Universidad Industrial de Santander

Master's Thesis

Generalization of post-stack seismic image enhancement through data augmentation using deep generative models

Author:
Geraldine Andrea Ariza Serrano

Advisor:
PhD. (c) Yesid Paul Goyes Peñafiel

Co-Advisor:
PhD. José David Sanabria Gómez

A thesis presented in partial fulfillment of the requirements for the degree of Master of Science in Geophysics

Universidad Industrial de Santander
Facultad de Ciencias
Escuela de Física
Maestría en Geofísica
Bucaramanga

2025

Dedicatoria

En memoria de mi nona Hermelinda,
por ser mi sostén durante toda mi vida.
Sus oraciones son hoy una realidad
y su amor me acompaña en este logro.
1934 - 2025

Generalización del mejoramiento de imágenes sísmicas post-apilado a partir del aumento de datos usando modelos generativos profundos

RESUMEN

Las imágenes sísmicas post-apiladas suelen contener ruido y distorsiones estructurales que dificultan la interpretación precisa del subsuelo, especialmente en contextos geológicamente complejos. Los métodos de mejoramiento sísmico incluyen técnicas clásicas de filtrado, que dependen del conocimiento de expertos y pueden comprometer la fidelidad estructural, así como enfoques de aprendizaje supervisado, más flexibles pero dependientes de datos representativos y etiquetados. Aunque estos métodos han mostrado resultados prometedores, persisten limitaciones en su capacidad de generalización cuando se aplican a datos de campo, debido a la discrepancia entre los conjuntos de entrenamiento y la variabilidad real del subsuelo. Para superar estas restricciones, se propone una metodología generativa–supervisada que integra síntesis estructural con entrenamiento dinámico. El enfoque aprovecha cinco complejidades estructurales observadas en cuencas emergentes colombianas (plano, inclinado, plegado, diapírico y fallado) para entrenar un generador progresivo condicional y un autoencoder variacional, con los cuales se producen parches limpios que sirven como referencia en la tarea de mejoramiento supervisado. Previo a esta tarea, los parches se degradan dinámicamente mediante operadores sintéticos, generando en cada iteración pares limpios–ruidosos diversos para robustecer el entrenamiento. Experimentos con datos de campo muestran que nuestro método atenúa el ruido de manera efectiva y preserva la geometría estructural, demostrando capacidad de generalización incluso cuando el entrenamiento se realiza exclusivamente con datos sintéticos.

Palabras clave: datos sísmicos, modelos generativos, geología estructural, aumento de datos, atenuación de ruido sísmico, aprendizaje supervisado.

OBJETIVOS

Objetivo general

Generalizar la tarea de mejoramiento de imágenes sísmicas post-apilado mediante el aumento de datos usando modelos generativos profundos que integren información de geología estructural dentro de un esquema de aprendizaje supervisado.

Objetivos específicos

1. Definir el conjunto de características estructurales y patrones presentes en las imágenes sísmicas post-apilado.
2. Implementar una red neuronal para la generación de imágenes sísmicas post-apilado considerando información de geología estructural.
3. Desarrollar un esquema de aprendizaje supervisado para resolver la tarea de mejoramiento de datos sísmicos integrando el modelo generativo de imágenes sísmicas post-apilado.
4. Validar la capacidad de generalización de la tarea de mejoramiento de imágenes sísmicas con experimentos de atenuación de ruido y corrección de daños estructurales.

PREGUNTA DE INVESTIGACIÓN

¿Cómo puede la integración de modelos generativos profundos que incorporan información de geología estructural mejorar la capacidad de generalización en la tarea de mejoramiento de imágenes sísmicas post-apilado, específicamente en la atenuación de ruido y la corrección de daños estructurales?

Generalization of post-stack seismic image enhancement through data augmentation using deep generative models

Geraldine Ariza-Serrano*, Paul Goyes-Peñafiel†, José Sanabria-Gómez‡

ABSTRACT

Post-stack seismic images often contain noise and structural distortions that hinder accurate subsurface interpretation, particularly in geologically complex settings. Seismic enhancement methods include classical filtering techniques, which rely on expert knowledge and may compromise structural fidelity, as well as supervised learning approaches, which are more flexible but depend on labeled representative data. Although these methods have shown promising results, their generalization capability remains limited when applied to field data, mainly due to discrepancies between training datasets and real subsurface variability. To address these limitations, we propose a supervised–generative methodology that integrates structural synthesis with dynamic training. The approach leverages five structural complexities observed in Colombian emerging basins (planar, dipping, folded, diapiric, and faulted) to train a conditional progressive generator and a variational autoencoder, producing clean patches that serve as references for the supervised enhancement task. Before this task, the patches are dynamically corrupted using synthetic degradations, thereby generating diverse clean–noisy pairs in each iteration. Experiments with field data demonstrate that our method effectively attenuates noise while preserving structural geometry, confirming its generalization capability even when training is performed exclusively with synthetic data.

INTRODUCTION

The quality of seismic data largely depends on the signal-to-noise ratio (SNR) present in the recordings, which may be affected by various sources of noise, including environmental factors (e.g., wind or traffic), instrumental interference, and electrical noise (Zhu et al., 2018). These unwanted signals can obscure important geological features, such as faults and folds, which are crucial for subsurface interpretation and hydrocarbon exploration (Wang et al., 2017).

To obtain reliable seismic images, raw field data must undergo a series of processing steps aimed at enhancing the signal and attenuating noise. This process, known as seismic data processing, includes data editing, amplitude recovery, deconvolution, filtering, normal moveout (NMO) correction, stacking, and migration (Alsadi, 2017; Liner, 2016). Among these steps, stacking of common midpoint (CMP) traces is particularly effective, as it increases the SNR by a factor of $\sqrt{N}$, where N is the number of stacked traces (Rashed, 2014). The output of this procedure is a so-called post-stack seismic image, which represents the subsurface after the stacking operation has enhanced coherent reflections and attenuated random noise.

Post-stack seismic images, although improved in SNR, may still contain residual noise and structural distortions. Traditional noise removal techniques often rely on spectral filtering, temporal smoothing, or deconvolution; however, these methods are typically dependent on expert knowledge and require careful parameter tuning, in addition to being time-consuming (Cheng et al., 2023; Al-Heety and Thabit, 2022). To overcome these challenges, deep learning-based methods—especially convolutional neural networks (CNNs)—have been proposed (Liu et al., 2023). CNNs are particularly well-suited for image processing tasks because they can learn to extract meaningful patterns from data (Zhang et al., 2021). These networks are typically trained within a supervised learning framework, which requires paired data, where noisy images serve as input and clean images serve as labels for training. Since clean seismic images are often scarce, due to limited public availability and data confidentiality, generative models have been employed to synthesize new examples based on a small set of available data (Goodfellow et al., 2020; Creswell et al., 2018; Marano et al., 2024; Oliveira et al., 2019; Grijalva et al., 2021). Meanwhile, degraded images are created by applying various noise operators to clean

*Department of Geophysics, Universidad Industrial de Santander, Bucaramanga, Colombia E-mail: geralariza99@gmail.com
† **Paul Goyes-Peñafiel (advisor)**. Department of Computer Science, Universidad Industrial de Santander, Colombia E-mail: goyes.yesid@gmail.com
‡ **José Sanabria-Gómez (co-advisor)**. Department of Geophysics, Universidad Industrial de Santander, Colombia E-mail: jsanabri@uis.edu.co

post-stack images (Fogat et al., 2023), producing synthetic training pairs that mimic field conditions and enable the network to learn effective denoising mapping (Birnie and Ravasi, 2024).

While machine learning and deep learning have demonstrated exemplary performance in seismic applications, their ability to generalize beyond the training domain remains a common challenge. Generalization refers to a model's ability to perform well on unseen data. In seismic applications, models are often trained on limited datasets, frequently restricted to specific survey areas or geological settings, which increases the risk of overfitting. As a result, when applied to new data with different structural patterns or acquisition conditions, the models may not achieve the same level of accuracy, potentially producing artifacts or suboptimal results (Park et al., 2025; Choi et al., 2021). Recent studies suggest that incorporating explicit structural information—for instance, through geological priors or class-conditional generative models—can help improve generalization and enable networks to adapt to a broader variety of subsurface scenarios (Wang et al., 2022).

To address this challenge and enable generalization to field data, this study proposes a framework for post-stack seismic image enhancement using deep learning models trained on synthetically generated datasets enriched with structural geological information. In this study, five structural complexity classes (planar, dip, folded, diapiric, and faulted) were defined based on observations from emerging Colombian sedimentary basins and used to guide the generative and supervised components of the framework. Clean synthetic images are first generated based on these structural classes and then degraded with multiple noise operators to produce paired training samples. The enhancement network learns a supervised mapping between degraded and clean patches, enabling denoising and structural recovery. Trained entirely on synthetic images capturing diverse structural patterns, the network improves its generalization capability and performs effectively on field seismic data. Therefore, this framework provides a systematic approach for post-stack seismic image enhancement while maintaining adaptability to geologically complex exploration environments.

THEORETICAL FRAMEWORK

Post-stack seismic images

The seismic method is an active technique that exploits elastic waves to investigate subsurface characteristics. A controlled source generates seismic waves that propagate through the ground, where they are reflected by subsurface layers and subsequently recorded by instruments known as geophones, which measure the resulting ground deformations. Seismic reflection analysis focuses on these reflected waves to characterize the structure of geological layers, with the recorded ground deformation over time forming a time series known as a seismic trace. The propagation path of the waves is determined by the physical properties of the subsurface materials (Liner, 2016; Dentith and Mudge, 2014).

The data acquired through this method requires specific processing steps to interpret the subsurface response properly. This includes applying static corrections to adjust variations in travel times and applying filters to eliminate unwanted features of the seismic signals. Subsequently, a stacking process is performed to convert the data into a zero-offset format and sum the seismic traces, which provides information on subsurface velocities while suppressing unwanted noise and non-reflected arrivals (Liner, 2016). As illustrated in Figure 1, stacking can cancel out some random noise, resulting in a clearer seismic signal.

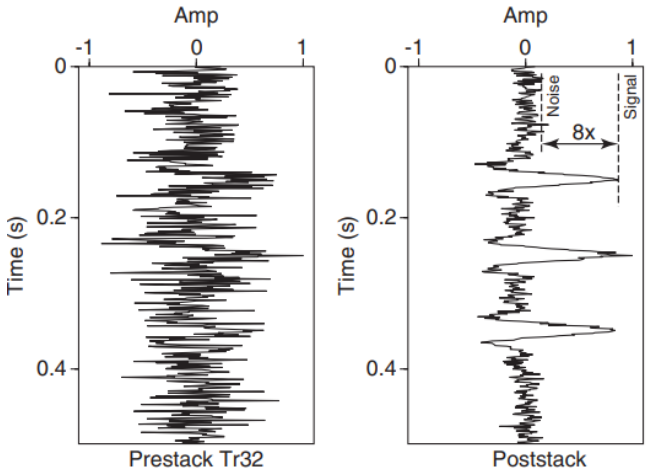

Figure 1: Example of stacking for random noise cancellation (Liner, 2016).

After the stacking process, it is common to represent the results using post-stack seismic images (Oren and Nowack, 2018). These images provide a representation of the stratigraphic and structural properties of the subsurface. This allows geological layers and other relevant subsurface features to be more clearly observed and analyzed, facilitating the interpretation of seismic information (Barnes, 2016; Du et al., 2022). For example, Figure 2 shows a seismic section and its equivalent interpreted geological model, where different faults can be identified as geological structures.

Seismic images obtained after stacking often contain various types of noise introduced during acquisition, which were not removed in previous processing steps. This noise originates from sources such as wind, traffic, electrical interference, and instrumentation (Zhu et al., 2018). In general, it can be classified as coherent or random noise (Fogat et al., 2023), each affecting the data differently. Coherent noise, which shows continuity across multiple seismic traces, includes direct arrivals, refractions, diffractions, surface waves, and multiples, and can obscure true

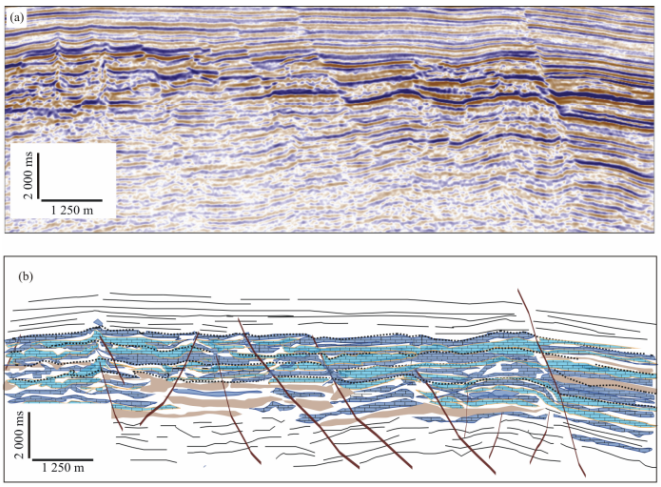

Figure 2: (a) Seismic image and (b) its interpreted geological model, highlighting the presence of faults (Jamaludin et al., 2018).

subsurface structures if not removed. Random noise, on the other hand, lacks a consistent spatial pattern and disrupts the structural coherence of reflectors unpredictably (Bormann and Wielandt, 2013). Therefore, after stacking, additional processing is carried out to remove this noise and clearly define seismic structures.

