

A Novel Diffusion Model for Pairwise Geoscience Data Generation with Unbalanced Training Dataset

Junhuan Yang¹, Yuzhou Zhang², Yi Sheng¹, Youzuo Lin³, Lei Yang¹

¹George Mason University

²Northeastern University

³The University of North Carolina at Chapel Hill

jyang71@gmu.com, zhang.yuzho@northeastern.edu, ysheng2@gmu.edu, yzlin@unc.edu, lyang29@gmu.edu

Abstract

Recently, the advent of generative AI technologies has made transformational impacts on our daily lives, yet its application in scientific applications remains in its early stages. Data scarcity is a major, well-known barrier in data-driven scientific computing, so physics-guided generative AI holds significant promise. In scientific computing, most tasks study the conversion of multiple data modalities to describe physical phenomena, for example, spatial and waveform in seismic imaging, time and frequency in signal processing, and temporal and spectral in climate modeling; as such, multi-modal pairwise data generation is highly required instead of single-modal data generation, which is usually used in natural images (e.g., faces, scenery). Moreover, in real-world applications, the unbalance of available data in terms of modalities commonly exists; for example, the spatial data (i.e., velocity maps) in seismic imaging can be easily simulated, but real-world seismic waveform is largely lacking. While the most recent efforts enable the powerful diffusion model to generate multi-modal data, how to leverage the unbalanced available data is still unclear. In this work, we use seismic imaging in subsurface geophysics as a vehicle to present “UB-Diff”, a novel diffusion model for multi-modal paired scientific data generation. One major innovation is a one-in-two-out encoder-decoder network structure, which can ensure pairwise data is obtained from a co-latent representation. Then, the co-latent representation will be used by the diffusion process for pairwise data generation. Experimental results on the OpenFWI dataset show that UB-Diff significantly outperforms existing techniques in terms of Fréchet Inception Distance (FID) score and pairwise evaluation, indicating the generation of reliable and useful multi-modal pairwise data.

Introduction

The breakthroughs in generative AI technology have dramatically transformed everyday life, exemplified by the emergence of ChatGPT (OpenAI 2023). In the realm of images, the advent of diffusion models (Sohl-Dickstein et al. 2015; Song and Ermon 2019; Ho, Jain, and Abbeel 2020) has made high-quality image generation a reality. This has significant implications for the real world, as it allows for creating images in various styles tailored to human needs, with applications in production, education, work, and artistic creation (Zhou and Lee 2024). However, the significance

of these advancements is somewhat limited when it comes to scientific data generation. Unlike natural images, generating standalone scientific data presents unique challenges, as such data typically serve specialized purposes and rely heavily on specific scientific contexts and applications.

Some scientific data types, such as seismic waveform data, are challenging for humans to interpret through direct observation (Alcalde et al. 2019). These data require complementary information, like velocity maps, to become meaningful and useful in geophysical research. In geophysics, Full-waveform Inversion (FWI) is a state-of-the-art approach in seismic data processing, designed to construct detailed subsurface models by leveraging the comprehensive information within seismic waveforms (Virieux and Operto 2009). Its ability to deliver high-resolution insights has made FWI invaluable across various subsurface applications, including subsurface energy exploration, earthquake early warning, and carbon capture and sequestration. Currently, widely used data-driven approaches employ machine learning to associate seismic data with subsurface structures, relying on comprehensive training datasets for accurate predictions (Zeng et al. 2022; Zhang and Lin 2020). The literature shows that data-driven methods typically achieve higher spatial resolutions than conventional physics-driven FWI approaches (Lin, Theiler, and Wohlberg 2023).

While data-driven models offer great potential for portable, real-time, and detailed subsurface imaging, they have significant limitations. Unlike computer vision, the subsurface geophysics field is challenged by data scarcity, mainly due to a prevalent culture of non-sharing data. What’s worse, in practical applications, data imbalance presents a significant challenge. Velocity maps, which are more intuitive and understandable for humans, can be more easily simulated through various physical methods. In contrast, seismic data—critical for understanding subsurface structures—is often more difficult and expensive to acquire. This imbalance results in an abundance of velocity maps, while the corresponding seismic data needed to create paired datasets remains limited. Therefore, efficiently generating paired multi-modal data is crucial for achieving accurate and comprehensive subsurface modeling, as it addresses the real-world scarcity of balanced, high-quality datasets.

Recent work has started exploring the simultaneous generation of multi-modal data (Chen et al. 2024). However,

	DDPM	LDMs	MT-Diff	UB-Diff (ours)
Good generation quality	✓	✓	✓	✓
Diffusion on latent	×	✓	✓	✓
Multi-modality data	×	×	✓	✓
Unbalanced data	×	×	×	✓

Table 1: The comparison among classical and SOTA diffusion-based generation approaches and our approach

these methods require a large amount of paired multi-modal data as input, which is uncommon in real-world scenarios, particularly in scientific fields like geophysical and biomedical imaging. Acquiring comprehensive paired datasets is costly and technically challenging due to factors like high data collection costs, the need for specialized equipment, and ethical considerations. Furthermore, existing approaches rely on extensive, well-annotated paired data, limiting their generalizability and applicability across diverse scientific fields. Consequently, the scarcity of paired data, where one modality is abundant but its counterpart is limited, presents a significant barrier in scientific data generation. Overcoming this challenge requires innovative strategies to effectively leverage unpaired or partially paired data to create accurate and valuable paired data.

In this work, we propose UB-Diff, designed to generate paired multi-modal geoscience data simultaneously, specifically seismic waveform and velocity maps. Inspired by a characteristic identified and validated by recent research (Feng et al. 2022a), we use abundant data to train a diffusion-based model with two independent decoders for generating paired data. UB-Diff outperforms state-of-the-art (SOTA) generation approaches on unbalanced data, where one data type is scarce – a common issue in real-world scenarios.

