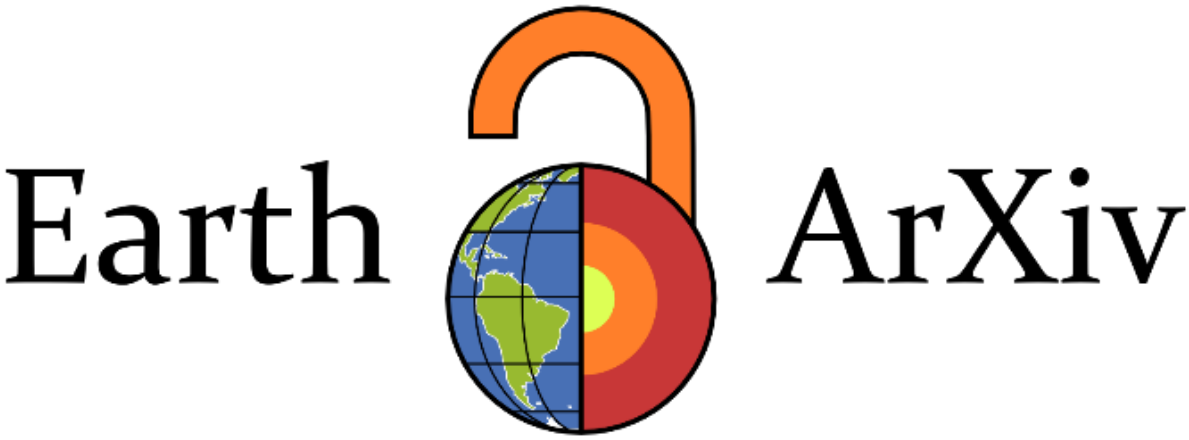

Geomodelling of multi-scenario non-stationary reservoirs with enhanced GANSim

Suihong Song[1,*], Tapan Mukerji[1], Celine Scheidt[1], Hisham Alqassab[2], Man Feng[2]

[1] Stanford University, 367 Panama St, Stanford, CA 94305, USA

[2] ExxonMobil Technology & Engineering Company, Spring, Texas, United States

Corresponding author: Suihong Song, suihong@stanford.edu

Abstract

Reservoir geomodelling is critical for groundwater management, CO_2 storage, geothermal exploitation, and hydrocarbon exploration, yet traditional geostatistical methods like multiple-point statistics (MPS) struggle with simulating complex geological patterns. GANSim, a Generative Adversarial Networks-based geomodelling method, has proven effective for single-scenario stationary reservoirs, but its performance on multi-scenario non-stationary systems remained unverified. Additionally, current GANSim may overlook single-pixel well facies data, causing local disconnections around wells. Therefore, this study proposes two workflows for multi-scenario reservoirs: one combining all scenarios together during training and another incorporating an explicit scenario falsification process before GANSim training. GANSim neural network architecture is further enhanced by proposing a local discriminator design to address the local disconnection problem of single-pixel well facies data. Validated on a multi-scenario non-stationary turbidite reservoir, both GANSim workflows generate realistic, conditional, and non-stationary facies models while falsifying incompatible scenarios. The local disconnection issue of single-pixel well facies data is effectively eliminated. Compared to MPS, GANSim demonstrates superior reproduction ability of expected geological patterns and computational efficiency, achieving simulations ~1000 times faster than MPS.

Plain Language Summary

Predicting the spatial distribution of subsurface reservoirs is essential for managing groundwater, storing CO_2, exploring for energy resources, and more. Scientists build digital models to represent reservoir distributions, but fully capturing their spatial patterns is challenging—especially when multiple geological scenarios are possible. Traditional geomodeling methods often struggle with such scenario uncertainty and pattern complexity. This study improves a deep learning-based method called GANSim, which uses neural networks to learn complex reservoir patterns from emulated geological models and generate realistic 3D subsurface models. Two new workflows are proposed to handle multiple geological scenarios, and a novel “local discriminator” design is introduced to better capture realistic local patterns around well locations. Using a complex non-stationary turbidite reservoir example, the enhanced GANSim workflows generate accurate models that honor input well facies and geophysical interpretation data. They also eliminate implausible geological scenarios using observed data—a process known as scenario falsification. These improvements make geomodelling faster, more reliable, and more realistic. The results show that the enhanced GANSim approach is up to 1,000 times more efficient than traditional methods. This offers a powerful tool for subsurface prediction, uncertainty reduction, and decision-making in geoscience and energy applications.

Key Words:

Reservoir geomodelling, Generative Adversarial Networks (GANs), scenario falsification, non-stationarity, turbidite

Key Points:

- Local discriminator design is proposed to address local disconnection problem of single-pixel well facies data existing in GANSim
- Two GANSim workflows are proposed for multi-scenario reservoirs and validated in a complex non-stationary turbidite reservoir
- Both workflows generate realistic, conditional, non-stationary facies models, falsify incompatible scenarios, and outperform MPS

1. Introduction

Reservoir geomodelling plays a critical role in various applications, including groundwater management, contamination disposal, geological storage of CO_2, hydrogen exploration, geothermal energy development, and hydrocarbon exploitation. Geomodelling is fundamentally a process of integrating diverse sources and forms of data and information, such as geological knowledge, borehole interpretations, geophysical data, and temporal well production data. Despite the availability of these data types, the data are often insufficient to fully capture the complexity, high-dimensionality, and multi-scale nature of subsurface systems. As a result, uncertainty persists even when all collected data sources are effectively integrated. In practice, multiple facies geomodel realizations are produced to represent the uncertainty of subsurface reservoirs, thus providing quantitative and scientific basis for analysis and decision-making.

Traditionally, geostatistical approaches have been widely used for reservoir facies geomodelling, such as variogram-based methods (e.g., sequential indicator simulation; Deutsch, 2002) and multiple-point statistics (MPS; Mariethoz & Caers, 2014). Variogram-based and MPS methods are effective at honoring well facies interpretations and geophysics-interpreted facies probability cubes. However, these methods often struggle to accurately reproduce complex spatial geological patterns (Song et al., 2025b). Recently, deep generative methods have demonstrated strong ability to capture and reproduce complex spatial patterns, e.g., Generative Adversarial Networks (GANs; Goodfellow et al., 2014) and diffusion models (Ho et al., 2020). GANs, in particular, have been extensively researched and applied in geomodelling (Alqassab et al., 2024; Song et al., 2022b, 2023; Zhang et al., 2019), while diffusion models, despite their potential, are still in the early research stage for geomodelling (Di Federico & Durlofsky, 2025; Lee et al., 2025; Xu et al., 2024).

In GANs, the generator maps a latent vector to an output, while the discriminator distinguishes between real and generated data. Through alternative training of the discriminator and generator, the generator eventually learns and can reproduce complete spatial patterns exhibited in the training dataset. Supporting Information S1 introduces GANs with more details. When applied to geomodelling, the generator learns geological patterns from training conceptual geomodels which are constructed based on geological knowledge, and can then generate reservoir geomodels that exhibit the learned patterns (Laloy et al., 2018; Song et al., 2021a; Zhang et al., 2019). However, vanilla GANs alone do not provide conditioning for geomodelling. To address this, post-GANs latent vector perturbation workflows have been proposed. These workflows search for appropriate latent vectors that, when fed into a pre-trained generator, produce geomodel realizations that are both realistic and consistent with given conditioning data (e.g., Laloy et al., 2018; Mo et al., 2020; Mosser et al., 2020; Nesvold & Mukerji, 2021; Zhang et al., 2019). Such latent vector search methods can include Markov chain Monte Carlo (McMC; Brooks et al. (2011)), iterative ensemble smoother (IES; Chen & Oliver (2013)), and gradient-based optimization techniques. However, once the conditioning data changes, the latent vector search process must be repeated, which is computationally expensive. Additionally, 100% conditioning accuracy cannot always be guaranteed for well facies data.

Therefore, Song et al. (2021b, 2022a) proposed GANSim, a GANs-based geomodelling workflow, for reservoir geomodelling directly conditioned on well facies data and geophysically interpreted facies probability cubes. In GANSim, in addition to the original adversarial loss, condition-based loss functions are introduced to train the generator to learn both conditioning rules (i.e., the relationship between

output facies models and input conditioning data) and geological patterns. With these two types of learned knowledge, the trained generator can quickly produce multiple realistic and conditional facies model realizations for any given conditioning data. Furthermore, due to the fully convolutional design of the generator architecture, although trained on smaller-scale conceptual geomodels or their patches, the generator can produce reservoir geomodel realizations of arbitrarily large scales once trained. GANSim has been successfully applied for conditional geomodelling of 3D field karst cave reservoirs (Song et al., 2022b) and meandering channel reservoirs (Hu et al., 2024). Supporting Information S2 describes GANSim workflow with more details. To incorporate well production data, Song et al. (2023) further proposed a GANSim-surrogate framework. In this framework, a deep learning-based surrogate model is trained in conjunction with the pre-trained generator, and appropriate input latent vectors are searched so that the output facies geomodels are consistent with the given production data through the trained surrogate as well as the given well facies, facies probability maps, and global feature values. To efficiently construct such a surrogate, a purely physics-informed neural operator approach can be employed (Song et al., 2025a).

However, the proposed GANSim workflow and its field applications still have limitations. First, in previous algorithm development research and field applications (Hu et al., 2024; Song et al., 2021b; Song et al., 2022a; Song et al., 2022b), GANSim has primarily been used for single-scenario reservoirs, i.e., reservoirs exhibiting unimodal distribution in geological pattern uncertainty. However, in reality, due to limited knowledge of the subsurface, the assumed reservoir scenarios can be diverse (scenario uncertainty). For example, in the case of a channel reservoir, the assumed geological scenarios may include meandering channels, braided channels, and a transitional form between them. Similarly, delta reservoirs may involve scenarios such as the jet-plume model (Wellner et al., 2005) as well as the classic model comprising a delta plain, delta front, and pro-delta. Turbidite reservoirs may involve scenarios with varying levels of confinement (McHargue et al., 2021). Given this diversity of possible geological scenarios, how can GANSim be adapted for such common multi-scenario cases?

Secondly, due to computational resource limitations, large-size conceptual geomodels are commonly cropped into smaller-size patches for training GANSim. As a result, the generator learns pattern fragments from these smaller conceptual geomodel patches. When applied to geomodelling of large-size reservoirs, the trained generator essentially samples the learned pattern fragments for different locations within the reservoir domain. This approach is theoretically reasonable and has proven successful in cases where the spatial patterns are statistically stationary, such as karst cave and meandering channel reservoirs (Hu et al., 2024; Song et al., 2022b). It works because pattern fragments across the entire reservoir domain are similar (from the same distribution) in stationary cases, and there is no strict ordering requirement for pattern fragments at different locations during generation. However, for non-stationary reservoirs, the pattern fragments vary significantly across the entire domain, and there is an implicit ordering requirement for local patterns at different locations. In this case, if GANSim is still trained with conceptual geomodel patches, the generated large-size reservoir geomodels may exhibit significant disorganization of patterns. For example, a proximal delta pattern might directly connect to a distal delta pattern without a transitional medial pattern. In the current GANSim algorithm, the conditioning on facies probability cubes and maps of global features (Song et al., 2025b) is considered, which might guide the generator to sample appropriate pattern fragments for different locations during geomodelling. In this sense, the trained generator may still produce well-organized non-stationary geological patterns. However, this requires verification with non-stationary complex reservoir cases.

Thirdly, a persistent challenge in current GANSim implementations is the discriminator's tendency to overlook single-pixel (single-voxel) well facies data, leading to severe disconnections between the single-pixel well facies data and surrounding regions of the same facies type (Hu et al., 2024; Song et al., 2021b). To mitigate this, studies have artificially expanded well facies into multi-pixel blocks (e.g., 4×4 pixels; Song et al. (2021b); Song et al., (2022b)). However, this deliberate expansion of well facies introduces local unrealistic artifacts around the wells and artificially reduces the local reservoir prediction uncertainty. Is it possible to improve GANSim so that the original single-pixel well facies data are not overlooked by the discriminator, eliminating the need for deliberate well facies expansion? Solving this challenge is crucial for all GANSim applications, regardless of reservoir scenario uncertainty (single vs. multi-scenario) and spatial pattern types (stationary vs. non-stationary).

Therefore, in this paper, we propose two GANSim workflows for multi-scenario cases in Section 2. Building on the workflows, the GANSim neural network architectures are enhanced by introducing local discriminators to address the single-pixel conditioning with local disconnection problem in Section 3. In Section 4, we define a multi-scenario problem related to geomodelling of a non-stationary turbidite reservoir. In Sections 5 and 6, the two GANSim workflows are applied to the turbidite geomodelling problem. Section 7 provides analyses and discussions about the results of the two workflows. A comparison to one MPS method is also rendered in this section. Finally, conclusions are presented in Section 8.

2. Two GANSim geomodelling workflows for multi-scenario reservoirs

We propose two GANSim geomodelling workflows for reservoirs with geological scenario uncertainty (Figure 1). In the first workflow, multiple geological scenarios are combined to train a single comprehensive generator that learns geological knowledge across all scenarios and conditioning rules (i.e., the relationship between output facies models and input conditioning data). Once trained, the generator can be applied to geomodelling any scenario. A prior falsification process is performed before GANSim training to refine the compiled geological knowledge and the conceptual geomodel creation process. The second workflow is similar, with the key difference being that some prior geological scenarios are falsified using available conditioning data. Separate GANSim models are then trained for each of the remaining (unfalsified) scenarios. GANSim training requires many more conceptual geomodels than the prior and scenario falsification processes. Compared to the first workflow, the trained generators in the second workflow are limited to geomodelling unfalsified scenarios only but is generally easier to train than the comprehensive generator in the first workflow. The two workflows are described in detail in the following sections.

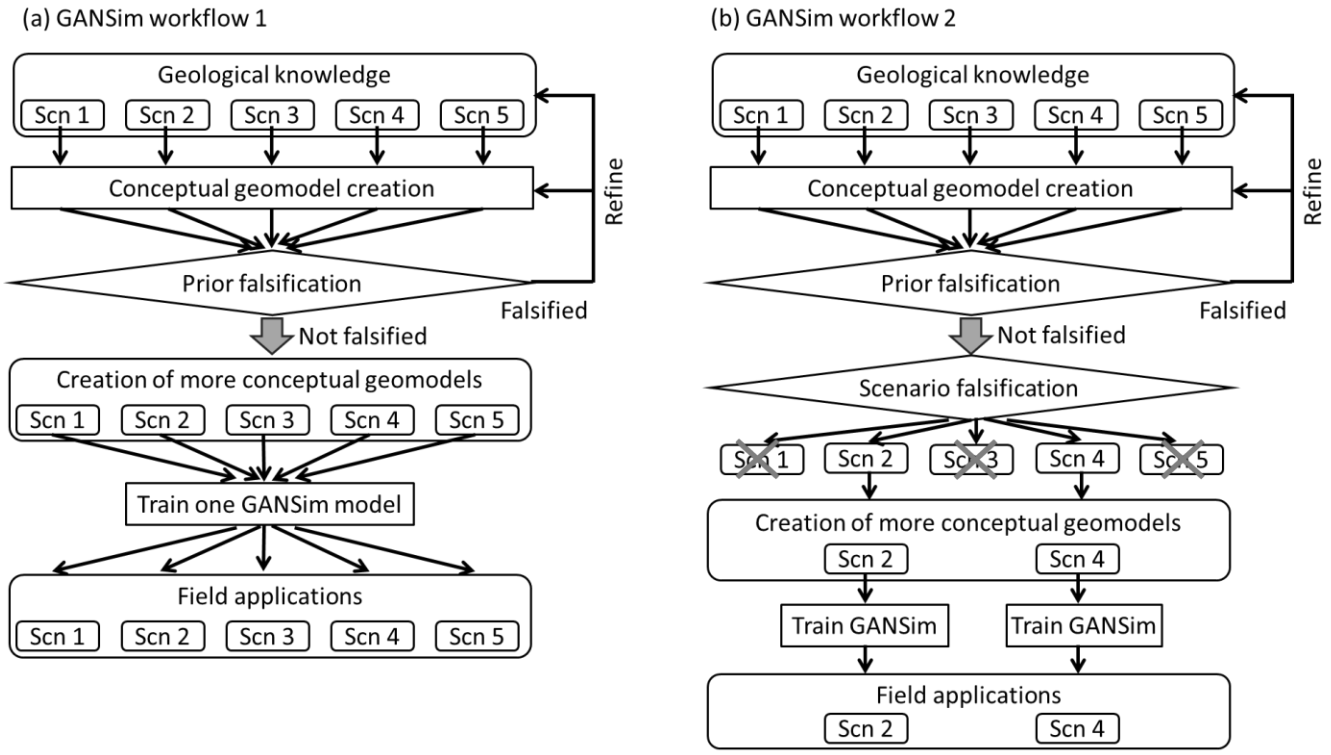

Figure 1. Major steps in the proposed two GANSim workflows. “Scn” denotes a geological scenario. Five scenarios are illustrated in the figure.

2.1. Steps of both GANSim workflows

Figure 1 compares the major procedures of the two workflows. The following provides more detailed explanations of each step:

1. *Geological knowledge compilation for different scenarios:* Compile geological knowledge for each possible geological scenario. This may include variations in geometrical shapes and topological relationships of reservoir geobodies, variations in sedimentary facies, reservoir formation processes, and the dimensional ranges of geobodies.
2. *Creation of conceptual geomodels for each scenario*: The complied geological knowledge is conceptualized into computer-readable formats such as logical rules, mathematical equations, and maps. Automated workflows are then developed based on this conceptualized knowledge to efficiently emulate conceptual geomodels for different scenarios. These workflows may be physics-based, process-mimicking, or object-based, depending on the desired balance between geological realism and computational efficiency (Pyrcz & Deutsch, 2014).
3. *Prior and scenario falsification using observed data*: Due to insufficient understanding and inaccuracies during the compilation and conceptualization of geological knowledge, the emulated conceptual geomodels may not reflect the uncertainty distribution of real subsurface reservoir geomodels. To address this, a falsification process is applied using

observed (conditioning) data to validate the compiled knowledge and the emulation workflows. In both workflows, the ensemble of conceptual geomodels of all scenarios (collectively referred to as the prior) are combined to go through the prior falsification process using observed actual data. If this prior ensemble is falsified, the geological knowledge or emulation workflow is refined. Otherwise, the workflow proceeds. In the second workflow, an additional scenario falsification step is performed using observed data after the prior falsification. The workflow proceeds with only the remaining unfalsified scenarios. Detailed procedures of the two falsification processes are described in Section 2.2.

4. *Building the training dataset for GANSim*: The training dataset includes an ensemble of diverse conceptual geomodels, well facies data, and facies probability cubes. Significantly more conceptual geomodels with sufficient diversity are required for training than for the prior and scenario falsification processes. To reduce the training burden, large-size conceptual geomodels are typically cropped into smaller-size patches, from which the training well facies and probability cubes can be obtained. The crop size represents a trade-off between training efficiency and the geological realism of the generated models (Song et al., 2022b).
5. *Design of neural network architectures*: Generator and discriminator architectures are designed with specific enhancements described in Section 3.
6. *Training of GANSim*: GANSim is trained using the training dataset and the loss functions described in Section 3. In the first workflow, a single GANSim model is trained across all scenarios, while in the second one, separate models are trained for each unfalsified scenario using their respective training datasets.
7. *Evaluation of the trained generator*: The generator is evaluated based on its ability to reproduce expected geological patterns and geobody dimensions as well as its conditioning capability to given well facies and probability cubes. Evaluation methods include visual inspection, multidimensional scaling (MDS), and other techniques, as detailed in Song et al. (2022a).
8. *Geomodelling for large field reservoirs*: The trained generator is used to model arbitrarily large field reservoirs by taking in field conditioning data and varying expanded latent cubes. This enables generation of multiple conditional realizations, which are then assessed for geological realism and conditioning accuracy. Note that under Workflow 1, the trained generator can be applied to any scenario, but in Workflow 2, the trained generators apply only to unfalsified scenarios.

2.2. Prior and scenario falsification process

In this study, the term prior refers to the ensemble of all geological scenarios and can be viewed as a “super scenario”. In the following, we describe the falsification process for multiple scenarios, which also applies to the prior considered as a single super scenario.

