Convolutional Neural Networks as a Tool for Seismic Data Interpolation

Maurílio Ferreira Salgado*[1,2], Carlos Alberto Moreno Chaves[2] and Roberto Hirata Júnior[2], [1] Petrobras, [2] USP

Abstract

Convolutional Neural Networks (CNNs) algorithms have been very successful in solving some image problems, such as some classification and object detection ones. Applying CNNs on seismic images is straight forward and academic research on this field is gaining momentum. Uses range from facies classification, semantic segmentation, and automatic faults and horizons interpretation to the improvement of seismic resolution, noise filtering, data reconstruction and others. In this work, we aim to solve a common seismic acquisition problem, which is to image the subsurface beneath an obstruction. Undershooting procedures and processing techniques are normally tried and usually can achieve good results in mitigating this problem, but on near offsets images are generally blank. We use GANs (Generative Adversarial Networks) to train a CNN model that interpolates seismic data, and thus, can be used to generate data in cases of platform obstruction with high likelihood and coherence.

Introduction

Seismic imaging of producing oil and gas reservoirs faces many challenges, one of them is to overcome the lack of subsurface sampling beneath platforms. Towed streamer vessels have restrictions to navigate the production structures vicinity due to safety and maneuverability reasons. The fold of these areas can reach zero value, with no information at all.

Undershooting (Hill, 1986) is a crucial technique to mitigate the probing of an area of difficult access, as reported by Johann et al. (2006), Sano et al. (2020), among others. It consists of using two separate vessels, one for the source only (air gun i.e.), that navigates on one side of the obstruction and the other vessel carries a reduced number of receptor cables (or cable) that navigates on the other side, getting closer to the platform. This arrangement is a very efficient technique, because it places the midpoint between source and receptor on the gap area recovering much data, albeit, for the near offsets, most information is completely lost.

In addition to this acquisition solution, processing procedures further mitigate the problem by using regularization, migration, seismic merging and matching, but as pointed by Nath and Vershuur (2020): “If we use conventional seismic imaging algorithms that rely on primary reflection data only, migration artefacts become unavoidable”. They propose to use the surface-related multiples to get around cases like this one, with large gaps of information.

From the current level of technology development, it is clear that no complete solution exists and non-conventional techniques must address these near offsets gaps.

Our approach uses recent CNNs architectures that successfully solve image problems and apply them to seismic data interpolation. Mandelli et al. (2019), for example, have performed great work using the U-net (Ronneberger et al., 2015) architecture, but only on synthetic data and have not addressed a large coherent gap, while Oliveira et al. (2018) have used the pix2pix (Isola et al., 2017) network, applying in seismic patches of 80x80 samples, further reshaped to 256x256, increasing the possible number of training images, but introducing an up-sampled training image. They showed that the larger the size of the gap, the lower the correlation of the interpolated image with the real image.

This work uses the ShiftNet (Yan et al. 2018), trained on a GAN scheme (Goodfellow et al., 2014). This network outperforms both the plain vanilla U-net, used in Mandelli et al. (2019), and the pix2pix network, used in Oliveira et al. (2018), in the inpainting task on the Paris StreetView dataset (Doersh et al. 2012). The network is trained and tested on real data from Brazil’s offshore, kindly handed over by Petrobras.

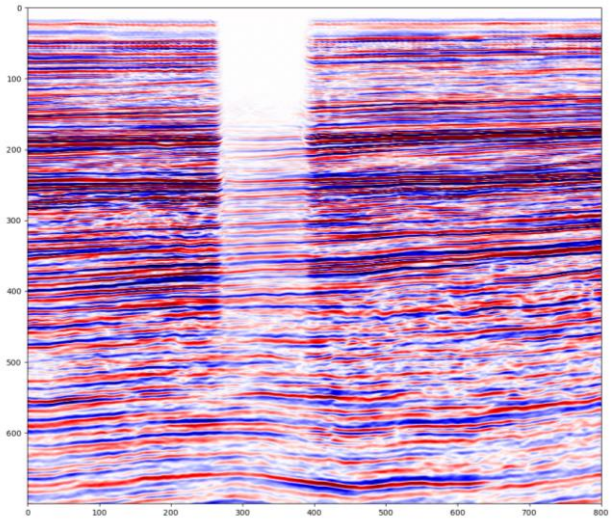

***Figure 1**- Migrated and stacked seismic session showing the information gap due to platform obstruction.*

Method

Model Description

Literature on Neural Networks architectures is extensive and exponentially growing. We chose the Shiftnet model because it outperforms other benchmark models in the task of inpainting RGB images.