Our main contributions are summarized as follows:

- **Multi-task paired data generation addressing unbalanced data issues:** This study tackles the prevalent problem of unbalanced data in geophysical applications, integrating this into the data generation process.
- **A simple yet effective framework with a matched training scheme:** The proposed diffusion-based model employs a matched training scheme to generate high-quality paired data from unbalanced data successfully.
- **Superior performance over SOTA models:** Experimental results demonstrate that the proposed framework, UB-Diff, outperforms current SOTA models in both paired and single data generation with unbalanced data.

Background and Related Work

Background on Full-waveform Inversion Task

Seismic FWI is a computational technique for imaging subsurface structures, which plays a pivotal role in Geophysics. It is powerful to obtain detailed subsurface structures by analyzing seismic waveforms. Figure 1(a) presents a real-world depiction of flat rock layers. To understand subsurface structures, a source generates waveforms, and a set of receivers receives the waveforms. Figure 1(b) showcases a velocity map reflecting the speed of seismic waves traveling through

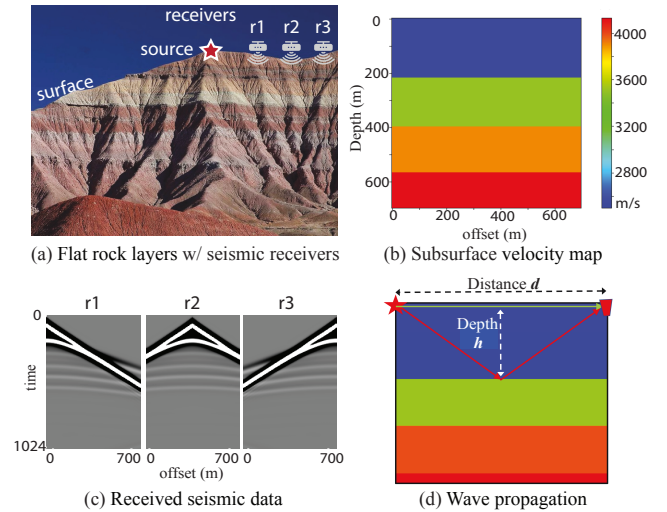


Figure 1: Illustration of FWI: (a) photo of a flat rock layer; (b) velocity map used to show the subsurface structure; (c) seismic waveform obtained from the receivers placed on the surface; and (d) wave propagation in the velocity map.

different subsurface media. However, the direct acquisition of velocity maps presents a significant challenge, as we can only deploy receivers on the ground. Receivers collect seismic data, as depicted in Figure 1(c), which can be used by FWI to construct the velocity maps. A more detailed illustration of wave propagation is shown in Figure 1(d), where the reflection time can be obtained by using distance d between the source and receiver and the depth of the subsurface interface to the surface h .

With the development of machine learning and deep learning and the relative computing capability, data-driven methods provide a promising solution for FWI, replacing the physics-driven approach. The data-driven method employs deep neural networks (DNNs), such as convolutional neural networks (CNNs), to directly learn the inversion operator (Jin et al. 2022). The process usually requires paired seismic data and corresponding velocity maps to train a DNN as supervised learning (Wu and Lin 2019). Recently, the authors in (Feng et al. 2022b) found a near-linear relationship between the seismic waveform and velocity map in the high dimensional space. This inspires us to use a common latent space to represent both data modalities.

Related Work on Generation with Diffusion

Diffusion models (Sohl-Dickstein et al. 2015; Song and Ermon 2019; Ho, Jain, and Abbeel 2020) have recently delivered impressive results in a number of applications (Dhariwal and Nichol 2021; Saharia et al. 2022; Yu et al. 2022; Ramesh et al. 2021; Ruiz et al. 2023; Li et al. 2024, 2022; Wu et al. 2023; Gong et al. 2022; Singh et al. 2023; Huang et al. 2023; Liu et al. 2023; Ramesh et al. 2022; Saharia et al. 2022) and become SOTA generative models.

Table 1 compares characteristics among classical diffusion-based approaches, SOTA multi-modal approaches,

and our method. The Denoising Diffusion Probabilistic Model (DDPM) (Ho, Jain, and Abbeel 2020) utilizes the diffusion process to achieve competitive generative abilities compared to traditional methods like GANs. The advancement of latent diffusion (LDMs) or Stable Diffusion (Rombach et al. 2022) has facilitated the application of the diffusion process to latent variables, achieving improved generation capabilities, enhanced efficiency, and enabling conditional generation. Recently, MT-Diffusion (Chen et al. 2024) was introduced to generate multi-modal data simultaneously by aggregating multiple modalities in the diffusion space. It requires the input data on multi-modality to be pairwise; however, the unbalance of data in real-world applications commonly exists. MT-Diffusion, thus, cannot fully exploit the unbalanced multi-modal data. On the other hand, UB-Diff presented in this work can fully leverage all available data to enhance the generation process, ensuring robust performance even with unbalanced data inputs.

UB-Diff Framework

Figure 2 shows the framework overview of our UB-Diff framework. UB-Diff can utilize all the data (m modal-1 and n modal-2) for training. By doing so, UB-Diff can learn more from the data. In Figure 2, we use a thicker arrow to represent the training process, which can benefit from the m modal-1 data. Although one decoder seems to only utilize n modal-2 data (represented by the thinner arrow), the better-trained co-latent also helps reconstruct the modal-2 data. We will introduce the details in the following subsections.