The idea of falsification in science is not new (Popper, 1959) and has been applied across various scientific domains, including geosciences (Scheidt et al., 2018; Tarantola, 2006). A simple understanding of the logic behind reservoir scenario falsification is as follows: if a geological scenario can potentially

produce the ground truth reservoir geomodel—meaning that the ground truth geomodel lies within the distribution of conceptual geomodel samples generated from the scenario—then the observed data (e.g., geophysical data) from the ground truth geomodel will also lie within the distribution of data forward-simulated from these conceptual geomodels, both in the original data space and in any reduced feature space. Conversely, if the observed data does not lie within the simulated data distribution in the original or any reduced feature space, the scenario cannot produce the ground truth geomodel and is therefore falsified. This follows the logical principle: if $a \Rightarrow b$, then $\neg b \Rightarrow \neg a$.

The common steps of scenario falsification involve generating numerous conceptual geomodels from each scenario, forward-simulating these conceptual geomodels to obtain corresponding data, extracting sensitive features from both the simulated data and the observed data, and projecting these data into a low-dimensional space using the Multidimensional Scaling (MDS) method (Borg & Groenen, 2005) based on the feature distances of every pair of data. Finally, the posterior probability of the scenarios given the features of the observed data (see Equation (3) below) is evaluated in the low-dimensional MDS space to determine whether the scenario is falsified. The posterior probability is calculated as

$$P(Sc_k|f(d)) \approx \frac{P\left(f(d)\middle|Sc_k\right)P(Sc_k)}{\sum_k P\left(f(d)\middle|Sc_k\right)P(Sc_k)}, \qquad (3)$$

where, Sc_k represents the *k*-th scenario ($1 \leq k \leq K$), $f(d)$ represents the feature of the observed data *d*. The likelihood $P(f(d)|Sc_k)$ is evaluated based on the data density at the point of the observed data in the MDS space (see e.g., Scheidt et al., 2015b). The "≈" symbol is used in the equation because the likelihood is evaluated in a low-dimensional MDS space rather than the original feature space, and it also accounts for all inaccuracies involved in the conceptual geomodel generation and forward simulation. Note that the evaluated $P(Sc_k|f(d))$ differs from the posterior $P(Sc_k|d)$, as only partial information of the observed data (i.e., its global statistical features) is used for conditioning here in the falsification process. Uncertainty in the forward simulation, such as uncertainty in the rock physics model for seismic data, should also be considered. The feature filter, $f(\cdot)$, should be sensitive to geological scenarios. An appropriate filter can differentiate scenarios, to a large extent.

The data *d* can be any observed data type related to reservoir facies geomodel, including flow simulation data (Park et al., 2013), well facies data (Scheidt et al., 2015a), seismic data (Scheidt et al., 2015b), or electrical resistivity tomography (Hermans et al., 2014). In this paper, we consider well facies interpretations at different locations and seismic data as *d*. As shown in Figure 2, for well facies interpretations, the multi-point histogram (MPH) analysis is used as the feature filter operation, denoted as $f_{MPH}(\cdot)$, to extract a histogram of different facies stacking patterns for each well profile, as described in Tan et al. (2014). MPH can be applied at multiple resolutions using a pyramid approach. When performing the MDS method to project histogram features into a reduced-dimensional space, the Jensen-Shannon (JS) divergence (Lin, 1991) is used to evaluate the histogram distance between every pair of well facies data. Then, in the reduced-dimensional MDS space, the probability density functions (pdf) of any data feature for different scenarios, $P(f_{MPH}(w)|Sc_k)$ where $f_{MPH}(w)$ refers to features of any well facies data, are evaluated after a kernel smoothing procedure on the point density. The likelihood $P(f_{MPH}(w_{obs})|Sc_k)$, where w_{obs} refers to observed well facies interpretations, is then obtained from the pdf. Figure 2 (bottom left) illustrates how likelihoods are read from pdf maps for two scenarios (adapted from Park et al., (2013)). Finally, using the Bayes' equation (Equation (3)), the posterior $P(Sc_k|f_{MPH}(w_{obs}))$ can be calculated. More details are provided in Scheidt et al. (2015b).

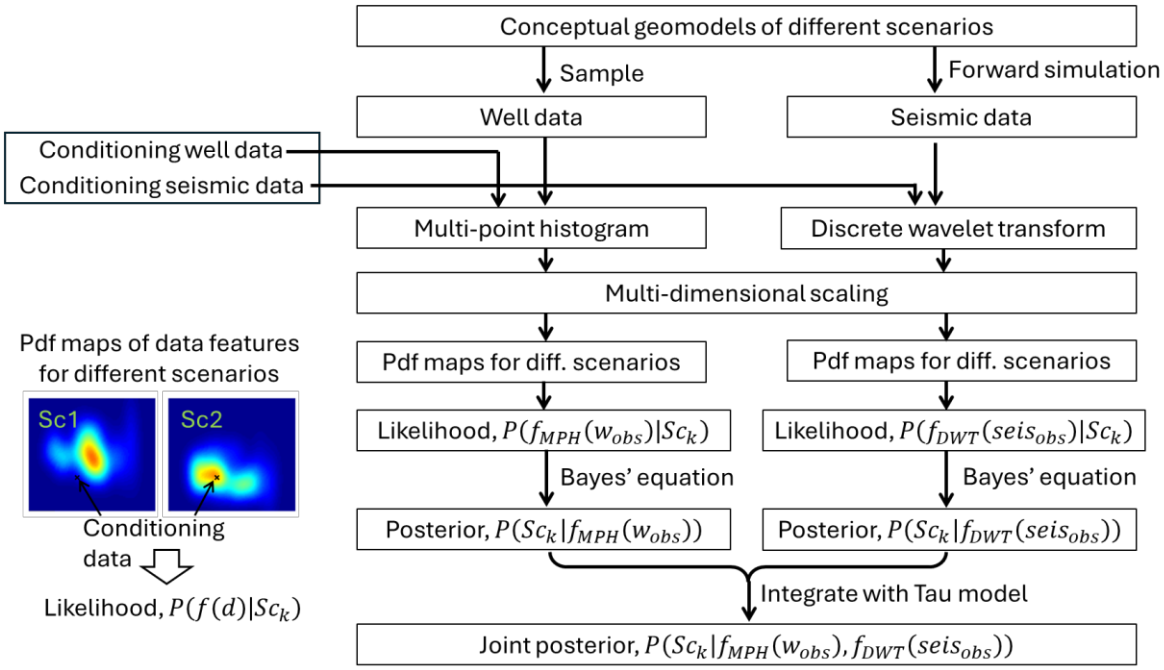

Figure 2. Workflow of scenario falsification used in this paper. The two probability density function (pdf) maps are adapted from Park et al. (2013).

For seismic data, a 2D discrete wavelet transform (DWT) approach is used as the feature filter, denoted as $f_{DWT}(\cdot)$, as described in detail in Scheidt et al. (2015b). The 2D wavelet transform decomposes the seismic image into a set of wavelet coefficients at different levels and directions, which effectively characterizes the spatial patterns of seismic data and is thus informative for geological scenarios. A similarity distance (L2-norm) based on the histograms of the wavelet coefficients is used when applying the MDS method. The approaches to obtain the likelihood and the posterior, $P(Sc_k|f_{DWT}(seis_{obs}))$ where $seis_{obs}$ refers to the conditioning seismic data, are similar to the well facies case.

Due to the potential information redundancy between the two posteriors obtained using observed well facies and seismic data—$P(Sc_k|f_{MPH}(w_{obs}))$ and $P(Sc_k|f_{DWT}(seis_{obs}))$—a probability combination scheme commonly used in geostatistics, the Tau model (Journel, 2002), is used to obtain the joint posterior, $P(Sc_k|f_{MPH}(w_{obs}), f_{DWT}(seis_{obs}))$, based on which scenarios with low posterior values are falsified.

3. GANSim enhancement with local discriminator design

Supporting Information S2 provides a detailed description of the previous version of GANSim. In the earlier implementation, the discriminator primarily focuses on assessing the global geological realism of the entire reservoir geomodel and often overlooked the sparsely distributed single-pixel well facies data, leading to significant disconnections between the input well facies pixels and the surrounding facies of the same type in facies geomodels produced by the generator (see Figure 9). This local

disconnection issue significantly affects flow behaviors of the simulated reservoir geomodels. To address this limitation, we propose to design local discriminators that specifically evaluate the local geological realism around well facies data, in addition to the original global discriminator. Multiple local discriminators can be designed to focus on different scales of locality around the wells. These local discriminators can either be separate from or integrated into the global discriminator. In the integrated design, all discriminators share the same shallow convolutional layers. Our experiments show that both separate and integrated architectures produce similar results. See Figure 3 as an example of the integrated local discriminator design. Note the well location indicator cube is considered as another input of the discriminator for easier calculation of local discriminator scores.

The Wasserstein loss function with gradient penalty (Gulrajani et al., 2017) is used in GANSim. We found that the generator achieves good local and global realism by summing up the Wasserstein losses of different discriminators, rather than summing their scores first and then calculating a single Wasserstein loss. Appropriate weights are assigned to the Wasserstein losses of the global and local discriminators during the summation. With this local discriminator design, the well facies data expansion approach (Song et al., 2021b; Song et al., 2022b) is no longer necessary, thereby alleviating artifacts associated with expanded well facies data.

The loss function of the generator becomes

$$L(G) = \mathbb{E}_{z\sim p_z,(w,p)\sim p_{(w,p)}} \sum_{i=0}^{n} \alpha_i[-D_i(G(z,w,p),I_{wloc})] + \beta_1 L(G)_w + \beta_2 L(G)_p, \qquad (1)$$

where, G refers to the generator, $D_0, D_1, \dots, D_n$ represent the global discriminator (D_0) and n local discriminators respectively, α_i are the corresponding weights for each discriminator; z, w, and p refer to the input latent vector, input training well facies data, and input training probability maps or cubes, p_z and $p_{(w,p)}$ are distributions of the latent vector and the training data; I_{wloc} is the well location indicator map; $L(G)_w$ and $L(G)_p$ are the condition loss terms for well facies and probability maps (described in detail in Supporting Information S2), β_1 and β_2 are their respective weights, and $\mathbb{E}$ is the expectation operator. The discriminator loss is defined as

$$L(D) = \mathbb{E}_{z\sim p_z,(w,p)\sim p_{(w,p)},x\sim p_x} \sum_{i=0}^{n} \alpha_i[D_i(G(z,w,p),I_{wloc}) - D_i(x,I_{wloc}) + \lambda\mathbb{E}(\| \nabla_{\hat{x}} D_i(\hat{x},I_{wloc}) \|_2 - 1)^2], \quad (2)$$

where, x refers to a training facies model and p_x is its distribution, $\hat{x}$ is sampled between an $x\sim p_x$ and an $x_G = G(z,w,p)$, i.e., $\hat{x} = tx + (1-t)x_G$, with $t\sim\text{Uniform}(0,1)$, and λ is a predefined weight for gradient penalty. Global feature conditioning is not explicitly included in the above loss functions. To incorporate global features, the corresponding loss terms (Equation (S2-2)) can be directly added to the above generator loss (Equation (1)).

To improve the reasonability of the neural network architectures in GANSim, we further revised the generator's output and the discriminator's input from a single facies geomodel to multiple facies indicator models (i.e., "one-hot" encoding of facies category) of all facies types (see Figure 3). During geomodelling for field reservoirs, these facies indicators are converted into a reservoir facies model using argmax or thresholding-related operations. This design is common in variogram-based facies geomodelling algorithms, such as sequential indicator simulation (Pyrcz & Deutsch, 2014).

Figure 3 illustrates the enhanced architecture of the generator and discriminator used for facies geomodelling of turbidite reservoirs in this study. The turbidite reservoir has three facies categories: channels, lobes, and background mud. The generator takes as inputs 8 latent cubes with a size of 4×4×4,

2 well facies cubes each with a size of 128×128×32 representing a well facies code cube and a well location indicator cube, and 2 probability cubes for channel and lobe facies. The output of the generator includes three facies indicator cubes for channel, lobe, and mud facies, respectively, each with a size of 128×128×32. The generator employs a fully convolutional layer design, enabling geomodelling of arbitrarily large reservoirs by the generator after training. This approach is not constrained by the size of the training geomodel patches, as proposed by Song et al. (2022b). ReLU activation functions are used for all layers of the generator except the last one, which utilizes a softmax activation function.

For the discriminator, the inputs consist of three facies indicator cubes and one well location indicator cube each with a size of 128×128×32. The architecture includes one global discriminator and three local discriminators. In each local discriminator, following several convolutional and downsampling layers—that are shared with the global discriminator—and three separate convolutional layers (one with the kernel size of 3×3×1 and two with the kernel size of 1×1×1), a small feature cube (e.g., the green cube with the size of 64×64×16×1 in Figure 3(b)) is produced. Each feature value in the feature cube corresponds to a receptive field — that is, a localized region of the input facies indicator cubes contributing to the feature value—and reflects the realism of that region. An element-wise multiplication is then performed between this feature cube and a downsampled well location indicator cube of the same size, effectively masking the feature cube to only focus on receptive fields near wells. The result is then averaged to compute a local score, which quantifies the local realism within receptive fields around all wells in the input facies indicator cubes. The size of the receptive fields can be analytically derived from the convolution and downsampling operations from the input facies indicators to the feature cube. Local discriminator 1 evaluates the realism within a local region of 18×18×14 cells (receptive size of the local score) around all wells, while local discriminator 2 corresponds to a locality of 40×40×32 cells, and local discriminator 3 corresponds to a locality of 84×84×32 cells around all wells. Since the Wasserstein loss is used, the final layer of the discriminator employs a linear activation function, while ReLU is used for all other layers.

A comparison of the generated realism around wells—with and without the proposed local discriminator design—is presented in Section 5.2 in the context of turbidite reservoir geomodeling (Figures 7, 8, and 9), demonstrating the effectiveness of the design.

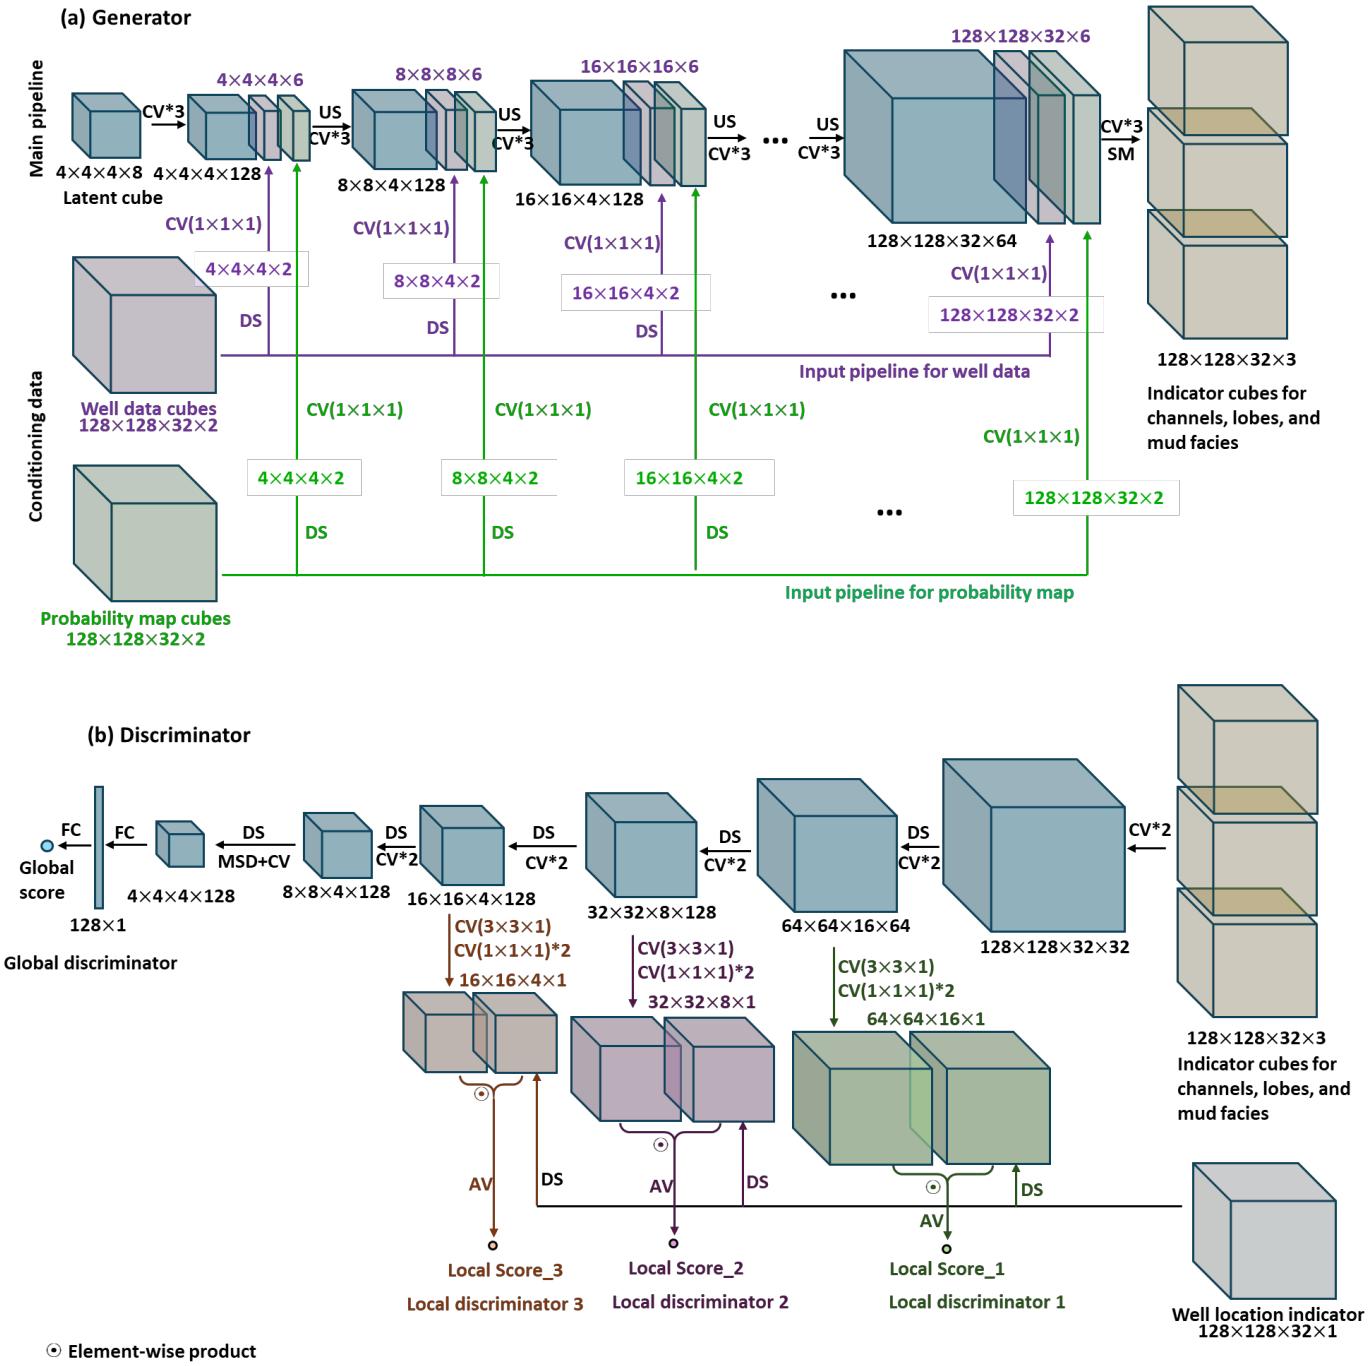

Figure 3. The enhanced generator and discriminator architectures used for geomodelling of turbidite reservoir in this study.

4. Definition of geomodelling problem for a multi-scenario non-stationary turbidite reservoir

Submarine turbidite channel-lobe systems are relatively sand-rich sediment deposits produced by gravity-driven turbidity currents within a predominantly muddy submarine environment (Normark, 1978). These systems are primarily characterized by sinuous connective channels—which act as conduits for sediment transport—and lobes—which form as sheet-like distributary deposits at the distal ends of channels (Mutti & Lucchi, 1978). Channel fills and lobate deposits are the major permeable sedimentary facies compared to muddy levees and the background submarine mud due to their coarse-grained nature. Turbidite deposits have been extensively studied in terms of classification, dimensions, architectures, evolution, hierarchies, controlling factors, and geomodelling (Chen et al., 2024; McHargue et al., 2021; Prélat et al., 2009, 2010; Spychala et al., 2020). Turbidite channel-lobe deposits are excellent reservoirs for CO_2, underground water, and hydrocarbons.

