

RoofDiffusion: Constructing Roofs from Severely Corrupted Point Data via Diffusion

Kyle Shih-Huang Lo¹, Jörg Peters¹, and Eric Spellman²

¹ University of Florida, Gainesville FL 32611, USA

² Meta Platforms, Inc.

kyleshihuanglo@ufl.edu, jorg@cise.ufl.edu, espellman@meta.com




Fig. 1: RoofDiffusion restores height maps of challenging roof geometry, even under conditions of extreme sparsity, regional incompleteness, and noise. The bookend columns of point clouds are 3D views, the other columns are top views of height maps.

Abstract. Accurate completion and denoising of roof height maps are crucial to reconstructing high-quality 3D buildings. Repairing sparse points can enhance low-cost sensor use and reduce UAV flight overlap. RoofDiffusion is a new end-to-end self-supervised diffusion technique for robustly completing, in particular difficult, roof height maps. RoofDiffusion leverages widely-available curated footprints and can so handle up to 99% point sparsity and 80% roof area occlusion (regional incompleteness). A variant, No-FP RoofDiffusion, simultaneously predicts building footprints and heights. Both quantitatively outperform state-of-the-art unguided depth completion and representative inpainting methods for Digital Elevation Models (DEM), on both a roof-specific benchmark and the BuildingNet dataset. Qualitative assessments show the effectiveness of RoofDiffusion for datasets with real-world scans including AHN3, Dales3D, and USGS 3DEP LiDAR. Tested with the leading City3D algorithm, preprocessing height maps with RoofDiffusion noticeably improves 3D building reconstruction. RoofDiffusion is complemented by a new dataset of 13k complex roof geometries, focusing on long-tail issues in remote sensing; a novel simulation of tree occlusion; and a wide variety of large-area roof cut-outs for data augmentation and benchmarking. Code and dataset³: github.com/kylelo/RoofDiffusion

³ Created and released by the University of Florida




Fig. 2: Types of corrupted roof height maps with real-world scan.

Fig. 3: RoofDiffusion focuses on the “long tail” of hard-to-handle complex, corrupted geometry (inside the dashed blue).

1 Introduction

Digital Surface Models (DSMs), a.k.a. *height maps*, are monochromatic images where the pixel value captures the elevation of features, both natural and artificial. Height maps can be created by rasterizing airborne Light Detection and Ranging (LiDAR) point clouds. Height maps serve as an indispensable data source for reconstructing 3D models of urban buildings. Numerous studies [27, 29, 53, 92] have explored methods for generating compact 3D building models from such data. 3D building models are crucial in a wide range of applications, including navigation [5], urban planning [45], and simulations [33, 55]. With 500M buildings labeled in OpenStreetMap [56], even a 1% failure rate implies that 5M buildings need to be repaired. A close scrutiny of USGS 3DEP LiDAR sampled over Wayne County, MI [78] and Cambridge, MA [79] reveals, respectively, as much as 34% and 50% of height maps corrupted by incompleteness or noise. While height maps are broadly accessible, the following challenges often impede compact and precise building reconstructions.

- **Sparsity.** Factors such as low sensor resolution and poor surface reflectance cause low point density [90]. (Fig. 2a)
- **Incompleteness.** Portions of the roof data can be missing, due to environmental interference, occlusions by taller surrounding objects, or roof substructure when the camera angle is not orthogonal to the ground [75]. (Fig. 2b)
- **Noise.** Intrusions on building footprints, such as trees, can lead to incorrect reconstruction, resulting in artifacts like non-existing dormers. (Fig. 2c). Environmental factors like light and dust can also introduce noise.

Since the reconstruction quality of the 3D models is heavily influenced by the precision and completeness of the underlying height maps, roof repair is crucial, but is under-explored.

RoofDiffusion addresses all three bulleted challenges by building on and adjusting reliable techniques. Unlike typical image inpainting tasks such as [1, 14, 41, 43, 48], available height map pixels can be noisy, and completing sparse data does not fit the standard super-resolution paradigms [21, 35, 44, 98], because missing pixels are usually unevenly distributed. Acknowledging these complexities,

we nevertheless, succeed in conceptualizing roof repair as an image restoration task.

Thanks to the efforts of contributors worldwide, OpenStreetMap [56] now increasingly provides vectorized high-precision footprints [24, 29]. Consequently, recent building reconstruction algorithms [2, 24, 29] have incorporated footprints into the reconstruction process. However, these approaches still struggle with sparsity, incompleteness and noise. Footprint-guided RoofDiffusion addresses these long-tail challenges (Fig. 3).