In this work, we apply UB-Diff to the geoscience application to generate the *paired* multi-modal seismic waveform data and velocity map simultaneously. By using a single modal of data (velocity map or seismic waveform data) from the majority group to train the UB-Diff framework, we can generate the paired two modalities of data (velocity map and seismic waveform). This idea comes from the near-linear relationship between the seismic latent space and the velocity latent space, validated in previous work (Feng et al. 2022b). We consider utilizing one co-latent variable to represent seismic waveform and velocity map once their latent spaces can be aligned to the co-latent space. This process can be done more easily based on this observation. After modeling the two-modality data into the co-latent space, we can do a diffusion process in such a co-latent space and reverse the co-latent variable back to each data space.

Encoder-Decoder Design and Optimizaiton

Figure 3 shows the proposed devised encoder-decoder, a 1-in-2-out network composed of one encoder and two decoders. In this work, we treat the velocity map as the majority group with more data and the seismic waveform as the minority group having less data.

Encoder-decoder design Encoder E is implemented by a CNN to encode the input data, such as a velocity map, which compresses the input to co-latent space to obtain latent \mathbf{z} . Inspired by InversionNet (Wu and Lin 2019), we down-sample the data to 1×1 size. This process is shown as Equation 1, where \mathbf{ma} is the input data from the majority group, such as

velocity map, E is the encoder used, \mathbf{z} is the encoded latent, and c is the channels of latent variable \mathbf{z} .

$$\mathbf{z} = E(\mathbf{ma}), \text{ where } \mathbf{z} \in R^{c \times 1 \times 1} \quad (1)$$

The obtained latent \mathbf{z} will be processed by two decoders: D_s and D_v . The designs of these two decoders are based on the physics property of the data type. First, as seismic data is temporal, we design D_s as a transformer-based decoder, which decodes the latent variable \mathbf{z} to seismic data $seis'$. We perform a linear transformation (i.e., fully connected layer) from \mathbf{z} to \mathbf{z}'_s before D_s . This enables the transformation from the co-latent space to the seismic latent space.

On the other hand, the velocity map is spatial data, and we design D_v using a CNN-based decoder, which decodes the latent variable \mathbf{z} to velocity map vel' . Before D_v , we also implement a fully connected layer to transform the co-latent variable into a velocity latent variable from \mathbf{z} to \mathbf{z}'_v .

Two-Step training optimization for encoder-decoder

The objective of the 1-in-2-out network is to output a pair of seismic data and velocity map simultaneously. To train the network, it is essential to have a pair of input velocity map and the corresponding seismic wave as the training labels. However, when one modality of data is scarce, the network would perform poorly. To optimize the 1-in-2-out network, we propose the two-step training optimization scheme. To fully utilize the majority group of data, we will devise the only encoder to accommodate the data from the majority group and use all the data from the majority group to train the network in the first step through self-supervised learning:

$$\theta_{\text{self}}^* = \arg \min_{\theta} \frac{1}{m} \sum_{i=1}^m \ell_{ma}(f_{ma}(\mathbf{ma}^{(i)}; \theta), \mathbf{ma}^{(i)}) \quad (2)$$

where $\mathbf{ma}^{(i)} \in \mathcal{G}_{ma}$, represents the data from majority group \mathcal{G}_{ma} , which is denoted as $\{\mathbf{ma}^{(i)}\}_{i=1}^m$, f_{ma} represents the network's output function for \mathcal{G}_{ma} , ℓ_{ma} is the loss function for the majority group, and m is the number of majority data. After the network is trained by the data from the majority group, we can then opt to freeze the encoder and decoder or not, and use the minority data to fine-tune the network:

$$\theta^* = \arg \min_{\theta} \frac{1}{n} \sum_{j=1}^n \ell_{mi}(f_{mi}(\mathbf{ma}^{(j)}; \theta_{\text{self}}^*), \mathbf{mi}^{(j)}) \quad (3)$$

where $\mathbf{mi}^{(j)} \in \mathcal{G}_{mi}$ represents data from minority group \mathcal{G}_{mi} , which is denoted as $\{\mathbf{mi}^{(j)}\}_{j=1}^n$, $\mathbf{mi}^{(j)}$ and $\mathbf{ma}^{(j)}$ mean the paired data from the minority group and majority group respectively, f_{mi} represents the network's output function for \mathcal{G}_{mi} , ℓ_{mi} is the loss function for the minority data, θ_{self}^* is the optimal model from the Equation 2 in the first self-supervised training step, and n is the number of minority data and $m \gg n$. Following the classical data-driven FWI work (Wu and Lin 2019), we design the loss as follows:

$$\begin{aligned} \ell_{ma} &= \gamma_1 \|f_{ma}(\mathbf{ma}; \theta) - \hat{\mathbf{ma}}\|_1 + \gamma_2 \|f_{ma}(\mathbf{ma}; \theta) - \hat{\mathbf{ma}}\|_2^2 \\ \ell_{mi} &= (1 - F) \times \ell_{ma} + (\gamma_3 \|f_{mi}(\mathbf{ma}; \theta_{ma}^*) - \hat{\mathbf{mi}}\|_1 + \\ &\gamma_4 \|f_{mi}(\mathbf{ma}; \theta_{ma}^*) - \hat{\mathbf{mi}}\|_2^2) \end{aligned} \quad (4)$$

