Dual-band generative learning for low-frequency extrapolation in seismic land data

Oleg Ovcharenko, Vladimir Kazei‡, Daniel Peter*, Ilya Silvestrov †, Andrey Bakulin †, Tariq Alkhalifah **
* *KAUST, Thuwal, Saudi Arabia*
† *EXPEC ARC, Saudi Aramco, Dhahran, Saudi Arabia*
‡ *Aramco Services Company: Aramco Research Center – Houston, USA*

SUMMARY

The presence of low-frequency energy in seismic data can help mitigate cycle-skipping problems in full-waveform inversion. Unfortunately, the generation and recording of low-frequency signals in seismic exploration remains a non-trivial task. Extrapolation of missing low-frequency content in field data might be addressed in a data-driven framework. In particular, deep learning models trained on synthetic data could be used for inference on the field data. Such an implementation of switching application domains remains challenging. We, therefore, propose the concept of generative dual-band learning to facilitate the knowledge transfer between synthetic and field seismic data applications of low-frequency data extrapolation. We first explain the two-step procedure for training a generative adversarial network (GAN) that extrapolates low frequencies. Then, we describe the workflow for synthetic dataset generation. Finally, we explore the feasibility of the dual-band learning concept on real near-surface land data acquired in Saudi Arabia.

INTRODUCTION

Imaging for deep subsurface targets is incredibly challenging when applied to land seismic data (Bakulin et al., 2018). Unlike a water layer in the marine acquisition, a land near-surface layer is naturally inhomogeneous. Land surveys feature variable topography and highly complex structures in the first few hundreds of meters below the surface. Also, the nature of the topmost layer affects the coupling of the sensors and sources with the medium. It thus contributes to higher noise levels in recorded data. Another distinct feature of land data is strong surface waves that dominate body-wave signals conventionally used for seismic imaging. Therefore, a successful inversion of land data requires building an accurate model of the near-surface (Baeten et al., 2013) since inaccuracies accumulated in the shallow subsurface considerably magnify at depth.

Data-driven methods can be found in a broad range of applications in geophysics (Alali et al., 2020; Sun and Alkhalifah, 2020; Song et al., 2021). We focus on initial velocity model building, which can be approached in both data and model domains. Model-domain approaches might predict low-wavenumber velocity models directly from data (Kazei et al., 2020b,a; Zwartjes, 2020; Plotnitskii et al., 2019), whereas data-domain approaches concentrate on extrapolation of low-frequency content of seismic data, subsequently used by classic imaging algorithms (Aharchaou et al., 2020; Ovcharenko et al., 2017, 2019, 2020; Fabien-Ouellet, 2020; Hu et al., 2020; Wang et al., 2020; Sun and Demanet, 2019, 2020).

The supervised learning framework is commonly used for solving inverse problems in a data-driven fashion. The bottleneck, however, is the lack of realistic training datasets, sufficient for direct inference on field data. Unlike seismic data interpolation, which might be addressed in an unsupervised fashion (Ovcharenko and Hou, 2020), the task of low-frequency extrapolation suffers from the lack of labeling data for training. Meaning that only high-frequency input data is known from seismic surveys, while the low-frequency label is a derivative of a solution of an ill-posed inverse problem of waveform inversion.

Usually, constructing a dataset for supervised learning using exclusively real-world data is not feasible. The workaround is to create a dataset of synthetic input-target pairs and leverage the concept of transfer learning to migrate the learned knowledge between datasets (Siahkoohi et al., 2019). However, the datasets should be statistically similar. Despite many approaches proposed for generating realistic seismic data (Kazei et al., 2019), the knowledge transfer between synthetic and field applications remains unresolved.

We propose the concept of dual-band learning designed to facilitate the knowledge transfer between training on synthetic data and application on field data. The key idea lies in injecting field data samples into training alongside with the synthetic data flow. This approach with modifications might also be used in other geophysical applications.

DUAL-BAND LEARNING

The generative dual-band learning for low-frequency extrapolation implies using two frequency bands of seismic data as input channels to the network. The implementation also requires a two-stage training of a GAN, with synthetic and field data being used jointly.

