Motion sensor noise attenuation using deep learning

Bagher Farmani[1]*, Yash Pal[1], Morten W. Pedersen[1] and Edwin Hodges[1] present a workflow consisting of deep learning noise attenuation and curvelet addback to attenuate the noise from motion sensor records.

Abstract

A method is proposed to attenuate both instrumental and environmental noise from motion sensor records in multisensor streamer acquisition. The main elements are two convolutional neural network models. The first model attenuates vertical narrow band high amplitude noise mainly generated by the instruments attached to the streamers. The second model attenuates widespread background noise mainly associated with environmental conditions. To reduce the risk of possible signal loss an addback flow in the curvelet domain is used. The motivation for the work presented here was to develop a fully automated noise attenuation method that eliminates the need for time-consuming and subjective user parameter testing. The method has been validated using seismic data from different parts of the world and shown to consistently produce superior results to other state-of-the-art noise attenuation processes.

Introduction

Multisensor streamers record pressure and particle motion using collocated sensors. These recordings can be combined to separate the wavefield into up- and down-going parts, thereby mitigating the effect of sea surface ghost reflections (e.g., Day *et al*., 2013). Noise attenuation remains an essential pre-conditioning step for optimal wavefield separation. Noise attenuation must be applied to data from all sensors, but the particle motion data presents particular challenges. Prior to the emergence of deep-learning-based algorithms, the attenuation of noise in seismic data was most often performed using statistical and mathematical tools. The most popular algorithms were developed using statistical methods for the noise detection and FX filters for the noise attenuation (e.g., Bekara and van der Baan, 2010; Chen and Sacchi, 2017). The challenge of using such algorithms is that depending on the level and extent of the noise, the user needs to test a number of parameters to ensure adequate noise attenuation without signal loss. Therefore, a time-consuming testing phase is required. Often, the degree of noise varies within a survey and the testing needs to be repeated. A key motivation of our work is to remove such subjective judgments on choice of parameters from the noise attenuation process. In addition, FX-based algorithms have some inherent limitations. For example, if the noise detection relies on an amplitude-based statistical method, local high amplitude signal could be identified as noise or spatially wide bands of noise identified as signal. In addition, FX filters cannot easily reconstruct the signal where the noise has a large lateral extent, and the signal-to-noise ratio is low.

In recent years, the application of machine learning technologies to seismic noise attenuation has received considerable interest. Such applications can be categorized into three main groups: quality control of the noise content and signal loss; detection of noise and signal loss for automatic guiding of the noise attenuation engine; and finally, direct noise attenuation. Inspired by the great success of machine learning for image classification, different machine learning methods have been used to classify seismic records based on their noise content (Bekara and Day, 2019; Farmani and Pedersen, 2020a; Walpole *et al*., 2020). Such applications are often used before noise attenuation to evaluate the noise content, after noise attenuation to evaluate the effectiveness of the noise attenuation and, sometimes, to detect possible signal loss. Machine learning-based noise classification can also be used inside the noise attenuation algorithm to perform the noise detection step (Farmani and Pedersen, 2020b; Farmani and Pedersen, 2022). Based on the noise level, type and extent, all samples in the seismic records are classified and passed to different noise attenuation processes which are not machine learning based. This process is ideally applied in an iterative fashion. Finally, more recent advances in image denoise using deep learning inspired our community to develop the third group of machine learning applications where seismic noise is directly attenuated using deep learning regression (Kumar *et al*., 2022; Valenciano *et al*. 2022).

We present a workflow that uses deep learning to directly attenuate the noise from motion sensor records. The heart of the workflow is a convolutional neural network called real image denoising network (RIDNet). RIDNet was originally designed to denoise photographic images (Anwar and Barnes, 2019). To make the workflow robust and generic, we use two RIDNet models that target different types of noise found in motion sensor records. Both models were trained using supervised learning and operate on the frequency band 19-95 Hz. FX filters are used to attenuate the noise at frequencies outside 19-95 Hz. We also use

