

Deep Learning for Seismic Data Reconstruction: Opportunities and Challenges

O. Ovcharenko^{1,2}, S. Hou^{1*}

¹ CGG; ² King Abdullah University of Science and Technology

Abstract

Natural and instrumental conditions during field seismic survey lead to noise and irregularities in acquired seismic data. In this work, we explore challenges and opportunities related to denoising and interpolation of seismic data by deep convolutional neural networks. In particular, we apply three network configurations to field data and match them with suitable applications. We show that U-Net is beneficial for denoising applications while adversarial generative networks (GAN) are superior in interpolation tasks. Enhanced interpolation capability of GANs, however, comes at cost of increased uncertainty in the results and we raise awareness about this observation. In the end, we consider the pitfalls of conventional metrics and outline the requirements for data-driven approaches to be suitable in production applications.

Introduction

Successful seismic imaging and inversion of subsurface structures heavily depend on quality of the acquired seismic data. However, the field seismic data are often corrupted by various types of random and coherent noise. Reconstructing the signal from noisy data as well as interpolation for the missing content are important steps prior to imaging. In recent years, deep convolutional neural networks (DCNNs), especially generative adversarial neural networks (GANs, Goodfellow et al., 2014) have grown in popularity in seismic data reconstruction due to their achievements in computer vision. Gadyshin et al. (2019) used a GAN to reconstruct local wavefront attributes for 3D prestack seismic data. Siahkoohi et al. (2018) used an adversarial model for interpolation of heavily under-sampled seismic data in frequency domain. Xie et al. (2018) also used GAN to address the task of land seismic data denoising. Non-adversarial approaches also demonstrate promising results in data reconstruction. Mandelli et al. (2019) applied the U-Net architecture (Ronneberger et al., 2015) to jointly denoise and interpolate individual common shot gathers. Richardson and Feller (2019) also used the U-Net with ResNet blocks on multiple adjacent gathers to jointly address the tasks of seismic data denoising and deblending.

In this paper, we applied DCNNs to seismic data reconstruction, denoising and/or interpolation, on simple synthetic data to study the potentials and challenges of DCNN approaches. We compare the performance of three DCNN architectures: U-Net (Ronneberger et al., 2015), U-Net with adversarial loss (UGAN) and GMCNN (Wang et al., 2018) and analyze the result in detail.