Shiftnet (Yan et al. 2018) is a U-net with an additional layer named shift-layer. The U-net architecture consists of an encoder-decoder scheme with skip-connections (He. et al., 2015) from the encoder part to the correspondent decoder one. We do not discuss its details in this paper, and we invite the readers to refer to Shiftnet and U-net papers. We will focus on the Shiftnet innovation, which is the shift-layer. Yan et al. (2018) were inspired by the deterministic technics classified as exemplar-based inpainting methods. As they describe: "the completion is conducted from the exterior to the interior of the missing part by searching and copying best matching patches from the known region."

Their insight was to make a similar approach, but in the convolutional feature space. They added a layer constructed by comparing the transpose convolutional feature vectors of the decoder layer on the receptive field of the missing image part (the green layer in Fig. 2) with the skip-connection features of the corresponding layer, but on the complementary receptive field (the salmon layer in Fig. 2). Then they copied the most similar feature vector of the skip-connection on the shift-layer, in the same place as the transpose connection (the light blue layer in Fig. 2).

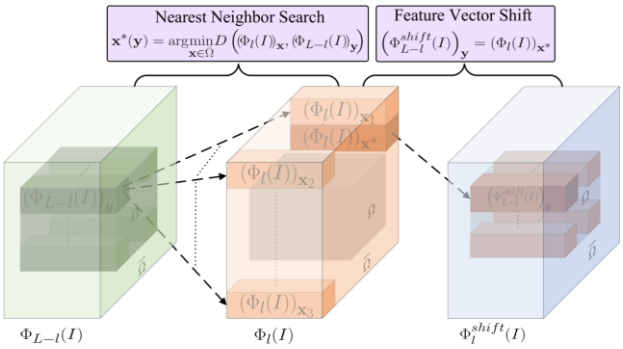

Figure 2 – *The shift-layer (light blue) is constructed with feature vectors of the skip-connection layer (light orange) that are most alike to the transpose convolutional layer feature vectors (light green). Image obtained from https://github.com/Zhaoyi-Yan/Shift-Net_pytorch.*

The name shift-layer is due to its feature vectors being the ones of the skip-connection layer shifted to the missing data slot. The layers are stacked and feed the next convolutional layer. Yan et al. (2018) experimented with the shift-layer loction and found that it was most effective (balancing time of computation and final result) in the L-3 layer, L being the last layer of the U-net.

Data description

The dataset comes from Brazil's offshore in Santos basin. It is a PSDM (post-stack depth migration) merge of two different acquisitions, with standard techniques applied to mitigate the obstruction problem, but the gap is still present (Fig. 1). The exact location of the data is omitted, due to a previous agreement with Petrobras.

Seismic data has a size of 800x1200x1500 samples, of which 800 are the inlines, 1200 are the xlines and 1500 are the depth samples. The training set was obtained from the first 700 inlines, excluding the inlines with an information gap, totalizing 563 usable inlines. From these inlines, 15000 2D images patches with 256x256 pixels were randomly picked, leaving 12000 for training and 3000 for validation. Statistics of the mean and maximum values of the training set is stored, and used later in training to normalize the images. We use the last 100 inlines to test the model, from which 1000 patches were randomly generated.

All those images are perfect, with no information gap. For training purposes, we introduce the missing samples artificially as a mask in order to have the ground truth from the uncorrupted image. At the present stage, we use only a central mask that covers 25% of the image. Figure 3 shows an example of artificial deleting.

