

Article

MBES-DDPM: Multibeam Echo Sounder Bathymetry Swath Gap Reconstruction Based on Denoising Diffusion Probability Model

Jianbing Chen ^{1,2}, Ziyin Wu ^{1,2,3,*}, Dineng Zhao ² , Xianhai Bu ^{1,4}, Jieqiong Zhou ², Jihong Shang ² ,
Mingwei Wang ² and Yang Liu ² 

¹ College of Geodesy and Geomatics, Shandong University of Science and Technology, Qingdao 266590, China; chen@sdust.edu.cn (J.C.); buxianhai0826@sust.edu.cn (X.B.)

² State Key Laboratory of Submarine Geoscience, Second Institute of Oceanography, Ministry of Natural Resources, Hangzhou 310012, China; zhaodineng@sio.org.cn (D.Z.); janezhou@sio.org.cn (J.Z.); jhshang@sio.org.cn (J.S.); wangmw@sio.org.cn (M.W.); liuy@sio.org.cn (Y.L.)

³ School of Oceanography, Shanghai Jiao Tong University, Shanghai 200030, China

⁴ Key Laboratory of Ocean Geomatics, Ministry of Natural Resources, Qingdao 266590, China

* Correspondence: ziywu@sio.org.cn

Highlights

What are the main findings?

- A novel generative AI model, MBES-DDPM, is proposed, which is the first application of a denoising diffusion probabilistic model (DDPM) for restoring swath gaps in multibeam echo sounder (MBES) bathymetric data.
- MBES-DDPM achieves superior performance, demonstrating an average reduction in RMSE of at least 34.21% and an average increase in PSNR of over 3.71 dB compared to baseline methods, while best preserving terrain slope accuracy.

What are the implications of the main findings?

- This research establishes a pioneering framework that demonstrates the significant potential of advanced generative AI paradigms, coupled with multisource geophysical data fusion, for solving key reconstruction challenges in underwater remote sensing.
- It will contribute to the advancement of the Seabed 2030 Project and enhance the quality of seamless global seafloor topography modeling.

Abstract

The multibeam echo sounder (MBES) is a key tool for acquiring high-precision seabed topographic data. However, measurement gaps resulting from its swath-based measurement mode are prevalent, severely compromising the completeness of seabed terrain modeling. To address this issue, this study first categorizes multibeam data gaps into two data degradation patterns with clear hydrographic survey backgrounds: “random degradation” and “rule-based degradation.” Based on this categorization, a highly realistic training dataset that closely matches actual conditions is constructed. To improve the reconstruction accuracy and topographic fidelity, a novel multibeam echo sounder data reconstruction model, the MBES-DDPM, is proposed. Based on the denoising diffusion probabilistic model (DDPM) framework, this model innovatively incorporates gravity anomaly data as prior knowledge. Then, with a designed multisource data fusion guidance mechanism, macro-topographic structural constraints are injected during the diffusion process. Furthermore, a targeted quantitative and qualitative evaluation system is established. The experimental results show that compared with the baseline methods, the MBES-DDPM achieves the best performance across various complex scenarios. Its restored results exhibit an average reduction in root mean square error of at least 34.21% and an average increase



Academic Editor: Francesco Serafino

Received: 9 January 2026

Revised: 30 January 2026

Accepted: 2 February 2026

Published: 3 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

in peak signal-to-noise ratio of more than 3.71 dB. Furthermore, it achieves the highest reconstruction fidelity in terms of the terrain slope accuracy metrics. Thus, this research provides a new and reliable solution for accurately restoring large-scale MBES data.

Keywords: multibeam echo sounder; seabed topography modeling; void area reconstruction; diffusion model

1. Introduction

The seabed digital elevation model (DEM) is a digital representation of the real-world seabed topography, providing an essential spatial framework for a series of marine scientific studies, including those that have revealed the dynamics of plate tectonics [1], elucidated marine geological activities [2–5], explored the Earth’s internal structure [6], and delineated seafloor substrates [7]. At the application level, the seafloor DEM directly support subsea pipeline routing planning and deployment [8], subsea oil and gas exploration [9], mineral resource prospecting [10], and maritime safety assurance [11]. Consequently, the resolution and completeness of the seabed DEM determine the credibility of scientific discoveries and the reliability of engineering decisions based upon them.

Modern seafloor topography detection primarily involves two technical approaches: indirect inversion and direct measurement. Indirect inversion techniques include two main methods: gravity anomaly-based [12] methods and remote sensing image-based [13] methods. Among these, gravity anomaly-based bathymetric inversion can rapidly outline the macrostructure of the seabed topography over large areas, but its resolution is limited to approximately 1 arc-minute, and it tends to lose a significant amount of high-frequency detail in the seabed terrain [14]. Moreover, remote sensing-based depth inversion demands higher water quality [15,16] and is limited to depths of tens of meters or shallower. Due to the long-range penetration capability of sound waves through water, acoustic sounding systems constitute the dominant direct measurement method [17]. The multibeam echo sounder (MBES) employs fan-shaped beam scanning, resulting in the formation of water depth swaths with a certain coverage width through which the survey vessel moves [18], thereby achieving high-precision, full-coverage sweep measurements of localized sea areas.

Crowdsourced bathymetry represents a novel paradigm for extending seabed mapping from local to global scales. Its core concept involves integrating real-time bathymetric data collected by ordinary vessels, fishing boats, and other marine participants during routine voyages to construct high-resolution seabed topography models [19]. To accelerate global seafloor mapping, GEBCO launched the Seabed 2030 Project [20] with support from the International Hydrographic Organization and the United Nations Educational, Scientific and Cultural Organization. This initiative has adopted the crowdsourced bathymetry mode and data, while also incorporating the GMRT [21] compiled dataset and MBES bathymetric data from various regions worldwide, with the goal of producing a complete map of the world’s ocean floor by 2030. Driven by the Seabed 2030 Project, the global coverage of sonar-derived bathymetric data has increased from less than 9% [20] to 27.3% [22].

Despite the annual increase in the global coverage of MBES data, the persistent issue of data voids in the final products remains significant. This problem primarily stems from two levels: (1) At the data integration level, global bathymetric compilations typically rely on multisource data from diverse origins and collected under varying standards. The absence of unified survey line planning for these datasets makes data voids virtually inevitable (Figure 1a). (2) At the survey implementation level, practical maritime operations are constrained by economic costs and operational efficiency, sometimes opting for non-full-coverage surveying strategies, which create gaps between bathymetric swaths (Figure 1b).

This incompleteness of the data poses challenges to the accuracy and reliability of seabed topographic models.




Figure 1. Representation of gaps in MBES swaths on the world seabed topographic map, where (a) represents voids resulting from mosaicking multisource data, while (b) represents gaps arising from non-full-coverage surveys.

To address the incompleteness of elevation data and achieve a continuous representation of topography, researchers initially relied primarily on various interpolation algorithms to address void areas in DEMs, such as inverse distance weighting (IDW) [23], kriging [24,25], tension spline [26], and triangulated irregular network (TIN) [27] interpolation. However, most interpolation algorithms depend on mathematical smoothing and spatial autocorrelation assumptions and lack consideration of the complexity and randomness of seabed topography. Consequently, they are applicable only when repairing DEMs with small-sized data voids [28] or in flat terrain areas [29]. When confronted with large-scale data gaps or regions with complex topographic structures, traditional methods often struggle to reconstruct authentic topographic features [30].

In recent years, generative deep learning methods have provided a new paradigm for topographic data reconstruction, gradually overcoming the limitations of traditional interpolation techniques. Among them, two core approaches—generative adversarial networks (GANs) and diffusion models (DMs)—have demonstrated remarkable effectiveness in data reconstruction. GANs learn the data distribution of real terrain via adversarial training between a generator and a discriminator, enabling the generation of visually coherent topographic structures and achieving promising results in the local reconstruction of terrain data [31–33]. Qiu et al. specifically applied an enhanced DCGAN to avoid filling DEMs of rugged mountainous terrain. The results demonstrated that DCGAN effectively adapts to complex terrain variations and achieves good visual perception and reconstruction accuracy across void filling tasks of varying sizes [31].

However, GANs still face challenges in terms of their training stability and generation diversity [34], and their reconstruction results may sometimes produce artifacts due to mode collapse [35–37]. DMs, as advanced probabilistic generative models, have recently demonstrated stronger capabilities in detail recovery and training stability [34,38]. The “forward noise addition” and “reverse denoising” diffusion framework proposed by the Denoising Diffusion Probabilistic Model (DDPM) [39] has become the core paradigm for current image generation. By simulating diffusion and reverse-diffusion processes, it gradually reconstructs data from noise, addressing issues such as training instability or blurred outputs in traditional generative models. They not only generate high-quality images with rich details and reasonable structures but also allow precise control over the generated

content through flexible conditional guidance mechanisms, demonstrating significant potential when handling incomplete information and content generation [40–42]. Inspired by these findings, DMs are being introduced into the field of scientific data restoration [43,44]. Recent related studies indicate that DMs outperform traditional interpolation methods and GANs in DEM restoration tasks. The generated terrain not only exhibits advantages in terms of accuracy and visual coherence but also demonstrates greater stability in terms of the restoration results [30,45,46].

