

ARTICLE

FaciesGAN: A conditional GAN framework for realistic facies scenario generation as an efficient alternative to multiple-point statistics

Yineth Viviana Camacho-De Angulo^{1*}, **Tiago Mazzutti²**,
Bruno B. Rodrigues³, and **Mauro Roisenberg¹**

¹Department of Informatics and Statistics, Universidade Federal de Santa Catarina, Florianópolis, Santa Catarina, Brazil

²Instituto Federal Catarinense, Campus Camboriú, Santa Catarina, Brazil

³PETROBRAS, Rio de Janeiro, Rio de Janeiro, Brazil

Abstract

Facies are rock bodies that reflect specific depositional environments and play a central role in reservoir characterization. Accurate facies modeling is a key challenge in generating realistic geological scenarios that honor sparse well data while capturing geological uncertainty. This study introduces FaciesGAN, a novel deep learning framework based on conditional generative adversarial networks (cGANs). The method employs a hierarchical structure of generators and discriminators that progressively refine coarse estimates into high-resolution facies models, ensuring consistency with well data and depositional patterns at each stage. FaciesGAN was validated using the limited Stanford Earth Science Data dataset, demonstrating strong performance even under data scarcity. The quantitative evaluation employed multidimensional scaling and yielded an intersection over union index of 99.96% relative to the conditioning well data. These results confirmed the model's ability to generate diverse scenarios with high fidelity while preserving statistical distributions. Compared with a traditional multiple-point statistics implementation, FaciesGAN produced more realistic and varied geological realizations with significantly greater computational efficiency. These results indicate that cGAN-based approaches, such as FaciesGAN, represent a promising direction for subsurface modeling, offering robust tools for data augmentation, improved uncertainty assessment, and enhanced reservoir characterization.

*Corresponding author:

Yineth Viviana Camacho-De Angulo
(y.camacho@posgrad.ufsc.br)

Citation: Angulo YVC, Mazzutti T, Rodrigues BB, Roisenberg M. FaciesGAN: A conditional GAN framework for realistic facies scenario generation as an efficient alternative to multiple-point statistics. *J Seismic Explor.* 2025;34(6):45-59. doi: 10.36922/JSE025370069

Received: September 9, 2025

Revised: October 29, 2025

Accepted: November 3, 2025

Published online: November 28, 2025

Copyright: © 2025 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Keywords: Conditional generative adversarial network; Facies; Hard data; Geostatistical simulations; Seismic inversion

1. Introduction

In the context of reservoir characterization, facies are defined as rock units with specific attributes that reflect the depositional environment and directly influence the petrophysical properties and heterogeneity of the reservoir.^{1,2} Facies are essential for understanding depositional environments, as they enable geoscientists to correlate these units with seismic and well data, thereby playing a crucial role in the seismic inversion process.^{3,4} For example, since sandy facies generally exhibit higher porosity than shale

facies, they help identify and distinguish productive from non-productive zones.⁵ From this perspective, the integration of facies into seismic inversion algorithms provides greater consistency between the models and the petrophysical properties obtained, generating more robust and reliable models.⁶⁻⁸

The generation of multiple facies scenarios that reproduce complex heterogeneous structures plays a key role in the characterization and modeling of geological reservoirs and the stochastic seismic inversion workflow.^{9,10} By enabling multiple plausible realizations of the subsurface, facies scenario modeling provides a rigorous framework to explicitly capture and quantify geological uncertainty.¹¹ This approach reduces biases arising from single deterministic interpretations and ensures that the resulting reservoir models remain consistent with both geological knowledge and observed field data.^{2,12} This ability is particularly critical in seismic inversion workflows, where the relationship between well log measurements, core analysis, and seismic responses must be established in a consistent and geologically meaningful way.^{13,14}

Facies scenario generation can be carried out using classical and modern methodologies that combine geology, statistics, and artificial intelligence.^{10,15} Classical methodologies include techniques used in geology and geostatistics. For example, sequential indicator simulation (SIS) uses binary indicators for each facies, generating scenarios conditioned on available data.¹⁶ SIS is useful for modeling facies, but has several limitations. It often produces loosely connected patterns and oversimplified geological structures,¹⁷ making it difficult to represent features such as channels or faults. SIS is sensitive to variogram fitting, complicating its use with sparse data. It allows the quantification of uncertainty; however, if not accurately calibrated, it can result in geologically inconsistent models.¹⁸

Modern methodologies include techniques that have revolutionized geological modeling by allowing the representation of complex patterns and advanced spatial relationships. Multiple-point statistics (MPS) represent a significant advance in this area, enabling the capture of spatial patterns in geological data and modeling of multi-location relationships.^{19,20} These techniques are especially useful for simulating facies distributions in regions with limited information. They adhere to spatial distributions observed in training data, such as geological maps and previous simulations.^{13,21} MPS may face difficulties in constructing representative training images, as it relies on the analyst's expertise. Furthermore, conditioning to real data may be complex to implement without breaking the continuity of the simulated patterns.^{13,22}

In this context, generative adversarial networks (GANs) emerge as an innovative methodology for the generation of facies scenarios. GANs offer significant advantages over traditional geostatistical methods and MPS-based simulation. They can learn directly from real data, preserving first-order statistical features (facies proportions) and second-order statistical features (spatial continuity and body geometry).^{23,24} GANs are capable of capturing complex spatial patterns and facies relationships, thereby producing more realistic realizations and reducing the subjectivity inherent in model design.²⁵⁻²⁷ Moreover, they open the possibility of training networks as a complement to stochastic facies simulation.^{28,29}

There are two competing networks in GANs: A generator network creates synthetic data, and a discriminator network assesses the authenticity of the generated data in relation to the real data.^{22,25} In the context of facies, GANs can be used to generate new synthetic records that preserve the statistical characteristics of real data, for example, facies distributions and the geophysical properties of wells. On the other hand, conditional GANs (cGANs) include a conditional layer in the data generation process. This conditional layer allows the generation of synthetic data based on specific previous information, such as the type of facies in a particular depth range, providing even more control over the generation process.^{30,31} This characteristic allows the assessment of large-scale geological scenarios and the validation of hypotheses about reservoir connectivity and quality.

In recent years, facies scenario generation has been studied through several case studies, showcasing the effectiveness of advanced technologies. For example, Liu *et al.*³² proposed an approach for generating 3D subsurface facies map models based on GAN. Miele *et al.*³³ proposed integrating a GAN with spatially-adaptive denormalization (SPADE) to predict realistic facies map patterns while adhering to local probabilities. It combines with geostatistical methods of sequential simulation to model facies-conditioned rock properties. Furthermore, Feng *et al.*²³ proposed a GAN-based method in which the network is trained on facies map images. Research has demonstrated excellent results using facies map data and statistical similarity. However, few studies have incorporated known hard information from the GAN training stage, such as observed facies sequences in wells at specific locations. Specifically, no applications have been published on facies data in 2D vertical sections.

Considering the current progress, this work aims to explore advanced techniques for generating facies scenarios, with a particular focus on cGANs. The objective is to evaluate the effectiveness of this technique on 2D

vertical section facies data conditioned on well data. Using a public dataset with a limited number of samples, this study aims to demonstrate that the proposed method can effectively address one of the main challenges in reservoir characterization: data scarcity. This approach leverages synthetic training models to enhance the integration of well logs, facies distributions, and seismic information, producing scenarios that adhere to geological conditions and maintain statistical and spatial consistency. Accordingly, this study demonstrates that incorporating conditioning information enables cGANs to generate more accurate and robust models for reservoir characterization. cGANs offer an innovative solution to overcome the limitations of traditional techniques, contributing to a more coherent and efficient reservoir modeling.

2. Methodology

This study followed the workflow provided in Figure 1. The methodology comprises several interconnected stages.

2.1. Data collection and preprocessing

The Stanford Earth Science Data dataset was chosen and downloaded from the GitHub repository (<https://github.com/SCRFPublic/Stanford-VI-E>).⁴ The database contains data from oil well logs, with detailed samples of the different facies found in the reservoirs. The facies data are stored in .dat format, facilitating preprocessing and analysis. From the dataset, the available facies classes were floodplain (0), point bar (1), channel (2), boundary (3), and deltaic system (Figure 2A). The dataset primarily represented meandering channel systems, emphasizing facies categories relevant to this study (reservoir and non-reservoir types). The remaining facies were reclassified to simplify the categories into “reservoir” (channel; 1) and “non-reservoir” (floodplain, point bar, and boundary; 0), as shown in Figure 2B.