[1] PGS

* Corresponding author, E-mail: bagher.farmani@pgs.com

DOI: 10.3997/1365-2397.fb2023010

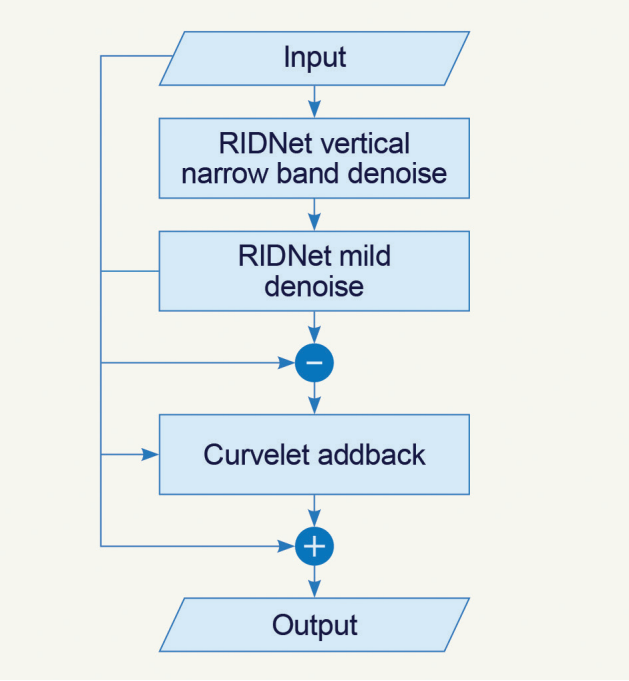

Figure 1 Schematic of the motion sensor noise attenuation workflow.

an addback process in the curvelet domain to minimise the risk of signal loss when the workflow is applied to any new survey which might have slightly different characteristics to the data used for training and where these differences might cause local weak signal loss. The workflow has been extensively verified using data from a wide variety of geographical locations and has proved to be robust and automated. The workflow has also been used in full-scale production.

Noise attenuation workflow

Our workflow comprises three main steps shown in Figure 1. First, motion sensor records are passed into a RIDNet model where vertical narrow band noise with high local amplitudes are attenuated. Then, the output from the first RIDNet model is passed into another RIDNet model where the background noise is attenuated. To ensure no signal loss due to the noise attenuation, the output of the second RIDNet model and attenuated energy are transformed to the curvelet domain where possible signal losses are recovered. The recovered signals are added back to the outputs of the second RIDNet model to generate the final output. In the following sections we explain these three steps in more detail.

Choice of network architecture

After initial evaluation of a suitable network for attenuating noise from motion sensor records, we chose RIDNet as a suitable candidate. To improve both runtime and quality performance we made some minor modifications to the original network. Figure 2 shows the architecture of the network. RIDNet is a modular network comprising three main modules: feature extraction, feature learning residual module, and reconstruction. Feature extraction consists of one 2D convolution layer that extracts initial features. These features are then passed to a sequence of modules called enhancement attention modules (EAM). The first part of EAM branches the input features and passes them through two dilated 2D convolutions. Dilated convolution is similar to convolution but it involves sample skipping to cover a larger area of the input. Then features are further passed through some 2D convolutions and local skip connections. By using dilation and local skip connections the network can learn both low-frequency and high-frequency features in the input data. Since EAM boxes are connected sequentially, it is easy to find the optimum number of EAM boxes needed for any specific task. The output of the last EAM is passed to a 2D convolution layer to reconstruct the noise with opposite polarity. Finally, the reconstructed noise is added back to the input and the final output is generated.

Training the noise attenuation models

The main noise on motion sensor records can be roughly categorised to two groups: noise generated by the devices attached to the streamer (e.g., depth and lateral control devices) and widespread background noise owing to the environmental conditions (e.g., sea state and barnacle growth). Therefore, we trained two separate RIDNet models with identical network structure as shown in Figure 2 to attenuate each of these two groups of the noise. We call the RIDNet models 'vertical narrow band denoise model' and 'mild denoise model'. Both models were trained using supervised learning.

