Deep learning model for full deghosting of Ultra-High Resolution Seismic data with a custom autocorrelation loss

Bagher Farmani and Mohammed Faouzi Zizi, TGS, Oslo*

Summary

Ultra-High Resolution Seismic data (UHRS) are used for exploration activities like wind farm characterization and development. Ghost reflections from the sea surface contaminate primaries and, therefore, should be attenuated in the early stage of processing. Due to the high variability of the sea state, the ghost attenuation of high-resolution data can be quite challenging. We present the use of a deep learning network to simultaneously attenuate both source and receiver side ghosts. This network is trained with the ground truth data obtained from an inversion-based deghosting. We show how an autocorrelation can be imbedded in the loss function to help it cope with any residual ghost energy remaining in the training data, if needed. Using a field data example, we demonstrate that the performance of the trained network on unseen data is robust, and the trained network can be used for both fast track and full integrity processing. Moreover, deghosting of seismic data using a deep learning network can significantly reduce both computation and turn-around times.

Introduction

3D Ultra-High Resolution Seismic (UHRS) data provide high-resolution 3D images together with additional attributes including velocity fields, diffraction images, and physical characteristics of the subsurface. Particularly, shallow subsurface images are widely used for wind farm site characterization. In offshore UHRS survey, seismic sources (e.g., sparkers, boomers) and receivers (hydrophones) are towed below the sea surface. Energy emitted from the seismic sources travels through the water and subsurface. Part of the energy is reflected upwards from the interfaces of the geological layers. There is also an undesired and very strong downward reflection from the sea surface called the ghost. Ghost reflections occur on both the receiver and source side. With every reflection, polarity of the energy changes. In comparison to conventional seismic surveys, UHRS surveys involve receivers towed at shallower depths (e.g. 3m) and sources are towed just below the sea surface (e.g. 0.5 m). To capture finer details of the subsurface, UHRS surveys have very fine temporal and spatial sampling intervals. As the sea state continuously varies during the acquisition, the actual source and receiver depths continuously fluctuate from the nominal depths, which causes a high variability in the ghost arrival times relative to the original reflection. Since receivers are towed deeper than the sources, they naturally exhibit much higher variability.

Geophysicists have developed various processing-based methods to attenuate the ghost and improve imaging results. Traditional deghosting techniques typically rely on plane-wave decomposition and are applied in domains such as the f-k domain (Day et al., 2013) or the τ-p domain (Masoomzadeh and Woodburn, 2013). However, these deghosting techniques often struggle to perform well on UHRS data due to the high variability of the source and receiver depths where the precise depths are often unknown (Provenzano et al., 2020). Inversion-based deghosting methods, such as the approach proposed by Bekara et al. (2024), estimate the depths and suppress ghost reflections using deconvolution in the τ-p domain. While effective, these methods require significant computational resources and precise parameter settings. Inaccurate source and receiver depth estimation can lead to suboptimal ghost removal.

Recent advancements in Deep Learning (DL) offer a faster and automated alternative for UHRS deghosting (e.g. Farmani et al. 2024, Van Borselen and Vasconcelos 2024, Farmani et al. 2025). Van Borselen and Vasconcelos (2024) showed on synthetic data with a source depth of 0.5 m and receiver depths in the range 0.3-1.1 m that a trained DL network can perform good source and receiver deghosting without any knowledge of the actual depth. Farmani et al. (2024) showed on real field data that a trained DL network can perform good receiver deghosting provided that the receivers' depth variation for the unseen data is close to the range of receiver depths of training data. Later, Farmani et al. (2025) showed on real field data that the same network can be used to simultaneously perform both source and receiver deghosting with the quality acceptable for both fast track and full integrity processing. A majority of the studies mentioned above rely on supervised training, meaning that ground truth data should be provided to train the network. However, preparing ground truth data can be challenging as residual ghosts may be present in the data due to imperfect parameterization or very high variability of the ghost. Therefore, it would be desired to partially ease the requirement of having perfect ground truth data for training the network.

