

Diffusion posterior sampling with 3D lateral constraints for post-stack impedance inversion

Zeyang Liu^{a,b}, Dawei Liu^{c,*}, Mauricio D. Sacchi^b, Jingye Li^a, Xiaohong Chen^a, Wei Zhang^b

^a State Key Laboratory of Petroleum Resources and Prospecting, National Engineering Laboratory for Offshore Oil Exploration, China University of Petroleum—Beijing, Changping, Beijing 102249, China

^b Department of Physics, University of Alberta, Edmonton, AB T6G 2E1, Canada

^c School of Information and Communication Engineering, Xi'an Jiaotong University, Xi'an, Shaanxi 710049, China

ARTICLE INFO

Keywords:

3D seismic inversion
Diffusion model

ABSTRACT

Traditional inversion methods rely on prior information to obtain a stable and unique solution to ill-posed problems. This stability is typically enforced through regularization terms derived from statistical priors, combined with geophysical data using Bayes' rule. Regularization leads to an optimization problem where the goal is to simultaneously fit the data while minimizing a prior-induced constraint that imposes certain characteristics on the estimated model. However, these constraints can introduce solution styles that may oversimplify geological features. For instance, quadratic regularization produces overly smooth solutions, potentially omitting critical details, while edge-preserving regularization can lead to unrealistic models. Recently, diffusion models have emerged as a flexible and robust approach for approximating complex, high-dimensional probability distributions, offering more expressive priors for solving inverse problems. Their ability to learn and reverse a noisy forward process makes them particularly useful for generating realistic geological models when trained appropriately. Inspired by this, we propose a diffusion model-based method for acoustic impedance (AI) inversion. The approach is not limited to impedance inversion but can be applied to any geophysical inverse problem where enough earth-representative synthetic models are available. Our method first trains a diffusion model on synthetic AI images to approximate their distribution. We then apply diffusion posterior sampling, conditioned on seismic data, to estimate an impedance model that aligns with the observed data. Specifically, the diffusion posterior sampling integrates the gradient of the error between observed and reconstructed data, guiding the reverse diffusion process toward solutions that are consistent with observed seismic traces. We incorporate 3D lateral constraints and gradient momentum during the reverse sampling process, further improving the continuity and stability of 3D AI inversion. We demonstrate the feasibility of using diffusion models as an alternative to traditional regularization techniques for AI inversion through synthetic and real data examples.

1. Introduction

Acoustic impedance (AI) profiles, derived from seismic data, are a fundamental tool for quantitatively interpreting the subsurface. Estimating AI has long been a key challenge in applied reflection seismology (Oldenburg et al., 1983; Walker and Ulrych, 1983; Lines and Treitel, 1984; Russell and Hampson, 1991; Das et al., 2019). AI inversion is inherently ill-posed because the seismic wavelet functions as a band-pass filter, attenuating both the seismic data's low- and high-frequency components. This filtering effect makes it difficult to recover low-frequency impedance trends and high-frequency fine-scale details, posing a significant challenge for AI inversion. Noise further complicates

the problem, making AI inversion prone to noise amplification if not properly regularized. Traditional AI inversion methods use techniques such as the generalized linear method (Cooke and Schneider, 1983), LI-inversion (Oldenburg et al., 1983), autoregressive spectral extrapolation (Walker and Ulrych, 1983; Bianchin et al., 2019). Modern techniques include deterministic least-squares inversion with lateral continuity constraints (Hamid and Pidlisecky, 2015; Lin et al., 2024) and total-variation inversion methods (Gholami, 2015; Gholami, 2016; Li et al., 2024a). While effective, they depend on accurate prior information, such as a good background impedance field.

Significant progress has been made in solving inverse problems using machine learning techniques. In the case of AI inversion, several

* Corresponding author at: the School of Information and Communication Engineering, Xi'an Jiaotong University, Xi'an, Shaanxi 710049, China
E-mail address: dawei.liu@xjtu.edu.cn (D. Liu).

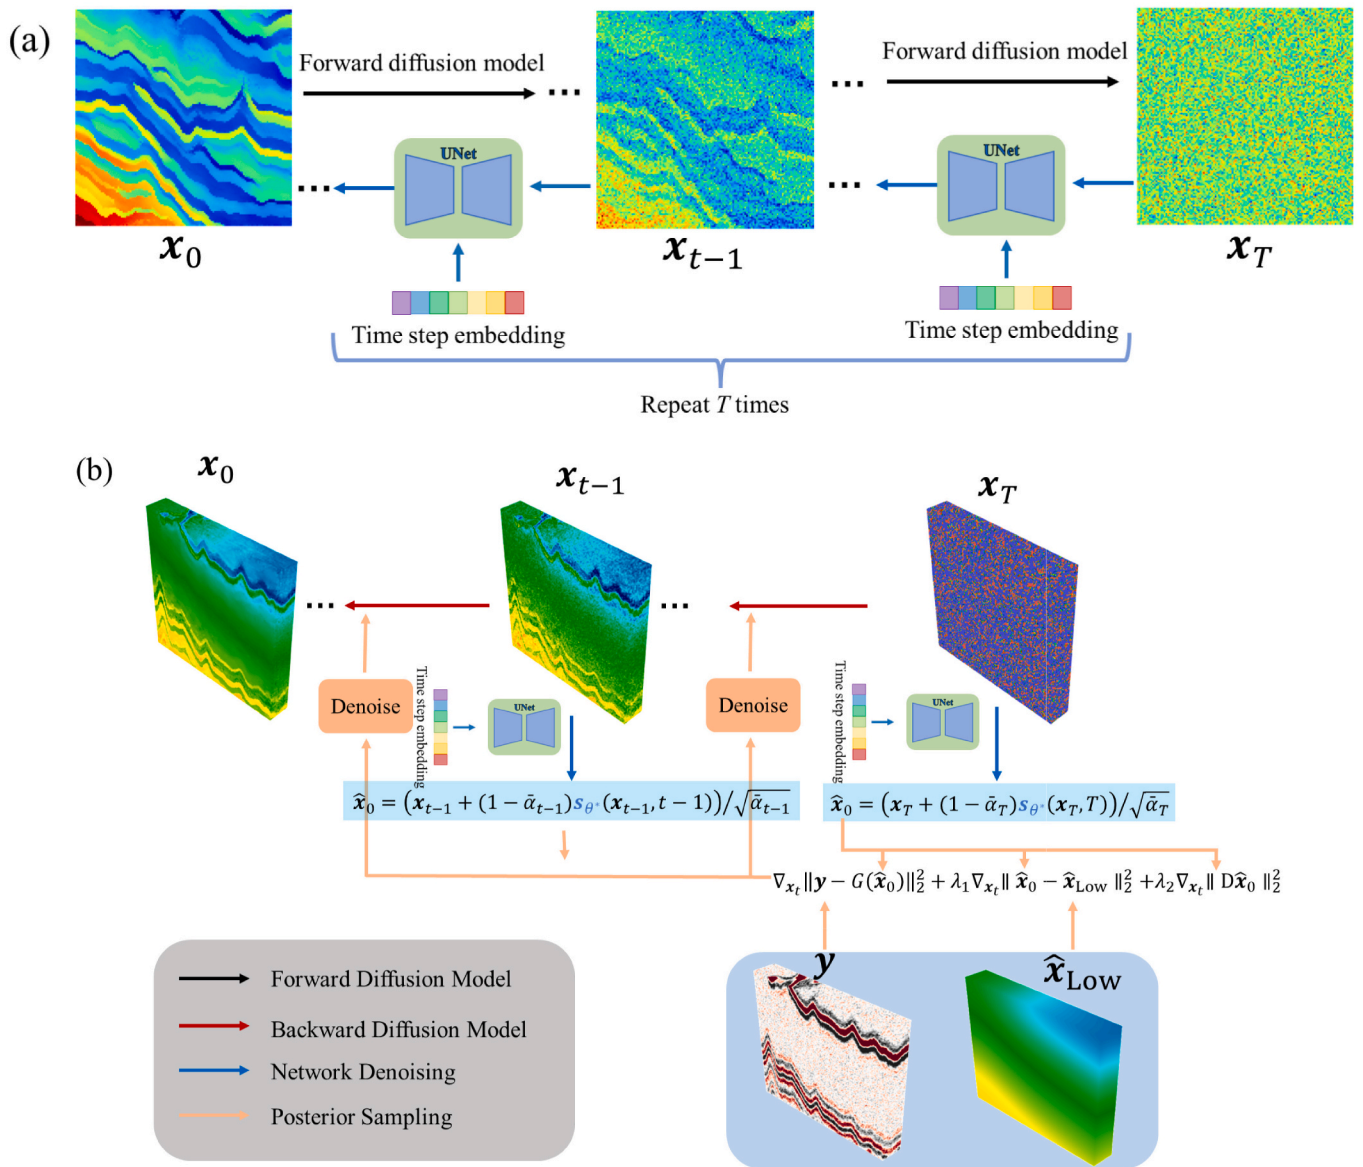


Fig. 1. Schematic diagram of the workflow structure. (a) The forward process of the diffusion model. (b) The sampling process, where y is seismic data, \hat{x}_{Low} is low-frequency model of AI, α_t represents the noise level, x_t is the process of adding noise to AI (x_0) until it becomes pure noise (x_T). And we define $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ to denote the product of all noise levels from $i = 1$ to t . $s_{\theta'}$ represents the network. The black arrows represent the forward noising process of the diffusion model. The blue arrows indicate the neural network learning to denoise. The red arrows depict the denoising steps during sampling, which correspond to the reverse process of the diffusion model. The orange arrows represent the posterior sampling steps, acting as intermediate stages within this reverse process. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

machine learning-inspired methods have been proposed, including Caianiello neural network (Fu, 2004) and data-model-driven networks (Liu et al., 2024). These methods can generally be categorized into supervised learning and unsupervised learning. While supervised learning achieves high accuracy, its effectiveness is constrained by the availability of large, well-labeled datasets (Zhang et al., 2021; Wu et al., 2021). In contrast, unsupervised learning leverages physical constraints to guide the network, but it still faces challenges related to model stability (Wang et al., 2022b; Li et al., 2024b).

Considering these challenges, generative models have gained significant attention for their ability to learn complex data distributions and generate realistic samples. Among these generative models, diffusion models (Sohl-Dickstein et al., 2015) have emerged as powerful tools in machine learning, with applications in denoising (Ho et al., 2020) and

image generation (Ho et al., 2022). The central idea of diffusion models is a generative process in which noise is progressively added to an image from an unknown distribution until it becomes pure noise. The model then learns to reverse this process, reconstructing an image from noise (Dhariwal and Nichol, 2021). This framework can be leveraged for a probabilistic approach to seismic AI inversion, where high-resolution AI models serve as the images injected into diffusion models. The reverse process then generates AI profiles that act as priors for inversion.