The underlying rationale for employing deep learning in repairing voids in DEMs stems from its ability to implicitly learn the complex patterns of terrain distribution from vast numbers of training samples. The technical approaches can be broadly categorized into two main types: One is the DEM restoration mode independent of prior data [45,47], which primarily relies on local contextual constraints. It extracts multiscale spatial features from known elevation data in the void neighborhood to capture terrain continuity within a local scope, thereby achieving unsupervised inference based on context. Diff-DEM [45] is a typical unsupervised DEM restoration method without prior constraints. Its primary approach involves excavating regular holes in land DEMs and utilizing DDPM for restoration, achieving state-of-the-art performance in both qualitative and quantitative evaluations. The other type is a restoration mode guided by prior data, which employs embedding auxiliary information such as remote sensing imagery [48], terrain shading [49], or terrain contour lines [31,50] into the generation process to constrain the model for more precise DEM void restoration. This learning paradigm, which combines local contextual constraints with global prior knowledge, enables the model to generate spatially coherent terrain surfaces that are geomorphologically plausible.

However, the current DEM void filling methods focus primarily on land terrain, where the restoration targets are typically voids caused by cloud cover or topographic shadows. In contrast, seabed topography faces more widespread and severe data gap issues due to the high costs of data acquisition and complex environmental interference factors, creating a more pressing demand for high-fidelity restoration. Furthermore, most current restoration models are built upon idealized degradation assumptions and are often trained and validated by using artificially “excavated” regular-shaped (e.g., rectangular) voids in complete data [32,45]. These conditions fundamentally differ in both morphology and origin from the irregular, randomly distributed complex gaps that actually exist in MBES data due to incomplete swath coverage. Consequently, the task of MBES bathymetric swath restoration presents new challenges regarding the generalization capability and robustness of models.

This paper addresses the special requirements of MBES bathymetric data restoration by innovatively proposing the MBES-DDPM method. The core innovation of this study lies in the first introduction of the DDPM into the field of multibeam bathymetric data restoration. Unlike traditional methods, the MBES-DDPM constructs training samples featuring data degradation that conforms to actual operational characteristics by simulating the MBES seabed surveying process. Concurrently, gravity anomaly models are utilized as prior knowledge, providing fine-grained topographic structural guidance for the diffusion process via a multisource data fusion mechanism. Finally, a systematic model evaluation framework is established. The experimental results demonstrate that this method significantly outperforms the existing baseline models in terms of both restoration accuracy and topographic fidelity, offering a novel technical approach for addressing data gap challenges in marine surveying and mapping.

2. Materials and Methods

This section first describes a random degradation simulation process for seabed DEMs that was designed based on the operational mode of the MBES to systematically generate training and test samples that replicated real-world data loss characteristics. Subsequently, a DDPM incorporating a terrain feature guidance mechanism was constructed. By adopting a conditional embedding approach, the surrounding valid bathymetric data and prior terrain features were integrated as constraints to achieve the intelligent restoration of data with gaps. Finally, a comprehensive evaluation system was established by using both quantitative and qualitative metrics to thoroughly assess the restoration results from the perspectives of numerical accuracy and terrain morphological fidelity. An overview of the proposed methodology is illustrated in Figure 2.




Figure 2. Schematic diagram of the MBES-DDPM.

2.1. Data and Study Area

In this study, the GEBCO_2025 dataset was selected as the foundational topographic model. GEBCO_2025 integrates multisource bathymetric data, including ship-based multi-beam and single-beam measurements, airborne LiDAR surveys, and bathymetric information derived from satellite gravity data inversion. This model provides globally continuous coverage of ocean depths and land elevations with a grid resolution of 15 arc-seconds. GEBCO also provides a corresponding type identifier (TID) grid, which labels the source data types for each grid cell. Supported by the Seabed 2030 Project, GEBCO is updated annually, ensuring high scientific rigor and data completeness.

In addition, this study employed the free-air anomaly model from SIO V32.1 as additional prior data to provide the model with general guidance on topographic trends. SIO V32.1 is a marine gravity anomaly model released by the Scripps Institution of Oceanography at the University of California, San Diego, with a spatial resolution of 1 arc-minute. Its free-air anomaly model is height corrected, preserving the signal variations that are primarily caused by the contrast in density between seawater and submarine rocks. Due to their mass surplus, positive topographic features such as seamounts and oceanic ridges generate significant positive free-air anomalies, while negative features such as trenches and troughs exhibit negative anomalies due to mass deficits. These characteristics provide clues for inferring topographic trends within void areas of the seabed.

As shown in Figure 3, two training areas and six independent test areas were selected globally. Through degradation simulation experiments, we simulated multibeam bathy-

metric swath gaps to create the low-quality data (LQ) used in this study, with GEBCO_2025 serving as the ground truth (GT). The test areas encompass typical seafloor topography features including trenches, seamounts, abyssal plains, and mid-ocean ridges. Among them, Tests 1–3 represent a random degradation pattern, simulating data voids resulting from multisource data integration. Tests 4–6 represent a rule-based degradation pattern, simulating inter-swath gaps formed by parallel survey lines in a planned survey. Furthermore, the hydrological analysis tools in ArcGIS (v10.5) were used to extract ridge and valley lines from the V32.1 free-air anomaly data, serving as prior knowledge for this study. Further details regarding the test area are summarized in Table 1.




Figure 3. Study area map.

Table 1. Test area information.

Test Set Name	Test 1	Test 2	Test 3	Test 4	Test 5	Test 6
Degenerative pattern	random degradation	random degradation	random degradation	rule-based degradation	rule-based degradation	rule-based degradation
Survey line azimuth	random	random	random	0°	45°	−45°
Percentage of void space	29.90%	29.64%	34.58%	50.00%	47.28%	47.28%

2.2. Simulation of the Seabed DEM Degradation Process

Unlike the regular and concentrated data gaps found in land DEMs, data gaps in seabed topography primarily result from the swath-scanning measurement characteristics of the MBES. These gaps typically manifest as unclosed strip-shaped voids between bathymetry swaths. As illustrated in the two scenarios depicted in Figure 1, such gaps exhibit inherent randomness while also displaying a degree of regularity. We categorize this degradation process into two types: random degradation and rule-based degradation. The primary distinction lies in whether the survey lines are randomly distributed.

To facilitate the mathematical modeling of the simulation process, we simplified the actual measurement process of the MBES into two aspects: (1) The MBES survey lines are modeled by using linear equations; and (2) the coverage width of the MBES swath is a predetermined fixed value that does not vary with the terrain. Based on the above assumptions, point $A(x_0, y_0)$ on the survey line and its azimuth angle $\theta \in [-\pi/2, \pi/2]$ are first determined within the survey area. The general equation of the survey line can then be expressed as follows:

$$Ax + By + C = 0, \quad (1)$$

where $A = -k$, $B = -1$, $C = -k \times x_0 + y_0$, and $k = \tan(\theta)$. For the random degradation simulation, the coordinates of point A and the azimuth angle θ are obtained through random sampling. For rule-based degradation simulation, the values of point A 's coordinates and azimuth angle θ are predetermined through prior calculation. Once the equation of the survey line is established, the scanning range of the seabed is determined based on the coverage width of the MBES. To obtain this information, it is necessary to traverse the grid points of the seabed topography and compute the planar distance d from each point to the survey line, expressed as follows:

$$d = \frac{|A \times x_i + B \times y_i + C|}{\sqrt{A^2 + B^2}}, \quad (2)$$

where x_i and y_i are the coordinates of any point within the survey area. If $d < w/2$, the point lies within the MBES swath coverage and data exist for that location; otherwise, the data for that point are void. As shown in Figure 4a, we employ a binary mask to record the calculation results, where pixels in the data area are valued as 1 and pixels in the data void area are valued as 0. The simulated LQ for seabed topography degradation is then as follows:

$$LQ = GT \times mask, \quad (3)$$




Figure 4. Schematic diagram of the seabed topography degradation simulation. (a) Simulation of MBES bathymetry and mask generation; (b) simulation results of random degradation, where the bathymetry swaths for the LQ are based on random sampling of survey lines; and (c) simulation results of rule-based degradation, where the survey line azimuths in the LQ are fixed at 45° .

2.3. Modeling the Diffusion Process in MBES Swath Restoration

2.3.1. Modeling of Bidirectional Diffusion Processes

In this study, we adopted the DDPM as the core mathematical framework to restore MBES swaths. The data restoration process is formulated as a bidirectional Markov process. The forward diffusion process progressively adds noise through a series of sequential steps, while the reverse diffusion process gradually reconstructs the target topographic data from the noise distribution through iterative denoising operations. Specifically, at each denoising step, the terrain feature lines extracted from the V32.1 free-air model are incorporated as conditional guidance to ensure that the data restoration aligns with the morphological consistency and continuity of the seabed topography.

