GAN-based priors for Bayesian inference of subsurface geology at large scale

Kevin B. Daly, Tuan A. Tran, Brent D. Wheelock, Grant J. Seastream
ExxonMobil Technology and Engineering Company

SUMMARY

In principle, Bayesian inference is a robust way to simultaneously incorporate prior knowledge and quantify uncertainty during a physics-based inversion like seismic inversion. In practice, Bayesian methods can be computationally costly, requiring many evaluations of the forward model and samples from the prior, the latter becoming significant when samples are large three-dimensional (3D) images (>1e7 voxels) containing complex sub-seismic geology. In our Bayesian approach to seismic inversion, we mitigate this computational cost by encapsulating prior knowledge within a Generative Adversarial Network (GAN; Goodfellow et al., 2014), a generative machine-learning (ML) model with fast sampling-time and high-quality output. We trained a GAN on realistic process stratigraphy data and implemented an architectural change that improves the quality of very large outputs. We also demonstrated a training strategy for decreasing the size of the GAN's latent space, which we expect will reduce the difficulty of the inverse problem by reducing its dimensionality. Finally, we augmented the stratigraphic training set, and therefore the GAN, with other types of prior knowledge, specifically sand-body orientation, which helps further constrain the output of seismic inversion.

INTRODUCTION

Relevant prior knowledge for seismic inversion includes plausible distributions of rock properties like Vclay and porosity, which convey their spatial organization at multiple length scales, especially scales <10 m in height that are poorly constrained by seismic data alone. One source of this prior knowledge is process stratigraphy models that simulate the deposition of sediment in deepwater environments (Wahab et al., 2022). Over time, sediment in certain regions accumulates into large bodies of sand and shale, known as submarine fans. While submarine fans generated by these process models are physically plausible, they do not by themselves precisely match a specific seismic dataset. However, they are a rich source of small-scale geologic information that can be learned by generative ML models. Such models can then interpolate between, and in some cases stitch together, the learned geologic patterns producing a much wider variety of plausible subsurface geologies, improving the chances of finding some that match observations like seismic data.

As part of training, generative ML models fit a highly parameterized, flexible probability distribution to the training data (Goodfellow, 2016). After training, the resulting model can efficiently generate new samples from that distribution. Provided their sampling time is fast enough, these models may function as prior probability distributions in Bayesian inference. In some generative models, sampling proceeds by first sampling a vector in a latent space, which is typically distributed as a standard multivariate normal. Then, that vector is mapped to pixel space via a trained neural network, generating the output image (or volume) representing the physical domain where observations take place. Often, the latent space has a lower dimension than the pixel space, a benefit for inversion applications (Fernández-Martínez & Fernández-Muñiz, 2020). For example, there are inversion frameworks where Bayesian inference happens entirely within the low-dimensional latent space, and the generative ML prior performs the task of mapping posterior samples from the latent space back to the pixel space (Tewari et al., 2022).

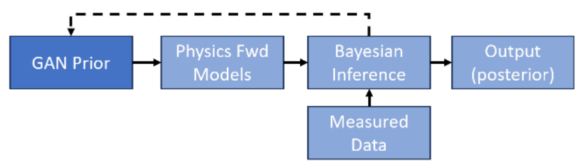

Figure 1: Bayesian inference workflow with a GAN-based prior.

The objective of the present study was to identify and train a generative ML model that meets the requirements of fast sampling-time and a highly compressed latent space intended for Bayesian inference (Fig. 1), specifically Bayesian seismic inversion at commercially relevant scales (>1e7 voxels). To that end, we have investigated different ML architectures, training strategies, and data augmentation approaches, some of which are described here.

METHODOLOGY

We source training data from completed process stratigraphy simulations (Wahab et al., 2022), ideally run under flow conditions relevant to the ultimate inversion use case. On one hand, the quantity of available simulation images (limited by computational cost), if treated as individual data points, would be insufficient to train a generative ML model. On the other hand, for seismic inversion, the generative ML model primarily needs to learn geologic patterns near or below a vertical length scale of 10 m. So, each simulation can be mined for many small overlapping training patches

(Fig. 2). The patch size is an important hyperparameter that should be selected based on 1) whether it produces sufficiently numerous/diverse patches from the simulation and 2) the length scale of geologic patterns that need to be learned. If the prior is trained on patches, the question remains of how to generate much larger posterior images for the workflow in Fig. 1.