Among various sampling strategies in diffusion models, Denoising Diffusion Probabilistic Models (DDPMs) have demonstrated exceptional effectiveness in capturing complex distributions. Their probabilistic nature enables the generation of highly realistic models (Luo, 2022; Chan, 2024). However, since DDPMs operate as unconditional generative processes, they cannot directly solve inverse problems where




Fig. 2. Examples of impedance samples for training the network.

observations must be honored. To overcome this limitation, they can be modified to incorporate conditioning. Readers seeking a comprehensive discussion on the application of diffusion models to inverse problems, such as inpainting, deconvolution, and denoising, are encouraged to refer to the review by Daras et al. (2024) and references therein. We highlight several key methodologies to provide a succinct introduction to the extensive research on diffusion models for inverse problems. The generative process in diffusion models has been adapted to produce inversion results that align with observed data through various approaches. These include iterative latent variable refinement (Choi et al., 2021), Diffusion Posterior Sampling (DPS) (Chung et al., 2022a), and the denoising diffusion null-space model (Wang et al., 2022a). Among these, DPS has gained recognition for its capability to address nonlinear inverse problems and effectively manage noisy observations (Daras et al., 2024).

DPS improves the accuracy and reliability of inversion results by iteratively updating the gradient of the cost function of the inverse problem based on available observed seismic data. Traditional inversion methods often depend on predefined priors, such as low-frequency models, which may be inadequate for capturing complex geological structures. In contrast, DPS leverages prior information learned during the generative process and dynamically integrates it with physical

constraints in the inversion workflow. This approach leads to more robust and reliable outcomes in seismic imaging and related applications.

We propose an AI inversion method based on DPS by utilizing the learned generative prior information. We first train the network on a synthetic AI dataset through noise addition and denoising, thereby obtaining a generative prior of plausible AI models. Then, we apply the DPS method to perform the AI inversion and obtain the estimated AI that honors seismic traces. To adapt to the AI inversion problem and the complex AI distribution, we add additional constraints, such as low-frequency model constraints and a 3D lateral constraint. These constraints help the model produce more geologically realistic solutions and improve lateral continuity in the inversion results. Additionally, we incorporate momentum estimation (Kingma and Ba, 2014) into the optimization stage of the diffusion model to ensure stability during gradient updates. The resulting method is a gradient-enhanced DPS with a 3D lateral constraint (GDPS). We perform tests on a 3D synthetic dataset and a 3D real dataset. The results are compared with a simple Bayesian Maximum A Posteriori method (MAP) (Tarantola, 2005), which in our case reduces to damped least-squares with a provided background AI field. The synthetic and real data results demonstrate that the GDPS method is more stable than the originally proposed DPS

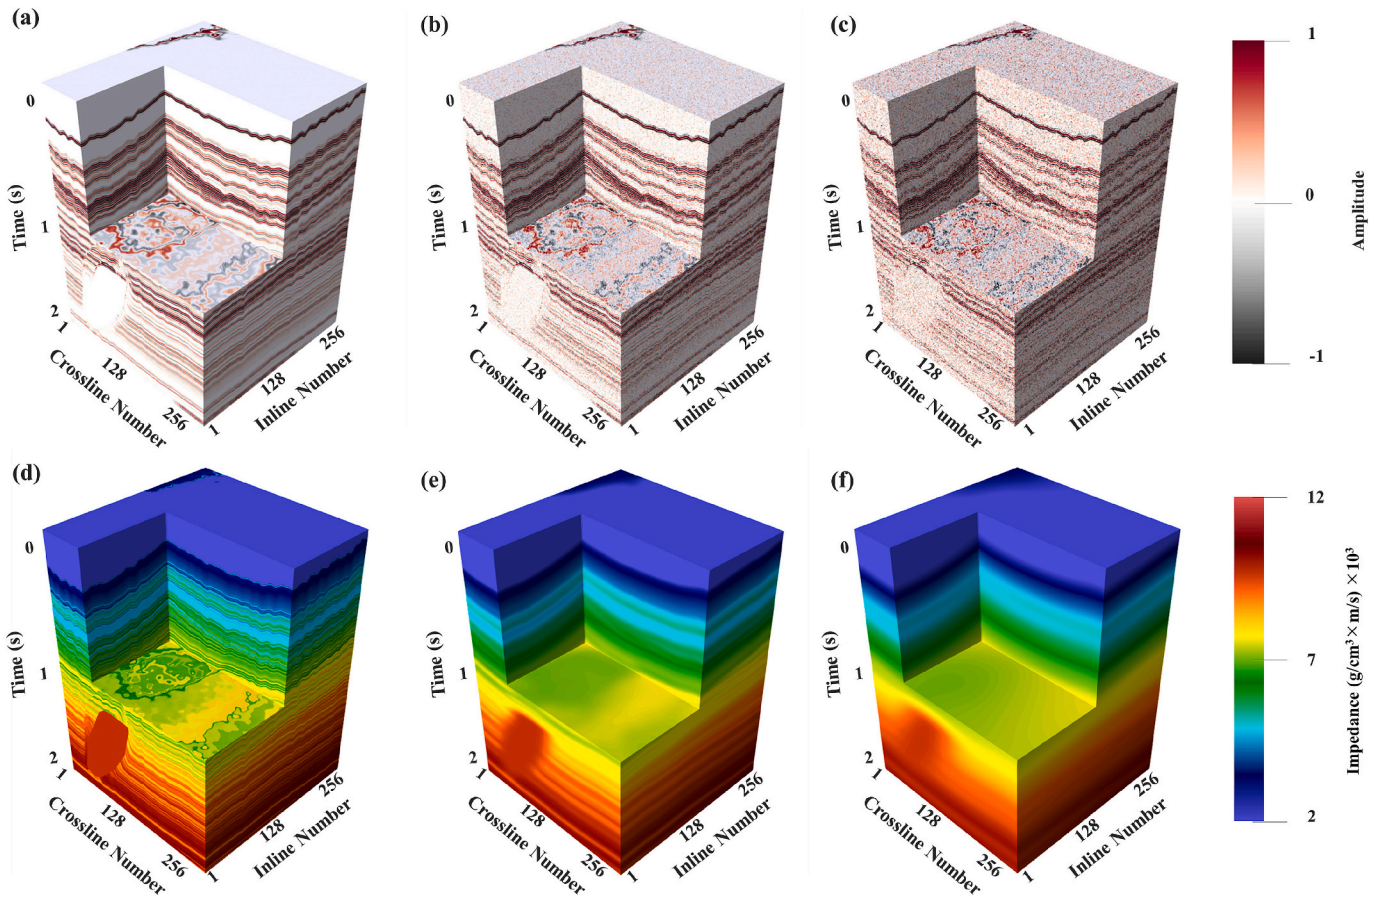


Fig. 3. Synthetic seismic records. (a) Clean poststack seismic data. (b) Seismic data with $S/N_{in} = 3$ dB. (c) Seismic data with $S/N_{in} = 0.5$ dB. (d) AI model parameters. (e) Low-frequency model with blur radius = 10. (f) Low-frequency model with blur radius = 30.

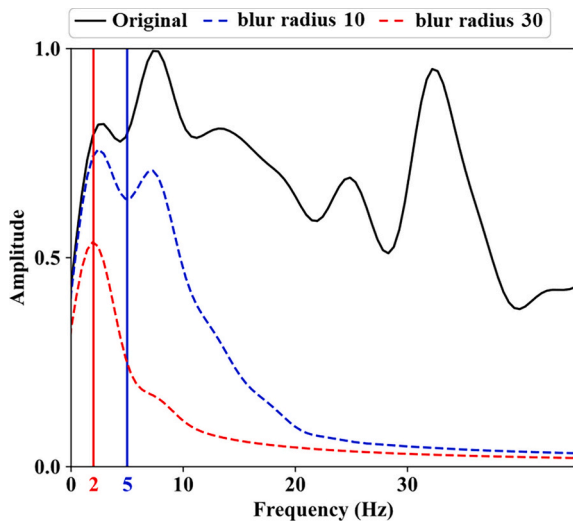


Fig. 4. Low-frequency model spectrum. The red dashed line represents the spectrum of the low-frequency model with a blur radius of 30, while the blue dashed line represents the spectrum of the low-frequency model with a blur radius of 10. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

(Chung et al., 2022a). Additionally, the accuracy of AI inversion results for the real data improves by 5% compared to the MAP solution. The contributions of our method are as follows.

First, we introduce 3D lateral constraints that go beyond standard 2D

diffusion methods, making it possible to perform full 3D inversion. This is a major improvement over earlier DPS methods.

Second, we suggest using momentum-based gradient updates to make the reverse diffusion process more stable. Adding momentum helps avoid unstable changes during sampling.

Finally, we include low-frequency models in the diffusion process, combining physics-based rules with data-driven detail generation. This mixed approach keeps both the large-scale impedance patterns and the small geological features.

2. Method

We first introduce the forward model employed in this paper. Then, we explain how the diffusion model is used for AI inversion. Finally, we describe how low-frequency constraints, 3D lateral constraints, and gradient optimization are incorporated into the sampling process.