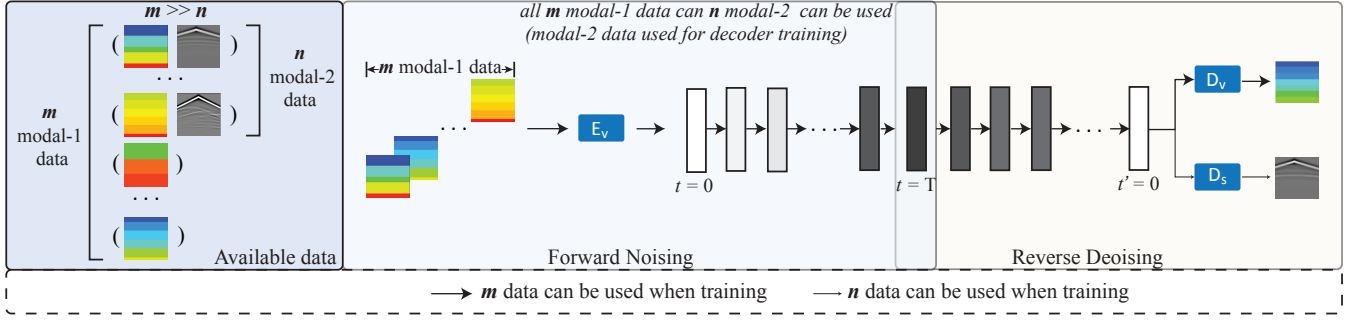


Figure 2: Overview of UB-Diff, which utilizes all available data, benefiting the whole process, especially the diffusion process.




Figure 3: 1-in-2-out network for seismic waveform and velocity map. Using the example with the velocity map as the majority data and the seismic waveform as the minority.




Figure 4: Diffusion process of UB-Diff, generating latent of \mathbf{ma} and \mathbf{mi} simultaneously.

where $F \in \{0, 1\}$ refers to the freeze flag, $\gamma_1 - \gamma_4$ are used to control the impact of L_1 and L_2 for each data modality.

Diffusion Process

In UB-Diff, the forward process will collect the data information in the co-latent space by collecting information from the majority group. Figure 4 shows the diffusion process of UB-Diff. By building the relationship between the data from the majority group and minority group and encoding the data from the majority group into the co-latent space, UB-Diff can model the single-modal data from the majority group but generate data both in the majority and minority groups.

Forward process The forward process of the UB-Diff will be conditioned on the majority group. The definition of the joint distribution at time t based on the paired data $\mathbf{X} =$

$(\mathbf{ma}, \mathbf{mi})$ and the diffusion latent variable \mathbf{z}_t at timestep t , conditioned on the data at timestep $t - 1$ can be shown as:

$$q(\mathbf{z}_t, \mathbf{X} | \mathbf{z}_{t-1}) = q(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{ma})q(\mathbf{ma}) \quad (5)$$

where $q(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{ma})$ represents the distribution of \mathbf{z}_t from \mathbf{z}_{t-1} and $q(\mathbf{ma})$ denotes prior distributions of data from the majority group. Similar to the classical DDPM (Ho, Jain, and Abbeel 2020), the posterior transition distribution can be shown as follows:

$$\begin{aligned} q(\mathbf{z}_t | \mathbf{z}_0) &= \mathcal{N}(\mathbf{z}_t; \alpha_t \mathbf{z}_0, \sigma_t I) \\ &= \mathcal{N}(\mathbf{z}_t; \alpha_t E(\mathbf{ma}_0), \sigma_t I) \end{aligned} \quad (6)$$

Reverse process To reverse the forward process, we define the reverse process as $p_\phi(\mathbf{z}_{t-1}, \mathbf{ma}, \mathbf{mi} | \mathbf{z}_t)$, where ϕ represents the reverse model parameters. Although the reverse transition at time t can be decomposed, in this geoscience application, the noise in the diffusion process may destroy the reconstructed velocity map and seismic waveform. Unlike natural images, e.g., the eyes can be big or small, a small error in the scientific data can mislead the physics. Thus, we give up to decompose the latent at time t (when $t > 0$), i.e., $p_\phi(\mathbf{ma}, \mathbf{mi} | \mathbf{z}_t)$, to avoid this issue. Thus, we still have a Gaussian distribution with mean and covariance denoted as $\mu_\phi(\mathbf{z}_t, \mathbf{ma}, t)$ and $\sigma_t^2 \mathbf{I}$. To enhance the reverse process and the reconstruction data quality, we follow the work (Salimans and Ho 2022) to parameterize the U-Net model $\mathbf{u}_\phi(\mathbf{z}_t, t)$ to predict an intermediate variable $\mathbf{u} = \alpha_t \epsilon - \sigma_t \mathbf{z}_0$. The variables \mathbf{ma} and \mathbf{mi} and in the $p_\phi(\mathbf{ma}, \mathbf{mi} | \mathbf{z}_0)$ will be mapped to their own latent space, \mathbf{z}'_{ma} and \mathbf{z}'_{mi} , according to two fully-connected layers. Two decoders D_{ma} and D_{mi} will be used to decode the \mathbf{z}'_{ma} and \mathbf{z}'_{mi} back to their own data space.

