

Dfilled: Repurposing Edge-Enhancing Diffusion for Guided DSM Void Filling

Daniel Panangian* Ksenia Bittner

The Remote Sensing Technology Institute

German Aerospace Center (DLR), Wessling, Germany

{daniel.panangian, ksenia.bittner}@dlr.de

Abstract

Digital Surface Models (DSMs) are essential for accurately representing Earth’s topography in geospatial analyses. DSMs capture detailed elevations of natural and man-made features, crucial for applications like urban planning, vegetation studies, and 3D reconstruction. However, DSMs derived from stereo satellite imagery often contain voids or missing data due to occlusions, shadows, and low-signal areas. Previous studies have primarily focused on void filling for digital elevation models (DEMs) and Digital Terrain Models (DTMs), employing methods such as inverse distance weighting (IDW), kriging, and spline interpolation. While effective for simpler terrains, these approaches often fail to handle the intricate structures present in DSMs. To overcome these limitations, we introduce Dfilled, a guided DSM void filling method that leverages optical remote sensing images through edge-enhancing diffusion. Dfilled repurposes deep anisotropic diffusion models, which originally designed for super-resolution tasks, to inpaint DSMs. Additionally, we utilize Perlin noise to create inpainting masks that mimic natural void patterns in DSMs. Experimental evaluations demonstrate that Dfilled surpasses traditional interpolation methods and deep learning approaches in DSM void filling tasks. Both quantitative and qualitative assessments highlight the method’s ability to manage complex features and deliver accurate, visually coherent results.

1. Introduction

Digital elevation models (DEMs)—encompassing both Digital Surface Models (DSMs) and Digital Terrain Models (DTMs)—are essential tools in geospatial analysis, representing the Earth’s topography. DSMs, in particular, include elevations of natural and man-made features such as vegetation and buildings, offering more detailed and complex information than DTMs. This richness makes DSMs




Figure 1. Overview of the proposed guided DSM void filling approach. The method utilizes a image guidance to fill voids in the DSM, resulting in a complete and accurate DSM reconstruction.

invaluable for applications requiring fine-scale surface features, including urban planning, vegetation analysis, and 3D reconstruction.

However, DSMs generated from stereo satellite imagery often suffer from voids or holes (areas with missing elevation data) due to mismatches in occluded, shadowed, or low-signal regions. These voids degrade the quality and reliability of DSMs, impacting critical tasks in photogrammetry and remote sensing like change detection, object recognition, and 3D modeling. Filling these voids is particularly challenging in DSMs compared to DTMs because of the additional complexity introduced by surface features such as buildings and trees.

Traditional void filling methods involve interpolation techniques estimate unknown values using spatially neighboring information [15]. While these methods may be sufficient for small gaps in DTMs, they often struggle with the complex features present in DSMs, especially when dealing with large voids or intricate urban landscapes. The reliability of these methods diminishes as the void size increases or

*Corresponding author

the terrain becomes more complex, leading to inaccurate or unrealistic surface representations.

To overcome these limitations, researchers have explored the use of auxiliary data sources, such as integrating multiple DEMs or utilizing remote sensing imagery. Techniques like fill-and-feather, delta surface fill, and moving window erosion have been developed to fuse data from various sources [4, 6]. However, these approaches may falter in areas with unreliable auxiliary data or fail to capture the detailed features inherent in DSMs.

The advent of deep learning has opened new avenues for DEM void filling. Generative Adversarial Networks (GANs), in particular, have shown promise due to their ability to generate realistic textures and patterns. Various GAN-based methods have been proposed, including conditional GANs [2], Wasserstein GANs with contextual attention mechanisms [5], and multi-attention GANs [24]. Some approaches have incorporated topographic features like slope and relief degree to enhance the training process [14]. An alternative strategy is to utilize optical remote sensing images as auxiliary data for void filling. Remote sensing imagery provides rich spectral and textural information that correlates with the surface features present in DSMs. For example, buildings, trees, and other structures visible in optical images correspond to features in DSMs, offering valuable information for reconstructing missing elevation data. Previous work has attempted to incorporate such imagery by using shadow maps or other features to guide DEM reconstruction [3]. Nonetheless, these methods often focus on DTMs or small missing areas and may not fully leverage the rich feature information inherent in remote sensing imagery. They may also lack the ability to preserve the fine-scale details of man-made structures and vegetation, which are critical in DSM applications.