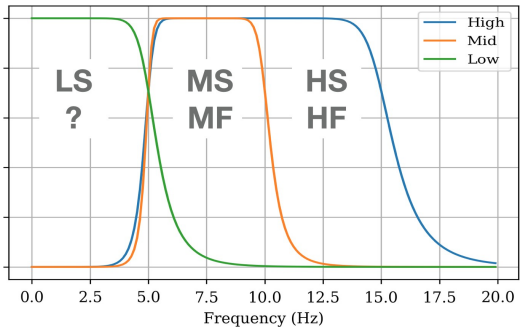

Figure 1: Bandpass filters to split data into high, middle and low frequency partitions. The high and middle partitions are available in both synthetic and field data, while the low-frequency target only exists in the synthetic dataset.

Assume there are two datasets of full-band seismic shot gathers. The training dataset, D_S, which consists of synthetic data, and the testing dataset, D_F, which comprises the recorded field data. Commonly, the elements of each dataset are split into high-frequency inputs and low-frequency targets by applying low-and high-pass filters. For the synthetic dataset, these are H_S and L_S, respectively, while for the field dataset, only the high-frequency partition H_F is available. The low-frequency pair of field data L_F is the ultimate target of the entire application. We also propose to extract the band of "medium" (denoted as "mid," for brevity) frequencies, M_S and M_F, from high-frequency bands of synthetic and field data, respectively (Figure 1). This additional band will be used to inject field data into synthetic training.

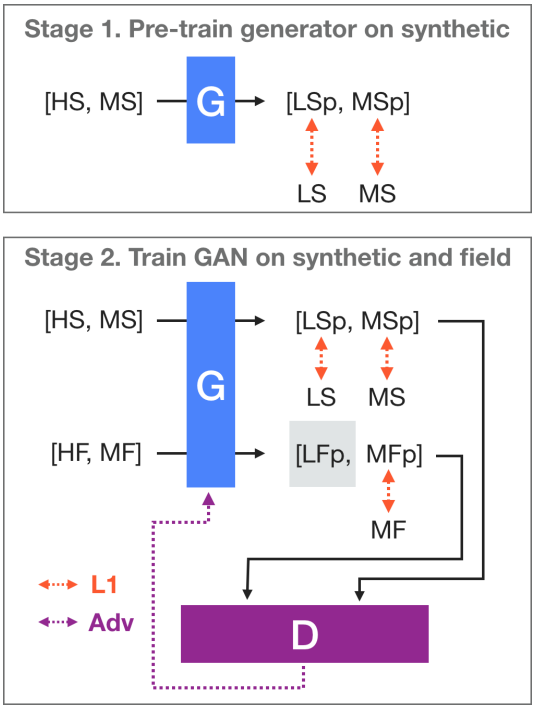

Figure 2: Two-stage workflow for dual-band learning. H_S, M_S, and L_S stand for High, Medium, and Low-frequency bands of synthetic data. H_F and M_F are the corresponding High and Medium frequency bands from the field dataset. The "p" subscript denotes predicted low-frequency and middle-frequency data. An accurate reconstruction of L_{Fp} is the ultimate goal.

The network architecture that implements the dual-band learning concept is a generative adversarial neural network (Goodfellow et al., 2014). The GAN consist of a generator G- and discriminator D-network. The G-network accepts the input volume of the data and outputs the predicted data. Meanwhile, the D-network operates simultaneously on reference and predicted data, attempting to tell whether these data came from the same or different distributions. Training of the GAN is highly unstable by design, where G- and D-networks are competing with each other. For this reason, we develop a two-stage strategy that would focus the training on predicting low-frequency data.