During training we randomly divided the data into 80% for training and 20% for validation. We tested different input sizes, number of convolution filters, kernel size, number of EAM

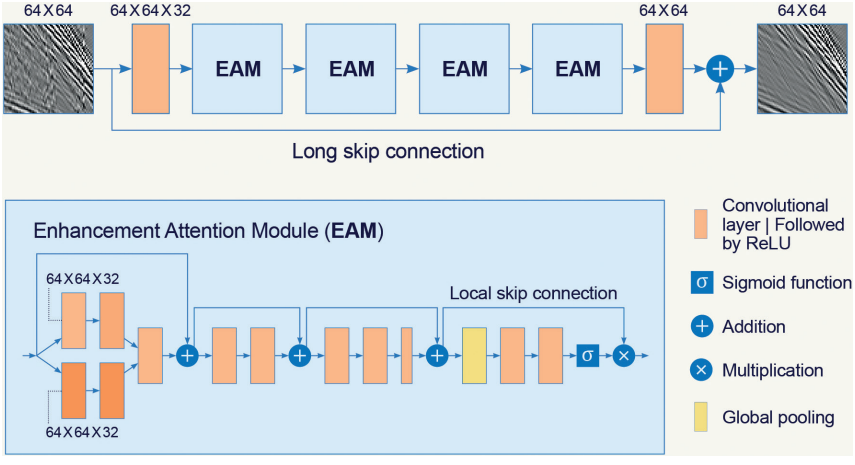

Figure 2 Schematic of the RIDNet convolutional neural network architecture.

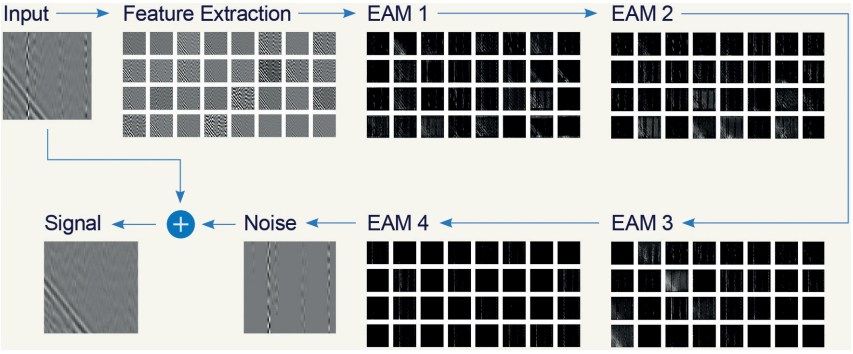

Figure 3 Input, feature map of the output of each module, and final output for the trained RIDNet vertical narrow band denoise model. Note that the vertical narrow band noise has been correctly identified as such by the network.

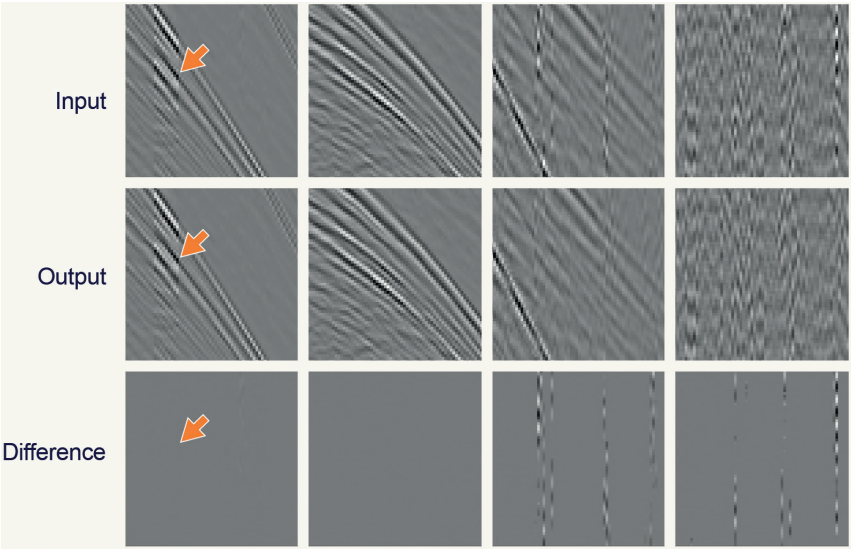

Figure 4 Noise attenuation using the trained RIDNet vertical narrow band denoise model on some example validation tiles. Note that the network has successfully identified areas of vertical narrow band noise whilst doing almost nothing where no such noise is present.

modules, loss function, optimiser and learning rate. In the end, we chose an input size of 64x64 samples, 32 convolutional filters, kernel size of 3x3, mean absolute error loss function, 4 EAM modules, Adam optimiser with dynamic learning rate starting from 0.001. The only normalisation we use is that the maximum amplitude of input tiles to the models is normalised to one. Hence, our RIDNet models are not sensitive to the amplitude of the input data.