By utilizing edge-enhancing diffusion techniques, our proposed method enhances edges with optical image’s guidance (see Fig. 1), which is critical for maintaining structural integrity in DSMs.

The key contributions of this work are:

1. We introduce a novel approach that adapts deep anisotropic diffusion models, which originally designed for super-resolution tasks, to the problem of DSMs void filling. By redefining the problem formulation using the heat equation and modifying the model to handle localized missing data, including local refinement strategies for coarse reconstruction, we effectively propagate contextual information into voids while preserving critical structural details.
2. We employ Perlin noise to generate inpainting masks that simulate the natural void patterns found in DSMs. This ensures that the model is trained on realistic missing data scenarios and enhances its ability to generalize

to real DSM voids.

3. We demonstrate the effectiveness of the proposed method through extensive experiments on various simulated and real DSM datasets. We propose the use of Perlin noise also for realistic evaluation. Our approach outperforms traditional interpolation techniques and state-of-the-art deep learning methods, particularly in handling complex features and providing accurate, visually realistic results in handling both small and large void filling in DSM.

2. Related Work

Voids in DEMs typically arise from limitations in data acquisition technologies such as radar or Light Detection and Ranging (LiDAR). Factors like water bodies, dense vegetation, low reflectivity surfaces, and complex terrain can impede the collection of accurate elevation data. For example, radar-based missions like the Shuttle Radar Topography Mission (SRTM) are prone to issues like shadowing and layover in steep or rugged terrains, leading to data gaps [6]. Additionally, atmospheric conditions and instrument limitations can contribute to missing data in DEMs. In contrast, voids in DSMs often result from challenges inherent in stereo image matching processes used to generate DSMs from optical imagery. Occlusions caused by tall structures, shadows cast by buildings or terrain features, and regions with low texture or homogenous surfaces hinder the matching algorithms [9]. Urban environments with dense infrastructure and areas with significant vegetation present particular difficulties, leading to more frequent and extensive voids in DSMs compared to DTMs.

2.1. Classical Methods for Void Filling

To address the issue of voids in DEMs and DSMs, a variety of traditional methods have been developed. In DEMs, interpolation techniques such as inverse distance weighting (IDW), kriging, and spline interpolation are commonly used [15]. These methods estimate missing elevation values based on the spatial correlation of surrounding data points. While effective for small voids in relatively flat and homogeneous terrains, their accuracy diminishes with increasing void size and terrain complexity. Many methods also have been proposed to leverage auxiliary data. The Fill and Feather (FF) [4] technique replaces missing data with values from an auxiliary DEM and applies smoothing at the edges to ensure seamless transitions. The Delta Surface Fill (DSF) [6] method creates a delta surface by computing the difference between the DEM and a resampled auxiliary surface, which is then used to adjust the fill surface for smooth integration. These methods rely heavily on the availability of high-quality auxiliary DEMs.




Figure 2. Summary of the proposed architecture for DSM void filling. The method comprises a two-step process: First, high-dimensional features are extracted from both the void-filled DSM and high-resolution optical imagery using a pre-trained feature extractor. Next, a refinement network integrates residual blocks and upsampling operations to reconstruct missing elevation values. This is followed by an edge-enhancing diffusion network that iteratively refines the DSM, leveraging edge features from the optical imagery to ensure accurate and realistic reconstruction of terrain and structural details.

In the context of DSMs, traditional interpolation methods often struggle due to the complexity introduced by surface features like buildings and vegetation. To overcome this, Krauß *et al.* [9] proposed a context-based approach for filling voids in DSMs generated from dense stereo matching. They categorized voids based on their characteristics and applied tailored strategies for each type, improving the accuracy of the filling process in complex urban areas. However, these methods may still be limited by the quality and availability of auxiliary data.