Training and generation A simple training loss for the UB-Diff is used following the work (Salimans and Ho 2022). The U-Net will be updated according to the L2 loss $\|u_t - u_\phi(\mathbf{z}_t, t)\|_2^2$, where $\mathbf{z}_t \sim q(\mathbf{z}_t | \mathbf{z}_0)$ and $\mathbf{z}_0 = E(\mathbf{ma})$. Thus, all the data from the majority group can be used to train the diffusion model. Similar to the classical DDPM, we also apply the stochastic optimization when training the UB-Diff. The data from the majority group \mathbf{G}_{ma} will be encoded to \mathbf{z}_0 , by the pre-trained encoder E , where the weight is obtained in Equation 3. In the training process, a mini-batch of randomly selected timesteps t will be sampled, and

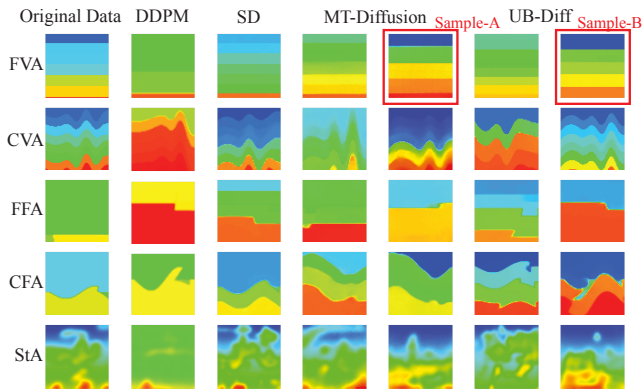


Figure 5: Generated velocity map by baselines and UB-Diff.

the corresponding mini-batch of \mathbf{z}_t will be calculated from the encoded \mathbf{z}_0 . The mini-batch of \mathbf{z}_t will be fed to the U-Net to predict the intermediate variable \mathbf{u} , and the U-Net parameters will be updated through gradient descent based on the L2 loss described above.

In the generation process, our UB-Diff aims to generate the paired data in both the majority domain and the minority domain. A pure Gaussian noise \mathbf{z}'_T with the same size as the \mathbf{z} will be randomly sampled. It will be denoised gradually and becomes latent (i.e., \mathbf{z}'_0) at the last diffusion step. The paired types of data can be generated through two fully-connected layers and two decoders.

Experiment

To evaluate our UB-Diff framework, we employ several commonly used geoscience datasets and apply the velocity maps as the data of majority group \mathcal{G}_{ma} and seismic wave as the data of minority group \mathcal{G}_{mi} . In this section, we will introduce the detailed experimental setup and results. As a work of paired data generation, we evaluate the generation quality for single modality data and pairwise quality from macro, pairwise, and micro perspectives.

Experimental Setup

Dataset: We employ five datasets: FlatVel-A, CurveVel-A, FlatFault-A, CurveFault-A, and Style-A, from the openFWI (Deng et al. 2022). To evaluate the UB-Diff, we follow the classical FWI work and the dataset to choose around 80% of velocity maps from each dataset (24,000 from FlatVel-A and CurveVel-A, 48000 from FlatFault-A and CurveFault-A, 60000 from Style-A). However, only 1,000 and 5,000 corresponding paired seismic data are used.

Metrics: Like other generation tasks, we employ the FID (Heusel et al. 2017) to evaluate the similarity between the generated and original data from a macro perspective. As a paired data generation framework, we still evaluate the pairwise quality of the generated data. Thus, we employ the classical InversionNet (Wu and Lin 2019) and train the InversionNet by the generated data and test on the original dataset (following the work of (Khader et al. 2023; Saragih, Hibi,

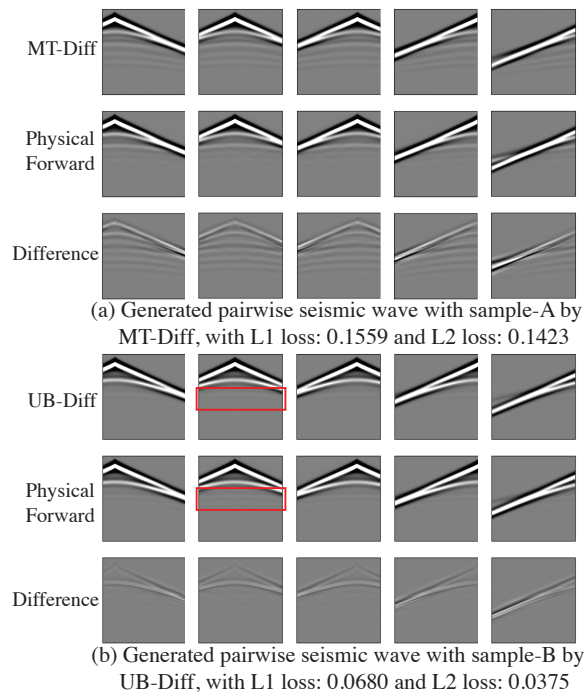


Figure 6: Generated seismic waveform samples by MT-Diffusion (sample-A in Figure 5) and UB-Diff (sample-B in Figure 5). Five columns refer to five channels of the data. Besides direct waves, UB-Diff can generate reflected waves more accurately than MT-Diffusion.

and Tyrrell 2024)). Thus, we employ three performance metrics to show the pairwise of the generated data: (1) Structural Similarity Index (SSIM), (2) Mean absolute error (MAE), and (3) Mean squared error (MSE). We also compare the seismic data generated by machine learning and physical forward modeling to evaluate from a micro perspective.

Competitors: For comparison, we reproduce the MT-Diffusion (Chen et al. 2024), the first multi-modal generation framework, for the paired data generation. We also compare with the classical DDPM for velocity map (Yang et al. 2024), SD (Rombach et al. 2022) (reproduced with same encoder and decoder) for both single modality data generation. To better evaluate our two-step training scheme, we also show the results without two-step training (shown as “UB-Diff w/o opt”), utilizing only 1,000 or 5,000 data when training the encoder-decoder.

Training Setting: The encoder-decoder will be trained 1000 epochs for both two steps (Equation 2 and 3). We both freeze or not for the second step of training and report a better result. The parameters (γ_1 to γ_4) for encoder and decoder training are all set to 1. The dimension of \mathbf{z} is set to 128. Timestep, T , is set as 256 for all diffusion-based models.

Experimental Results

Paired multi-task data generation In the first set of experiments, we evaluate the performance of our framework, UB-Diff, in pairwise multi-task data generation.