We used 600,000 tiles from selected shot gathers from a few selected surveys acquired in different parts of the world to train the vertical narrow band denoise model. The input tiles contained energy in the frequency band 19-95 Hz. The desired output for this model was generated using the hybrid workflow presented by Farmani and Pedersen (2022). The hybrid workflow uses models with U-Net architecture (Ronneberger *et al*. 2015) to automatically classify the samples and pass them to the appropriate noise attenuation steps. With targeted noise detection in the hybrid workflow, a better-targeted noise attenuation with appropriate levels of harshness can be achieved compared to statistical detection and FX filtering using a fixed threshold. To force the model not to mistake local high-amplitude signal with noise, we intentionally boosted the amplitude of signal on some random narrow band traces with a random ratio on both input and desired output. Our training could reach peak signal-to-noise ratio (PSNR) of 37.1 on validation data. Figure 3 shows the input, feature map of the output of each module, and final output for the trained RIDNet vertical narrow band denoise model. Note how the model progressively and effectively separates the signal from the noise. Figure 4 shows input, output, and attenuated noise for some validation tiles. The arrow in the first tile shows the traces with intentionally boosted amplitude. The RIDNet model learnt to recognise these traces as signal and kept them untouched. The second tile contains only signal and the RIDNet model does not remove any meaningful energy from the input as expected. The narrow band noise on the third tile is attenuated without any visible signal loss. The fourth tile contains both background noise and local high-amplitude noise. Note that the model only attenuates the noise where it stands out from the neighbouring samples. Background noise is left in the output and is targeted with the subsequent RIDNet mild denoise model.

We also used 300,000 tiles from selected shot gathers containing energy in the frequency band 19-95 Hz to train the RIDNet mild denoise model. The desired output of this model was generated by a workflow consisting of intermediate RIDNet models, FX filters and signal addback. The desired output was intentionally created to mildly attenuate the background noise such that signal-to-noise ratio is improved but all the noise is not necessarily removed. In the workflow (Figure 1), the RIDNet mild denoise model is applied to the output of the RIDNet vertical narrow band denoise model. Therefore, most of the stronger noise associated with devices attached to the streamer is already attenuated in the input to this model. The task of this model is to reduce the background noise arising from environmental conditions such as barnacle noise or swell noise.

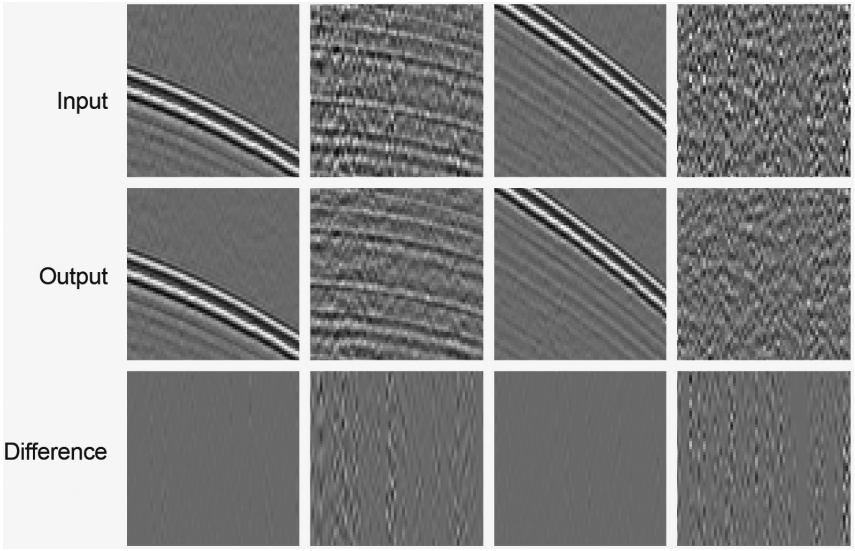

Figure 5 Noise attenuation using the trained RIDNet mild denoise model on some validation tiles. Note that this network has successfully suppressed widespread background noise whilst preserving the signal.

