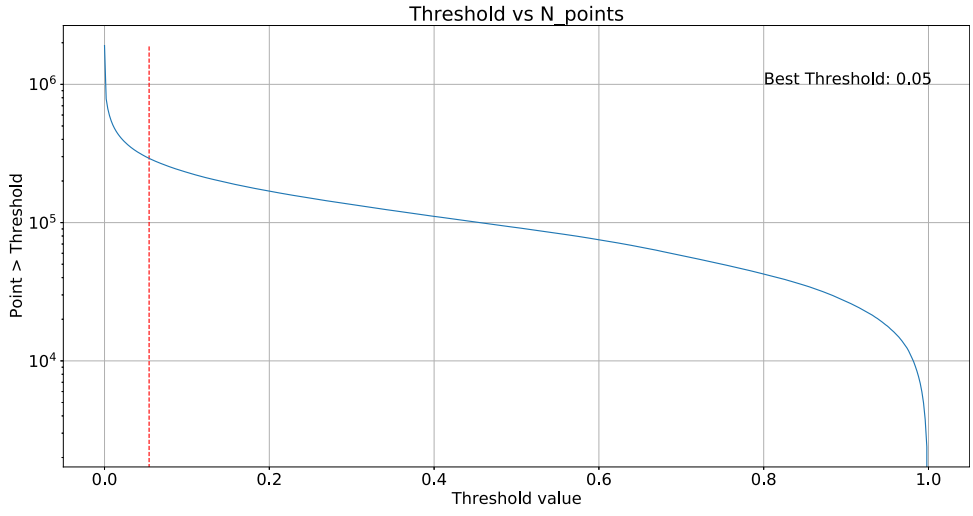

Fig. 6. Number of points classified as reflectors (vertical axis) vs. Threshold value: optimal threshold indicated in red at the knee point. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

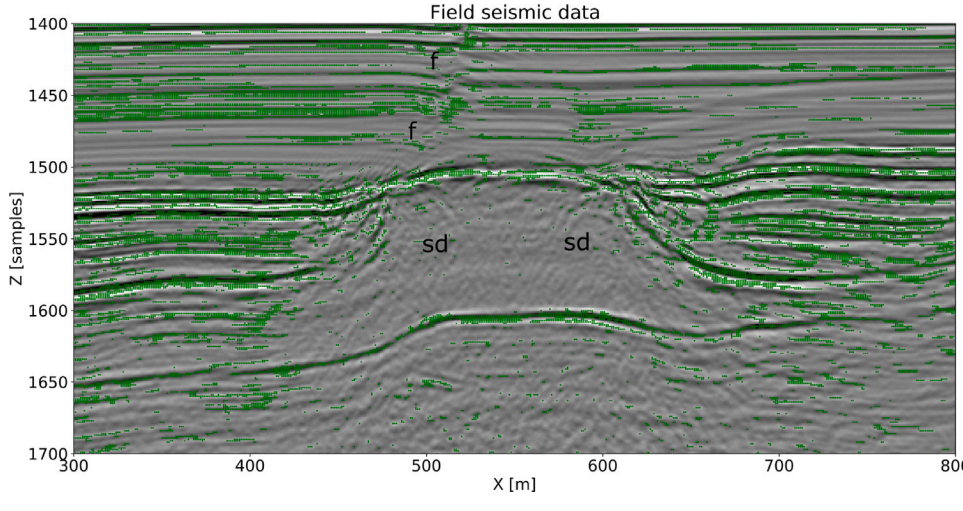

Fig. 7. Example of application of the NN to field seismic data: the data are part of a seismic line from the western Sardinian margin of the WS10 exploration project [see text and Geletti et al. (2014) for details]. The prediction results are shown in green: "sd", salt dome; "f", main fault. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