In the first stage (Figure 2), generator G is trained exclusively on synthetic data. In particular, G learns the mapping of the combination of high- and mid-frequency data into low- and mid-frequency data, $[H_S, M_S] \rightarrow [L_S, M_S]$, by optimizing the L1-misfit between those. The network also learns the relation between predicted middle and low frequencies. In the second stage, the training on synthetic data continues, alongside feeding the input pair $[H_F, M_F]$ into G. For the field data pair, only gradients based on the medium frequency M_F's prediction are backpropagated. At the same time, we trained the discriminator D to distinguish between the true pair of synthetic $[L_S, M_S]$ and predicted $[L_{Sp}, M_{Sp}]$, and $[L_{Fp}, M_{Fp}]$. The adversarial loss attempts to compensate for the amplitude and phase mismatch between sources in synthetic and field data. Without training the discriminator, the generator produces low-frequency data that looks like synthetics even when given an input pair from the field dataset.

FIELD AND SYNTHETIC DATA

In this study, we generate a synthetic dataset that pivots on field data and then applies the dual-band learning approach to transfer knowledge from training on synthetics to a field data application.

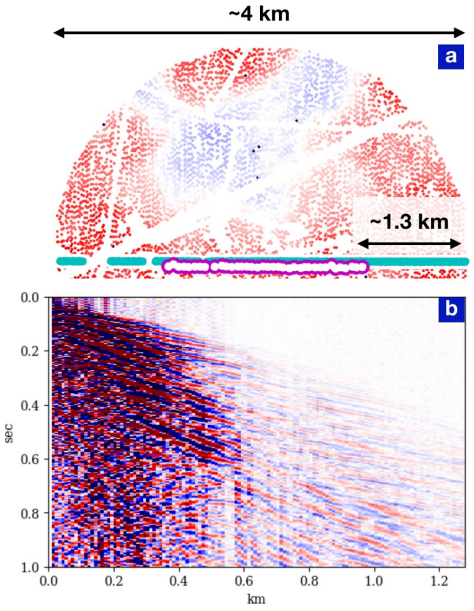

Figure 3: Portion of a land seismic survey with buried receivers and surface vibrators (a), where every dot represents ashot location colored according to its elevation. The selected line of buried receivers (magenta) and sources (green) is highlighted. A limited-offset common-receiver gather (b) centeredin one of the receivers on the line.

The field dataset from the desert environment contains data extracted from a 3D land survey performed with buried receivers and surface vibrators (Jervis et al., 2018; Smith et al., 2018). Data are recorded with 2 ms sampling by a limited number of receivers buried at depths of 50-80 m. Vibrator sources are placed on a 10 by 10 m grid (Figure 3a). Unlike in marine streamer surveys, the shot grid is distorted by surface obstructions, and data have variable maximum

offset. Note, that deep learning applications usually require data of constant size. To meet this requirement, we extract the limited-offset data measuring 128 shots in the offset direction (Figure 3b). We also apply a low-pass filter to the data below 15 Hz and limit the duration of recording to 1 s. This turns out to be sufficient for a proof-of-concept study in the near-surface setup.

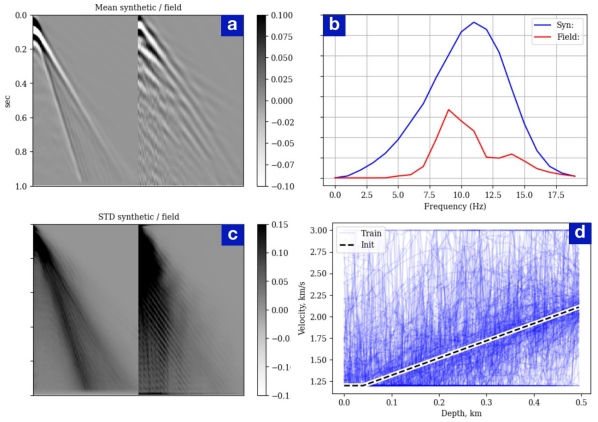

Figure 4: Comparison of the mean time-domain records (a), mean frequency power spectrum (b), and standard deviation (c) between synthetic and field datasets. Vertical profiles froma set of generated random models (d).