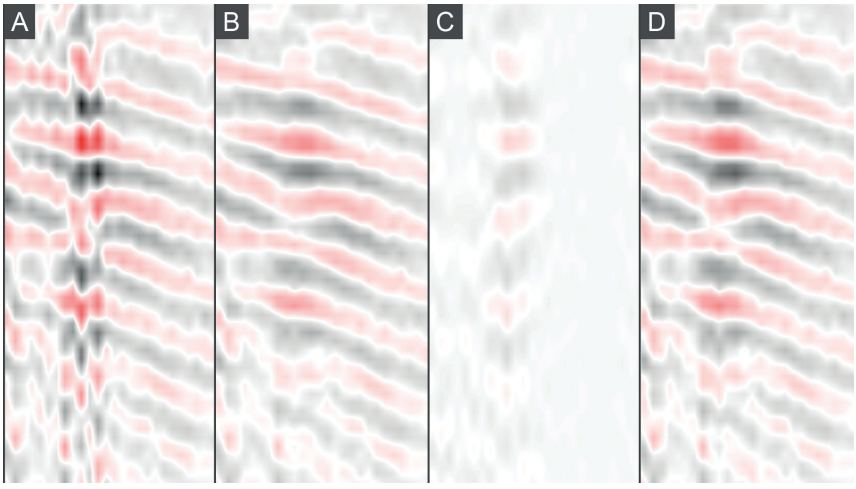

Figure 6 Input to RIDNet models (a), output of the RIDNet models (b), signal estimated by the curvelet addback flow (c) and final output (d). Note that the energy identified by the curvelet addback flow in panel (c) correlates well with the initial signal estimate in panel (b).

Our experience showed that, by forcing the model to perform only mild denoise, the risk of signal loss was greatly reduced. Our training could reach PSNR of 35.2 on validation data. Figure 5 shows input, output, and attenuated noise for some validation tiles. Note that the model tries to reduce the noise for all samples in all tiles. When the noise is stronger as in the second and fourth tiles, the model also attenuates more noise. However, the outputs of all tiles do not have a similar level of signal-to-noise ratio. After inference the second and fourth tiles are still noisier than the first and third tiles, but overall the total signal is preserved and variance of the noise is reduced.

Curvelet domain process for combating signal loss

During the extensive validation process using many datasets from different parts of the world, we observed very weak local signal leakage on two lines from different surveys. As the intention is to provide a workflow that is as generic and automated as possible, we used an addback flow to recover the signal loss if it happens. The inputs to the addback flow are the output from the RIDNet mild denoise model and the attenuated energy. First, both inputs are transformed into the curvelet domain. The curvelet coefficients with highest signal amplitude are identified and used to create a mask (Nguyen and Liu, 2017). This mask is then used to select the corresponding curvelet coefficients in the energy removed, which are then scaled and then transformed back to the TX domain to provide an estimate of signal that is present in the attenuated energy. The residual noise in the estimated signal loss is further attenuated, and the remaining energy is added back to the output of RIDNet mild denoise model.

Figures 6a and 6b show a field data example of the input to and output from the two RIDNet models, respectively. The local signal loss estimated by the addback flow is shown in Figure 6c, which is added to Figure 6b to create the final noise attenuated output (Figure 6d). On field data it is not easy to say for sure which one of Figures 6b and 6d are the closest to the earth response by inspection of the shot gathers and the removed energy. However, since the energy added back by the addback flow are from the same curvelet coefficients as signal, signal correlations between the final removed and preserved energy are naturally reduced. Therefore, the workflow will be more signal friendly. The best domain for evaluating the signal loss is common midpoint stack. In the stack domain, the energy removed should not show any correlation with the energy preserved on visual inspection. We use this quality control (QC) product in the following examples to demonstrate the robustness of the workflow.