Influenced by mud proportion in turbidity currents and paleo-geomorphology, turbidite channels on the abyssal plain may present a distributary character and are confined to varying degrees (McHargue et al., 2021). For example, different periods of channel fills can be strongly confined by surrounding muddy levees in a channel-levee system (Deptuck et al., 2007; McHargue, 1991; Posamentier & Kolla, 2003), confined by a previously eroded valley or paleo-geomorphological highs (Deptuck et al., 2007; McHargue et al., 2011), or completely unconfined. Based on the level of channel confinement, we establish five geologic scenarios for turbidite sediments on the abyssal plain, each containing a channel facies, a lobe facies, and a background mud facies. The five scenarios, Sc_k $(1 \leq k \leq 5)$, each assumed to have equal prior probability of 1/5, are (1) strongly confined channels and lobes, (2) transitional scenario between the strongly confined and the valley-confined scenario, (3) valley-confined channels and lobes, (4) weakly confined channels and lobes, and (5) unconfined channels and lobes. In each scenario, lobate deposits can be eroded by later channel fills, and both channels and lobes diverge from the source to the distal part of the system. Due to this divergent style of channels and lobes, the turbidite reservoir on the abyssal plain exhibits typical non-stationary characteristics. The combination of multi-scenario and non-stationary features makes the turbidite reservoir an ideal case to validate the GANSim workflows with local discriminator enhancement proposed in previous sections.

Building on the above turbidite geological knowledge, we employed an object-based workflow, implemented in Petrel software, to emulate 3D turbidite conceptual geomodels. It consists of five major steps: (1) construction of a probability map to sample channel endpoints, (2) sampling of a specified number of channel endpoints, (3) simulation of channels from the end to the source (i.e., from the south to the north in Figure 5), (4) simulation of lobes, and (5) combination of simulated channels and lobes. By varying the probability map construction methods and adjusting the number of channel endpoints, conceptual geomodels for all five scenarios can be emulated. Key hyperparameters of the workflow include global proportions of channels and lobes, their orientation maps, distribution trends, and dimension ranges. Adjusting these hyperparameters allows for production of diverse geomodel realizations for each scenario. Figure 4 displays one emulated conceptual geomodel example for each scenario. In these conceptual geomodels of all the five scenarios, lobes and channels are distributed in a divergent style along the flow direction, and channels may cut through preexisting lobes. In the strongly confined scenario, channels are strictly confined, leading to vertical stacking of different periods of channels. In the valley-confined scenario, channels are distributed within a valley. Confinement becomes looser in the weakly confined scenario, and there is no confinement in the unconfined scenario. The transitional scenario exhibits channel confinement levels between the strongly confined and valley-

confined scenarios. These observations align with expected geological knowledge, demonstrating the effectiveness of the proposed object-based emulation workflow.

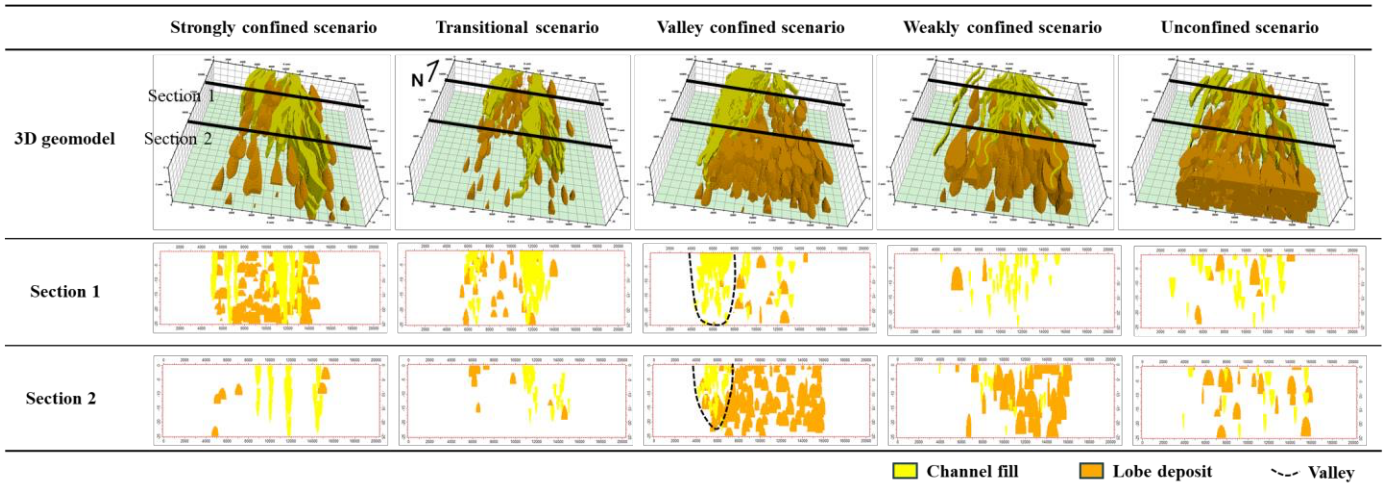

Figure 4. Emulated conceptual geomodel examples, each including a 3D geomodel and two sections, of the five turbidite reservoir scenarios. The background mud facies is set to be invisible for better visualization. Each conceptual geomodel consists of 400×400×50 cells, with each cell measuring 50×50×0.5 m.

The object-based workflow is used to produce a thick conceptual geomodel from the strongly confined scenario as the reference ground truth geomodel in this study (Figure 5). It comprises 400×400×100 cells, with each cell measuring 50×50×0.5 m. A detailed description of the working steps and hyperparameter selection for the ground truth geomodel is provided in Supporting Information S3. Next, we sampled 18 wells from the ground truth facies geomodel. A forward simulation—comprising reservoir property simulation, statistical rock physics modeling, and normal-incidence convolution seismic simulation—was performed to generate 3D seismic data for the ground truth geomodel. Details of the forward simulation process are provided in Supporting Information S4. Theoretically, facies probability cubes should be interpreted from the seismic data. However, to simplify the procedure, we obtained the probability cubes for channel and lobe facies by applying a Gaussian smoothing filter to the facies indicator models of the ground truth geomodel. This approximates the results of a statistical geophysical interpretation process (Avseth et al., 2005). A certain percentage of noise was introduced during the Gaussian smoothing operation to make the probability cubes more realistic. The well locations, well facies interpretations, seismic data, and facies probability cubes are shown in Figure 5.

The proposed two GANSim workflows are subsequently applied for geomodelling of the referenced reservoir using these conditioning data and the geological knowledge (or the object-based workflow), without knowing the true scenario.

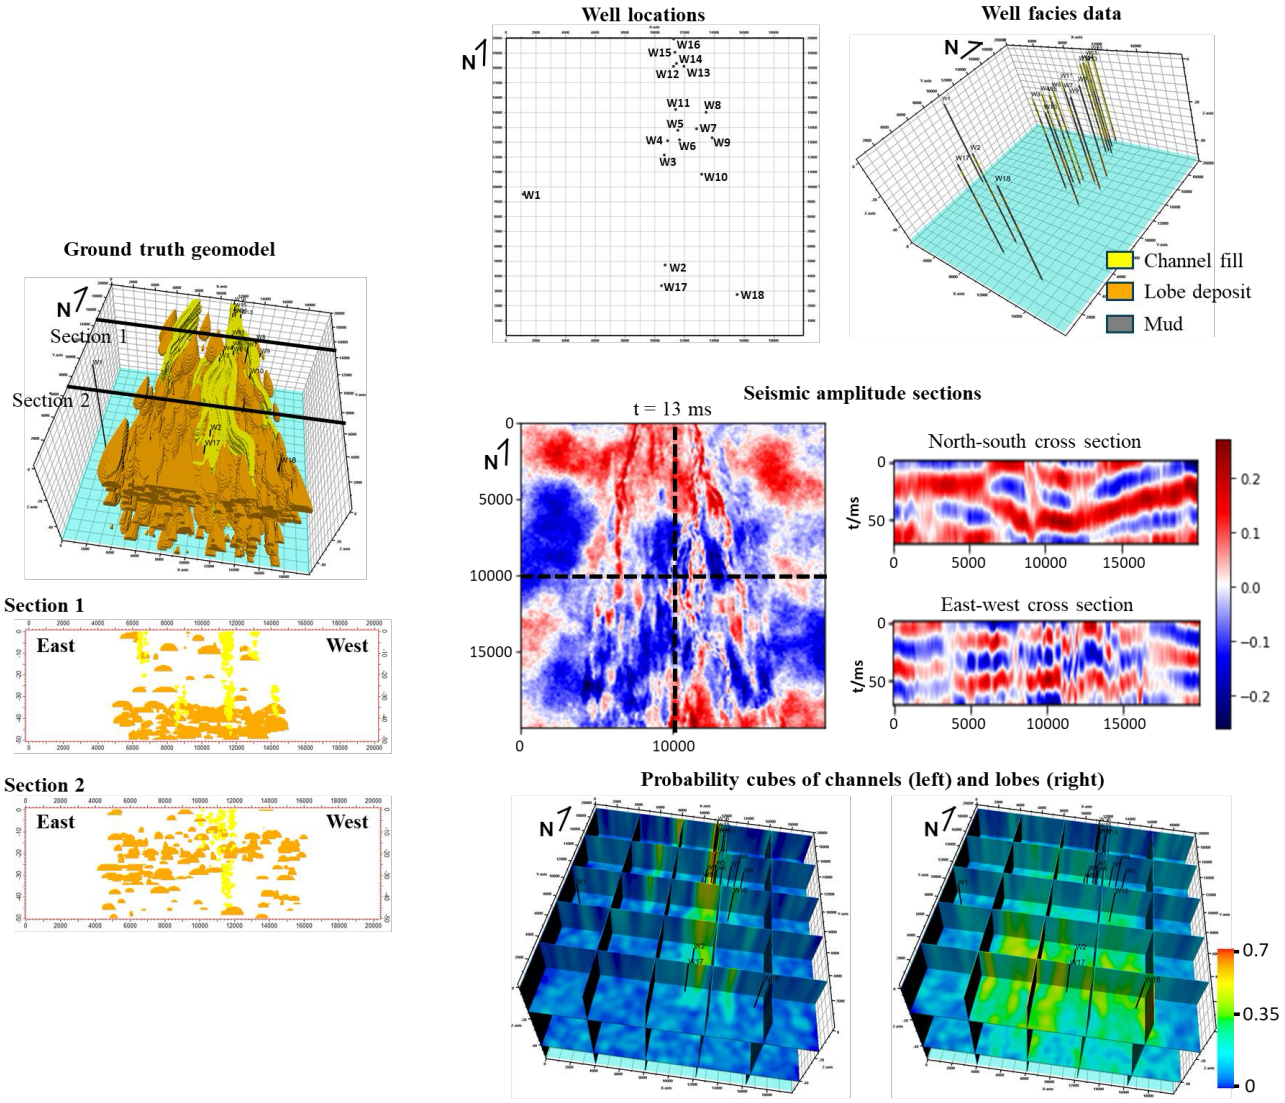

Figure 5. The ground truth reservoir geomodel and its two sections, along with the corresponding 18 well facies data, seismic amplitude data, and probability cubes of channels and lobes. The background mud facies is set invisible in the ground truth geomodel and its sections. The locations of the two vertical seismic sections are marked in the horizontal seismic section.

5. Geomodelling of turbidite reservoir with GANSim workflow 1

The first step in the GANSim workflow is the compilation of geological knowledge about the target reservoir, with particular attention to different possible scenarios. For the studied turbidite reservoir in abyssal plains, the relevant geological knowledge has been detailed in Section 4. The specifications for the remaining steps are outlined in the following sections.

5.1. Construction of conceptual geomodels and the training dataset

We apply the object-based workflow described in Section 4 to emulate 640 large-scale conceptual geomodel realizations for each of the five geological scenarios. Figure 4 displays an emulated conceptual geomodel example for each scenario. Each conceptual geomodel consists of 400×400×50 cells (half the vertical cell number of the ground truth geomodel), with cell dimensions of 50×50×0.5 m (matching the ground truth geomodel's cell size). This results in a total thickness for each conceptual geomodel half that of the ground truth geomodel.

Following step 3 of GANSim workflow 1, these prior conceptual geomodels should be validated through a prior falsification process using the given conditioning data. However, in this synthetic study, since the conceptual geomodels are emulated using the same procedure as the ground truth geomodel, these prior conceptual geomodels of all scenarios, as a whole, would not be falsified by the conditioning data. Therefore, we skip step 3 and proceed directly to preparing the training dataset. Note that in GANSim workflow 2, we still carry out the scenario falsification step, described later in Section 6.

To save computational resources, the emulated conceptual geomodels are randomly cropped along channels and lobes to form smaller-size training facies models. According to Song et al. (2022b), for stationary reservoirs, a rule of thumb for the crop size is that it should be larger than the largest repetitive pattern unit, such as a meander in a meandering channel reservoir. In this non-stationary reservoir study, we maximize the crop size into 128×128×32 cells within the limits of GPU memory resources (during training) to enable the generator to learn as complete non-stationary geological patterns as possible. Finally, 38,400 cropped facies model patches are obtained (12 from each conceptual geomodel), with 35,000 used as training facies models and the remaining reserved for test. Because the progressive training technique (Karras et al., 2017) is applied in GANSim, these facies model patches are further downsampled into coarser resolutions from 64×64×16, 32×32×8, …, to 4×4×4. Since facies indicator models are required for training (see Figure 3), we further decompose these patches of different resolutions, into indicator models of channel, lobe, and mud facies. From each facies model patch of the finest resolution (128×128×32), we randomly sample 1 to 10 vertical wells—each occupying 1×1 cell horizontally—as the training well facies data. The Gaussian smoothing method, with added noise, is applied to the decomposed indicator patches of the finest resolution to mimic geophysically interpreted probability cubes of channels and lobes, following a similar approach as in Song et al. (2022b). Some examples of training facies model patches and the corresponding training well facies data and facies probability cubes are shown in Figure 6.

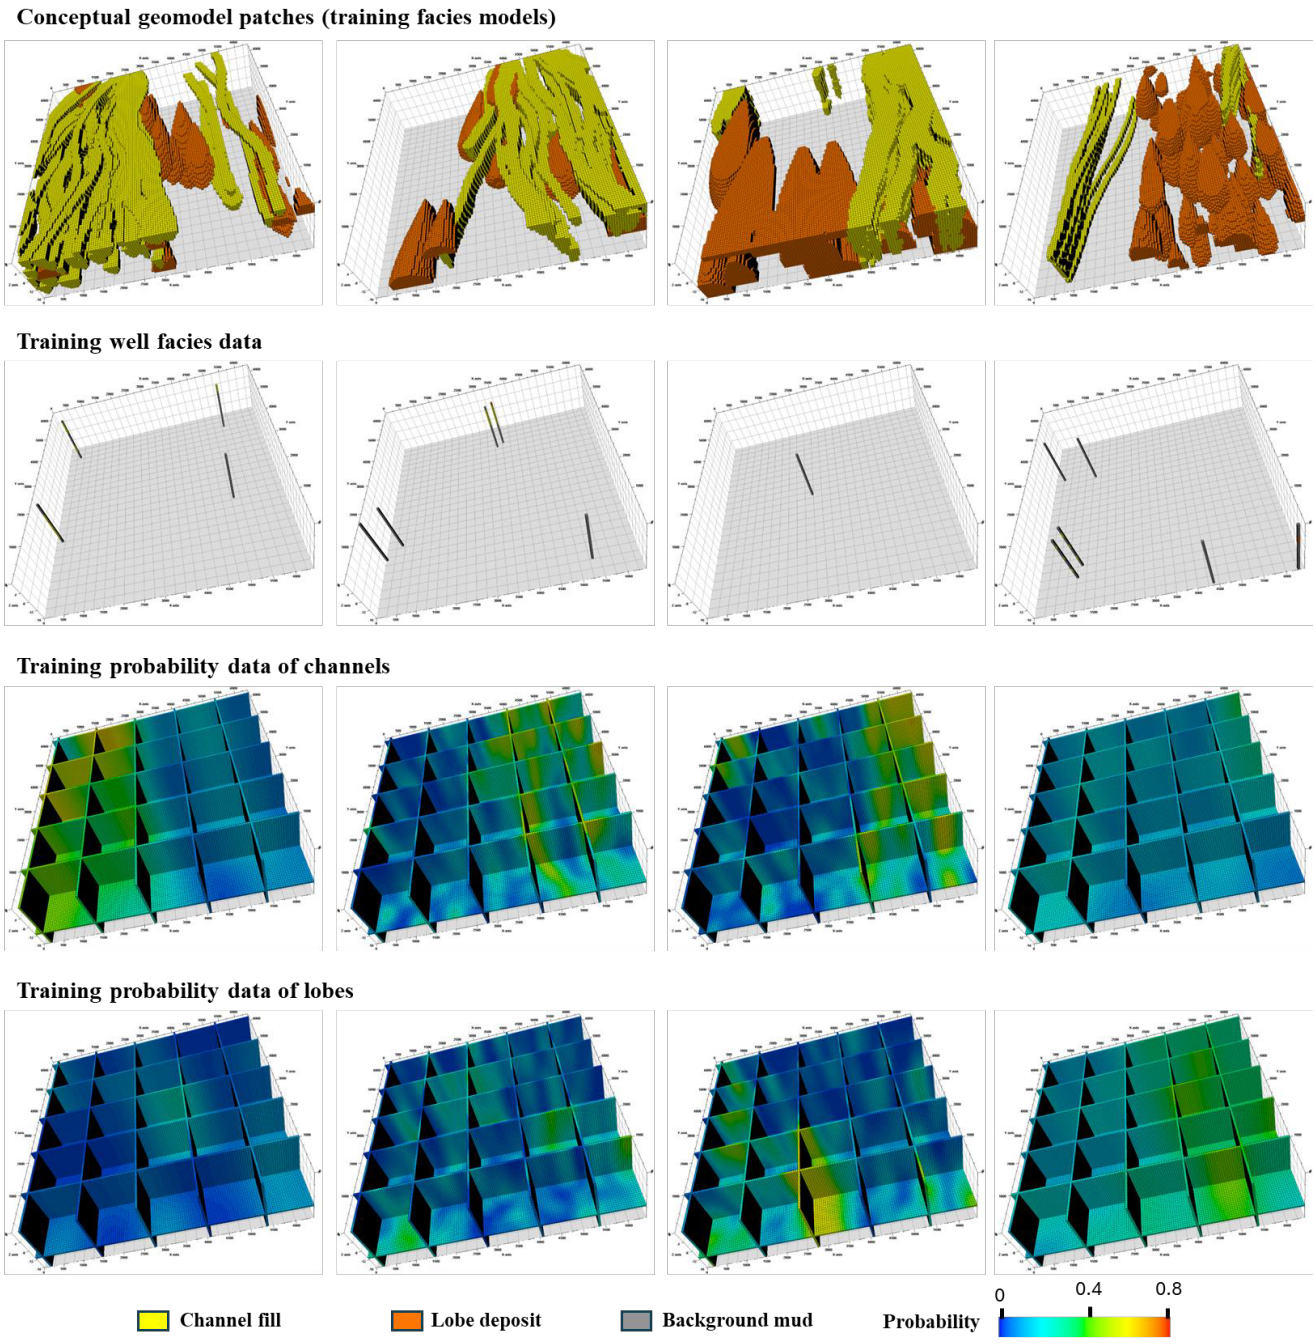

Figure 6. Random examples of training conceptual facies model patches and the corresponding training well facies data and training probability cubes of channels and lobes.

5.2. Neural network architecture design, training, and evaluation of GANSim

The neural network architectures of the generator and discriminator are described in Section 3.

We use the minibatch gradient descent method and the Adam optimizer with default parameters (Kingma & Ba, 2014) for training. The minibatch size is set to 128, which is appropriate for the available

GPU memory. Since a progressive training approach is used in GANSim, the layers of the generator and the discriminator are trained progressively from shallow to deep. The generator and discriminator are alternately trained, each with a single minibatch. The weights for the condition-based loss terms of well facies and probability cubes in Equation (1) (β_1 and β_2) are set as 2 and 0.08, respectively, after trial-and-error experiments. There are one global discriminator and three local discriminators (see Figure 3). Their weights in the loss function (i.e., α_0, α_1, α_2, and α_3 in Equation (1) and Equation (2)) are set as 0.1, 0.9, 9, and 90 for the global discriminator, local discriminator 3, local discriminator 2, and local discriminator 1, respectively. The weight λ in Equation (2) is set as 10 by default. The enhanced GANSim is trained on 8 GPUs (A100-40G) in parallel for 72 hours, until the generated facies models are realistic (e.g., exhibiting good shapes of channels and lobes), diverse, and consistent with the input conditioning data.