Specifically, we learn a strong prior that accurately approximates the actual distribution of the roof height map. Taking inspiration from the restoration approach in [70] of Joint Photographic Experts Group (JPEG) images, we frame the roof height map restoration challenge as a denoising diffusion process. While JPEG restoration focuses on reverting compressed images to their original state, our strategy seeks to fill in roof heights and remove noise. The following are the major contributions of this paper:

- We introduce RoofDiffusion, a novel method based on conditional Diffusion Probabilistic Model (DPM) [70]. Both footprint-guided RoofDiffusion and a variant, No-FP RoofDiffusion, trained without footprint, allow robust roof height map completion under extreme conditions, surpassing the state-of-the-art depth completion methods [15, 85] and representative DEM inpainting algorithms [3, 32, 74]. The footprint-guided version handles up to 99% missing data points and 80% regional incompleteness while remaining resilient to tree occlusion noise (for some applications, construction from very few data might be flagged for hallucination potential). No-FP RoofDiffusion predicts both footprint and height.
- We propose a novel “tree planting” method for simulating tree occlusion noise, and we introduce multi-Gaussian masking for synthesizing incompleteness in roof height maps. These techniques enable data augmentation for self-supervised learning and benchmark creation for quantitative comparisons.
- We unveil the PoznanRD (Poznan Roof Dataset) with 13k Level of Detail (LoD) 2.2 [4] noise-free roof meshes and height maps. When treated with our noise and incompleteness algorithm, our complex roofs dataset can effectively generate ample training data to address “long tail” challenges, see Fig. 3.

2 Related Works

Reviewing Digital Elevation Model (DEM) inpainting from remote sensing is most relevant, followed by unguided depth completion methods for restoring dense, clean depth maps from sparse LiDAR points. Also, we review denoising diffusion models in computer vision tasks and roof datasets.

DEM Inpainting restores terrain height data missing caused by occlusions, such as mountains obstructing their own opposite sides and areas covered by water. Researchers commonly use Inverse Distance Weighting (IDW) [74],

Kriging [50, 65], and Spline fitting [32] for inpainting voids in DEMs. These methods work well for small areas but can fail in larger, complex regions due to the lack of terrain geometry knowledge. To address larger inpainting regions, Delta surface-based approaches [20, 47] fill voids with auxiliary DEM patches. Learning-based methods have gained prominence in DEM inpainting due to their superior feature learning capabilities. In particular, several methods [12, 13, 18, 40, 61, 89, 97, 102, 103] utilize Generative Adversarial Nets (GAN) [19, 52] for inpainting by conditioning models on voids to predict filled areas.

Compared to DEM inpainting, roof DSM inpainting deals with higher sparsity, more noise from elements such as trees, and larger areas of incompleteness.

Unguided Depth Completion aims to restore the dense depth map solely from sparse depth data, mainly for urban driving data [81]. The traditional Convolutional Neural Networks (CNNs) [36] often suffer from a mosaic effect of images with sparse pixel values. This issue is addressed by applying convolutions solely to valid pixels using a binary mask [30, 81]. However, a binary mask uniformly weights each pixel, conflicting with the reality that pixels have varying importance. To counter this, a continuous confidence mask was proposed in [16]. Furthermore, a learnable mask is introduced in [15] to address the issue that directly inferring a confidence mask from data may be problematic due to noise. To incorporate semantic learning, the works [46, 93] train the network to additionally predict auxiliary images, including RGB and normal maps. CU-Net [85] employs a two-stage U-Net for coarse-to-fine depth completion.

Most depth completion methods target evenly distributed, sparse depth maps without large missing areas. RoofDiffusion can handle height maps missing more than 80% of the area and still generate inpainting harmonious with the existing roof structure.

Diffusion Models have been proposed in pioneering works [11, 25, 77] showing remarkable image generation capability. Subsequent research has built upon this foundation by introducing text-conditioned models, enabling text-guided image generation [54, 62, 67, 69, 71, 99]. The technology has further been employed to significantly improve image enhancement tasks like super-resolution [17, 26, 72, 87], inpainting [48, 70, 88], and translation [38, 51, 58, 70]. There are also a few diffusion model-based methods for RGB-guided depth completion [37, 63].

To the best of our knowledge, our work pioneers the use of the diffusion model for DSM completion. We formulate our problem as an image restoration task, leveraging a conditional DPM. Inspired by [31], our approach eliminates the need for a confidence map [15, 16, 81], directly conditioning the model on the sparse image.