2.2. Image generation and data labeling

A Python 3.12 environment was configured using image processing and visualization libraries to generate visual representations of the filtered and categorized facies. The tabular data were subsequently converted into images. The 3D Stanford VI reservoir model was employed as training

data for the deep-learning workflow. This reference model was defined on a $150 \times 200 \times 200$ cell grid, with cell dimensions of 25 m in the horizontal (X and Y) directions and 1 m in the vertical (Z) direction. This resulted in a total physical size of 3,750 m (X-axis) \times 5,000 m (Y-axis) \times 200 m (Z-axis/depth). The 200 m vertical thickness was composed of three distinct layers (80 m, 40 m, and 80 m). To generate the 2D training images, 200 vertical slices (representing X–Z planes) were extracted, corresponding to one slice for each of the 200 cell positions along the Y-axis. Each slice represented the full horizontal (X-axis) distance of 3,750 m and the top 80 m layer (Layer 1). Subsequently, this physical section of 3,750 m \times 80 m was resampled to a 256×256 -pixel matrix. This process resulted in final images with a resolution of approximately 14.65 m/pixel in the horizontal direction and 0.31 m/pixel in the vertical (depth) direction.

A total of 200 divisions in 2D vertical slices were generated and extracted from the 3D facies model. These were used as training images, with 256×256 pixels, and categorized according to the corresponding facies class (Figure 2C). Each image was annotated with the depth condition and used as an external label to guide the process. The annotations delineating vertical polygons indicated the different facies represented, based on the 2D section facies found in the Stanford Earth Science Data dataset, and were used as conditioning data for the cGAN.

2.3. Facies scenario generation with the proposed cGAN

The proposed cGAN, termed FaciesGAN, features a multistage architecture designed to generate geologically realistic facies realizations conditioned on well data.

The FaciesGAN model is an adaptation of SinGAN²⁸ and WGAN-GP.³⁴ SinGAN is a generative model that can learn from a single natural image.²⁸ It consists of a pyramid of fully convolutional GANs, each modeling the distribution of image patches at a distinct spatial scale. This allows for generating new samples of arbitrary size and proportion. Although the generated samples exhibit considerable variability, they retain the overall structure and fine textures of the training image.

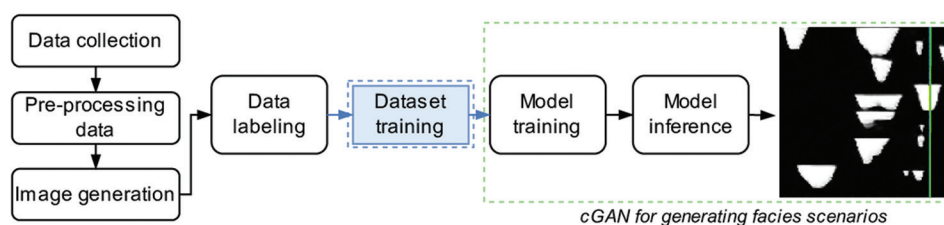


Figure 1. General proposed methodology

The FaciesGAN model is structured as a hierarchy of generators and discriminators operating at progressively

higher resolutions, as shown in Figure 3. The process begins with a low-resolution generator that produces

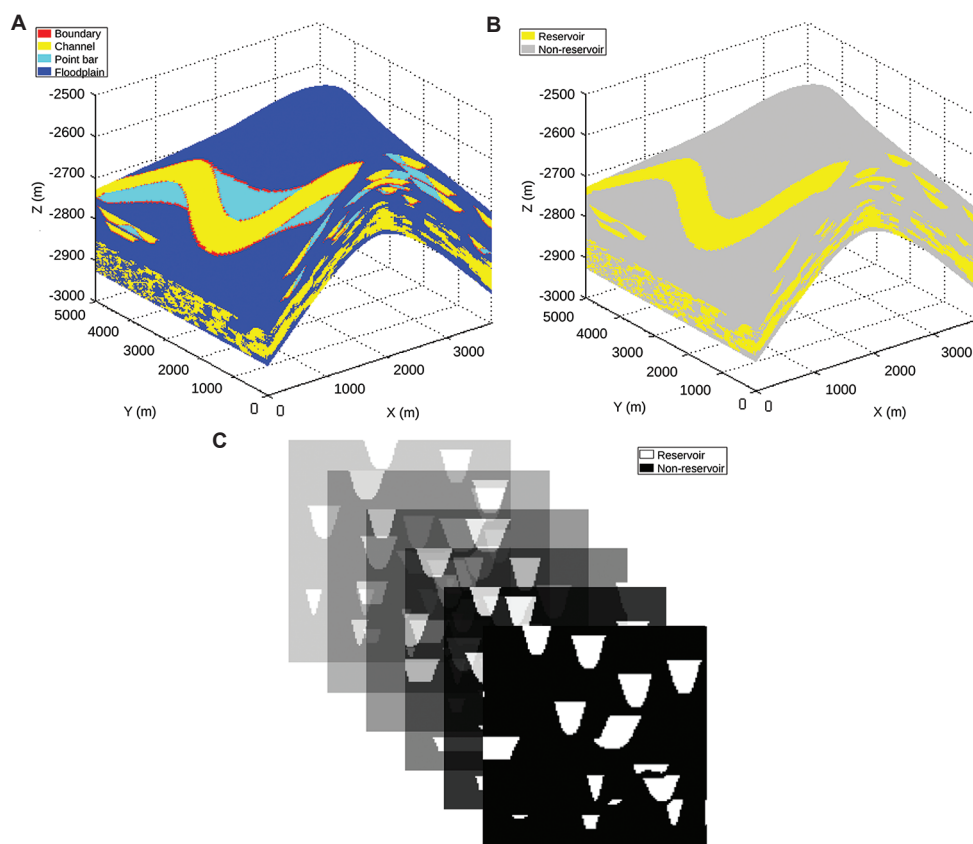


Figure 2. The Stanford Earth Science Data dataset. (A) Multiple sedimentary facies visualization. (B) Facies classification into reservoir (yellow) and non-reservoir (gray). (C) 2D slices examples of projection along the depth, differentiating the reservoir (white) from the non-reservoir (black). Image reproduced and adapted with permission from Lee and Mukerji.⁴

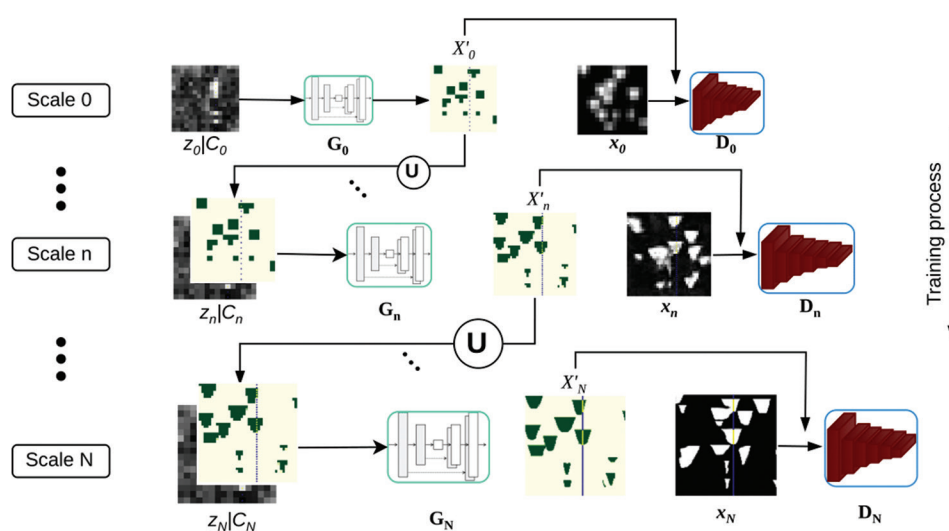


Figure 3. Schematic representation of the proposed cGAN for generating facies scenarios. The generator and discriminator are trained from coarse to fine scale (0 to N).