2.2. Deep Learning-Based Methods

The rise of deep learning has led to the development of more advanced void filling techniques for DEMs. GANs have been employed due to their ability to model complex data distributions and generate realistic outputs. Dong *et al.* [2] introduced a conditional GAN for recovering missing elevation data, demonstrating improved performance over traditional interpolation methods. Gavriil *et al.* [5] proposed a Wasserstein GAN with a contextual attention mechanism to enhance texture generation in void regions of DEMs. For DSMs, deep learning-based void filling methods are less prevalent but have shown potential. The complexity of DSM data, which includes detailed representations of surface features, poses significant challenges for modeling. Existing models often require extensive training data and may not fully exploit available information, such as high-resolution optical imagery, to improve void filling. Moreover, deep learning models may struggle to preserve fine-scale details crucial in DSM applications, such as edges of buildings and vegetation structures. Further enhancements have been achieved by integrating attention mechanisms into a conditional GAN, such as in the context-aware models proposed by Zhang *et al.* [22] and Zhou *et al.* [24], which

leverage multi-attention mechanisms to improve void filling performance. Additionally, domain-specific constraints have been introduced to guide restoration. For instance, Shadow-constrained GAN (SCGAN) incorporates terrain shadow geometry into its loss function [18], while the Topographic Knowledge-Constrained GAN (TKCGAN) penalizes incorrect valley and ridge predictions [23]. Recently, Lo *et al.* [11] proposed a novel approach using a conditional Denoising Diffusion Probabilistic Model (DDPM) for void filling in DEMs. The Diff-DEM framework leverages the iterative refinement capabilities of DDPMs to generate high-quality reconstructions. This method represents a shift from traditional GANs to probabilistic diffusion models, offering promising results in both accuracy and robustness.

2.3. Use of Guidance in Void Filling Methods

Incorporating guidance from auxiliary data modalities has proven effective in enhancing void filling techniques. For DEMs, remote sensing imagery provides valuable contextual information that aids in improving reconstruction quality. Dong *et al.* [3] utilized shadow maps extracted from optical images as constraints within a convolutional neural network, resulting in improved accuracy in void filling for DEMs. Similarly, Qiu *et al.* [14] introduced a terrain texture generation model that integrates topographic features, such as slope and relief degree, into a GAN framework, enhancing the realism of generated terrain in void regions of DEMs. Building on these advancements, Yue *et al.* [21] proposed a terrain feature-guided transfer learning approach assisted by remote sensing images. This method leverages terrain features and auxiliary data to guide the generative process, achieving more accurate and realistic void filling in DEMs. In the case of DSMs, the use of auxiliary guidance remains a question. Optical imagery, rich in spectral and

textural information, could serve as a valuable data source to capture surface feature details and improve reconstruction quality in DSM void filling tasks.

3. Methodology

3.1. Repurposing Deep Anisotropic Diffusion for Inpainting

Deep anisotropic diffusion models have demonstrated significant success in super-resolution tasks by enhancing spatial resolution through the reconstruction of high-frequency details from low-resolution images [12]. These models utilize anisotropic diffusion processes to propagate information directionally, effectively enhancing edges and fine details while mitigating noise and artifacts. Recognizing their potential in handling structural complexities, we repurpose deep anisotropic diffusion models for the task of image inpainting.

In image inpainting, the objective is to fill in missing or corrupted regions within an image in a manner that is coherent with the surrounding content. By adapting deep anisotropic diffusion models, we aim to effectively propagate contextual information from known regions into voids, resulting in seamless and visually plausible inpainted images (see Fig. 2).

3.2. Problem Formulation

The inpainting problem can be formulated using the heat equation, which models the diffusion of heat (or, analogously, information) over time. The classical isotropic heat equation is given by:

$$\frac{\partial u}{\partial t} = \nabla^2 u, \quad (1)$$

where $u(x, y, t)$ represents the image intensity at position (x, y) and time t , and ∇^2 is the Laplacian operator.

For image inpainting, we are interested in finding a steady-state solution ($\frac{\partial u}{\partial t} = 0$) to the anisotropic diffusion equation:

$$\frac{\partial u}{\partial t} = \nabla \cdot (D(x, y) \nabla u), \quad (2)$$

where $D(x, y)$ is the diffusion tensor that controls the rate and direction of diffusion at each point in the image. The diffusion tensor is designed to encourage diffusion along the edges and inhibit it across edges, preserving important structural details.