Dataset	Data type	All Vel + 1k Seis			All Vel + 5k Seis		
		MT-Diff	UB-Diff w/o opt	UB-Diff	MT-Diff	UB-Diff w/o opt	UB-Diff
FlatVel-A	FID of V. ↓	336.5417	339.6236	16.4475	99.8874	63.2182	15.6393
	FID of S. ↓	318.2893	158.7187	97.0339	71.5330	96.7443	9.2542
CurveVel-A	FID of V. ↓	485.0098	449.9948	103.2765	365.0310	222.7844	65.5289
	FID of S. ↓	214.9118	174.1814	18.9828	64.4392	9.7614	8.7620
FlatFault-A	FID of V. ↓	32.2887	27.3071	16.5089	27.9260	21.8450	21.4411
	FID of S. ↓	53.2955	30.7746	26.8663	49.3511	14.8987	14.8630
CurveFault-A	FID of V. ↓	87.8209	32.1489	14.5018	51.2559	29.7906	21.6959
	FID of S. ↓	206.0490	123.2481	98.9439	57.7593	31.5317	31.8391
Style-A	FID of V. ↓	11.8603	4.9723	0.5452	10.6906	4.6962	0.5394
	FID of S. ↓	3221.2707	53.1792	91.2248	1623.8443	207.8604	39.1894

Table 2: FID score of multi-task generation compared with MT-Diffusion (shown as MT-Diff) and UB-Diff.

Dataset	Metrics	All Vel + 1k Seis			All Vel + 5k Seis		
		MT-Diff	UB-Diff w/o opt	UB-Diff	MT-Diff	UB-Diff w/o opt	UB-Diff
FlatVel-A	MAE ↓	0.2650	0.2153	0.1291	0.1677	0.1173	0.0316
	MSE ↓	0.1154	0.0814	0.0358	0.0579	0.0344	0.0040
	SSIM ↑	0.6511	0.6944	0.7881	0.7394	0.7997	0.9467
CurveVel-A	MAE ↓	0.2530	0.2655	0.1737	0.2415	0.1568	0.1412
	MSE ↓	0.1097	0.1260	0.0609	0.1147	0.0548	0.0480
	SSIM ↑	0.5949	0.5936	0.6680	0.5919	0.6860	0.6922
FlatFault-A	MAE ↓	0.2357	0.2121	0.1902	0.1155	0.0908	0.0861
	MSE ↓	0.0998	0.0910	0.0747	0.0397	0.0260	0.0237
	SSIM ↑	0.7224	0.7384	0.7670	0.8383	0.8783	0.8763
CurveFault-A	MAE ↓	0.2729	0.1761	0.1645	0.1831	0.1020	0.1036
	MSE ↓	0.1331	0.0709	0.0643	0.0810	0.0333	0.0336
	SSIM ↑	0.5801	0.7744	0.7807	0.7619	0.8516	0.8552
Style-A	MAE ↓	0.1716	0.1539	0.1568	0.1849	0.1847	0.1494
	MSE ↓	0.0483	0.0422	0.0446	0.0601	0.0621	0.0418
	SSIM ↑	0.6749	0.7072	0.6909	0.6547	0.6494	0.7027

Table 3: Evaluation of the pairwise of generated multi-task data, through training InversionNet (Wu and Lin 2019) by the 10,000 generated paired data and testing on the original dataset.

Table 2 shows the FID scores for simultaneously generated velocity maps and seismic waveforms. The column “All Vel + 1k Seis” in this table reports the results based on all velocity maps (details in Sec. Experimental Setup) and 1,000 paired seismic waveforms. In comparison, “All Vel + 5k Seis” reports results for all velocity maps and 5,000 paired seismic waveforms. Under each setting, “MT-Diff” refers to the SOTA multi-modal generation approach, MT-Diffusion. “UB-Diff” and “UB-Diff w/o opt” represent our approach with and without the proposed training scheme. The “FID of V.” and “FID of S.” indicate the FID scores for the generated velocity maps and seismic waveforms, respectively.

For the first group with very limited seismic waveforms available, UB-Diff w/o opt outperforms MT-Diff in all generations, except for a close FID score in the velocity generation of FlatVel-A. When we apply the matched training scheme, shown as UB-Diff, the FID scores for velocity maps decrease dramatically across all datasets. The FID scores for seismic data continue to decrease significantly on four datasets, except for Style-A, but remain much lower than those achieved by MT-Diff. For example, on FlatVel-A, UB-Diff achieves FID scores of 16.45 and 97.03 for the velocity map and seismic data, respectively, which are much lower than MT-Diff. In the second group, with 5,000 seismic wave-

form data available, UB-Diff achieves the lowest FID scores across all datasets for velocity map and seismic waveform generation. MT-diffusion seems to perform poorly on Style A. This is because the decomposition loss may destroy the diffusion loss, causing poor generation performance.