Abbreviations: cGAN: Conditional generative adversarial network; U: Upsampling.

an initial facies image conditioned on well information (e.g., facies at well locations). Subsequent stages refine this output by adding progressively finer geological details. Each generator stage is paired with a corresponding discriminator that evaluates the realism of the generated facies at its specific resolution while enforcing consistency with the conditioning data. Conditioning is maintained throughout all stages of the generation pipeline, ensuring that the final high-resolution outputs honor well constraints. This progressive refinement strategy allows the model to capture both large-scale geological structures and small-scale heterogeneities, resulting in high-quality, data-consistent facies simulations.

The pyramid of generators $G_0, \dots, G_n, \dots, G_N$ is a multi-scale, fully convolutional architecture, as shown in Figure 4A. At each scale, the generator considers a resized version of the previous output x'_{n-1} and a condition z_n , which are concatenated channel-wise. These are then passed through a series of 2D convolutional layers with leaky rectified linear unit (LeakyReLU) to produce a residual output. This is added to the up-sampled input to generate the new 2D section facies map x'_n . Each generator G_n is trained to learn the internal structure of the training images at different scales. G_n finer details from the training images are learned sequentially.²³

The discriminators, $D_0, \dots, D_n, \dots, D_N$, are implemented as a convolutional PatchGAN classifier, which assesses the realism of local image patches rather than making a single global prediction (Figure 4B). The architecture consists of a sequence of convolutional blocks, each comprising a 2D convolutional layer followed by a LeakyReLU activation function.²⁸ The number of feature channels is progressively reduced across layers (e.g., from 64 to 1), enabling hierarchical feature extraction at multiple spatial resolutions. Notably, normalization layers (e.g., batch normalization) are applied to preserve the raw feature dynamics and stabilize the training process. The final output is a single-channel feature map in which each spatial location corresponds to the discriminator's assessment of whether a specific image patch is real or synthetic.^{23,34}

At the n^{th} scale level, an adversarial training process is performed separately: the generator G_n tries to generate fake images x_n to fool the discriminator D_n . The discriminator D_n attempts to distinguish the real images x_n from the fake ones.²³ This multi-scale approach captures the large-scale structures present in the geologic models of interest.³⁵ The formulation for generating an image sample at the n^{th} level is expressed as follows:

$$x'_n = \begin{cases} G_0(z_0), & n=0 \\ G_n(z_n, \cup(x'_{n-1})), & 0 < n \leq N \end{cases} \quad (1)$$

where \cup represents the upsampling based on interpolators in the 2D and 3D cases.

The loss function at the n^{th} scale level for G_n and D_n is formulated as:²⁸

$$\min_{G_n} \max_{D_n} \mathcal{L}(G_n, D_n) = \mathcal{L}_{adv}(G_n, D_n) + \alpha \mathcal{L}_{rec}(G_n) \quad (2)$$

where \mathcal{L}_{adv} is the adversarial loss for penalizing the distribution distance between the down-sampled images x_n and the generated images x'_n , α is a weighting factor to balance the two loss functions, and \mathcal{L}_{rec} is the reconstruction loss to ensure that x_n can be reproduced given a specific set of random noise maps.

The generator G_n and discriminator D_n at each pyramid scale n are trained with a combined objective inspired by WGAN-GP³⁴ and SinGAN.²⁸ The goal is to simultaneously enforce adversarial learning and faithful reconstruction of the image at multiple resolutions.

The discriminator D_n is optimized using the Wasserstein loss with gradient penalty, ensuring Lipschitz continuity. The discriminator loss is formulated as:

$$\begin{aligned} \mathcal{L}_D^{(n)} = & -\mathbb{E}_{x_n \sim P_{data}} [D_n(x_n)] + \mathbb{E}_{z_n} \left[D_n \left(G_n \left(z_n, x_{n+1}^{\uparrow} \right) \right) \right] \\ & + \lambda_{gp} \mathbb{E}_{\hat{x}_n} \left[\left(\nabla_{\hat{x}_n} D_n(\hat{x}_n)_2 - 1 \right)^2 \right] \end{aligned} \quad (3)$$

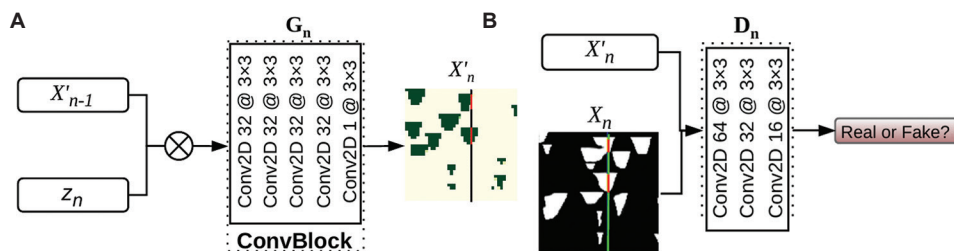


Figure 4. Network architecture at the n^{th} scale level. (A) The generator. (B) The discriminator.

Where x_n denotes a real image at scale n , $G_n \left(z_n, x_{n+1}^{\uparrow} \right)$

is the generated image conditioned on noise z_n and the upsampled output from the next coarser scale, and $\hat{x}_n = \varepsilon x_n + (1 - \varepsilon) x'_n$ with $\varepsilon \sim u[0,1]$ is the interpolated sample used for the gradient penalty.

The generator G_n is trained with two complementary objectives: (i) An adversarial loss that encourages generated samples to be indistinguishable from real ones at scale n , and (ii) a reconstruction loss that ensures faithful reproduction of the reference image when a fixed noise map z_n^* is used. The generator loss is:

$$\mathcal{L}_G^{(n)} = -\mathbb{E}_{z_n} \left[D_n \left(G_n \left(z_n, x_{n+1}^{\uparrow} \right) \right) \right] + \alpha_n \mathcal{L}_{rec}^{(n)} \quad (4)$$

With the reconstruction loss defined as:

$$\mathcal{L}_{rec}^{(n)} = x_n - G_n \left(z_n^*, \hat{x}_{n+1}^{\uparrow} \right)_2 \quad (5)$$

Where α_n is a scale-dependent weighting factor balancing adversarial and reconstruction objectives, and \hat{x}_{n+1}^{\uparrow} denotes the generated sample from the scale $n+1$, upsampled to match the resolution of scale n .

This hierarchical optimization scheme allows the generator to progressively capture global structure at coarse scales and fine details at higher resolutions, while the reconstruction term stabilizes training and preserves fidelity.

2.4. Algorithm and implementation

The FaciesGAN model is an architecture designed to generate geologically consistent facies images from a multi-scale noise pyramid. It uses an improved adversarial training framework. FaciesGAN's training loop comprises two alternating main stages, which involve updating the discriminator and generator parameters. The algorithm incorporates additional mechanisms, such as gradient penalty, reconstruction, and masking losses that contribute to improving training stability and fidelity of the generated images. The algorithm and the core procedure for training the FaciesGAN model at a single resolution scale are presented in Algorithm 1. FaciesGAN core training loop (at a single scale).

The training hyperparameters were determined based on the original WGAN-GP and SinGAN models, with empirical adjustments for our specific application. The gradient penalty weight λ_{gp} (referred to as λ in WGAN-GP) was set to 0.115, a value that we found stabilized training

Algorithm 1. FaciesGAN core training loop (at a single scale)

Input:

$x_{real} \leftarrow$ Real data
 $M \leftarrow$ Mask
 $x_{rec_in} \leftarrow$ Reconstruction input

Models:

$G \leftarrow$ Generator
 $D \leftarrow$ Discriminator

Hyperparameters:

$\lambda_{gp}, \alpha_{rec}, k_d \leftarrow$ Discriminator steps
 $k_g \leftarrow$ Generator steps

/* Step 1: Train Discriminator */

```
1: for  $j = 1$  to  $k_d$  do
2:   Sample noise pyramid  $Z \leftarrow \text{GETNOISE}()$ 
3:   Generate fake images  $x_{fake} \leftarrow G(Z)$  // Forward pass-through Generator
4:    $L_{real} \leftarrow -\mathbb{E} [D(x_{real})]$  // Loss for real data
5:    $L_{fake} \leftarrow \mathbb{E} [D(x_{fake})]$  // Loss for fake data
6:    $L_{gp} \leftarrow \lambda_{gp} \cdot \text{CALCULATEGRADIENTPENALTY}(D, x_{real}, x_{fake})$ 
7:    $L_D \leftarrow L_{real} + L_{fake} + L_{gp}$  // Total Discriminator loss
8:   Update  $D$ 's parameters  $\theta_d$  by ascending the gradient of  $L_D$ 
9: end for
```