The forward diffusion process gradually transforms the original complete terrain data (x_0) into an isotropic Gaussian distribution (x_T) by adding noise. This transformation follows a predefined noise scheduling strategy, enabling the model to complete the full transition from ordered terrain data to disordered noise. Each step in which random noise is added is independent. Thus, the forward diffusion process is defined as a Markov chain:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t|\sqrt{1-\beta_t}x_{t-1},\beta_t I), t = 1, 2, \dots, T, \quad (4)$$

where β_t is a predefined noise scheduling parameter. Equation (4) primarily achieves two objectives: on the one hand, it scales x_{t-1} with $\sqrt{1-\beta_t}$; on the other hand, it controls the intensity of the added Gaussian noise with β_t . By utilizing the properties of the Gaussian distribution, the distribution of x_t at any given time can be derived as follows:

$$q(x_t|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}) = \mathcal{N}(x_t|\sqrt{\bar{\alpha}_t}x_0, (1-\bar{\alpha}_t)I) \quad t = 1, 2, \dots, T, \quad (5)$$

where $\bar{\alpha}_t = \prod_{t=1}^T (1-\beta_t)$. Reparameterizing Equation (5) yields:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon_t, \quad (6)$$

where $\epsilon_t \sim \mathcal{N}(0, I)$ represents noise randomly sampled from a standard Gaussian distribution. Equation (6) indicates that the transition from x_0 to any time point x_t can be achieved through a single noise addition.

The reverse diffusion process begins at the terminal point x_T of the forward diffusion process. Its core objective is to progressively reconstruct multibeam bathymetric data from pure noise by learning a parameterized transition function. This process is fundamentally an iterative probabilistic reconstruction problem under conditional constraints. As demonstrated by Sohl-Dickstein et al. [38], the reverse diffusion process still adheres to Gaussian distribution characteristics. Therefore, it can be expressed as follows:

$$P_\theta(x_{t-1}|x_t, LQ, C) = \mathcal{N}(x_{t-1}|\tilde{\mu}_\theta(x_t, LQ, C, t), \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t I), \quad (7)$$

where C denotes the terrain feature line extracted from the V32.1 free-air model, which is used to constrain the terrain trend generated in the data void area. $\tilde{\mu}_\theta(\bullet)$ represents the estimated mean of this conditional distribution, and its specific form is expressed as follows:

$$\tilde{\mu}_\theta(x_t, LQ, C) = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}})\epsilon_\theta(x_t, LQ, C, t), \quad (8)$$

where $\alpha_t = 1 - \beta_t$, and $\epsilon_\theta(\bullet)$ denotes the noise component predicted by the network. The mean function of the reverse diffusion in Equation (8) is solved via neural network-based

noise prediction. The model is trained by optimizing the following weighted variational lower bound, with the loss function being defined as follows:

$$L = E_{x_t, LQ, C, t} \left[\left\| \epsilon_t - \epsilon_\theta \left(\sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon_t, LQ, C, t \right) \right\|^2 \right], \quad (9)$$

where ϵ_t represents the noise added during the forward process according to Equation (6).

2.3.2. Eliminating Elevation Jumps Through Cycling Sampling

During the model inference stage, a key methodological challenge lies in addressing the integration between known data and generated data. Unlike common image generation tasks, the central objective of MBES swath restoration is to utilize measured valid bathymetric data to predict and fill values in void areas. The final requirement involves seamlessly integrating the model's restoration results with the original known data, to form a complete and continuous seafloor topographic surface. However, the restoration results generated by the diffusion model via multistep iterative sampling often exhibit discrepancies with the original measured values at the boundaries of known regions. If a simple direct replacement is used to stitch the restored areas with the known data, elevation jumps occur between these two types of data, adversely affecting accuracy. Related studies have shown that repeated predictions via neural networks can mitigate or even eliminate such elevation discontinuities within the patch [30,51,52].

As shown in Figure 5, for the sampling process at each time step, we introduce a multistep iterative cycling sampling strategy. First, the prediction of x_{t-1} from x_t is obtained through the conventional reverse diffusion process. By multiplying it with the binary mask image (1—mask) that records the positions of measured points, the unknown portion of x_{t-1} is extracted. Simultaneously, forward diffusion is applied to the LQ up to time step x_{t-1} , and by multiplying with the mask, the known portion of x_{t-1} is obtained. These two parts are then summed to yield the single-step sampling result for x_{t-1} . However, since the known and unknown parts of x_{t-1} are derived through different computational methods, directly concatenating them may leave visible seams. To address this issue, single-step forward diffusion is performed on x_{t-1} , and noise is added with Equation (4) to obtain x_t . The cyclic process of “ $x_t \rightarrow x_{t-1} \rightarrow x_t$ ” is repeated N times until the elevation discontinuity between the two parts data is eliminated. The pseudocode for the cycling sampling strategy is provided in Algorithm 1.

Algorithm 1: Cycling sampling

```

1:  $x_T \sim \mathcal{N}(0, I)$ 
2: for  $t = T, \dots, 1$  do
3:   for  $k = 1, \dots, N$  do
4:      $P_\theta(x_{t-1}|x_t, LQ, C) \rightarrow x_{t-1} \odot (1 - \text{mask}) \rightarrow x_{t-1}^{\text{unknow}}$ 
5:      $q(x_{t-1}|x_0) \rightarrow x_{t-1} \odot \text{mask} \rightarrow x_{t-1}^{\text{know}}$ 
6:      $x_{t-1} = x_{t-1}^{\text{know}} + x_{t-1}^{\text{unknow}}$ 
7:     if  $k < N$ 
8:        $q(x_t|x_{t-1}) \rightarrow x_t$ 
9:     end if
10:  end for
11: end for
12: return  $x_0$ 
```




Figure 5. MBES-DDPM cycling sampling strategy.

2.4. Noise Prediction Network of MBES-DDPM

The reverse diffusion process introduced in Section 2.3.1 is achieved by predicting the noise at each timestep. For this purpose, we specifically designed a noise prediction network based on the U-Net architecture [53]. As shown in Figure 6, the overall network is structured with five layers. First, a convolutional layer increases the channel count to 32. Then, through progressive downsampling, the original input image size of 128×128 is reduced to 8×8 , while the channel count is gradually expanded to 256. Subsequently, via progressive upsampling, the image size is restored to 128×128 , and the channel count is progressively reduced to 1, ultimately outputting the predicted noise. As described in Equation (7), this model incorporates terrain feature lines C as prior knowledge to guide the training process. Accordingly, we designed a dedicated conditional encoder for C, which injects the encoded conditional information at the bottom of the model. This information is fused with the deep features extracted by the model to jointly predict the noise.

The integration of transformer with diffusion models, which can leverage the self-attention mechanism to enhance the performance of diffusion models, holds significant research and application potential. For instance, the notable diffusion transformer (DiT) model has surpassed traditional stable diffusion models [54]. However, the computational complexity of transformer is substantial, and further increases when it is applied to the image domain, resulting in high training costs [55]. CNNs excel at capturing local image features, but their limited receptive field makes it challenging to model complex relationships between distant components within an image. In contrast, Transformers possess a strong capacity for global modeling. However, their application in the image domain leads to an explosion in computational complexity, and their handling of local features may not be as fine-grained as that of CNNs. Therefore, this hybrid CNN-Transformer architecture design has been extensively utilized and validated within the noise prediction process of diffusion models [56–58]. In our noise reduction network, the transformer layer incorporates multi-head self-attention layers to extract global features, while densely con-

nected residual blocks [36,59] serve as the foundational modules with which the CNN can capture local features. These two components work in tandem, not only consuming fewer computational resources but also leveraging their respective strengths during multiscale feature extraction.




Figure 6. Noise reduction network and its submodule structure diagram.

3. Results

3.1. Experimental Setting

The experiments were conducted on a computing platform equipped with an NVIDIA RTX 4090 GPU, utilizing the PaddlePaddle 3.0 deep learning framework. The MBES-DDPM adopts the DDPM framework, with a total diffusion timestep $T = 500$. The noise schedule parameter β_t is linearly interpolated between 0.0001 and 0.02, and the number of cycling sampling iterations $N = 5$. During training, the Adam optimizer was employed with a learning rate of 1×10^{-5} and a batch size of 8. A total of 1000 epochs were completed to ensure sufficient model convergence. To comprehensively validate the model's effectiveness, this study selected three interpolation algorithms—IDW, Kriging, and Spline—along with two generative deep learning models: DCGAN [60] and Diff-DEM [45]. All methods were tested under identical datasets and training strategies to ensure fair comparison.

Furthermore, this study also designed both qualitative and quantitative evaluation systems. First, the restoration results were displayed and analyzed through multi-level comparison. Subsequently, the results were evaluated using relevant metrics across two dimensions: accuracy and similarity. Accuracy was assessed using the root mean square error (RMSE) and mean absolute error (MAE), while similarity was evaluated using the peak signal-to-noise ratio (PSNR) and the structural similarity index measure (SSIM).

3.2. MBES Swath Restoration Results

To systematically evaluate the generalizability and robustness of the proposed method across different real-world data acquisition scenarios, in this study, two sets of experiments

were designed and conducted. These experiments were based on distinct data degradation mechanisms that simulated two typical types of data gap issues encountered in multibeam bathymetry operations: (1) random degradation restoration experiments to simulate incomplete terrain with randomly distributed gaps in size and location resulting from multisource data stitching without survey line planning (the restoration results are shown in Figure 7), and (2) rule-based degradation restoration experiments to simulate the strip-like missing data characterized by uniformly spaced gaps resulting from systematic survey line intervals (the restoration results are shown in Figure 8). These two experimental scenarios address the challenges associated with random and regular degradation patterns in seabed topography restoration, with the goal of comprehensively validating the overall performance of the model in terms of addressing complex real-world problems. In addition to the comparative presentation of restoration effects across large regions, detailed local restoration results are provided.