Post-stack seismic image enhancement

Post-stack seismic images often contain noise that hinders the interpretation of subsurface features. Traditional enhancement techniques, such as filtering and deconvolution, aim to reduce noise and improve structural continuity but typically require manual parameter tuning and may fail in complex geological settings (Ewida and Sarhan, 2023a). Deep learning has emerged as a promising alternative, leveraging neural networks to automate the enhancement process by attenuating the present noise. The success of these deep learning methods depends heavily on the availability of sufficient and representative training data—a significant limitation in seismic imaging, where high-quality post-stack datasets are expensive to acquire, often not publicly available, and consequently, only a limited number of examples can be utilized (Sheng et al., 2025). To overcome this limitation, state-of-the-art approaches employ generative neural networks, such as variational autoencoders (VAEs) and generative adversarial networks (GANs), which learn the underlying data distribution and can generate realistic synthetic seismic images consistent with the observed data.

Generative models for data augmentation

The scarcity of high-quality post-stack seismic data, which are costly to acquire and often not publicly available, limits the amount of data that can be used for training deep learning models. Data augmentation strategies are, therefore, essential to increase the diversity and volume of training examples. Generative models, such as variational autoencoders (VAEs) and generative adversarial networks (GANs), are particularly well-suited for this task, as they learn the underlying distribution of the training data and can produce synthetic samples that realistically mimic the characteristics of the original dataset (Suzuki and Matsuo, 2022).

Variational Autoencoders (VAEs) enable the extraction of a latent representation from input data, allowing the generation of new samples with similar characteristics (Hou et al., 2016; Akkem et al., 2024). Unlike traditional autoencoders that merely compress and reconstruct data, VAEs learn the underlying data distribution, enabling probabilistic sampling from the latent space to generate new instances (Ko et al., 2024). Specifically, the encoder maps an input $\mathbf{x}$ to a latent vector $\mathbf{z}$, parameterized by ϕ, and $\mathbf{z}$ is sampled from a Gaussian distribution defined by these parameters. The decoder then reconstructs the output $\mathbf{x}'$ from $\mathbf{z}$, parameterized by θ, approximating the original data distribution. This variational inference framework allows VAEs to produce diverse samples, avoid mode collapse common in GANs, and efficiently cover the data distribution. However, the generated outputs may be slightly less sharp (Kebaili et al., 2023). Figure 3 illustrates this encoding-decoding process and the probabilistic nature of the latent space.

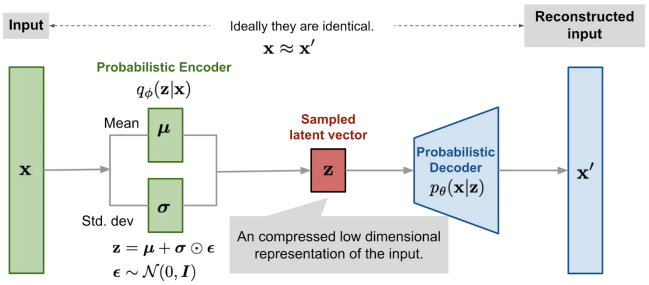

Figure 3: Schematic of a Variational Autoencoder (VAE) (Zhu, 2020).

GANs, in contrast, consist of two neural networks—a generator and a discriminator—that are trained interactively. The generator aims to produce synthetic data capable of deceiving the discriminator, while the discriminator strives to distinguish real from generated data. This zero-sum game continues until the discriminator can no longer reliably distinguish between real and synthetic samples (Creswell et al., 2018; Li et al., 2020). GANs have proven effective for generating high-resolution synthetic data, increasing both the diversity and quantity of samples for model training. This has been particularly beneficial in seismic image processing, where GAN-generated data enrich limited datasets (Goodfellow et al., 2020; Marano

et al., 2024; Oliveira et al., 2019).

Variants such as Progressive Growing GANs (PGGAN) progressively increase network resolution during training, allowing the model first to learn general patterns and then refine details (Karras et al., 2017). The PGGAN generates images from a latent vector $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ where $\mathbf{I}$ is the identity matrix, the synthetic image can be generated by:

$$\mathbf{X} = G(\mathbf{z}), \tag{1}$$

where G is the progressively trained generator network, increasing the resolution from $K \times L$ to $M \times N$ (Torres-Quintero et al., 2025). There are cases where these networks can be conditioned. For example, Conditional PGGAN extends the standard PGGAN framework by incorporating additional information $\mathbf{c}$, which can guide the image generation process. This condition can represent class labels or any auxiliary information (Golfe et al., 2023). The generator then produces images based on both a latent vector $\mathbf{z}$ and the condition $\mathbf{c}$:

$$\mathbf{X} = G(\mathbf{z}|\mathbf{c}). \tag{2}$$

Degradation process

In traditional seismic data processing workflows, degradation modules are commonly employed to emulate acquisition and processing artifacts. These modules apply noise operators to clean seismic images, generating corrupted counterparts that are then paired with the original clean images for supervised training. This process enables enhancement networks to learn how to suppress noise and restore structural continuity. Reported strategies incorporate up to sixteen degradation functions, such as Gaussian noise, impulse noise, speckle, and wave-like patterns, defined from the analysis of field post-stack seismic images (Torres-Quintero et al., 2025). These operators are specifically designed to reproduce both random and coherent noise typically encountered in seismic datasets (Fogat et al., 2023; Al-Shuhail and Al-Dossary, 2020).

This degradation process can be described mathematically as:

$$\mathbf{Y} = P(\mathbf{X}), \tag{3}$$

where $\mathbf{X}$ denotes the clean seismic image, P is a corruption operator, and $\mathbf{Y}$ represents the degraded image. This formulation corresponds to the forward problem, in which synthetic degradation simulates realistic acquisition or processing artifacts. Figure 4 illustrates this process, where clean seismic images ($\mathbf{X}$) are corrupted by applying degradation operators, producing degraded outputs ($\mathbf{Y}$).

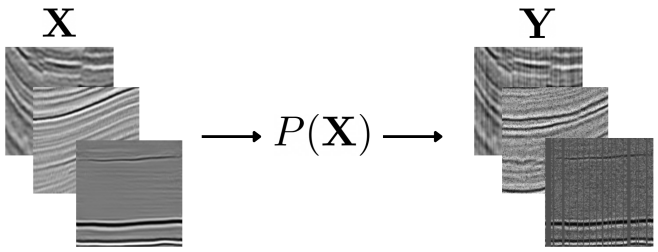

Figure 4: Forward problem applied to seismic images, where a degradation operator P simulates noise and structural distortions. Clean images $\mathbf{X}$ are transformed into degraded images $\mathbf{Y}$.

Supervised Enhancement Framework with CNNs

The enhancement of seismic images through noise attenuation has frequently relied on convolutional neural networks (CNNs), which have demonstrated strong performance in preserving reflector continuity while reducing both coherent and random noise (Li et al., 2022; Wang and Chen, 2019; Yang et al., 2021; Liu et al., 2020). These models follow a supervised learning paradigm, requiring paired data for training. In this context, clean synthetic images are generated using deep generative models. In contrast, the corresponding degraded images are obtained by applying realistic degradation operators, such as Gaussian, Poisson, and speckle noise, to emulate field acquisition conditions (Yuqing et al., 2019; Fogat et al., 2023). This framework has proven effective in seismic image enhancement tasks and has direct implications for hydrocarbon exploration, as clearer seismic sections enable more accurate delineation of subsurface traps and reservoirs.

Building on this approach for generating paired clean–noisy data, a supervised enhancement network can be formulated to learn the mapping between degraded and clean seismic images. Such a network, denoted as $\mathcal{M}_\theta$, is trained using dynamically generated paired data $(\mathbf{X}, \mathbf{Y})$, where $\mathbf{X}$ represents clean synthetic patches and $\mathbf{Y}$ their degraded counterparts, obtained through the degradation model described in Equation 3. The training process aims to optimize the network parameters θ by minimizing an appropriate loss function (Torres-Quintero et al., 2025):

$$\theta^* = \arg\min_\theta \mathcal{L}(\theta), \tag{4}$$

where the cost function $\mathcal{L}(\theta)$ is formally defined as:

$$\mathcal{L}(\theta) = \frac{1}{B} \left\| \mathbf{X} - \mathbf{X}^* \right\|_1 + \left(1 - \mathrm{SSIM}(\mathbf{X}, \mathbf{X}^*)\right). \tag{5}$$

The loss function employed in this work is inspired by methodologies from related studies on seismic data enhancement (Goyes-Peñafiel et al., 2024). It captures the similarity between the original clean seismic data $\mathbf{X}$ and the enhanced data $\mathbf{X}^* = \mathcal{M}_{\theta^*}(\mathbf{Y})$. The function combines the

Mean Absolute Error (MAE) and the Structural Similarity Index (SSIM), striking a balance between pixel-level fidelity and perceptual image quality.

Generalization of the enhancement model

Generalization refers to a model's ability to recognize patterns or features beyond the training data, i.e., to interpret structures not previously observed correctly. In the context of supervised seismic image enhancement, several factors have been shown to influence generalization, including the training strategy, the optimization procedure, the choice of network architecture, and the use of data augmentation to increase the diversity of input samples (Zhu et al., 2020). Table 1 summarizes approaches commonly employed in seismic modeling that have been demonstrated to promote better generalization .

As summarized in Table 1, using data augmentation focused on structural features improves the generalization capabilities of deep learning models applied to seismic image enhancement. By incorporating information about folds or faults and including both synthetic and field data during training, dataset diversity is substantially increased while also controlling the latent space of the generative network. This enables the model to recognize and process a broader range of features, thereby reducing the gap between training and field data. Although overfitting problems may persist because not all structural features are represented, this approach provides a promising solution to improve the generalization capabilities of these models, which is crucial for their practical application in the processing and analysis of complex seismic images.

METHODOLOGY

To address the challenge of limited generalization in deep learning–based seismic image enhancement, this work proposes a methodology aimed at improving performance across different structural settings. The framework integrates deep generative models with supervised learning, relying exclusively on synthetic data. The proposed strategy is structured around three key components:

1. The definition of structural complexity classes based on patterns observed in Colombian emerging basins.
2. The synthetic generation of structurally conditioned clean seismic images using generative models.
3. The application of degradation operators to create corrupted-clean training pairs.

These elements are integrated into a dynamic training framework, where clean synthetic images are generated and degraded on-the-fly while the enhancement network learns to restore them. This workflow enables simultaneous data generation and model training, improving variability and reducing memory usage.

Definition of structural complexities

To ensure that the synthetic dataset reflected realistic geological conditions, this study focused on the structural patterns observed in Colombia's emerging sedimentary basins as the basis for defining the research objectives. These basins are regions where a petroleum system has been recognized, but which remain underexplored or only partially studied, making them promising targets for future research and hydrocarbon development (Ministerio de Minas y Energía and ANH, 2017). According to the National Hydrocarbons Agency (ANH), Colombia's emerging basins include Guajira (GUA), Guajira Offshore (GUAOFF), Cesar–Ranchería (CR), Sinú–San Jacinto (SSJ), and the Lower Magdalena Valley (VIM) as shown in Figure 5.

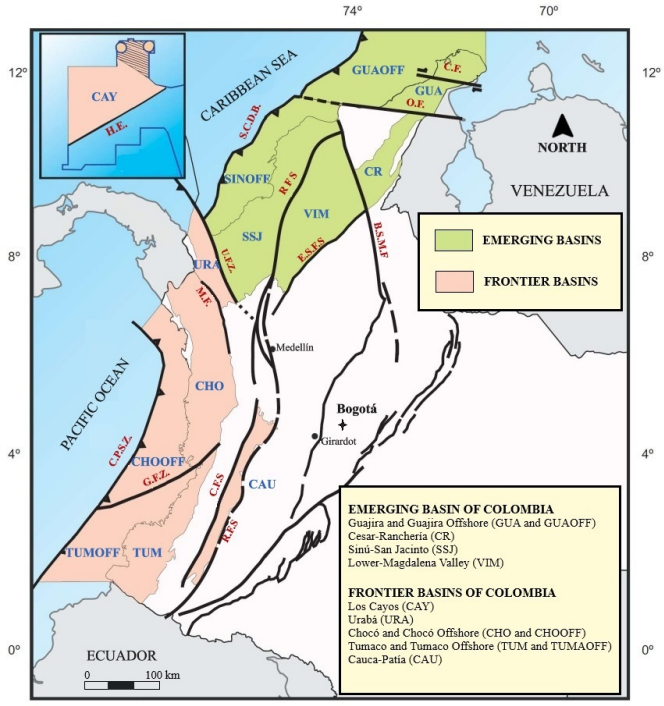

Figure 5: Emerging and frontier basins of Colombia (Ministerio de Minas y Energía and ANH, 2017).