/* Step 2: Train Generator */

```
10: for  $j = 1$  to  $k_d$  do
11:   Sample noise pyramid  $Z \leftarrow \text{GETNOISE}()$ 
12:   Generate fake images  $x_{fake} \leftarrow G(Z)$ 
13:    $L_{adv} \leftarrow -\mathbb{E} [D(x_{fake})]$  // Adversarial loss
14:   Sample reconstruction noise  $Z_{rec} \leftarrow \text{GETNOISE}(rec=True)$ 
15:    $x_{rec} \leftarrow G(Z_{rec}, in\_facie = x_{rec\_in})$  // Reconstruction pass
16:    $L_{rec} \leftarrow \alpha_{rec} \cdot \text{MSE}(x_{rec}, x_{real})$  // Reconstruction loss
17:    $L_{mask} \leftarrow 100 \cdot \alpha_{rec} \cdot \text{MSE}(x_{fake} \odot M, x_{real} \odot M)$  // Masked loss
18:    $L_G \leftarrow L_{adv} + L_{rec} + L_{mask}$  // Total Generator Loss
19:   Update  $G$ 's parameters  $\theta_g$  by ascending the gradient of  $L_G$ 
20: end for
```

effectively for the facies data (in contrast to the $\lambda = 10$ used in the original WGAN-GP). The reconstruction weight α_{rec} (referred to as α in Equation [2] and α_n in Equation [4]) was set to 10, a value commonly used in SinGAN-based models that provided an optimal balance between adherence to geological structure and training stability.

3. Results

This study evaluated the capability of generative models to generate geological facies scenarios. For this purpose, FaciesGAN was trained and validated through visual inspection and multidimensional scaling (MDS) to determine the consistency and representativeness of the generated scenarios with the original facies. For comparison purposes, the same data were modeled using an MPS method, specifically, the single normal equation simulation (SNESIM).¹³ The scenarios generated using

the two methodologies were compared through visual inspection of the spatial continuity of the patterns and the facies proportion histogram. This allowed for a qualitative and quantitative analysis of the representativeness and consistency of the simulated models.

Specifically, FaciesGAN was trained to generate facies scenarios using the Stanford Earth Science dataset. The model was developed with a limited training set of 200 samples. During the inference stage, some samples were analyzed using metrics such as visual inspection, average facies proportion, and MDS to determine the consistency and representativeness of the generated scenarios with the original facies.

For FaciesGAN training, appropriate labels were required for each image. The labeling process is shown in Figure 5. These labels are important because they provide information about the characteristics of each image, allowing the model to learn to generate coherent and realistic images based on specific conditions. In this context, the labels corresponded to hard data derived from a simulated well, representing known subsurface information used to condition the facies generation process. The correctly labeled images were integrated into the dataset and associated specifically with each corresponding image. The model used the labels as conditioning input to generate facies scenarios consistent with the characteristics and structures defined by the labels.

The FaciesGAN model was trained for 100 epochs per scale across 10 scales, with a gradient penalty weight λ_{gp} of 0.1, using the Adam optimizer with a learning rate of $5e^{-5}$ and β of 0.5. The kernel size for 2D filters was 3×3 , with a stride step of 1×1 . At the coarser scales, image resolution ranged from 16 to 128 pixels. The model was trained on the complete dataset using a workstation with an Intel i7-8700K CPU (6 cores, 3.7 GHz), an NVIDIA GeForce GTX 1080Ti GPU, and 64 GB of RAM.

3.1. Global model evaluation

In the first test, the model generated 1,000 facies scenarios in 20 s; twenty randomly selected conditioned realizations

are shown in Figure 6. It was observed that the facies configuration of the conditioning trace, highlighted in green to simulate a real drilled and analyzed well, was closely reproduced in the images generated by FaciesGAN.

The results are promising considering the limited training set, highlighting the applicability of the proposed approach in characterizing oil reservoirs, where well log and facies data are often scarce, costly to obtain, and subject to privacy restrictions. Nevertheless, the model showed remarkable consistency in reproducing the facies spatial distributions. These findings provide insight into the model's capability to produce images that consistently reflect the expected facies proportions. Visual comparisons with real distributions confirmed that the model captured key features of the input data while generating consistent variations. Furthermore, the time required to generate facies scenarios was short, highlighting the computational efficiency of the proposed approach. The short generation time allows for practical integration into workflows that require multiple simulations.

Next, MDS was applied to quantitatively evaluate the trained model and to compare patterns of spatial variability. MDS is a technique commonly used in data analysis and visualization. It represents high-dimensional data in a lower-dimensional space, usually 2D or 3D, while preserving the relative distances (or dissimilarities) between data points and the potential differences between them.³⁶ The generated facies overlapped closely with the training images in 2D space, demonstrating excellent similarity, as shown in Figure 7. Regions where blue and red overlap indicate highly agreement between generated and real images, suggesting robust model generalization.

The generated facies (red) effectively covered the space of the real facies (blue), indicating the diversity and quality of the generator. The real facies (blue) were closely surrounded by the generated facies, suggesting that the generator interpolates well within the known domain. This indicates a high degree of spatial consistency.

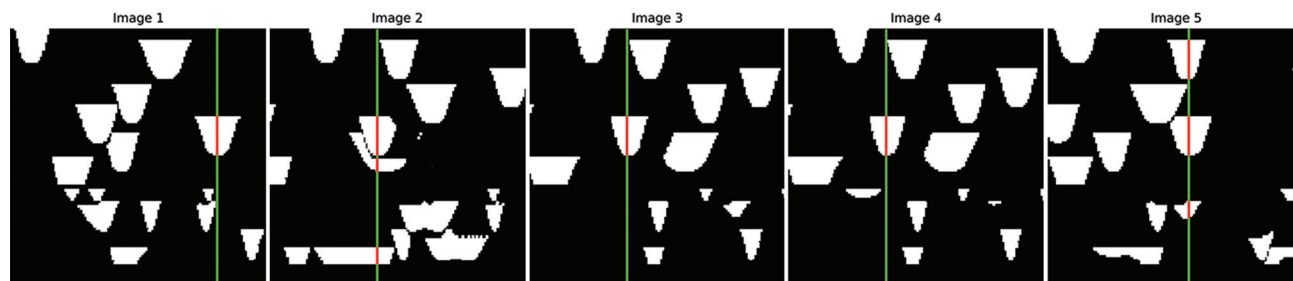


Figure 5. Reservoir (in white) and non-reservoir (in black) with drilled-well conditioning. Conditional traces are highlighted in red (reservoir) and green (non-reservoir). Note: Each image corresponds to a 2D crossline section represented in the pixel domain (256×256 pixels) to an 80 m (depth) \times 3,750 m (width) vertical section.