In addition to evaluating the generation of multi-task data, we also validated the pairwise quality of the generated data. We generated 10,000 data pairs for each dataset to train InversionNet, a classic neural network for performing FWI tasks. The entire original dataset was used as the test set, and the results are shown in Table 3. We used metrics from previous FWI work (Wu and Lin 2019), including MAE, MSE, and SSIM, to evaluate the FWI performance. These metrics assess the quality and fidelity of the generated data when trained on InversionNet and tested against the original datasets. For the “All Vel + 1k Seis” scenario, UB-Diff achieved the best results across almost all metrics, with significant improvements compared to MT-Diff, especially in datasets like FlatVel-A and CurveFault-A. Only in the Style-A dataset does UB-Diff w/o opt to achieve slightly better performance than UB-Diff. When more seismic data was available (“All Vel + 5k Seis”), UB-Diff continued to lead in performance, maintaining the lowest error rates and highest structural similarity in most datasets, except for CurveFault-

Datasets	DDPM	SD	UB-Diff-1k	UB-Diff-5k
FlatVel-A	126.5761	20.3334	16.4475	15.6393
CurveVel-A	430.3813	216.1318	103.2765	65.5289
FlatFault-A	74.8551	20.0942	16.5089	21.4411
CurveFault-A	210.5555	27.7611	14.5018	21.6959
Style-A	98.1559	1.1186	0.5452	0.5394

Table 4: FID score of single-task generation for velocity map

A, where results were similar between UB-Diff w/o opt and UB-Diff. Figure 5 shows the generated velocity maps by DDPM, SD, MT-Diffusion, and UB-Diff. The generated velocity maps based on FlatVel-A (FVA), CurveVel-A (CVA), FlatFault-A (FFA), CurveFault-A (CFA), and Style-A (StA) are shown in the first to fifth row, respectively. We also visualize the generated seismic waveform from two similar velocity maps (sample-A and sample-B in Figure 5 by MT-Diffusion and UB-Diff) in Figure 6. The physical forward modeling is employed to show how accurate the generated seismic wave is. Specifically, Figure 6 shows the seismic waveform generated by machine learning and physical forward modeling, and the difference between them. Figure 6 (a) and (b) show the result of MT-Diffusion and UB-Diff. We can easily observe that, while both approaches can generate the dominant direct wave, UB-Diff more accurately captures the minor reflected wave (in the red box). This indicates that UB-Diff is better at generating more reliable data. More generated samples are shown in the Appendix.

Comparison with single task generation approach In the second set of experiments, we evaluate our method by comparing it with classical single-task generation methods. Competitors can use all available velocity maps to generate velocity maps. Similarly, our method also uses all velocity maps and 1,000 and 5,000 seismic waveforms to enhance the generation process. Competitors use 1,000 and 5,000 seismic data points to generate seismic data. Our method, however, can also utilize velocity maps to aid the entire process. Since the process for UB-Diff is the same as the first set of experiments, we continue to use the results, and for convenience, we put it in Table 4. Table 4 presents the FID scores for velocity map generation across different methods and datasets. Competitors include DDPM and SD, compared with our method, UB-Diff, using 1,000 and 5,000 seismic data points (UB-Diff-1k and UB-Diff-5k, respectively). For the FlatVel-A dataset, UB-Diff achieves FID scores of 16.45 (UB-Diff-1k) and 15.64 (UB-Diff-5k), outperforming DDPM (126.58) and SD (20.33). In CurveVel-A, our method shows substantial improvement with scores of 103.28 (1k) and 65.53 (5k), compared to DDPM (430.38) and SD (216.13). In FlatFault-A, UB-Diff-1k excels with an FID of 16.51; in CurveFault-A, it achieves the lowest FID of 14.51. For Style-A, UB-Diff-5k slightly outperforms UB-Diff-1k, achieving an FID of 0.54. Overall, UB-Diff effectively generates high-quality velocity maps, leveraging seismic data to surpass traditional single-task methods.

We compare our approach with SD for seismic waveform generation since the size of seismic waveforms is significantly larger than velocity maps, making SD more suitable

Datasets	1k Seis + all Vel		5k Seis + all Vel	
	SD	UB-Diff	SD	UB-Diff
FlatVel-A	122.7463	70.2228	23.1423	5.4164
CurveVel-A	94.5071	18.9828	33.4856	8.7620
FlatFault-A	19.8758	26.3492	9.6232	14.8630
CurveFault-A	46.7977	67.5801	21.8252	21.8139
Style-A	64.3933	53.1792	10.7206	7.5269

Table 5: FID score of single-task generation for seismic waveform

for this task. Our UB-Diff framework focuses on optimizing the generation of seismic waveforms (the minority) first, followed by velocity maps (the majority). This involves performing Equation 3 initially, followed by Equation 2.

For the “1k Seis + all Vel” scenario, UB-Diff consistently achieves better performance than SD in 3 datasets, including FlatVel-A, CurveVel-A, and Style-A, with FID scores of 70.22, 18.98, and 53.18. In the “5k Seis + all Vel” scenario, UB-Diff leads in most datasets except for FlatFault-A. The FID scores continue to showcase our UB-Diff’s superior capability in generating high-quality seismic waveforms.

Overall, for the single-task generation, UB-Diff performs better than classical generation approaches in 5 tasks of velocity map generation (5 datasets) and 7 out of 10 tasks (utilizing 1,000 and 5,000 seismic waveforms in 5 datasets). These results highlight the robustness and adaptability of our framework in handling various data scenarios, demonstrating its potential for broader applications in scientific data generation where data availability and quality are often challenging constraints. UB-Diff effectively leverages available data to produce superior outcomes, making it a valuable tool in fields requiring high-quality multi-modal data synthesis.

Conclusion

In this work, we introduced UB-Diff, a novel diffusion model designed for multi-modal paired scientific data generation, specifically addressing the challenge of data imbalance in multi-modal scientific applications. By leveraging a one-in-two-out network, UB-Diff effectively generates a colatent representation from a single modality of data, which is then utilized in the diffusion process. Our experimental results on the OpenFWI dataset demonstrate that UB-Diff significantly outperforms existing techniques in generating reliable and useful pairwise data, as evidenced by improved FID scores and better performance in FWI tasks. This advancement highlights the potential of UB-Diff in overcoming data scarcity and modality imbalance, paving the way for more accurate and comprehensive scientific modeling.