Figure 7. Random degradation seabed topography restoration results.




Figure 8. Rule-based degradation of seabed topography restoration results.

3.2.1. Overall Restoration Results

To visually evaluate the restoration performance of the different methods, a qualitative comparative analysis was conducted on randomly degraded seabed restoration results across three representative regions (Areas 1–3) with distinct topographic characteristics. As shown in Figure 7, the mid-ocean ridge terrain in Area2 is characterized by complex folds, with a large void zone present in its upper left corner. In this region, IDW and kriging introduced noticeable topographic distortions and generated apparent interpolation artifacts, whereas the spline method produced overly intricate and unrealistic textures within the void area. These results highlight the inherent limitations of traditional interpolation methods when addressing extensive data gaps in complex terrains. Area 3 is a highly rugged submarine mountainous region, where traditional interpolation methods resulted in unnatural “terrace-like” artifacts, significantly compromising topographic details. Although DCGAN restored some details, it exhibited pronounced elevation discontinuities along the boundaries of the original data. In contrast, the predictions from Diff-DEM contained numerous unreasonable anomalous elevation protrusions. Overall, across all of the complex scenarios that were tested, the proposed MBES-DDPM demonstrated superior performance in terms of visual fidelity and topographic detail restoration, and its results maintain high consistency with those of the GT.

In the rule-based degradation restoration experiment illustrated in Figure 8, we further evaluated the performance of the different methods across three seabed areas with representative survey line azimuths (Area 4: 0° ; Area 5: 45° ; and Area 6: -45°). The experimental results indicate that the restoration outputs of traditional interpolation methods (IDW, kriging, spline) all exhibit systematic stripe-like artifacts that are highly aligned with the original survey line directions. In regions with pronounced topographic relief, such as the ridge edges in Area5, these artifacts were particularly prominent, leading to noticeable sawtooth-like distortions in the terrain contours. Although the DCGAN method was able to reconstruct the basic macro-morphology of the terrain, visible discontinuities persisted at the junctions between known data and restored regions. The Diff-DEM method continued to demonstrate issues with overestimated elevations, producing scattered protrusion artifacts. In contrast, our proposed MBES-DDPM method effectively suppressed directional artifacts and outliers across all orientation-based gap tests. The generated terrains appeared visually continuous and natural, without significant topographic structural distortions, demonstrating the proposed method's robust ability to restore rule-based degradation patterns.

3.2.2. Local Detail Restoration Effect

To further evaluate the effectiveness of different methods in terms of restoring terrain details, we present a local magnification of the feature-rich area (Figure 9) and extracted topographic profile lines (Figure 10). Figure 9 displays a local enlargement of Area3 from Figure 7, where the LQ exhibits random degradation with significant data gaps in the central area. The GT data indicate substantial terrain variation in this area, making it a representative case. The IDW, kriging, and spline interpolation methods produced noticeable “step-like” elevation discontinuities in the void areas. In the DCGAN restoration results, discontinuous elevation transitions are observed between the predicted and measured regions, with the restored terrain in the gap area generally lower than that in the GT. The Diff-DEM restoration avoids the elevation jumps seen in interpolation methods and DCGAN, presenting more natural terrain changes and realistic visual effects. However, most terrain details within the restored area are inaccurate and inconsistent with the GT. In contrast, our proposed MBES-DDPM method yields more authentic restoration results, accurately recovering the basic topographic trends compared to the GT, although some high-frequency details are not fully restored. Overall, the restoration performance of the MBES-DDPM surpasses that of the other baseline models.




Figure 9. Local detail comparison in Area 3.




Figure 10. Profiles of the seabed topography restoration results, the arrows in LQ indicate the position and direction of the profiles.

As shown in Figure 10, Profiles 1–3 traverse a narrow ridge, whereas Profiles 4–6 are located in a mid-ocean ridge region. The analysis reveals that the profile curves generated by the IDW and kriging methods exhibit sharp fluctuations, which correspond precisely to the step-like interpolation artifacts observed in Figures 7 and 8, confirming the inherent limitations of such methods in restoring continuous terrain. The Diff-DEM predictions exhibit systematic bias, significantly underestimating true elevations across multiple profiles (e.g., Profiles 1, 2, 4, 5) while overestimating them excessively in Profile 6. Overall, it shows a tendency to deviate from the actual terrain. In contrast, the profile curves of MBES-DDPM closely align with the GT in most cases, particularly in Profiles 4–6 under rule-based degradation scenarios, where the alignment with GT is even more pronounced. These terrain profiles further confirm that MBES-DDPM outperforms the comparative methods in both restoring terrain details and maintaining elevation accuracy, demonstrating significant advantages especially when handling rule-based degradation patterns.

3.3. Accuracy Evaluation of MBES Swath Restoration

Figure 11 displays the residual error distributions of the restoration results in two representative regions. The results show that in Area 1, which features a prominent mountainous ridge with significant topographic variation, the residuals of all of the methods are generally greater near the ridge. In contrast, over relatively flat areas, even when there are extensive data gaps, the residuals remain effectively controlled, indicating that terrain complexity is a key factor influencing restoration accuracy. Specifically, traditional interpolation methods such as IDW, kriging, and spline exhibit similar residual distribution patterns, which are characterized by fine-grained, high-frequency errors, which reflects their limitations in recovering fine-scale terrain structures and continuous variations. Diff-DEM results in widespread negative residuals in both regions, which is consistent with the presence of anomalous elevation protrusions in its predictions. In comparison, the proposed MBES-DDPM method demonstrates the most spatially uniform residual distribution

with the smallest overall magnitude. This observation visually confirms the significant advantage of MBES-DDPM in enhancing both the overall accuracy and stability of the restoration results.




Figure 11. Residual distribution map of the seabed topography restoration results. Negative residuals indicate GT values lower than the restoration results, whereas positive residuals indicate GT values higher than the restoration results.

Tables 2 and 3 present the error and similarity quantification metrics of each method across the six test regions, respectively. The quantitative results reveal a consistent phenomenon: for rule-based degraded data (Tests 4–6), the errors for all methods are significantly lower, and the similarity is notably higher compared to those for random degradation data (Tests 1–3), confirming that random degradation data pose greater challenges to restoration algorithms. Specifically, IDW and kriging exhibit comparable error levels across all regions, demonstrating stable performance. The spline method yields higher errors in random degradation scenarios (Tests 1–3) but achieves lower errors in rule-based degradation scenarios (Tests 4–6), indicating its varying sensitivity to different data gap patterns. Across all test regions and every evaluation metric, the proposed MBES-DDPM method consistently achieves optimal performance. Statistically, compared with the second-best baseline model, the MBES-DDPM reduces the RMSE of the restoration results by an average of at least 34.21% and improves the PSNR by more than 3.71 dB on average. This significant quantitative improvement confirms that the proposed method offers higher accuracy and robustness when addressing different types of data gaps.

Table 2. Accuracy statistics for seabed topography restoration results (the best results are highlighted in bold blue font).

Indicator	Methods	Test 1	Test 2	Test 3	Test 4	Test 5	Test 6
RMSE (m)	IDW	119.73	137.21	91.38	100.40	94.14	68.91
	Kriging	111.47	137.15	91.40	99.53	94.16	68.56
	Spline	137.36	149.08	114.86	88.83	66.16	46.14
	DCGAN	144.37	154.79	102.33	147.69	138.76	110.61
	Diff-DEM	232.58	217.70	151.43	160.27	168.49	168.89
	MBES-DDPM	89.35	111.52	68.37	51.63	43.21	32.20
MAE (m)	IDW	32.31	40.68	31.35	42.90	28.22	23.54
	Kriging	28.72	40.90	31.54	43.56	28.46	23.58
	spline	32.95	35.27	35.11	36.42	18.88	14.87
	DCGAN	45.51	42.47	36.39	58.35	60.30	42.62
	Diff-DEM	85.01	70.97	60.43	76.36	73.68	72.55
	MBES-DDPM	26.49	28.03	24.23	19.76	13.39	11.38

Table 3. Statistical table of the similarity in the seabed topography restoration results (the best results are highlighted in bold blue font).