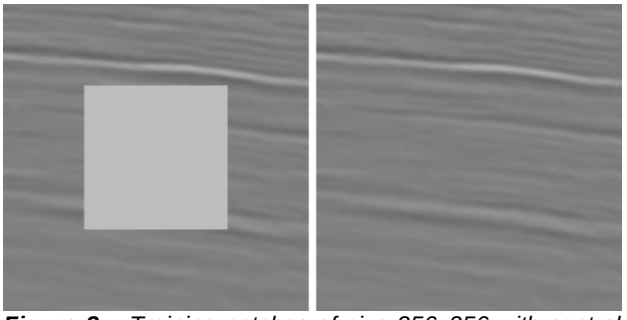

Figure 3 – *Training patches of size 256x256 with central mask on the left and with the ground truth on the right.*

Training Phase

The model was trained in an adversarial manner (GAN). Shiftnet is the generative network G and the discriminator network D is constructed as a sequential convolutional neural network. It has five convolutional layers with leaky ReLU activation. Instance normalization is also applied after convolutional layers 2, 3 and 4. The discriminator output is obtained after a sigmoid function and yields the probability of the image be a real image (ground truth).

Adversarial training is a competition between the Generative network and the Discriminative network. While D is trained to uncover if an image is authentic or fake (generated by G), G generates increasingly better images to be tested by the discriminator. G fills the image gap generating fake images with detailed near realistic seismic image. Besides the adversarial loss, an L1 norm is used to ensure that the final fake image should be an approximation of the ground truth image.

In summary, we get a 256x256 image sample and normalize it with the mean and maximum values stored from the whole training set. Also, we do a random horizontal flip with a 50% probability in order to increase data variability (data augmentation). Then, a mask is added in order to simulate an image gap. This masked image is fed to the generative model, which outputs a filled fake image. The discriminator is fed with both fake and ground truth, outputting probabilities for each one. D

is updated (with adversarial loss), G is updated (with adversarial loss and L1 norm), and the cycle is repeated.

We trained the network for 250 epochs with a batch size of one (SGD – stochastic gradient descent), and within each epoch, the model sees all training set. A mean L1 norm is computed for each epoch for both the training and validation sets and evolution can be observed in Fig. 4. We see that the validation L1 norm reached a plateau roughly after epoch 50, oscillating around 0.79, but the network was still learning. The plateau can be explained because the network does not recover the same data. Therefore, the L1 norm cannot improve, but on the training set, the model is still learning to generate realistic data. Overfitting is discarded based on validation performance.

We run the model training on GPU, a GeForce GTX 1080 ti with 6.1 computing capability. The code was written in Python, modified from the Shift-Net PyTorch repository. The total time of training was 93h.

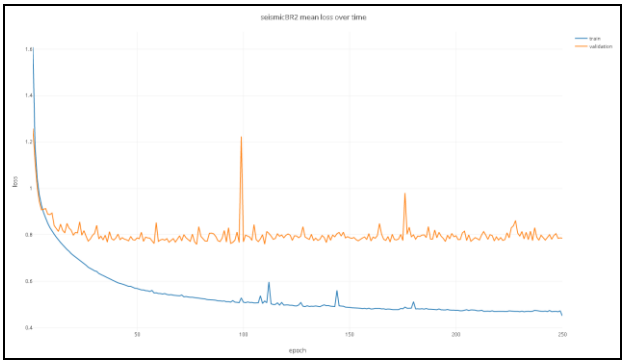

***Figure 4** – Mean L1 loss evolution for the training set (blue) and the validation set (orange).*

Results

After training, the resulting model was applied to the test set with 1000 patch images. Results are displayed for three different depths to show the effect of varying frequency content. We show three images for each selected depth: (1) the leftmost image is the true patch with the artificial mask (real_A); (2) the center image is the fake one generated by the trained Shift-Net model (fake_B); (3) rightmost image is the ground truth (real_B).