Given the importance of structural characterization in these emerging basins, it is essential to define what is meant by a geological structure. A geological structure is defined as the geometric configuration of rocks in a specific area, emphasizing their geometry, distribution, and formation. In the context of Colombia's emerging basins, the structural features identified were grouped into a set of structural complexities relevant for seismic image analysis, as they directly control the trapping and accumulation of hydrocarbons. Table 2 summarizes these main structural complexities observed in post-stack seismic images, including different types of faulting, folding, and diapirism. To better illustrate these patterns, Figures 6 to 10 present representative seismic examples that were selected to reflect the structural complexities summarized in Table 2,

Table 1: Comparison of proposed approaches for generalization.

Proposed Approach	Advantages	Disadvantages	Reference
A deep learning fault identification method using an attention mechanism to focus on target features	Training the model with the augmented dataset improved robustness, learning capacity, and accuracy in fault identification, thus enhancing generalization.	Requires a representative seismic dataset with fault labels for training; may not capture all variations and complexities present in field seismic data, which can affect performance in extreme cases.	Wang et al. (2023)
A structural data augmentation technique to improve dataset diversity	Incorporating information such as folds and faults, and combining synthetic with field data during training, narrows the gap between training and field datasets, improving seismic data generalization.	Structural data augmentation does not increase waveform diversity or noise types, which may lead to overfitting.	Wang et al. (2022)
GAN-guided supervised learning for seismic data reconstruction	Improves generalization by controlling the generation of seismic patterns from GAN's latent space, reducing dependence on large seismic databases.	Generalization may be limited if field data are not included during network training.	Goyes-Peñafiel et al. (2024)

highlighting key geological features such as normal faults, reverse faults, and mud diapirism.

Table 2: Structural complexities identified in Colombia's emerging sedimentary basins.

Structural Complexity	GUA	GUAOFF	CR	SSJ	VIM
Normal faulting	X	X		X	X
Reverse faults	X	X	X	X	X
Strike-slip faults	X	X		X	X
Mud diapirism	X	X			X
Propagation folds			X	X	
Growth folds				X	
Elongated folds				X	
Inverted faults	X	X	X		
Thrust faults	X	X			

Due to similarities among certain structural complexities, such as normal (Figure 6) and reverse faulting (Figure 7), which both represent layer discontinuities, the identified features were grouped into five generalized structural categories. Although mixed structural patterns are common in field seismic sections (e.g., faults combined with horizontal layers or folds), in this work, each patch was assigned to the category corresponding to its dominant or highest level of complexity. This criterion ensured consistent labeling while still capturing the representative structural information. The resulting categories are:

1. Multiple planar layers
2. Multiple dipping layers
3. Folded layers

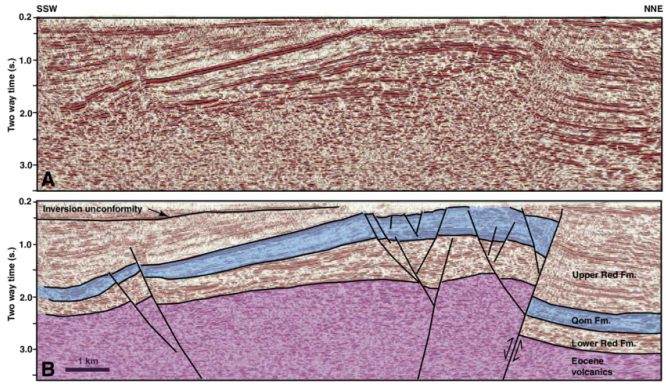

Figure 6: Two-dimensional seismic section across an inverted west-northwest-east-southeast-striking normal or transtensional fault. (A) Uninterpreted seismic line. (B) Interpreted line. (Morley et al., 2009b).

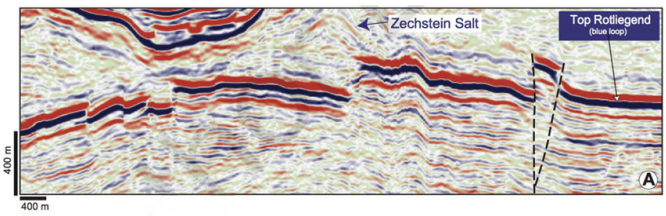

Figure 7: Seismic section showing minor reverse faults (De Jager and Visser, 2017).

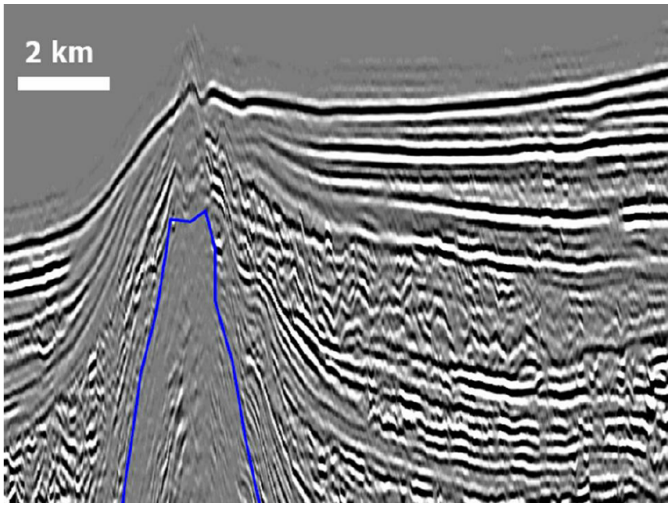

Figure 8: Seismic image of a mud diapiric (Folle, 2008).

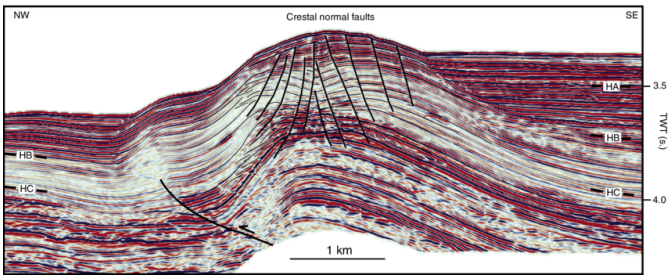

Figure 9: Seismic section showing a fault-propagation fold forming an anticline, accompanied by normal faults at the crest (Morley et al., 2009a).

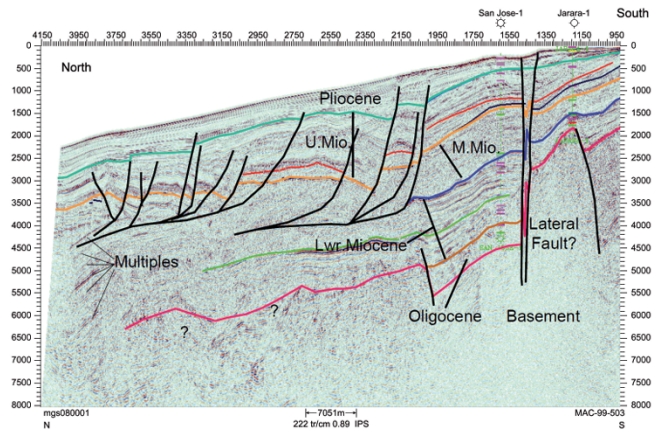

Figure 10: Seismic profile showing growth folds in the northern Guajira basin (Aguilera, 2011).

4. Diapirism

5. Faults

A current limitation in Colombian emerging basins is the scarcity of high-quality, publicly accessible seismic images. Because seismic datasets from these basins are rarely available in open repositories—being subject to commercial confidentiality and restricted access—representative datasets from analogous geological settings were instead employed to emulate the five structural complexities defined above. These datasets were selected to ensure that the extracted patches preserved structural patterns consistent with those observed in Colombian basins. In this study, exclusively synthetic seismic data were used, including TGS Salt Identification (Howard et al., 2018), SEAM Phase 1 (Fehler and Keliher, 2011), 1994 BP Migration from Topography (Sam and Kurt, 1995), and AGL Elastic Marmousi (Martin et al., 2006).

Synthetic image generation

To construct a sufficiently large dataset for the enhancement network, and based on the initial dataset of Colombian structural complexities, three generative models were evaluated: Progressive Growing GAN (PGGAN) (Karras et al., 2017), a variational autoencoder-based GAN (VAE/GAN) (Larsen et al., 2015), and Wasserstein GAN (WGAN) (Arjovsky et al., 2017). Each model was trained on 2,500 images (500 per class), without explicit labels, to assess their ability to capture intrinsic seismic patterns in an unsupervised manner. Following the methodology proposed by Torres-Quintero et al. (2025), the selection of these models was guided by the generative trilemma framework, which considers sampling speed, diversity, and quality. Given that the priority in this study was high-quality and fast image generation, GAN-based architectures were deemed optimal for synthesizing seismic images.

Training followed the main principles of each architecture. For the PGGAN, both generator and discriminator progressively increased in resolution, starting from 32×32 pixels (since smaller resolutions proved uninformative for seismic patches) and growing to 128×128 pixels, which matched the size of the extracted dataset patches. The VAE/GAN combined a variational autoencoder with an adversarial discriminator to enforce perceptual similarity, while the WGAN employed Wasserstein distance to improve training stability and mitigate mode collapse, a common issue where the generator produces only a limited subset of outputs when diversity is low.

Based on the comparative evaluation (see subsection Generative model performance), the PGGAN was identified as the most suitable model for seismic image synthesis and was subsequently adapted to include structural conditioning. Specifically, class labels corresponding to the five generalized structural complexities were integrated into a latent vector and the generator, allowing the model to synthesize images based on the selected class. On the discriminator side, the labels were also integrated and concatenated with the input image, ensuring that both networks were conditioned together during training.

The conditional PGGAN (CPGGAN) followed a progressive training strategy, where both the generator and discriminator increased their resolution throughout training. The process was carried out for a total of 1,300 epochs, distributed across 100 epochs at a resolution of 32×32, 500 at a resolution of 64×64, and 700 at a resolution of 128×128 pixels. Each growth stage consisted of convolutional, normalization, and activation layers, allowing the model to capture global seismic patterns and progressively refine structural details at higher resolutions.

During CPGGAN training, a limitation was observed in the synthesis of images from complexity class 5 (see subsection Generative model performance for a detailed discussion of this behavior). Unlike the other structural classes, the generated fault images frequently lacked well-defined discontinuities, suggesting that the model struggled to capture the structural complexity of fault configurations. To address this, a Variational Autoencoder (VAE) was trained exclusively for that class, using a symmetric encoder–decoder architecture with three convolutional and three transposed-convolutional layers, and a 64-dimensional latent space. Subsequently, both models were integrated into a hybrid generative framework that ensured the generation of all five structural complexities.

To automatically balance the dataset across classes, an iterative procedure was implemented to generate an equal number of synthetic images per category. During inference, the system routes each request by class: non-fault classes are sampled from the CPGGAN. In contrast, the fault class is generated by decoding latent representations extracted from field fault images with the VAE. This hybrid design provides class balancing, class conditioning, and structural realism, while also supplying diversity. This hybrid scheme can be observed in the Figure 11.

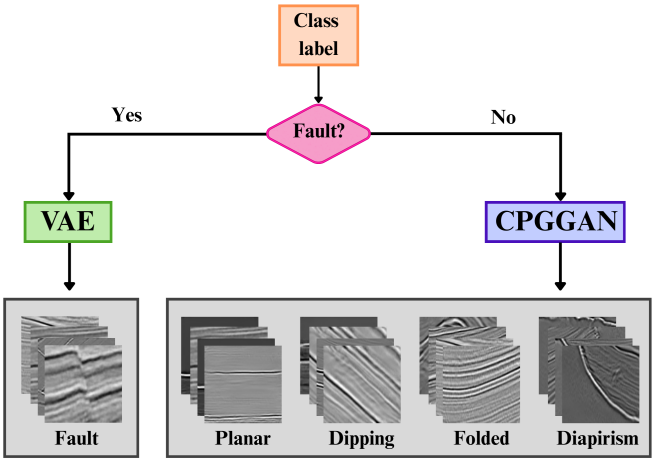

Figure 11: Hybrid framework combining CPGGAN and VAE for seismic image generation.

Degradation of seismic images

To simulate acquisition and processing artifacts typically encountered in field seismic surveys, the cleaned synthetic seismic images were degraded using a set of predefined noise operators. This procedure was designed to emulate random and coherent noise patterns commonly present in post-stack seismic data, such as Gaussian noise, salt-and-pepper noise, acquisition fringes, frequency perturbations, and amplitude distortions.