2.1. Forward model

Seismic poststack records result from the convolution of a series of reflection coefficients with the wavelet. Typically, the series of reflection coefficients is modeled as a white (uncorrelated) sequence of spikes. The wavelet can be estimated deterministically using sonic logs or statistically by assuming its phase. The model assumes that the seismic data has undergone appropriate processing, allowing the wavelet to be considered stationary and known (Yilmaz, 2001). Under these assumptions, the forward model is generally expressed as follows

$$y = G(x_0) + n \tag{1}$$

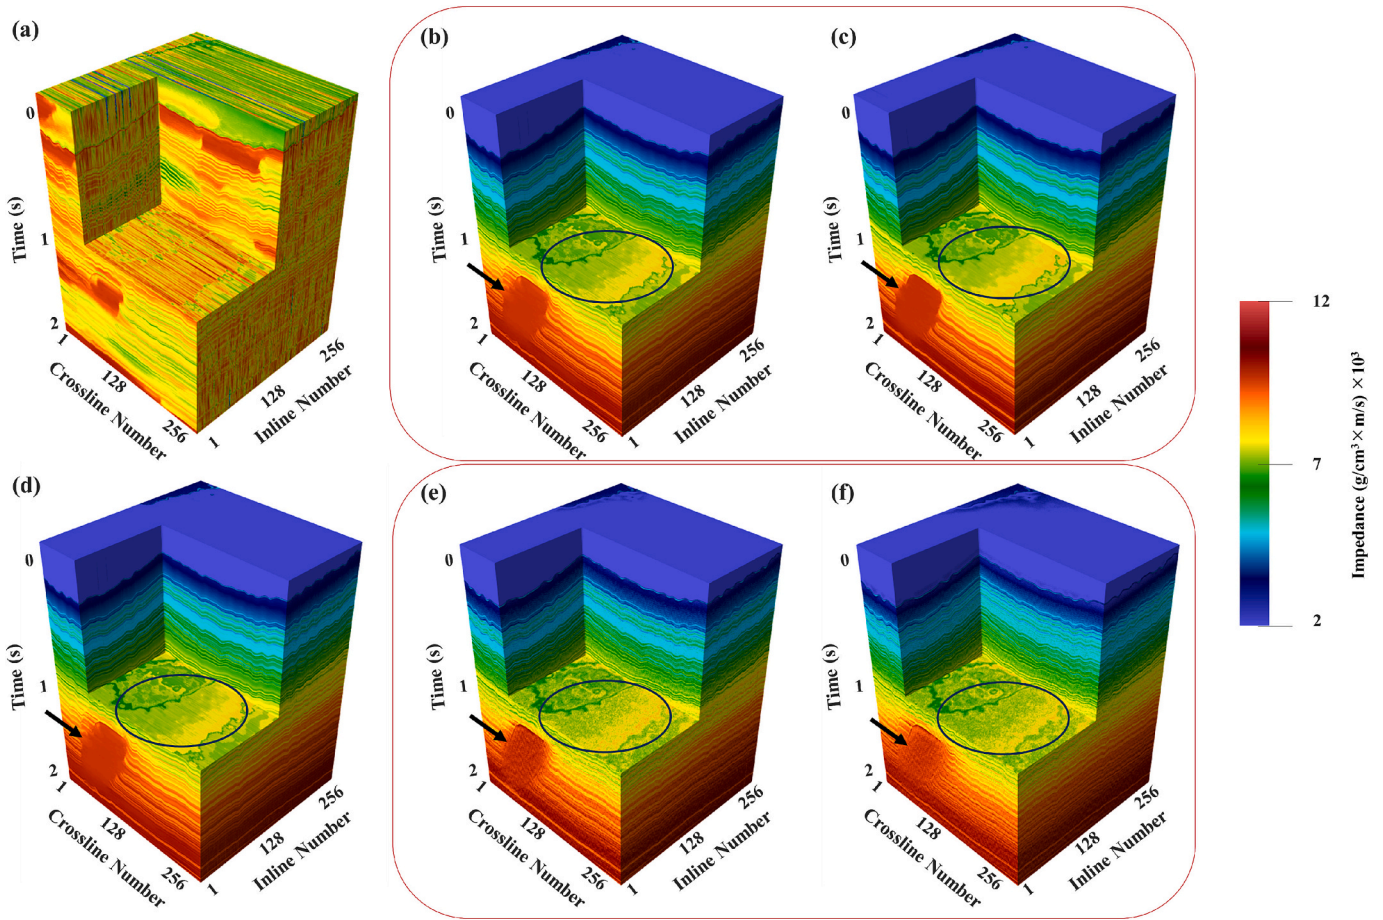


Fig. 5. Inversion results. (a) DPS without low-frequency constraints for $S/N_{in} = 3$ dB. (b) GDPS for $S/N_{in} = 3$ dB, blur radius = 10. (c) GDPS for $S/N_{in} = 3$ dB, blur radius = 30. (d) GDPS for $S/N_{in} = 0.5$ dB, blur radius = 10. (e) MAP for $S/N_{in} = 3$ dB, blur radius = 10. (f) MAP for $S/N_{in} = 3$ dB, blur radius = 30. The black arrows indicate salt dome, and the black circles indicate sandstone reservoir boundaries. The red box is used to make it easier to compare the GDPS and MAP methods. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

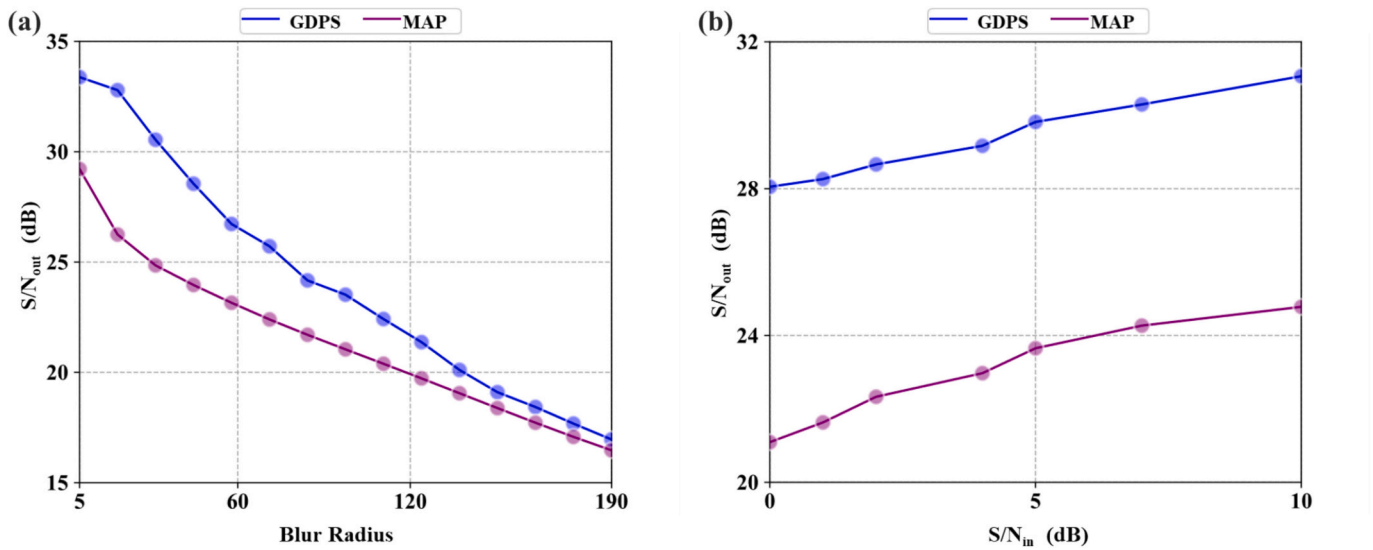


Fig. 6. Compare the impact of the GDPS and MAP methods on low-frequency models and S/N_{in} . (a) Different blur radii. (b) Different S/N_{in} . A larger blur radius indicates a poorer low-frequency model, while a larger S/N_{out} indicates a better prediction result.

where \mathbf{G} , the forward operator, represents a conversion from impedance (\mathbf{x}_0) to reflectivity and convolution with a known wavelet. The observations \mathbf{y} correspond to an ensemble of seismic traces typically stored as

a 3D volume $y(i, j, k)$ where i , a discrete time index, corresponds to the two-way travel time of the P-wave and j, k corresponds to inline and crossline Common-Mid-Point (CMP) coordinates. Similarly, \mathbf{n} represents

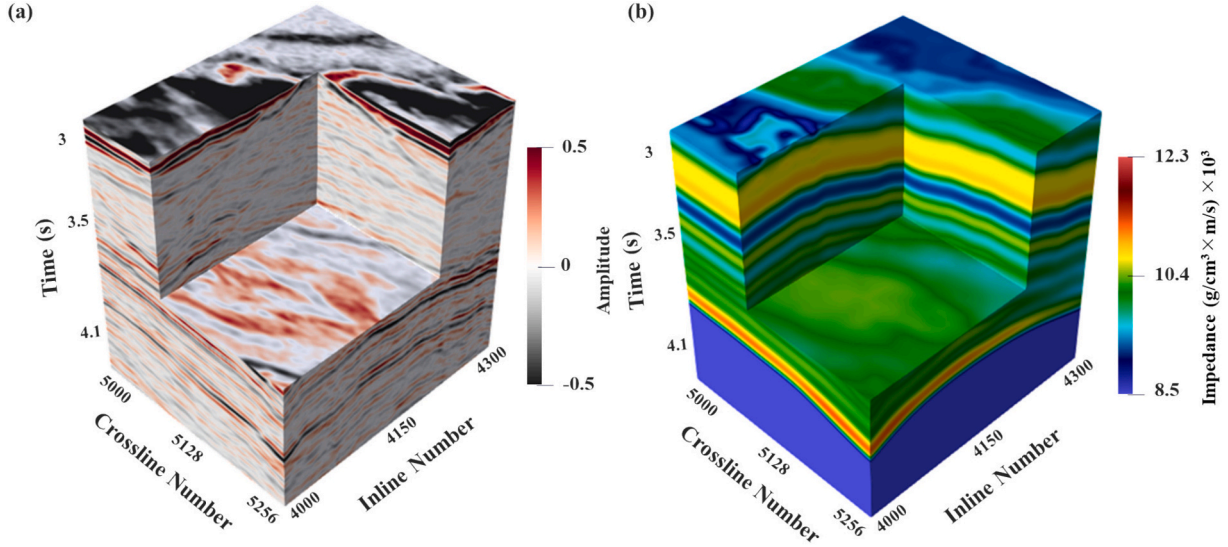


Fig. 7. 3D real data for impedance inversion. (a) Poststack seismic records. (b) Low-frequency model.

additive noise that is considered spatially incoherent and of known variance. The expression of the reflection coefficient (r) is as follows

$$r(i, j, k) = \frac{x_0(i+1, j, k) - x_0(i, j, k)}{x_0(i+1, j, k) + x_0(i, j, k)} \quad (2)$$

2.2. Diffusion model with DPS for AI inversion

Diffusion models define the generative process as the reverse of the noising process. It will provide a good way to capture the prior information from the training datasets. The well-known DDPM is formulated in the discrete form, which can be proved to the equivalence of scored-based generative modeling in continuous form (Ho et al., 2020; Song et al., 2020). In Appendix A, we provide a review about the equivalence of DDPM and scored-based generative. In this subsection, we will introduce the stochastic differential equation (SDE) for DDPM and describe how to solve our non-linear inverse problem using the process of diffusion posterior sampling.