Acknowledgments

We gratefully acknowledge the support of the startup funding (NO.170662) from George Mason University. Y. Lin acknowledges support from the University of North Carolina’s School of Data Science and Society through a faculty startup grant. This project was also supported by resources from the Office of Research Computing at George Mason University and partially funded by NSF grants (NO.2018631).

References

- Alcalde, J.; Bond, C. E.; Johnson, G.; Kloppenburg, A.; Ferrer, O.; Bell, R.; and Ayarza, P. 2019. Fault interpretation in seismic reflection data: an experiment analysing the impact of conceptual model anchoring and vertical exaggeration. *Solid earth*, 10(5): 1651–1662.
- Chen, C.; Ding, H.; Sisman, B.; Xu, Y.; Xie, O.; Yao, B. Z.; Tran, S. D.; and Zeng, B. 2024. Diffusion Models for Multi-Task Generative Modeling. *arXiv preprint arXiv:2407.17571*.
- Deng, C.; Feng, Y.; Feng, S.; Jin, P.; Zhang, X.; Zeng, Q.; and Lin, Y. 2022. OpenFWI: Benchmark Seismic Datasets for Machine Learning-Based Full Waveform Inversion. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022) Datasets and Benchmarks Track*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Feng, Y.; Chen, Y.; Feng, S.; Jin, P.; Liu, Z.; and Lin, Y. 2022a. An Intriguing Property of Geophysics Inversion. In *The Thirty-ninth International Conference on Machine Learning (ICML)*.
- Feng, Y.; Chen, Y.; Feng, S.; Jin, P.; Liu, Z.; and Lin, Y. 2022b. An intriguing property of geophysics inversion. In *International Conference on Machine Learning*, 6434–6446. PMLR.
- Gong, S.; Li, M.; Feng, J.; Wu, Z.; and Kong, L. 2022. Dif-fuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, R.; Huang, J.; Yang, D.; Ren, Y.; Liu, L.; Li, M.; Ye, Z.; Liu, J.; Yin, X.; and Zhao, Z. 2023. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *International Conference on Machine Learning*, 13916–13932. PMLR.
- Jin, P.; Zhang, X.; Chen, Y.; Huang, S.; Liu, Z.; and Lin, Y. 2022. Unsupervised Learning of Full-Waveform Inversion: Connecting CNN and Partial Differential Equation in a Loop. In *The Tenth International Conference on Learning Representations (ICLR)*.
- Khader, F.; Müller-Franzes, G.; Tayebi Arasteh, S.; Han, T.; Haarburger, C.; Schulze-Hagen, M.; Schad, P.; Engelhardt, S.; Baeßler, B.; Foersch, S.; et al. 2023. Denoising diffusion probabilistic models for 3D medical image generation. *Scientific Reports*, 13(1): 7303.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35: 4328–4343.
- Li, Y.; Wang, H.; Jin, Q.; Hu, J.; Chemerys, P.; Fu, Y.; Wang, Y.; Tulyakov, S.; and Ren, J. 2024. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds. *Advances in Neural Information Processing Systems*, 36.
- Lin, Y.; Theiler, J.; and Wohlberg, B. 2023. Physics-Guided Data-Driven Seismic Inversion: Recent Progress and Future Opportunities in Full Waveform Inversion. *IEEE Signal Processing Magazine*, 40: 115–133.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv preprint arXiv:2301.12503*.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*, 8821–8831. Pmlr.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22500–22510.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Saragih, D. G.; Hibi, A.; and Tyrrell, P. N. 2024. Using diffusion models to generate synthetic labeled data for medical image segmentation. *International Journal of Computer Assisted Radiology and Surgery*, 1–11.
- Singh, M.; Cambronero, J.; Gulwani, S.; Le, V.; Negreanu, C.; and Verbruggen, G. 2023. Codefusion: A pre-trained diffusion model for code generation. *arXiv preprint arXiv:2310.17680*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.

- Virieux, J.; and Operto, S. 2009. An Overview of Full-waveform Inversion in Exploration Geophysics. *Geophysics*, 74(6): WCC1–WCC26.
- Wu, T.; Fan, Z.; Liu, X.; Zheng, H.-T.; Gong, Y.; Jiao, J.; Li, J.; Guo, J.; Duan, N.; Chen, W.; et al. 2023. Ar-diffusion: Auto-regressive diffusion model for text generation. *Advances in Neural Information Processing Systems*, 36: 39957–39974.
- Wu, Y.; and Lin, Y. 2019. InversionNet: An Efficient and Accurate Data-Driven Full Waveform Inversion. *IEEE Transactions on Computational Imaging*, 6: 419–433.
- Yang, J.; Wang, H.; Sheng, Y.; Lin, Y.; and Yang, L. 2024. EdGeo: A Physics-guided Generative AI Toolkit for Geophysical Monitoring on Edge Devices. In *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 1–6.
- Yu, J.; Xu, Y.; Koh, J. Y.; Luong, T.; Baid, G.; Wang, Z.; Vasudevan, V.; Ku, A.; Yang, Y.; Ayan, B. K.; et al. 2022. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3): 5.
- Zeng, Q.; Feng, S.; Wohlberg, B.; and Lin, Y. 2022. InversionNet3D: Efficient and Scalable Learning for 3-D Full-Waveform Inversion. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–16.
- Zhang, Z.; and Lin, Y. 2020. Data-driven seismic waveform inversion: A study on the robustness and generalization. *IEEE Transactions on Geoscience and Remote Sensing*, 58(10): 6900–6913.
- Zhou, E.; and Lee, D. 2024. Generative artificial intelligence, human creativity, and art. *PNAS nexus*, 3(3): pgae052.