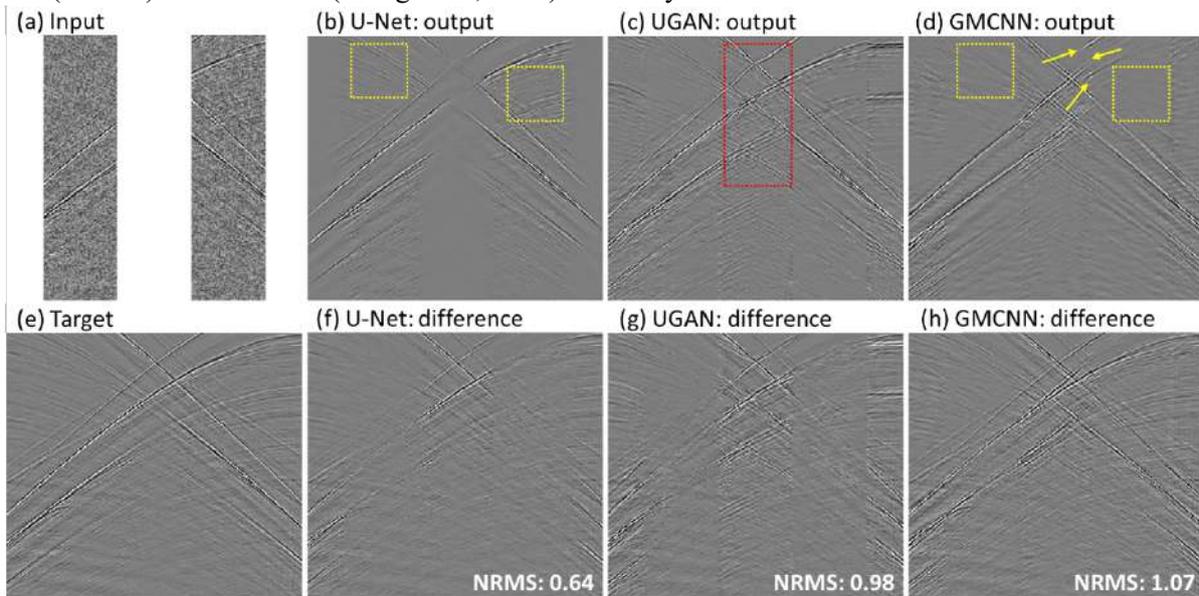


Figure 1 Reconstruction of seismic with random noise and large gaps via U-Net, UGAN and GMCNN architectures. The NRMS score measures the mismatch between entire output and target images.

Seismic data reconstruction via DCNNs

Three models of U-Net, UGAN and GMCNN are trained in the identical experimental setup in terms of the training dataset and training strategy: 5000 shot gathers are selected from field datasets and randomly cropped to images of 256 by 256 pixels (Trace interval of 6.25m and sample rate of 8ms in time). Each image is then used to produce an input and target pair for training. The original crop from seismic data becomes the target while the input data is created by adding random noise and random muting masks to the target (figure 1(a)). This way, we train the models for image-to-image mapping, where the same patch of seismic data serves both as input and target. Each model is trained for 400 epochs with L1 loss (plus adversarial loss for UGAN and GMCNN). Note that, on the inference stage, we test the extreme cases of high noise level and large gaps in order to easily identify the problems. The result shows that the trained GMCNN struggles to recover the low amplitude events (highlighted by the yellow boxes in figure 1(d)) as well as interpolating strong events with curvatures (highlighted by the

yellow arrows). This is a negative example, which teaches us that the state-of-art GAN architectures for computer vision tasks might not be optimal for seismic reconstruction when used out of the box. U-Net (figure 1(b) and (f)) achieves the best accuracy in noise removal and signal preservation (highlighted by the yellow boxes in figure (b)) whilst UGAN performs best in terms of interpolating between large gaps by visual evaluation (highlighted by the red boxes). Note that the gap in the input (figure 1(a)) is extensive and far beyond the spatial Nyquist. For such ill-posed problems, conventional interpolation algorithms add a physics-based a-priori constraint; the discriminator in the UGAN works similarly to an a-priori, but it is learnt from the data rather than derived from the physics. In summary, this example has demonstrated the ability of both U-Net and UGAN for seismic reconstruction, but it also reveals the challenges and pitfalls, with the large differences when comparing to the ground truth in figure 1(e). To illustrate the challenges, we further analyze the performance of these pre-trained models for denoising (via U-Net in figure 2) and interpolation (via UGAN in figure 4), separately.