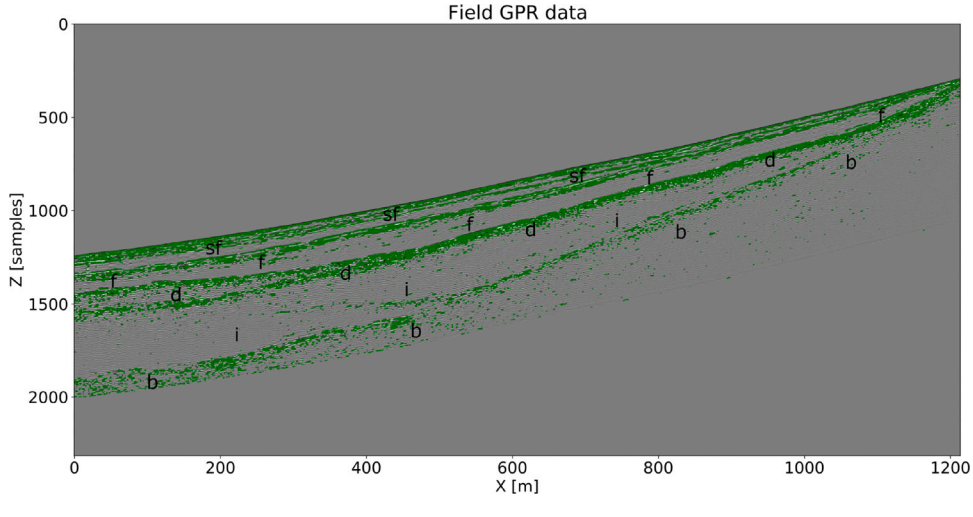

Fig. 8. Example of application of the NN to field GPR data: the GPR profile crosses a glacier in the Eastern Alps [see text and Colucci et al. (2014) for details]. The results of the prediction are shown in green: "sf", snow and firn; "f", firn; "d", debris; "i", ice; "b", bedrock. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

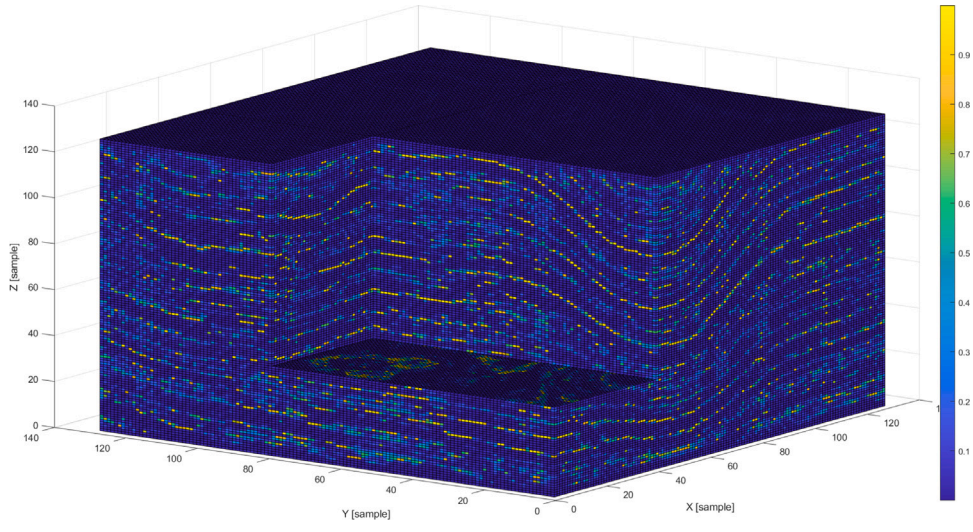

Fig. 9. Example of results of NN prediction on a 3-D synthetic dataset generated from a subsurface model corresponding to a faulted environment (Wu et al., 2019): the color scale is associated to the probability level (1 = reflection, see text for details). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

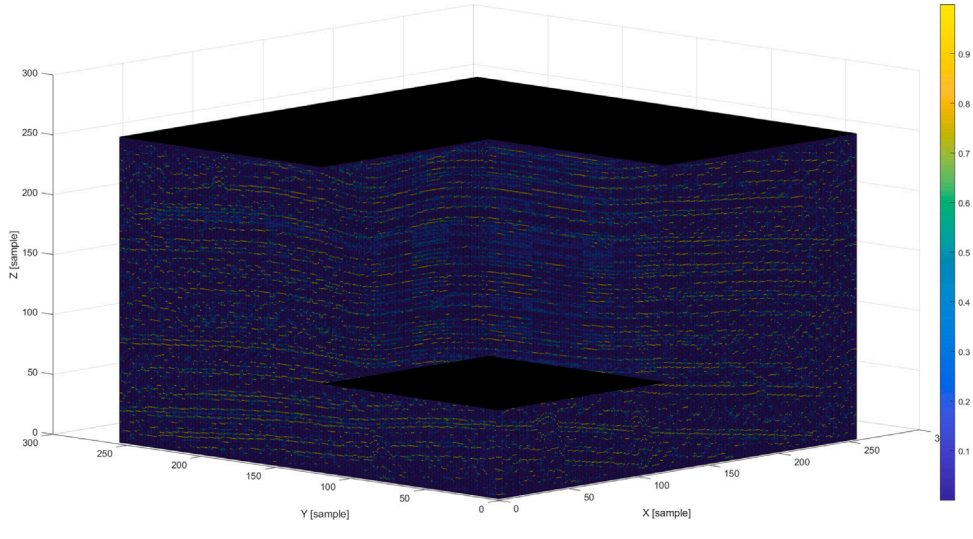

Fig. 10. Example of results of NN prediction on a 3-D synthetic dataset generated from a subsurface model corresponding to a paleokarst environment (Wu et al., 2020): the color scale is associated to the probability level (1 = reflection, see text for details). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in which a portion of the available data is used to train the NN and subsequently applied to the entire data set, has the obvious inherent constraint that the previous data is excluded from the extraction procedure. An example of such an approach can be found in Tschannen et al. (2020), where the authors trained a Convolutive Neural Network on a small portion of the dataset after a manual interpretation, then applying the trained model on the remaining data.

