3D Conditional Generative Adversarial Networks to enable large-scale seismic image enhancement

Praneet Dutta¹, Bruce Power², Adam Halpert², Carlos Ezequiel¹, Aravind Subramanian³ Chanchal Chatterjee¹, Sindhu Hari³, Kenton Prindle¹, Vishal Vaddina³, Andrew Leach¹ Raj Domala³, Laura Bandura², Massimo Mascaro¹

¹Google AI

²Chevron Energy Technology Company

³ Quantiphi

{aravind.subramanian, sindhu.hari, vishal.vaddina, raj.domala}@quantiphi.com

Abstract

We propose GAN-based image enhancement models for frequency enhancement of 2D and 3D seismic images. Seismic imagery is used to understand and characterize the Earth's subsurface for energy exploration. Because these images often suffer from resolution limitations and noise contamination, our proposed method performs large-scale seismic volume frequency enhancement and denoising. The enhanced images reduce uncertainty and improve decisions about issues, such as optimal well placement, that often rely on low signal-to-noise ratio (SNR) seismic volumes. We explored the impact of adding lithology class information to the models, resulting in improved performance on PSNR and SSIM metrics over a baseline model with no conditional information.

1 Introduction

In geophysical imaging, resolution limitations of the seismic migration methods are wellknown [Beylkin et al. 1985; Vermeer 1999]. At depths typical of modern exploration targets, the seismic wavelength can be in excess of 250 meters, meaning that geo-scientists may be unable to resolve individual rock formations less than 60 meters thick, which is needed to be successful in exploration settings. In addition, variations in lithology and features such as faults can cause further disruption and attenuation of the acoustic wave energy. As a result, interpreting the underlying geological model from these image volumes has high uncertainty and has been the focus of much of the current and ongoing energy resource exploration. These limitations make purely data-driven approaches such as image super-resolution techniques to enhance seismic images more attractive.

Image super-resolution is the process of converting low-resolution images into high-resolution ones. The use of deep neural network approaches to image super-resolution is a growing research area with real-world applications in various fields [Wang et al. 2019]. Promising attempts at image super-resolution through deep neural networks include use of a pixel-recursive method [Dahl et al. 2017] and zero-shot learning [Shocher et al. 2017]. A survey of approaches can be found in [Wang et al. 2019]. Early attempts to use deep learning and GANs for seismic image enhancement have also shown promise [Halpert 2018].

Our proposed seismic image enhancement approach expounds on the work of [Ledig et al. 2016] that describes a super-resolution generative adversarial network (SRGAN) model, which we adapt to support both 2D slices of a 3D seismic cube and 3D cube partitions. Additionally, we employ pixel-level class conditional information based on geological lithology to improve our performance.

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

2 Image enhancement approach

We add conditional information in the form of additional image channels that determine the lithology class associated with each pixel. We considered two methods in representing lithology class information: deterministic and probabilistic. In the deterministic method, the classes are represented as a one-hot encoded vector for each corresponding image pixel with a total of thirty one possible classes. In the probabilistic method, the class channel represents the binary probability of a pixel belonging to the salt class.

To incorporate the class conditional information into SRGAN, we considered multiple fusion locations, such as early, mid and late fusion [Snoek et al. 2005] and fusion types, such as concatenation and dot product of the conditional information with ground truth (depicted in Figure 1). Early fusion meant augmenting the class conditional information at the generator/discriminator input layer. Mid fusion meant adding the information at the first residual block of the generator, and before the first residual block of the discriminator. Late fusion happened after all the repeating residual blocks of the generator, and in the same place as mid fusion for the discriminator. Since input and output images needed to have the same dimensions, we omitted an upsampling layer in the generator.

2.1 Architecture



Figure 1: SRGAN-based model architecture for enhancing seismic images. In this work, the depth of the generator and discriminator is represented as a function of the number of repeating residual blocks.



Figure 2: Geological lithology classes for conditioning GAN prediction.

2.2 Loss function

We base our loss function on the work of [Ledig et al. 2016] that defines a perceptual loss function as a weighted sum of content loss (based on MSE loss) and adversarial loss respectively.

$$l_{MSE} = \frac{1}{w \cdot h} \sum_{i=1}^{i=w} \sum_{j=1}^{j=h} (x_{i,j} - G(z|c)_{i,j}))^2$$
(1)

In Equation 1, w and h refer to image pixel width and height respectively, G is the generator function that takes as input both the noisy image z and conditional information c, and x is the ground truth image. Minimizing MSE loss maximizes peak signal-to-noise ratio (PSNR), which is a commonly used image quality estimation. The adversarial loss is based on [Mirza and Osindero 2014] that adds extra conditional information to the two-player minimax game with value function V(D,G) originally proposed in [Goodfellow et al. 2014].

$$\min_{G} \max_{D} V(D,G) = \mathbb{E}_p[log(D(x|c)) + log(1 - D(G(z|c)))]$$
(2)

where p is the joint distribution of x, y, and c. Note that this differs from standard GAN formulations in that the noise is not explicitly sampled, but is a result of the observation process of the ground truth x that yields z and c. D refers to the discriminator function, which pushes the generator to output images in the enhanced seismic image manifold. We further direct the generation process of G by conditioning both G and D with additional information c that refers to the lithology classes in our study.