Examples

The first data example is from the Outer Moray Firth survey acquired using a multisensor streamer northeast of Aberdeen in 2018. The water depth in the survey area was between 80 m and 140 m. In seismic streamer acquisition, such water depths are considered shallow. Attenuating noise from shallow water data has been always proved challenging in the past. Due to the shallow seabed, water bottom reflections and refractions as well as their short period multiples usually contain high amplitude energy with high moveout. Traditional algorithms such as FX filters usually rely on thresholding to detect the noise in the application window. It is obvious that there is a high risk that high amplitude energy with high moveout will be erroneously identified as noise if thresholding methods are used. To avoid this, geophysicists usually design a protection area around the water bottom and its short period multiples where different noise attenuation techniques or parameters are used. This is a manual process that is time-consuming and subjective, especially in cases where energy from adjacent shots is present in the same record. Figure 7 shows some shot gathers before and after noise attenuation with the proposed workflow. The workflow was able to remove almost all the noise from the shot gathers. The attenuated energy is shown in Figure 7c. Note how the RIDNet models have learnt to naturally protect the areas with high-amplitude signal around the water bottom and its short period multiples. The workflow does not remove any meaningful noise where signal-to-noise ratio is high. If one wants to remove more noise from the naturally protected areas, this can be done with, for example, a simple FK filter applied to energy outside the signal cone. Figure 8 shows a 2D QC stack of part of the line contaminated by strong environmental noise before and after noise attenuation as well as the attenuated energy from the records. The noise is effectively attenuated, and the stack of the attenuated energy does not show any meaningful correlation with the signal in the noise-attenuated output.

The second example is from a survey acquired offshore Guinea using a multisensor streamer in 2019. The water depth for this survey changed from very shallow to very deep. Therefore, this example contains almost all possible water depths in streamer acquisition. Figure 9 shows input, output, and attenuated energy from the records. The main noise in the motion sensor records for this line was from the devices attached to the streamer. However, as acquisition of the line progressed from the shallow to the deep part, the level of the background noise increased on the records due to the gradual attachment and growth of barnacles. Therefore, the first input shot is noisier than the other two (Figure 9a). As mentioned previously, our workflow applies mild background noise attenuation. Hence, the first output shot is still noisier than the other two output shots after noise attenuation (Figure 9b) but much cleaner than the input shot. The level of remaining noise was considered acceptable. Note also how strong signals are

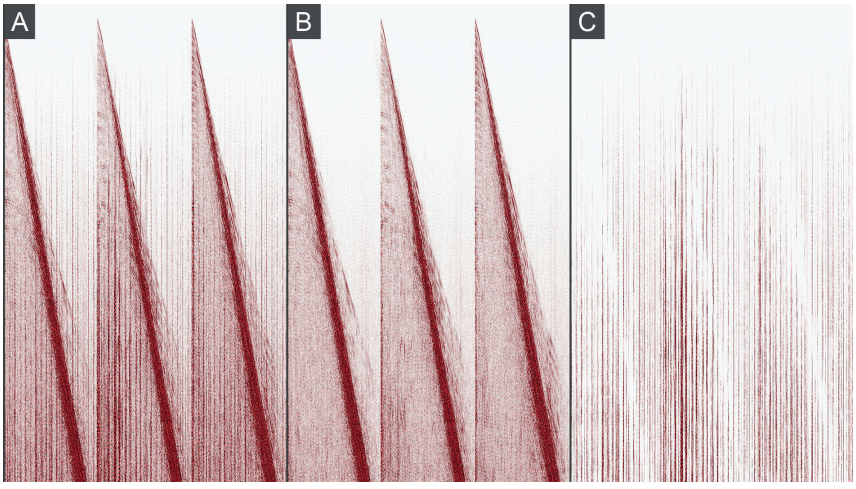

Figure 7 Example motion sensor shot gathers from field data acquired offshore Scotland: input (a), output (b) and attenuated energy (c) using the proposed noise attenuation workflow. The displays have a 20-25 Hz Ormsby lowcut filter to reflect the typical bandwidth for which motion sensor data are used in wavefield separation.