To address this challenge, we propose using a custom loss function that consists of two terms: the Mean Absolute Error (MAE) - commonly used in conventional DL methods, and an autocorrelation penalty term based on the autocorrelation function (ACF). The autocorrelation penalty term can effectively suppress unwanted ghost patterns, reducing the DL model's dependence on the ground truth parameterization. By penalizing residual ghost patterns during the training, the model produces cleaner outputs and achieves superior deghosting results. This approach maintains the speed and automation benefits of DL while improving robustness for real field data. Results demonstrate that our autocorrelation loss-based DL method (ACL-DL) can outperform conventional DL methods when there is residual ghost energy in the training data.

Method

First, a small portion of the data in the acquired survey (e.g. one early sequence) is fully deghosted using the method proposed by Bekara et al. (2024). Although a perfect deghosting of the training data is ideal, a small amount of residual ghost is considered acceptable in order to reduce the testing time. A perfect parameterization for each shot gather is time-consuming and difficult to achieve. Then, we train a convolutional neural network called real image denoising network (RIDNet) (Anwar and Barnes, 2019) to reproduce the fully deghosted output. RIDNet is a modular network comprising three main modules: feature extraction, feature learning residual module, and reconstruction. It has a sequence of modules called enhancement attention modules (EAM) which are sequentially connected to each other. Our RIDNet uses 4 EAM boxes. Compared to many other DL networks, RIDNet is a relatively small network with fewer parameters to learn. Therefore, training the model neither requires large hardware resources nor a large amount of training data. It is important that the training data encompasses the expected variation in source and receiver depths. We have observed that the network can handle small deviations of source and receiver depths from the training data. However, as the deviation increases, the network performance becomes more and more suboptimal.

To enable the model to minimize ghost residuals, we designed a custom loss function by incorporating an autocorrelation penalty term informed by expected receiver ghost periods. We compute the ghost period for each trace of the input data based on the estimated receiver depth and acquisition geometry with the following equation:

$$\tau = \frac{2\, r_{depth} \cos(\alpha)}{v} \qquad (1)$$

Here, r_{depth} is the estimated receiver depth, v is the water velocity, and α is the angle of incidence. These ghost periods indicate the expected arrival times of ghost reflections in the input data, allowing us to generate masks m that highlight regions affected by ghost patterns. Let f_θ represent the RIDNet network parameterized by θ, and x be the input seismic data containing ghost energy. During the training stage, input patches are processed by the network to produce a predicted ghost-free signal $\hat{s} = f_\theta(x)$. Our custom loss function $\mathcal{L}$ is defined as the sum of the MAE and an autocorrelation penalty term:

$$\mathcal{L} = \frac{1}{N}\sum_1^N |s_i - \hat{s}_i| + \lambda\, \frac{1}{N}\sum_1^N \left|\hat{s}_i^{autocorr} . m_i\right| \qquad (2)$$

where: s is the ground truth signal, N is the number of data points, $\hat{s}_i^{autocorr}$ = ACF($\hat{s}$) is the autocorrelation of the predicted signal, m_i is a sample mask based on ghost sample locations, and λ is a scaling factor controlling the contribution of the autocorrelation penalty term. The first MAE term measures the errors between the ground truth and the predicted patches. However, if the ground truth is obtained by using a parameterization that does not fully attenuate ghosts in all shot gathers, for example due to ambiguity in receiver depths, the network may learn to retain residual ghost energy. To address this issue, the second term of our loss function applies an autocorrelation penalty to the predicted patches, weighted by the ghost arrival times. High autocorrelation in these regions indicates the presence of ghost residuals. By penalizing this autocorrelation, the model learns to suppress ghost residuals effectively. This approach allows the model to reduce ghost patterns independently of the ground truth parameterization, as the second term of the loss function (equation 2) does not rely on ground truth data. This combination of the MAE term, which ensures signal accuracy, with the autocorrelation penalty term, which targets ghost suppression, maintains a balance between ghost removal and signal preservation. During the training, the network parameters are updated via gradient descent, enhancing RIDNet's ability to outperform conventional loss methods in ghost removal. Finally, the trained RIDNet network is used to quickly and automatically attenuate the source and receiver ghost from all sequences in the survey.

