Revealing the Diffracted Wavefield in Ultra High-Resolution Seismic Data with Deep Learning

Introduction

Diffractions are generated when a wave interacts with small subsurface features such as a tip, point, or edge discontinuity. To image diffractions is to image those fine details that naturally occur in the Earth. This property finds immediate application in the detection and mapping of point diffracting subsurface boulders, especially those buried tens to hundreds of meters beneath the sea floor, which pose substantial risk to the installation of wind turbines (Wenau et al., 2022). Extracting the diffractions from seismic data is notoriously difficult because diffractions are weak and can easily be overpowered by reflections or noise. Nevertheless, strong and even weak diffractions are easily visible to the human eye. This suggests that a convolutional neural network (CNN), which excels at computer vision problems, may do well at this problem.

Traditional methods for extracting diffractions fall into two broad categories. On the one hand, there are time domain, coherency-based methods that capture and subtract the reflected energy to reveal the diffracted energy plus noise (e.g. plane wave destruction, Taner et al., 2006). The quality of the extracted diffractions is dependent on the residual noise level. On the other hand, there are image domain methods that target the diffractions directly during migration (Schwarz, 2019). These methods can be computationally demanding and require a good velocity model. In either case, these traditional approaches suffer from requiring extensive parameter testing.

More recent approaches involve the use of machine learning models, specifically neural networks, to detect diffractions within seismic data (Lowney et al., 2021; Bauer et al., 2025). These models are promising as they can learn the features associated with diffractions directly, resulting in better signal to noise ratio. Once a good model has been trained these models offer fast, parameter free solutions. However, the burden of work falls back to training, of which the major component is preparing the training data. In this study, we explore the impact of data labelling approaches to train a DUCK-Net (Dumitru et al. 2023) for extraction of the diffracted wavefield.

Data

We use Ultra High Resolution (UHR) seismic data from the Doordewind Wind Farm Zone offshore Netherlands acquired in 2024. Sampled every 0.125 ms, this sparker source data has been processed through a comprehensive sequence incorporating designature, deghosting, statics, denoise, demultiple, and 4D regularisation, output to a trace spaced every 1.5625 m. Manual investigation confirms that this processing sequence has preserved the diffraction energy. This dataset poses a challenge in that the diffractions are especially weak, with diffractions only visible within a limited aperture.

Method

To prepare training data we first create a cleaned section by removing diffraction energy plus noise from the stacked data using a plane wave destruction technique. We then add synthetic diffractions to the cleaned section and train the neural network to predict the synthetic diffraction labels (Fig 1).

We use Kirchhoff time demigration to rapidly generate synthetic training data. Demigrating a dense field of randomly placed diffraction generators (spikes) creates a rich field of overlapping and interfering energy that is intended to resemble the true complexity of the diffracted wavefield. The actual density of diffractions (i.e. the number of spikes per trace) is randomized on a line-by-line basis. The strength of individual diffractions is varied by random scaling using a log-normal distribution to produce few strong and many weak diffractions. The spatial size of diffractions is varied by running separate demigration processes with different aperture parameters and summing the results.

We experiment with two approaches to selecting the modelling wavelet. One uses a wavelet derived from a local window in the data itself; thereby adding diversity to the wavelet and capturing real data

characteristics (e.g. frequency content). The other uses an idealised wavelet representative of the processed data.

For simplicity we limit this study to 2D common offset profiles. We do expect uplift if exploiting more redundancy at higher dimensions, but the models would take a lot longer to run, and the data patch sizes would be more tightly constrained by memory.

Each training run used data from 180 inlines while the validation dataset had 10 inlines. The combined volume for both the training and validation datasets was approximately 36 GBs. This struck a balance between size and training expediency allowing for fast iteration. The network used a patch size 512 x 512, 16 filters in the first layer, leaky ReLU activation functions, and the AdamW optimizer. The loss function consisted of Huber and SSIM components weighed in the ratio of 20:1. Training was done on a V100 GPU over 160 epochs.