Sixteen noise types previously developed by Torres-Quintero et al. (2025) were implemented; these noises are listed in Table 3 and were randomly applied during training, ensuring that each cleaned synthetic image produced multiple corrupted versions. While structural balance across classes was explicitly enforced, noise balance was not strictly guaranteed. In each iteration, three operators were randomly chosen and applied independently, followed by an additional corruption generated by combining three different operators. This strategy emphasized variability in the degradation process, exposing the network to a wide spectrum of corruption patterns without enforcing equal representation of each noise type.

Table 3: List of 16 noise related to post-stack seismic data (Torres, 2024).

Noise	Type	Cause
Gaussian	Random	Human activities that generate vibrations affecting the geophones
Poisson	Random	
Speckle	Random	
Salt and pepper	Random	
s_1	Random	
s_1 blur	Random	
Linear	Random	Faulty sensors
Linear v_2	Random	
Stripes	Random	
Blur	Random	Rayleigh-type surface wave that masks relevant seismic events
Waves	Coherent	Over-migration and under-migration of the data
Waves$_2$	Coherent	
Correlated g_1	Coherent	Surface data acquisition geometry
Correlated g_2	Coherent	
Correlated g_{v1}	Coherent	
Correlated g_{v1I}	Coherent	

Supervised enhancement network

To restore the quality of seismic images after applying synthetic degradations, a supervised learning scheme was implemented. The objective of this stage was to train a neural network capable of mapping degraded seismic images to their corresponding clean versions, thereby improving structural continuity and reducing noise artifacts.

For this purpose, a U-Net architecture with attention

gates (AttU-Net) was adopted, given its encoder–decoder structure with skip connections, which are designed to preserve high-frequency details and focus on relevant reflector regions. The encoder consists of five convolutional stages, while the decoder mirrors this hierarchy through upsampling layers. Attention gates modulate the skip connections, enabling the network to prioritize structural patterns such as layer continuity or fault discontinuities during denoising. Further implementation details and a repository ensuring the reproducibility of the neural network are provided in the original work (Oktay et al., 2018).

Unlike the generative models, the denoising network was not explicitly conditioned on structural labels. Instead, class balance was achieved at the dataset level: for each iteration, an equal number of clean patches was generated per structural category (500 per class, five classes, i.e., 2,500 total), ensuring that the model was exposed to all structural complexities during training. Clean samples were obtained from the hybrid generative framework—conditional PGGAN for planar, dipping, folded, and diapiric structures, and a fault-specialized VAE for discontinuities—and subsequently degraded using a custom module of sixteen operators, including Gaussian, Poisson, speckle, salt-and-pepper, and others. Two corruption strategies were applied:

- **Simple degradation**: three noise types were randomly selected and applied independently to the clean sample.
- **Composite degradation**: three noise types were applied consecutively, producing more complexity and realistic distortions.

This procedure ensured that each clean image produced multiple corrupted variants, significantly expanding variability and forcing the model to generalize across diverse noise conditions. Up to 10,000 degraded–clean pairs were generated per iteration, and datasets were dynamically regenerated with new random seeds to guarantee diversity. A total of 10 iterations were performed, each trained for 100 epochs, resulting in a cumulative total of 1000 epochs. By integrating generative models, degradation operators, and the AttU-Net enhancement network, the framework delivered supervised mapping of degraded to clean seismic images. Figure 12 shows a schematic description of how this denoising process would work.

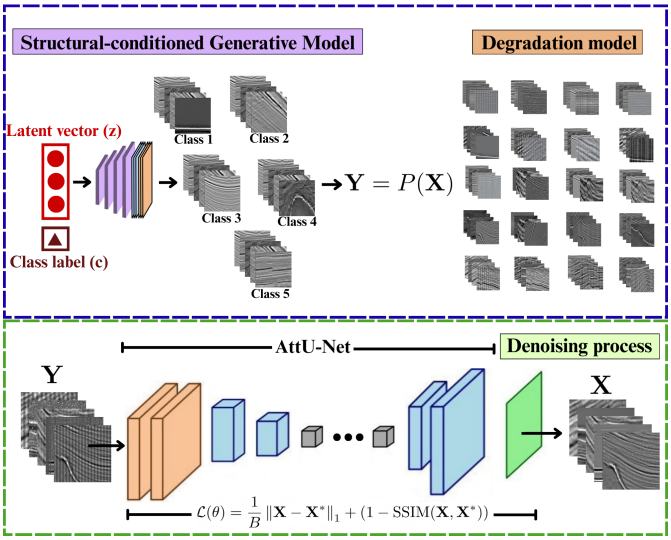

Figure 12: Workflow of the supervised enhancement framework. Clean seismic patches ($\mathbf{X}$) are synthesized using a structural-conditioned generative model, where a latent vector (z) is combined with a structural class label (c) to control the generation process. CPGGAN is used for planar, dipping, folded, and diapiric structures, while a VAE is employed for faults. The clean images are degraded ($\mathbf{Y}$) by a custom module with 16 noise operators. The resulting degraded–clean pairs are then used to train the supervised enhancement network (AttU-Net), which learns to recover structural continuity and attenuate noise.

EXPERIMENTS AND RESULTS

Characterization of structural patterns

As mentioned in the methodology, due to the scarcity of public seismic data from Colombian emerging basins, freely accessible datasets were used to construct the structural dataset. From these sources, patches, each sized 128×128 pixels, were extracted and grouped into five structural categories: planar, dip, fold, diapiric, and faulted. This process resulted in a total of 8,850 clean patches. To ensure structural consistency, corrupted or noisy samples were manually discarded. Exclusion criteria included: (i) patches dominated by random noise with very low signal-to-noise ratio, (ii) sections where acquisition artifacts masked the geological features, and (iii) cases where horizons or discontinuities could not be reliably identified. Representative examples of excluded images are shown in Figure 13, highlighting cases where noise or acquisition artifacts predominated over geological patterns. A summary of the number of patches per class is presented in Table 4.

Table 4: Number of synthetic post-stack seismic patches (128×128) per structural category extracted from publicly available datasets selected based on structural similarity with Colombian emerging basins.

Structural Complexity	Number of Patches
1) Planar layers	539
2) Dipping layers	1,187
3) Folded layers	3,584
4) Diapirism	3,234
5) Faults	306
Total	**8,850**

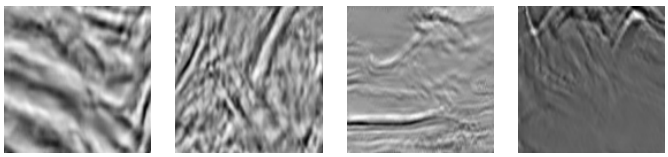

Figure 13: Examples of discarded patches removed from the dataset due to excessive noise, acquisition artifacts, or lack of clear structural information.

In the case of structural complexity 5 (failed), only 153 valid patches could be collected. To increase the representation of this class and preserve geological consistency, a reflection strategy was applied. Figure 14 illustrates this process, which resulted in a total of 306 patches, as shown in Table 4.

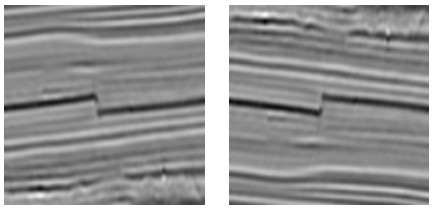

(a) Original patch (b) Reflected patch

Figure 14: Patch augmentation by reflection for fault class. (a) Original patch, and (b) a reflected version. This procedure was applied to increase the number of samples in the fault category while preserving geological consistency.

Figure 15 shows representative examples of the final clustered patches that do not contain noise and illustrate the typical structural features associated with each category. These patterns respect the geological conditions observed in Colombia's emerging sedimentary basins and serve as the basis for generative modeling.

Generative model performance

Following the methodology described above, different generative models were initially evaluated to identify the most suitable for seismic image synthesis. Deep learning models were implemented in Python using the PyTorch library, and all training was performed on an NVIDIA GeForce RTX 3060 GPU with 12 GB of memory. Table 5 summarizes the performance of the reference architectures (PGGAN, WGAN, and VAE/GAN), while Figure 16 presents representative examples of the images generated after training.

The comparative evaluation of different generative architectures revealed that not all models are equally applicable to the various seismic image generation scenarios. Firstly, the PGGAN network showed a remarkable correlation with the patterns present in the training data, which consisted

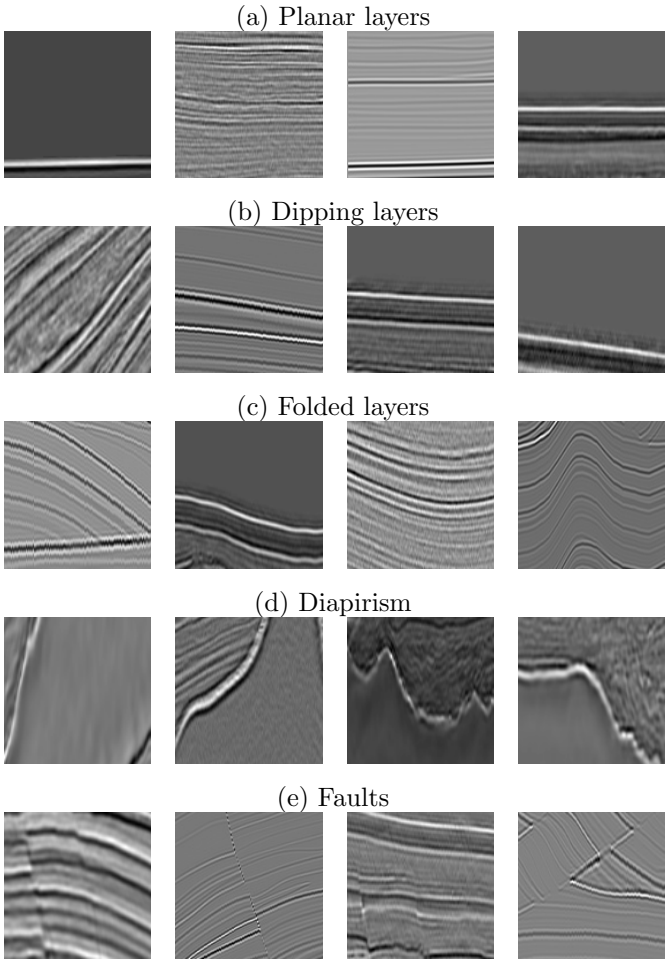

Figure 15: Examples of 128×128 post-stack seismic image patches for each structural category.

Table 5: Comparison of generative models for post-stack seismic patch generation.

Model	Number of Patches	CPU Time (s)	Metric		
			FID	IS	CMMD
PGGAN	1000	112	**62.844**	**3.321±0.110**	**0.871**
VAE/GAN	1000	**30**	228.668	1.921±0.040	1.114
WGAN	1000	264.95	153.662	2.857±0.122	0.872

of clean post-stack seismic images grouped by structural characteristics. This behavior suggests that the progressive resolution growth approach significantly contributes to capturing fine structural details, as previously reported in the literature (Karras et al., 2017). Furthermore, in comparison to other models, PGGAN achieved higher visual quality and better quantitative metrics.

On the other hand, training the VAE/GAN model revealed stagnation in learning capacity, despite increasing the number of epochs; image reconstruction became more computationally expensive and structurally less coherent. Finally, although the WGAN model was able to generate samples with certain similarities to the input data, its

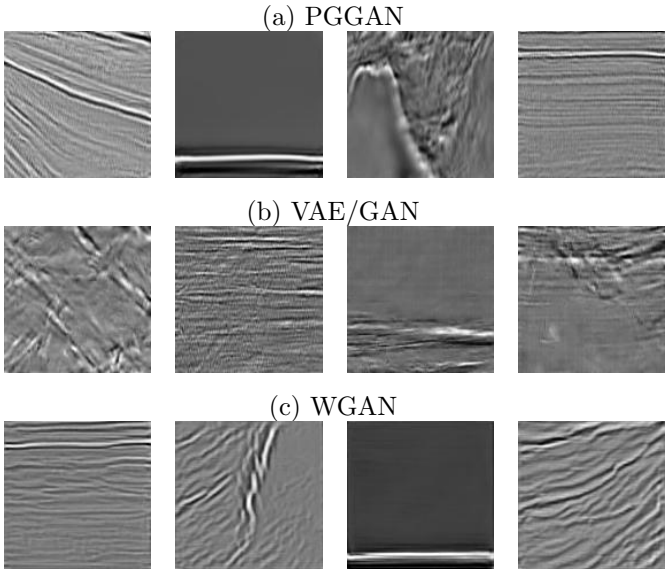

Figure 16: Representative 128×128 post-stack seismic image patches from the dataset compiled in this work, illustrating each of the structural categories.

performance was inferior in terms of structural fidelity and visual realism when compared to PGGAN.