Example

The data used for this study are from a 3D UHRS survey acquired in 2024. The survey was conducted using 10 streamers, each 150 m long. The distance between receivers and streamers were 3.125 m and 12.5 m respectively. Four dual-stacked sparkers were used as seismic sources. The nominal receiver depth was 3 m, and source depth was 0.45 m. Temporal sampling was 0.125 milliseconds corresponding to a Nyquist frequency of 4 kHz. We first chose one of the early acquired sequences. The chosen sequence was fully deghosted using the inversion-based method for generating the desired output for supervised training. There was occasional residual receiver ghost left in the ground truth data used to train the models. Apart from excluding the direct arrival, no additional pre-processing was applied to prepare the data for the training. The RIDNet network was then trained to learn the full deghosting task once with MAE loss function and then, with ACL-DL loss function. Both trained networks were then ready to remove the source and receiver ghost from the other sequences in the survey.

Figure 1 shows 2D QC stacks from an unseen sequence by the network for an outer cable. There is a strong reflector right below the seabed which makes the deghosting task quite challenging. High variability of the receiver ghost is visible in Figure 1a. The inversion-based method was applied with the same parametrization used for the training. It effectively attenuates both source and receiver ghosts (1b). However, it leaves more residual ghost of the seabed and the reflector right below it than the RIDNet models particularly at the left side of the image, pointed by the arrows. In addition, this method has slightly damaged the primary reflector highlighted by oval. RIDNet with MAE loss more successfully attenuates the ghost and better preserves the primary reflector mentioned above (1c). However, residual ghost from the last strong reflector is more visible than the other two methods. RIDNet with ACL-DL loss shows the best balance in primary preservation and the ghost attenuation (1d). Figure 2 shows 2D QC stacks of an inner cable from the same sequence. Inversion- based deghosting attenuates most of the ghost energy but leaves a weak amplitude of the last strong reflector pointed by arrows (2b). Similar partial damage to the primary reflector highlighted by oval is seen in figure 2b. RIDNet with MAE loss leaves more residual ghost of the strong reflector right below the seabed as well as the last strong reflector (2c). RIDNet with ACL-DL loss shows again the best performance with persevering the primaries and also effectively attenuating the ghost of the reflector right below the seabed (2d). Weak residual ghost of the last strong reflector is still visible in the data after the deghosting (2d). Figure 3 shows amplitude spectra of all QC stacks in Figures 1 and 2. In frequency domain, amplitude of all three methods produce comparable results. However, RIDNet with MAE loss shows slightly lower amplitude for frequencies below 500 Hz. This can also be visually observed by looking at reflector around 70 milliseconds in Figures 1 and 2.