Roof Datasets commonly provide pairs of point clouds and ground truth mesh for 3D reconstruction research. UrbanScene3D [42] and STPLS3D [6] offer real-world LiDAR point clouds along with reconstructed dense triangular meshes. City3D [29] and Building3D [84] provide substantial datasets featuring more compact meshes for Computer-Aided Design (CAD)-like building-reconstruction research. However, in these datasets where ground truth is constructed via 3D reconstruction algorithms, mesh accuracy can be suboptimal due




Fig. 4: Training the RoofDiffusion reverse transition kernel. GT denotes ground truth, \otimes is transformation (3), \oplus^4 removes predicted noise from x_t by (4). The masked loss is defined by Eq. (5).

to algorithmic limitations and real-world scanning noise. Government datasets like [9, 10] provide compact noise-free meshes but exhibit misalignment with point clouds, often due to modeling simplifications. To ensure alignment between compact meshes and point clouds, BuildingNet [73] generates point clouds directly from compact meshes crafted by artists. However, BuildingNet [73] provides only 2k buildings and contains non-building elements such as people and vehicles. Similarly, we can sample point clouds from datasets with compact meshes, [60, 64, 86], but these have limited roof types.

Considering the limitations in quantity and variety of compact meshes, we propose the PoznanRD, featuring 13k noise-free complex roofs. Besides, existing datasets do not focus on corrupted data, so we introduce a corruption synthesizing approach to generate challenging data.

3 Problem Statement

Given a noisy, sparse, and incomplete roof height map z along with a building footprint m , the goal of RoofDiffusion is to estimate a complete and noise-free height map \hat{z} approximating to the ground truth z_{gt} . Measured from ground level, z , \hat{z} , and z_{gt} are single-channel images whose pixel values indicate building heights in meters. In z , zero-valued pixels signify either missing measurements or ground level. For \hat{z} and z_{gt} , zero values represent ground level. Generally, the number of non-zero pixels in \hat{z} and z_{gt} is significantly greater than in z . The mask m delineates the footprint, i.e. equals 1 for pixels belonging to the roof, and 0 otherwise.

Accurate reconstruction of real-world height maps is challenging due to a wide range of building heights, noise, and missing data. While the typical normalization for images or depth scenes of autonomous driving can employ fixed min-max values, the normalization required for roof heights needs to adapt to the wide range of roof heights for each building.

⁴ \hat{x}_{t-1} calculation can be ignored during training (the arrows colored as gray).

4 RoofDiffusion

We introduce a diffusion model-based roof completion method to reconstruct $\hat{\mathbf{z}}$ from \mathbf{z} so that $\hat{\mathbf{z}}$ is close to \mathbf{z}_{gt} . Section 4.1 presents the conversion of input height maps \mathbf{z} to roof-focused height images \mathbf{x} in the range $[-1, 1]$ suitable for diffusion models. Section 4.2 applies a diffusion model, conditioned on corrupted normalized height maps \mathbf{x} within their building footprint \mathbf{m} , to predict complete and noise-free height maps $\hat{\mathbf{x}}_0$.

4.1 Roof-Focused Height Map Normalization

Diverse building heights and the fixed value range of diffusion models require a careful normalization approach. To focus solely on the roof structure, we identify the lowest roof pixel and subtract this value from the entire height map. Let δ_i be the difference between the maximum and minimum height of the roof structure for i -th building. Then \bar{z} is the maximum of the δ_i after removing the largest 1% of δ_i . The normalized input height image is

$$\mathbf{x} := \frac{2}{\bar{z}} \left(\mathbf{z} - \frac{1}{2} (\mathbf{z}_{\max} + \mathbf{z}_{\min}) \right) \in [-1, 1], \quad (1)$$

where \mathbf{z}_{\min} denotes the smallest nonzero value in \mathbf{z} and \mathbf{z}_{\max} the largest. Our analysis of 13k buildings shows the cut-off value for \bar{z} is 10 meters.

4.2 Height Completion based on a Diffusion Model

Given \mathbf{x} , our goal is to generate a predicted image, $\hat{\mathbf{x}}_0$, that closely approximates the normalized ground truth image, \mathbf{x}_{gt} . We conceptualize this height completion task as a diffusion process. Drawing on the work of [25], our approach employs both a *forward* and a *reverse* Markov chain. The forward chain perturbs the ground truth data through noise injection, while the reverse chain utilizes a learnable model, conditioned on the input, to restore the data, see Fig. 4.

The Forward Process transforms a complex distribution $q(\mathbf{x}_{\text{gt}})$, the ground truth height maps, into the Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Starting with $\mathbf{x}_0 = \mathbf{x}_{\text{gt}} \sim q(\mathbf{x}_{\text{gt}})$, a transition kernel $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ generates a series of random variables \mathbf{x}_t , $t \in \{1, \dots, T\}$. The joint distribution of $q(\mathbf{x}_1, \dots, \mathbf{x}_T | \mathbf{x}_0)$, namely $\prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$, is marginalized based on [76] to derive closed-form expression

$$\mathbf{x}_t := f_{\text{forward}}(t, \mathbf{x}_0, \boldsymbol{\epsilon}, \mathbf{m}) \quad (2)$$

$$= \mathbf{m} \odot (\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}) - \mathbf{m}', \quad (3)$$

where \odot denotes element-wise multiplication, α_i is a hyperparameter, $\bar{\alpha}_t = \prod_{i=0}^t \alpha_i$, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Subtracting \mathbf{m}' , the complement mask of \mathbf{m} , sets pixel values outside the footprint to -1 for two advantages. First, the model

can infer the building footprint directly from the noise-injected areas in \mathbf{x}_t , obviating the need for an additional channel to represent the footprint. Second, non-building heights are prevented from influencing the prediction.