Indicator	Methods	Test 1	Test 2	Test 3	Test 4	Test 5	Test 6
PSNR (dB)	IDW	35.08	33.16	35.00	34.03	37.33	37.12
	Kriging	35.70	33.17	35.00	34.11	37.33	37.17
	Spline	33.89	32.44	33.01	35.10	40.39	40.61
	DCGAN	33.46	32.12	34.02	30.68	33.96	33.01
	Diff-DEM	29.32	29.15	30.61	29.97	32.27	29.34
	MBES-DDPM	37.62	34.97	37.52	39.81	44.09	43.73
SSIM	IDW	0.98	0.98	0.98	0.96	0.98	0.98
	Kriging	0.98	0.98	0.98	0.96	0.98	0.98
	spline	0.97	0.97	0.96	0.96	0.99	0.99
	DCGAN	0.97	0.97	0.97	0.94	0.97	0.96
	Diff-DEM	0.96	0.96	0.96	0.94	0.96	0.94
	MBES-DDPM	0.99	0.99	0.99	0.99	0.99	0.99

3.4. Accuracy Evaluation of the Slope in the MBES Swath Restoration Results

To further evaluate the morphological fidelity of the restored terrain, in this study, the residual errors between the slopes of the restored results and the GT slopes were analyzed, with results shown in Figure 12 and Table 4. The quantitative metrics in Table 4 show that the IDW and kriging methods produce relatively high slope residuals (in terms of both the RMSE and the MAE) across all test regions. Notably, while the slope residuals of the DCGAN and Diff-DEM are lower than those of the traditional interpolation methods, their corresponding absolute elevation errors (Table 2) are higher. This seemingly contradictory phenomenon indicates that although such generative methods can produce rich high-frequency textures, these details are not necessarily accurate. In contrast, traditional interpolation methods produce overly smooth results that approximate the average elevation but completely lose the critical slope information reflecting terrain structure. Therefore, generative methods demonstrate a relative advantage in capturing slope metrics that characterize terrain morphology. Among all of the compared methods, the proposed MBES-DDPM in this paper achieved the best results for the majority of the tested metrics. This finding proves that the MBES-DDPM can not only achieve high-precision elevation reconstruction but also most accurately restore the first-order derivative features of the

terrain, realizing a higher-order restoration of the real seabed topography in terms of morphological structure.




Figure 12. Spatial distribution of the slope in the MBES swath restoration results.

Table 4. Statistical table of the slope accuracy for the MBES swath restoration results (the best results are highlighted in bold blue font).

Indicator	Methods	Test 1	Test 2	Test 3	Test 4	Test 5	Test 6
RMSE (m)	IDW	12.71	10.66	13.93	13.37	16.10	15.29
	Kriging	11.25	10.58	13.91	13.63	16.14	15.29
	Spline	9.13	7.19	10.43	8.66	13.38	11.67
	DCGAN	10.43	6.36	10.79	8.19	14.96	12.15
	Diff-DEM	11.42	7.58	13.32	10.80	15.27	14.13
	MBES-DDPM	8.97	5.60	9.94	6.65	10.89	9.55
MAE (m)	IDW	4.31	3.15	5.29	5.28	7.24	6.86
	Kriging	3.69	3.14	5.31	5.62	7.31	6.91
	spline	2.67	1.66	3.44	3.47	5.61	4.84
	DCGAN	3.43	1.78	3.90	3.18	6.74	5.20
	Diff-DEM	3.91	2.34	5.23	4.57	6.96	6.45
	MBES-DDPM	2.83	1.53	3.59	2.29	4.39	3.77

3.5. Analysis of the Gap Restoration Effects for Different Sizes

To quantitatively evaluate the impact of the data gap size on restoration effectiveness, we systematically simulated four different swath gap widths (7.5 km, 12 km, 16.5 km and

21 km) in the Test 5 area and performed restoration by using the MBES-DDPM. A comparison of the simulated low-quality data (LQ), restoration results, and residual distributions for the local region (Area5) is shown in Figure 13, and the corresponding quantitative accuracy assessment results are summarized in Table 5.




Figure 13. Restoration results of the MBES-DDPM on gaps of varying sizes. The numbers below the LQ correspond to the widths of the gaps between adjacent swaths, specifically 7.5 km, 12 km, 16.5 km, 21 km, 25.5 km and 30 km.

Table 5. Accuracy of the restoration results for gaps of different sizes.

Gap Widths (km)	RMSE (m)	MAE (m)
7.5	22.76	5.81
12	43.21	13.39
16.5	71.54	25.01
21	101.95	39.19
25.5	124.28	51.76
30	159.14	70.47

The experimental results indicate a clear negative correlation between the restoration accuracy and gap width. When the gap width is 7.5 km, the original data remain relatively intact, and the narrow gaps result in overall smaller residual errors in the model's restoration. As the gap width increases, the range of residual errors in the restoration results expands significantly, and the ability of the model to recover topographic details tends to decrease. This visual pattern was also rigorously validated with quantitative metrics. As shown in Table 5, as the gap width increases from 7.5 km to 30 km, both the RMSE and the MAE of the restoration results increase systematically. This finding demonstrates that, under the same topographic conditions, larger geographic data gaps correspond to increased challenges and uncertainties in the restoration task.

Furthermore, a detailed examination of the spatial distribution of the residuals (Figure 13) reveals that larger residual values are not uniformly distributed but are notably concentrated in regions with significant topographic variation, such as trenches and their flanking ridges. In contrast, over relatively flat seabed areas, residuals remain at low levels even when the gaps are wider. These observations clearly indicate that the restoration accuracy is constrained by both the gap width and local terrain complexity. Areas with complex topography, which are characterized by drastic elevation changes and rich structural features, impose greater demands on the model's contextual inference and detail generation capabilities, thereby becoming the primary source of restoration errors.

Furthermore, quantitative experimental results further reveal that the decline in restoration accuracy with increasing gap size is not linear but exhibits a distinct non-linear, accelerating downward trend. Analysis of the data in Table 5 shows that as the gap size increases in fixed increments (4.5 km), the corresponding degradation in accuracy increases significantly. For example, when the gap width expanded from 7.5 km to 12 km, the RMSE and MAE of the restoration results increased by 20.45 m and 7.58 m, respectively. However, starting from a larger base gap of 25.5 km to 30 km, the increases in RMSE and MAE reached 34.85 m and 18.71 m, respectively. This phenomenon indicates that larger gap sizes are associated with higher risks and greater uncertainty regarding accuracy loss from further expansion.

The underlying mechanism of this phenomenon is that restoring large gaps fundamentally and critically depends on the model's ability to extract and infer global topographic features from the limited surrounding known areas. When the gap is relatively small, the model has sufficient contextual information to constrain the generation process. As gap size increases, the available effective contextual information decreases drastically. The model is then forced to perform feature inference and synthesis over longer distances with higher uncertainty. This exponentially intensifies the demands on the model's structural priors and feature extraction capabilities, leading to accelerated error reduction.

3.6. Ablation Study

3.6.1. The Effectiveness of Cycling Sampling

To validate the effectiveness of the proposed cycling sampling mechanism in enhancing the internal consistency of the restored results, we designed a dedicated analytical experiment. This mechanism aims to mitigate local elevation discontinuities caused by the direct embedding of LQ into the predicted output. As shown in Figure 14, when the number of cycling sampling steps $N = 1$, distinct boundaries from the embedded LQ is clearly visible in the restoration result, forming discontinuous artifacts. As N increases, these discontinuous pseudo-artifacts are progressively eliminated through iteration. At $N = 5$, the restored topography exhibits a smooth and natural transition, demonstrating high conformity with the reference ground truth (GT). Quantitative analysis further supports this observation: as indicated in Table 6, the RMSE of the restoration result is 131.26 m at $N = 1$, which decreases significantly to 89.35 m at $N = 5$, representing a reduction of 31.92%. This result demonstrates that, even without introducing any additional information, optimizing the sampling strategy and increasing the number of cycling sampling can effectively promote the seamless integration of topography and elevation between restored areas and LQ, thereby significantly enhancing the overall consistency and geometric stability of the results.




Figure 14. Seabed topography restored using different numbers of cycling sampling.

Table 6. Accuracy of restoration results at different numbers of sampling cycles.

N	RMSE (m)	MAE (m)
1	131.26	44.33
2	102.39	32.05
3	94.00	28.33
4	89.89	26.57
5	89.35	26.49

3.6.2. The Effectiveness of Feature Line

To validate the effectiveness of using topographic feature lines derived from marine gravity anomaly data as prior constraints, we designed a set of ablation experiments. In the comparative model without feature line constraints, we removed the conditional encoder from the noise prediction network and adjusted the input-output structure of the relevant U-Net layers accordingly. This model was then trained using the identical dataset and training strategy as the complete MBES-DDPM. A visual comparison of the restoration results is shown in Figure 15, with quantitative metrics summarized in Table 7. The experimental results indicate that with the introduction of feature line constraints, the restoration results show significantly higher agreement with the GT, exhibiting lower overall error and a more spatially uniform error distribution. In contrast, the unconstrained model produced results with greater deviation and more scattered errors. Concurrently, the results in Table 7 confirm this pattern: after applying the topographic feature line constraints to the model, the RMSE of the restoration results decreased by 37.21%. This clearly demonstrates the important role of external prior topographic knowledge in enhancing restoration accuracy.


**Figure 15.** Comparison of restoration effects with and without terrain feature lines.

Table 7. Accuracy of restoration results with and without feature line guidance.

	RMSE (m)	MAE (m)
with feature line	68.37	24.23
without feature line	108.88	36.07

The restoration of gaps in MBES bathymetry swaths is inherently an ill-posed problem, with solutions possessing a high degree of uncertainty. The inherent diversity in diffusion model generation can produce seemingly plausible but “deceptive” details that violate macro-scale terrain structures. The topographic feature lines introduced in this study, serving as prior knowledge describing regional topographic trends, provide macro-structural constraints for the diffusion process. They effectively restrict and guide the model’s reasoning and generation, forcing the restoration process to adhere to true topographic structures. This suppresses the generation of implausible details and ensures that the restoration results remain consistent with the actual seabed topography in both overall structure and local features. Therefore, the constraint mechanism of topographic feature lines fundamentally enhances the model’s capability to address this type of ill-posed problem.