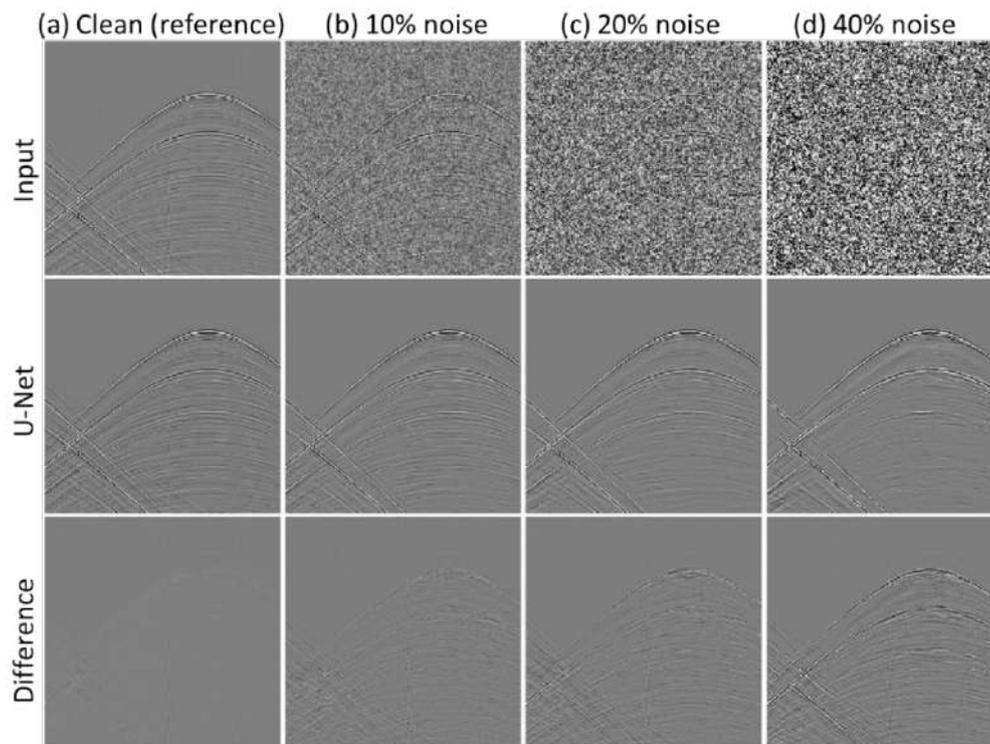


Figure 2 Seismic denoising via U-Net: (top row) Inputs with 0%, 10%, 20% and 40% of random noise; (middle row) Denoised outputs via U-Net; (bottom row) Residuals between denoised and reference data.

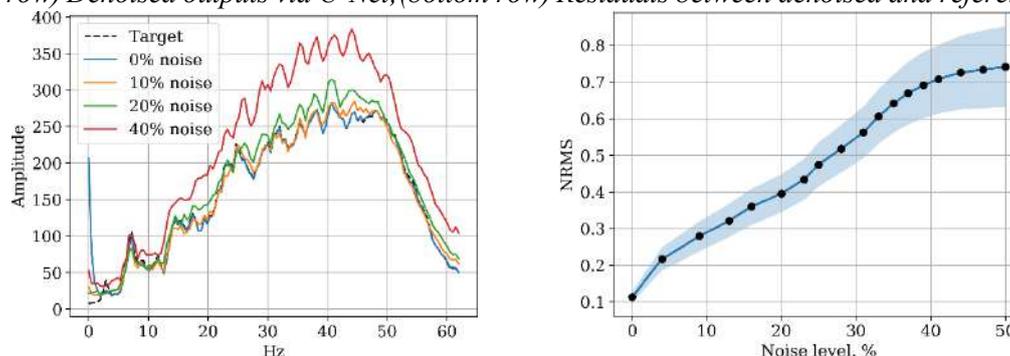


Figure 3 Seismic denoising via U-Net: (a) Left: amplitude spectrum of figure 2; (b) Right: NRMS of residuals increasing with the noise level. (Note that noise level of 50% corresponds to SNR of 1.)

Discussion: Challenges for using DCNNs for seismic data reconstruction

Signal fidelity. Signal fidelity is critical for modern seismic processing. Like most denoising algorithms, the signal fidelity of U-Net model decreases in line with the signal-to-noise ratio (SNR) of the input data (figures 2 and 3(b)). While succeeding in recovery of strong signal, U-Net cannot fully

recover the weak events. We attribute this failure to the use of the metrics of success, e.g. L1 and L2 loss functions, which are biased to strong events and the dominant frequency. For the same reason, U-Net also produces large errors at low frequencies (see figure 3(a)). On the other hand, U-Net also overestimates strong events when noise levels are high, by reconstructing wrong signal energy from the random noise. The comparison of frequency spectra in figure 3(a) clearly reveals this correlation between the amplitude enhancement and SNR. It appears critical to pay close attention to this misreconstruction of random noise, as preservation of amplitudes is a key component of seismic processing. Further studies on the metrics of success will be necessary before implementing DCNNs in real production of seismic denoising.