Based on the results obtained, the PGGAN model was selected as the primary generative architecture due to its superior performance in terms of visual fidelity, preservation of structural patterns, and quantitative metrics. Subsequently, the network was conditioned to learn the predefined structural complexity classes. Training was performed using five structural categories, with the number of post-stack seismic image patches per class summarized in Table 4. The generative performance of each structural category is illustrated in Figure 17.

As shown in Figure 17, the conditional PGGAN (CPPGAN) is capable of synthesizing structurally consistent patterns aligned with planar, dipping, folded, and diapiric classes. However, for the faulted class, the model failed to reproduce the expected discontinuities. This limitation is attributed to the fact that, when training begins at a low resolution, the distinctive characteristics of faults—mainly abrupt reflector interruptions—tend to be lost. As a result, the network fails to capture this pattern adequately and, in later stages of training, tends to misinterpret it as continuous geometries such as folds. As shown in Figure 18, reducing the original image to a lower dimension leads the network to perceive the discontinuity as a feature similar to other complexities, resulting in an incorrect structural representation.

To address this limitation, a specialized generative model based on a Variational Autoencoder (VAE) was trained, focusing exclusively on the fault class. As illustrated in Figure 19, the VAE was able to more consistently represent

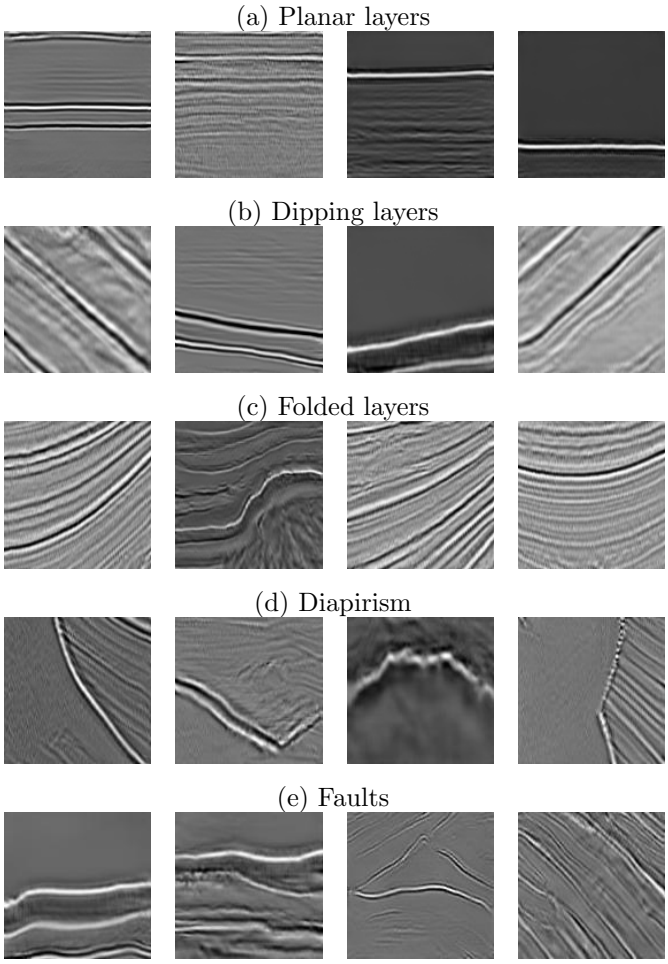

Figure 17: Examples of synthetic post-stack seismic images generated for each structural category using the conditional PGGAN.

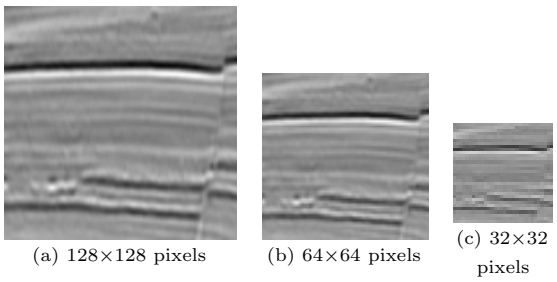

Figure 18: Progressive loss of resolution in seismic faults: (a) original patch at 128×128 pixels, (b) downscale to 64×64 pixels, and (c) further reduced to 32×32 pixels.

the discontinuous geometry of faults, demonstrating its effectiveness as a complementary model to CPGGAN for the synthetic generation of this specific structural category.

In this way, both generative networks were integrated into a conditional framework, resulting in a hybrid strategy that guaranteed the generation of all five structural complexity classes. Qualitative assessment confirmed the

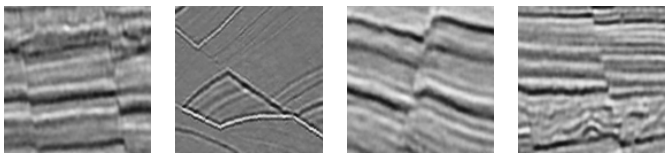

Figure 19: Synthetic post-stack seismic images generated by the VAE trained exclusively on faulted structures (complexity 5).

structural similarity between the generated and training images, while quantitative evaluation was performed using the Fréchet Inception Distance (FID). As summarized in Table 6, the results reflect a correlation between structural complexity and model performance: lower FID values were obtained for simpler geometries (planar: 52.59, folded: 52.78), indicating a closer match between generated and field data distributions and thus higher visual fidelity. In contrast, faults exhibited the highest FID score (92.25), which is attributed to their geometric heterogeneity and potential overlap with features of other classes, highlighting the increasing challenge of reproducing discontinuous seismic features.

Table 6: Fréchet Inception Distance (FID) per structural category.

Structural Complexity	FID Score
1) Planar layers	52.592
2) Dipping layers	97.116
3) Folded layers	52.782
4) Diapirism	64.788
5) Faults	92.253

Moreover, the results obtained support the concept of the generative model trilemma, which poses three fundamental objectives: sample quality, sample diversity, and sampling rate as illustrated in Figure 20 (Xiao et al., 2021). In this context, it is observed that GANs excel in generating high-quality visual images, while VAEs excel in offering greater structural variability.

Complementing the quantitative analysis based on FID, a t-SNE (t-Distributed Stochastic Neighbor Embedding) projection was used to map the data distribution in a low-dimensional space (Costa et al., 2021). As shown in Figure 21, the results reveal coherent clustering by structural class: compact clusters with near-complete overlap between real (•) and generated (×) samples for simpler categories (e.g., planar), and a multimodal distribution for faults, consistent with their intrinsic variability.

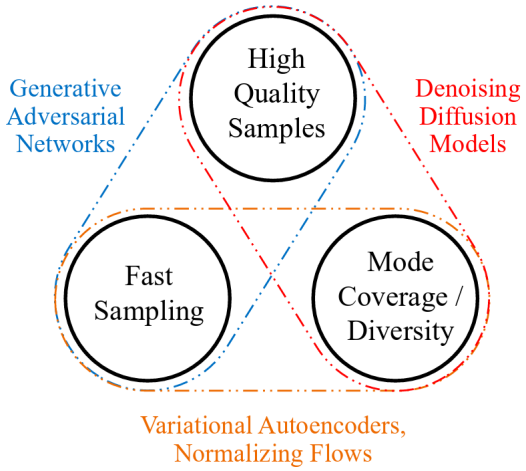

Figure 20: Generative learning trilemma (Xiao et al., 2021).

Supervised enhancement scheme

Denoising is a crucial step in seismic image processing, as the quality of geological interpretation depends mainly on the fidelity of the signal. However, in many cases, noise not only degrades the image but also directly interferes with the continuity of seismic events, generating visual artifacts and even false discontinuities that can lead to misinterpretations. In this context, effective noise attenuation processes not only improve the visual clarity of the image but also facilitate the accurate identification of key subsurface structures (Chirtu and Radoi, 2022).

To simulate realistic field acquisition and processing conditions, a synthetic degradation module composed of sixteen noise operators was implemented, previously developed by Torres-Quintero et al. (2025) from the analysis of field post-stack seismic images. These operators were designed to emulate both random noise and Gaussian or impulsive noise, as well as coherent artifacts such as wave-like patterns. These noise types are summarized in Table 3, and representative examples of these degradation types are shown in Figure 22.

Since the generative models employed in this study produce clean images by default, these outputs were integrated into the degradation workflow to automatically create the paired clean–damaged samples required for training the enhancement network. To implement the proposed workflow, a neural network based on the U-Net architecture was deployed to dynamically learn supervised mappings between degraded seismic images and their corresponding clean versions. One advantage of this dynamic training strategy is the significant reduction in memory consumption. Because degraded images are generated on-the-fly, they do not need to be pre-stored, thereby avoiding unnecessary disk usage. As the goal is to obtain a large and diverse set of images, multiple samples are required.

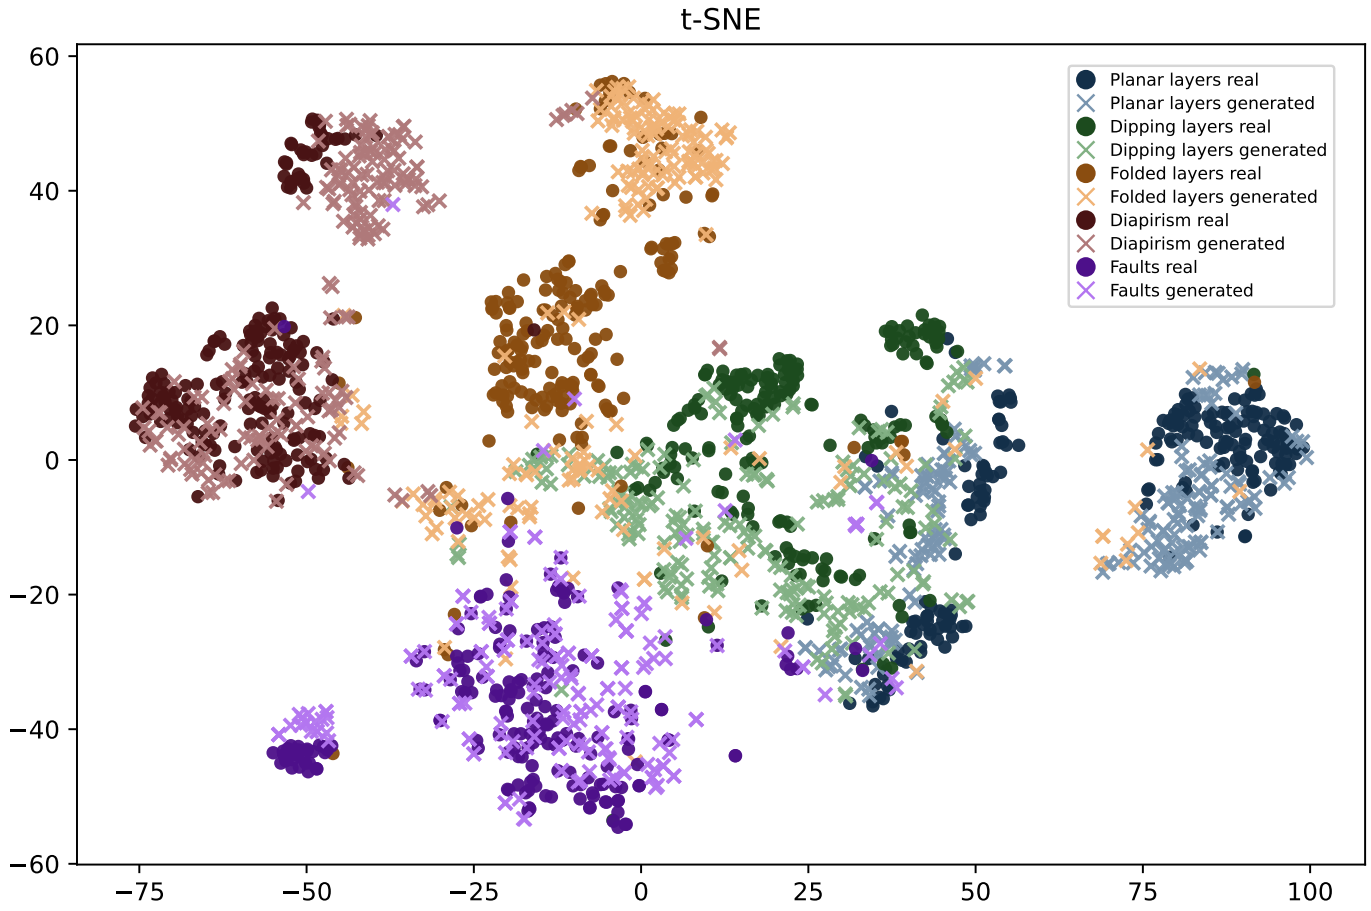

Figure 21: t-SNE of real (•) and generated (×) seismic images.