Song and Ermon (2019) define SDE that gradually transforms the input natural image into pure noise, that is forward (noising) process of diffusion model. Our aim is to incorporate the noising and denoising process of the diffusion models for inverting the AI of the subsurface. The AI noising process is described by the following continuous form

$$d\mathbf{x} = -\frac{\beta(t)}{2}\mathbf{x}dt + \sqrt{\beta(t)}d\mathbf{w} \quad (3)$$

where $\beta(t)$ the noise schedule of the process (Ho et al., 2020) and \mathbf{w} denotes a Wiener process (a continuous-time stochastic process modeling random noise with independent, normally distributed increments). When $t = 0$, the AI distribution is defined as $\mathbf{x}_0 \sim P_{AI}$, and when $t = T$, it reaches the Gaussian distribution, i.e. $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$. The forward process is illustrated in Fig. 1(a). In addition, we define $\alpha_i = 1 - \beta_i = \beta(t + \Delta t)\Delta t$, where $\Delta t = 1/T$. And we define $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ to denote

the product of all noise levels from $i = 1$ to t . The relationship between \mathbf{x}_{t-1} and \mathbf{x}_t (with noise added) is presented in Eq. (A-1) in Appendix A. The goal of the reverse process is to recover the AI distribution from

the Gaussian distribution, that is, starting from the sample \mathbf{x}_T to obtain \mathbf{x}_0 . This step can be implemented by the reverse-time SDE (Friedman, 1975)

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)\nabla_{\mathbf{x}}\log p_t(\mathbf{x}) \right] dt + \sqrt{\beta(t)}d\bar{\mathbf{w}} \quad (4)$$

where t corresponds to the reverse time and $\bar{\mathbf{w}}$ is the time reversed process of \mathbf{w} . In other words, $d\bar{\mathbf{w}} = \mathbf{w}(t - \Delta t) - \mathbf{w}(t)$. The symbol $p_t(\mathbf{x})$ is the probability distribution of \mathbf{x} at time t . The relationship between \mathbf{x}_t and \mathbf{x}_{t-1} (with denoising) is presented in Eq. (A-8) in Appendix A. We can approximate t -dependent score function $\nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t)$ using a denoising neural network \mathbf{s}_θ (Vincent, 2011) by minimizing the following loss function

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \|\mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t)\|_2^2 \quad (5)$$

Once θ^* is obtained from Eq. (5), we can use the trained score function $\mathbf{s}_{\theta^*}(\mathbf{x}_t, t)$ as an estimate of score function to substitute $\nabla_{\mathbf{x}}\log p_t(\mathbf{x})$ into Eq. (4). Solving Eq. (4) corresponds to sampling from the AI distribution, which is the goal of generative models (Chung et al., 2022a).

For the inverse problem, the goal is to predict the AI model parameter \mathbf{x}_0 from observed data \mathbf{y} . However, the original reverse process in Eq. (4) only allows for unconditional updates (Song et al., 2020; Chung et al., 2022b), which may generate AI models that do not match the observed data. To guarantee accurate inversion results, we need to incorporate the diffusion model to solve the more general AI inversion problem. To leverage the diffusion model as a prior for solving the inverse problem, we can directly modify Eq. (4) to obtain the reverse diffusion sampler for sampling from the posterior distribution

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)\nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t|\mathbf{y}) \right] dt + \sqrt{\beta(t)}d\bar{\mathbf{w}} \quad (6)$$

where $\nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t|\mathbf{y})$ can be computed from the Bayes' rule

$$\nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}}\log p(\mathbf{y}|\mathbf{x}_t) \quad (7)$$

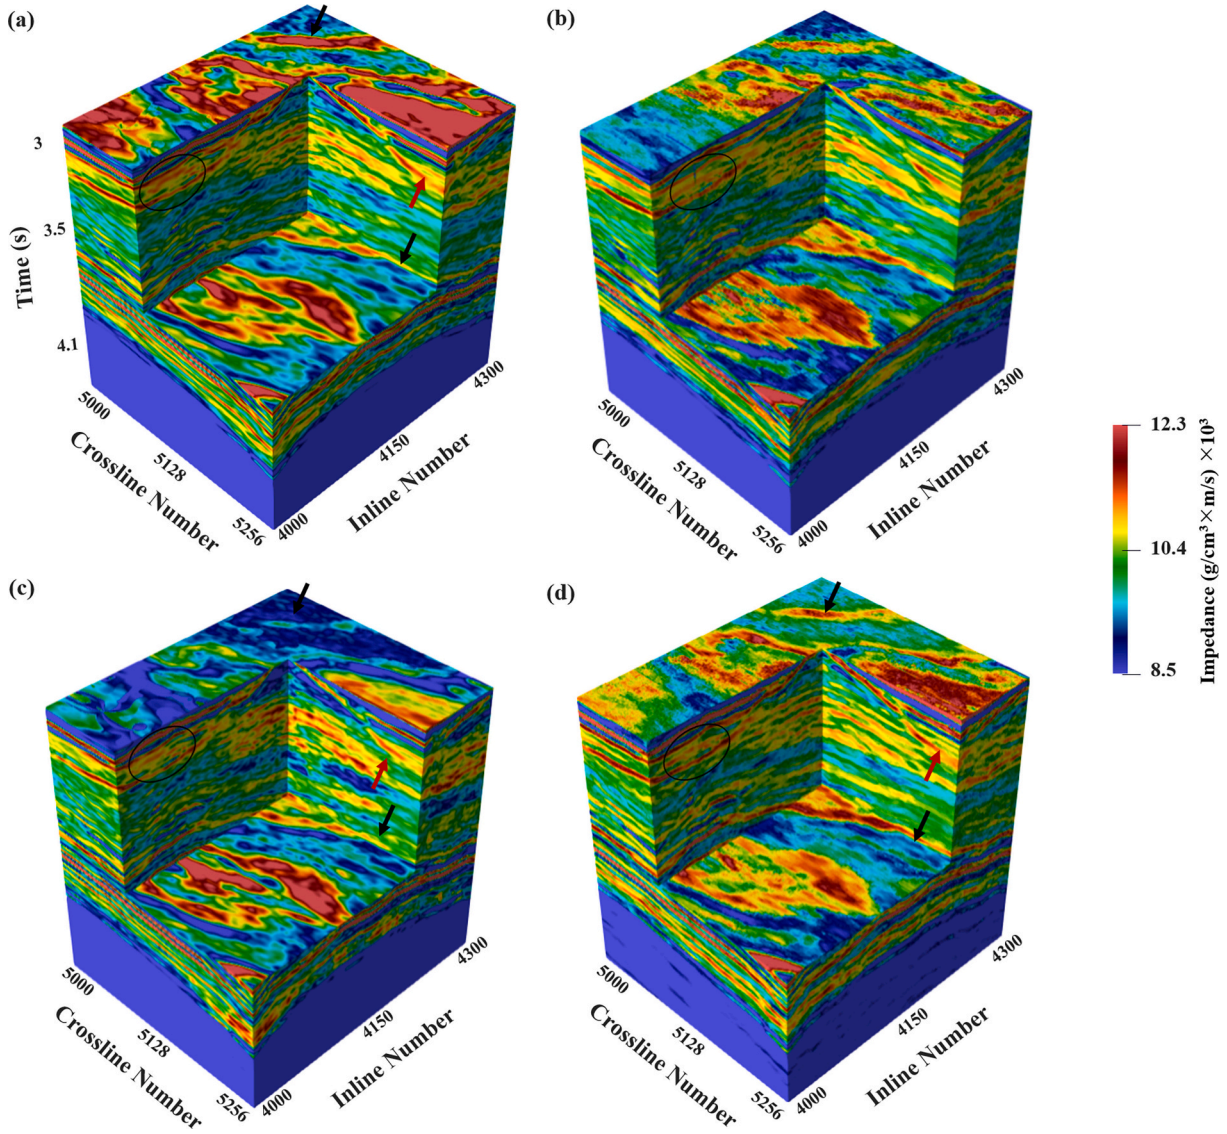


Fig. 8. Inversion results of different methods. (a) MAP. (b) GDPS without 3D lateral constraint. (c) Unsupervised learning, (d) GDPS with 3D lateral constraint. The black circles emphasize that GDPS provides higher vertical resolution and lateral continuity. The black and red arrows highlight that GDPS provides higher resolution. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The first term of the score function in Eq. (7) can be computed from the trained score function. It is hard to acquire the closed-form equation for the second term $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$, due to its dependence on t . The reason is that there is only explicit dependence between \mathbf{y} and \mathbf{x}_0 . We use the approximation $p(\mathbf{y}|\mathbf{x}_t) \approx p(\mathbf{y}|\hat{\mathbf{x}}_0(\mathbf{x}_t))$, where $\hat{\mathbf{x}}_0$ is the posterior mean of $p(\mathbf{x}_0|\mathbf{x}_t)$ (Chung et al., 2022b). The relationship between \mathbf{x}_t and \mathbf{x}_0 is presented in Eq. (A-4) in Appendix A. The likelihood function can be written as

$$p(\mathbf{y}|\hat{\mathbf{x}}_0(\mathbf{x}_t)) = \frac{1}{\sqrt{(2\pi)^N \sigma^{2N}}} \exp \left[-\frac{\|\mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|_2^2}{2\sigma^2} \right] \quad (8)$$

where N represents the dimension of \mathbf{y} . Taking the gradient of likelihood function with respect to \mathbf{x}_t , will produce

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t) \approx -\frac{1}{2\sigma^2} \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|_2^2 \quad (9)$$

Finally, by iteratively applying the reverse diffusion sampler (Eq. (10)), we generate samples from the posterior distribution

$$d\mathbf{x} \approx \left[-\frac{\beta(t)}{2} \mathbf{x} - \beta(t) \left(\mathbf{s}_{\theta'}(\mathbf{x}_t, t) - \frac{1}{2\sigma^2} \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|_2^2 \right) \right] dt + \sqrt{\beta(t)} d\bar{\mathbf{w}} \quad (10)$$

For convenience, we denote $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ as $\hat{\mathbf{x}}_0$ in the following sections. See Appendix A for details about how to estimate $\hat{\mathbf{x}}_0$ from \mathbf{x}_t .

2.3. Improved DPS for AI inversion

In Eq. (9), the right side behaves as the gradient of the waveform-matching objective function of the traditional inversion method. Just like we do in the traditional inversion, we can incorporate more prior information for the conditional sampling process of the diffusion model, e.g., the low-frequency model and lateral constraint. The lateral constraint will improve the inversion accuracy in terms of spatial continuity. The gradient of the posterior probability in the proposed GDPS method can be expressed as




Fig. 9. Comparison of single-trace inversion results. The black solid line is the well-log data, the purple solid line is the inversion result of MAP, the cyan solid line is the inversion result of the GDPS, the red solid line is the inversion result of the unsupervised learning, and the yellow dotted line is the low-frequency model. The red arrow indicates that the GDPS prediction results are more reliable. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t | \mathbf{y}) \approx \mathbf{s}_{\theta'}(\mathbf{x}_t, t) - l_r (\nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0)\|_2^2 + \lambda_1 \nabla_{\mathbf{x}_t} \|\hat{\mathbf{x}}_0 - \hat{\mathbf{x}}_{\text{Low}}\|_2^2 + \lambda_2 \nabla_{\mathbf{x}_t} \|\mathbf{D}\hat{\mathbf{x}}_0\|_2^2) \quad (11)$$

the sampling process is performed using a full 2D profile to generate 3D results. Consequently, it is necessary to impose a smoothness constraint along the crossline direction. In other words, we can consider the three dimensions of $\hat{\mathbf{x}}_0$ as crossline, time, and inline, respectively, so the difference operator can be expressed as $\hat{\mathbf{x}}_0(j, i, k) - \hat{\mathbf{x}}_0(j, i + 1, k)$.