For validation purposes, we tested the proposed procedure on the synthetic Marmousi dataset from Martin (2004). For this dataset all the model details are known, including reflector's geometry, velocity model and seismic wavelet. We can evaluate our results considering both the seismic section (Fig. 11a) and the real velocity model (Fig. 11b).

The shallowest part of the profile, from 0 to about 1500 m, is well predicted and matches very well both the seismic section and the velocity model; also the faults are correctly identified.

The analysis of the results of prediction in the deep part of the section gives useful indications about the performance of the network as a function also of the noise level. The complexity of Marmousi's velocity model gives rise to several non-primary events, such as multiple reflections, which appear more and more at increasing depths. This is particularly evident within the wedge located between 2250–2500 m and 4000–7000 m in vertical and horizontal coordinates, respectively.

The comparison between the seismic velocity model and the seismic section reveals that several coherent events (similar to reflectors) occur in the section although the wedge is actually homogeneous and characterized by a single velocity value. Despite the increasing number of primary and non-primary interfering events, the NN makes accurate predictions and correctly retrieves the main elements of the structural model.

A remarkable result of the proposed approach is that the NN is able to deal also with the wavelet stretching due to the well known low-pass filtering effect of the geological materials (Yilmaz, 2001). The data about the computational load and performance of the NN refers to a machine with a 2 core Intel(R) Xeon(R) CPU, 2.20 GHz, 12 Gb Ram and a Nvidia Tesla T4 GPU.

We performed 100 tests to test the generation time of the algorithm: each test produced a dataset of 10,000 traces 256 samples long. The total time for each test was 2.80 s $\pm$ 0.13 s.