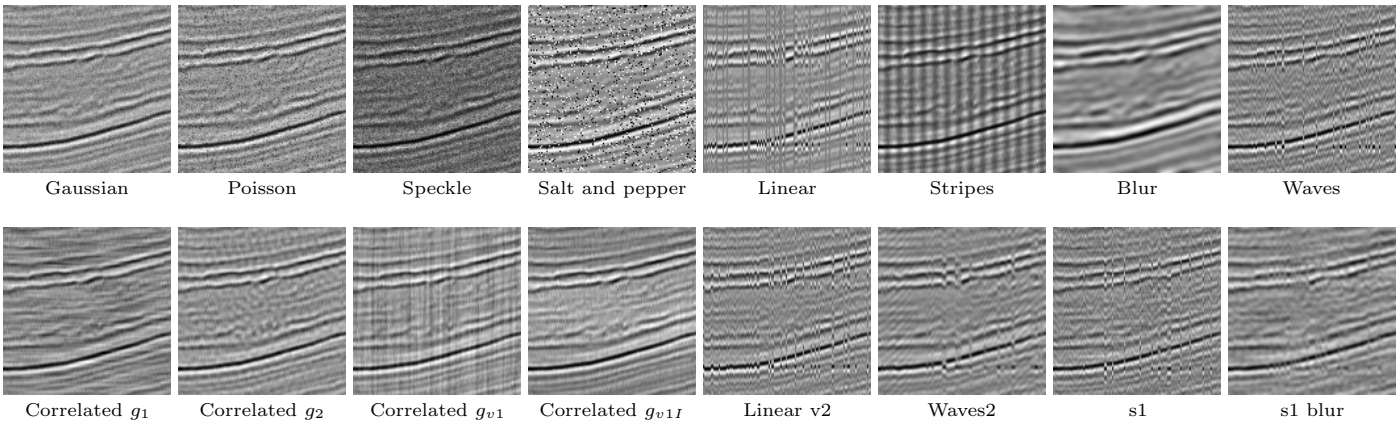

Figure 22: Visual examples of synthetic seismic patches affected by sixteen noise operators, including random and coherent noise.

Additionally, the dynamic generation of new samples at each iteration enhances the diversity of the data to which the network is exposed. Unlike static datasets (where the same examples are reused in every epoch), this strategy ensures greater pattern variability. This continuous variation acts as an inherent data augmentation process that broadens the exploration of the input space, allowing the model to more effectively capture the underlying distribution of seismic data and reducing the risk of overfitting to specific examples. Previous works (Cerqueira et al., 2024; Yang et al., 2025) have shown that such strategies positively contribute to the generalization capability of deep learning models, especially in contexts with limited training data or high structural variability. Furthermore, the use of a conditional generator provides explicit control over the

number of images generated for each structural complexity class, ensuring a balanced distribution across the training set. This prevents under-representation or omission of any class during a given iteration, thereby reducing the risk of bias and further reinforcing the model's ability to generalize.

An initial test scenario applied all sixteen noise operators sequentially, generating up to sixteen corrupted variants per clean sample. Although this approach maximized dataset variability, it proved computationally prohibitive, requiring around 18 hours to prepare data for 10 iterations. To overcome this limitation, a more efficient trade-off strategy was adopted. In each iteration, three noise operators were randomly selected and independently applied to each clean image. Additionally, a composite corruption was introduced, consisting of a sequence of three randomly chosen operators. This configuration preserved variability in the training data while keeping computational requirements manageable.

The denoising capability of the network was evaluated on a representative synthetic patch. As shown in Figure 23, the original clean image (a) was corrupted with controlled pepper noise (b). The enhanced output (c) demonstrates the ability of the network to attenuate noise while preserving reflector continuity effectively, and the difference map (d) highlights the removed noise components.

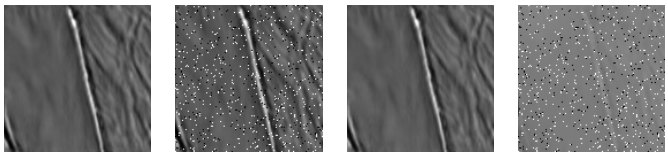

(a) Original (a) Noisy (b) Enhanced (d) Difference

Figure 23: Example of a synthetic seismic patch affected by pepper noise: (a) original clean patch, (b) noisy input with pepper noise, (c) enhanced output, and (d) difference map highlighting the attenuated noise.

In a second experiment, a complete post-stack seismic section was used as a reference to evaluate the model's generalization capability on previously unseen field data. Although this image does not belong to the emerging Colombian basins, it shares similar structural characteristics. It therefore provides a benchmark to validate the model's performance in field scenarios (Keys and Foster, 1998). Specifically, this section corresponds to the Mobil AVO dataset, which was previously analyzed by Torres-Quintero et al. (2025). The seismic line was divided into 128×128 patches, processed through the enhancement network, and subsequently reassembled.

Figure 24 illustrates the comparison between the original noisy seismic section and the enhanced outputs produced by the proposed dynamic framework, alongside the results obtained using the method by Torres-Quintero et al. (2025). The main difference between the two approaches lies in the generative model. While the method by Torres-Quintero et al. (2025) employed a non-conditional generator trained with a mixed dataset of synthetic and field data, our framework relies on a conditional generator trained exclusively on synthetic data, ensuring that both the input patches and the generated images originate from a fully synthetic setting. When analyzing the highlighted areas, it can be observed that the high-frequency noise masking the signal is attenuated by both methods, thus improving reflector visibility. However, in the region marked with the red box, corresponding to a horizontal reflector, the dynamic framework achieves a more effective suppression of deeper spurious events. In contrast, the method proposed by Torres-Quintero et al. (2025) tends to emphasize these signals, as indicated by the blue arrow. Visually, the proposed approach produces smoother transitions and more localized enhancements, which is consistent with expectations for post-stack images, where improvements are generally more subtle compared to pre-stack processing (Ewida and Sarhan, 2023b). Nevertheless, these results demonstrate the model's ability to generalize to field data not seen during training and suggest that the incorporation of structural conditioning strengthens noise attenuation.

Generalization and validation on field data

Generalization refers to the network's ability to enhance images that were not seen during any stage of training. To evaluate this capability, several experiments were conducted using field-acquired seismic images. It is important to emphasize that in this work, the network was trained exclusively with synthetic data; therefore, any successful application to field data inherently validates its generalization capacity, as no field information was introduced during training. The analyzed data correspond to field seismic records from emerging sedimentary basins in Colombia; however, their specific locations are not disclosed due to confidentiality restrictions.

For additional comparison, experiments were also conducted using a static scheme in which synthetic images were pre-generated and stored locally for retrieval during training. For each structural class, 500 clean reference images were generated and subsequently degraded using four different noise operators, producing 2,000 noisy variants. This resulted in a total of 2,500 images per class and 12,500 images across the five classes.

Figures 25–26 illustrate the visual results, comparing (a) the original noisy section, (b) the enhanced output obtained with the proposed dynamic method, (c) the enhanced result from the static scheme, and (d) the baseline method by Torres-Quintero et al. (2025).

In the case shown in Figure 25, a significant improve-

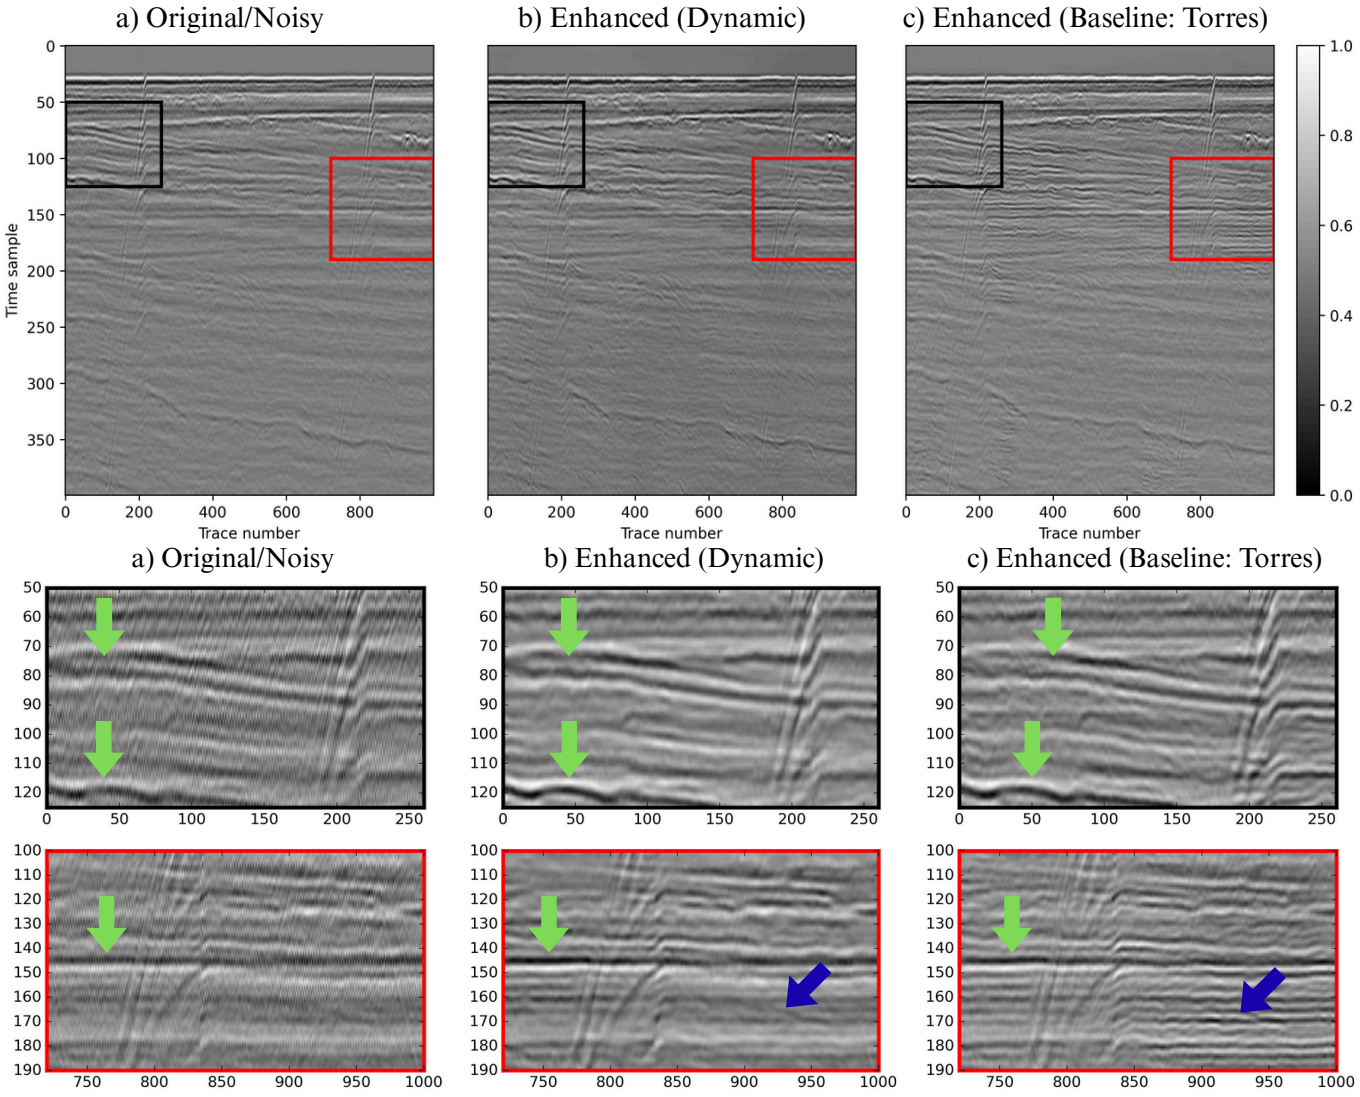

Figure 24: Comparison of seismic enhancement methods. Top row: (a) original noisy section, (b) enhanced section using the dynamic approach, and (c) result obtained with Torres' method. Zoom areas show the most prominent image enhancements.

ment is observed in the central part of the image. When focusing on the selected patch, it becomes evident that all methods attenuate background noise and enhance the main reflector, even revealing events in the lower part of the section that were previously difficult to identify. Both the dynamic and static approaches produce smoother transitions, which are effective in relatively simple structural scenarios. Conversely, Figure 26 presents a case with pronounced distortions in the upper part of the section. In this scenario, the proposed dynamic framework delivers the best enhancement, successfully reconstructing missing data and achieving greater continuity, which the static model does not achieve. While the method by Torres-Quintero et al. (2025) also reconstructs some missing information, it simultaneously introduces background noise, as indicated by the red arrow, thereby adding undesired artifacts. These results highlight that the proposed approach is more effective in handling structurally complex cases.

As an additional test, an image with folded structures and predominantly random noise was evaluated (Figure 27). The proposed dynamic method not only effectively suppressed noise but also preserved the lateral continuity of the reflectors and maintained the geometry of the fold, producing a sharper section. In contrast, the static and Torres-Quintero et al. (2025) methods partially attenuated the noise but tended to introduce artifacts and loss of resolution.