One solution is to select a generative ML model with a fully convolutional architecture, where the latent coordinate is now a tensor having the same rank as the output images, and the size of the output can be increased by adding more "voxels" to the latent coordinate. This architecture has been demonstrated in GAN's (Jetchev et al., 2016) and applied to geologic images (Laloy et al., 2018; Song et al., 2022; Zheng & Zhang, 2022). A useful feature of this architecture is that by varying the training-patch size along with other hyperparameters of the network, the "projective field" (Jetchev et al., 2016), or neighborhood of influence of a latent voxel, can be varied, providing control over the length-scale of correlations in the output image. In the present work, we adopt this fully convolutional GAN approach; later we describe an architectural change, specifically to the convolutional padding, that improves the quality of large images generated from a GAN that was trained on small patches.

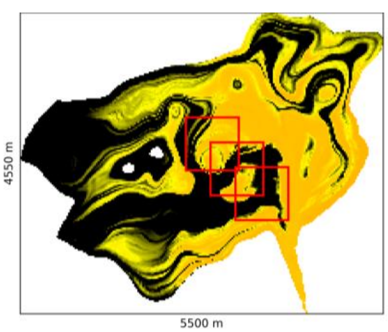

Figure 2: A map view of a 3D image generated by a process stratigraphy model, with feeder channel shown in the lower right. Color indicates average Vclay (black: shale; yellow: sand). Red squares are example patches extracted for training the GAN.

For our seismic inversion approach, the output images at a minimum contain an attribute for Vclay; porosity could be populated by a deterministic correlation, a separate distribution or learned from the training set. Other attributes could be added that incorporate other types of prior information; we experimented with adding two attributes (a sine and cosine) that encode the azimuthal angle formed between a vector parallel to the feeder channel and another vector connecting the channel mouth to the voxel location. It is observed in our simulations that channelized sand bodies preferentially align in this direction. By adding these attributes, we give the inversion the option to control regional alignment of sand bodies in its results, while preserving the local freedom of sand bodies to deviate from this direction, as seen in the process stratigraphy simulations.

RESULTS & DISCUSSION

GAN Image Quality, Statistics, and Sampling Time

We trained a GAN on 32 m × 800 m × 800 m overlapping patches extracted from process stratigraphy simulations like the one shown in Fig. 2. For this training set, the chosen patch size turned out to be a good compromise between quantity of training patches and continuity of GAN output images, while, importantly, containing the <10 m features needed for seismic inversion. Based on visual comparison of randomly selected training patches (Fig. 3a) and GAN outputs (Fig. 3b-d), the quality was deemed sufficient for our applications.

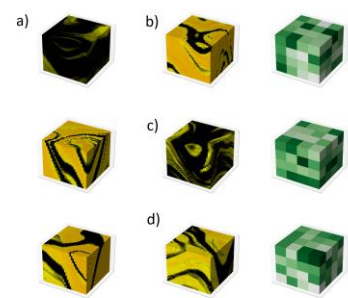

Figure 3: a) Randomly selected training patches of height/width 32 m/800 m. b-d) Random samples from a trained GAN (left) together with the corresponding latent space coordinate (right), padding voxels omitted for clarity.

When extrapolating to larger output sizes (Fig. 4a), quality remains sufficient, although at larger length scales, we observed more randomness than would be typical in a real submarine fan. This randomness is to be expected, because spatial correlations only extend as far as the projective field of the latent voxels, whereas in real submarine fans, the correlations can be much longer-range. Fortunately, in a full Bayesian workflow, the GAN latent space would not be sampled at random but would be focused, according to the Bayesian inference algorithm, by the measured seismic and well data that constrain the larger scales of geology.

To illustrate what a more correlated sample from the GAN might look like, we constructed a valid latent-space coordinate by tiling the latent coordinates of smaller patches, one having mostly sand and the other mostly shale (Fig. 3b & 3c, respectively). The top four rows of latent voxels are constructed from 16 copies of Fig. 3b (with adjustments made to the latent space padding). The next four rows consist of 16 copies of Fig. 3c, followed by 16 copies of Fig. 3b, and finally 16 of Fig. 3c again. This periodic latent space coordinate generates a periodic output image (Fig. 4b), with alternating sand/shale layers roughly corresponding to the locations of the latent patches belonging to the sand/shale patches in Fig. 3b-c. Some blending and transformation of the patterns also occurs due to the convolutions and nonlinear activation functions in the GAN.