To ensure stability during gradient updates, we introduce the first-moment estimate m_t and second-moment estimate v_t of the gradient $\mathbf{g}_t = \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0)\|_2^2 + \lambda_1 \nabla_{\mathbf{x}_t} \|\hat{\mathbf{x}}_0 - \hat{\mathbf{x}}_{\text{Low}}\|_2^2 + \lambda_2 \nabla_{\mathbf{x}_t} \|\mathbf{D}\hat{\mathbf{x}}_0\|_2^2$ defined as

$$m_t = \beta_1 m_{t+1} + (1 - \beta_1) \mathbf{g}_t \quad (12a)$$

$$v_t = \beta_2 v_{t+1} + (1 - \beta_2) \mathbf{g}_t^2 \quad (12b)$$

where β_1 and β_2 are parameters controlling the decay rates of the moment estimates (Kingma and Ba, 2014). To correct the bias in the moment estimates, we use

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^{T-t+1}} \quad (13a)$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^{T-t+1}} \quad (13b)$$

Finally, the gradient update is given by

$$\hat{\mathbf{g}}_t = \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}} \quad (14)$$

where ϵ helps to avoid numerical instability by ensuring the denominator does not become too small or zero. Finally, the gradient of the posterior probability in the GDPS method can be updated by

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) \approx \mathbf{s}_{\theta'}(\mathbf{x}_t, t) - l_r \hat{\mathbf{g}}_t \quad (15)$$

where l_r is the learning rate.

We compare the sampling workflows of the DPS and GDPS methods as shown in Algorithm 1. The input data of GDPS is the observed seismic records, and the initial value is \mathbf{x}_T . We can then use Eqs. (10) and (15) to apply the DPS and GDPS methods. In other words, we sample \mathbf{x}_T using the learned prior information to obtain \mathbf{x}_0 . The specific process is shown in Fig. 1(b).

Algorithm 1. Comparison of the sampling workflows of the DPS and GDPS methods.

Algorithm 1 DPS	Algorithm 1 GDPS
Require: $T, \mathbf{y}, \Sigma_t, l_r$	Require: $T, \mathbf{y}, \Sigma_t, l_r, \lambda_1, \lambda_2, \hat{\mathbf{x}}_{\text{Low}}$
1 : $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$	1 : $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I}), m_{T+1} = \mathbf{0}, v_{T+1} = \mathbf{0}$
2 : for $t = T$ to 1 do	2 : for $t = T$ to 1 do
3 : $\hat{\mathbf{s}} \leftarrow \mathbf{s}_{\theta'}(\mathbf{x}_t, t)$	3 : $\hat{\mathbf{s}} \leftarrow \mathbf{s}_{\theta'}(\mathbf{x}_t, t)$
4 : $\hat{\mathbf{x}}_0 \leftarrow (\mathbf{x}_t + (1 - \alpha_t) \hat{\mathbf{s}}) / \sqrt{\alpha_t}$	4 : $\hat{\mathbf{x}}_0 \leftarrow (\mathbf{x}_t + (1 - \alpha_t) \hat{\mathbf{s}}) / \sqrt{\alpha_t}$
5 : $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$	5 : $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$
6 : $\mathbf{x}'_{t-1} \leftarrow (\mathbf{x}_t + (1 - \alpha_t) \hat{\mathbf{s}}) / \sqrt{\alpha_t} + \Sigma_t \mathbf{z}$	6 : $\mathbf{x}'_{t-1} \leftarrow (\mathbf{x}_t + (1 - \alpha_t) \hat{\mathbf{s}}) / \sqrt{\alpha_t} + \Sigma_t \mathbf{z}$
7 : $\mathbf{x}_{t-1} \leftarrow \mathbf{x}'_{t-1} - l_r \nabla_{\mathbf{x}_t} \ \mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0)\ _2^2$	7 : $\mathbf{g}_t \leftarrow \nabla_{\mathbf{x}_t} (\ \mathbf{y} - \mathbf{G}(\hat{\mathbf{x}}_0)\ _2^2 + \lambda_1 \ \hat{\mathbf{x}}_0 - \hat{\mathbf{x}}_{\text{Low}}\ _2^2 + \lambda_2 \ \mathbf{D}\hat{\mathbf{x}}_0\ _2^2)$
8 : return $\hat{\mathbf{x}}_0$	8 : $\hat{m}_t = (\beta_1 m_{t+1} + (1 - \beta_1) \mathbf{g}_t) / (1 - \beta_1^{T-t+1})$
	$\hat{v}_t = \sqrt{(\beta_2 v_{t+1} + (1 - \beta_2) \mathbf{g}_t^2) / (1 - \beta_2^{T-t+1})}$
	9 : $\mathbf{x}_{t-1} \leftarrow \mathbf{x}'_{t-1} - l_r \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon)$
	10 : return $\hat{\mathbf{x}}_0$

where \mathbf{D} is the first-order difference operator of the first derivative along the crossline direction (batch size dimension) between adjacent 2D profiles and $\hat{\mathbf{x}}_{\text{Low}}$ is the low-frequency model of the AI. Due to the use of the 2D patching scheme during the training phase, as shown in Fig. 1(a),

2.4. Network framework

We employed a simplified U-Net architecture to train the reverse process, which learns to denoise the data during the generative phase. The network primarily consists of a time embedding module,

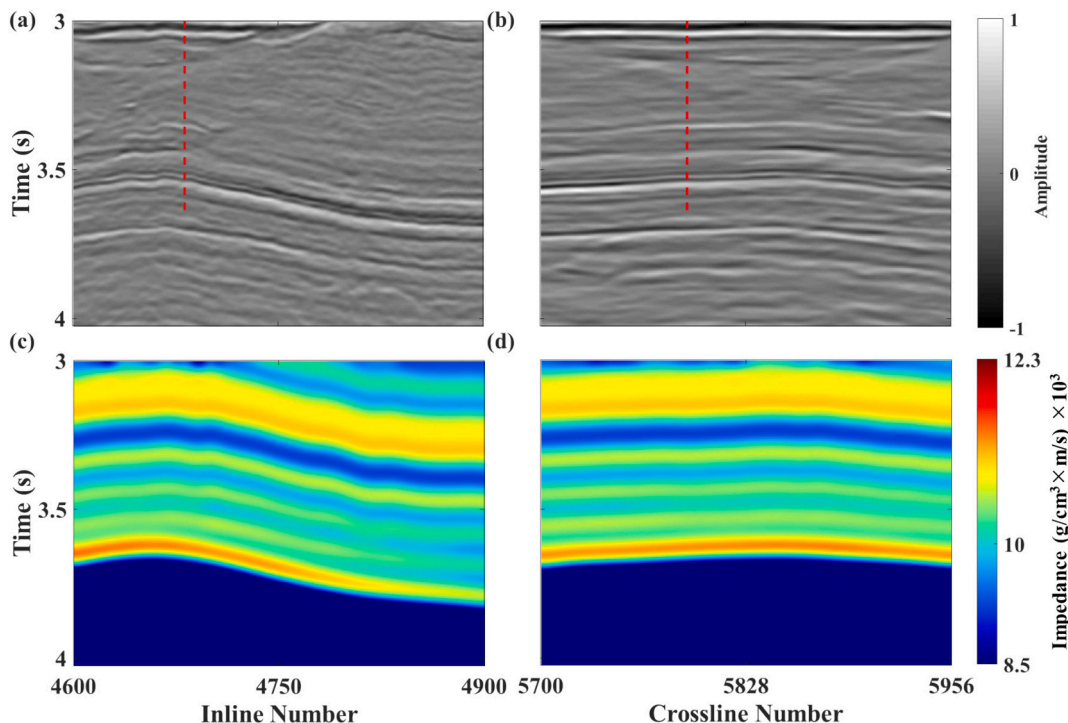


Fig. 10. Real data seismic records in (a) crossline 4083 and (b) inline 5093. Low-frequency model in (c) crossline 4083 and (d) inline 5093. The red dashed line represents the location of well 1. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

convolutional layers, and sampling layers. The time embedding module maps the time input to a combination of sine and cosine functions, ensuring periodicity and continuity in the embedding. The time information is processed through a linear layer and ReLU activation function, ensuring the deeper layers of the network can effectively utilize it. The encoding and decoding stages follow the same structure as those in the traditional U-Net architecture. The network progressively reduces the image's spatial resolution while increasing the number of feature channels through down-sampling, enabling the model to learn higher-level features. Afterward, the network gradually restores the image resolution during the up-sampling stages. The input image is first passed through an initial convolutional layer, which maps it to a feature map with 64 channels. The image then undergoes multiple down-sampling modules containing two convolutional layers, batch normalization, and a ReLU activation. The time embedding vector is added to the feature map at each convolutional layer, allowing the network to leverage time information during the feature extraction process. The down-sampling modules progressively reduce the spatial resolution while increasing the number of channels in steps: first 64 channels, then 128, and finally 256 channels. In the up-sampling stages, the network gradually restores the spatial resolution of the image through deconvolution operations and merges lower-level details with higher-level semantic information by concatenating the feature maps from the down-sampling stages with the current feature maps. Finally, after up-sampling, the feature maps are passed through a 1×1 convolutional layer to generate the output image. This model architecture is particularly suitable for the diffusion model framework, where the time steps control the denoising process, and the U-Net structure helps the network capture detailed information during the progressive generation process, thus enabling the generation of high-quality images. Therefore, this network effectively combines temporal and spatial features, generating clear and coherent images, especially in diffusion model-based image generation tasks.

The lightweight network structure used in this study is designed to be compatible with most GPU processors. While it is possible to increase the number of convolutional channels to 1024 or incorporate self-attention

mechanisms to improve performance, these changes would significantly increase computational costs. This study mainly focuses on the validation of the DPS method, so the network is not explored in depth and is tested using only a simple, general architecture.