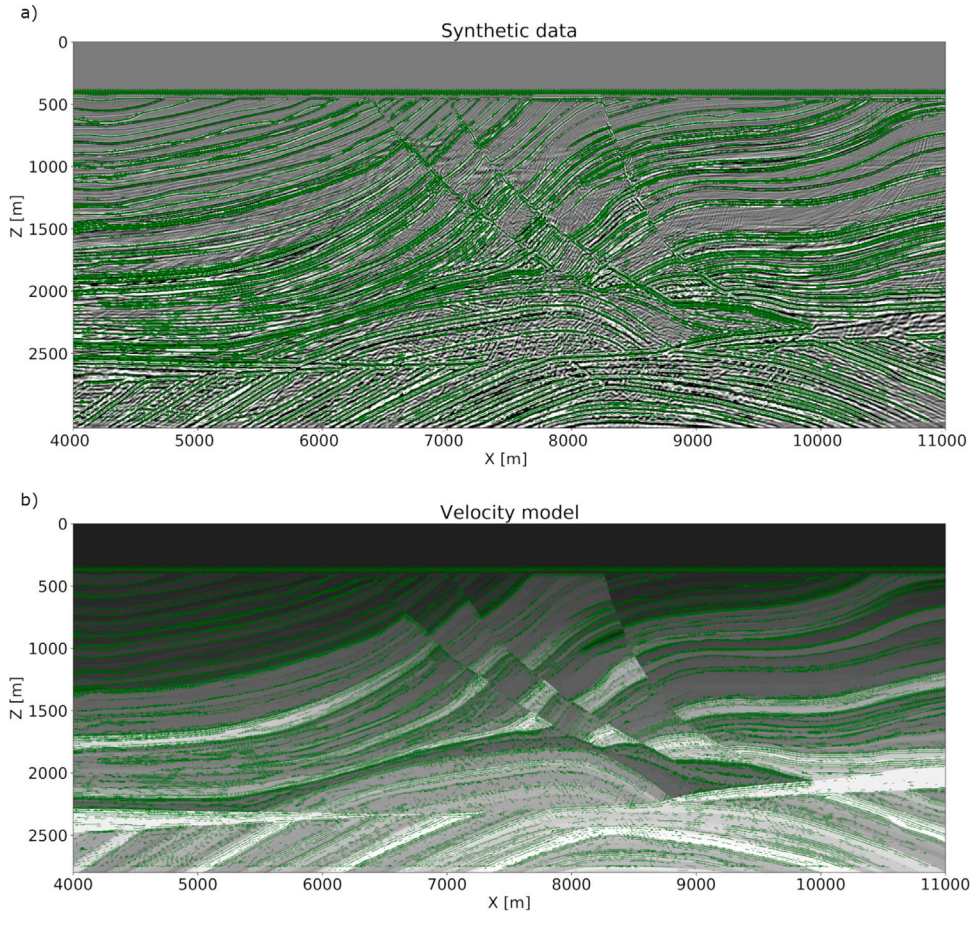

Fig. 11. Marmousi dataset: example of prediction (a) on the seismic profile and (b) on the true velocity model.

Training time is not an issue, as the NN does not need to be retrained for each specific task, as above highlighted, but can be applied directly to most field datasets.

The first training, without noise, lasts 100 epochs on a training dataset of 100,000 × 256 samples and this takes less than 1 h. The training with noise is longer, due to the increasing complexity of the task: it lasts 4000 epochs for a total training time of about 26 h.

A great strength of this method is the significantly shorter prediction time than traditional approaches. We tested it on 1000 datasets with variable time dimensions from 4 to 4096 samples.

LSTM is able to predict any temporal length and we analyzed also the change in prediction time depending on the record length, resulting, as expected, in a linear relationship between temporal length and time of prediction.

5. Conclusions

We propose a Deep Learning-based methodology that is fast and accurate in extracting different types of reflectors from reflection seismic dataset. The results obtained from its application to synthetic and field datasets show that it can be an effective tool for applications that require accurate automated recovery of reflection signals from large datasets, such as e.g. velocity analysis, tomography, interpretation.

One of the main advantages introduced is the prediction with a correlated probability: this is crucial in the recovery of horizons (i.e. reflecting surfaces) from both 2-D and 3-D data because this parameter can weigh the estimate of the optimum surface, decreasing the level of subjectivity of the whole procedure.

Other possible applications of this methodology could be the phase assessment of reflections, which in turn represents a very important task for velocity estimation. In particular, in the field of velocity modeling it may offer soft constrains for velocity model geometry in all the inversion methodologies that require an accurate input model, like for instance Full Waveform inversion.

The proposed algorithm is fast, versatile, and does not require any input from the user, as the threshold is estimated automatically using a Knee function (see Section 3.1 for details) while the training phase is done on synthetic data thus avoiding possible bias due to manual horizons picking and intrinsic limitations and peculiarities of a specific field dataset. Further researches will focus on horizons patching and automated phase recognition.

CRediT authorship contribution statement

G. Roncoroni: Conceptualization, Methodology, Software, Writing – original draft. **E. Forte:** Conceptualization, Supervision, Software, Writing – review & editing. **L. Bortolussi:** Conceptualization, Software, Writing – review & editing. **M. Pipan:** Conceptualization, Supervision, Software, Writing – review & editing, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Computer code availability

Codes used in this article are available on github: https://github.com/Giacomo-Roncoroni/Efficient_horizons_extraction

Acknowledgments

This research was supported by PNRA projects IPECA (PNRA18_0-0186) and CRIOVEG (PNRA18_00288) and by the project "Dipartimento di Eccellenza" of the Department of Mathematics and Geosciences of the University of Trieste. We gratefully acknowledge the support of Shearwater and Halliburton Landmark through their academic grants.

References

Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., Shelhamer, E., 2014. cuDNN: Efficient primitives for deep learning. CoRR, abs/1410.0759. arXiv:1410.0759.

Chollet, F., et al., 2015. Keras. https://keras.io.

Colucci, R.R., Forte, E., Boccali, C., Dossi, M., Lanza, L., Pipan, M., Guglielmin, M., 2014. Evaluation of internal structure, volume and mass of glacial bodies by integrated LiDAR and ground penetrating radar surveys: The case study of Canin Eastern Glacieret (Julian Alps, Italy). Surv. Geophys. 36 (2), 231–252. http://dx.doi.org/10.1007/s10712-014-9311-1.

Cubizolle, F., Valding, T., Lacaze, S., Pauget, F., 2015. Global method for seismic-well tie based on real time synthetic model. In: SEG Technical Program Expanded Abstracts 2015. pp. 1776–1781. http://dx.doi.org/10.1190/segam2015-5862834.1.

Dorn, G.A., 1998. Modern 3-D seismic interpretation. Lead. Edge 17 (9), 1262. http://dx.doi.org/10.1190/1.1438121.