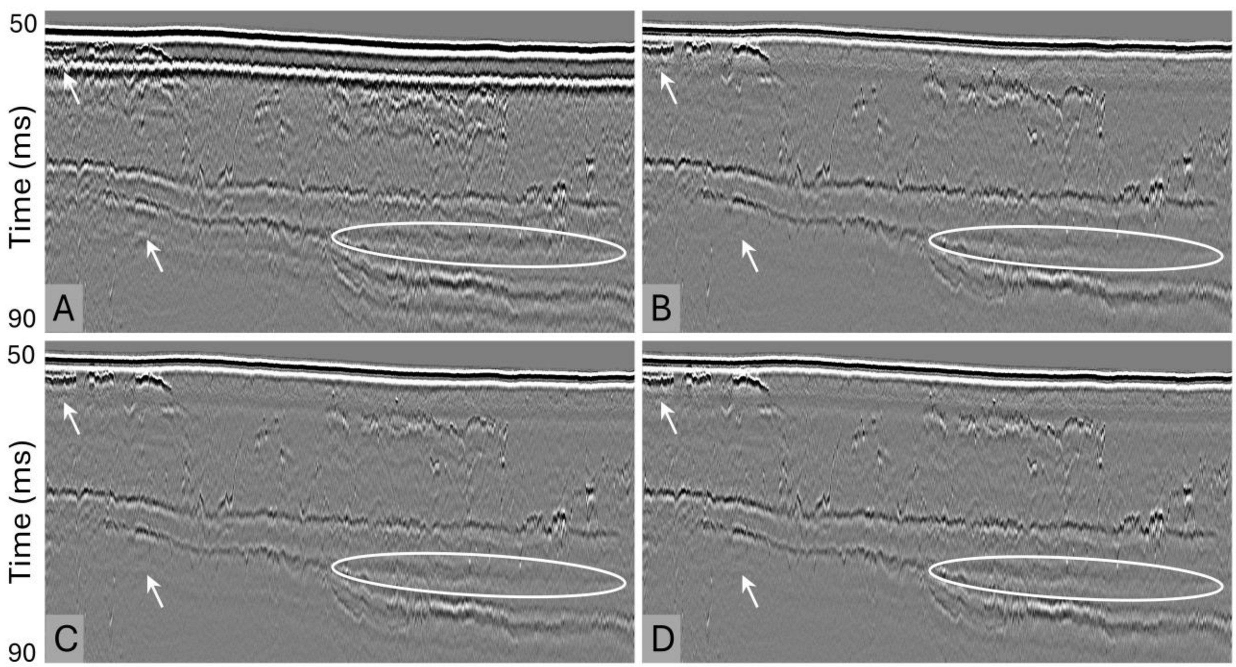

Figure 1: Data example of an outter cable from an unseen sequence by the RIDNet network. 2D QC stack a) before deghosting, b) after inverstion based deghosting, c) after RIDNet with MAE loss deghosting, and d) after RIDNet with ACL-DL loss deghosting. Arrows below the seabed point to the ghosts of the seabed and reflector right below it. Deeper arrows point to the ghost of the last strong reflector. RIDNet with ACL-DL loss deghosting shows the best consistency and performance. Ovals point to a primary reflector better preserved by RIDNet models.

Conclusions

We present UHRS source and receiver deghosting using a convolutional neural network called RIDNet. The RIDNet model was trained using data deghosted by a data-driven inversion-based deghosting method. We also showed that with a custom loss function incorporating an autocorrelation penalty term, we can train a DL network with ground truth data containing a small amount of residual receiver ghost and still teach the network to effectively attenuate the ghost of unseen data. When perfect training data is not available, this method effectively suppresses ghost residuals and reduces reliance on a precise ground truth parametrization. Results show that ACL-DL maintains the speed of DL while delivering cleaner outputs, outperforming conventional DL loss methods across various conditions. Our methodology has been applied to both fast track and full integrity processing with successful performance. The advantages of using this method include removing the need for user parametrization, relaxing the need for perfect ground truth data, maintaining consistency in the performance and significantly reducing both hardware usage and turn-around.

Acknowledgments

We would like to thank RVO and TGS for the permission to publish these results.