Our goal is to reconstruct the missing regions Ω_{missing} in the image $u(x, y)$ such that the inpainted image satisfies the anisotropic diffusion equation within Ω_{missing} and agrees with the known pixel values in the known regions Ω_{known} :

$$\begin{cases} \frac{\partial u}{\partial t} = \nabla \cdot (D(x, y) \nabla u), & \text{for } (x, y) \in \Omega_{\text{missing}}, \\ u(x, y) = u_0(x, y), & \text{for } (x, y) \in \Omega_{\text{known}}, \end{cases} \quad (3)$$

where $u_0(x, y)$ denotes the known pixel values.

By formulating the problem in terms of the heat equation, we can employ the deep anisotropic diffusion model to iteratively solve for $u(x, y)$, effectively diffusing information from known regions into the missing areas.

3.3. Proposed Model

To tailor the deep anisotropic diffusion model for inpainting, we introduce the following key extensions:

3.3.1 Coarse Reconstruction with Local Refinement Network

The initialization of missing regions significantly impacts the diffusion performance. Inadequate initialization can lead to poor convergence or unrealistic artifacts in the final output [13]. We implement the local refinement network for coarse reconstruction of the missing regions. This local refinement provides an informed starting point for the diffusion process, enabling the model to focus on refining details rather than constructing basic structures from scratch. To achieve this we implement the same network used in [13] by removing the residual connection.

3.3.2 Perlin Noise for Mask Generation

An essential aspect of training and evaluating inpainting models is the design of masks that simulate missing regions [17]. Previous methods often employ masks composed of irregular and rectangular shapes. While effective for general images, these masks may not accurately represent the void patterns found in DEMs and DSMs, which often exhibit natural, continuous missing regions due to occlusions or sensor limitations.

To address this, we utilize Perlin noise to generate masks that more closely resemble the voids in DEMs and DSMs, as seen in Fig. 3. Perlin noise is a gradient noise function producing natural-looking textures with continuous gradients, widely used in procedural texture generation. Algorithm 1 details the mask generation procedure.

4. Experiments

4.1. Datasets and Implementation

Imagery and Study Area We use DSMs and corresponding RGB orthoimages acquired over Ho Chi Minh City, Vietnam; Zurich and Bern, Switzerland; and Dushanbe, Tajikistan, all with a Ground Sampling Distance (GSD) of 0.5 m. The DSMs for Ho Chi Minh City and Dushanbe




Figure 3. Real and synthetic void masks for DSM void filling. (a) Real Voids illustrate naturally occurring, complex patterns in DSMs. (b) LaMa masks are structured synthetic voids often used in model training but may not accurately reflect real void distributions. (c) Perlin masks, generated with procedural noise, better mimic the irregularity and complexity of real voids.

Algorithm 1: PERLINMASK

Input: s : image size

Output: $mask$

Randomly select parameters: $scale(sc)$, $octaves(o)$, $persistence(p)$, $lacunarity(l)$, $base(b)$

Initialize an $s \times s$ array, $noise$

for $i = 0$ **to** $s - 1$ **do**

for $j = 0$ **to** $s - 1$ **do**

$noise[i, j] \leftarrow$
 PerlinNoise($i/sc, j/sc, o, p, l, b$)

Normalize $noise$ to $[0, 1]$

$threshold \leftarrow \text{Uniform}(0, 1)$

Initialize $mask$ as $s \times s$

for $i = 0$ **to** $s - 1$ **do**

for $j = 0$ **to** $s - 1$ **do**

$mask[i, j] \leftarrow \mathbf{1}[noise[i, j] > threshold]$

return $mask$

are generated from Pleiades 1B satellite data using a single triplet stereo acquisition. The high-resolution DSMs for Switzerland are LiDAR DSMs provided by The Federal

Office of Topography on the Swisstopo Portal¹ and are the same as those proposed in [13].

The dataset consists of approximately 4000 patches for training, 400 for validation, and 1300 for testing, each of size (256 px, 256 px). We use Ho Chi Minh City and Switzerland data for training and validation, and use Dushanbe data for testing. The study areas encompass diverse urban environments, including widely spaced, detached residential buildings, allotments, and high commercial buildings.