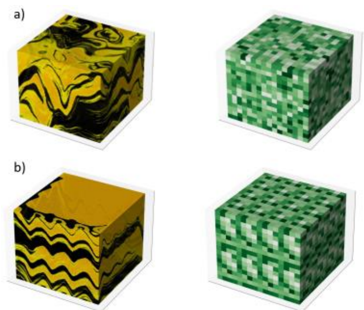

Figure 4: a) A random 128 m × 3200 m × 3200 m sample from a trained GAN (left) together with the corresponding latent space coordinate (right), padding voxels omitted for clarity. b) A GAN sample generated from a latent coordinate manually constructed by tiling the smaller latent coordinates from Fig. 3b-c.

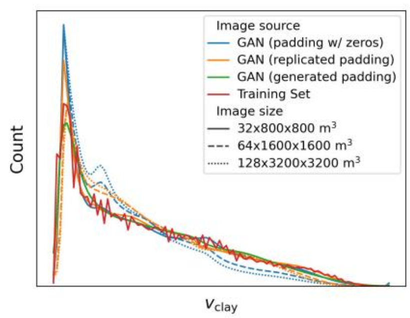

Figure 5: Histograms of Vclay values in voxels collected from multiple images from either trained GAN's or the training set.

For a more quantitative comparison between GAN outputs and training data, we computed histograms of individual voxels collected from multiple images (Fig. 5), as well as pooled averages of 32 × 32 × 32 voxel neighborhoods (Fig. 6) to examine the statistics at a larger length scale. The statistics of the trained GAN (green curves) agree well with the training set (red curve), both at the original training-set-patch size as well as larger, extrapolated output sizes.

As a final test of whether the trained GAN is fit for use, we measured the sampling time per image with a constant minibatch size of two on an NVIDIA® V100 GPU using PyTorch v2.1 (Fig. 7; Paszke et al., 2019). To improve accuracy, we increased the total number of generated images so that the total wallclock time is at least 1 second for each data point. As shown in Fig. 7, sampling time approaches linear scaling with voxel count at large sizes, where images containing 3e7 voxels still take less than a second.

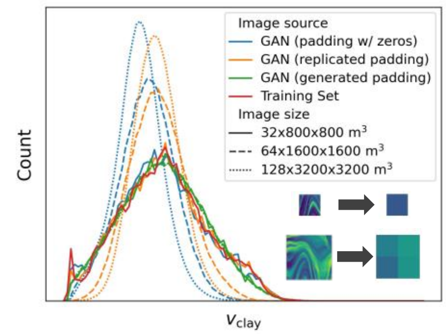

Figure 6: Histograms of Vclay values in 32 × 32 × 32 pooled averages of neighboring voxels (pooling example in lower right) collected from multiple images, either sampled from the training set or from GAN's trained with different types of padding.

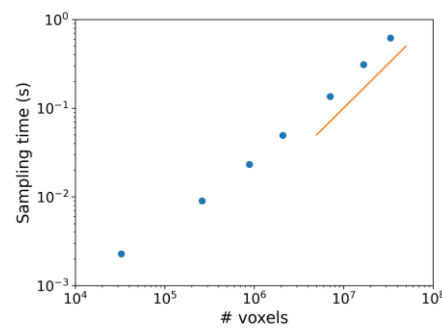

Figure 7: Time to sample a single image from the GAN as a function of voxel count. Orange line indicates linear scaling.

Improved Convolutional Padding for Better Quality

With any convolutional neural network (CNN), a decision needs to be made about how to treat pixels/voxels at the boundary of the input, which lacks a full neighborhood for the convolutional kernel to operate on. In computer vision CNNs, it is common to pad the boundary with zeros, an approach adopted for a previous version of the fully convolutional GAN (Jetchev et al., 2016; Song et al., 2022). For a fixed output size, we would expect the network to adapt to this artifact during training, however when extrapolating to larger output sizes, the ratio of boundary to interior voxels decreases, leading to unpredictable output and statistics, as shown by the blue curves in Figs 5 and 6 that deviate from the training set distribution. Switching to "replicated" padding gives similar results. To avoid these artifacts, we switched all convolutions to "valid" mode, meaning they only operate on full neighborhoods, shrinking the size of the convolution output. To compensate for the shrinking, we padded the latent space with multiple shrouds of voxels, accounting for the number of convolutions, up-sampling steps, and kernel size. In other words, by padding the latent space, the GAN is generating its own padding for the downstream convolutions, rather than assuming a fixed value. This alternative padding produces statistics that are invariant with output size, leading to better quality and more predictable behavior, as shown by the green curves in Figs 5 and 6 that agree well with the training set at multiple sizes and length-scales.