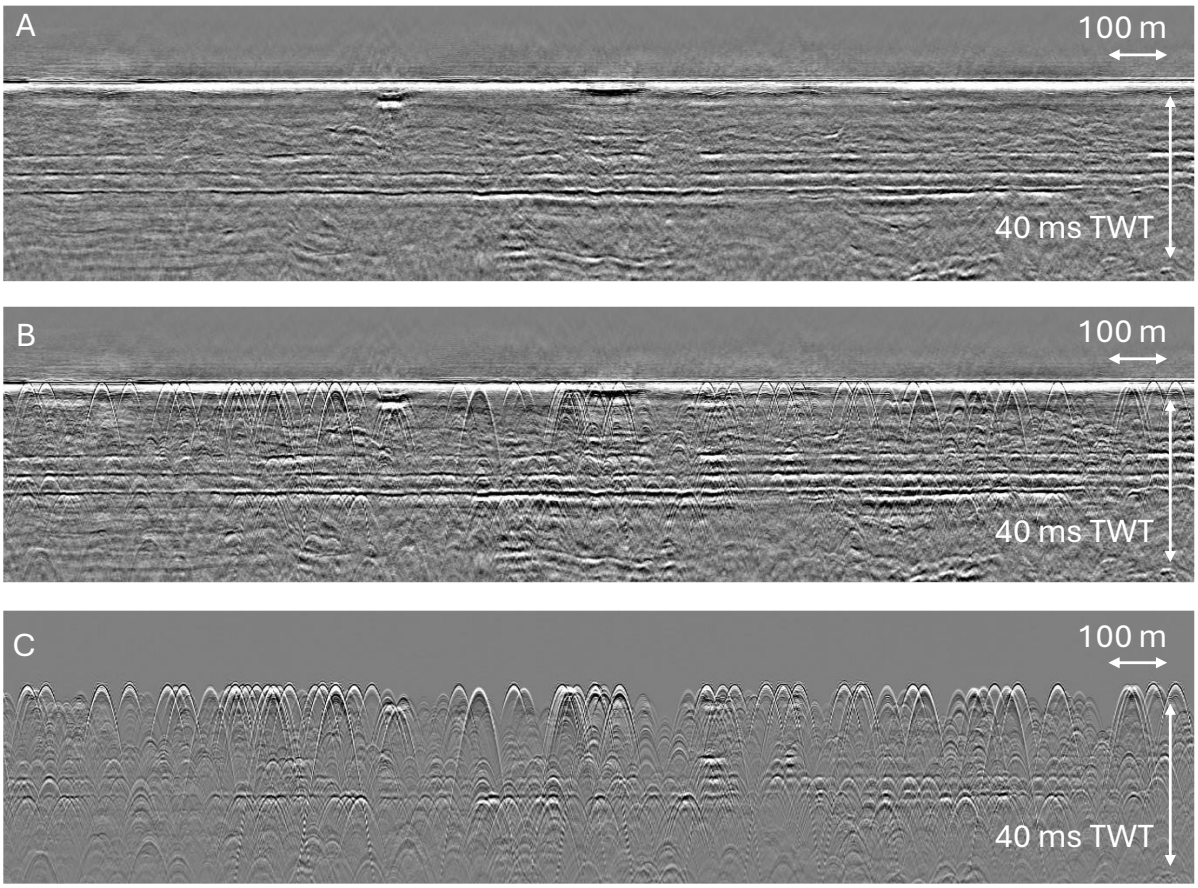

Figure 1 Training data: (A) Cleaned seismic section, supposedly diffraction free; (B) Cleaned section with added synthetic diffractions (input to training); (C) Synthetic diffractions (training labels); these labels were produced using a wavelet derived from the data itself, which is why they resemble the cleaned section. Two-way time (TWT) of 40 ms corresponds to about 40 m in depth.

Results

ML predicted diffractions from a stacked inline are shown in Figure 2. The prediction contains point and edge diffractions emanating from key geological structures; complexity is exemplified by the cluster of overlapping diffractions towards the right of the image. There is minimal leakage of reflection energy. Furthermore, the prediction is relatively free of noise. The data character appears realistic with broadband frequency content and excellent kinematic fit; however, we observe that the amplitudes tend to be underpredicted. The predictions have a somewhat incomplete appearance in some areas, with diffractions occasionally only partially revealed, and some areas appearing blank.