4. Discussion

4.1. Analysis of the Restoration Effects for Different Data Degradation Patterns

In this study, the incompleteness of the MBES data is categorized into two patterns, random degradation and rule-based degradation, and a systematic evaluation framework was thus established accordingly. Both the quantitative and qualitative experimental results collectively reveal an important finding: Although our proposed MBES-DDPM achieves the best restoration accuracy on both types of degraded data, the margin of its performance improvement over the second-best model differs substantially between the two degradation modes.

Specifically, in the random degradation scenario (Test 1–3, average data void ratio ~30%), compared with the second-best baseline method, the MBES-DDPM reduced the average RMSE by 21.24%. In contrast, in the rule-based degradation scenario (Tests 4–6, with an average data void ratio as high as ~50%), the average reduction in the RMSE reached 35.59% compared with the second-best baseline method. These results indicate that the complexity of the data gap pattern may pose a greater challenge to restoration algorithms than the data void ratio alone. Although the overall void ratio is lower in randomly degraded data, its gaps vary in size and exhibit highly random spatial distributions and shapes; thus, this pattern is particularly prone to forming isolated large-area data voids. As analyzed in Section 3.5, the restoration accuracy is generally lower for larger gaps. This random degradation pattern disrupts the spatial autocorrelation of topographic data, increasing the difficulty for models in terms of accurately inferring the terrain structure in unknown regions.

In contrast, although rule-based degradation data have a higher void ratio, their gaps exhibit a uniform, directional strip-like distribution with predictable spatial patterns. Such systematic gaps partially preserve the continuity of terrain features along specific directions, making it easier for models—including certain traditional interpolation methods—to learn and perform restoration. Consequently, the more pronounced performance advantage demonstrated by the MBES-DDPM on rule-based degradation data, precisely demonstrates its robust ability to capture and exploit the latent regular structures within the data.

In summary, this study identifies and analyzes key characteristics of MBES data restoration tasks: The difficulty of restoration depends not only on the total amount of missing data but also, more critically, on the randomness and structural nature of the

missing patterns. The MBES-DDPM remains stable even in the more challenging random degradation scenario, further validating its enhanced applicability and robustness when addressing complex, unstructured real-world marine surveying data gaps.

4.2. Analysis of the Restoration Effect of MBES Measured Data Sources

To further validate the model's generalization capability in real-world scenarios, we conducted swath restoration experiments on actual MBES survey data. Specifically, we utilized the type identifier (TID) grid of GEBCO_2025 to extract a section of seafloor topography from GEBCO_2025 that is entirely based on MBES measurement sources, as shown in the GT in Figure 16. This region features a seamount with complex, rugged terrain and a water depth range of 12 to 5797 m. We applied a random degradation simulation to this data, resulting in an LQ version where approximately 35% of the area consists of data voids.




Figure 16. The repair effect of MBES actual measured data.

Baseline models and the proposed MBES-DDPM were employed to restore the LQ data. The results are shown in Figure 16, with the quantitative accuracy summarized in Table 8. Analysis reveals that traditional interpolation methods (IDW, Kriging, Spline) continue to exhibit significant limitations on the real measured data. Their restoration results contain systematic step-like elevation artifacts and generate a substantial amount of non-physical, distorted topographic textures. The results from deep learning methods, DCGAN and Diff-DEM, show some improvement in visual coherence. However, the topographic details they generate deviate significantly from the GT. As indicated in Table 8, their accuracy metrics are even lower than some interpolation methods, highlighting the risk of these methods producing “plausible but erroneous” details when lacking strong constraints. In contrast, the proposed MBES-DDPM maintains high visual consistency with the GT and achieves the best quantitative accuracy among all baseline models, reducing RMSE by approximately 27.95% compared to the second-best model. This result confirms that MBES-DDPM exhibits no strong dependency on the data source, demonstrating excellent potential for practical application.

Table 8. Accuracy of MBES measured data restoration.

Methods	RMSE (m)	MAE (m)
IDW	84.23	29.96
Kriging	76.54	26.34
Spline	96.27	29.12
DCGAN	118.80	42.46
Diff-DEM	132.42	52.89
MBES-DDPM	55.15	20.48

The strong generalization capability demonstrated by MBES-DDPM fundamentally stems from two key technical strengths: first, a context feature extraction and modeling mechanism that enables the model to effectively learn from and infer the topographic structure of missing regions based on the complex spatial relationships present in known areas; and second, the prior knowledge constraint mechanism, which provides stable macro-topographic guidance for the diffusion generation process by incorporating external knowledge such as gravity anomalies. These dual capabilities work in concert, allowing the model to learn the intrinsic distribution patterns of seabed topography from the training data, rather than simply memorizing or fitting the superficial characteristics of specific data sources.

4.3. Limitations and Future Prospects

The model proposed in this study performs well when handling typical MBES data gaps, yet its performance remains constrained by the inherent limitations of the current deep learning restoration paradigms. The primary limitation lies in the fact that, to suit the model input requirements and computational efficiency, large-area seabed topography must be cropped into fixed-size image patches (e.g., 128×128 pixels) for independent processing. Due to the randomness of actual survey line distributions, the density of valid data varies significantly across different patches. Some patches may contain very few valid measurements, resulting in a high proportion of “open-type” gaps (i.e., gap regions that contact or approach the patch boundaries), as illustrated by the extreme case shown in Figure 17. Under such conditions, the model suffers from a notable decline in restoration performance due to the lack of necessary surrounding topographic context. The root cause of this issue is that the local processing scheme based on fixed-size patches possesses a limited receptive field and is unable to capture the spatial correlations and constraints required for restoration from a broader geographic scope.

To overcome these limitations mentioned, future research could explore next-generation restoration architectures that go beyond fixed local receptive fields. A highly promising direction is to draw inspiration from the Latent Diffusion Model (LDM) framework in the natural image processing domain [40,54]. This approach adopts a multi-stage generation pipeline: First, a large-scale seabed topography map is compressed into a low-dimensional latent space via an encoder; diffusion-based restoration is then performed within this latent space; and finally, a decoder reconstructs the restored latent features back to the original geographic spatial scale. Such a paradigm allows the model to capture and leverage global topographic structures and contextual relationships over broader spatial extents through the compressed latent representation, thereby offering a potential solution for restoring “open-type” gaps located in extremely data-sparse regions that lack local neighborhood information. The incorporation of such mechanisms is expected to significantly increase the model’s robustness and restoration accuracy when handling extremely sparse and uneven data distributions, thereby advancing generative artificial intelligence toward deeper and more practical applications in complex marine surveying scenarios.




Figure 17. “Open-type” data void formed under conditions of extreme sparsity and its restoration results.

Additionally, as mentioned in Section 2.1, when constructing the data degradation model in this study, the coverage width of the MBES bathymetric swath was assumed to be a fixed value, primarily to simplify mathematical expressions. However, in actual operations, swath coverage width varies dynamically with factors such as water depth and topographic relief. For instance, insufficient coverage may occur in areas with steep slopes, such as seamounts, leading to localized data gaps. This simplifying assumption introduces a discrepancy between the simulated data and the complexity of real-world conditions. It is worth noting that irregular gaps caused by topographic variations can be categorized as a type of non-systematic gap pattern. The random degradation simulation method employed in this study, to some extent, captures the characteristics of such irregular gaps. Therefore, the proposed method remains effective in addressing certain gap issues arising from real topographic variations. In future research, it will be necessary to consider more realistic application scenarios. On one hand, the degradation algorithm should be refined to incorporate mechanisms that account for the dynamic influence of terrain on swath coverage width. On the other hand, efforts should be made to accumulate higher-quality, multi-scenario measured datasets. By enhancing the realism and diversity of training samples, it is anticipated that the model’s generalization capability and restoration accuracy in complex real-world scenarios can be further strengthened.

5. Conclusions

This study addresses the core challenge of restoring gaps between MBES bathymetric swaths by proposing a complete solution framework that encompasses a problem definition, data simulation, and model construction. First, based on practical marine surveying operations, two data degradation models, random degradation and rule-based degradation—are explicitly defined, both of which have clear real-world implications. This categorization provides a highly realistic and quantifiable data foundation for model training and evaluation. Furthermore, the MBES-DDPM is introduced; its core innovation is the development of a multisource data-fusion-guided denoising diffusion probabilistic framework. For the first time, gravity anomaly data are incorporated as a macro-topographic prior. With a carefully designed conditional injection mechanism, the diffusion process is guided to generate seabed topography consistent with regional geophysical princi-

ples, thereby advancing the paradigm from “data-driven interpolation” to “multisource knowledge-guided generation”.

The experimental results demonstrate that the MBES-DDPM exhibits outstanding and stable restoration performance across various typical and complex seabed topographic scenarios. Compared with mainstream interpolation methods and state-of-the-art deep learning models, the proposed approach achieves significant quantitative advantages: the RMSE of the restored results is reduced by an average of at least 34.21%, and the PSNR is improved by more than 3.71 dB on average. More importantly, in qualitative assessments and analyses of the terrain morphological metrics (such as the slope fidelity), the MBES-DDPM successfully recovers more continuous and reasonable macro-topographic structures as well as finer micro-geomorphic details, effectively avoiding common issues such as over-smoothing distortions, structural artifacts, or spurious feature generation. These outcomes fully validate the effectiveness and superiority of the multisource information fusion guidance mechanism.