After training, we evaluate the generator using test facies model patches. As shown in Figure 7 and Figure 8, two test facies models are used as validation. The corresponding conditioning well facies interpretations (128×128×32) and probability cubes of channels and lobes (128×128×32) are taken into the trained generator, along with random latent cubes, to produce different facies model realizations (128×128×32). Each realization takes an average of 0.02 seconds on one GPU (A100). From these realizations and their cross-sections, we observe that most of the simulated channels are connected (although a few disconnected channels are present, e.g., the channel patch indicated by the green arrow in realization 3 of Figure 7) and exhibit a downward concave shape. The simulated lobes show the expected fan-like upward concave shape. The channels may cut through lobes. These features align with expected geological knowledge.

In addition to the geological realism of the produced models, we calculated the well facies reproduction accuracies, which are 100% and 99% for the two cases shown in Figure 7 and Figure 8. When comparing the generated realizations with and without (Figure 9) the proposed local discriminator design, unrealistic disconnection artifacts are clearly visible in the cross-well sections of the latter case (highlighted by blue circles in Figure 9). In contrast, such artifacts are not observed in Figures 7 and 8, demonstrating the effectiveness of the local discriminator design. Frequency cubes of channels and lobes were calculated based on 100 randomly generated geomodel realizations. From a comparison among the geomodel realizations, the calculated facies frequency cubes, and the input facies probability cubes in 3D cubes or 2D sections, we observe a distinct consistency between the generated channels and lobes and the input probability cubes of the two facies, indicating a clear conditioning effect of the input probability cubes. Furthermore, compared to the input probability cubes, the distributions of high-frequency values are more concentrated, reflecting a reduction in reservoir prediction uncertainty with GANSim. This reduced uncertainty primarily results from the assimilation of conditioning well facies data, probability cubes, and geological knowledge.

In summary, the trained generator can efficiently produce realistic and conditional facies model realizations. Song et al. (2022b) observed similar results in stationary karst cave cases. However, they expanded the well facies data to ensure that the input well facies interpretations were not overlooked by GANSim. With the enhanced GANSim, such expansion is no longer necessary, while the generator can still produce strong local realism around the input well facies data.

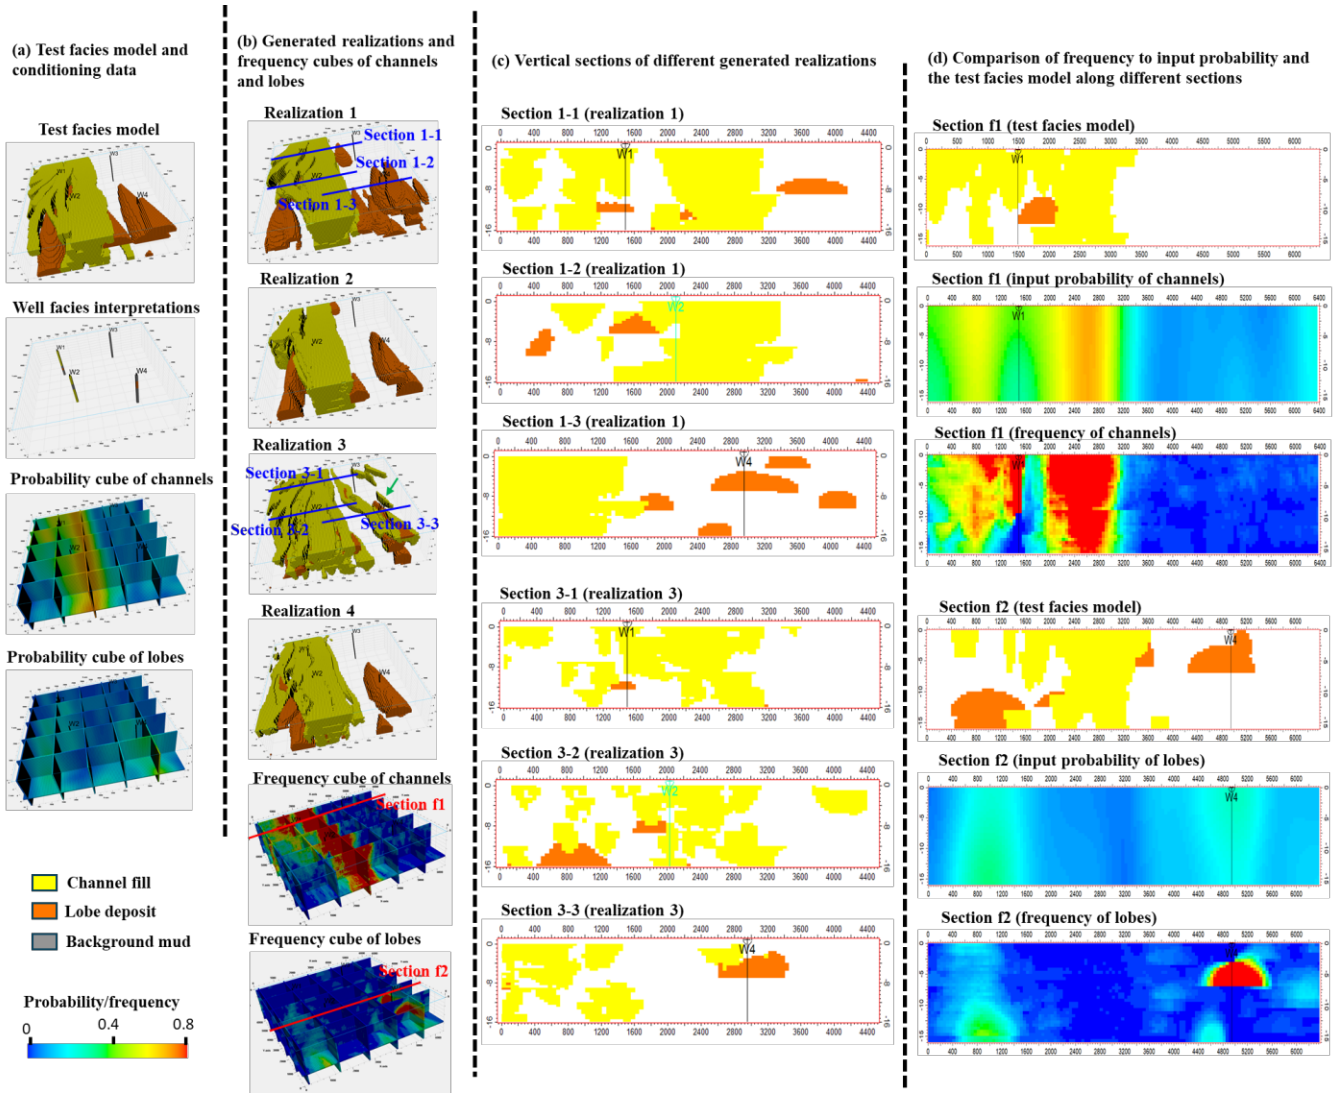

Figure 7. Evaluation of the trained generator based on one test facies model patch. (a) The test facies model patch (from the test dataset) is shown, along with its corresponding conditioning well facies interpretations and probability cubes of channels and lobes. (b) Four randomly generated facies model realizations are displayed, along with the frequency cubes of channels and lobes. (c) Cross-well sections of different realizations are presented. (d) Frequency distributions are compared to the input probability distributions and the test facies models across different cross-well sections.

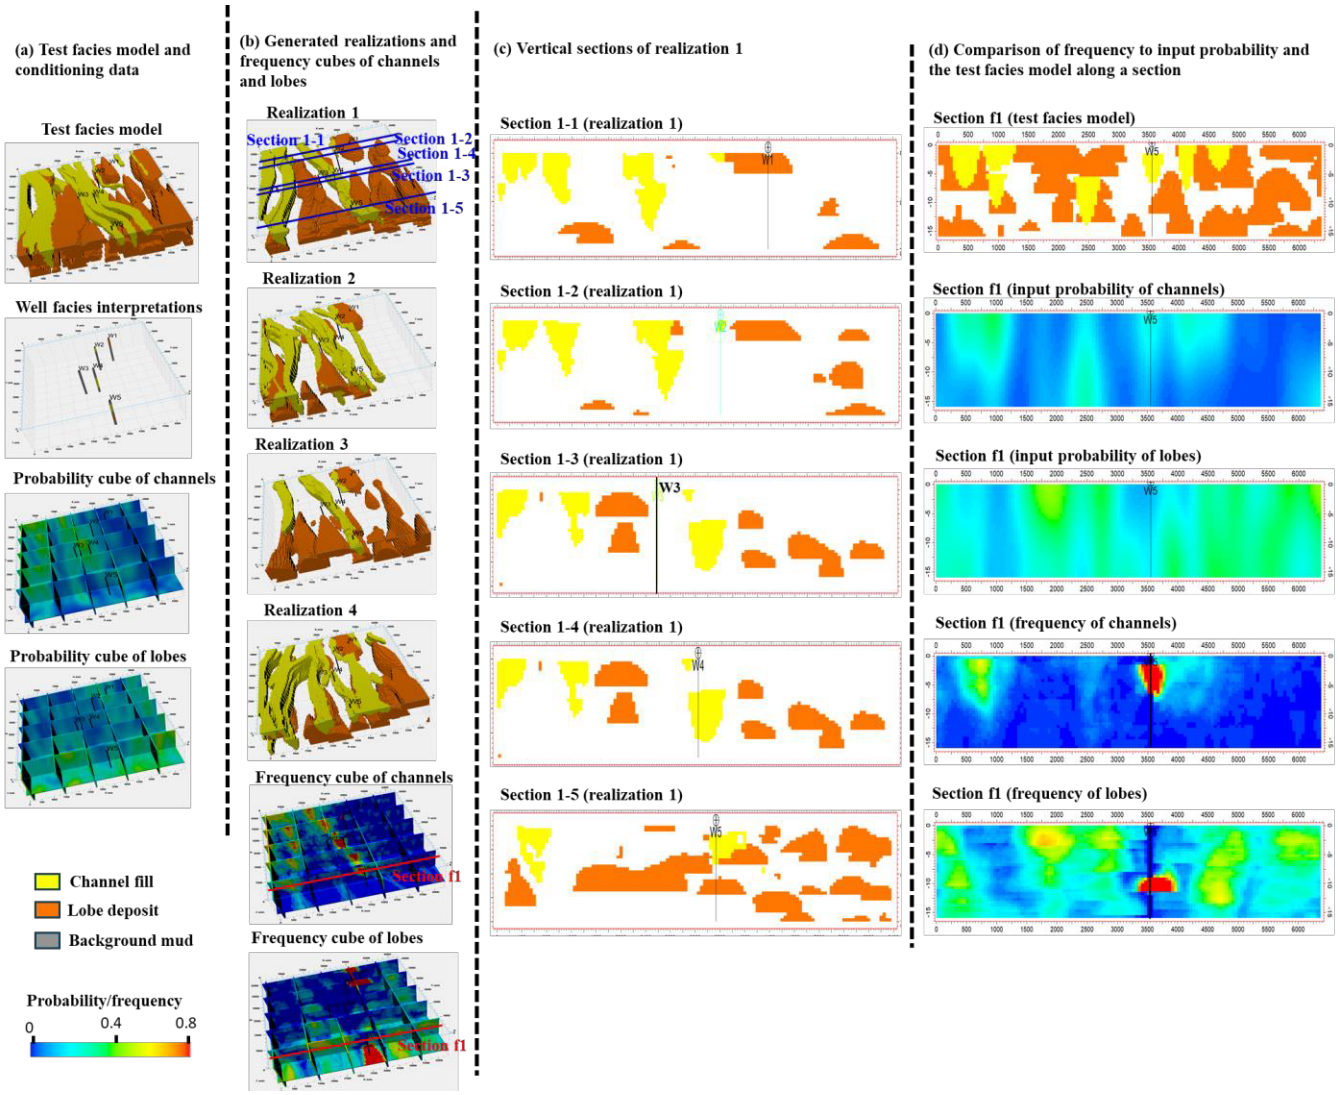

Figure 8. Same as Figure 7, but for a different test facies model. Section f1 is equivalent to section 1-5.

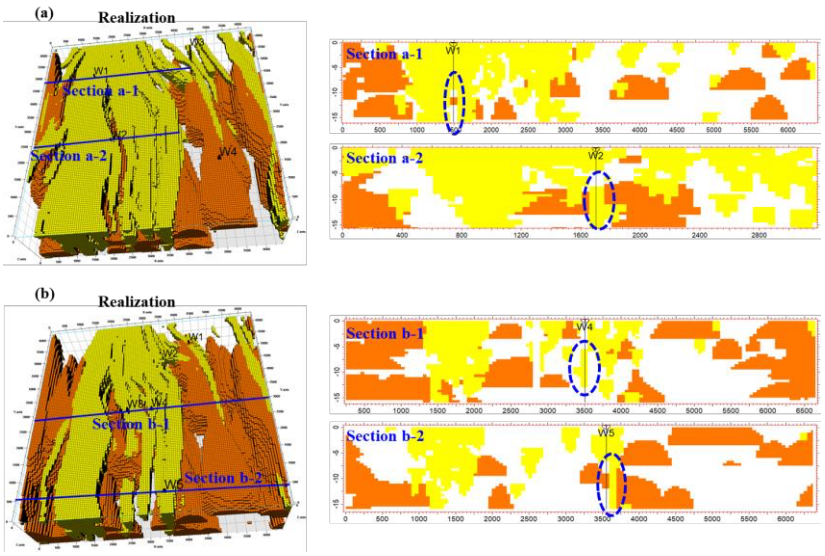

Figure 9. (a) and (b) show the realizations (3D views and 2D vertical sections) produced by a GANSim generator trained without the local discriminator design, using the conditioning data from Figures 7 and 8. Blue circles highlight unrealistic disconnection artifacts around the well facies data.

5.3. Geomodelling for the large reference reservoir

After evaluation, the trained generator is used for geomodelling of the large referenced reservoir described in Section 4 without knowing its ground truth facies model. Note that the reservoir, with dimensions of 400×400×100 cells, is over 30 times larger than the training facies model patches, which have dimensions of 128×128×32 cells. The conditioning well facies interpretations, probability cubes of channels and lobes, and random latent cubes are taken as inputs into the trained generator to produce different facies model realizations with a size of 400×400×100. Frequency cubes of channels and lobes are calculated based on 50 random realizations. Figure 10 (a) shows three randomly simulated facies model realizations and the calculated facies frequency cubes. Figure 11 and Figure 12 (the second row) illustrate one realization and the facies frequency distributions in two cross-well sections. The locations of the two cross-well sections are marked in Figure 10 (a).

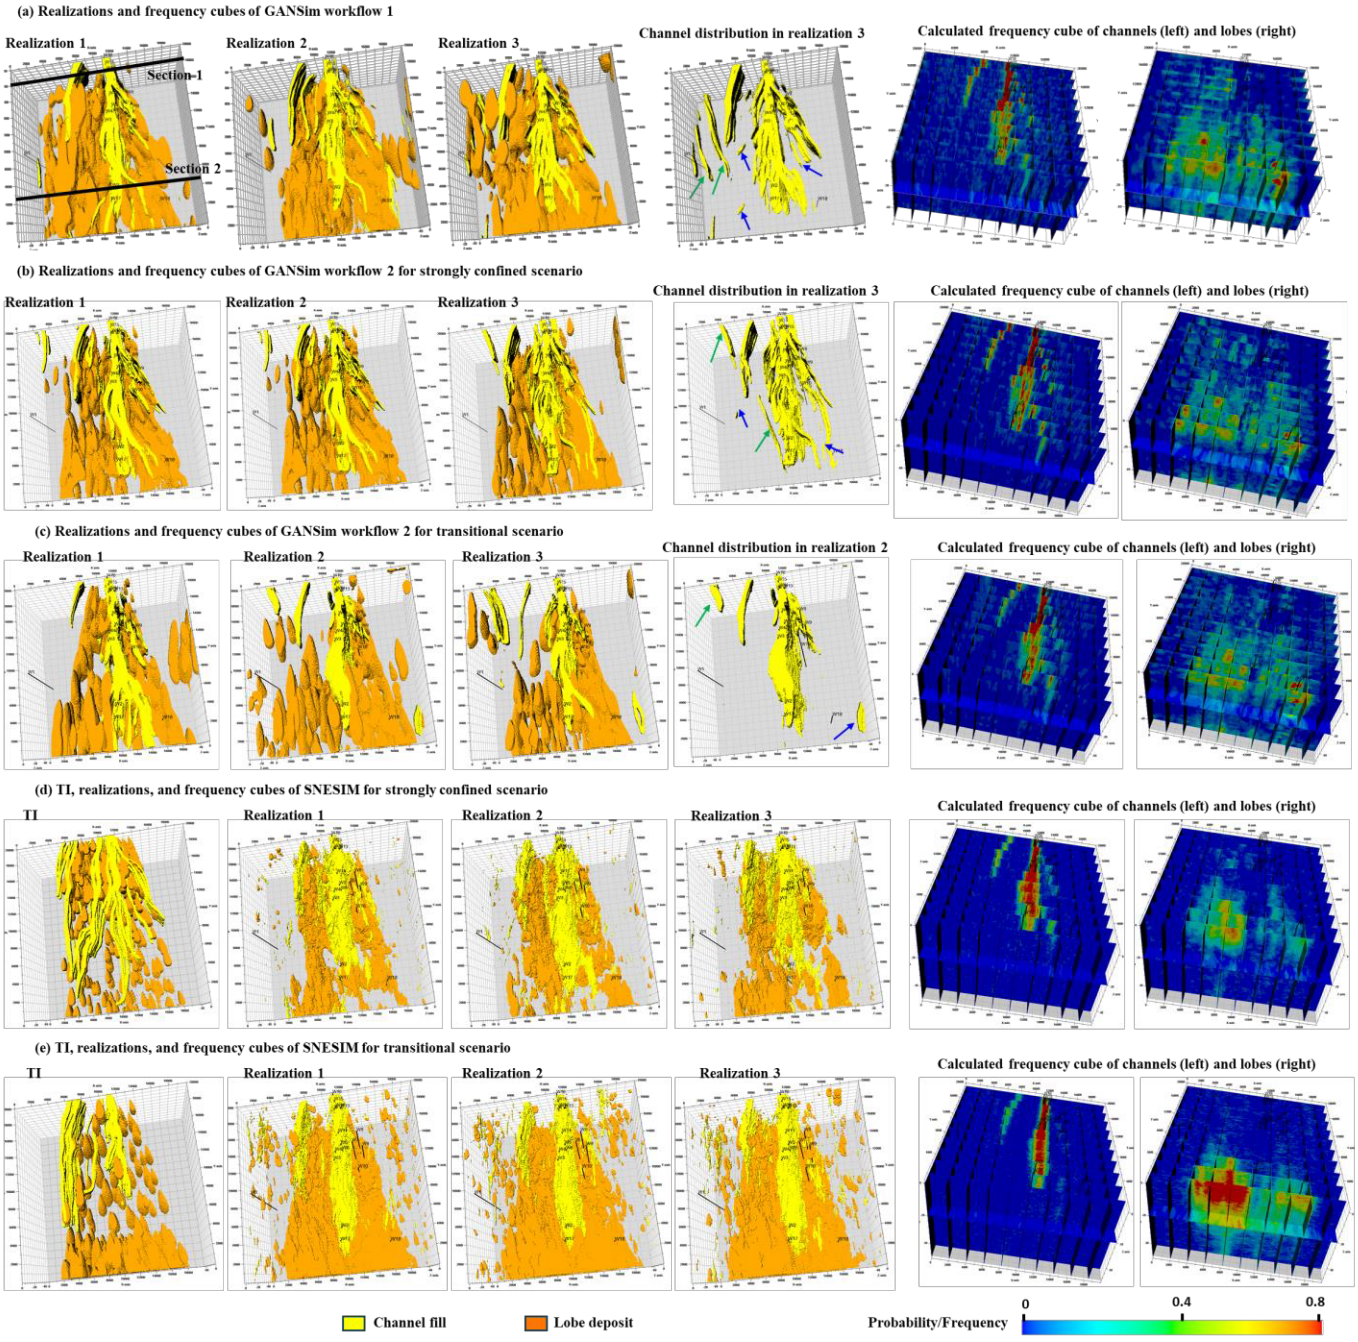

Figure 10. Simulated turbidite facies model realizations and the corresponding facies frequency cubes using (a) GANSim workflow 1, (b) GANSim workflow 2 with strongly confined scenario, (c) GANSim workflow 2 with transitional scenario, (d) SNESIM method with strongly confined scenario, and (e) SNESIM method with transitional scenario. Training images (TIs) are shown for SNESIM method. Details from section 1 and section 2 marked in (a) are shown in following figures 9 and 10. Blue arrows indicate examples of locally disconnected channel fragments, while green arrows highlight unintended local patterns that contradict the global styles.