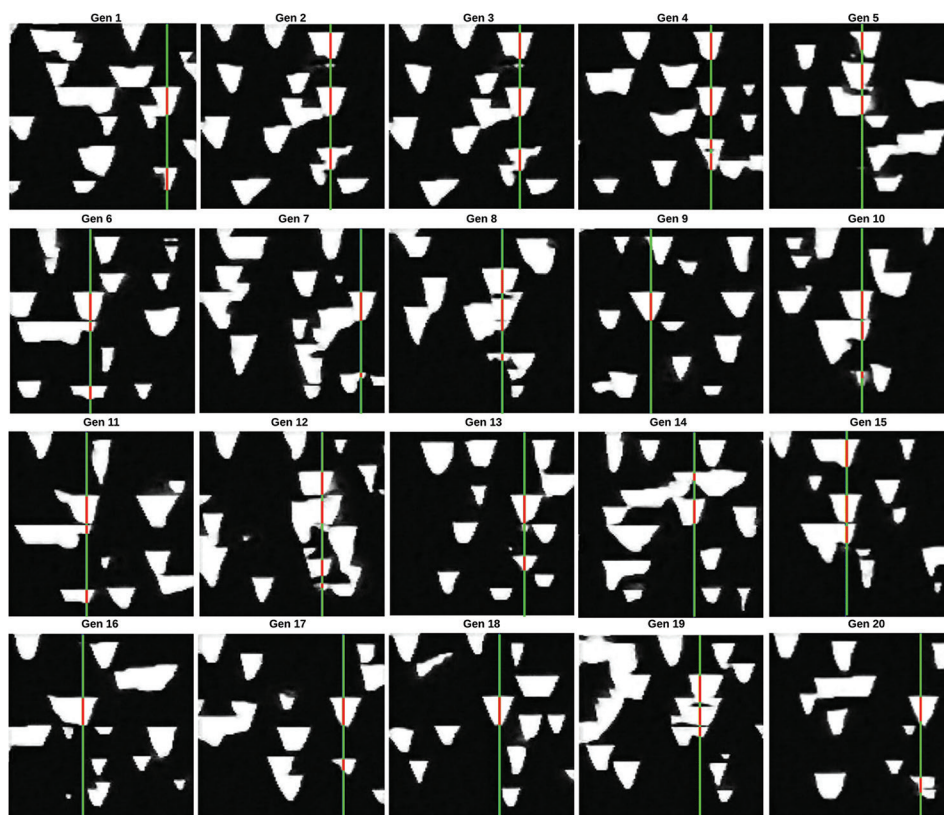


Figure 6. Twenty randomly selected realizations generated by the proposed cGAN. The generated facies are shown in white, while the conditioning is in red (reservoir) and green (non-reservoir). Note: Each image represents a 2D crossline section in the pixel domain (256×256 pixels), corresponding to a vertical section 80 m deep and 3,750 m wide.

Abbreviation: cGAN: Conditional generative adversarial network.

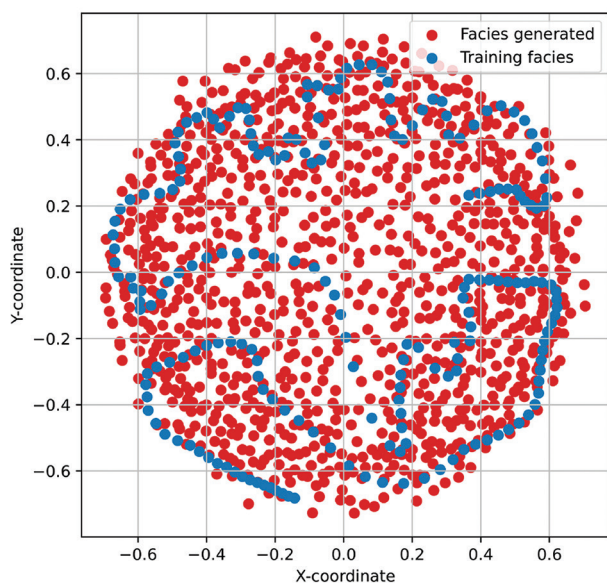


Figure 7. Multidimensional scaling plot of the training images with conditional realizations

In addition, to validate how effectively the conditioning well information was honored by the FaciesGAN model,

a procedure was designed that reformulates the problem as a spatial classification task. Specifically, the generated scenarios were compared with their respective original images to evaluate how accurately the location and shape of the facies were reproduced around the actual wells. The intersection over union (IoU) index was used as the evaluation metric. This index is defined as the ratio between the intersection area and the union area of the predicted and reference data:

$$IoU = \frac{GT \cap PD}{GT \cup PD} \quad (6)$$

where PD is the prediction mask and GT is the ground truth. In this evaluation, the prediction mask corresponded to the pixels generated under hard conditioning by FaciesGAN, while the hard-conditioning reference data from the original facies image served as the ground truth. A total of 1,000 images generated from a set of 200 original images were analyzed. The IoU was calculated for each pair of images, yielding an overall mean IoU of 99.96%. This result indicates exceptionally high fidelity in preserving the well-conditioning constraints and demonstrates

that FaciesGAN generates stochastic images that almost perfectly adhere to the geological information observed in the well. These findings validate its effectiveness as a geological conditioning tool.

The effectiveness of the FaciesGAN model was validated by comparing its results with synthetic facies scenarios generated using MPS. Specifically, SNESIM, an improved and scalable extension of the extended normal equation simulation (ENESIM) algorithm for multipoint simulation, was used to generate 1,000 facies realizations based on the 200 training images. The results of 20 representative simulations are presented in Figure 8. These results were generated in approximately 137 min.

The facies scenarios generated by FaciesGAN (Figure 6) exhibited distributions consistent with the expected geology. The scenarios accurately respected the conditions (in green). These realizations reflect the remarkable ability of the model to capture complex spatial patterns with high diversity among simulations. In comparison, the realizations generated by the MPS SNESIM method also preserved the spatial continuity of the facies; however, they

exhibited less structural variability than those produced by FaciesGAN. Visual comparison suggests that FaciesGAN accurately reproduced the input conditioning and provided greater structural diversity in its realizations. This demonstrates that the proposed methodology is a robust alternative for generating complex geological scenarios.

Next, we evaluated the overall distribution and class balance within the dataset. The histogram of reservoir facies proportions is shown in Figure 9. In addition to the dataset distribution, the histogram also includes the facies proportion results obtained from the FaciesGAN and MPS SNESIM simulations. This enables a comparative analysis of class balance between the original data and the synthetic realizations produced by the two methods.

The distribution of the dataset (in red) showed a primary peak near 0.16, representing the dominant facies ratio in the real data. The dispersion is moderate, with most realizations concentrated between 0.12 and 0.20. The distribution generated by FaciesGAN (in green) showed a similar behavior, with values concentrated in the same range. However, a slight deviation was observed

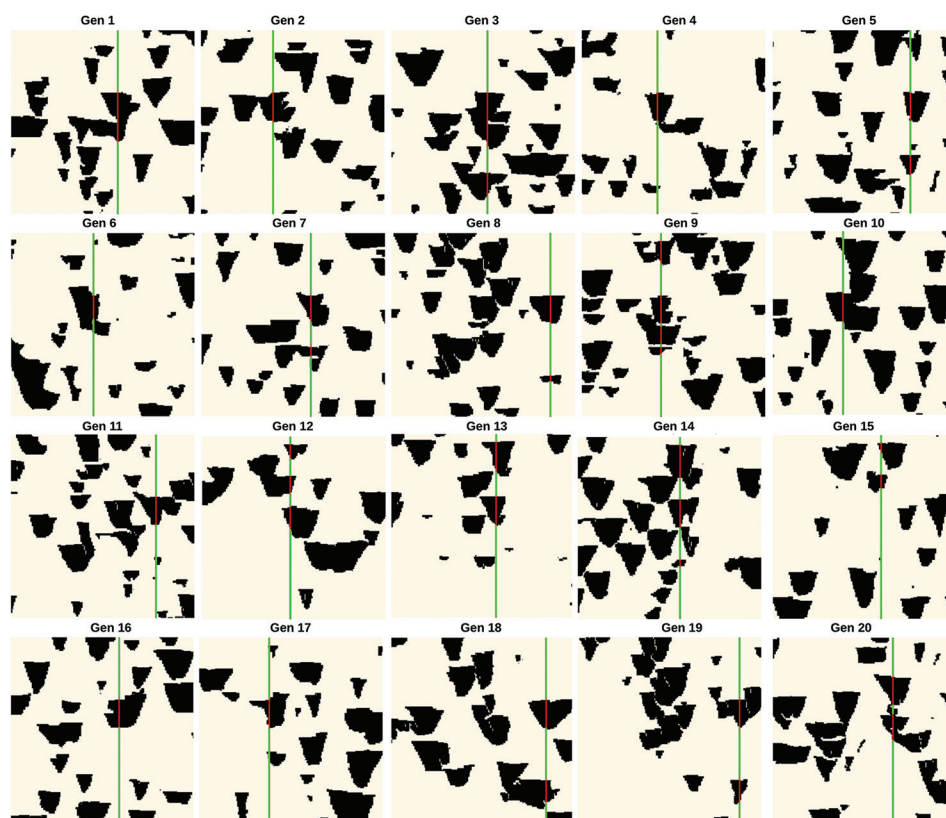


Figure 8. Twenty randomly selected realizations generated by the MPS SNESIM algorithm. The generated facies are shown in black, while the conditioning is in red (reservoir) and green (non-reservoir). Note: Each image represents a 2D crossline section in the pixel domain (256×256 pixels), corresponding to a vertical section 80 m deep and 3,750 m wide.