3 Experiments

3.1 Seismic Dataset

More than 100 000 training images were extracted from the SEAM I dataset [Fehler and Keliher 2011]. This seismic data is a result of a finite difference forward model where a simulation of an acoustic wave field is propagated through the earth model volume computationally. The original earth model is also a synthetic model that was designed to effectively reproduce the actual lithology and structure found in the earth's subsurface in sedimentary basins where energy exploration and production occurs. As a result, the lithologies and labels of the rock properties are known down to the pixel scale. The lithology classes considered for the study are shown in Figure 2.

To generate the degraded seismic image input, we used a 5Hz low-pass filter and added 50% uniform random noise.



(a) Ground truth seismic data (full (b) Degraded input (filtered and bandwidth). (b) Degraded input (filtered and additive noise).



(c) Generated output with no conditional information.

(d) Generated output with conditional information.



3.2 Metrics

To validate the model, we considered the following objective image quality metrics: peak signalto-noise ratio (PSNR), structural similarity index (SSIM) and multi-scale SSIM (MS-SSIM). The multi-scale method (MS-SSIM) provides more flexibility in that it can incorporate image details at different resolutions.

3.3 Results for 2D

The outputs of our 2D models were visually examined by domain experts, who verified the efficacy of the approach (see Figure 3). The results were visually examined to determine if reflection amplitude, phase, and coherence were consistent with the high frequency image as well as the underlying earth model. Table 1 summarizes the best results we obtained from the different 2D image enhancement models. We achieved the best result using a model trained with probabilistic conditional information (second row of Table 1). In our experiments, conditional models appear to have performed better than Vanilla SRGAN (baseline) in most cases, regardless of fusion strategy.



(a) 3D image ground truth. (b) Degraded 3D image used as (c) Enhanced 3D image. input data.

Figure 4: Cross-sections of an image volume for ground truth (a), degraded input (b) and model output (c) respectively. The image in (a) identifies the three different transects through the cube - plane view *a*, x-y cross-section *b*, and z-y cross-section *c*.

Table 1. Summary of Dest results for 2D

Model	Generator Depth	Fusion Type	Fusion Pos	MS-SSIM	SSIM	%SSIM Gain	PSNR	% PSNR Gain
Vanilla SRGAN	16	-	-	-	0.549	-	20.47	-
Probabilistic	32	Concat	Early	0.784	0.656	19.48	22.97	12.21
Deterministic	32	Dot	Mid	0.785	0.642	16.93	22.96	12.16
Deterministic	8	Dot	Late	0.760	0.600	09.28	22.13	08.11
Probabilistic	8	Concat	Early	0.741	0.575	04.73	20.87	01.95

Table 2: Summary of results for 3D

Model	Fusion Type	Fusion Pos	SSIM	SSIM % Gain	PSNR	% PSNR Gain
Vanilla SRGAN	-	-	0.94	-	29.95	-
Deterministic	Concat	Late	0.98	4.25	33.22	10.91
Deterministic	Concat	Early	0.95	1.06	30.40	1.5

3.4 Results for 3D

Table 2 summarizes the results of the different 3D image enhancement models, with the best result obtained using a model trained with deterministic information and late fusion. Even without conditional information in the Vanilla SRGAN model, we were still able to achieve an SSIM of 0.94. Figure 4 provides a visual comparison of the 3D model outputs by showing orthogonal cross-sections through the 3D image volume. For models trained with deterministic information, late fusion with concatenation provided the best results.