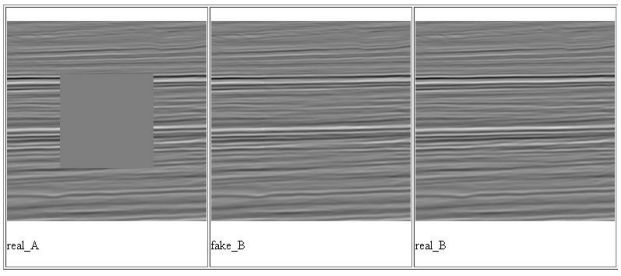

***Figure 5** – Test images sampled from 550m. On the left is the real masked image (real_A), on the center the fake generated (fake_B) and, on the right, the real uncorrupted image (real_B).*

The first sampled depth is around 550m and it is shown in Fig. 5. One can observe that most energetic reflectors are well recovered and finely detailed information is also present. If the fake image is not presented as a fake one, the interpreter could be fooled by it. Yet, with a closer look, it is possible to see a slight difference of energy in the center of the fake image, and some border artifacts, but nothing too flashy.

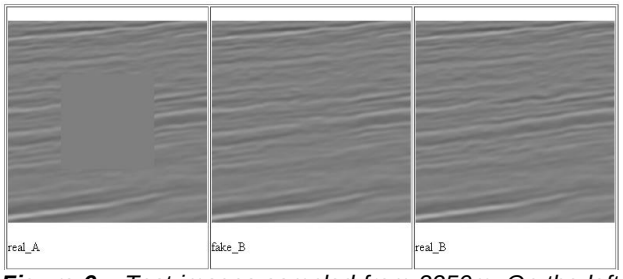

***Figure 6** – Test images sampled from 2250m. On the left is the real masked image (real_A), on the center the fake generated (fake_B) and, on the right, the real uncorrupted image (real_B).*

The second depth sample is close to 2250m and can be seen in Fig. 6. The loss of frequency content is clear and the reflectors are inclined. Performance is a little inferior as compared to the shallower counterpart, but still very realistic.

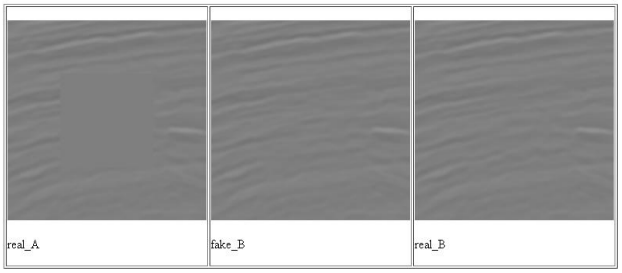

***Figure 7** - Test images sampled from 4000m. On the left is the real masked image (real_A), on the center the fake generated (fake_B) and, on the right, the real uncorrupted image (real_B).*

The third depth sample is from near 4000m with even less frequency content, and can be seen in Fig. 7. Performance is close to the depth of 2250m.

We also analyze the coherence of the generated images between adjacent inlines. This is important to reconstruct seismic data because if one inline is completely different from the next one, it is impossible to interpret a horizon. Generative models can produce images completely different from inputs slightly distinct. In our case, the Shift-layer serves as a contour condition, because it uses the information from the uncorrupted image, producing coherent images for sequential inlines. This kind of analysis was not done in previous related work and is the first to our knowledge.

Below is possible to apprize the coherence between three generated patches from sequential inlines. We show images for three different depths, the shallower is shown in Fig. 8, the middle depth is shown in Fig. 9 and the deeper one is shown in Fig. 10. The high similarity between these sequential inline makes this trained network suitable for reconstructing large gap information seismic data. We reason that this depends on the data itself, if the inline distance is too high so that images from

one inline to another are too different, this level of coherence will not be achieved.