The synthetic dataset should resemble the field dataset as much as possible. Ideally, it would be modeled with the source wavelet extracted from field data. We assume that the source imprint is unknown, and thus we create the synthetic dataset using a generic bandlimited spike source function. The synthetic dataset construction starts from putting together an assembly of random subsurface models. We build 256 models measuring [150 x 500] m with 5 m spacing. The random initializations are sampled around the 1D velocity trend matching approximately the direct arrivals in field gathers (Figure 4d). The mean of generated models should follow the selected 1D trend while delivering a broad variance around it. We also limit the velocity range of generated models within realistic box conditions. Finally, we simulate the elastic wavefield in each of these models using bandlimited spikes with a corner frequency of 10 Hz as a source. The survey design for generation of synthetic data is unimportant as long as the source-receiver configuration matches data in the field dataset. We placed three sources in a streamer-like setup, recording the wavefield by 128 trailing receivers spaced by 10 m. This approximates a single-side common-receiver-gather from the land data. The elastic forward modeling is powered by a 2D time-domain isotropic elastic finite-difference solver (Köhn, 2011). The mean and standard deviation of both synthetic and field datasets are shown in Figure 4a and Figure 4c. The power spectrum (Figure 4b) shows boosted amplitudes of synthetic data compared to field data. The reason is that we intentionally do not use an accurate source wavelet and want to explore whether dual-band learning can balance these at the inference stage.

DEEP LEARNING FRAMEWORK

Inputs and targets. We select the field common-receiver gathers from 729 receiver locations, which leads to the volume of field dataset measuring [729, 128, 500]. The synthetic dataset contains 768 shot gathers, after modeling of 3 sources in each of 256 random models and, thus, measures [768, 128, 500]. We split each dataset into high, 5 Hz < $H_{S,F}$ < 15 Hz, mid, 5 Hz < $M_{S,F}$ < 10 Hz, and low, $L_{S,F}$ < 5 Hz, frequency data partitions (Figure 5). Finally, we downsample the data along time dimension by a factor of four and pad it with three zeros to reach the square dimensions. This leads to the final shape of the training and testing data of [768, 128, 128] and [729, 128, 128], respectively.

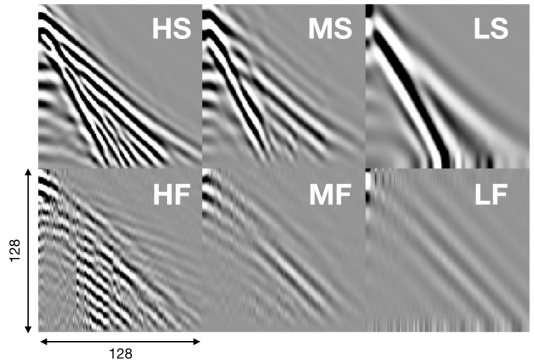

Figure 5: Comparison of synthetic (top) and field data (bottom) samples. The columns stand for high-, mid- and low-frequency data partitions.

Pre-processing. Amplitudes of high- and mid-frequency data are generally an order of magnitude larger than those of low-frequency data. This is not the case when approaching bandwidth extrapolation as a direct mapping of high to low frequencies. When jointly predicting mid and low frequencies, the latter might be neglected compared to large amplitudes in the mid-range. We follow a rather simple approach to balance contributions of low and mid frequencies in the output. First, we divide every shot gather by the maximum of its absolute value. This maps data values to the range [-1, 1]. Then, we divide the low-frequency data by the maximum of its absolute value and multiply it by the one of the mid-range data. This way, the two components of the network outputs contribute to the training more evenly.

Architecture. The GAN architecture consists of a generator, G, and a discriminator, D. The generator is the UNet (Ronneberger et al., 2015), with [32, 64, 128, 256, 512], 3 x 3 kernels in each convolutional layer of the encoder. The decoder branch is symmetric to the encoder. The patch-discriminator (Isola et al., 2017) is built as a stack of 4 convolutional layers with the following combinations of number of kernels per layer, kernel size, padding, and stride: [16, 11, 5, 4], [32, 5, 2, 2], [64, 3, 1, 2], [1, 3, 0, 1]. The output is a 6 x 6 matrix of fidelity estimates for respective partitions of the output data. We then take a mean estimate and use it for training. For the training of the discriminator D we optimize the LSGAN loss

(Mao et al., 2017), which aims to produce values > 1 for samples drawn from the "true" distribution, and values < −1 for those drawn from the "fake" distribution.