Thus, this study provides a reliable new method for the high-precision restoration of large-scale MBES bathymetric swaths, contributing to the advancement of full-coverage, high-resolution seabed topographic modeling. However, the current approach remains constrained by its patch-based local processing paradigm when handling extremely large-area scenes, experiencing challenges when addressing gaps that are highly sparse or unevenly distributed. In the future, we will explore global restoration frameworks based on architectures such as Latent Diffusion Models to overcome the local receptive-field limitations and leverage broader spatial contextual information, thereby further enhancing the model’s performance in large-scale restoration tasks.

Author Contributions: Conceptualization, J.C., Z.W. and X.B.; methodology, J.C., Z.W. and D.Z.; software, D.Z. and J.Z.; validation, D.Z. and J.S.; formal analysis, X.B.; investigation, Z.W., X.B., J.Z. and M.W.; resources, Z.W.; data curation, J.C. and D.Z.; writing—original draft preparation, J.C.; writing—review and editing, Z.W. and D.Z.; visualization, J.C. and D.Z.; supervision, Z.W. and Y.L.; project administration, J.C.; funding acquisition, Z.W., X.B., J.Z. and M.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Specialized Scientific Research Fund of State Key Laboratory of Submarine Geoscience [grant number SGLabZZKT2025-02], the National Key Research and Development Program of China [grant numbers 2022YFC2806600, 2022YFC2806605], the National Natural Science Foundation of China [grant numbers 41830540, 42204049, 42306210], the Oceanic Interdisciplinary Program of Shanghai JiaoTong University [grant numbers SL2020ZD204, SL2023ZD203, SL2023ZD102], the National Natural Science Foundation of Shandong Province [grant number ZR2025MS593], the Qingdao Natural Science Foundation [grant numbers 24-8-4-zrjj-2-jch], and the Scientific Research Fund of the Second Institute of Oceanography, MNR [grant numbers QNYC2403, JG2503, XRJH2505].

Data Availability Statement: The Gebco_2025 model is accessible via: <https://www.gebco.net/data-products/gridded-bathymetry-data#global> (accessed on 30 January 2026). Meanwhile, the SIO V32.1 gravity anomaly model is accessible via: https://topex.ucsd.edu/pub/global_grav_1min (accessed on 30 January 2026). Most of the data presented in this paper is included in the main manuscript, and additional data are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Heine, C.; Müller, R.D.; Steinberger, B.; DiCaprio, L. Integrating Deep Earth Dynamics in Paleogeographic Reconstructions of Australia. *Tectonophysics* **2010**, *483*, 135–150. [CrossRef]
2. Zhang, K.; Wang, Y.; Luo, Y.; Zhao, D.; Wang, M.; Yang, F.; Wu, Z. Complex Tsunamigenic Near-Trench Seafloor Deformation during the 2011 Tohoku–Oki Earthquake. *Nat. Commun.* **2023**, *14*, 3260. [CrossRef]

3. Zhou, J.; Wu, Z.; Jin, X.; Zhao, D.; Cao, Z.; Guan, W. Observations and Analysis of Giant Sand Wave Fields on the Taiwan Banks, Northern South China Sea. *Mar. Geol.* **2018**, *406*, 132–141. [\[CrossRef\]](#)
4. Wu, Z.; Jin, X.; Li, J.; Zheng, Y.; Wang, X. Linear sand ridges on the outer shelf of the East China Sea. *Chin. Sci. Bull.* **2005**, *50*, 2517–2528. [\[CrossRef\]](#)
5. Zhou, J.; Wu, Z.; Zhao, D.; Guan, W.; Cao, Z.; Wang, M. Effect of Topographic Background on Sand Wave Migration on the Eastern Taiwan Banks. *Geomorphology* **2022**, *398*, 108030. [\[CrossRef\]](#)
6. Hu, H.; Niu, X.; Wang, W.; Xu, W.; Ruan, A.; Gong, W.; Wei, X.; Xu, M.; Li, T. Ocean Bottom Seismograph Relocation and Time Correction Using the MCMC Algorithm. *Acta Oceanol. Sin.* **2025**, *44*, 101–111. [\[CrossRef\]](#)
7. Zhu, Z.; Tao, C.; Wu, T.; Schneider Von Deimling, J.; Zhang, J.; Zhang, G. Seafloor Classification by Fusing AUV Acoustic and Magnetic Data: Toward Complex Deep-Sea Environments. *IEEE Trans. Geosci. Remote Sens.* **2025**, *63*, 4203215. [\[CrossRef\]](#)
8. Kang, J.Y.; Lee, B.S. Optimisation of Pipeline Route in the Presence of Obstacles Based on a Least Cost Path Algorithm and Laplacian Smoothing. *Int. J. Nav. Archit. Ocean Eng.* **2017**, *9*, 492–498. [\[CrossRef\]](#)
9. Serov, P.; Mattingdal, R.; Winsborrow, M.; Patton, H.; Andreassen, K. Widespread Natural Methane and Oil Leakage from Sub-Marine Arctic Reservoirs. *Nat. Commun.* **2023**, *14*, 1782. [\[CrossRef\]](#) [\[PubMed\]](#)
10. Wang, M.; Wu, Z.; Best, J.; Yang, F.; Li, X.; Zhao, D.; Zhou, J. Using Multibeam Backscatter Strength to Analyze the Distribution of Manganese Nodules: A Case Study of Seamounts in the Western Pacific Ocean. *Appl. Acoust.* **2021**, *173*, 107729. [\[CrossRef\]](#)
11. Sederlin, M.; Flötteröd, G. Estimation of Near-Coastal Bathymetry Using AIS Ship Movements. *WMU J. Marit. Aff.* **2024**, *23*, 437–455. [\[CrossRef\]](#)
12. Yu, Y.; Sandwell, D.T.; Dibarboure, G. Abyssal Marine Tectonics from the SWOT Mission. *Science* **2024**, *386*, 1251–1256. [\[CrossRef\]](#)
13. Qin, X.; Wu, Z.; Luo, X.; Shang, J.; Zhao, D.; Zhou, J.; Cui, J.; Wan, H.; Xu, G. MuSRFM: Multiple Scale Resolution Fusion Based Precise and Robust Satellite Derived Bathymetry Model for Island Nearshore Shallow Water Regions Using Sentinel-2 Multi-Spectral Imagery. *J. Photogramm. Remote Sens.* **2024**, *218*, 150–169. [\[CrossRef\]](#)
14. Becker, J.J.; Sandwell, D.T. Global Estimates of Seafloor Slope from Single-beam Ship Soundings. *J. Geophys. Res.* **2008**, *113*, 2006JC003879. [\[CrossRef\]](#)
15. Caballero, I.; Stumpf, R.P. Confronting Turbidity, the Major Challenge for Satellite-Derived Coastal Bathymetry. *Sci. Total Environ.* **2023**, *870*, 161898. [\[CrossRef\]](#) [\[PubMed\]](#)
16. Liu, Z.; Liu, H.; Ma, Y.; Ma, X.; Yang, J.; Jiang, Y.; Li, S. Exploring the Most Effective Information for Satellite-Derived Bathymetry Models in Different Water Qualities. *Remote Sens.* **2024**, *16*, 2371. [\[CrossRef\]](#)
17. Drennon, H.; Ferrini, V.; Morton, J.J. Multibeam Data Curation and Access through the Global Multi-Resolution Topography (GMRT) Synthesis. In Proceedings of the OCEANS 2023—MTS/IEEE U.S. Gulf Coast, Biloxi, MS, USA, 25–28 September 2023; pp. 1–4. [\[CrossRef\]](#)
18. Lu, Y.; Xu, J.; Zhong, Y.; Lin, H. Enhancing Marine Topography Mapping: A Geometrically Optimized Algorithm for Multibeam Echosounder Survey Efficiency and Accuracy. *Appl. Sci.* **2024**, *14*, 8875. [\[CrossRef\]](#)
19. Taak, M.; Yasrab, M.; Jamshed, A. Ocean Bathymetry: Decadal Advances, Persistent Challenges, and Future Horizons. *Int. Hydrogr. Rev.* **2025**, *31*, 196–204. [\[CrossRef\]](#)
20. Mayer, L.; Jakobsson, M.; Allen, G.; Dorschel, B.; Falconer, R.; Ferrini, V.; Lamarche, G.; Snaith, H.; Weatherall, P. The Nippon Foundation—GEBCO Seabed 2030 Project: The Quest to See the World’s Oceans Completely Mapped by 2030. *Geosciences* **2018**, *8*, 63. [\[CrossRef\]](#)
21. Ryan, W.B.F.; Carbotte, S.M.; Coplan, J.O.; O’Hara, S.; Melkonian, A.; Arko, R.; Weissel, R.A.; Ferrini, V.; Goodwillie, A.; Nitsche, F.; et al. Global Multi-Resolution Topography Synthesis. *Geochem. Geophys. Geosystems* **2009**, *10*, 2332. [\[CrossRef\]](#)
22. GEBCO Compilation Group. GEBCO 2025 Grid. 2025. Available online: https://www.bodc.ac.uk/data/published_data_library/catalogue/10.5285/37c52e96-24ea-67ce-e063-7086abc05f29 (accessed on 30 January 2026).
23. Shepard, D. A Two-Dimensional Interpolation Function for Irregularly-Spaced Data. In Proceedings of the 1968 23rd ACM National Conference, New York, NY, USA, 27–29 August 1968; pp. 517–524. [\[CrossRef\]](#)
24. Reuter, H.I.; Nelson, A.; Jarvis, A. An Evaluation of Void-filling Interpolation Methods for SRTM Data. *Int. J. Geogr. Inf. Sci.* **2007**, *21*, 983–1008. [\[CrossRef\]](#)
25. Jassim, F.A.; Altaany, F.H. Image Interpolation Using Kriging Technique for Spatial Data. *arXiv* **2013**, arXiv:1302.1294. [\[CrossRef\]](#)
26. Yue, Y.; Luo, M.; Zhang, B. Study on Spatial Distribution of DEM Interpolation Errors in the Gully of Dry-Hot Valley. *Geomatics and Information Science of Wuhan University* **2018**, *43*, 1122–1128. [\[CrossRef\]](#)
27. Luedeling, E.; Siebert, S.; Buerkert, A. Filling the Voids in the SRTM Elevation Model—A TIN-Based Delta Surface Approach. *J. Photogramm. Remote Sens.* **2007**, *62*, 283–294. [\[CrossRef\]](#)
28. Tao, P.; Li, Z.; Hu, Z.; Xiong, Q.; Chen, Q. A Simple Adaptive DEM Void Filling Method Based on Raster Distance Transform. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 6012805. [\[CrossRef\]](#)