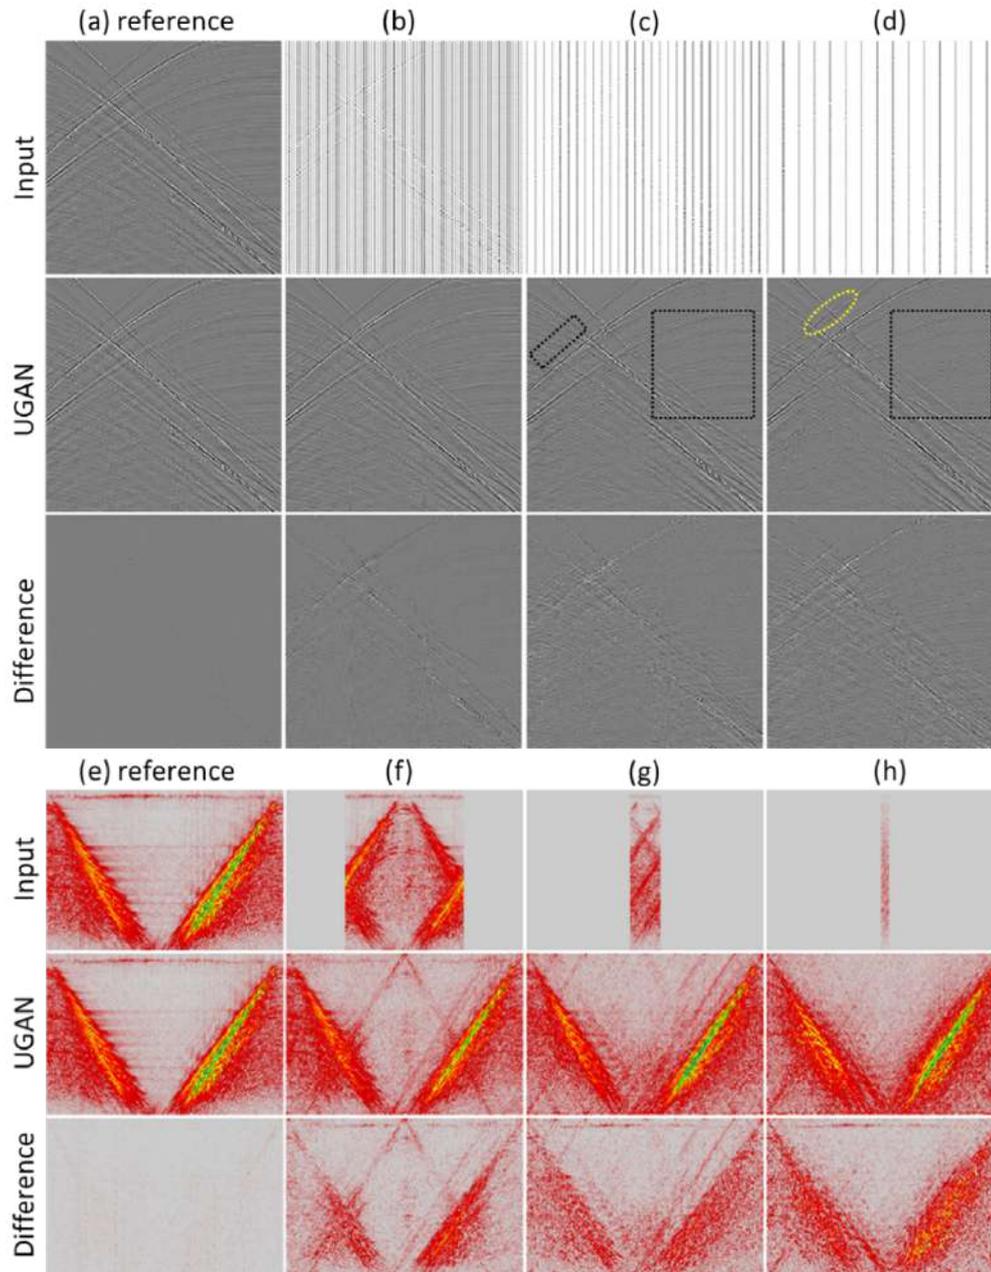


Figure 4 UGAN for seismic reconstruction using: (a) all traces; (b) 1/2 traces; (c) 1/8 traces; (d) 1/16 traces; and their f - k spectrum (e-h).