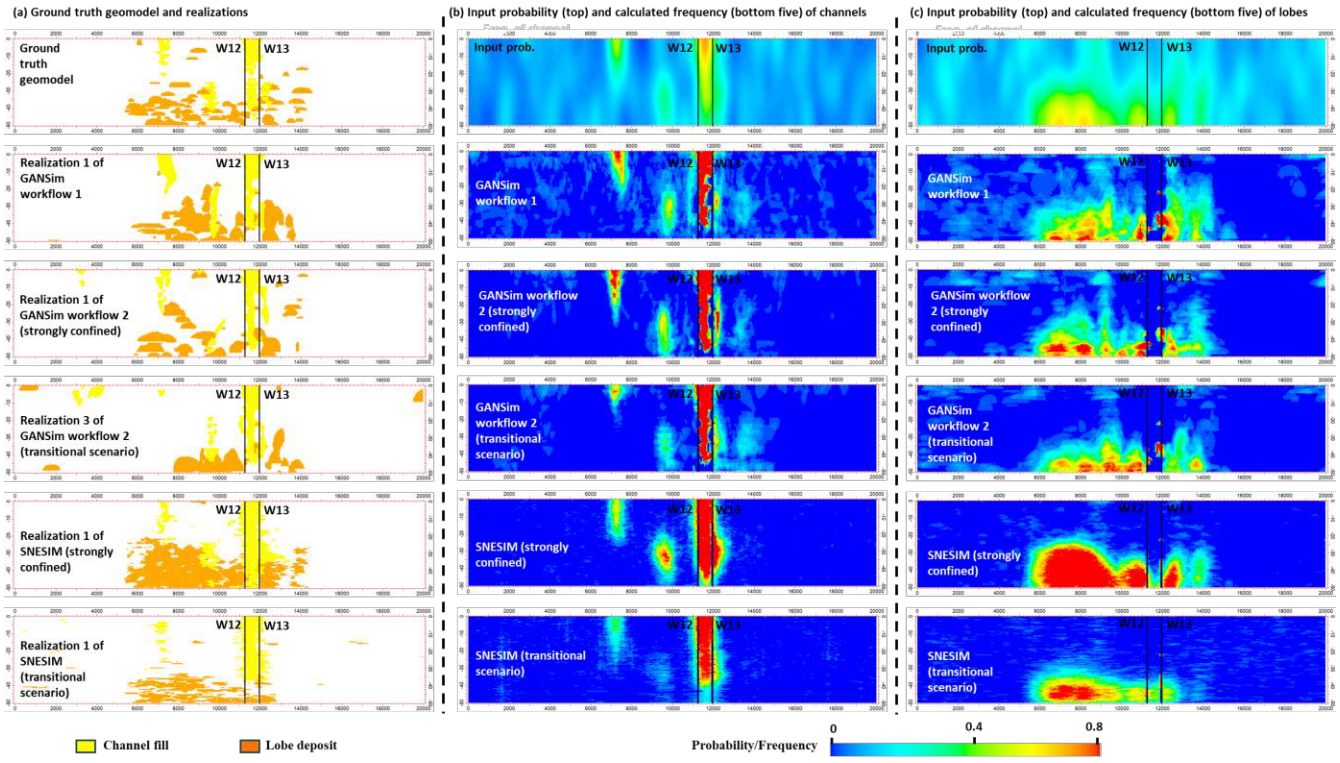

Figure 11. (a) Comparison of the ground truth facies model and random facies model realizations produced by different approaches in section 1 crossing Well 12 and Well 13, which is marked in Figure 10 (a). (b) Comparison of the input probability cube of channels and the calculated frequency cubes of channels resulting from different approaches in the same section. (c) The same as (b), but for lobes.

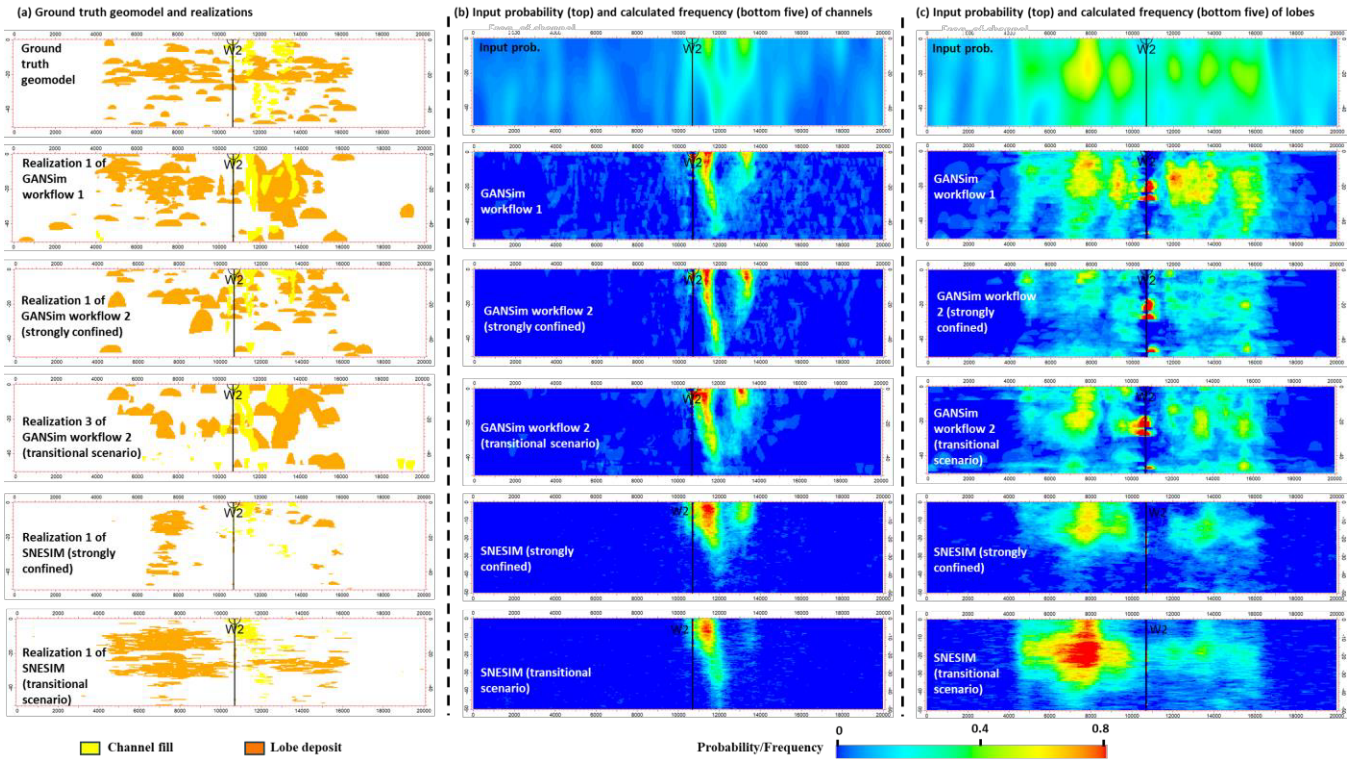

Figure 12. The same as Figure 11, but for section 2 crossing Well 2, which is marked in Figure 10 (a).

6. Geomodelling of turbidite reservoir with GANSim workflow 2

Unlike GANSim workflow 1, in workflow 2, after making sure that the prior conceptual geomodels of all scenarios as a whole are not falsified (which is not required for this synthetic case), an additional scenario falsification process is performed to test the five scenarios using the given conditioning well facies and seismic data, based on the procedures outlined in Section 2.2. In the following sections, we perform the falsification process for the five independent scenarios. Then, for each unfalsified scenario, a GANSim is trained, and the trained generator is used for geomodelling of the large reference reservoir.

6.1. Falsification of the five geological scenarios

Since the conceptual geomodels (25 m thick) are half the thickness of the reference reservoir (50 m thick), two conceptual geomodels of the same scenario are stacked to match the thickness of the reference reservoir. We then sampled well facies profiles from 120 stacked conceptual geomodels of each scenario, using the same well locations as the reference conditioning well data in Figure 5. Following the workflow in Figure 2, after multi-point histogram (MPH) analysis, multi-dimensional scaling (MDS), and likelihood evaluation, we calculated the posterior probability $P\left(Sc_k \middle| f_{MPH}\left(w_{obs}^i\right)\right)$, where Sc_k $(1 \le k \le 5)$ refers to one of the five scenarios and w_{obs}^i $(1 \le i \le 18)$ refers to one of the 18 conditioning well facies data. The calculated posterior probability values for each conditioning well are listed in Table S5-1 of Supporting Information S5. To identify which wells are more informative, we computed the Kullback-Leibler (KL) divergence (Kullback & Leibler, 1951) between the prior ($P(Sc_k)$ =1/5 for each scenario) and the posterior probabilities for each conditioning well. A higher KL

divergence indicates that the posterior probability deviates more from the prior probability, meaning the well is more informative. Well 8, 13, 16, and 17 have a large KL divergence (above 1), indicating that these four wells are the most sensitive to geological scenarios (see Table S5-1). We combined the posterior probability values of the four most informative wells into a single posterior probability value for each scenario using the Tau model (Journel, 2002). A brief description of Tau model is shown in Supporting Information S6. Since the four wells are located far from each other, independence is assumed (i.e., Tau = 1) when applying the Tau model. The posterior probabilities for the four wells and their combined posterior probabilities are shown in Table 1.

Table 1. Posterior probabilities of different reservoir scenarios for Well 8, 13, 16, and 17 and the combined posterior probability values.

	Strongly confined	Transitional	Valley confined	Weakly confined	Unconfined
Well 8	0.24	0.55	0.1	0.1	0.01
Well 13	0.35	0.42	0.01	0.19	0.03
Well 16	0.49	0.24	0.1	0.17	0
Well 17	0.18	0.57	0.15	0.1	0
Combined	0.41	0.56	0.01	0.02	0

Following the workflow in Figure 2, a forward simulation procedure similar to that used for the ground truth geomodel was performed for 160 stacked conceptual geomodels of each scenario to obtain the corresponding seismic data. Supporting Information S4 describes this procedure and settings of related parameters with more details. Unlike the application in Scheidt et al. (2015b), where the seismic data was 2D, here we have 3D seismic data. While 3D discrete wavelet transform (DWT) analysis could capture features related to 3D spatial patterns, it would be memory- and CPU-intensive. To make the DWT procedure more efficient, we limited the seismic cube to a smaller cube (with 200×200 cells horizontally; see Figure S4-1) around the center of the domain and used only 2D east-west vertical slices (perpendicular to channel directions). These seismic slices preserve effective spatial patterns related to different geological scenarios. We extracted 25 equally spaced seismic slices for 2D DWT analysis. Experiments showed that using more slices did not result in significant changes in the output posterior. By evaluating the likelihood in MDS space, the posterior probability of each scenario given the DWT features of the conditioning seismic data, $P(Sc_k|f_{DWT}(seis_{obs}))$, was obtained. The posterior probabilities are 0.32, 0.18, 0.25, 0.17, and 0.08 for the strongly confined, transitional, valley-confined, weakly confined, and unconfined scenarios, respectively.

Finally, the Tau model was used to combine the posterior probability values resulting from the conditioning well facies (combined) and seismic data. Independence between the MPH features of well facies data and the DWT features of seismic data was assumed in the Tau model (i.e., Tau = 1), as the former reveals local vertical facies stacking patterns around wells, while the latter represents the spatial patterns of seismic data. The joint posterior probabilities for the five scenarios, from strongly confined to unconfined scenarios, conditioned on the features of the conditioning well facies and seismic data,

are 0.51, 0.47, 0.01, 0.01, and 0, respectively. Thus, the latter three scenarios—valley-confined, weakly confined, and unconfined—are falsified, leaving only the strongly confined and transitional scenarios.

6.2. GANSim training and geomodelling for unfalsified scenarios

We train GANSim model for each of the two unfalsified scenarios using the training dataset of corresponding scenarios prepared in Section 5. The GANSim architectures and other hyperparameter settings are the same as in workflow 1. Training the generator for each scenario took 72 hours.

After evaluating the trained generators, we take the conditioning data of the large reference reservoir as inputs into the trained generators to produce different facies model realizations for each of the two unfalsified scenarios (the strongly confined scenario and the transitional scenario). Facies frequency cubes were calculated from 50 random realizations for each scenario. Figure 10 (b) and (c) shows three random realizations and facies frequency cubes for each scenario. Figure 11 and Figure 12 (the third and the forth row) illustrate one realization and the frequency distributions along two cross-well sections for each of the two scenarios.

7. Results analyses and discussions

7.1. GANSim results analyses

Both GANSim workflows achieve 100% accuracy in reproducing the conditioning well facies data. The simulated channels and lobes are consistent with the corresponding input probability cubes, which is confirmed by visual comparisons between the simulated realizations (or calculated frequency cubes) and the input probability cubes in both 3D (Figure 5 with Figure 10) and 2D (Figure 11 and Figure 12). The simulation of each realization (400×400×100 cells) takes 0.7 seconds on one GPU (A100) using the pretrained generator in both workflows.

Regarding the reproduction of geological patterns and scenarios, we can see from the comparison of the simulated realizations with the ground truth facies model or conceptual geomodels (Figure 4) that the GANSim results from both workflows are generally realistic in terms of the shape of channels and lobes, the connectivity of channels, the divergent distribution style of channels and lobes along the sedimentary direction (southward), and the cutting relationships between channels and lobes (channels may cut through lobes), although a few disconnected channel fragments appear in regions with locally higher input channel probabilities, e.g., the channel fragments pointed out by the blue arrows in Figure 10. Furthermore, the channels in all realizations exhibit near-vertical stacking pattern, suggesting either the strongly confined or the transitional scenario for these realizations (Figures 8, Figure 11, and Figure 12). Note that the ground truth geomodel corresponds to the strongly confined scenario. Such a reduction of scenario uncertainty from the prior five scenarios to the posterior two scenarios stems primarily the conditioning effect of the given data in both workflows. The reduced posterior scenario types are the same for both workflows, indicating the same scenario falsification effect. Additionally, the improved local discriminator design in GANSim eliminates the disconnection issue of single-pixel well facies data (see Figure 12 for example), which was a significant issue in previous GANSim versions (Hu et al., 2024; Song et al., 2021b). GANSim with a local discriminator design has also been applied to the geomodelling of a field-scale 3D delta reservoir for practical CO_2 storage (Alqassab et al., 2024). It achieved the expected geological realism—significantly better than the results obtained using MPS—and demonstrated effective conditioning.

In stationary reservoirs, geological patterns repeat spatially, so GANSim training with cropped conceptual geomodel patches was successful in stationary cases (Song et al., 2022b). In non-stationary reservoirs, geological patterns vary across the domain. However, directly training on large conceptual geomodels (each containing the full non-stationary pattern) would be computationally prohibitive. Thus, in this study, we continue to use cropped conceptual geomodel patches for training in both workflows. Although the generator can only learn fragments of the entire non-stationary pattern from cropped patches, when geomodelling with the trained generator for large reservoirs, the conditioning effect of the spatially distributed probability cube serves as a guide to "stitch" the learned pattern fragments into appropriate locations. As a result, the simulated realizations globally present realistic and complete non-stationary patterns, as we can see from Figure 10. A few pattern fragments are occasionally stitched into unintended locations, e.g., the channels indicated by the green arrows in Figure 10, which contradict with the globally divergent style and should have been developed in the eastern part of the domain. In addition to conditioning of probability cubes, when global features (e.g., channel direction and facies proportion) are treated as another type of conditioning data, their spatial distributions can also guide the trained generator to produce non-stationary geological patterns, as proved by Song et al. (2025b). This is similar to the idea in SNESIM where a map of varying affinity and rotation values is taken as a constraint to produce nonstationary geomodels (Mariethoz & Caers, 2014).

Figures 8–11 show that the high-value regions in the frequency cubes from both GANSim workflows are more concentrated than in the input probability cubes, indicating a reduced reservoir prediction uncertainty and a mitigation effect for noises in the probability cubes possibly inherited from geophysical data. These effects in uncertainty reduction and noises mitigation mainly results from the integration of geological patterns (e.g., connectivity features and stacking patterns of channels) of the trained generator.

7.2. Comparison of two GANSim workflows

For the majority of results in workflow 1, geological patterns across one realization are consistent, i.e., they belonging to one fixed scenario. However, in very few realizations, patterns at different locations may be sampled from different scenarios. For example, in Figure 13, the channels in the central region of the realization present the characteristics of strongly confined or transitional scenario, but the channels in the west region, indicated by the blue dashed curve in both the 3D geomodel and the 2D section, basically present the characteristics of the valley-confined or weakly confined scenario (i.e., channels are loosely confined), leading to mixed patterns from different scenarios in one geomodel. This occurs because the generator in workflow 1 learns pattern fragments from all scenarios, and during geomodelling, the scenario refinement is based on constraints of conditioning data. Thus, in regions where conditioning effect is weaker (e.g., the western region of the realization in Figure 13 without any well facies data), pattern fragments from more scenarios may be generated; in contrast, in regions with stronger conditioning effect (e.g., the middle region of the realization in Figure 13 containing 17 wells), pattern fragments from fewer scenarios may be generated. Such occasional scenario mixing issue is also due to training with cropped patches. In contrast, workflow 2 avoids the scenario mixing issue since its generator was trained exclusively on a single unfalsified scenario.

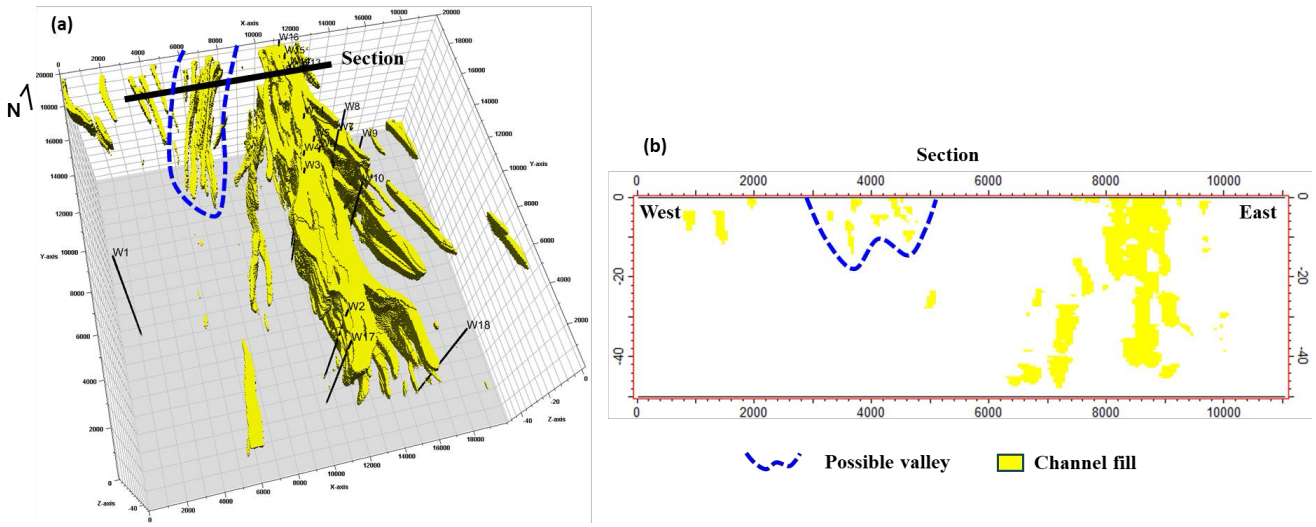

Figure 13. One 3D facies model realization resulting from GANSim workflow 1 (a) and a vertical section of the 3D realization (b), showing an example of scenario mixing in workflow 1. Lobe and mud facies are set invisible. The location of the section is marked in the 3D geomodel. The channels in the central region of the realization present the characteristic of strongly confined or transitional scenario, but the channels in the west region, indicated by the blue dashed curve in both the 3D geomodel and the 2D section, present characteristics of the valley-confined or weakly confined scenario (i.e., channels are loosely confined).

Besides occasionally generating mixed-scenario geomodels, workflow 1 also requires training data for all scenarios, but once trained, the generator can be used for geomodelling of all scenarios of reservoirs. In comparison, GANSim workflow 2 only requires training data of unfalsified scenarios, but the trained generator can only be used for geomodelling of these specific scenarios. Indeed, a certain number of conceptual geomodels of each scenario (though significantly fewer than in workflow 1) are needed for the scenario falsification process in workflow 2, and a sensitive feature filter for each type of conditioning data is also necessary in the process. If the feature filter fails to adequately constrain the scenario type, many scenarios still persist. Training separate generators for each unfalsified scenario would be impractical, but resorting to a single comprehensive generator of all unfalsified scenarios would reintroduce limitations similar to those in workflow 1.