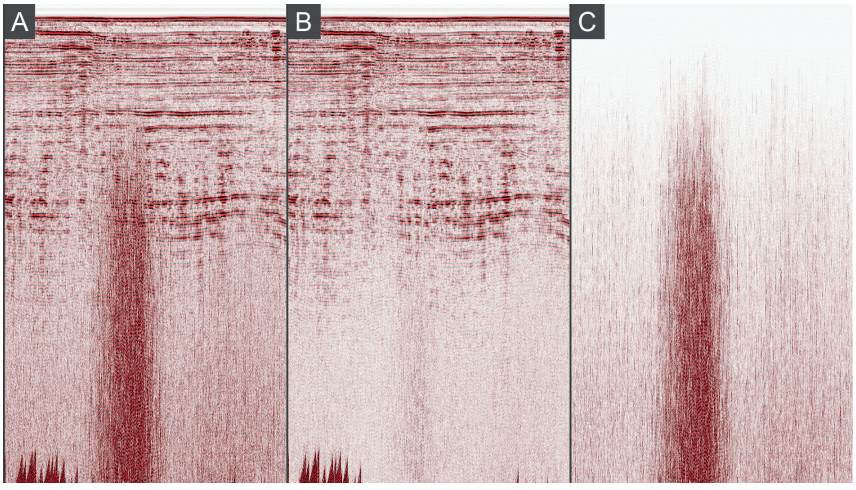

Figure 8 Motion sensor 2D QC stack: input (a), output (b) and attenuated energies (c) using the proposed noise attenuation workflow. The displays have 20-25 Hz Ormsby lowcut filter to reflect the typical bandwidth for which motion sensor data are used in wavefield separation.

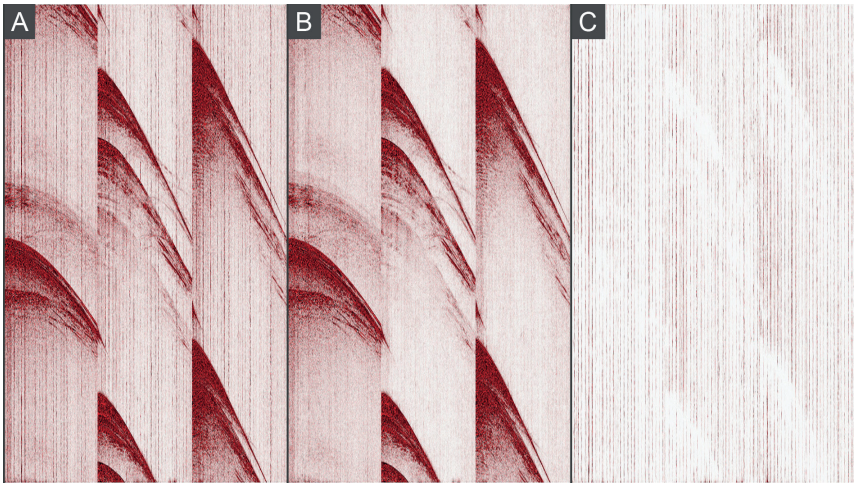

Figure 9 Example motion sensor shot gathers from field data acquired offshore Guinea: input (a), output (b) and attenuated energy (c) using the proposed noise attenuation workflow. The displays have a 24-30 Hz Ormsby lowcut filter to reflect the typical bandwidth for which motion sensor data is used in wavefield separation.

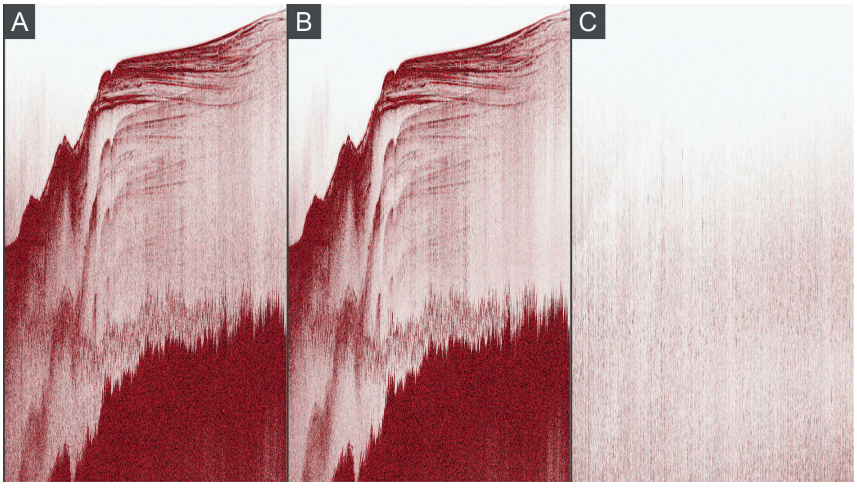

Figure 10 Motion sensor 2D QC stack for data acquired offshore Guinea: input (a), output (b) and attenuated energy (c) using the proposed noise attenuation workflow. Displays have a 24-30 Hz Ormsby lowcut filter to reflect the typical bandwidth for which motion sensor data is used in wavefield separation.

again naturally protected. Figure 10 shows a 2D QC stack of the whole line before and after the noise attenuation as well as the attenuated energy from the records. It is clear in Figure 10c that more noise is removed from the deepwater part of the line, where the input has more background noise.