Abbreviations: MPS: Multiple-point statistics; SNESIM: Single normal equation simulation.

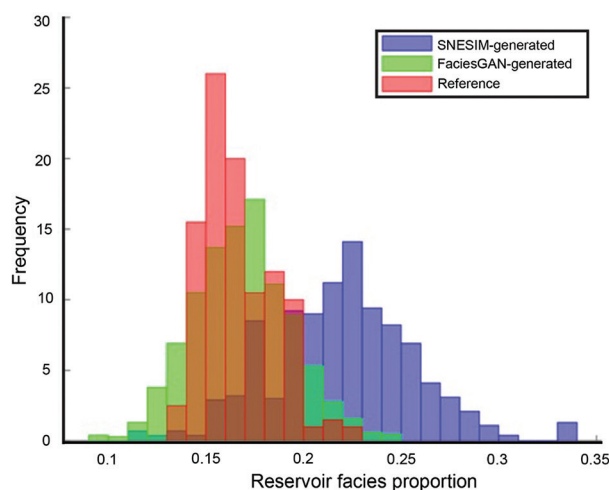


Figure 9. Comparative histogram of reservoir facies proportions generated by FaciesGAN and SNESIM

Abbreviations: GAN: Generative adversarial network; SNESIM: Single normal equation simulation.

toward higher values, indicating a minor overestimation of the proportion in some simulations. In contrast, the distribution generated by SNESIM (in blue) was significantly dispersed, covering a broader range from 0.10 to 0.35. A clear trend toward higher proportions implies lower statistical fidelity compared to the real data. In addition, SNESIM-generated results demonstrated higher variability than those by FaciesGAN and the reference, but were in a controlled interval.

These results indicate that FaciesGAN provides a closer approximation of the observed facies ratios in the reference data compared to SNESIM. The higher variability of SNESIM results in deviations from the true statistical behavior, which can be a major limitation when accurate preservation of facies proportions is required. In addition, the computational efficiency of FaciesGAN is notably superior: while SNESIM took approximately 137 min to generate 1,000 realizations, FaciesGAN produced the same number in only 20 s. This highlights that FaciesGAN has a greater ability to learn and reproduce the distributions observed in real data, enabled by its deep learning-based generative process.

3.2. Well-specific conditioning results

In the second test, five conditioning images distributed in 2D space were selected. For each image, the trained model generated 100 samples at approximately 6 s per image. From each image, five generated facies scenarios are presented in Figure 10. The figure presents the real facies (left column, in green, with the depth condition shown in black) and multiple random model-generated realizations (five columns per well, in black and white) for five different wells, with the depth condition highlighted in green.

The condition incorporated during FaciesGAN training was held constant. The generated facies closely surround or overlap with the real facies, indicating strong spatial consistency between the model realizations and the true data. Although variability was present among generated facies, most realizations maintained structural patterns consistent with the real facies, suggesting that the model adequately learned the underlying spatial patterns.

The real facies were closely surrounded by the generated facies, demonstrating the model's capability to preserve spatial structures across different realizations. This consistency highlights the model's reliability in reproducing subsurface geologic patterns, even under stochastic variability. In such cases, the generation of facies scenarios was fast.

The percentages of pixels corresponding to each facies were compared to evaluate statistical consistency between the real images and those generated by FaciesGAN. The comparison between the percentages observed in the real images and the averages obtained from 100 generated scenarios for the five wells shown in Figure 10 is presented in Table 1. This comparison analyzed the model's ability to reproduce facies distributions realistically, ensuring that the simulations preserve the original geologic characteristics.

A strong correspondence was observed between the facies percentages of the real and generated images. Across all wells, differences between the real values and the generated averages were <4%, indicating that FaciesGAN maintains high fidelity in reproducing facies proportions. For example, in well 181, facies 0 accounted for 72.66% in the real image and 72.71% in the generated average—a practically insignificant difference. Similar cases were observed in the other wells, with the largest deviation occurring in well 63, where facies 1 decreased by approximately 3%. This minor underestimation remains within acceptable ranges for stochastic simulations.

In addition, the IoU metric was calculated for each well to further evaluate the fidelity of facies preservation within the conditioned zones. Table 1 presents the average IoU obtained from each case. The average IoU values ranged from 99.37% to 99.66%, indicating an extremely high agreement between the generated and original facies in the hard-conditioned wells. The results suggest that the FaciesGAN maintains near-identical facies proportions and effectively reproduces the stratigraphic continuity observed in the real data.

The generated scenarios were encouraging, as the model demonstrated a high degree of consistency and realism in reproducing the spatial distribution of facies from a

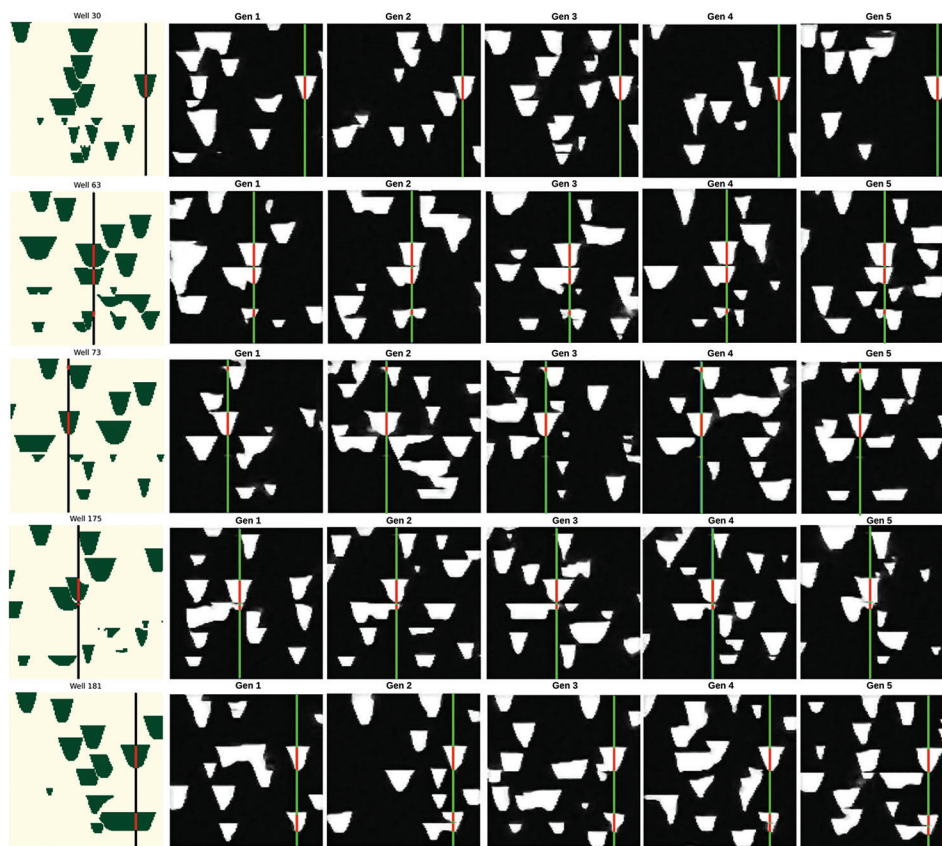


Figure 10. Real and generated facies for five different wells. The first column shows the real 2D facies logs (in green) with the conditioning shown in black. The five columns (Gen 1 to Gen 5) display different random realizations generated by the model, with facies shown in white. The conditioning is highlighted in red (reservoir) and green (non-reservoir).

Note: Each image represents a 2D crossline section in the pixel domain (256×256 pixels), corresponding to a vertical section 80 m deep and 3,750 m wide.

Table 1. Averaged facies percentages in real and generated scenarios across five selected wells

Well number	Real images		Generated images		IoU (%)
	Facies 0 (%)	Facies 1 (%)	Facies 0 (%)	Facies 1 (%)	
30	84.77	15.23	85.16	14.84	99.66
63	69.92	30.08	72.35	27.65	99.42
73	79.30	20.70	80.47	19.53	99.38
175	79.30	20.70	81.35	18.65	99.37
181	72.66	27.34	72.71	27.29	99.43

Abbreviation: IoU: Intersection over union.

limited dataset. Furthermore, visual and quantitative comparisons with the actual facies distributions confirmed the model's ability to capture key features of the input data, integrate conditional information, and generate significant variability across realizations. These results suggest that the FaciesGAN model generalizes effectively and serves as a robust tool for generating 2D facies scenarios in petroleum applications, even with limited training

datasets. This capability can significantly enhance reservoir characterization and support operations planning.