EXAMPLE

We evaluate the low-frequency extrapolation capability of neural networks in the following configurations denoted according to the corresponding input data composition:

1. H_S: training on synthetic data, direct mapping $[H_S] \rightarrow [L_S]$ by UNet.
2. H_S+M_S: training on synthetic data using mid-band mapping $[H_S, M_S] \rightarrow [L_S, M_S]$ by UNet.
3. H_S+M_S+M_F: training on synthetic and field data, dual-band learning with $[H_S, M_S] \rightarrow [L_S, M_S]$ and $[H_F, M_F] \rightarrow [L_F, M_F]$ by GAN.

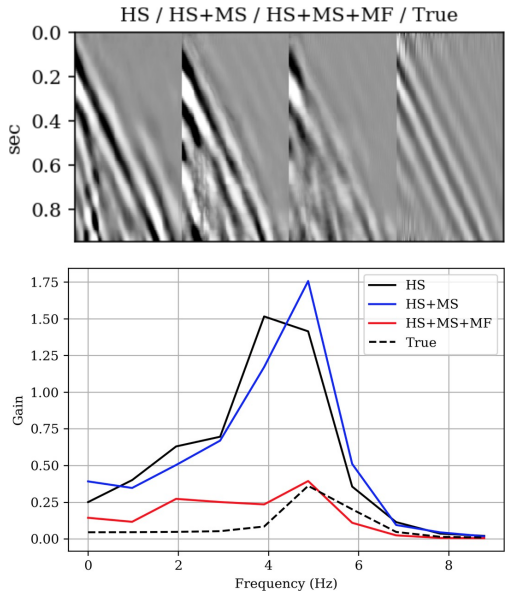

Figure 6: Predicted low-frequency data (< 5 Hz), produced from different input data compositions: H_S (high-frequency synthetic only), H_S+M_S (added medium-frequency synthetic), H_S+M_S+M_F (added medium-frequency field data and GAN). Observed data is on the right, the corresponding power spectrum is shown at the bottom.

Assuming unknown source wavelets from the field survey, we created the dataset of synthetic data using a bandlimited spike as a source. This, expectedly causes boosted amplitudes of the generated waveforms compared to field observations. Because of that, when UNet is trained on synthetic data exclusively and applied to field data, the predicted low-frequency data (Figure 6, H_S) shares the amplitude and overall appearance of samples from the synthetic dataset. In particular, the footprints of surface waves remain clearly visible. Alternatively, when adding the mid-range data as a second channel of the input data and running the first stage of dual-band approach training, the imprint of ground-roll in the prediction becomes less prominent (Figure 6, H_S+M_S). The amplitude of predicted data remains overestimated compared to field observations (Figure 6, True). Finally, we launch the second stage of dual-band training end enable the discriminator to evaluate the fidelity of predicted combinations of low- and mid-frequency data. Extrapolated low-frequency data in this case (Figure 6, H_S+M_S+MF) perceptually resemble the target.

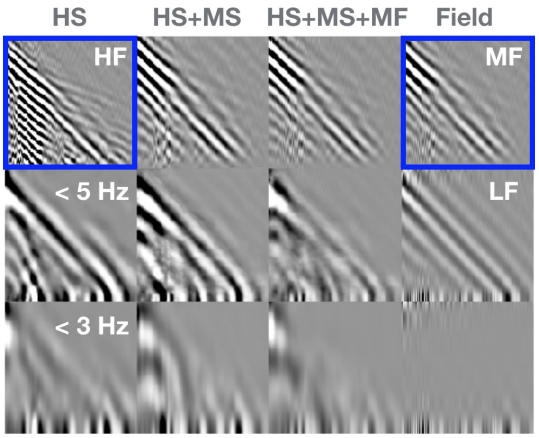

Figure 7: Reconstructed mid-frequency data (first row) and predicted low-frequency data (second row) produced from different input data compositions, same as in Figure 6. Predictions low-passed below 3 Hz (third row). Blue boxes highlight the high- and mid-frequency field data used as input to the network.