Forte, E., Dossi, M., Pipan, M., Ben, A.D., 2016. Automated phase attribute-based picking applied to reflection seismics. Geophysics 81 (2), V141–V150. http://dx.doi.org/10.1190/GEO2015-0333.1.

Geletti, R., Zgur, F., Ben, A.D., Buriola, F., Fais, S., Fedi, M., Forte, E., Mocnik, A., Paoletti, V., Pipan, M., Ramella, R., Romeo, R., Romi, A., 2014. The Messinian salinity crisis: New seismic evidence in the west-Sardinian margin and eastern Sardo-Provençal basin (west Mediterranean sea). Mar. Geol. 351, 76–90.

Guo, R., Zhang, J., Liu, D., Zhang, Y., Zhang, D., 2019. Application of Bi-directional Long Short-Term Memory Recurrent Neural Network for Seismic Impedance Inversion, Vol. 2019. (1), European Association of Geoscientists & Engineers, pp. 1–5. http://dx.doi.org/10.3997/2214-4609.201901386.

Herron, D., 2014. Pitfalls in horizon autopicking. Interpretation 3, SB1–SB4. http://dx.doi.org/10.1190/INT-2014-0062.1.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Comput. 9 (8), 1735–1780. http://dx.doi.org/10.1162/neco.1997.9.8.1735.

Hoyes, J., Cheret, T., 2011. A review of "global" interpretation methods for automated 3D horizon picking. Lead. Edge 30 (1), 38–47. http://dx.doi.org/10.1190/1.3535431.

Hughes, T.W., Williamson, I.A.D., Minkov, M., Fan, S., 2019. Wave physics as an analog recurrent neural network. Sci. Adv. 5 (12), http://dx.doi.org/10.1126/sciadv.aay6946.

Kavzoglu, T., 2009. Increasing the accuracy of neural network classification using refined training data. Environ. Model. Softw. 24, 850–858.

Kingma, D., Ba, J., 2014. Adam: A method for stochastic optimization. In: International Conference on Learning Representations.

Lomask, J., Guitton, A., 2007. Volumetric flattening: an interpretation tool. Lead. Edge 26 (7), 888–897. http://dx.doi.org/10.1190/1.2756869.

Mannor, S., Peleg, D., Rubinstein, R., 2005. The cross entropy method for classification. In: Proceedings of the 22nd International Conference on Machine Learning. In: ICML '05, pp. 561–568. http://dx.doi.org/10.1145/1102351.1102422.

Martin, G.S., 2004. The Marmousi2 model, elastic synthetic data, and an analysis of imaging and AVO in a structurally complex environment. Master's thesis. University of Houston.

Mendes-Moreira, J., Jorge, A.M., Soares, C., de Sousa, J.F., 2009. Nsemble learning: A study on different variants of the dynamic selection approach. In: Machine Learning and Data Mining in Pattern Recognition. Springer Berlin Heidelberg, pp. 191–205. http://dx.doi.org/10.1007/978-3-642-03070-3_15.

Satopaa, V., Albrecht, J., Irwin, D., Raghavan, B., 2011. Finding a "Kneedle" in a Haystack: Detecting Knee points in system behaviour. In: 2011 31st International Conference on Distributed Computing Systems Workshops. pp. 166–171.

Stark, T., 2004. Relative geologic time (age) volumes—Relating every seismic sample to a geologically reasonable horizon. Geophysics 23, http://dx.doi.org/10.1190/1.1803505.

Tschannen, V., Delescluse, M., Ettrich, N., Keuper, J., 2020. Extracting horizon surfaces from 3D seismic data using deep learning. Geophysics 85 (3), N17–N26. http://dx.doi.org/10.1190/geo2019-0569.1.

Ursin, B., 1983. Review of elastic and electromagnetic wave propagation in horizontally layered media. Geophysics 48 (8), 1063–1081. http://dx.doi.org/10.1190/1.1441529.

Wang, Y., 2015. Frequencies of the Ricker wavelet. Geophysics 80, A31–A37.

Wu, X., Liang, L., Shi, Y., Fomel, S., 2019. FaultSeg3D: using synthetic datasets to train an end-to-end convolutional neural network for 3D seismic fault segmentation. Geophysics 48 (1), 1063–1081.

Wu, X., Yan, S., Qi, J., Zeng, H., 2020. Deep learning for characterizing paleokarst collapse features in 3-D seismic images. J. Geophys. Res.: Solid Earth 125.

Yilmaz, O., 2001. Seismic Data Analysis, second ed. In: Investigations in Geophysics, vol. 10, Society Of Exploration Geophysicists, http://dx.doi.org/10.1190/1.9781560801580.