4. Discussion

The results demonstrate that FaciesGAN produces highly coherent and realistic geological facies scenarios, even when trained on a limited dataset. The generated realizations captured key geological patterns and spatial continuity, showing high fidelity to the conditioning information.

The stochastic simulation methodology employed does not aim to identify a single, “optimal” scenario but rather to quantify geological uncertainty by producing multiple realistic representations of the subsurface that respect the conditioning data. This approach is crucial for risk-based decision-making, as the true facies distribution is unknown. Accordingly, the workflow uses the full ensemble (e.g., as inputs to flow models) to estimate outcome ranges (e.g., P10, P50, and P90) rather than a single “best” result. The statistical consistency of this ensemble with the reference model was evaluated using MDS (Figure 7) and

facies proportion histograms (Figure 9), ensuring that the generated scenarios were statistically representative and effectively explored the geological uncertainty space, as detailed in Section 3.1.

Visual inspection confirmed that FaciesGAN accurately reproduced facies structures and variability, outperforming SNESIM in terms of structural diversity. Quantitative metrics further reinforced this result: the mean IoU (99.96%) indicates near-perfect preservation of well conditioning, and facies-percentage distributions closely aligned with the real data, with deviations generally <4%.

The comparative histogram highlighted that FaciesGAN better approximated the observed facies proportions than SNESIM, which tended to generate realizations with greater variability and statistical deviation. In addition, FaciesGAN demonstrated superior computational efficiency, producing 1,000 scenarios in 20 s compared to SNESIM's 137 min.

These results suggest that deep generative approaches, such as FaciesGAN, offer a robust alternative for simulating geological facies. The model's ability to learn spatial patterns and accurately reproduce conditioning information makes it promising for reservoir modeling workflows that demand both accuracy and efficiency. The results also demonstrate FaciesGAN's potential for generalization, as the model maintained high consistency across different conditioning wells, with stochastic variability remaining within acceptable geostatistical ranges.

The FaciesGAN model was trained and validated exclusively on the Stanford Earth Science Data dataset, which represents a particular meandering-channel system. While the results are promising, especially under conditions of scarce data, further research is necessary to confirm the model's applicability to various geological contexts, including turbidite systems, deltaic environments, and carbonate platforms, which exhibit distinct spatial patterns and heterogeneities.

Furthermore, transfer learning is a promising area for future research. A model pre-trained on a large and diverse set of public geological models could be fine-tuned using smaller, field-specific datasets. This approach could enhance the practicality of FaciesGAN in real-world reservoir characterization projects, where data availability is always limited.

Overall, the tests confirm that FaciesGAN can generate realistic, diverse, and conditionally consistent facies realizations, offering advantages in terms of speed and statistical performance compared with traditional MPS methods. While this study primarily compared

FaciesGAN with the conventional MPS SNESIM method, subsequent research should evaluate its performance against additional deep generative models. For example, variational autoencoders could be explored, although cGANs have already demonstrated key advantages in generating scenarios with sharper geological boundaries and closer alignment to the true statistical distribution of reservoir properties³⁷—key attributes for realistic facies modeling.

5. Conclusion

The use of artificial intelligence-based techniques to generate facies scenarios is an innovative area aimed at improving the accuracy and robustness of machine learning models in oil exploration and production. cGANs are gaining prominence due to their capacity to generate high-quality synthetic data that preserves known geological characteristics. FaciesGAN was successfully trained with a limited number of 2D facies images, demonstrating strong performance in data-scarce scenarios and offering an effective approach for data augmentation with small datasets. The main advantage of the model lies in its ability to generate synthetic 2D facies scenarios while honoring known conditional information, ensuring consistency with real conditions derived from drilled and analyzed wells. MDS and facies-proportion statistics produced favorable results, highlighting the ability of this network to consistently reproduce the conditioning data. The generated realizations preserve the frequency distributions and spatial correlations characteristic of the original images, ensuring visual consistency and maintaining statistical and geological integrity. Furthermore, a comparative analysis with SNESIM demonstrated that FaciesGAN provides a more accurate representation of facies proportions, with reduced dispersion and skewness relative to the reference data, while maintaining higher spatial and statistical fidelity. Another significant advantage of FaciesGAN is its computational efficiency, which enables the rapid generation of multiple scenarios compared with sequential simulation methods. These synthetic realizations can be integrated into reservoir characterization workflows to support uncertainty estimation and enhance the quality of results. Finally, the proposed methodology can be extended to subsequent workflow steps, such as incorporating facies with acoustic and/or elastic properties, generating synthetic seismic data, and evaluating consistency with actual seismic observations. In particular, future work will focus on applying the approach to real-world field datasets—a logical and most important next step. This extension will enable synthetic seismic generation to be combined with seismic inversion. This integration will

establish a direct and useful link among facies modeling, seismic inversion, and reservoir characterization.

Acknowledgments

None.

Funding

This research was funded by the project “4D Integrated Quantitative Characterization of Reservoirs” (Contract No. FEESC/PETROBRAS/4600671602) and the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior–Brasil (CAPES)–Finance Code 001.

Conflict of interest

The authors declare that they have no competing interests.

Author contributions

Conceptualization: Yineth Viviana Camacho-De Angulo, Mauro Roisenberg

Data curation: Yineth Viviana Camacho-De Angulo

Formal analysis: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti, Mauro Roisenberg

Funding acquisition: Mauro Roisenberg

Investigation: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti, Mauro Roisenberg

Methodology: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti, Mauro Roisenberg

Project administration: Bruno B. Rodrigues, Mauro Roisenberg

Resources: Bruno B. Rodrigues

Software: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti

Supervision: Mauro Roisenberg

Validation: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti, Mauro Roisenberg

Visualization: Yineth Viviana Camacho-De Angulo

Writing–original draft: Yineth Viviana Camacho-De Angulo

Writing–review & editing: Yineth Viviana Camacho-De Angulo, Tiago Mazzutti, Bruno B. Rodrigues, Mauro Roisenberg

Availability of data

This study employed a publicly available facies dataset from the Stanford VI-E Reservoir: A Synthetic Data Set for Joint Seismic-EM Time-lapse Monitoring Algorithms provided by the Stanford Center for Reservoir Forecasting (SCRF), accessible at <https://github.com/SCRFpublic/Stanford-VI-E/tree/master>.

The code availability for the FaciesGAN model is as follows:

FaciesGAN

Contact: y.camacho@posgrad.ufsc.br

Hardware requirements: No hardware requirement.

Program language: Python

Software required: Python 3.1x

Program size: 62.9 MB

The source codes are available for downloading at the link: <https://github.com/mazzutti/FaciesGAN>

References

1. Sheriff RE, Geldart LP. *Exploration Seismology*. Cambridge: Cambridge University Press; 1995. Available from: https://books.google.com.br/books?hl=es&lr=&id=wRYgAwAAQBAJ&oi=fnd&pg=PA1901&dq=Sheriff+RE,+Geldart+LP.+Exploration+Seismology.+Cambridge:+Cambridge+University+Press%3B+1995&ots=wCZUKyTUNq&sig=ELLBeg-IVBcYuxFxazP0aboszHY&redir_esc=y#v=onepage&q=Sheriff%20RE%2C%20Geldart%20LP.%20Exploration%20Seismology.%20Cambridge%3A%20Cambridge%20University%20Press%3B%201995&f=false [Last accessed on 2025 Nov 06].
2. Grana D, Mukerji T, Doyen P. *Seismic Reservoir Modeling: Theory, Examples, and Algorithms*. United States: John Wiley and Sons; 2021.
3. Sangree JB, Widmier JM. Interpretation of depositional facies from seismic data. *Geophysics*. 1979;44(2):131-160. doi: 10.1190/1.1440957
4. Lee J, Mukerji T. *The Stanford VI-E Reservoir: A Synthetic Data Set for Joint Seismic-EM Time-Lapse Monitoring Algorithms*. 25th Annual Report, Stanford Center for Reservoir Forecasting, Stanford, CA: Stanford University; 2012.
5. Selley RC. *Elements of Petroleum Geology*. United States: Gulf Professional Publishing; 1998.
6. Veeken PCH. *Seismic Stratigraphy, Basin Analysis and Reservoir Characterisation*. Vol. 37. Netherlands: Elsevier; 2006.
7. Grana D, Mukerji T, Dvorkin J, Mavko G. Stochastic inversion of facies from seismic data based on sequential simulations and probability perturbation method. *Geophysics*. 2012;77(4):M53-M72. doi: 10.1190/geo2011-0417.1
8. Kemper M, Gunning J. Joint impedance and facies inversion-seismic inversion redefined. *First Break*. 2014;32(9):89-95. doi: 10.3997/1365-2397.32.9.77968
9. Pyrcz MJ, Deutsch CV. *Geostatistical Reservoir Modeling*. Oxford: Oxford University Press; 2014.