29. Yue, L.; Yu, W.; Shen, H.; Zhang, L.; He, Y. Accuracy assessment of SRTM V4.1 and ASTER GDEM V2 in high-altitude mountainous areas: A case study in yulong snow mountain, China. In Proceedings of the 2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Milan, Italy, 26–31 July 2015; pp. 5011–5014. [\[CrossRef\]](#)
30. Zhao, J.; Yuan, Y.; Dong, Y.; Li, Y.; Shao, C.; Yang, H. Void Filling of Digital Elevation Models Based on Terrain Feature-Guided Diffusion Model. *Remote Sens. Environ.* **2024**, *315*, 114432. [\[CrossRef\]](#)
31. Qiu, Z.; Yue, L.; Liu, X. Void Filling of Digital Elevation Models with a Terrain Texture Learning Model Based on Generative Adversarial Networks. *Remote Sens.* **2019**, *11*, 2829. [\[CrossRef\]](#)
32. Zhou, G.; Song, B.; Liang, P.; Xu, J.; Yue, T. Voids Filling of DEM with Multiattention Generative Adversarial Network Model. *Remote Sens.* **2022**, *14*, 1206. [\[CrossRef\]](#)
33. Zhang, C.; Shi, S.; Ge, Y.; Liu, H.; Cui, W. DEM Void Filling Based on Context Attention Generation Model. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 734. [\[CrossRef\]](#)
34. Dhariwal, P.; Nichol, A. Diffusion Models Beat GANs on Image Synthesis. In Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS), Virtual Conference, 6–14 December 2021; Curran Associates Inc.: Red Hook, NY, USA, 2021; pp. 8780–8794.
35. Bau, D.; Zhu, J.-Y.; Wulff, J.; Peebles, W.; Zhou, B.; Strobel, H.; Torralba, A. Seeing What a GAN Cannot Generate. In Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea, 29 October–1 November 2019; pp. 4501–4510. [\[CrossRef\]](#)
36. Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; Change Loy, C. Esrgan: Enhanced Super-Resolution Generative Adversarial Networks. In Proceedings of the European Conference on Computer Vision (ECCV) Workshops, Munich, Germany, 8–14 September 2018; pp. 63–79. [\[CrossRef\]](#)
37. Ye, J.; Deng, Y.; Shao, Z.; Xiang, N.; Ou, Y. A DEM Image Superresolution Reconstruction Method Based on the Texture Transfer of High-Resolution Remote Sensing Images. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2024**, *17*, 11536–11549. [\[CrossRef\]](#)
38. Sohl-Dickstein, J.; Weiss, E.A.; Maheswaranathan, N.; Ganguli, S. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In Proceedings of the 32nd International Conference on International Conference on Machine Learning (ICML), Lille, France, 6–11 July 2015; pp. 2256–2265.
39. Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. In Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS), Virtual Conference, 6–12 December 2020; Curran Associates Inc.: Red Hook, NY, USA, 2020; pp. 6840–6851.
40. Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. In Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 19–24 June 2022; pp. 10674–10685. [\[CrossRef\]](#)
41. Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D.J.; Norouzi, M. Image Super-Resolution Via Iterative Refinement. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *45*, 4713–4726. [\[CrossRef\]](#) [\[PubMed\]](#)
42. Grechka, A.; Couairon, G.; Cord, M. GradPaint: Gradient-Guided Inpainting with Diffusion Models. *Comput. Vis. Image Underst.* **2024**, *240*, 103928. [\[CrossRef\]](#)
43. Zhang, W.; Zhang, X.; Dong, J.; Song, X.; Pang, R. CIDM: A Comprehensive Inpainting Diffusion Model for Missing Weather Radar Data with Knowledge Guidance. *J. Photogramm. Remote Sens.* **2025**, *221*, 299–309. [\[CrossRef\]](#)
44. Panangian, D.; Bittner, K. Dfilled: Repurposing Edge-Enhancing Diffusion for Guided DSM Void Filling. In Proceedings of the 2025 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW), Tucson, AZ, USA, 28 February–4 March 2025; pp. 526–534. [\[CrossRef\]](#)
45. Shih-Huang Lo, K.; Peters, J. Diff-DEM: A Diffusion Probabilistic Approach to Digital Elevation Model Void Filling. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 1–5. [\[CrossRef\]](#)
46. Yuan, Y.; Zhao, J.; Shao, C.; Dong, Y. A Terrain Feature Guided-Diffusion Model for Void Filling of Digital Elevation Models. In Proceedings of the 2024 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Brisbane, Australia, 3–8 August 2024; pp. 7007–7010. [\[CrossRef\]](#)
47. Park, Y.-J.; Choi, H.-L. A Neural Process Approach for Probabilistic Reconstruction of No-Data Gaps in Lunar Digital Elevation Maps. *Aerosp. Sci. Technol.* **2021**, *113*, 106672. [\[CrossRef\]](#)
48. Yue, L.; Gao, B.; Zheng, X. Generative DEM Void Filling with Terrain Feature-Guided Transfer Learning Assisted by Remote Sensing Images. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 6011105. [\[CrossRef\]](#)
49. Dong, G.; Huang, W.; Smith, W.A.P.; Ren, P. Filling Voids in Elevation Models Using a Shadow-Constrained Convolutional Neural Network. *IEEE Geosci. Remote Sens. Lett.* **2020**, *17*, 592–596. [\[CrossRef\]](#)
50. Li, S.; Hu, G.; Cheng, X.; Xiong, L.; Tang, G.; Strobl, J. Integrating Topographic Knowledge into Deep Learning for the Void-Filling of Digital Elevation Models. *Remote Sens. Environ.* **2022**, *269*, 112818. [\[CrossRef\]](#)

51. Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; Van Gool, L. RePaint: Inpainting Using Denoising Diffusion Probabilistic Models. In Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 19–24 June 2022; pp. 11451–11461. [\[CrossRef\]](#)
52. Han, Q.; Huang, B.; Li, Y. SAR-Conditioned Consistency Model for Effective Cloud Removal in Remote Sensing Images. *Remote Sens.* **2025**, *17*, 3721. [\[CrossRef\]](#)
53. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Munich, Germany, 6–8 October 2015; pp. 234–241. [\[CrossRef\]](#)
54. Peebles, W.; Xie, S. Scalable Diffusion Models with Transformers. In Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, 2–6 October 2023; pp. 4172–4182. [\[CrossRef\]](#)
55. Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; Guo, B. Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. In Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 11–17 October 2021; pp. 9992–10002. [\[CrossRef\]](#)
56. Han, L.; Zhao, Y.; Lv, H.; Zhang, Y.; Liu, H.; Bi, G.; Han, Q. Enhancing Remote Sensing Image Super-Resolution with Efficient Hybrid Conditional Diffusion Model. *Remote Sens.* **2023**, *15*, 3452. [\[CrossRef\]](#)
57. Yue, Z.; Wang, J.; Loy, C.C. ResShift: Efficient Diffusion Model for Image Super-Resolution by Residual Shifting. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NeurIPS), New Orleans, USA, 10–16 December 2023*; Curran Associates Inc.: Red Hook, NY, USA, 2023; pp. 13294–13307.
58. Yu, X.; Li, Y.; Ma, J.; Li, C.; Wu, H. Diffusion-RSCC: Diffusion Probabilistic Model for Change Captioning in Remote Sensing Images. *IEEE Trans. Geosci. Remote Sens.* **2025**, *63*, 5617013. [\[CrossRef\]](#)
59. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely Connected Convolutional Networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 2261–2269. [\[CrossRef\]](#)
60. Mehralian, M.; Karasfi, B. RDCGAN: Unsupervised Representation Learning with Regularized Deep Convolutional Generative Adversarial Networks. In Proceedings of the 9th Conference on Artificial Intelligence and Robotics and 2nd Asia-Pacific International Symposium, Kish Island, Iran, 10 December 2018; pp. 31–38. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.