In addition, from the perspective of probability, workflow 1 expresses a straightforward process of "$P(m) \rightarrow P(m|d)$", where m and d refer to geomodel and conditioning data, while workflow 2 can be expressed as "$P(m) \rightarrow P\left(Sc_k \middle| f(d)\right) = P\left(\{m\}_k \middle| f(d)\right) \rightarrow P(m|d)$", where Sc_k refers to the k-th scenario and $f(\cdot)$ is a feature filter. The workflow 2 is a process of gradual conditioning, with the first subprocess being scenario falsification and the second one being generator conditioning. Note that the unfalsified scenarios from the first subprocess are not necessarily the final scenarios, and an implicit scenario falsification is still involved in the second generator conditioning subprocess. In comparison, scenario falsification is completely fulfilled implicitly in the workflow 1. There may be slight differences between the results of the two workflows, because workflow 1 relies only on facies probability cubes interpreted from geophysical data, while workflow 2 also uses the global patterns from the geophysical data in the initial scenario falsification subprocess.

7.3. Comparison of geomodelling results between GANSim and MPS approaches

We also use an MPS method (i.e., SNESIM in Petrel software) for geomodelling of the reference reservoir, given the conditioning well facies and facies probability cubes. Assuming the geological scenario falsification process (Section 6.1) is also performed and the strongly confined and transitional scenarios are inferred, we take a random conceptual geomodel of each of the two unfalsified scenarios as the training image (TI), along with the conditioning well facies and facies probability cubes, into the SNESIM algorithm to simulate different facies model realizations. The input hyperparameters—e.g., data search parameters, probability weight, and target fraction tolerance—are carefully tuned to produce the most realistic facies models. The TI, three random realizations, and facies frequency cubes calculated from 50 realizations of each scenario are shown in Figure 10 (d) and (e). Two cross-well sections of one realization and the facies frequency distributions are shown in Figure 11 and Figure 12 (the last two rows) for each scenario.

The SNESIM results are consistent with the conditioning data, as this is inherently designed in the algorithm. However, compared to both GANSim workflows, the SNESIM algorithm is much less effective in reproducing expected geological realism, as shown by the shapes of simulated channels and lobes in GANSim and SNESIM results. Additionally, though the SNESIM algorithm does not require extensive training on GPUs, each realization takes ~10 minutes on an 8-core i7 laptop, which is nearly 1000 times slower than geomodelling with the trained generator in GANSim. While GANSim training may require days (less with more computational resources), the trained generator enables rapid geomodelling for arbitrarily large reservoirs with similar geological patterns.

8. Conclusions

This study proposes two GANSim workflows for geomodelling multi-scenario non-stationary reservoirs. The first workflow combines all scenarios into a single “super scenario” to train a comprehensive generator, while the second incorporates an explicit scenario falsification step to train generators exclusively for unfalsified scenarios. Building on this improvement, to addresses the local disconnection problem around single-pixel well facies data existing in previous GANSim, we propose a local discriminator to scrutinize local geological realism around wells.

The proposed GANSim workflows and local discriminator design are validated using a large synthetic dataset with five geological scenarios. The results demonstrate high conditioning accuracy, with 100% reproduction of well facies data and consistency with input geophysics-interpreted probability cubes. Despite training on cropped patches, the generator produces globally realistic non-stationary patterns by leveraging probability cubes as spatial guides. GANSim with a local discriminator has also been successfully applied to a 3D field delta reservoir geomodelling for CO_2 storage, achieving superior geological realism and conditioning effects (Alqassab et al., 2024). Maps of global features, such as facies proportion or channel direction, could serve as another type of potential guidance for non-stationarity, as demonstrated by Song et al. (2025b). The introduction of the local discriminator design eliminates the local disconnection issue of single-pixel well facies data. Both GANSim workflows effectively integrate geological knowledge with diverse conditioning data types. This integration achieves two key outcomes: (1) reduction of scenario uncertainty through conditioning data constraints, and (2) decreased reservoir prediction uncertainty compared to the input probability cubes though learned geological patterns and well facies constraint. The simulation of each realization with

400×400×100 cells takes 0.7 seconds on an A100 GPU, highlighting the method's scalability for field applications.

A comparison of the two GANSim workflows reveals distinct advantages. GANSim workflow 1, which requires training data for all scenarios, enables geomodelling across all scenarios after GANSim training but may occasionally produce mixed patterns of different scenarios in one realization. In contrast, the workflow 2 only requires training data of unfalsified scenarios, but the trained generator can only be used for geomodelling of these unfalsified scenarios. While the workflow 2 avoids the pattern mixing issue, it depends on robust scenario falsification, especially scenario-sensitive feature filters for conditioning data.

When compared to traditional multi-point statistics (MPS) methods, GANSim workflows demonstrate superior geological realism in produced realizations, and they also outperform MPS in computational efficiency, i.e., nearly 1000 times faster after training. A caveat is that GANSim requires many thousands of diverse conceptual geomodels for training. This requires efficient, automated workflows to generate these conceptual geomodels.

Acknowledgments

We acknowledge the sponsors of the Stanford Center for Earth Resources Forecasting (SCERF). Some of the computing for this project was performed on the Sherlock cluster at Stanford University. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results.

References

Alqassab, H. M., Feng, M., Becker, J. A., Song, S., & Mukerji, T. (2024). MAGCS: Machine Assisted Geologic Carbon Storage (Vol. ADIPEC, p. D021S044R001). https://doi.org/10.2118/222120-MS

Avseth, P., Mukerji, T., & Mavko, G. (2005). *Quantitative seismic interpretation: Applying rock physics tools to reduce interpretation risk*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511600074

Borg, I., & Groenen, P. J. F. (2005). *Modern multidimensional scaling: Theory and applications, 2nd ed. Modern multidimensional scaling: Theory and applications, 2nd ed.*

Brooks, S., Gelman, A., Jones, G. L., & Meng, X. L. (2011). *Handbook of Markov Chain Monte Carlo*. *Handbook of Markov Chain Monte Carlo*. https://doi.org/10.1201/b10905

Chen, M., Wu, S. H., Wang, R. F., Zhang, J. J., Xie, P. F., Wang, M., et al. (2024). Sedimentary architecture of submarine channel-lobe systems under different seafloor topography: Insights from the Rovuma Basin offshore East Africa. *Petroleum Science*. https://doi.org/10.1016/j.petsci.2023.11.021

Chen, Y., & Oliver, D. S. (2013). Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. *Computational Geosciences*. https://doi.org/10.1007/s10596-013-9351-5

Deptuck, M. E., Sylvester, Z., Pirmez, C., & O'Byrne, C. (2007). Migration-aggradation history and 3-D seismic geomorphology of submarine channels in the Pleistocene Benin-major Canyon, western Niger Delta slope. *Marine and Petroleum Geology*. https://doi.org/10.1016/j.marpetgeo.2007.01.005

Deutsch, C. V. (2002). *Geoestatistical reservoir modeling*. Oxford: Oxford University Press.

Di Federico, G., & Durlofsky, L. J. (2025). Latent diffusion models for parameterization of facies-based geomodels and their use in data assimilation. *Computers & Geosciences*, *194*, 105755. https://doi.org/https://doi.org/10.1016/j.cageo.2024.105755

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 27* (pp. 10–23). https://doi.org/10.1001/jamainternmed.2016.8245

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. (2017). Improved training of Wasserstein GANs. In *In Proceedings of the 31st International Conference on Neural Information Processing Systems* (pp. 5769–5779).

Hermans, T., Caers, J., & Nguyen, F. (2014). Assessing the probability of training image-based geological scenarios using geophysical data. *Lecture Notes in Earth System Sciences*. https://doi.org/10.1007/978-3-642-32408-6_147

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*.

Hu, X., Song, S., Hou, J., Yin, Y., Hou, M., & Azevedo, L. (2024). Stochastic Modelling of Thin Mud Drapes inside Point Bar Reservoirs with ALLUVSIM-GANSim. *Water Resources Research*.

https://doi.org/10.1029/2023WR035989

Journel, A. G. (2002). Combining knowledge from diverse sources: An alternative to traditional data independence hypotheses. *Mathematical Geology*, *34*(5), 573–596. https://doi.org/10.1023/A:1016047012594

Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of GANs for improved quality, stability, and variation. *ArXiv Preprint*, arXiv: 1710.10196.

Kingma, D., & Ba, J. (2014). Adam: a method for stochastic optimization. *ArXiv Preprint*, arXiv:1412.6980. https://doi.org/10.1109/ICCE.2017.7889386

Kullback, S., & Leibler, R. A. (1951). On Information and Sufficiency. *The Annals of Mathematical Statistics*. https://doi.org/10.1214/aoms/1177729694

Laloy, E., Hérault, R., Jacques, D., & Linde, N. (2018). Training-image based geostatistical inversion using a spatial Generative Adversarial Neural Network. *Water Resources Research*, *54*(1), 381–406. https://doi.org/10.1002/2017WR022148

Lee, D., Ovanger, O., Eidsvik, J., Aune, E., Skauvold, J., & Hauge, R. (2025). Latent diffusion model for conditional reservoir facies generation. *Computers & Geosciences*, *194*, 105750. https://doi.org/https://doi.org/10.1016/j.cageo.2024.105750

Lin, J. (1991). Divergence Measures Based on the Shannon Entropy. *IEEE Transactions on Information Theory*. https://doi.org/10.1109/18.61115

Mariethoz, G., & Caers, J. (2014). *Multiple-point geostatistics: stochastic modeling with training images*. New York: John Wiley & Sons. https://doi.org/10.1002/9781118662953

McHargue, T., Pyrcz, M. J., Sullivan, M. D., Clark, J. D., Fildani, A., Romans, B. W., et al. (2011). Architecture of turbidite channel systems on the continental slope: Patterns and predictions. *Marine and Petroleum Geology*. https://doi.org/10.1016/j.marpetgeo.2010.07.008

McHargue, T. R. (1991). Seismic Facies, Processes, and Evolution of Miocene Inner Fan Channels, Indus Submarine Fan. https://doi.org/10.1007/978-1-4684-8276-8_22

McHargue, T. R., Hodgson, D. M., & Shelef, E. (2021). Architectural Diversity of Submarine Lobate Deposits. *Frontiers in Earth Science*. https://doi.org/10.3389/feart.2021.697170

Mo, S., Zabaras, N., Shi, X., & Wu, J. (2020). Integration of Adversarial Autoencoders With Residual Dense Convolutional Networks for Estimation of Non-Gaussian Hydraulic Conductivities. *Water Resources Research*. https://doi.org/10.1029/2019WR026082

Mosser, L., Dubrule, O., & Blunt, M. J. (2020). Stochastic seismic waveform inversion using Generative Adversarial Networks as a geological prior. *Mathematical Geosciences*, *52*(1), 53–79. https://doi.org/10.1007/s11004-019-09832-6

Mutti, E., & Lucchi, F. R. (1978). Turbidites of the Northern Apennines: Introduction to facies analysis. *International Geology Review*. https://doi.org/10.1080/00206817809471524

Nesvold, E., & Mukerji, T. (2021). Simulation of fluvial patterns with GANs trained on a data set of satellite imagery. *Water Resources Research*, *57*(22). https://doi.org/10.1029/2019WR025787.

Park, H., Scheidt, C., Fenwick, D., Boucher, A., & Caers, J. (2013). History matching and uncertainty quantification of facies models with multiple geological interpretations. *Computational Geosciences*. https://doi.org/10.1007/s10596-013-9343-5

Popper, K. (1959). *The logic of scientific discovery*.

Posamentier, H. W., & Kolla, V. (2003). Seismic geomorphology and stratigraphy of depositional elements in deep-water settings. *Journal of Sedimentary Research*. https://doi.org/10.1306/111302730367

Prélat, A., Hodgson, D. M., & Flint, S. S. (2009). Evolution, architecture and hierarchy of distributary deep-water deposits: a high-resolution outcrop investigation from the Permian Karoo Basin, South Africa. *Sedimentology*. https://doi.org/10.1111/j.1365-3091.2009.01073.x

Prélat, A., Covault, J. A., Hodgson, D. M., Fildani, A., & Flint, S. S. (2010). Intrinsic controls on the range of volumes, morphologies, and dimensions of submarine lobes. *Sedimentary Geology*. https://doi.org/10.1016/j.sedgeo.2010.09.010

Pyrcz, M. J., & Deutsch, C. V. (2014). *Geostatistical Reservoir Modeling* (2nd ed.). Oxford University Press.

Scheidt, C., Jeong, C., Mukerji, T., & Caers, J. (2015). Probabilistic falsification of prior geologic uncertainty with seismic amplitude data: Application to a turbidite reservoir case. *Geophysics*. https://doi.org/10.1190/geo2015-0084.1

Scheidt, C., Tahmasebi, P., Pontiggia, M., Da Pra, A., & Caers, J. (2015). Updating joint uncertainty in trend and depositional scenario for reservoir exploration and early appraisal. *Computational Geosciences*. https://doi.org/10.1007/s10596-015-9491-x

Scheidt, C., Li, L., & Caers, J. (2018). *Quantifying uncertainty in subsurface systems*. John Wiley & Sons.

Song, S., Mukerji, T., & Hou, J. (2021a). Geological Facies modeling based on progressive growing of generative adversarial networks (GANs). *Computational Geosciences*. https://doi.org/10.1007/s10596-021-10059-w

Song, S., Mukerji, T., & Hou, J. (2021b). GANSim: Conditional Facies Simulation Using an Improved Progressive Growing of Generative Adversarial Networks (GANs). *Mathematical Geosciences*. https://doi.org/10.1007/s11004-021-09934-0

Song, S., Mukerji, T., & Hou, J. (2022a). Bridging the Gap between Geophysics and Geology with Generative Adversarial Networks. *IEEE Transactions on Geoscience and Remote Sensing*. https://doi.org/10.1109/TGRS.2021.3066975

Song, S., Mukerji, T., Hou, J., Zhang, D., & Lyu, X. (2022b). GANSim-3D for conditional geomodelling: theory and field application. *Water Resources Research*, *58*. https://doi.org/doi.org/10.1029/2021WR031865

Song, S., Zhang, D., Mukerji, T., & Wang, N. (2023). GANSim-surrogate: An integrated framework for stochastic conditional geomodelling. *Journal of Hydrology*. https://doi.org/10.1016/j.jhydrol.2023.129493

Song, S., Mukerji, T., & Zhang, D. (2025a). Physics-informed multi-grid neural operator: Theory and an

application to porous flow simulation. *Journal of Computational Physics*, *520*, 113438. https://doi.org/https://doi.org/10.1016/j.jcp.2024.113438

Song, S., Huang, J., & Mukerji, T. (2025b). Generative geomodelling: Deep Learning vs. Geostatistics. Under review.

Spychala, Y. T., Eggenhuisen, J. T., Tilston, M., & Pohl, F. (2020). The influence of basin setting and turbidity current properties on the dimensions of submarine lobe elements. *Sedimentology*. https://doi.org/10.1111/sed.12751

Tan, X., Tahmasebi, P., & Caers, J. (2014). Comparing training-image based algorithms using an analysis of distance. *Mathematical Geosciences*. https://doi.org/10.1007/s11004-013-9482-1

Tarantola, A. (2006). Popper, Bayes and the inverse problem. *Nature Physics*. https://doi.org/10.1038/nphys375

Wellner, R., Beaubouef, R., Van Wagoner, J., Roberts, H. H., & Sun, T. (2005). Jet-plume depositional bodies; the primary building blocks of Wax Lake Delta. *Transactions - Gulf Coast Association of Geological Societies*, *55*, 867–909.

William R. Normark. (1978). Fan Valleys, Channels, and Depositional Lobes on Modern Submarine Fans: Characters for Recognition of Sandy Turbidite Environments. *AAPG Bulletin*. https://doi.org/10.1306/c1ea4f72-16c9-11d7-8645000102c1865d

Xu, M., Song, S., & Mukerji, T. (2024). DiffSim: Denoising diffusion probabilistic models for generative facies geomodeling. In *Fourth International Meeting for Applied Geoscience & Energy* (pp. 1660–1664). https://doi.org/10.1190/image2024-4081304.1

Zhang, T., Switzer, P., & Journel, A. (2006). Filter-based classification of training image patterns for spatial simulation. *Mathematical Geology*, *38*(1), 63–80. https://doi.org/10.1007/s11004-005-9004-x

Zhang, T. F., Tilke, P., Dupont, E., Zhu, L. C., Liang, L., & Bailey, W. (2019). Generating geologically realistic 3D reservoir facies models using deep learning of sedimentary architecture with generative adversarial networks. *Petroleum Science*, *16*(3), 541–549. https://doi.org/10.1007/s12182-019-0328-4

Supporting Information for the paper of

Geomodelling of multi-scenario non-stationary reservoirs with enhanced GANSim

Suihong Song[1,*], Tapan Mukerji[1], Celine Scheidt[1], Hisham Alqassab[2], Man Feng[2]

[1] Stanford University, 367 Panama St, Stanford, CA 94305, USA

[2] ExxonMobil Technology & Engineering Company, Spring, Texas, United States

Corresponding author: Suihong Song, suihong@stanford.edu

Supporting Information S1

Basics of GANs

The Generative Adversarial Networks (GANs) framework was first introduced by Goodfellow in 2014. It consists of two neural networks: the generator (*G*) and the discriminator (*D*). The generator is designed to map a low-dimensional latent vector into a sample, while the discriminator maps a generated (fake) or real sample into a value representing the probability of the input being real. This framework is trained using a dataset of observed real samples.

The loss function for the original (vanilla) GANs (Goodfellow et al., 2014) is defined as:

$$L(G,D) = \mathbb{E}_{x_r \sim p_{data}}[logD(x_r)] + \mathbb{E}_{z \sim p_z}\left[\log\left(1 - D\left(G(z)\right)\right)\right], \quad \text{(S1)}$$

where, *L* is the GANs loss, *G* and *D* represent the generator and discriminator neural networks, respectively, p_{data} is the distribution of real samples, x_r is a real sample from the training dataset, z is a latent vector sampled from the distribution of p_z, and $\mathbb{E}$ denotes the expectation operator.

The discriminator's final activation function is typically a sigmoid function, which outputs a probability value between 0 and 1. During training, the discriminator and generator are optimized alternately: the discriminator maximizes the loss function to better distinguish real from fake samples, while the generator minimizes the loss to produce more realistic samples. This adversarial process forces the generator to learn the underlying patterns in the data. Training concludes when the generator produces samples so realistic that the discriminator can no longer reliably differentiate between real and fake samples. Once trained, the generator is retained for practical generative tasks. While the vanilla GAN loss function has been widely used, several alternative loss functions have been proposed in recent years (Lucic et al., 2017). Among these, the Wasserstein loss with gradient penalty (Gulrajani et al., 2017) has demonstrated superior performance, offering improved stability and convergence.

Traditionally, all layers of the generator and discriminator are trained simultaneously, without considering the scale of the patterns being learned. However, Karras (2017) introduced the progressive GANs training approach (also known as progressive growing of GANs), where the layers of the generator and discriminator are trained incrementally, allowing the model to learn patterns gradually from coarse to fine scales. This progressive approach has been shown to outperform traditional training methods in terms of training speed, stability, and the quality of generated results.

GANs have been applied to geomodelling of subsurface reservoirs using both traditional training approaches (Chan & Elsheikh, 2017, 2019; Dupont et al., 2018; Laloy et al., 2018; Mosser et al., 2020; Nesvold & Mukerji, 2021; Zhang et al., 2019) and the progressive training approach (Song et al., 2021). Regardless of the method, the generator learns geological patterns from the provided training facies models. Once trained, the generator can produce facies models that are consistent with the learned patterns, enabling unconditional geomodelling.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative Adversarial Networks. In Advances in Neural Information Processing Systems 27 (pp. 10–23). https://doi.org/10.1001/jamainternmed.2016.8245

Lucic, M., Kurach, K., Michalski, M., Gelly, S., & Bousquet, O. (2017). Are GANs created equal? A large-scale study. In *In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2467–2476).

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. (2017). Improved training of Wasserstein GANs. In In Proceedings of the 31st International Conference on Neural Information Processing Systems (pp. 5769–5779).

Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of GANs for improved quality, stability, and variation. ArXiv Preprint, arXiv: 1710.10196.