Pitfalls in GAN. DCNNs with adversarial loss, such as UGAN in figure 1(c), produce better results of seismic interpolation than U-net. We further examine the accuracy of UGAN for interpolating seismic data from all, 1/2, 1/8 and 1/16 of the total traces (figure 4), respectively. At first glance, UGAN produces visually reasonable results in reconstructing the spatially-aliased seismic events. Also, as expected, the NRMS increases as the trace interval becomes larger. The f - k spectrum in figure 4(f-g)

further demonstrates its ability of reconstructing heavily-aliased signals. However, a closer examination reveals two major pitfalls. First, as the input sampling rate goes sparser, the weak events are less preserved. This is a result of the use of L1 loss functions which is biased to the strong events, as discussed in previous paragraph (highlighted by the black boxes in figures 4(c) and (d)). Even more problematic, UGAN tends to reconstruct “fake” events, which are difficult to identify visually (highlighted by the yellow circle in figure (d)). In other words, the model is trained to “hide” the error from visual inspection. This is a consequence of using the adversarial loss, which describes the perceptual realism of the output data as a whole and has to be jointly optimized with a pixel-wise metric to produce a target-constrained result. The generative strength of adversarial training also leads to significant uncertainty in the final seismic images and makes their interpretation more risky. This contradicts the ultimate goal of using seismic to de-risk oil and gas production. It teaches a lesson that perceptual fidelity of the generated seismic data does not guaranty the realism of underlying physics. Further work needs to be done, possibly by incorporating a-priori knowledge on physical phenomena into training of data-driven models.

Conclusions

We studied the possibility of the application of DCNN on the seismic data reconstruction problem. We then discussed advantages and pitfalls of conventional metrics and demonstrate that adversarial models are prone to producing artificial events. To sum up, prior to production applications, denoising and interpolation results produced by DCNNs should meet high standards set by conventional data processing workflow. In particular, a suitable data-driven solution should ensure preservation of the original data quality while addressing the target task. Thus, further algorithmic developments and analysis is necessary to improve the signal fidelity preservation and to reduce the uncertainty in the results delivered by deep learning models.

Acknowledgements

The authors thank CGG for permission to publish. Special thanks to Henning Hoerber, Alex Clowes, Igor Mikhalev, Ewa Kaszycka, Gordon Poole, Sharon Howe, Jeremie Messud and our colleagues in CGG processing and imaging for the discussions and suggestions.

References

- Gadylyshin, K., Silvestrov, I., & Bakulin, A. [2019] Inpainting of local wavefront attributes using artificial intelligence. *89th SEG Conference and Exhibition*, Extended abstract.
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C. and Bengio, Y. [2014] Generative Adversarial Nets. *NIPS*.
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Reynolds, J., Melnikov, A., Lunova, N. and Reblitz-Richardson, O. [2019] PyTorch Captum. <https://github.com/pytorch/captum>.
- Li, X. R., Mitsakos, N., Lu, P., Xiao, Y., Zhan, C., & Zhao, X. [2019]. Generative Inpainting Network Applications on Seismic Image Compression and Non-Uniform Sampling. *Workshop on Solving Inverse Problems with Deep Networks*, NeurIPS
- Mandelli, S., Lipari, V., Bestagini, P. and Tubaro, S. [2019] Interpolation and Denoising of Seismic Data using Convolutional Neural Networks. *arXiv, abs/1901.07927*.
- Richardson, A. and Feller, C. [2019] Seismic data denoising and deblending using deep learning. *arXiv, abs/1907.01497*.
- Ronneberger, O., Fischer, P. and Brox, T. [2015] U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv, abs/1505.04597*.
- Siahkoobi, A., Kumar, R., & Herrmann, F. [2018] Seismic data reconstruction with generative adversarial networks. *80th EAGE Conference and Exhibition*, Extended abstract.
- Wang, Y., Tao, X., Qi, X., Shen, X. and Jia, J. [2018] Image Inpainting via Generative Multi-column Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, 331-340.
- Xie, P., Boelle, J. L., & Puntous, H. [2018] Generative adversarial network based fast noise removal on land seismic data. *88th SEG Conference and Exhibition*, Extended abstract.