Implementation Details We randomly load training patches during training. We initialize the voids in the DSMs with the median value of the known regions. We normalize the data using min-max normalization: for the refinement network, each DSM is normalized to the range $[-1, 1]$; for the diffusion network, DSMs are normalized to the range $[0, 1]$. Optical images are normalized using ImageNet statistics. In the training data, voids masks are introduced similarly to the method presented in [10], generating irregular mask shapes to simulate voids in the DSMs. The testing data includes real voids masks and corresponding ground truth.

In all experiments, we use a hidden feature dimension of 64 for the feature extractor and the refinement decoder. A ResNet-50 backbone [7] pretrained on ImageNet [1] is used as the feature extractor. For training, we employ the L1 loss across all methods, including our own. For the diffusion network, we adopt the same setup and strategy outlined in [13]. This adjustment decreases smoothing and makes the filled DSM more coherent.

We employ the ADAM optimizer [8] with a base learning rate of 5×10^{-5} and no weight decay. The batch size is set to 8 for training, using an NVIDIA A100 GPU. We stop training once the root mean square error (RMSE) on the validation set has converged. Our model is implemented in PyTorch.

4.2. Baselines

- **Spline Interpolation:** A traditional interpolation technique that estimates missing DSMs values by fitting spline functions to the available data points
- **Diff-DEM [11]:** An adaptation of the Palette diffusion model [16] for DSM void filling. The model uses a U-Net architecture with attention mechanisms in deeper layers. Inputs are processed as dual-channel $2 \times 128 \times 128$ images, incorporating the DSM containing voids and the step t approximation of the target DSM.
- **RSAGAN [21]:** A GAN-based framework, where generator architecture includes two encoders with gated

¹<https://www.swisstopo.admin.ch/en/geodata.html>




Figure 4. Visual comparison of void filling results for small masks, where voids cover a relatively small portion of the DSM. The proposed Dfilled method produces more regularized and smoother results compared to RSAGAN, aligning closely with the ground truth while preserving fine-scale structural details. Green boxes indicate regions of interest.

Table 1. Comparison of models on real and synthetic datasets using RMSE, NMAD, and MedAE metrics. The table lists the model name, whether guide image guidance is incorporated and the training mask used. The performance metrics are evaluated on three different datasets: Real, Small Synthetic, and Large Synthetic.

Model	Image	Train	Real			Small Synthetic			Large Synthetic		
	Guidance	Mask	RMSE	NMAD	MedAE	RMSE	NMAD	MedAE	RMSE	NMAD	MedAE
Spline	-	-	16.69	3.85	16.88	18.80	2.45	18.55	20.56	2.98	18.93
Diff-DEM	-	LaMa	3.87	0.98	0.67	3.49	0.76	0.54	13.59	5.82	12.89
			3.02	0.94	0.67	3.50	0.76	0.54	13.20	5.82	12.90
RSAGAN	✓	LaMa	4.25	1.66	1.18	2.37	0.87	0.63	6.60	5.82	3.86
			3.71	0.83	0.57	1.28	0.27	0.19	3.62	1.57	1.27
Dfilled (Ours)	✓	LaMa	2.91	0.30	0.20	1.14	0.20	0.14	2.49	0.78	0.54
			3.17	0.49	0.33	1.23	0.20	0.14	2.53	0.97	0.70

convolutions [20] to extract features while considering void masks. A Pyramid, Cascading, and Deformable Convolution (PCD) alignment module [19] aligns image and DEM features. To address large-scale voids, RSAGAN employs the image context attention module (ICAM) for preliminary void filling and the terrain feature-guided residual pixel attention block (TFG-RPAB) to refine the features by transferring image textures to topographic attributes.

4.3. Evaluation Metrics

We evaluate the models’ performance by examining the RMSE, the normalized median absolute deviation (NMAD), and the median absolute error (MedAE), which are all derived from per-pixel differences between prediction and ground truth.