Chan, S., & Elsheikh, A. H. (2017). Parametrization and generation of geological models with generative adversarial networks. ArXiv Preprint, arXiv: 1708.01810.

Chan, S., & Elsheikh, A. H. (2019). Parametric generation of conditional geological realizations using generative neural networks. Computational Geosciences, 23(5), 925–952. https://doi.org/10.1007/s10596-019-09850-7

Dupont, E., Zhang, T., Tilke, P., Liang, L., & Bailey, W. (2018). Generating realistic geology conditioned on physical measurements with Generative Adversarial Networks. In Stat (pp. 8–16).

Laloy, E., Hérault, R., Jacques, D., & Linde, N. (2018). Training-image based geostatistical inversion using a spatial Generative Adversarial Neural Network. Water Resources Research, 54(1), 381–406. https://doi.org/10.1002/2017WR022148

Mosser, L., Dubrule, O., & Blunt, M. J. (2020). Stochastic seismic waveform inversion using Generative Adversarial Networks as a geological prior. Mathematical Geosciences, 52(1), 53–79. https://doi.org/10.1007/s11004-019-09832-6

Nesvold, E., & Mukerji, T. (2021). Simulation of fluvial patterns with GANs trained on a data set of satellite imagery. Water Resources Research, 57(22). https://doi.org/10.1029/2019WR025787.

Song, S., Mukerji, T., & Hou, J. (2021). Geological Facies modeling based on progressive growing of generative adversarial networks (GANs). *Computational Geosciences*. https://doi.org/10.1007/s10596-021-10059-w

Zhang, T. F., Tilke, P., Dupont, E., Zhu, L. C., Liang, L., & Bailey, W. (2019). Generating geologically realistic 3D reservoir facies models using deep learning of sedimentary architecture with generative adversarial networks. Petroleum Science, 16(3), 541–549. https://doi.org/10.1007/s12182-019-0328-4

Supporting Information S2

GANSim: a direct conditional simulation approach based on improved GANs

To achieve conditional geomodels—generated facies models that are consistent with both the expected geological patterns and the given conditioning data (e.g., well data)—the generator must accomplish two key tasks. First, it must incorporate various types of conditioning data as inputs. Second, it must learn the relationship between the input conditioning data and the output facies model, in addition to capturing the underlying geological patterns. This relationship, termed conditioning ability, is critical for achieving conditional geomodelling. Song et al. (2021b, 2022a) proposed the GANSim framework to directly train such a generator using a progressive training method.

In GANSim, three types of conditioning data are considered: non-spatial global features of reservoirs (e.g., facies proportion and channel sinuosity), sparse well facies interpretations, and spatially distributed probability maps of all facies obtained from geophysical interpretation or geological understanding. Figure S2-1 illustrates the input pipelines for these three types of conditioning data in an example of producing 2D facies models of sinuous channels. The conditioning data related to global features are concatenated with the input latent vector and passed through all layers of the generator. Well facies data are first downsampled into various progressive resolutions (e.g., 8×8) and then converted into feature cubes (e.g., a feature cube of size 8×8×16) using 1×1 convolutional layers. These feature cubes are then concatenated with the feature cubes at the same resolution (e.g., a feature cube of size 8×8×128) in the main pipeline of the generator. The input pipeline for probability maps follows the same design as that for well data.

The design of concatenating global feature values with the input latent vector proved effective for simple global features such as mud proportion, channel width, and channel sinuosity in Song et al. (2021). However, Song et al. (2025) found that this approach becomes challenging when dealing with complex global features such as channel direction. Through experimentation, they discovered that a design similar to the input pipeline for well facies and probability maps—where global features are concatenated into multiple feature cubes—is significantly more efficient. This improvement is attributed to the flexible information transmission mechanism from shallow to deep layers, as validated in Song et al. (2025).

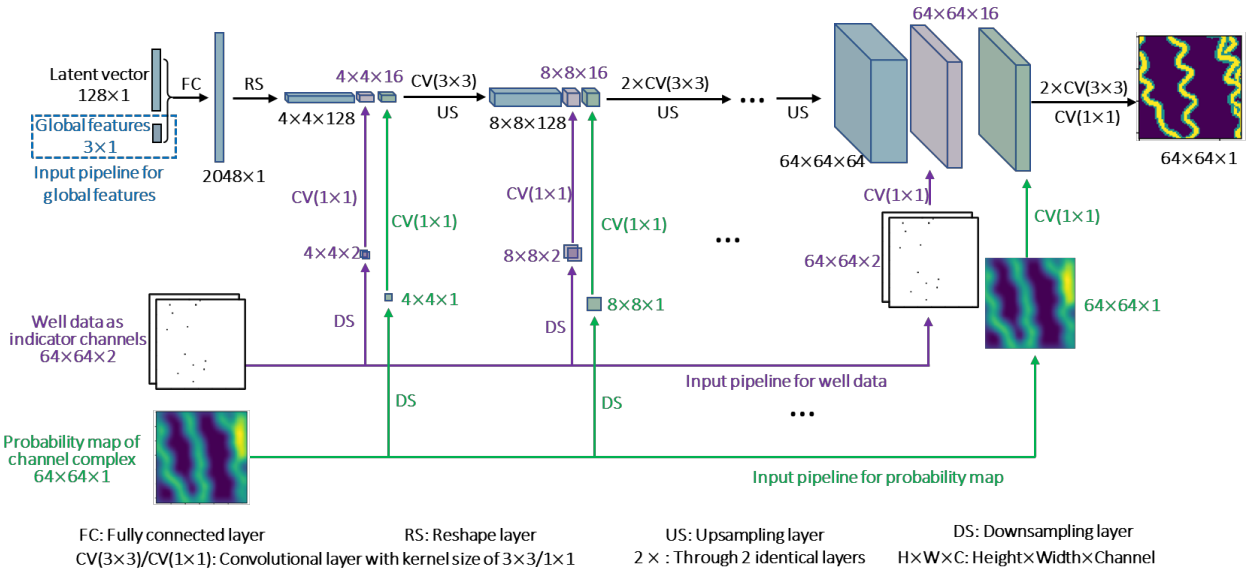

Figure S2-1. Input pipeline setup of generator for three types of common conditioning data (non-spatial global features, sparse well facies data, and probability maps) in an example of producing 2D facies models of sinuous channels (modified from Song et al., 2022). Global features (e.g., facies proportion) are concatenated with the input latent vector. Well data are downsampled into various resolutions and converted into feature cubes using 1×1 convolutional layers, which are then concatenated with feature cubes of the same resolution in the main pipeline. The input pipeline for probability maps follows the same design as that for well data.

To enable the generator to learn the conditioning ability, GANSim introduces a specially designed condition-based loss function, while retaining the original GAN loss (Equation (S1)) to ensure the learning of geological patterns. The general form of the condition-based loss function is:

$$L(G)_{con} = \mathbb{E}_{z\sim p_z, con_{in}\sim p_{con}} Dist(f_{con}[G(z, con_{in})], con_{in}). \quad \text{(S2-1)}$$

Here, con_{in} is the input condition (e.g., well facies data), p_{con} is the distribution of con_{in}, and $Dist$ is some type of distance function. f_{con} is a predefined inversion function that maps the generated facies model $G(z, con_{in})$ to the correct condition values. The condition-based loss $L(G)_{con}$ quantifies the inconsistency between the input condition con_{in} and the generated facies model $G(z, con_{in})$. By minimizing $L(G)_{con}$, the generator is forced to learn the relationship between the input condition and the output facies model, thereby developing conditioning ability.

According to the general form (Equation (S2-1)), the condition-based loss function of global features is specified as

$$L(G)_g = \mathbb{E}_{z\sim p_z, g\sim p_g} \parallel f_g[G(z,\ g)] - g \parallel_2, \quad \text{(S2-2)}$$

where, g is input (one or multiple types of) global features, p_g is the distribution of all possible g, and f_g maps generated facies models into the corresponding real global features. In cases where f_g is difficult to define, an additional neural network can be trained to approximate it (see Song et al., 2021b). The distance function $Dist$ is specified as the Euclidean L2 distance, i.e., $\|\cdot\|_2$.

The condition-based loss of well facies data is specified as

$$L(G)_w = \mathbb{E}_{z\sim p_z, w\sim p_w} \parallel I_{wloc} \odot [G(z,\ w)] - w \parallel_2, \tag{S2-3}$$

where w represents input well data, p_w is the distribution of w, I_{wloc} is the indicator of well locations, and $\odot$ is the element-wise product operator.

The condition-based loss of the probability map is specified as

$$L(G)_p = \mathbb{E}_{z_1, z_2, \ldots, z_m \sim p_z, p\sim p_p} \parallel f_p[G(z_1, p), G(z_2, p), \cdots, G(z_m, p)] - p \parallel_2, \tag{S2-4}$$

where, $z_1, z_2, \ldots, z_m$ are random samples of the latent vector z, p represents input probability maps for all facies types, p_p is the distribution of p, and f_p calculates the frequency map (in fraction) for each facies type from m generated facies models. Parameter m is a predefined hyperparameter. The frequency map for a facies type *F* is computed as:

$$f_p = \frac{\sum_{i=1}^{i=m} \mathbb{I}[G(z_i,\ p)=C(F)]}{m}. \tag{S2-5}$$

Here, $C(F)$ represents the code of facies type *F*, and $\mathbb{I}(\cdot)$ is an indicator function that equals 1 if the condition is satisfied and 0 otherwise.

These three types of conditioning data are not necessarily all included depending on specific cases of observed conditioning data. The total loss is a weighted combination of the condition-based loss and the original GANs loss:

$$L(G, D)_{total} = L(G, D) + \beta_1 L(G)_w + \beta_2 L(G)_p + \beta_3 L(G)_g. \tag{S2-6}$$

Here, β_1, β_2, and β_3 are predefined weights. During training, the generator minimizes both the condition-based loss and the GAN loss, while the discriminator only maximizes the GAN loss. GANSim employs the progressive training approach. For example, in Figure S2-1, the first convolutional layer (and the fully connected layer) at the resolution of 4×4 is initially activated and trained using 4×4 conditioning well data and probability maps. Subsequently, the following two convolutional layers at the resolution of 8×8 are further activated and trained together with the previous layers with the 8×8-conditioning data. This progressive process allows the generator to learn geological patterns and conditioning ability from coarse to fine scales. After training, the generator can produce facies models that are consistent with both the geological patterns and the input conditioning data. By varying the input latent vector, multiple realizations of conditional facies models can be generated.

Song, S., Mukerji, T., & Hou, J. (2021). GANSim: Conditional facies simulation using an improved progressive growing of Generative Adversarial Networks (GANs). Mathematical Geosciences. https://doi.org/https://doi.org/10.1007/s11004-021-09934-0

Song, S., Mukerji, T., & Hou, J. (2022). Bridging the Gap between Geophysics and Geology with Generative Adversarial Networks. *IEEE Transactions on Geoscience and Remote Sensing*. https://doi.org/10.1109/TGRS.2021.3066975

Song, S., Huang, J., & Mukerji, T. (2025). Generative geomodelling: Deep Learning vs. Geostatistics. Under review.

Supporting Information S3

Emulation workflow of the ground truth geomodel and conceptual geomodels

The emulation workflow for ground truth or conceptual geomodels is implemented in Petrel software. The first part below introduces the emulation steps tailored to different scenarios, including how hyperparameters are configured for constructing conceptual geomodels. The construction of the ground truth geomodel follows the same emulation workflow but with different hyperparameter settings, which are described in the second part.

S3.1. The emulation workflow and hyperparameters for conceptual geomodel construction

(1) Strongly confined scenario

A divergent-shaped "prior" probability map, as shown in Figure S3-1(b), is used to sample channel endpoints. Then, 4 to 20 channel endpoints are randomly sampled from the prior probability map. The number of endpoints is much smaller than the number of channels to be simulated, so the sampled endpoints are shared by different channels at different depths during the simulation process, resulting in vertically stacked channel deposits. The orientation map, shown in Figure S3-1(a), ensures a divergent shape of simulated channels from north to south. We use Petrel software to program this workflow. Although the channels flow from north to south, in Petrel, we use the sampled endpoints as “source points” to simulate channels flowing from south to north. When building conceptual geomodels, the distribution ranges of some hyperparameters are shown in Figure S3-1(c)–(h). For each conceptual geomodel, one hyperparameter is randomly sampled from each of these distribution ranges and input into the geomodelling workflow.

For lobes, a fixed orientation map, as shown in Figure S3-2(a), ensures a divergent pattern of simulated lobes. One of the four probability maps (Figure S3-2(b)) is randomly selected for each conceptual geomodel to locate lobes. The distribution ranges of other hyperparameters are illustrated in Figure S3-2(c)–(g).

Finally, simulated channels and lobes are combined in each conceptual geomodel. The channel facies has a higher priority than the lobe facies if a cell is assigned both facies codes.

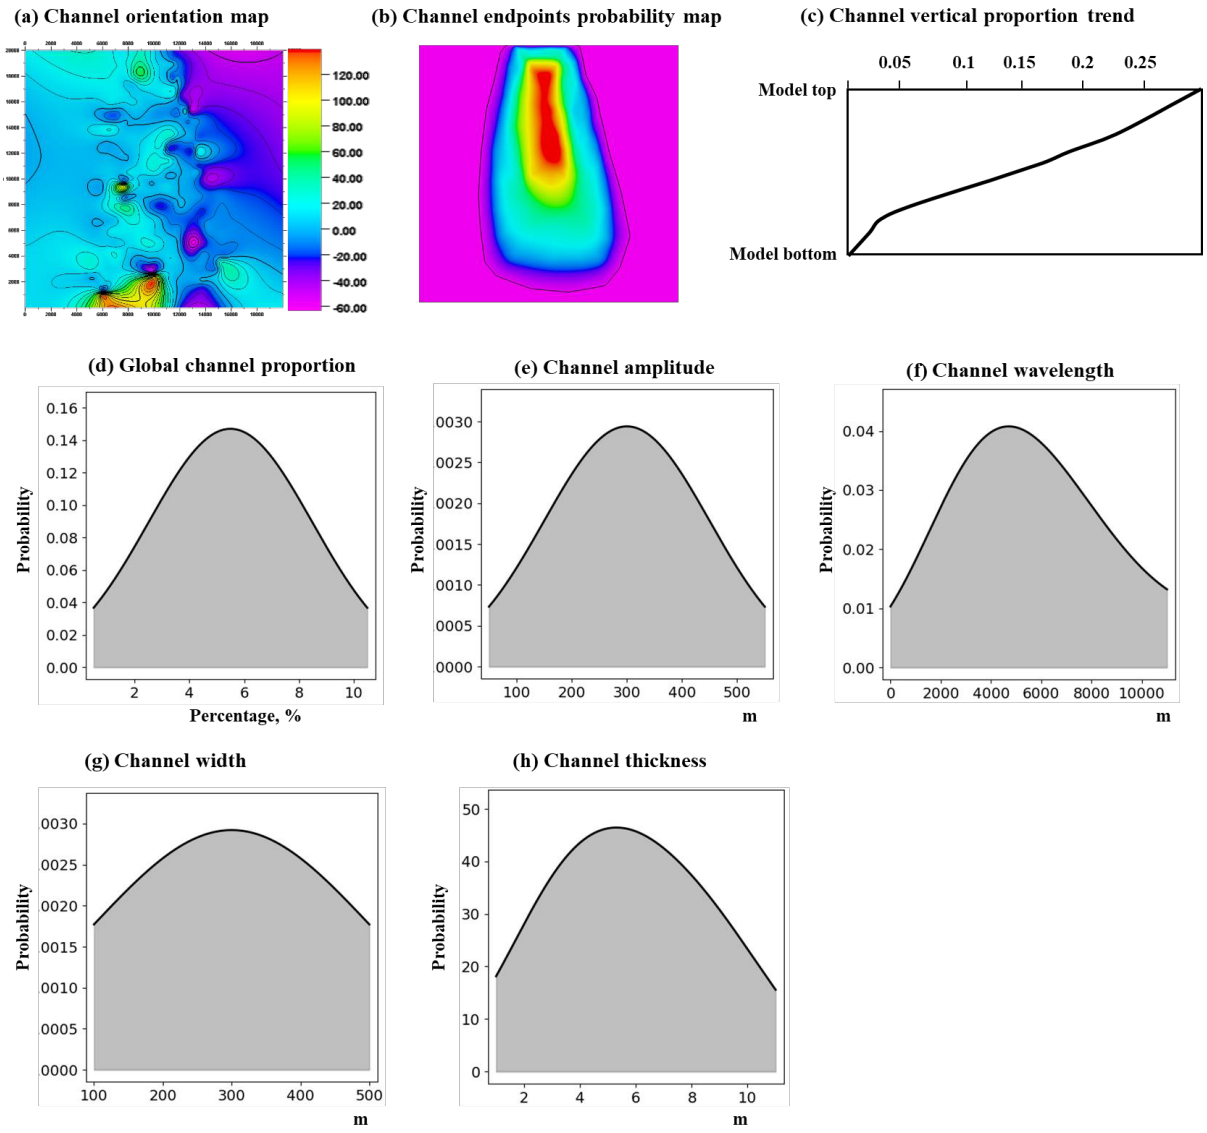

Figure S3-1. Input hyperparameters (ranges) related to channel simulation in the geomodelling workflow for the strongly confined scenario.

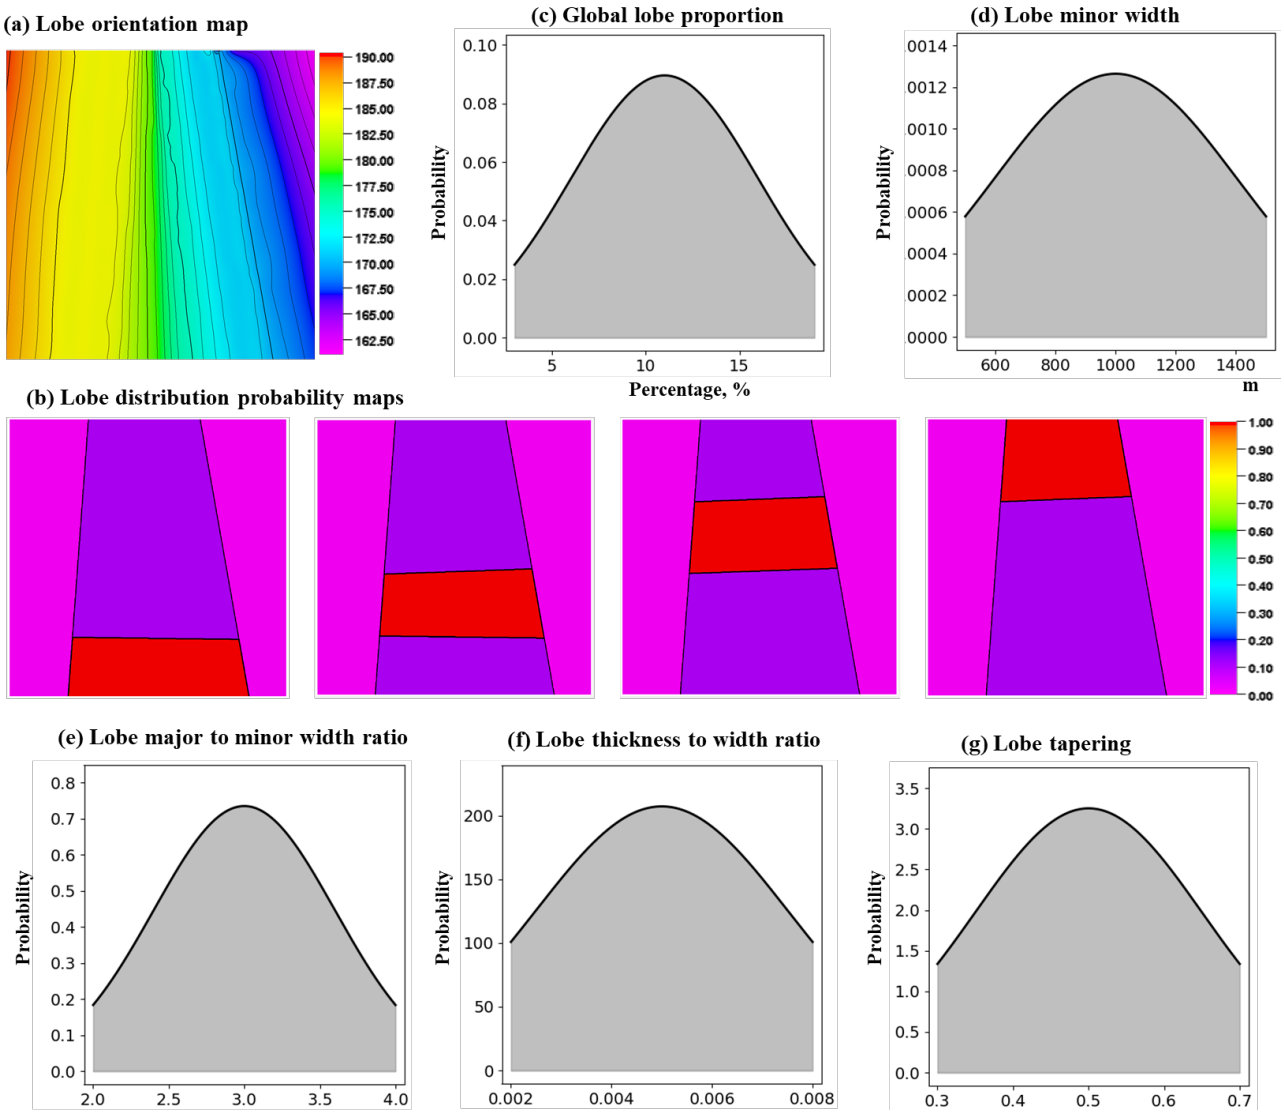

Figure S3-2. Input hyperparameters (ranges) related to lobe simulation in the geomodelling workflow for the strongly confined scenario.