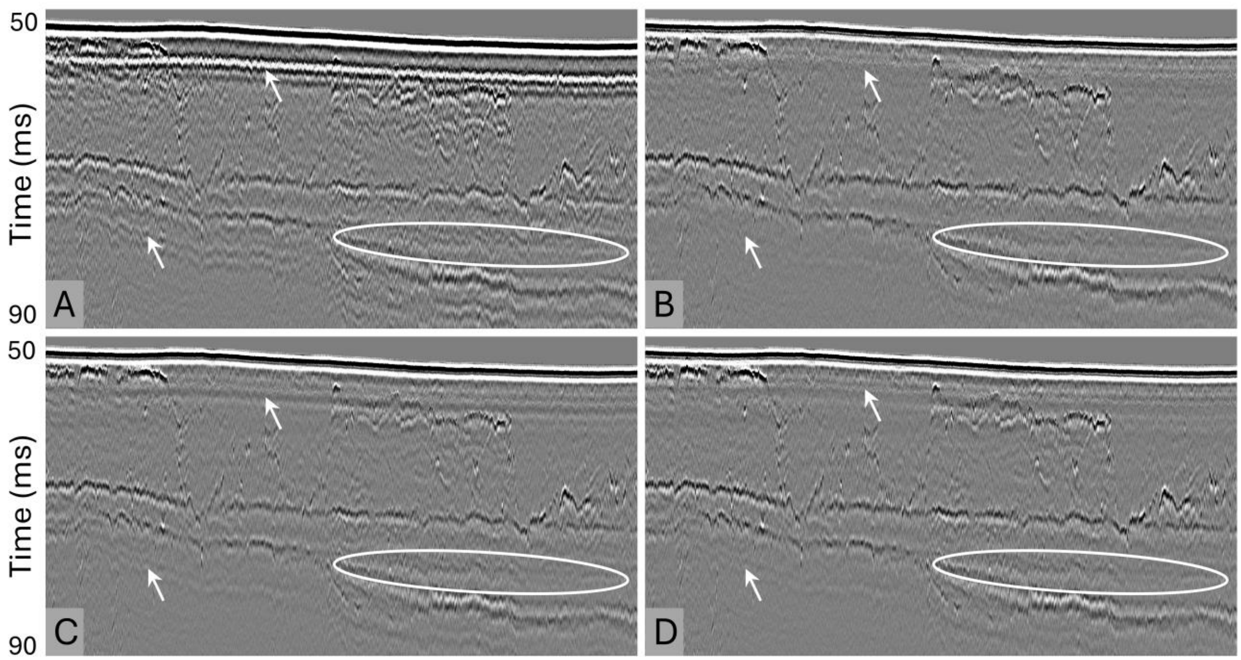

Figure 2: Data example of an inner cable from an unseen sequence by the RIDNet network. 2D QC stack a) before deghosting, b) after inverstion based deghosting, c) after RIDNet with MAE loss deghosting, and d) after RIDNet with ACL-DL loss deghosting. Arrows below the seabed point to the ghosts of the seabed and reflector right below it. Deeper arrows point to the ghost of the last strong reflector. RIDNet with ACL-DL loss deghosting shows the best consistency and performance. Ovals point to a primary reflector better preserved by RIDNet models.

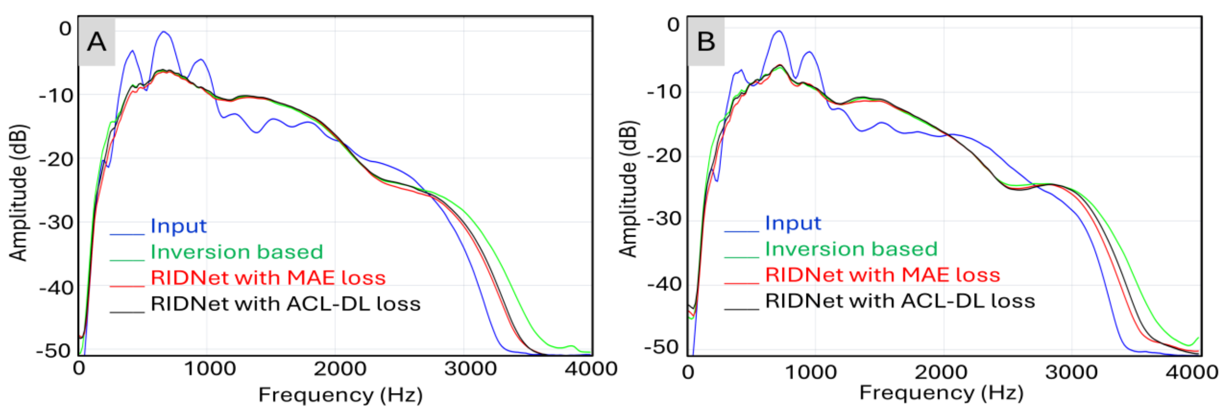

Figure 3: Amplitude spectra of a) outer cable, and b) inner cable, both before and after full deghost using an inversion-based method and two DL methods. This comparison shows that the quality of our DL output is comparable with the inversion-based method.