5. Results

5.1. Comparisons with Prior Works

Our proposed method, Dfilled, demonstrates superior performance in filling voids within DSMs, as evidenced by both quantitative and qualitative results. In Table 1, we compare four methods, Spline, Diff-DEM, RSAGAN, and Dfilled, across three different datasets (Real, Small Synthetic, and Large Synthetic), each trained on LaMa or Perlin masks. Dfilled obtains significantly lower RMSE, NMAD, and MedAE values compared to traditional interpolation (Spline) and state-of-the-art learning-based approaches (RSAGAN and Diff-DEM).

Interestingly, Diff-DEM, a recent single DSM inpainting approach, outperforms the guided RSAGAN on the Real dataset. However, Dfilled surpasses both methods across all three datasets. This underscores the robustness of our guided approach and its ability to leverage auxiliary information effectively during void filling. Although many methods benefit from being trained with our newly introduced Perlin mask, Dfilled achieves strong and consistent

results irrespective of whether the LaMa or Perlin mask is used in training, illustrating that its performance does not heavily depend on the specific mask type and can readily adapt to varying void patterns.

Complementing these quantitative findings, Fig. 4 and Fig. 5 provide visual comparisons for voids of different sizes. In scenarios with smaller voids (see Fig. 4), Dfilled not only reconstructs the missing elevation data more accurately than RSAGAN but also produces smoother, more regularized outputs, minimizing artifacts and ensuring better continuity of terrain and structural details. For larger voids covering up to 60–80% of the image (Fig. 5), Dfilled leverages high-resolution guide images to recover complex terrain and fine-scale features effectively. By contrast, RSAGAN frequently introduces artifacts and struggles to preserve fine details in both small and large void cases. Overall, these results highlight Dfilled’s adaptability, robustness, and ability to consistently produce realistic DSM reconstructions across a wide range of void sizes, making it a superior choice for demanding high-resolution applications.

5.2. Ablation Study

Table 2. Contribution of each component in our network

	Refinement	Diffusion	RMSE [m]
1		✓	3.49
2	✓		3.69
3	✓	✓	2.91

To assess the contribution of each component in our model, we conducted an ablation study by systematically enabling the Refinement and Diffusion modules. As shown in Tab. 2, using only the Diffusion module (row 1) achieves an RMSE of 3.49 m, while using only the Refinement module (row 2) yields 3.69 m. Enabling both modules (row 3)




Figure 5. Visual comparison of void filling results for large masks, where 60–80% of the DSM area is void-filled. The Dfilled method effectively utilizes high-resolution guide images to reconstruct missing elevation data, outperforming RSAGAN in recovering complex terrain features and structural details.




Figure 6. Visual comparison of ablation settings. (a) Using only Diffusion. (b) Using only Refinement. (c) Using both Refinement and Diffusion. Combining Refinement and Diffusion removes artifacts and smooths discontinuities.

significantly reduces the RMSE to 2.91 m, demonstrating the complementary nature of these components.

Figure 6 visually highlights the differences. Outputs with only the Refinement module (b) show noisy, dot-like artifacts, whereas adding the Diffusion module (c) removes these artifacts, producing smoother and more consistent results. This underscores the critical role of the Diffusion module in enhancing both accuracy and visual quality. Additionally, when combining both Refinement and Diffusion modules the resulted DSM is more regularized than the one produced without Refinement module (a).

6. Conclusion

In this paper, we introduced a novel approach for void filling in DSMs by adapting deep anisotropic diffusion models originally designed for super-resolution tasks. Our method redefines the problem using the heat equation and modifies the model to handle localized missing data through local refinement strategies for void initialization. This enables effective propagation of contextual information into voids while preserving critical structural details such as sharp building edges and smooth transitions in natural terrain.

To train our model on realistic missing data scenarios, we employed Perlin noise to generate inpainting masks that simulate natural void patterns commonly found in Digital Surface Models (DSMs). Our method demonstrates robustness against large masks and various types of masks, making it effective across diverse landscapes, including complex urban environments. Extensive experiments on both simulated and real DSM datasets show that our approach outperforms traditional interpolation techniques and state-of-the-art deep learning methods. Specifically, our method excels in handling complex features and provides accurate, visually realistic results where conventional methods often fail. By utilizing edge-enhancing diffusion techniques, our proposed method enhances edge and structural information critical for maintaining terrain integrity in DSMs.