The power spectrum in Figure 6 shows the amplitude match between predicted and observed data. There, the non-zero values below the cut-off frequency of 5 Hz imply the presence of generated low-frequency data. This might be visually confirmed by exploring the outcomes of low-pass filtering (< 3 Hz) of the predicted data (Figure 7).

Finally, the additional mid-range band together with the GAN architecture guides the extrapolation towards a more realistic prediction. This proof-of-concept experiment shows promise in the dual-band approach for low-frequency extrapolation.

CONCLUSIONS

We introduce the concept of dual-band generative learning designed to facilitate the knowledge transfer from training on synthetic seismic data to applications on field data. When applied in the framework of low-frequency extrapolation, it allows to introduce a domain-specific imprint of land seismic data into training on synthetic data. We show that using the additional frequency band in both inputs and outputs of the network leads to improved predictions, while the amplitude mismatch remains prominent. The GAN architecture addresses this issue and balances amplitudes, ensuring that predicted low-frequency data are consistent with the available mid-frequency range. This dual-band generative learning concept shows considerable promise and might be studied in further seismic applications, such as velocity model building.

ACKNOWLEDGMENTS

The research reported in this publication was supported by funding from Saudi Aramco and King Abdullah University of Science and Technology (KAUST).

REFERENCES

Aharchaou, M., A. Baumstein, T. Vdovina, R. Lu, and E. Neumann, 2020, Deep learning of bandwidth extension from seabed Seismic: 82nd EAGE Annual Conference & Exhibition, Expanded Abstracts, 1–5, doi: https://doi.org/10.3997/2214-4609.202011111.

Alali, A., V. Kazei, B. Sun, R. Smith, P. Nivlet, A. Bakulin, and T. Alkalifah, 2020, Cross-equalization of time-lapse seismic data using recurrent neural networks: SEG Technical Program, Expanded Abstracts, 1506–1510, doi: https://doi.org/10.1190/segam2020-3424773.1.

Baeten, G., J. W. de Maag, R.-E. Plessix, R. Klaassen, T. Qureshi, M. Kleemeyer, F. ten Kroode, and Z. Rujie, 2013, The use of low frequencies in a full-waveform inversion and impedance inversion land seismic case study: Geophysical Prospecting, **61**, 701–711, doi: https://doi.org/10.1111/1365-2478.12010.

Bakulin, A., M. Dmitriev, and I. Silvestrov, 2018, Enhancement of challenging prestack land data for improved processing and imaging: Presented at the SPE Kingdom of Saudi Arabia Annual Technical Symposium and Exhibition, SPE, doi: https://doi.org/10.2118/192313-MS.

Fabien-Ouellet, G., 2020, Low-frequency generation and denoising with recursive convolutional neural networks: SEG Technical Program, Expanded Abstracts, 870–874, doi: https://doi.org/10.1190/segam2020-3428270.1.

Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, 2014, Generative adversarial networks: arXiv preprint arXiv:1406.2661.

Hu, W., Y. Jin, X. Wu, and J. Chen, 2020, Physics-guided self-supervised learning for low-frequency data prediction in FWI: SEG Technical Program, Expanded Abstracts, 875–879, doi: https://doi.org/10.1190/segam2020-3423396.1.

Isola, P., J.-Y. Zhu, T. Zhou, and A. A. Efros, 2017, Image-to-image translation with conditional adversarial networks: CVPR.

Jervis, M., A. Bakulin, and R. Smith, 2018, Making time-lapse seismic work in a complex desert environment for CO_2 EOR monitoring—design and acquisition: The Leading Edge, **37**, 598–606, doi: https://doi.org/10.1190/tle37080598.1.

Kazei, V., O. Ovcharenko, and T. Alkhalifah, 2020a, Velocity model building by deep learning: From general synthetics to field data application: SEG Technical Program, Expanded Abstracts, 1561–1565, doi: https://doi.org/10.1190/segam2020-3428324.1.

Kazei, V., O. Ovcharenko, T. Alkhalifah, and F. Simons, 2019, Realistically textured random velocity models for deep learning applications: Presented at the 81st EAGE Conference and Exhibition.