3. Example

We generated two sets of synthetic AI models using the simulation method from Merrifield et al. (2022). One dataset of size $300 \times 300 \times 800$ (Inline \times Crossline \times Time) is used to construct the training dataset in the forward process, in which noise is progressively added to clean AI images, and the other dataset with size $256 \times 256 \times 1024$ (Inline \times Crossline \times Time) is employed for our inversion testing during the reverse sampling. We randomly cropped 5000 AI profile patches of size 128×128 (CDP \times Time) from the first dataset, as exemplified in Fig. 2. The Adam optimizer is adopted for training with a learning rate of $1e-3$, a batch size of 128, and a total of 60,000 iterations. The parameter T is set to 1000. The trained network is applied for posterior sampling on the second synthetic dataset and real data. For comparison, we employed the MAP inversion approach incorporating a background trend and smoothness constraints. While more sophisticated MAP formulations exist such as those leveraging model covariance matrices we deliberately assumed no access to such covariances to ensure a simple comparison. Consequently, our MAP solution for a Gaussian model simplifies to a damped least-squares solution with a known background trend model. This simplification does not overly restrict the analysis, as our primary objective was to explore the applicability of diffusion models in generating suitable priors for practical impedance inversions. Details pertaining to the MAP estimator can be found in Appendix B.

We define the input signal-to-noise ratio as $S/N_{in} = 20\log_{10}\|\mathbf{y} - \mathbf{n}\|_2/\|\mathbf{n}\|_2$, where $\mathbf{y} - \mathbf{n}$ represents the clean synthetic trace data and \mathbf{n} is the noise term. Similarly, one defines the equation for the output signal-to-noise ratio as $S/N_{out} = 20\log_{10}\|\mathbf{x}_{true}\|_2/\|\mathbf{x}_0 - \mathbf{x}_{true}\|_2$, where \mathbf{x}_0 is the estimated AI, and \mathbf{x}_{true} is the ground-truth AI.

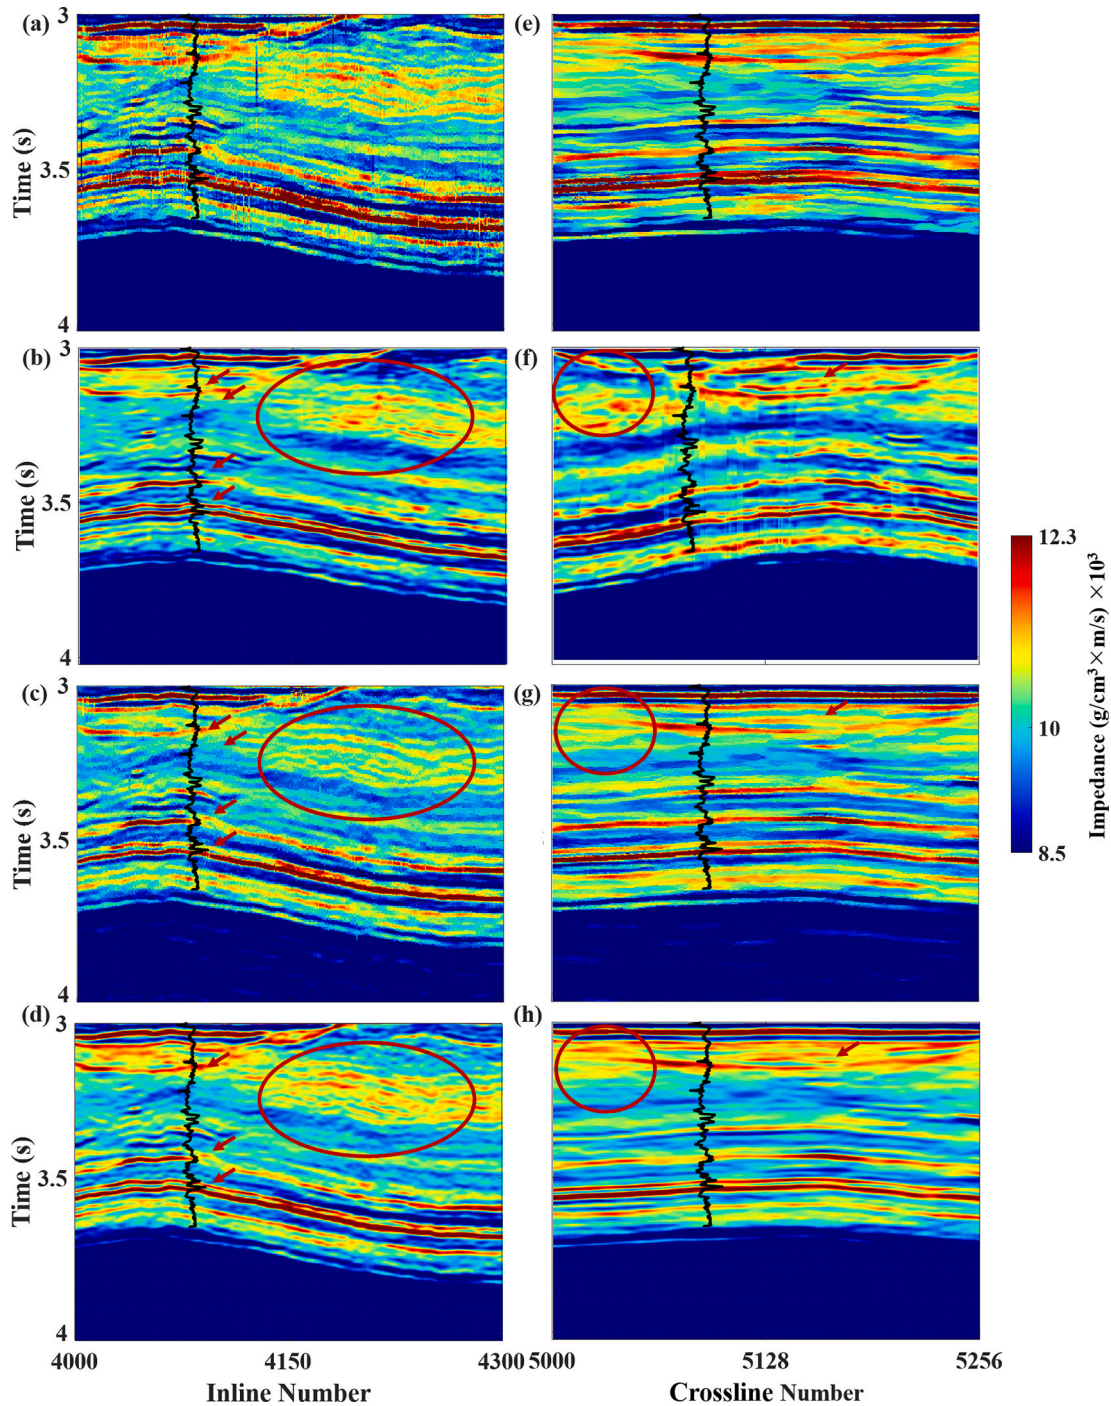


Fig. 11. Comparison of inversion results from different methods. The top row shows GDPS without 3D lateral constraints in (a) crossline 4083 and (e) inline 5093. The second row shows GDPS with 3D lateral constraints in (b) crossline 4083 and (f) inline 5093. The third row shows the results of unsupervised learning in (c) crossline 4083 and (g) inline 5093. The bottom row shows MAP in (d) crossline 4083 and (h) inline 5093. The black curve represents the well-log data, and the red arrows and circles highlight that GDPS provides higher resolution than MAP. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.1. Numerical experiments

We first examine synthetic results. The learning rate l_r is set to 0.4 for the sampling process, and the regularization parameters λ_1, λ_2 are set to 0.4 and 0.2, respectively. The regularization parameters λ_1 and λ_2 are selected through sensitivity analysis on synthetic validation data, optimizing for geological plausibility and data fit. For example, higher λ_1 strengthens low-frequency constraints, while λ_2 controls lateral smoothness. Fig. 3 (a) ~ (f) shows synthetic AI images used for testing

the inversion and the low-frequency models obtained from the true AI model through Gaussian filtering. We define a Gaussian filter function $G(x, t) = \exp\left(-\left(x^2 + (at)^2\right)/2(a\sigma_t)^2\right)/2\pi a\sigma_t^2$ where a is the spatio-temporal eccentricity to indicate filtering intensity. The larger the value a , the stronger the filtering effect in the x -dimension, while the filtering effect in the t -dimension becomes weaker, and vice versa. The parameter σ_t represents the blur radius, and the larger the value of σ_t , the stronger the filtering effect. The synthetic seismic records are generated using Eq.

(1) with different noise levels. A Ricker wavelet is used for generating the synthetic data. The central frequency of Ricker wavelet is 30 Hz with a sampling interval of 2 ms.

To further illustrate this, we plot the frequency spectrum of the low-frequency model along the time dimension $a = 1$, and $\sigma_t = 10, 30$, as shown in Fig. 4. The corresponding central frequencies are approximately 5 Hz and 2 Hz, respectively. The AI model represents a lacustrine shale deposition, where the sedimentary bodies are relatively continuous and have a uniform thickness. Additionally, we correct the reference surface at -200 ms. Between 400 ms and 800 ms, there is a sandstone reservoir with good connectivity. Therefore, the primary focus of our study is the interval between 400 ms and 800 ms. In the intermediate to deep layers, a slight uplift occurs due to the salt dome. The basin center exhibits a lower impedance, and the impedance increases radially outward with the increasing content of sandstone and conglomerate. Thus, for this synthetic AI model data, our main objective is to characterize the sandstone reservoir in the target interval, particularly the layers with good connectivity.

We first use the DPS method to predict AI. The inversion result is shown in Fig. 5(a). The diffusion model can effectively learn prior knowledge in terms of geological structure information, thus providing a useful correction for the gradient updates. However, the inverted result is characterized by poor spatial continuity and cannot provide an accurate impedance value. Therefore, for the reverse sampling process, we introduce low-frequency and lateral constraints to guide the gradient update. Additionally, we employ the first momentum estimate and second raw moment estimate of the gradient to ensure stability during gradient updates.

The inversion result of GDPS is shown in Fig. 5(b), where the input seismic data have a $S/N_{in} = 3$ dB, and the low-frequency model is obtained by applying a filter with a blur radius of 10 to the true AI model. Compared to the DPS method, the GDPS method not only retains its efficiency in handling noise robustness but also improves stability and geological interpretability for poststack seismic inversion. Specifically, GDPS not only preserves the ability to handle complex data distributions effectively but also retains better prior information. Thus, it reduces dependency on low-frequency models. This allows the GDPS method to strike a balance between data-driven prior learning and the use of low-frequency models, resulting in more accurate and geologically meaningful inversion outcomes.