In summary, these findings confirm that a framework based on structural conditioning with purely synthetic data can achieve competitive results when applied to field seismic sections, providing effective enhancement of reflectors while enabling more accurate representations and

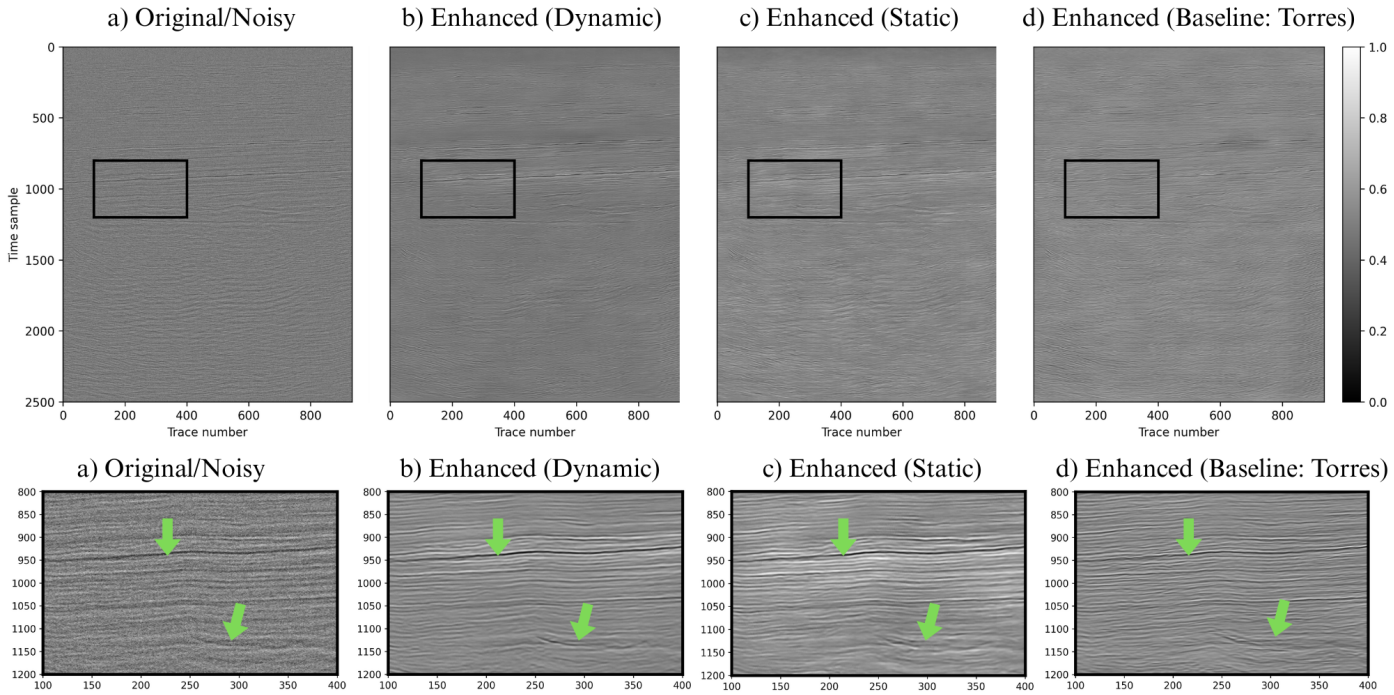

Figure 25: Comparison of seismic enhancement methods. Top row: (a) original noisy section, (b) enhanced section using the dynamic approach, (c) enhanced section using the static approach, and (d) result obtained with the method by Torres-Quintero et al. (2025). Bottom row: zoomed views of the highlighted region, where green arrows mark areas that highlight differences between the methods in terms of structural continuity, noise attenuation, and reflector preservation.

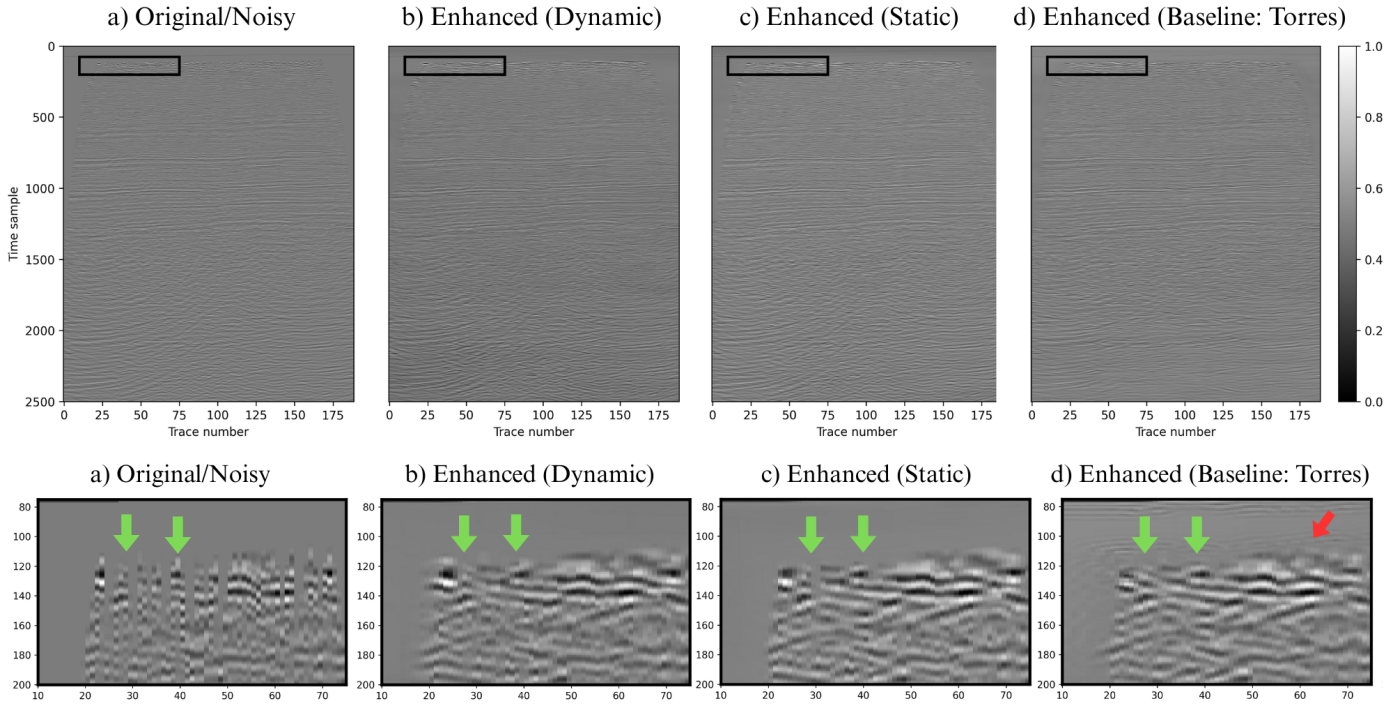

Figure 26: Comparison of seismic enhancement methods. Top row: (a) original noisy section, (b) enhanced section using the dynamic approach, (c) enhanced section using the static approach, and (d) result obtained with the method by Torres-Quintero et al. (2025). Bottom row: zoomed views of the highlighted region, where green arrows mark areas that highlight differences between the methods in terms of structural continuity, noise attenuation, and reflector preservation.

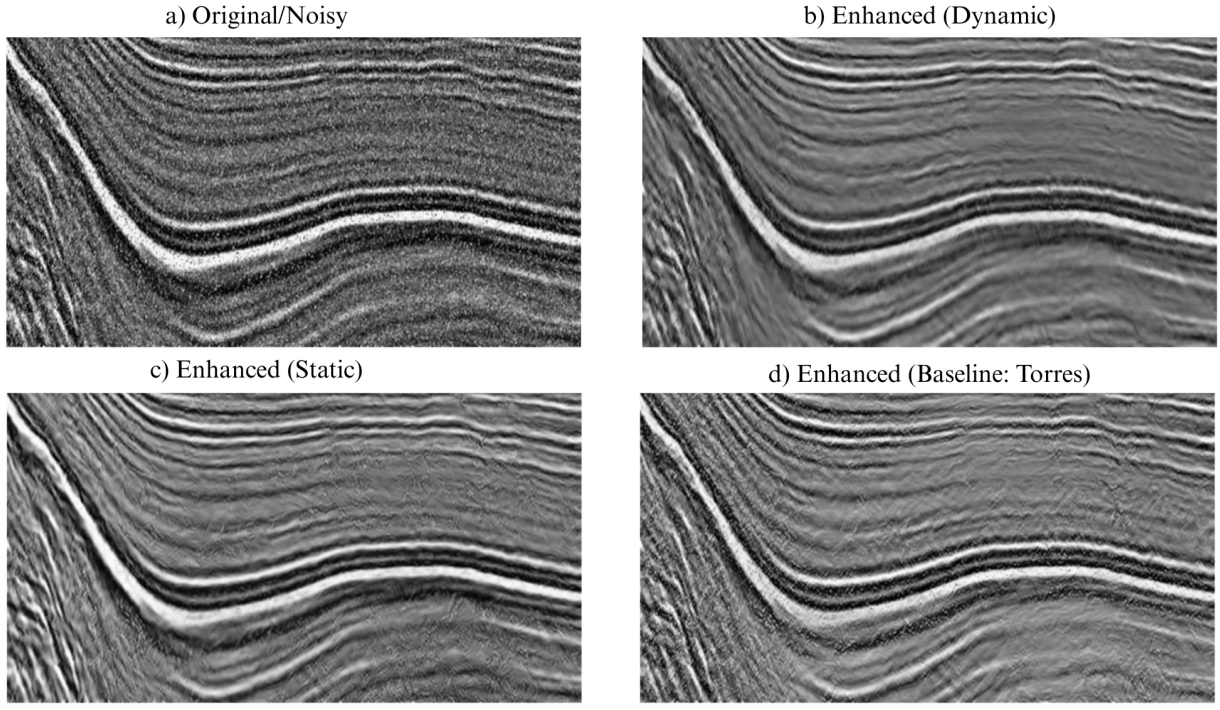

Figure 27: Example of seismic enhancement on field data with predominantly folded structures. (a) Original noisy section, (b) section enhanced with the dynamic approach, (c) section enhanced with the static approach, and (d) result obtained with the Torres-Quintero et al. (2025) method.

supporting more reliable geological interpretations.

CONCLUSIONS

This thesis introduced a data-driven framework for post-stack seismic image enhancement by integrating deep generative models with supervised learning. It shows that a framework trained entirely on structurally conditioned synthetic data can generalize effectively to field seismic sections, offering a robust and geologically coherent solution for seismic image enhancement in complex exploration settings. The four main contributions are:

1. Structural complexities commonly observed in post-stack seismic data (planar, tilted, folded, diapiric, and faulted layers) were systematically identified and organized to construct a representative dataset, providing a geological basis for conditioning generative models and ensuring structural plausibility in synthetic results.

2. A conditional generative network and a variational autoencoder, trained exclusively on synthetic datasets, successfully produced geologically consistent sections while addressing the scarcity of labeled field data. It is remarked that no field information was required during training.

3. The framework combined generative augmentation with supervised learning through dynamic degradation, generating clean–noisy pairs on the fly. This strategy improved robustness to diverse noise conditions, enhanced structural recovery, and reduced storage demands.

4. Validation with Colombian field seismic data demonstrated effective noise attenuation, improved reflector continuity, and smoother structural transitions. Compared with existing methods, the dynamic framework achieved superior results while avoiding artifacts.

ACADEMIC DISSEMINATION

International Conference

1. Ariza, G., Goyes, P., & Sanabria, J. D. (2025). Supervised Seismic Image Enhancement from Structurally Conditioned Synthetic Data. In 2025 XXV Symposium of Image, Signal Processing, and Artificial Vision (STSIVA) (pp. 1–5). 2025 XXV Symposium of Image, Signal Processing, and Artificial Vision (STSIVA). IEEE. https://doi.org/10.1109/stsiva66383.2025.11156479

ACKNOWLEDGMENTS

This work was supported by the Vicerrectoría de Investigación y Extensión from Universidad Industrial de Santander under Projects 4619 and 3925.

DATA AND MATERIALS AVAILABILITY

The implementation of the proposed framework, including the generative models, data degradation operators, and enhancement network, is openly available on GitHub to ensure reproducibility and foster future research: GitHub-seismic-image-enhancement.

REFERENCES

Aguilera, R., 2011, Petroleum geology of Colombia, Guajira and Cayos basins: Technical report, Universidad Eafit-AN.

Akkem, Y., S. K. Biswas, and A. Varanasi, 2024, A comprehensive review of synthetic data generation in smart farming by using variational autoencoder and generative adversarial network: Engineering Applications of Artificial Intelligence, **131**, 107881.

Al-Heety, A. J. R., and H. A. Thabit, 2022, Random and coherent noise attenuation for 2D land seismic reflection line acquired in Iraq: NRIAG Journal of Astronomy and Geophysics, **11**, 337–354.