(2) Valley-confined scenario

The workflow and hyperparameters for the valley-confined scenario are similar to those of the strongly confined scenario, except for the construction method of the channel endpoint probability map and the number of endpoints sampled. Figure S3-3 illustrates the workflow for constructing the channel endpoint probability map for the valley-confined scenario. For each conceptual geomodel, 3 to 6 endpoints are randomly sampled based on the prior probability map. Using these endpoints and the channel orientation map previously defined in Figure S3-1, several channel trajectories are simulated, representing the central curves of valleys. Finally, the Gaussian smoothing method with a kernel size ranging from 6 to 10 is applied to the central curve map (used as an indicator map, where cells around the central curve are assigned 1 and other cells are assigned 0) to construct the final valley-like channel endpoint probability map. Unlike the strongly confined scenario, there is no limit to the number of

channel endpoints in this scenario. Thus, simulated channels may not be stacked vertically but are confined within the simulated valleys.

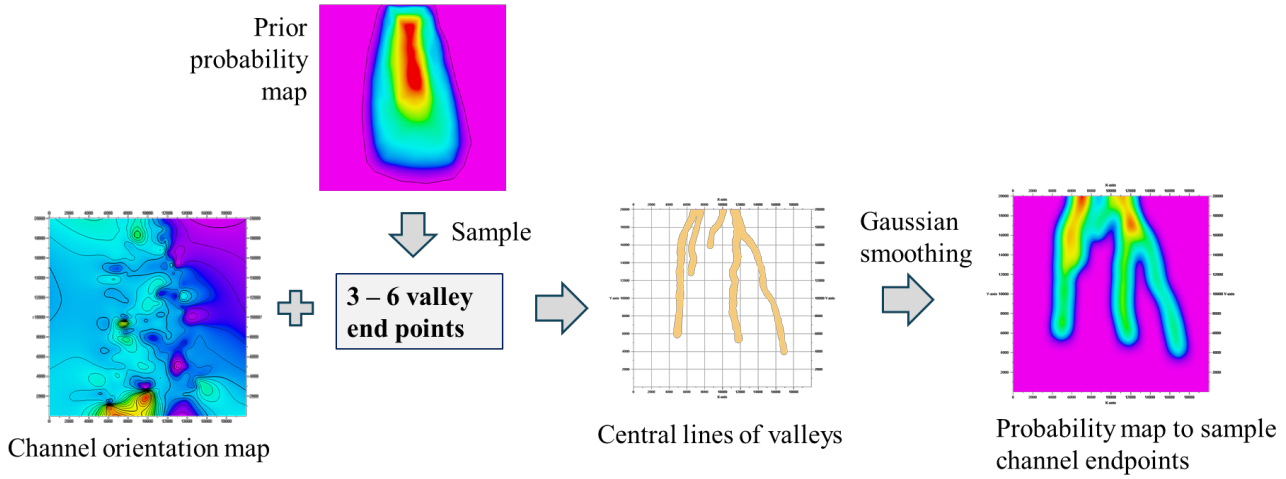

Figure S3-3. Workflow to construct the valley-like probability map of channel endpoints in the valley-confined scenario.

(3) Transitional scenario

Compared to the valley-confined scenario, the differences include: First, the Gaussian kernel size used in the construction of the channel endpoint probability map is smaller, ranging from 1 to 6; Second, the number of channel endpoints is gradually reduced to 20.

(4) Weakly confined scenario

Compared to the valley-confined scenario, the only difference lies in the construction of the channel endpoint probability map. In this scenario, the Gaussian kernel size ranges from 10 to 30 to widen the “valleys”. Additionally, the prior probability map is gradually incorporated into the probability map resulting from Gaussian smoothing to weaken the influence of valleys. This is achieved using the following equation:

$$\alpha \times prob_{Gaussian} + (1 - \alpha) \times prob_{prior} = Prob_{final}$$

where, α is an index ranging from 0 to 1 controlling the confinement effect of valleys, $prob_{Gaussian}$, $prob_{prior}$, $Prob_{final}$ and are the probability map resulting from Gaussian smoothing, the prior probability map, and the final probability map, respectively. Figure S3-4 shows an example of this process.

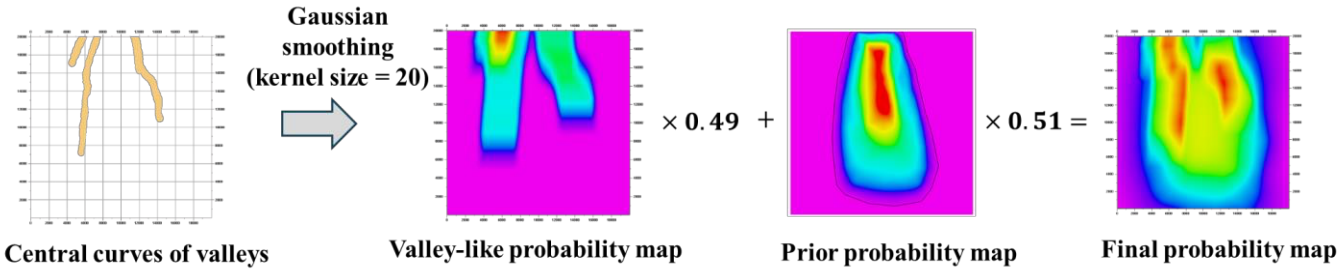

Figure S3-4. An example showing the construction process of the channel endpoint probability map in the weakly confined scenario.

(5) Unconfined scenario

In the unconfined scenario, the prior probability map is used to sample as many channel endpoints as needed. The remaining steps are the same as in the workflow of the strongly confined or valley-confined scenario.

S3.2. Hyperparameters used for construction of the ground truth geomodel

The ground truth geomodel is constructed using the emulation workflow of the strongly confined scenario. It is twice as thick as the previously emulated conceptual geomodels and is divided into two equal-thickness zones.

The channel orientation map, channel endpoint probability map (the prior map), channel vertical proportion trend, and lobe orientation map shown in Figure S3-1 and Figure S3-2 are used in this emulation process. For the lobe distribution probability maps in Figure S3-2(b), the ground truth geomodel uses the four maps from left to right for the bottom half of the bottom zone, the top half of the bottom zone, the bottom half of the top zone, and the top half of the top zone. Other hyperparameters are listed in Table S3-1.

Table S3-1. Some hyperparameters used for construction of the ground truth geomodel.

	Channel					Lobe				
	Global proportion	Amplitude	Wavelength	Width	Thickness	Global proportion	Minor width	Major-minor ratio	Thickness-width ratio	Tapering
Top zone	4%	200 m	7000 m	250 m	2 m	10%	700 m	3.5	0.0025	0.5
Bottom zone	3%	200 m	7000 m	250 m	2 m	15%	700 m	3.5	0.0025	0.5

Supporting Information S4

Forward simulation from facies geomodel to seismic data

Forward simulation from a facies geomodel to seismic data comprises three sub-processes: reservoir property simulation, statistical rock physics modeling, and normal-incidence convolution seismic simulation. These sub-processes are described below. The input facies geomodel, either the ground truth geomodel or the stacked conceptual geomodel, is 50 m thick, spanning from -5000 m to -5050 m.

(1) Reservoir property simulation

For the purpose of seismic simulation, the relevant reservoir properties include net-to-gross (NtG), porosity, permeability, and saturation. The simulation of NtG and porosity is based on the sequential Gaussian simulation method (see e.g., (Pyrcz & Deutsch, 2014)). The NtG of channel and lobe facies ranges from 0.3 to 0.9, and the porosity of these sandy facies generally ranges from 0.1 to 0.3. For permeability, an empirical relationship between permeability and porosity is used for the two sandy facies (channels and lobes): $per = 4.3944 \times e^{(19.737 * por)}$, where per and por refer to permeability and porosity, respectively. The oil-water contact (OWC) of the ground truth geomodel is set at -5040 m, while for conceptual geomodels, the OWC is sampled from a range of -5045 m to -5025 m. Based on the OWC, a semi-empirical relation involving permeability and OWC, accounting for capillary effects, is used to calculate the water saturation distribution (S_w). The form of that equation is

$$S_w = S_{wir} + \left(\frac{1-P_{nom}}{1+a \cdot P_{nom}}\right)^{\beta} \cdot (1 - S_{wir}),$$

where, P_{nom} refers to the distance above the OWC, a represents a permeability-dependent property, β is an empirical exponent, and S_{wir} refers to irreducible water saturation, calculated from the permeability property.

(2) statistical rock physics modeling

Based on the previously simulated porosity, NtG, and water saturation models, the density property is calculated using a weighted summation of different constituents, including sand, clay, water, and oil.

For channel and lobe facies, the constant-cement model (Avseth et al., 2005; Mavko et al., 2020) is used to randomly simulate the bulk modulus, shear modulus, and P-wave velocity (V_p). For mud facies, Gardner et al. (1974)'s empirical relation between P-wave velocity and density is used to calculate V_p. Initially, the simulated modulus properties and V_p are based on an assumption of 100% brine saturation. Subsequently, Gassmann's (1951) relation is applied for fluid substitution, replacing 100% brine with a mixture of brine and oil.

(3) Seismic simulation

Based on the simulated V_p and density, a normal-incidence convolution model is used to simulate post-stack seismic data. During the simulation process, we assume a signal-to-noise ratio (SNR) of 10, randomly sample 1/10 of the reflection coefficient signal as noise, and add this noise to the reflection coefficient model. A Ricker wavelet with a frequency of 40 Hz is applied. Finally, a horizontal box filter

with 3 × 3 pixels (approximately twice the wavelength) is used to smooth the simulated seismic data. The simulated seismic data comprise 400 × 400 × 72 cells, with each cell representing 0.5 ms vertically. Figure S4-1 shows examples of simulated seismic data for different scenarios.

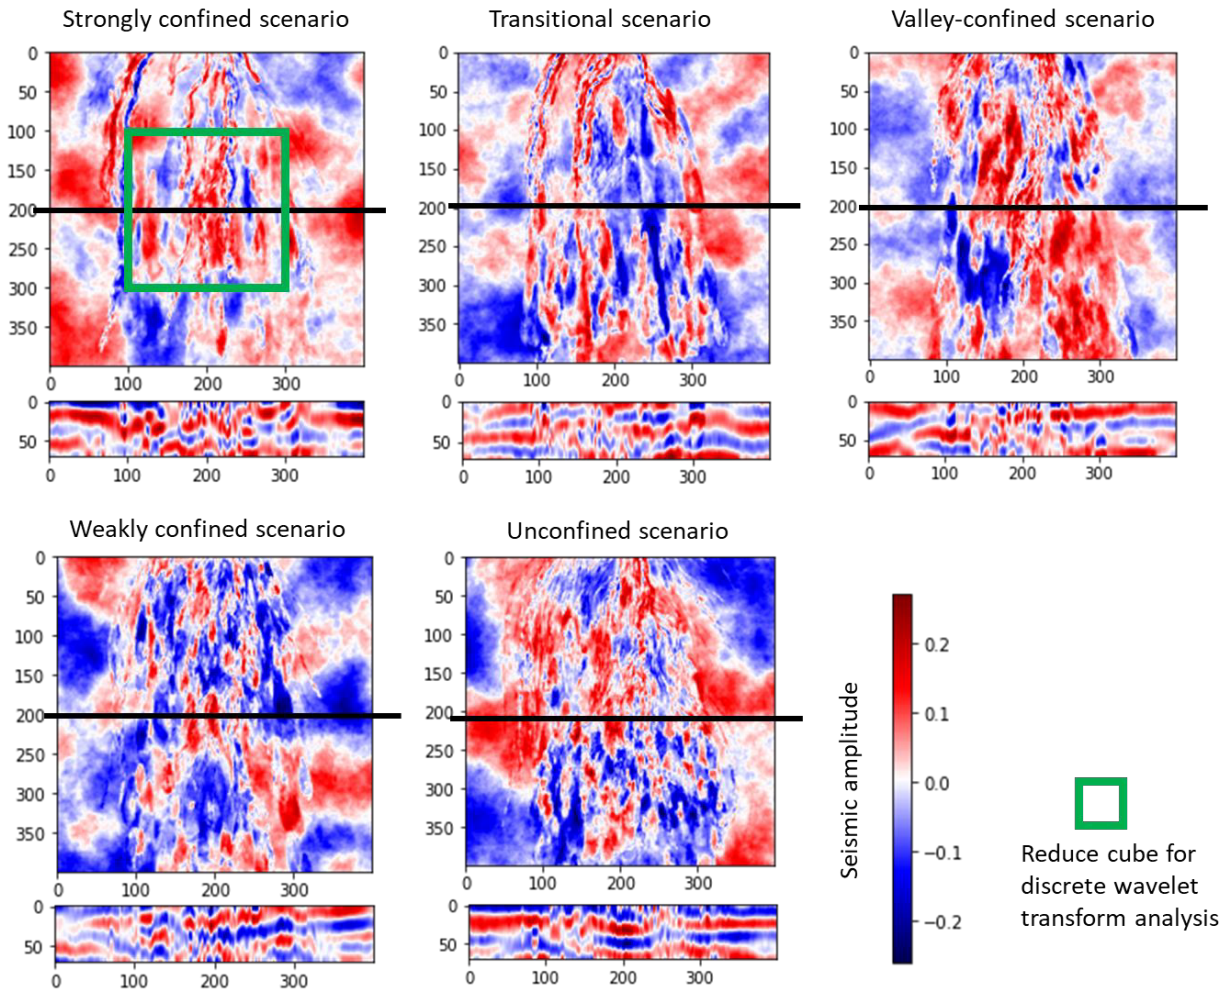

Figure S4-1. Simulated seismic data examples (a horizontal section and a vertical cross-section) for different scenarios. The horizontal section is taken at t = 13 ms. The locations of the cross-sections are marked in the horizontal sections.

Avseth, P., Mukerji, T., & Mavko, G. (2005). *Quantitative seismic interpretation: Applying rock physics tools to reduce interpretation risk*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511600074

Gardner, G. H. F., Gardner, L. W., & Gregory, A. R. (1974). FORMATION VELOCITY AND DENSITY - THE DIAGNOSTIC BASICS FOR STRATIGRAPHIC TRAPS. *Geophysics*. https://doi.org/10.1190/1.1440465

Mavko, G., Mukerji, T., & Dvorkin, J. (2020). *The Rock Physics Handbook*. *The Rock Physics Handbook*. https://doi.org/10.1017/9781108333016

Pyrcz, M. J., & Deutsch, C. V. (2014). *Geostatistical Reservoir Modeling* (2nd ed.). Oxford University Press.

Supporting Information S5

Posterior probability of different scenarios given the 18 conditioning well facies data

Following the workflow in Figure 2, we calculate the posterior probability $P\left(Sc_k \middle| f_{MPH}\left(w_{obs}^i\right)\right)$, where Sc_k $(1 \leq k \leq 5)$ refers to one of the five scenarios, w_{obs}^i $(1 \leq i \leq 18)$ refers to one of the 18 conditioning well facies data, and $f_{MPH}(\cdot)$ represents the multi-point histogram facies analysis (Tan et al., 2014). The calculated posterior values are listed in Table S5-1. To identify which wells are more informative, we computed the Kullback-Leibler (KL) divergence (Kullback & Leibler, 1951) between the prior ($P(Sc_k) =1/5$ for each scenario) and the posterior probabilities. A higher KL divergence indicates that the posterior probability deviates more from the prior probability, meaning the well is more informative. Well 8, 13, 16, and 17 have a large KL divergence (above 1), suggesting they are the most informative for scenario falsification.

Table S5-1. Posterior probabilities of different reservoir scenarios given the multi-point histogram feature of each conditioning well. The rows with KL divergence larger than 1 are highlighted.

	Strongly confined	Transitional	Valley confined	Weakly confined	Unconfined	KL div.
Well 1	0.2	0.2	0.2	0.2	0.2	0
Well 2	0.13	0.07	0.13	0.21	0.46	0.31
Well 3	0.1	0.14	0.13	0.34	0.29	0.16
Well 4	0.21	0.17	0.19	0.14	0.29	0.04
Well 5	0.19	0.28	0.12	0.31	0.1	0.15
Well 6	0.21	0.02	0.09	0.1	0.58	0.82
Well 7	0.19	0.12	0.24	0.28	0.17	0.06
Well 8	**0.24**	**0.55**	**0.1**	**0.1**	**0.01**	**1.28**
Well 9	0.37	0.22	0.17	0.09	0.15	0.15
Well 10	0.25	0.26	0.24	0.13	0.12	0.07
Well 11	0.21	0.14	0.21	0.24	0.2	0.03
Well 12	0.23	0.05	0.16	0.26	0.3	0.25
Well 13	**0.35**	**0.42**	**0.01**	**0.19**	**0.03**	**1.05**
Well 14	0.16	0.25	0.24	0.2	0.15	0.03
Well 15	0.3	0.02	0.34	0.1	0.24	0.51
Well 16	**0.49**	**0.24**	**0.1**	**0.17**	**0**	**1.25**
Well 17	**0.18**	**0.57**	**0.15**	**0.1**	**0**	**1.6**
Well 18	0.28	0.08	0.28	0.21	0.15	0.14

Kullback, S., & Leibler, R. A. (1951). On Information and Sufficiency. *The Annals of Mathematical Statistics*. https://doi.org/10.1214/aoms/1177729694

Tan, X., Tahmasebi, P., & Caers, J. (2014). Comparing training-image based algorithms using an analysis

of distance. *Mathematical Geosciences*. https://doi.org/10.1007/s11004-013-9482-1

Supporting Information S6

Tau model

The Tau model, proposed by Journel (2002), is a method for combining conditional probabilities while accounting for potential information redundancy among different data sources. For an event A, let $D_1, \dots, D_n$ represents observed data about A obtained from different measurements or sources. These data sources may contain partially overlapping information. For example, in the context of predicting reservoir distribution A, D_1 could represent borehole interpretation data, while D_2 could represent geophysical data about the subsurface, and they may have data redundancy in this context.

Given the prior probability $P(A)$ and the conditional probabilities $P(A|D_1)$, ..., $P(A|D_n)$, the goal is to estimate the joint conditional probability $P(A|D_1, \dots, D_n)$. The Tau model achieves this by defining probability ratios:

$x_0 = \frac{1-P(A)}{P(A)}$ (prior odds),

$x_i = \frac{1-P(A|D_i)}{P(A|D_i)}$ for *i = 1, …, n* (conditional odds).

The joint conditional probability is then computed as:

$P(A|D_1, \dots, D_n) = \frac{1}{1+x}$,

where x satisfies:

$\frac{x}{x_0} = \prod_{i=1}^{n} \left(\frac{x_i}{x_0}\right)^{\tau_i}$, $\tau_i \in (-\infty, +\infty)$.

Here, τ_i are Tau weights that account for the redundancy or dependency among the data sources $D_1, \dots, D_n$ in predicting event *A*. If τ_i = 1 for all *i*, the model assumes independence among the data sources. Values of τ_i greater than 1 indicate positive redundancy (overlapping information), while values less than 1 suggest negative redundancy or complementary information.

Journel, A. G. (2002). Combining knowledge from diverse sources: An alternative to traditional data independence hypotheses. *Mathematical Geology*, *34*(5), 573–596. https://doi.org/10.1023/A:1016047012594