4 Conclusion and Future Work

We were able to achieve improved performance on seismic image enhancement tasks using an SRGAN-based model with conditional information over one that did not have such information for 2D and 3D cases. Models trained with probabilistic information performed better than models trained on deterministic information on all metrics used in this study. In spite of this, the Vanilla SRGAN model was still able to produce enhanced seismic images that were visually similar to ground truth.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www.tensorflow.org/. Software available from tensorflow.org.
- G. Beylkin, M. Oristaglio, and D. Miller. Spatial Resolution of Migration Algorithms, pages 155–168.
 Springer US, Boston, MA, 1985. ISBN 978-1-4613-2523-9. doi: 10.1007/978-1-4613-2523-9_15.
 URL https://doi.org/10.1007/978-1-4613-2523-9_15.
- Ryan Dahl, Mohammad Norouzi, and Jonathon Shlens. Pixel recursive super resolution. *CoRR*, abs/1702.00783, 2017. URL http://arxiv.org/abs/1702.00783.
- Hao Dong, Akara Supratak, Luo Mai, Fangde Liu, Axel Oehmichen, Simiao Yu, and Yike Guo. TensorLayer: A Versatile Library for Efficient Deep Learning Development. *ACM Multimedia*, 2017. URL http://tensorlayer.org.
- Michael Fehler and P. Joseph Keliher. SEAM Phase 1: Challenges of Subsalt Imaging in Tertiary Basins, with Emphasis on Deepwater Gulf of Mexico. Society of Exploration Geophysicists, 2011. doi: 10.1190/1.9781560802945. URL https://library.seg.org/doi/abs/10.1190/ 1.9781560802945.
- Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Elliot Karro, and D. Sculley, editors. *Google Vizier: A Service for Black-Box Optimization*, 2017. URL http://www.kdd.org/kdd2017/papers/view/ google-vizier-a-service-for-black-box-optimization.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. 2014.
- Adam D. Halpert. *Deep learning-enabled seismic image enhancement*, pages 2081–2085. 2018. doi: 10.1190/segam2018-2996943.1. URL https://library.seg.org/doi/abs/10.1190/segam2018-2996943.1.
- Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. *CoRR*, abs/1609.04802, 2016. URL http://arxiv.org/abs/1609.04802.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *CoRR*, abs/1411.1784, 2014. URL http://arxiv.org/abs/1411.1784.
- Assaf Shocher, Nadav Cohen, and Michal Irani. "zero-shot" super-resolution using deep internal learning. *CoRR*, abs/1712.06087, 2017. URL http://arxiv.org/abs/1712.06087.
- Cees G. M. Snoek, Marcel Worring, and Arnold W. M. Smeulders. Early versus late fusion in semantic video analysis. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, MULTIMEDIA '05, pages 399–402, New York, NY, USA, 2005. ACM. ISBN 1-59593-044-2. doi: 10.1145/1101149.1101236. URL http://doi.acm.org/10.1145/1101149.1101236.
- Gijs J. O. Vermeer. Factors affecting spatial resolution. *GEOPHYSICS*, 64(3):942–953, 1999. doi: 10.1190/1.1444602. URL https://doi.org/10.1190/1.1444602.
- Zhihao Wang, Jian Chen, and Steven C. H. Hoi. Deep learning for image super-resolution: A survey. *CoRR*, abs/1902.06068, 2019. URL http://arxiv.org/abs/1902.06068.

5 Appendix

5.1 Performance Comparisons

Some of our performance improvements result from training and tuning the models on Google Cloud AI Platform. We utilized the TF-GAN library for all of our experiments. A summary of performance improvement results are shown in Table 3.

	Metric	GPU	TPU
Performance	Evaluation Set - MSE	0.022	0.00494
Batch Size	Batch Size	1	16
Processing HW	Processing Hardware	7 Tesla P100s	Basic TPU
Speed	Samples/s	4 samples/sec	11 samples/sec
Cost	ML Units per 1k samples	1.9	0.7

Table 3: Performance C	Comparison on	GPUs vs TPUs
------------------------	---------------	--------------

Table 3 summarizes our results on the box noise dataset with a model trained on deterministic information. This was run for our baseline model with 2 repeating residual blocks in the generator and 2 blocks in the discriminator. The conditional information was added in via concatenation.

Table 4 details results for a batch prediction run for 5000 3D Cubes.

Table 4: Speed up comparison using GPU's for Batch Prediction

CPU Time(minutes)	GPU Time(minutes)	GPU Speedup
208	15	13.8X

Reproducibility: The variance in SSIM is 0.05 on evaluation as observed after multiple runs at the end of full training.

5.2 Additional information

Table 5 elaborates more on the information on hyperparameters tuned. We list a more specific range of parameters used for both 2D and 3D datasets.

Parameter	Min Value	Max Val	Scale
2D Generator Depth	16	40	Linear
2D Discriminator Depth	6	18	Linear
2D Batch Size	6	12	Integer
3D Generator Depth	1	7	Linear
3D Discriminator Depth	1	6	Linear
3D Batch Size	1	4	Integer

Table 5: Specific range of values used for hyperparameter tuning in 2D and 3D

Figure 5 offers a visual comparison of outputs of the 2D models we trained with the low-resolution input images and ground truth.

5.3 Experiment setup

We used Google Cloud AI Platform to run model training and hyperparameter tuning in the cloud, enabling us to leverage Tesla P100 GPUs and Google TPUs for training and Tesla P100 GPUs for inference. We used the TensorFlow GAN Estimator framework [Abadi et al. 2015] to implement the model.



Figure 5: Sample output of 2D models: Vanilla SRGAN (b), probabilistic (c) and deterministic (d) conditional SRGANs compared to the degraded input data (a) and ground truth (e). Both conditional SRGAN outputs show a marked decrease in coherent noise artifacts (diagonal vertical striping) compared to the ground truth image.

The models were trained for 100 000 steps with a batch size (hyperparameter) of 8 examples. We used AI Platform Hyperparameter Tuning, which is based on Google Vizier [Golovin et al. 2017], to optimize model performance.

5.4 Open Source Implementation

An open source implementation of the work of [Ledig et al. 2016] is publicly available [Dong et al. 2017]. This provides a similar base implementation of what we used for our Vanilla SRGAN model.