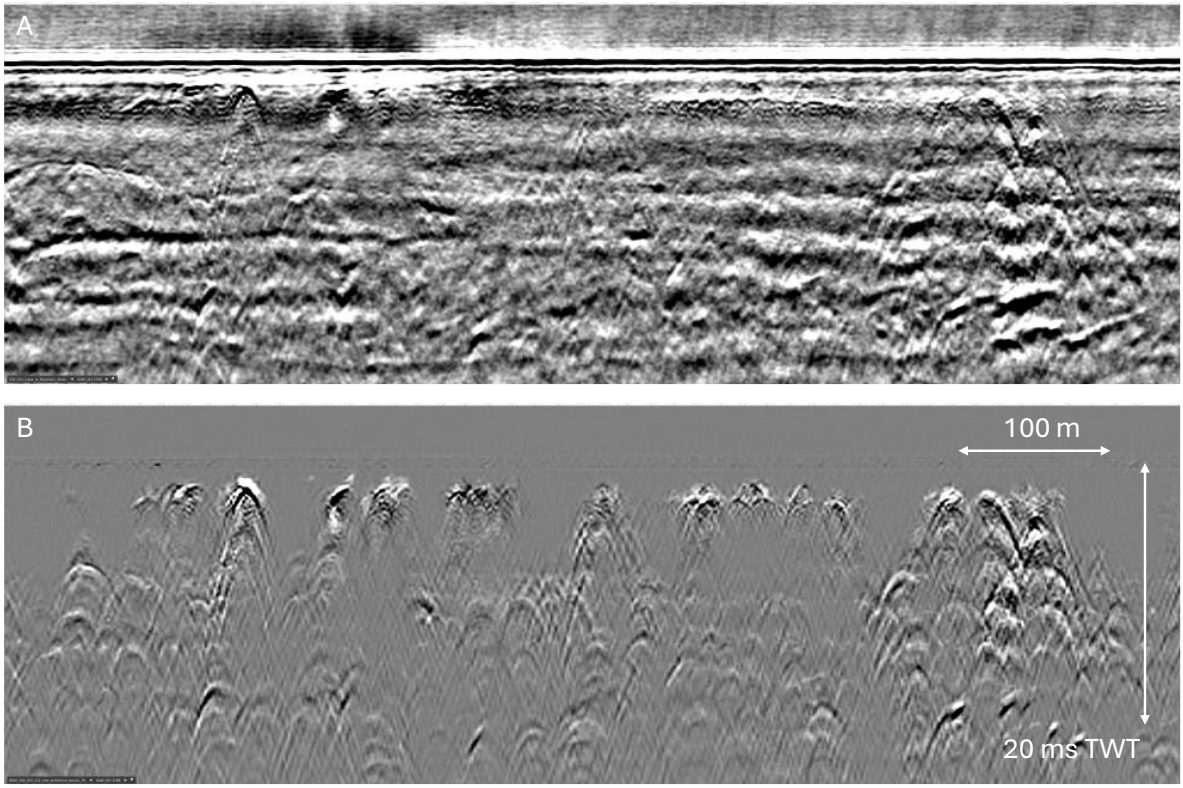

Figure 2 Diffraction extraction shown on a stacked inline in time: (A) input data; (B) extracted diffractions. Two-way time (TWT) of 20 ms corresponds to about 20 m in depth.

To gain further insight we image the diffractions using a 3D Kirchhoff depth migration. A slice through the migrated volume at ~20 m depth beneath the seabed is shown in Figure 3. The image reveals elongated edge structures corresponding to paleo-channel edges and point diffractions possibly from boulders. Even though this network is implemented in 2D it is remarkably robust at detecting consistent diffraction events in the 3D data.

Discussion

We experiment with training models with lower or higher diffraction population density and find this is a strong control on the "aggressiveness" of the resultant network. Too few diffractions and the network is very conservative (extracting very few diffractions), too many and the network becomes too aggressive (extracting some reflections). This fits with Huygen's principle: once the density of diffractions (multiple secondary sources) is high enough they will cohere to form a reflection (single source).

Experiments with the wavelet reveal that using a local window of data leads to more natural looking diffraction models, with reverberations and more data "character". However, this approach may be the cause of the amplitude underprediction problem. Using an idealized spikey wavelet results in much sharper models, with a tendency to capture the amplitude better.

Conclusions

We trained a neural network to extract diffractions from real UHR data. The network successfully extracts a rich diffraction field with minimal leakage of reflections and noise. The results still suffer from a slightly broken up appearance, and a tendency to underestimate the amplitude, but improvements in the synthetics could remedy this. The quality of the extracted diffractions is strongly influenced by the quality of the synthetic labels and therefore better synthetics, with less artefacts and more physics,

should lead to better results. The "aggressiveness" of the network can be controlled by the density of diffractions in the training set. The sharpness of the output can be controlled by the wavelet on the synthetic labels.