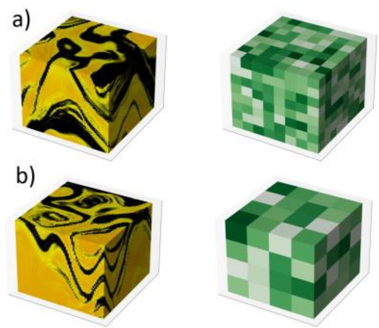

Figure 8: 64 m × 1600 m × 1600 m outputs and corresponding latent coordinates from a) the original GAN (Figs 3-4), and b) a more compressed GAN trained on output from the original GAN.

Training Strategy for Shrinking the Latent Space

Seismic inversion consists of searching the space of earth models (rock/fluid property distributions) for the model or distribution of models most consistent with the measured seismic and well data. We expect that this search would be easier in a lower-dimensional search space less impacted by the "curse of dimensionality" (Fernández-Martínez & Fernández-Muñiz, 2020). The GAN shown in Fig. 4 already achieves significant compression (500x, excluding latent padding). One way to push compression higher is to train on larger patches, while keeping the latent-space training size fixed. A side effect of this change is that the projective field also increases, leading to greater continuity at large scale at the expense of expressivity (diversity of output). Moreover, in training sets derived from process stratigraphy models (see Fig. 2), larger training patches would also decrease the number of unique patches that can extracted for training.

We tested a training strategy that increases compression while preserving the expressivity and training-set size of the original GAN. We used the trained GAN from Figs. 3-4 to synthesize training images of size 64 m × 1600 m × 1600 m, 8x larger than the original training patch size. We then trained another GAN on these synthetic patches, keeping the latent space size the same, achieving an ~8x increase in compression, for a total rate of 4000x (excluding latent padding). The output quality (Fig. 8) and statistics of the derived GAN are comparable to the original GAN.

Data Augmentation with Additional Prior Knowledge

As previously described, we experimented with adding two attributes that encode the relative azimuthal angle between the feeder channel axis and voxel location within the submarine fan. This angle roughly correlates with the orientation of sand bodies in our training set. The inclusion of these extra attributes would give a Bayesian inference algorithm more control over channel orientation in posterior samples, should such knowledge of regional orientation exist *a priori*. With the addition of these attributes, we initially observed lower quality in GAN output, a result we speculate to be caused by very different spatial correlations between these new attributes and the original ones (Vclay & porosity), versus correlations between the original ones themselves. We later observed much better quality after training with a GAN containing two separate generators connected to the same latent space (Fig. 9).

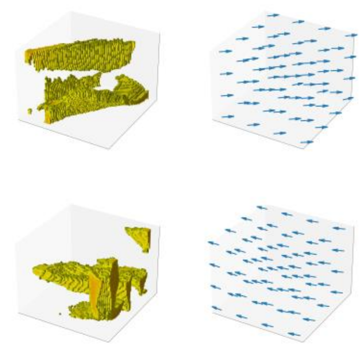

Figure 9: Outputs from a GAN trained on patches from a highly channelized submarine fan, where patches carry information about channel orientation. Vclay attribute with shale voxels made invisible for clarity (left). Vector field showing regional channel-orientation attribute (right).

CONCLUSIONS

We identified and trained a generative ML model that meets the requirements of fast sampling time and a highly compressed latent space, ideal for Bayesian inference, specifically, that involving seismic data at commercially relevant scales (>1e7 voxels). As part of this work, we increased the quality of very large outputs by improving how convolutions are padded, and demonstrated how compression could be increased by using a low-compression GAN to train a high-compression one. Finally, we augmented training data with attributes representing sand-body orientation, allowing a Bayesian inference algorithm to control channel orientation in posterior samples more directly.