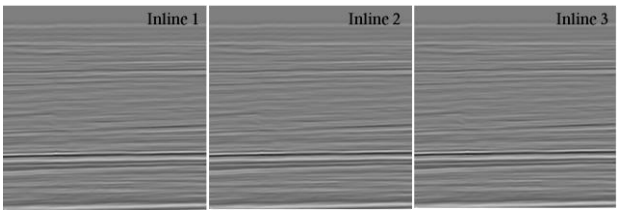

Figure 8 – *Fake (generated) images from sequential inlines for shallow depth.*

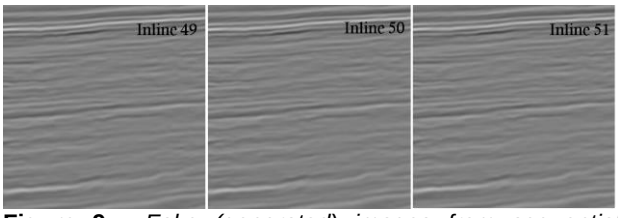

Figure 9 - *Fake (generated) images from sequential inlines for middle range depth.*

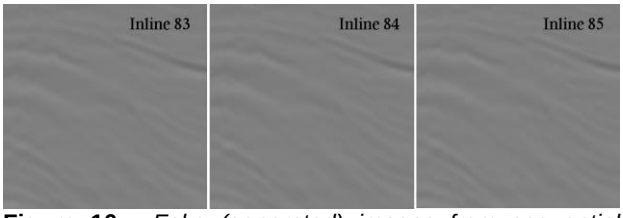

Figure 10 - *Fake (generated) images from sequential inlines for deeper depth.*

Conclusions

Generative Adversarial Networks are applied in this work to fill large gaps of information in real seismic data. The sources of these gaps vary and they are not rare in real images. The proposed method achieved good performance by filling the images with realistic and fine-detailed information using a real dataset and artificially generated gaps. We showed that this novel neural network provides a contour condition for generating coherent images across different adjacent inlines. Future work will address comparison metrics and the reconstruction of the data itself.

Acknowledgments

The authors would like to thank Petrobras for handing over the data and for permission to publish it. We thank FAPESP, project #2018/10767-0. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

References

DOERSCH, C. et al. What makes Paris look like Paris? Communications of the ACM, v. 58, n. 12, p. 103–110, 2015.

GOODFELLOW, I. J. et al. Generative adversarial networks. Advances in Neural Information Processing Systems, v. 3, n. January, p. 2672–2680, 2014.

HE, K. et al. Deep residual learning for image recognition. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, p. 770–778, 2015.

HILL, A. W. Two-ship undershooting of a production platform: The magnus high-resolution seismic project. Oceanology: Proceedings of an International Conference, p. 191–202, 1986.

ISOLA, P. et al. Image-to-image translation with conditional adversarial networks. Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, (CVPR). Honolulu, HI, USA: IEEE, 2017

JOHANN, P. R. S. et al. 4D Seismic in Brazil: Experiences in Reservoir Monitoring. Paper presented at the Offshore Technology Conference, Houston, Texas, USA, 2006.

MANDELLI, S. et al. Seismic data interpolation through convolutional autoencoder. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, 2018

NATH, A.; VERSCHUUR, D. J. Imaging with surface-related multiples to overcome large acquisition gaps. Journal of Geophysics and Engineering, p. 742–758, 2020.

OLIVEIRA, D. A. B. et al. Interpolating Seismic Data With Conditional Generative Adversarial Networks. IEEE GEOSCIENCE AND REMOTE SENSING LETTERS, v. 15, n. 12, p. 1952–1956, 2018.

RONNEBERGER, O.; FISCHER, P.; BROX, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. (N. Navab et al., Eds.)Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2015. Lecture Notes in Computer Science, vol 9351. Springer, Cham. 2015

SANO, S.; TAN, T. Q.; JO, G. High resolution 3D seismic undershooting acquisition over platforms and seismic processing challenges in a gas producing eld. Journal of the Japanese Association for Petroleum Technology, v. 85, n. 1, p. 44–53, 2020.

YAN, Z. et al. Shift-net: Image inpainting via deep feature rearrangement. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), v. 11218 LNCS, p. 3–19, 2018.