To compare with the MAP solution, we evaluate the noise robustness and the dependence on the low-frequency model for both methods. Fig. 6(a) shows the impact of different blur radii on both methods. When the low-frequency model is of poor quality, both methods struggle to produce reliable results. This aligns with our expectations, as poststack seismic data itself is not well-suited to reflect low-frequency trends. However, as the quality of the low-frequency model improves, the GDPS method demonstrates less dependence on the low-frequency model, and its prediction reliability increasingly diverges from the MAP solution. Fig. 6(b) illustrates the effect of input data S/N on both methods. Both methods show some noise robustness, but since the MAP solution is more reliant on the low-frequency, its overall prediction results are less reliable than those of the GDPS method. Furthermore, as the diffusion model is an excellent denoiser, it still performs well when the noise level exceeds the signal in the seismic data.

To further compare the methods, we present 3D inversion results for both approaches. Fig. 5(c) shows the inversion result of the GDPS method for $S/N_{in} = 3$ dB and a blur radius of 30, while Fig. 5(d) shows the inversion result of the GDPS method for $S/N_{in} = 0.5$ dB and a blur radius of 10. Fig. 5(e) displays the MAP result for $S/N_{in} = 3$ dB and a blur radius of 10, and Fig. 5(f) shows the MAP result for $S/N_{in} = 3$ dB and a blur radius of 30. Both methods can achieve stable inversion results while adhering to geological principles. However, GDPS outperforms MAP in both structural detail and accuracy, exhibits stronger noise robustness, and reduces dependence on the low-frequency model. For the target interval, GDPS more accurately characterizes the distribution

of the sand bodies (black circles in Fig. 5) and provides a more stable identification of salt dome boundaries (black arrows in Fig. 5), which is beneficial for subsequent reservoir prediction. However, neither method can address the issue of inaccurate physical models, and as such, they share common challenges in the inversion of real data. Additionally, due to the inconsistency in input size during training and sampling, we analyzed the impact of input size on GDPS. The results show that our algorithm can handle data larger than the training window, demonstrating its scalability. Due to the high memory requirements of 3D data, we did not conduct further tests on 3D data in this experiment.

3.2. Real data

We then perform predictions for 3D real data. Here, no training data or labels are needed, allowing for direct prediction on real data, which helps improve the generalization ability of deep learning. This eliminates the need for large labeled datasets, making the approach more efficient and scalable. As a result, the method significantly enhances the generalization capability of deep learning models. It can adapt to real data more effectively, demonstrating superior generalization compared to traditional discriminative learning strategies, which rely heavily on labeled training data. Fig. 7(a) shows the seismic records of the area, with the time range spanning from 3000 ms to 4024 ms and a sampling interval of 2 ms. The inline range is from 4000 to 4300, and the crossline range is from 5000 to 5256. The area includes a well (Well 1) with a length of 662 ms. Based on Well 1 and layer information, we constructed a low-frequency model, as shown in Fig. 7(b). Since there is no well-logging data available beyond 3662 ms, an effective low-frequency model could not be established. Additionally, the primary reservoir interval is from 3400 ms to 3600 ms. The key issues in this real data are similar to those in the synthetic data, so we adopted the same approach for prediction.

Fig. 8(a) presents the inversion result obtained using MAP. Overall, the results demonstrate good continuity and stability, aligning well with the geological structure. From the profile view, the resolution appears relatively high. However, when we extract the single-trace results and compare them with well-log data, we can identify some spurious features that caused these results, as shown in Fig. 9, which presents the single-trace prediction. Next, following the approach used for the model data, we applied the GDPS method to predict the results for the real data. Since the low-frequency model is essential, we did not compare the influence of the low-frequency model here. We also compare our method with a physics-guided data-driven approach, as shown in Fig. 8(c). To ensure a fair comparison, none of the methods use labeled well-log data. While the unsupervised method shows stable overall results, single-trace and fine-detail comparisons reveal lower inversion reliability. We do not include supervised methods, as they typically perform well near wells but generalize poorly in regions farther away. Fig. 8(b) shows the result of the GDPS method without the 3D lateral constraint, meaning the prediction is carried out along the crossline direction without any control in the inline direction. As can be seen, the continuity along the inline direction is poor, and there are many geological artifacts that are undesirable (black circles in Fig. 8). After introducing the 3D lateral constraint, the inversion result using the GDPS method is shown in Fig. 8(d), where the lateral continuity has been improved (black circles in Fig. 8). The GDPS method demonstrates better resolution and captures more detailed information (black and red arrows in Fig. 8), while MAP shows more overall continuity but less detailed characterization. Through the 3D time slice direction, the GDPS method demonstrates higher resolution, enabling a more effective description of the sweet spot area and providing reliable support for channel identification or sand body prediction.

To further compare the details of the predictions from the GDPS and MAP, we present the well-logging data to validate the accuracy of our predictions. As shown in Fig. 9, both methods have a high correlation with the well-log data to some extent. Still, the GDPS results show higher

consistency, particularly at the location marked with red arrows. This provides a more solid foundation for reservoir prediction. The GDPS method achieves a S/N_{out} of 23.74 dB, compared to 22.56 dB for MAP, corresponding to a 5% improvement in accuracy relative to the well-log data.

To further analyze the impact of the 3D lateral constraint, we present the results for two profiles: inline 4083 and crossline 5093. Fig. 10 shows the seismic records for both profiles and the low-frequency models for inline 4083 and crossline 5093. The red dashed line indicates the location of Well 1. Specifically, Fig. 10(a) displays the seismic record for inline 4083, Fig. 10(b) shows the seismic record for crossline 5093, Fig. 10(c) presents the low-frequency model for inline 4083, and Fig. 10(d) shows the low-frequency model for crossline 5093.

Fig. 11 presents the inversion results of the GDPS method, both without and with the lateral constraint. We also overlay Well 1 onto the profiles. The GDPS method without the lateral constraint already demonstrates higher resolution, but when the lateral constraint is added, the inversion results show significantly improved lateral continuity and stability. This further illustrates that including the 3D lateral constraint enhances the stability of the 3D inversion results for real data, making the predictions more reliable. We can easily observe that the 3D lateral constraints significantly impact the inversion results.

Additionally, by comparing the areas marked with red arrows and circles, the GDPS method offers more reliable and higher-resolution results. Due to the limited well-log data and layer information, we cannot construct a detailed low-frequency model. Seismic records show many thin layers with unclear reflection characteristics, and MAP faces difficulties in predicting these regions (red circles in Fig. 11). Although MAP demonstrates more continuous geological structures, the GDPS method provides more reliable outcomes when considering both the single-trace and profile results.

4. Discuss

While the proposed GDSP framework shows strong performance, there are several limitations and areas for improvement.

First, we did not include advanced regularization techniques such as structural or geological constraints. This was a deliberate choice to ensure fair comparison, as the baseline methods also omit such constraints. Still, future work could explore incorporating domain-specific regularization to enhance robustness, especially in complex geological settings.

From a computational standpoint, the main cost lies in the training phase, particularly in learning the noise prediction model, which requires many iterations. Our method is trained using 8 NVIDIA H100 GPUs for approximately three days. However, the sampling process is fast once training is complete. Recent methods can also help accelerate diffusion-based sampling. The trained model generalizes well across datasets from different regions, though fine-tuning may be needed when

Appendix A. Diffusion model posterior sampling

In this appendix, we will review the connection between SDEs and the denoising diffusion probabilistic model (DDPM). The forward iteration of DDPM for $t = 1, 2, \dots, T$ can be expressed as (Chan, 2024)

$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_{t-1} + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}_{t-1}, \boldsymbol{\epsilon}_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (\text{A-1})$$

where α_t represents the noise level. As a result, we obtain the adding noise transition step between \mathbf{x}_t and \mathbf{x}_{t-1} in the DDPM, which is same with Eq. (3). Using the closure property of the Gaussian distribution, we obtain the relationship between \mathbf{x}_t and \mathbf{x}_0

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_0, \mathbf{x}_t \sim \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}) \quad (\text{A-2})$$

where we define $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ as the product of all noise levels from $i = 1$ to t . We can also obtain

applying to areas with very different geology.

We compared GDSP primarily with the MAP method and a physics-driven unsupervised neural network, both of which, like GDSP, do not require labeled well-log data. Overall, GDSP delivers higher inversion accuracy, as shown by improved S/N_{out} on both synthetic and real datasets. However, MAP has advantages in lateral continuity and computational speed in simpler cases.

These findings suggest that GDSP is a strong candidate for data-driven seismic inversion, especially when accuracy and generalization are key. Future efforts will focus on improving training efficiency and integrating more effective constraints.

5. Conclusion

This paper presents an impedance inversion method based on a diffusion model, utilizing posterior sampling approximation to mitigate inaccuracies in prior information and enhance performance in noisy environments. The model effectively learns improved prior information by progressively adding and removing noise. Additionally, we integrate 3D lateral constraint inversion techniques and low-frequency information into the sampling process to stabilize the inversion. The inclusion of gradient moment estimation further enhances the algorithm's robustness. Experimental results show that, compared to MAP, our approach performs similarly under high S/N_{in} conditions but has a significant advantage under low S/N_{in} conditions. Our method improves inversion accuracy by 5% compared to the well-logging data, demonstrating its advantage in noisy environments and reducing its dependency on low-frequency models. Furthermore, the proposed approach naturally includes lateral continuity, incorporating lateral continuity without the need to provide local dips before the inversion. Notably, GDPS proves well-suited for the nonlinear AI problem as we did not have to use the logarithmic approximation, making it a promising technique for practical applications in seismic imaging and beyond.

CRedit authorship contribution statement

Zeyang Liu: Writing – original draft. **Dawei Liu:** Writing – review & editing. **Mauricio D. Sacchi:** Writing – review & editing. **Jingye Li:** Writing – review & editing. **Xiaohong Chen:** Supervision. **Wei Zhang:** Writing – review & editing.

Declaration of competing interest

There is no conflict of interest.