Conclusions

A motion sensor noise attenuation workflow has been designed where the heart of the workflow consists of two convolutional neural network models with RIDNet architecture. These models were trained using motion sensor data recorded in the field with supervised learning where the desired outputs were produced by alternative workflows. The first RIDNet model attenuates locally high amplitude vertical narrow band noise. Then, data are passed to the second RIDNet model where widespread background noise is attenuated. The second model was trained to perform mild noise attenuation to reduce the risk of possible loss of signals with weak amplitude. To make the workflow as generic and automated as possible, we also designed an addback process in the curvelet domain to recover any possible signal loss from the attenuated energy. The workflow has been extensively validated using records from different surveys acquired in different parts of the world. Throughout the validation process, there was no need for any workflow modification and, therefore, the workflow is automated and does not require any user interaction. The parameter testing phase that is needed for conventional noise attenuation approaches has been eliminated. Compared to the original RIDNet and many other convolutional neural network architectures used for seismic noise attenuation, our final RIDNet architecture has fewer parameters and, therefore, uses less hardware resources for inference. The workflow was recently used for a full survey production. Our experience so far shows that the proposed workflow can attenuate more noise from the motion sensor records than what could be achieved with the previous workflows and at the same time it has a better signal preservation without any user interaction.

Acknowledgements

We would like to thank PGS for the permission to publish these results. We also thank our colleagues in PGS for their support and feedback.

References

Anwar, S. and Barnes, N. [2019]. Real Image Denoising with Feature Attention. IEEE International Conference on Computer Vision (ICCV).

Bekara, M. and van der Baan, M., [2010]. High-amplitude noise detection by the expectation-maximization algorithm with application to swell-noise attenuation. *Geophysics*, **75**(3), pp.V39-V49.

Bekara, M. and A. Day. [2019]. Automatic QC of denoise processing using a machine learning classification. *First Break*, **37**(9), 51-58.

Chen, K. and Sacchi, M.D. [2017]. Robust f-x projection filtering for simultaneous random and erratic seismic noise attenuation. *Geophysical Prospecting*, **65**(3), pp.650-668.

Day, A., Klüver, T., Söllner, W., Tabti, H. and Carlson, D. [2013]. Wavefield-separation methods for dual-sensor towed-streamer data. *Geophysics,* **78**(2), pp.WA55-WA70.

Farmani, B. and Pedersen, M.W. [2020a]. Application of a convolutional neural network to classification of swell noise attenuation. SEG Technical Program Expanded Abstracts: 2868-2872.

Farmani, B. and Pedersen, M.W. [2020b]. Application of Convolutional Neural Network in Automated Swell Noise Attenuation. EAGE 2020 Annual Conference & Exhibition Online, Dec 2020, Volume 2020, p.1-5.

Farmani, B. and Pedersen, M.W. [2022]. Stepping Towards Automated Multisensor Noise Attenuation Guided by Deep Learning. 83rd EAGE Annual Conference & Exhibition, Jun 2022, Volume 2022, p.1-5.

Kumar, A., Dancer, K., Rayment, T., Hampson, G. and Burgess, T. [2022]. Deep Learning Swell Noise Estimation. *First Break*, **40**(9), 31-36.

Nguyen, T. and Liu, Y.J. [2017]. Seismic Noise Attenuation Using Curvelet Transform and Dip Map Data Structure. 79th EAGE Conference & Exhibition 2017, Jun 2017, Volume 2017, p.1-5.

Ronneberger, O., Fischer, P. and Brox, T. [2015]. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, LNCS, Vol.9351: 234-241.

Valenciano, A., Brusova, O. and Cheng, C. [2022]. Efficient Swell Noise Removal Using a Global Deep Neural Network Model. *First Break*, **40**(2), 51-55.

Walpole, J., Hallett, K., Brown, E. and Brittan, J. [2020]. Visual Identification of Noisy Seismic Records with Machine Learning. EAGE 2020 Annual Conference & Exhibition Online, Dec 2020, Volume 2020, p.1-5.