ACKNOWLEDGEMENTS

Special thanks are given to the GAF AG for the provision of the Pleiades 1B data.

References

- [1] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5

- [2] G. Dong, F. Chen, and P. Ren. Filling srtm void data via conditional adversarial networks. In *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, pages 7441–7443, 2017. 2, 3
- [3] G. Dong, W. Huang, W. A. P. Smith, and P. Ren. Filling voids in elevation models using a shadow-constrained convolutional neural network. *IEEE Geosci. Remote Sens. Lett.*, 17(4):592–596, 2020. 2, 3
- [4] S. Dowding, T. Kuuskivi, and L. Xiaopeng. Void fill of srtm elevation data - principles, processes and performance. In *ASPRS Images to Decision: Remote Sensing Foundation for GIS Applications*, Kansas City, Missouri, 2004. 2
- [5] K. Gavril, G. Muntingh, and O. J. D. Barrowclough. Void filling of digital elevation models with deep generative models. *IEEE Geosci. Remote Sens. Lett.*, 16(10):1645–1649, 2018. 2, 3
- [6] G. Grohman, G. Kroenung, and J. Strebeck. Filling srtm voids: The delta surface fill method. *Photogrammetric Engineering and Remote Sensing*, 72:213–216, 2006. 2
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [8] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2014. 5
- [9] T. Krauß and P. d’Angelo. Morphological filling of digital elevation models. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVIII-4/W19:165–172, 2011. 2, 3
- [10] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, pages 85–100, 2018. 5
- [11] Shih-Huang Kyle Lo and Jörg Peters. Diff-dem: A diffusion probabilistic approach to digital elevation model void filling. *IEEE Geoscience and Remote Sensing Letters*, 21:1–5, 2024. 3, 5
- [12] Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Guided depth super-resolution by deep anisotropic diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18237–18246, June 2023. 4
- [13] D. Panangian and K. Bittner. Real-gdsr: Real-world guided dsm super-resolution via edge-enhancing residual network. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-2-2024:185–192, 2024. 4, 5
- [14] Z. Qiu, L. Yue, and X. Liu. Void filling of digital elevation models with a terrain texture learning model based on generative adversarial networks. *Remote Sens.*, 11(23):2829, 2019. 2, 3
- [15] H. I. Reuter, A. Nelson, and A. Jarvis. An evaluation of void-filling interpolation methods for srtm data. *International Journal of Geographical Information Science*, 21:983–1008, 2007. 1, 2
- [16] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH ’22, New York, NY, USA, 2022. Association for Computing Machinery. 5
- [17] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. *arXiv preprint arXiv:2109.07161*, 2021. 4, 5
- [18] Jun Wang, Yi Zhang, and Hao Li. Shadow-constrained gan (scgan) for dem void filling using terrain shadow geometry. *IEEE Transactions on Geoscience and Remote Sensing*, 59:4536–4548, 2021. 3
- [19] Xintao Wang, Kelvin C. K. Chan, K. Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. pages 1954–1963, 2019. 7
- [20] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas Huang. Free-form image inpainting with gated convolution. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4470–4479, 2019. 7
- [21] Linwei Yue, Bing Gao, and Xianwei Zheng. Generative dem void filling with terrain feature-guided transfer learning assisted by remote sensing images. *IEEE Geoscience and Remote Sensing Letters*, PP:1–1, 01 2024. 3, 5
- [22] C. Zhang, S. Shi, Y. Ge, H. Liu, and W. Cui. Dem void filling based on context attention generation model. *ISPRS International Journal of Geo-Information*, 9(12):734, 2020. 3
- [23] G. Zhou, B. Song, P. Liang, J. Xu, and T. Yue. Voids filling of dem with multiattention generative adversarial networks. *Remote Sensing*, 13(2):234, 2021. 3
- [24] G. Zhou, B. Song, P. Liang, J. Xu, and T. Yue. Voids filling of dem with multiattention generative adversarial network model. *Remote Sens.*, 14(5):1206, 2022. 2, 3