Al-Shuhail, A., and S. Al-Dossary, 2020, Attenuation of Incoherent Seismic Noise: Springer International Publishing. Advances in Oil and Gas Exploration & Production.

Alsadi, H. N., 2017, Processing of Seismic Reflection Data, *in* Advances in Oil and Gas Exploration and Production: Springer, 245–290.

Arjovsky, M., S. Chintala, and L. Bottou, 2017, Wasserstein GAN.

Barnes, A. E., 2016, Chapter 1: Overview of Poststack Seismic Attributes, *in* Handbook of Poststack Seismic Attributes: Society of Exploration Geophysicists, 1–12.

Birnie, C., and M. Ravasi, 2024, Explainable artificial intelligence-driven mask design for self-supervised seismic denoising: Geophysical Prospecting, **72**, 1729–1744.

Bormann, P., and E. Wielandt, 2013, Seismic signals and noise, *in* New manual of seismological observatory practice 2 (NMSOP2): Deutsches GeoForschungsZentrum GFZ, 1–62.

Cerqueira, V., M. Santos, L. Roque, Y. Baghoussi, and C. Soares, 2024, Online Data Augmentation for Forecasting with Deep Learning.

Cheng, S., Z. Cheng, C. Jiang, W. Mao, and Q. Zhang, 2023, An effective self-supervised learning method for various seismic noise attenuation.

Chirtu, M.-A., and A. Radoi, 2022, Seismic Signal Denoising using U-Net in the Time-Frequency Domain: 2022 45th International Conference on Telecommunications and Signal Processing (TSP), IEEE, 6–10.

Choi, Y., Y. Jo, S. J. Seol, J. Byun, and Y. Kim, 2021, Deep learning spectral enhancement considering features of seismic field data: GEOPHYSICS, **86**, V389–V408.

Costa, V., N. Lourenço, J. Correia, and P. Machado, 2021, Demonstrating the Evolution of GANs Through t-SNE, *in* Applications of Evolutionary Computation: Springer International Publishing, 618–633.

Creswell, A., T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, 2018, Generative adversarial networks: An overview: IEEE signal processing magazine, **35**, 53–65.

De Jager, J., and C. Visser, 2017, Geology of the Groningen field – an overview: Netherlands Journal of Geosciences, **96**, s3–s15.

Dentith, M., and S. T. Mudge, 2014, Geophysics for the mineral exploration geoscientist: Cambridge University Press.

Du, H., Y. An, Q. Ye, J. Guo, L. Liu, D. Zhu, C. Childs, J. Walsh, and R. Dong, 2022, Disentangling Noise Patterns From Seismic Images: Noise Reduction and Style Transfer: IEEE Transactions on Geoscience and Remote Sensing, **60**, 1–14.

Ewida, H., and M. Sarhan, 2023a, Seismic noise attenuation using post-stack processing: a case study of Rabeh East Oil Field, Gulf of Suez Basin, Egypt: Euro-Mediterranean Journal for Environmental Integration, **8**, 645–664.

Ewida, H. F., and M. A. Sarhan, 2023b, Efficiency of post-stack processing in enhancing seismic data quality: a case study of Southwest Qarun-Field, Gindi Basin, Egypt: Euro-Mediterranean Journal for Environmental Integration, **8**, 949–967.

Fehler, M., and P. J. Keliher, 2011, SEAM Phase 1: Challenges of Subsalt Imaging in Tertiary Basins, with Emphasis on Deepwater Gulf of Mexico: Society of Exploration Geophysicists.

Fogat, M., S. Roy, V. Ferreira, and S. Singh, 2023, A Comparative Analysis of Convolutional Neural Networks for Seismic Noise Attenuation: Presented at the Day 3 Wed, June 07, 2023, SPE.

Folle, S., 2008, Salt Structures - Exploration and Limits of Interpretation.

Golfe, A., R. del Amor, A. Colomer, M. A. Sales, L. Terradez, and V. Naranjo, 2023, ProGleason-GAN: Conditional progressive growing GAN for prostatic cancer Gleason grade patch synthesis: Computer Methods and Programs in Biomedicine, **240**, 107695.

Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, 2020, Generative adversarial networks: Communications of the ACM, **63**, 139–144.

Goyes-Peñafiel, P., L. Suárez-Rodríguez, C. V. Correa, and H. Arguello, 2024, GAN Supervised Seismic Data Reconstruction: An Enhanced Learning for Improved Generalization: IEEE Transactions on Geoscience and Remote Sensing, **62**, 1–10.

Grijalva, F., W. Ramos, N. Perez, D. Benitez, R. Lara, and M. Ruiz, 2021, ESeismic-GAN: A Generative Model for Seismic Events From Cotopaxi Volcano: IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, **14**, 7111–7120.

Hou, X., L. Shen, K. Sun, and G. Qiu, 2016, Deep Feature Consistent Variational Autoencoder: 2017 IEEE Winter Conference on Applications of Computer Vision

(WACV), 1133–1141.

Howard, A., A. Sharma, J. Lenamond, J. Adamck, M. McDonals, S. Kainkaryam, and W. Cukierski, 2018, TGS Salt Identification Challenge — Kaggle.

Jamaludin, S. N. F., M. Pubellier, and D. Menier, 2018, Structural Restoration of Carbonate Platform in the Southern Part of Central Luconia, Malaysia: Journal of Earth Science, **29**, 155–168.

Karras, T., T. Aila, S. Laine, and J. Lehtinen, 2017, Progressive Growing of GANs for Improved Quality, Stability, and Variation: 6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings.

Kebaili, A., J. Lapuyade-Lahorgue, and S. Ruan, 2023, Deep Learning Approaches for Data Augmentation in Medical Imaging: A Review: Journal of Imaging, **9**, 81.

Keys, R. G., and D. J. Foster, 1998, 1. A Data Set for Evaluating and Comparing Seismic Inversion Methods, *in* Comparison of Seismic Inversion Methods on a Single Real Data Set: Society of Exploration Geophysicists, 1–12.

Ko, Y., S. Ko, and Y. Kim, 2024, Generative autoencoder to prevent overregularization of variational autoencoder: ETRI Journal.

Larsen, A. B. L., S. K. Sønderby, H. Larochelle, and O. Winther, 2015, Autoencoding beyond pixels using a learned similarity metric: 33rd International Conference on Machine Learning, ICML 2016, **4**, 2341–2349.

Li, Y., B. Ku, S. Zhang, J. K. Ahn, and H. Ko, 2020, Seismic Data Augmentation Based on Conditional Generative Adversarial Networks: Sensors 2020, Vol. 20, Page 6850, **20**, 6850.

Li, Z., F. Liu, W. Yang, S. Peng, and J. Zhou, 2022, A Survey of Convolutional Neural Networks: Analysis, Applications, and Prospects: IEEE Transactions on Neural Networks and Learning Systems, **33**, 6999–7019.

Liner, C. L., 2016, Elements of 3D Seismology: Society of Exploration Geophysicists.

Liu, D., W. Wang, X. Wang, C. Wang, J. Pei, and W. Chen, 2020, Poststack Seismic Data Denoising Based on 3-D Convolutional Neural Network: IEEE Transactions on Geoscience and Remote Sensing, **58**, 1598–1629.

Liu, S., C. Birnie, and T. Alkhalifah, 2023, Trace-wise coherent noise suppression via a self-supervised blind-trace deep-learning scheme: Geophysics, **88**, V459–V472.

Marano, G. C., M. M. Rosso, A. Aloisio, and G. Cirrincione, 2024, Generative adversarial networks review in earthquake-related engineering fields: Bulletin of Earthquake Engineering, **22**, 3511–3562.

Martin, G. S., R. Wiley, and K. J. Marfurt, 2006, Marmousi2: An elastic upgrade for Marmousi: The Leading Edge, **25**, 156–166.

Ministerio de Minas y Energía, and ANH, 2017, Acuerdo No. 02 de 2017.

Morley, C. K., U. B. Darussalam, T. Link, S. Begawan, and B. Darussalam, 2009a, Growth of folds in a deep-water setting: Geosphere, **5**, 59–89.

Morley, C. K., B. Kongwung, A. A. Julapour, M. Abdolghafourian, M. Hajian, D. Waples, J. Warren, H. Otterdoom, K. Srisuriyon, and H. Kazemi, 2009b, Structural development of a major late Cenozoic basin and transpressional belt in central Iran: The Central Basin in the Qom-Saveh area: Geosphere, **5**, 325–362.

Oktay, O., J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, and B. Kainz, 2018, Attention u-net: Learning where to look for the pancreas: arXiv preprint arXiv:1804.03999.

Oliveira, D. A. B., R. S. Ferreira, R. Silva, and E. V. Brazil, 2019, Improving Seismic Data Resolution With Deep Generative Networks: IEEE Geoscience and Remote Sensing Letters, **16**, 1929–1933.

Oren, C., and R. L. Nowack, 2018, An overview of reproducible 3D seismic data processing and imaging using Madagascar: Geophysics, **83**, F9–F20.

Park, J., S. Kim, S. J. Seol, and J. Byun, 2025, Improving generalization performance of deep learning–based seismic data interpolation: Geophysical Prospecting, **73**, 1534–1551.

Rashed, M., 2014, Fifty years of stacking: Acta Geophysica, **62**, 505–528.

Sam, G., and M. Kurt, 1995, Migration from topography: Improving the near-surface image.

Sheng, H., X. Wu, H. Gao, H. Di, S. Fomel, J. Li, and X. Si, 2025, On the workflow, opportunities and challenges of developing foundation model in geophysics.

Suzuki, M., and Y. Matsuo, 2022, A survey of multimodal deep generative models: Advanced Robotics, **36**, 261–278.

Torres, J., 2024, Aprendizaje supervisado y guiado por modelos generativos para la atenuación de ruido y corrección de daños estructurales en imágenes sísmicas post-apilado: Technical report, Universidad Industrial de Santander.

Torres-Quintero, J., P. Goyes-Peñafiel, A. Mantilla-Dulcey, L. Rodríguez-López, J. Sanabria-Gómez, and H. Arguello, 2025, Poststack seismic data denoising via dynamic guided learning: The Leading Edge, **44**, 692–704.

Wang, C., Y. Lyu, Q. Wang, G. Fu, Y. Wang, Y. Sun, Z. Huo, and J. Liu, 2017, Evaluations of oil and gas lateral migration across faults: A case study of Shigezhuang nose structure of Wen'an slope in Baxian sag, Jizhong depression, Bohai Bay Basin, East China: Petroleum Exploration and Development, **44**, 932–940.

Wang, F., and S. Chen, 2019, Residual Learning of Deep Convolutional Neural Network for Seismic Random Noise Attenuation: IEEE Geoscience and Remote Sensing Letters, **16**, 1314–1318.

Wang, S., X. Si, Z. Cai, and Y. Cui, 2022, Structural Augmentation in Seismic Data for Fault Prediction: Applied Sciences, **12**, 9796.

Wang, X.-W., S.-Z. Shi, X.-J. Yao, J.-B. Pei, Y.-F. Wang, H.-B. Yang, and D.-Q. Liu, 2023, Automatic identification of seismic faults via integrating Residual Network-50 residual blocks and convolutional block attention mod-

ules: Applied Geophysics, **20**, 20–35.

Xiao, Z., K. Kreis, and A. Vahdat, 2021, Tackling the Generative Learning Trilemma with Denoising Diffusion GANs: ICLR 2022 - 10th International Conference on Learning Representations.

Yang, L., W. Chen, H. Wang, and Y. Chen, 2021, Deep Learning Seismic Random Noise Attenuation via Improved Residual Convolutional Neural Network: IEEE Transactions on Geoscience and Remote Sensing, **59**, 7968–7981.

Yang, S., P. Ye, F. Shen, and D. Zhou, 2025, When Dynamic Data Selection Meets Data Augmentation.

Yuqing, W., L. U. Wenkai, L. JinLin, Z. Meng, and M. YongKang, 2019, Random seismic noise attenuation based on data augmentation and CNN: Chinese Journal of Geophysics, **62**, 421–433.

Zhang, C., X. Wei, and S.-W. Kim, 2021, Empirical Evaluation on Utilizing CNN-Features for Seismic Patch Classification: Applied Sciences, **12**, 197.

Zhu, D., 2020, Generate Images Using Variational Autoencoder (VAE).

Zhu, W., S. Mousavi, and G. Beroza, 2020, Chapter Four - Seismic signal augmentation to improve generalization of deep neural networks, *in* Machine Learning in Geosciences: Elsevier, volume **61** *of* Advances in Geophysics, 151–177.

Zhu, W., S. M. Mousavi, and G. C. Beroza, 2018, Seismic Signal Denoising and Decomposition Using Deep Neural Networks: IEEE Transactions on Geoscience and Remote Sensing, **57**, 9476–9488.