Kazei, V., O. Ovcharenko, P. Plotnitskii, D. Peter, X. Zhang, and T. Alkhalifah, 2020, Deep learning tomography by mapping full seismic waveforms to vertical velocity profiles: 82nd EAGE Conference and Exhibition, 1–5.

Köhn, D., 2011, Time-domain 2D elastic full-waveform tomography: Ph.D. thesis, Christian-Albrechts Universitat.

Kiel, M. X., Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, 2017, Least squares generative adversarial networks: Proceedings of the IEEE international conference on computer vision, 2794–2802.

Ovcharenko, O., and S. Hou, 2020, Deep learning for seismic data reconstruction: opportunities and challenges: First EAGE Digitalization Conference and Exhibition, 1–5, doi: https://doi.org/10.3997/2214-4609.202032054.

Ovcharenko, O., V. Kazei, M. Kalita, D. Peter, and T. Alkhalifah, 2019, Deep learning for low-frequency extrapolation from multioffset seismic data: Geophysics, **84**, R989–R1001, doi: https://doi.org/10.1190/geo2018-0884.1.

Ovcharenko, O., V. Kazei, D. Peter, and T. Alkhalifah, 2017, Neural network-based low-frequency data extrapolation: Presented at the 3rd SEG FWI workshop: What are we getting.

Ovcharenko, O., V. Kazei, P. Plotnitskiy, D. Peter, I. Silvestrov, A. Bakulin, and T. Alkhalifah, 2020, Extrapolating low-frequency prestack land data with deep learning: SEG Technical Program, Expanded Abstracts, 1546–1550, doi: https://doi.org/10.1190/segam2020-3427522.1.

Plotnitskii, P., T. Alkhalifah, O. Ovcharenko, and V. Kazei, 2019, Seismic model low wavenumber extrapolation by a deep convolutional neural network: ASEG, Extended Abstracts, 1–5, doi: https://doi.org/10.1080/22020586.2019.12073206.

Ronneberger, O., P. Fischer, and T. Brox, 2015, U-Net: Convolutional networks for biomedical image segmentation: International Conference on Medical image computing and computer-assisted intervention, Springer, 234–241.

Siahkoohi, A., M. Louboutin, and F. J. Herrmann, 2019, The importance of transfer learning in seismic modeling and imaging: Geophysics, **84**, A47–A52, doi: https://doi.org/10.1190/geo2019-0056.1.

Smith, R., A. Bakulin, M. Jervis, E. Hemyari, A. Alramadhan, and K. Erickson, 2018, 4D Seismic monitoring of a CO_2- EOR demonstration project in a desert environment: Acquisition, processing and initial results: Presented at the SPE Kingdom of Saudi Arabia Annual Technical Symposium and Exhibition.

Song, C., T. Alkhalifah, and U. B. Waheed, 2021, Solving the frequency-domain acoustic VTI wave equation using physics-informed neural networks: Geophysical Journal International.

Sun, B., and T. Alkhalifah, 2020, ML-misfit: Learning a robust misfit function for full-waveform inversion using machine learning: 82nd EAGE Annual Conference & Exhibition, 1–5, doi: https://doi.org/10.3997/2214-4609.202010466.

Sun, H., and L. Demanet, 2019, Extrapolated full-waveform inversion with deep learning: arXiv preprint arXiv:1909.11536.

Sun, H., and L. Demanet, 2020, Deep learning for low-frequency extrapolation of multicomponent data in elastic full- waveform inversion: arXiv preprint arXiv:2101.00099.

Wang, M., S. Xu, and H. Zhou, 2020, Self-supervised learning for low frequency extension of seismic data: SEG Technical Program, Expanded Abstracts, 1501–1505, doi: https://doi.org/10.1190/segam2020-3427086.1.

Zwartjes, P., 2020, Near-surface velocity estimation from phase velocity-frequency panels with deep learning: 82nd EAGE Conference and Exhibition, 1–5, doi: https://doi.org/10.3997/2214-4609.202010253.