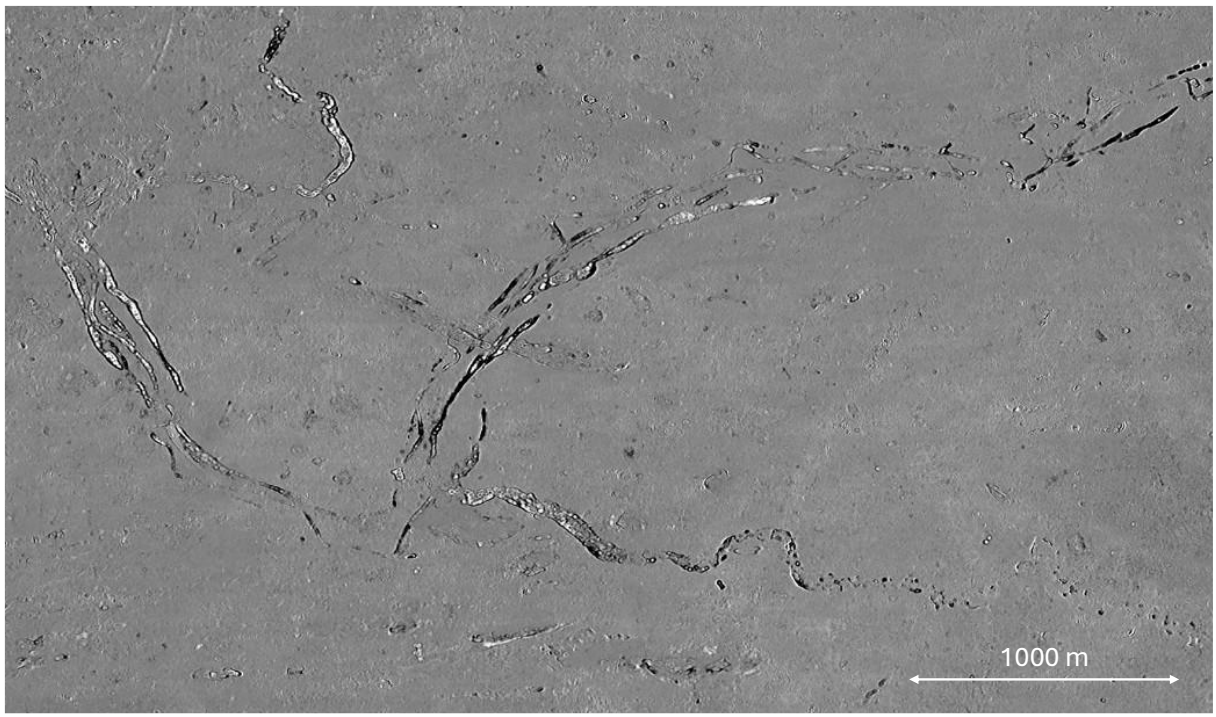

Figure 3 Diffraction image depth slice ~20 m below the seabed.

Acknowledgements

We thank the Netherlands Enterprise Agency RVO for data show rights; TGS for permission to present this work; and our colleagues, in particular Mathieu Lange for help with the dataset and Olga Brusova for help with machine learning.

References

Bauer, A., Schwarz, B., Walda, J., & Gajewski, D. (2025). DiffractioNet: Deep-learning seismic and GPR diffraction separation. *Geophysics*, *90*(3), V255-V264.

Dumitru, R. G., Peteleaza, D., & Craciun, C. (2023). Using DUCK-Net for polyp image segmentation. *Scientific reports*, *13*(1), 9803.

Lowney, B., Lokmer, I., & O'Brien, G. S. (2021). Multi-domain diffraction identification: A supervised deep learning technique for seismic diffraction classification. *Computers & Geosciences*, *155*, 104845.

Schwarz, B. (2019). An introduction to seismic diffraction. In *Advances in geophysics* (Vol. 60, pp. 1-64). Elsevier.

Taner, M.T., Fomel, S., & Landa, E. (2006). Separation and imaging of seismic diffractions using plane-wave decomposition. SEG International Exposition and 76th Annual Meeting 2401-2405.

Wenau, S., Schwarz, B., Bihler, V., Boyer, E. & Preu, B. (2022). Derisking Offshore Windfarm Installation by Sub-Seafloor Boulder Detection Based on Dedicated Seismic Diffraction Imaging. First Break. 40. 41-45.