10. Slatt RM. Geologic controls on reservoir quality. *Dev Petrol Sci.* 2013;61:229-281.
doi: 10.1016/B978-0-444-56365-1.00006-7
11. Petrovic A, Aigner T, Pontiggia M. Facies heterogeneities in a ramp carbonate reservoir analogue: A new high-resolution approach for 3D facies modelling. *J Petrol Geol.* 2018;41(2):155-174.
doi: 10.1111/jpg.12698
12. Slatt RM. Basic principles and applications of reservoir characterization. *Dev Petrol Sci.* 2013;61:1-38.
doi: 10.1016/B978-0-444-56365-1.00001-8
13. Strebelle S. Conditional simulation of complex geological structures using multiple-point statistics. *Math Geol.* 2002;34:1-21.
doi: 10.1023/A:1014009426274
14. Roksandić MM. Seismic facies analysis concepts. *Geophys Prospect.* 1978;26(2):383-398.
doi: 10.1111/j.1365-2478.1978.tb01600.x
15. Dembicki H Jr. *Practical Petroleum Geochemistry for Exploration and Production.* Netherlands: Elsevier; 2017.
doi: 10.1016/C2014-0-03244-3
16. Journel AG. *Fundamentals of Geostatistics in Five Lessons.* Vol. 8. Washington, DC: American Geophysical Union; 1989.
17. Al-Mudhafar WJ. Geostatistical Simulation of Facies and Petrophysical Properties for Hetero-Genicity Modeling in a Tidal Depositional Environment: A Case Study from Upper Shale Member in A Southern Iraqi oil Field. In: *Unconventional Resources Technology Conference, 26-28 July 2021*; 2021. p. 2627-2642.
doi: 10.15530/urtec-2021-5551
18. Journel AG. Nonparametric estimation of spatial distributions. *J Int Assoc Math Geol.* 1983;15:445-468.
doi: 10.1007/BF01031292
19. Strebelle SB, Journel AG. Reservoir Modeling using Multiple-Point Statistics. In: *SPE Annual Technical Conference and Exhibition?*; 2001. p. SPE-71324.
doi: 10.2118/71324-MS
20. Mariethoz G, Caers J. *Multiple-Point Geostatistics: Stochastic Modeling with Training Images.* Wiley-Blackwell; 2014. Available from: https://books.google.com.br/books?hl=es&lr=&id=Fe7sBQAAQBAJ&oi=fnd&pg=PA85&dq=Multiple%E2%80%90Point+Geostatistics:+Stochastic+Modeling+with+Training+Images&ots=P9LR5qUq9B&sig=98MN9Mk5KPtKMGfniK_ECGVnUSQ&redir_esc=y#v=onepage&q=Multiple%E2%80%90Point%20Geostatistics%3A%20Stochastic%20Modeling%20with%20Training%20Images&f=false [Last accessed 2025 Nov 06].
21. Harding A, Strebelle S, Levy M, *et al.* Reservoir facies modelling: New advances in MPS. In: *Geostatistics Banff 2004.* Berlin: Springer; 2004. p. 559-568.
doi: 10.1007/978-1-4020-3610-1_57
22. Zhang TF, Tilke P, Dupont E, Zhu LC, Liang L, Bailey W. Generating geologically realistic 3D reservoir facies models using deep learning of sedimentary architecture with generative adversarial networks. *Pet Sci.* 2019;16(3):541-549.
doi: 10.2523/IPTC-19454-MS
23. Feng R, Grana D, Mosegaard K. Geostatistical facies simulation based on training image using generative networks and gradual deformation. *Math Geosci.* 2025;57:1021-44.
doi: 10.1007/s11004-024-10169-y
24. Tian L, Wang Z, Liu W, Cheng Y, Alsaadi FE, Liu X. A new GAN-based approach to data augmentation and image segmentation for crack detection in thermal imaging tests. *Cognit Comput.* 2021;13(5):1263-1273.
doi: 10.1007/s12559-021-09922-w
25. Goodfellow IJ, Pouget-Abadie J, Mirza M, *et al.* Generative adversarial nets. *Adversarial Neural Information Processing Systems*; 2014. p. 27.
26. Bowles C, Chen L, Guerrero R, *et al.* *Gan Augmentation: Augmenting Training Data using Generative Adversarial Networks.* [arXiv preprint]; 2018.
doi: 10.48550/arXiv.1810.10863
27. Patel M, Wang X, Mao S. Data Augmentation with Conditional GAN for Automatic Modulation Classification. In: *Proceedings of the 2nd ACM Workshop on Wireless Security and Machine Learning*; 2020. p. 31-36.
doi: 10.1145/3395352.3402622
28. Shaham TR, Dekel T, Michaeli T. Singan: Learning a Generative Model from a Single Natural Image. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision.* 2019. p. 4570-4580.
doi: 10.48550/arXiv.1905.01164
29. Mosser L, Kimman W, Dramsch J, Purves S, La Fuente Briceño A, Ganssle G. Rapid Seismic Domain Transfer: Seismic Velocity Inversion and Modeling using Deep Generative Neural Networks. In: *80th Eage Conference and Exhibition 2018.* Vol. 2018. p. 1-5.
doi: 10.3997/2214-4609.201800734
30. Creswell A, White T, Dumoulin V, Arulkumaran K, Sengupta B, Bharath AA. Generative adversarial networks: An overview. *IEEE Signal Process Mag.* 2018;35(1):53-65.
doi: 10.1109/MSP.2017.2765202
31. Douzas G, Bacao F. Effective data generation for imbalanced learning using conditional generative adversarial networks. *Expert Syst Appl.* 2018;91:464-471.

- doi: 10.1016/j.eswa.2017.09.030
32. Liu M, Nivlet P, Smith R, BenHasan N, Grana D. Recurrent neural network for seismic reservoir characterization. In: Bhattacharya S, Di H, editors. *Advances in Subsurface Data Analytics*. Ch. 4. Netherlands: Elsevier; 2022. p. 95-116.
doi: 10.1016/B978-0-12-822295-9.00010-8
33. Miele R, Azevedo L, Grana D. Iterative Stochastic Seismic Inversion Method using GAN with Spatiallyadaptive Normalization. In: *Third EAGE Conference on Seismic Inversion*; 2024. p. 1-5.
doi: 10.3997/2214-4609.202438056
34. Gulrajani I, Ahmed F, Arjovsky M, Dumoulin V, Courville AC. Improved training of wasserstein gans. *Advance Neural Information Processing Systems*; 2017. p. 30.
doi: 10.1016/j.cageo.2024.105622
35. Tran TT. Improving variogram reproduction on dense simulation grids. *Comput Geosci*. 1994;20(7-8):1161-1168.
doi: 10.1016/0098-3004(94)90069-8
36. Cox Michael AA, Cox TF. Multidimensional scaling. In: *Handbook of Data Visualization*. Berlin: Springer Berlin Heidelberg; 2008. p. 315-347.
doi: 10.1007/978-3-540-33037-0_14
37. Miele R, Levy S, Linde N, Soares A, Azevedo L. Deep generative networks for multivariate fullstack seismic data inversion using inverse autoregressive flows. *Comput Geosci*. 2024;188:105622.
doi: 10.1016/j.cageo.2024.105622