Acknowledgements

This work was supported by the Oil & Gas Major Project (2025ZD1405502).

$$\mathbf{x}_0 = \left(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\varepsilon}_0 \right) / \sqrt{\bar{\alpha}_t} \tag{A-3}$$

Then, we can obtain that $\boldsymbol{\varepsilon}_0$ is equal to the t -dependent scaled version of $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ from the Tweedie's Formula (Chung et al., 2022a)

$$\hat{\mathbf{x}}_0 = \left(\mathbf{x}_t + (1 - \bar{\alpha}_t) \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right) / \sqrt{\bar{\alpha}_t} \tag{A-4}$$

By comparing Eqs. (A-3) and (A-4) we obtain

$$\boldsymbol{\varepsilon}_0 = -\sqrt{1 - \bar{\alpha}_t} \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \tag{A-5}$$

Using Bayes' theorem, we obtain

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0) q(\mathbf{x}_{t-1} | \mathbf{x}_0)}{q(\mathbf{x}_t | \mathbf{x}_0)} = \frac{\mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_{t-1}, (1 - \alpha_t) \mathbf{I}) \mathcal{N}(\mathbf{x}_{t-1}; \sqrt{\bar{\alpha}} \mathbf{x}_0, (1 - \bar{\alpha}) \mathbf{I})}{\mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}} \mathbf{x}_0, (1 - \bar{\alpha}) \mathbf{I})} \tag{A-6}$$

It can be found from this equation that the denoising transition step follows a Gaussian distribution with mean ($u_q(\mathbf{x}_t, \mathbf{x}_0)$) and variance ($\Sigma_q(t)$) as follows

$$u_q(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}) \mathbf{x}_t + \sqrt{\bar{\alpha}} (1 - \alpha_t) \mathbf{x}_0}{1 - \bar{\alpha}} \tag{A-7a}$$

$$\Sigma_q^2(t) = \frac{(1 - \alpha_t) (1 - \bar{\alpha})}{1 - \bar{\alpha}} \tag{A-7b}$$

By inserting Eqs. (A-4) to (A-6) we obtain

$$\mathbf{x}_{t-1} = \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}) \mathbf{x}_t + \sqrt{\bar{\alpha}} (1 - \alpha_t) \hat{\mathbf{x}}_0}{1 - \bar{\alpha}} + \Sigma_q(t) \boldsymbol{\varepsilon} \tag{A-8}$$

where

$$\hat{\mathbf{x}}_0 = \left(\mathbf{x}_t + (1 - \bar{\alpha}_t) \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right) / \sqrt{\bar{\alpha}_t} \tag{A-9}$$

By simplifying we obtain

$$\mathbf{x}_{t-1} = \left(\mathbf{x}_t + (1 - \alpha_t) \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right) / \sqrt{\bar{\alpha}_t} + \Sigma_q(t) \boldsymbol{\varepsilon} \tag{A-10}$$

As a result, we obtain the denoising transition step between \mathbf{x}_{t-1} and \mathbf{x}_t in the DDPM, which is same with Eq. (4).

Thus, the reversal unconditional sampling step in DDPM can be implemented by the following steps: (1) Using the trained score function $\mathbf{s}_{\theta'}(\mathbf{x}_t, t)$ in Eq. (5) to approximate $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$, (2) then computing $\hat{\mathbf{x}}_0$ using Eq. (A-9), (3) finally we can compute the \mathbf{x}_{t-1} using Eq. (A-8) and repeat step (1)–(3) until $t = 1$. We have obtained the unconditional calculation method for $\hat{\mathbf{x}}_0$.

Appendix B. Map method

We use the MAP approach for the traditional seismic inversion method. We define $\boldsymbol{\xi} = \ln \mathbf{x}_0$ and $\boldsymbol{\xi}_{\text{Low}} = \ln \hat{\mathbf{x}}_{\text{Low}}$. Then, we can linearize Eq. (1) as

$$\mathbf{y} = \tilde{\mathbf{G}} \boldsymbol{\xi} + \mathbf{n} \tag{B-1}$$

where $\tilde{\mathbf{G}}$ is the linear forward operator. According to Bayes' theorem, the posterior probability after linearization is

$$p(\boldsymbol{\xi} | \mathbf{y}) = C \exp \left[-\frac{1}{2} (\mathbf{y} - \tilde{\mathbf{G}} \boldsymbol{\xi})^T \mathbf{C}_y^{-1} (\mathbf{y} - \tilde{\mathbf{G}} \boldsymbol{\xi}) \right] \exp \left[-\frac{1}{2} (\boldsymbol{\xi} - \boldsymbol{\xi}_{\text{Low}})^T \mathbf{C}_{\boldsymbol{\xi}}^{-1} (\boldsymbol{\xi} - \boldsymbol{\xi}_{\text{Low}}) \right] \tag{B-2}$$

where C is a constant, and \mathbf{C}^{-1} denotes the covariance matrix. Since the conjugacy property of Gaussian distributions, we can obtain the maximum a posteriori estimates as

$$\boldsymbol{\xi}_m = \boldsymbol{\xi}_{\text{Low}} + \left(\tilde{\mathbf{G}}^T \mathbf{C}_y^{-1} \tilde{\mathbf{G}} + \mathbf{C}_{\boldsymbol{\xi}}^{-1} \right)^{-1} \tilde{\mathbf{G}}^T \mathbf{C}_y^{-1} (\mathbf{y} - \tilde{\mathbf{G}} \boldsymbol{\xi}_{\text{Low}}) \tag{B-3}$$

We deliberately assumed no access to such covariances to ensure a simple comparison. In addition, we perform a simple two-trace synchronous inversion in this case.

Data availability

The field data associated with this research are confidential and cannot be released.

References

- Bianchini, L., Forte, E., Pipan, M., 2019. Acoustic impedance estimation from combined harmonic reconstruction and interval velocity. *Geophysics* 84 (3), R385–R400.
- Chan, S., 2024. Tutorial on diffusion models for imaging and vision[J]. *Foundations and Trends in Computer Graphics and Vision* 16 (4), 322–471.
- Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S., 2021. ILVR: Conditioning method for denoising diffusion probabilistic models. *Proceedings of the IEEE/CVF International Conference on Computer Vision* 2021, 14367–14376.
- Chung, H., Kim, J., McCann, M.T., Klasky, M.L., Ye, J.C., 2022a. Diffusion posterior sampling for general noisy inverse problems. *ArXiv preprint arXiv: 2209.14687*.
- Chung, H., Sim, B., Ryu, D., Ye, J.C., 2022b. Improving diffusion models for inverse problems using manifold constraints. *Adv. Neural Inf. Proces. Syst.* 35, 25683–25696.
- Cooke, D.A., Schneider, W.A., 1983. Generalized linear inversion of reflection seismic data. *Geophysics* 48, 665–676.
- Daras, G., Chung, H., Lai, C.H., Mitsufuji, Y., Ye, J.C., Milanfar, P., Delbracio, M., 2024. A survey on diffusion models for inverse problems. *ArXiv preprint arXiv: 2410.00083*.
- Das, V., Pollack, A., Wollner, U., Mukerji, T., 2019. Convolutional neural network for seismic impedance inversion. *Geophysics* 84 (6), R869–R880.
- Dhariwal, P., Nichol, A., 2021. Diffusion models beat Gans on image synthesis. *Adv. Neural Inf. Proces. Syst.* 34, 8780–8794.
- Friedman, A., 1975. *Stochastic differential equations and applications*. In: *Stochastic differential equations*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Fu, L.Y., 2004. Joint inversion of seismic data for acoustic impedance. *Geophysics* 69, 994–1004.
- Gholami, A., 2015. Nonlinear multichannel impedance inversion by total-variation regularization. *Geophysics* 80 (5), R217–R224.
- Gholami, A., 2016. A fast automatic multichannel blind seismic inversion for high-resolution impedance recovery. *Geophysics* 81 (5), V357–V364.
- Hamid, H., Pidlisecky, A., 2015. Multitrace impedance inversion with lateral constraints. *Geophysics* 80 (6), M101–M111.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *Adv. Neural Inf. Proces. Syst.* 33, 6840–6851.
- Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T., 2022. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.* 23, 1–33.
- Kingma, D.P., Ba, J., 2014. Adam: A Method for Stochastic Optimization. *Computer Science*.
- Li, K., Dou, Y., Xiao, Y., Jing, R., Zhu, J., Ma, C., 2024a. TransInver: 3D data-driven seismic inversion based on self-attention. *Geophysics* 89, no. 1, WA127–WA141.
- Li, C., Medina, M.R., Warren, M., Tasrianto, R., Macy, B., Bowling, J., Sena, A., 2024b. Seismic super-resolution method: enhancing reservoir delineation and characterization through high-resolution seismic data. *Lead. Edge* 43, 843–851.
- Lin, Y., Zhang, G., Wang, B., Huang, M., Chen, S., Zhao, C., 2024. An improved stochastic inversion method for 3D elastic impedance under the prior constraints of random medium parameters. *Geoenergy Sci. Eng.* 233, 212421.
- Lines, L.R., Treitel, S., 1984. A review of least-squares inversion and its application to geophysical problems. *Geophys. Prospect.* 32, 159–186.
- Liu, Z., Song, W., Chen, X., Li, W., Li, Z., Liu, G., 2024. High-resolution reservoir prediction method based on data-driven and model-based approaches. *Geophys. Prospect.* 72, 1971–1984.
- Luo, C., 2022. Understanding diffusion models: a unified perspective. *ArXiv preprint arXiv: 2208.11970*.
- Merrifield, T.P., Griffith, D.P., Zamanian, S.A., Gesbert, S., Sen, S., De La Torre Guzman, J., Kuehl, H., 2022. Synthetic seismic data for training deep learning networks. *Interpretation* 10 (3), SE31–SE39.
- Oldenburg, D.W., Scheuer, T., Levy, S., 1983. Recovery of the acoustic impedance from reflection seismograms. *Geophysics* 48, 1318–1337.
- Russell, B., Hampson, D., 1991. Comparison of Poststack Seismic Inversion Methods: SEG Technical Program Expanded Abstracts, 1991, pp. 876–878.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., 2015. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics: International Conference on Machine Learning.
- Song, Y., Ermon, S., 2019. Generative Modeling by Estimating Gradients of the Data Distribution: Advances in Neural Information Processing Systems.
- Song, J., Meng, C., Ermon, S., 2020. Denoising diffusion implicit models. *ArXiv preprint arXiv:2010.02502*.
- Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics.
- Vincent, P., 2011. A connection between score matching and denoising autoencoders. *Neural Comput.* 23, 1661–1674.
- Walker, C., Ulrych, T.J., 1983. Autoregressive recovery of the acoustic impedance. *Geophysics* 48, 1338–1350.
- Wang, Y., Yu, J., Zhang, J., 2022a. Zero-shot image restoration using denoising diffusion null-space model. *ArXiv preprint arXiv:2212.00490*.
- Wang, Y.Q., Wang, Q., Lu, W.K., Ge, Q., Yan, X.F., 2022b. Seismic impedance inversion based on cycle-consistent generative adversarial network. *Pet. Sci.* 19, 147–161.
- Wu, X., Yan, S., Bi, Z., Zhang, S., Si, H., 2021. Deep learning for multidimensional seismic impedance inversion. *Geophysics* 86 (5), R735–R745.
- Yilmaz, O., 2001. *Seismic Data Analysis: Processing, Inversion and Interpretation of Seismic Data*. Society of Exploration Geophysicists.
- Zhang, J., Li, J., Chen, X., Li, Y., Chen, Y., 2021. Robust deep learning seismic inversion with a priori initial model constraint. *Geophys. J. Int.* 225, 2001–2019.