The Reverse Process sequentially removes noise from the data to yield a complete, noise-free version. Specifically, we can draw a noisy image \mathbf{x}_T from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ at time $t = T$. Then use a reverse transition kernel to recursively eliminate noise until $t = 0$, resulting in the restored height image, $\hat{\mathbf{x}}_0$. This transition kernel is parameterized by a learnable model ϵ_θ . While $\hat{\mathbf{x}}_0$ conforms to $q(\mathbf{x}_{\text{gt}})$, it may not always correspond to the normalized corrupted height image, \mathbf{x} . In our case, the objective is to ensure that the repaired roofs seamlessly integrate with pre-existing structures. To ensure a close association between \mathbf{x} and $\hat{\mathbf{x}}_0$, we condition ϵ_θ on \mathbf{x}_t , \mathbf{x} , and $\bar{\alpha}_t$, in the spirit of [70]. Each step of the reverse process can be formulated as

$$\hat{\mathbf{x}}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta \right) + \sqrt{1 - \alpha_t} \epsilon_t, \quad (4)$$

where α_t and $\bar{\alpha}_t$ are hyperparameters for variance scheduling, and $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Loss Function. For accurate noise prediction at various time steps t we train ϵ_θ with L_1 loss restricted to the footprint \mathbf{m} . For $t \sim \mathcal{U}(1, T)$ and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$,

$$L := \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}, \mathbf{m}), t, \epsilon} \|\mathbf{m} \odot (\epsilon - \tilde{\epsilon}_\theta)\|_1 \quad (5)$$

$$\text{where } \tilde{\epsilon}_\theta := \epsilon_\theta(f_{\text{forward}}(t, \mathbf{x}_0, \epsilon, \mathbf{m}), \mathbf{x}, \bar{\alpha}_t). \quad (6)$$

To restore the original scaling in meters for the minimum and maximum extent of a roof, we reverse (1), i.e. multiply $\hat{\mathbf{x}}_0$ by $\frac{1}{2}\underline{z}$ and add $\frac{1}{2}(z_{\min} + z_{\max})$.

5 Datasets & Benchmarks

Simulating the real-world requires noisy, sparse, and incomplete height maps. Ground truth requires a high-resolution, noise-free, complete height map; and to tackle long-tail geometric complexity illustrated in Fig. 3, the buildings should be more complex than existing datasets with a LoD 2.2 [4] (LoD 2.2 captures intricate roof details like dormers and gables, that LoD 2.0 overlooks).

To aid machine learning research, we curated a new dataset that not only serves our research needs, but future research needs in the field. We started with 16k compact and high-detail LoD 2.2 roof meshes from the city of Poznan, Poland [10]. To match our focus on complex roof structures (Fig. 3), we re-balanced the dataset by reducing the number of flat roof from 5k to 2k resulting in a new dataset of 13k buildings.

To address the long-tail corruption depicted in Fig. 3, collecting diverse corrupted height maps is essential. However, acquiring real-world corrupted data is challenging due to the difficulty in obtaining ground truths and identifying corrupted height maps. Therefore, we developed a procedure for synthesizing intentionally corrupted height maps suitable for training or benchmarking, as




Fig. 5: Examples of synthesizing corrupted height maps from the PoznanRD.




Fig. 6: Sparsity of the real-world USGS 3DEP LiDAR dataset [78, 79].

illustrated in Fig. 5. We first rasterize the roof mesh triangles into ground truth height maps. This process provides complete height maps without missing points and also reconciles real-world captured height maps and artist-generated meshes. Next, footprints are inferred from non-zero pixels in height maps. Then, we simulate real-world conditions by adding synthetic tree noise⁵, creating sparsity through random point removal, and injecting incompleteness.

In this paper, “s” denotes Sparsity (%), the ratio of randomly removed to total pixels in the footprint. Similarly, “i” signifies the Incompleteness (%), the proportion of pixels removed due to incompleteness in the footprint.

5.1 Tree Points

In analyzing height maps within a footprint, tree crowns may intrude, causing LiDAR to capture the tree canopy instead of the roof, see Fig. 2c. Thus, we simulate tree intrusion by planting virtual trees around the roof. We collected a real-world database with 1k tree height maps from [10]. Small trees are randomly placed outside the footprint, with random rotations and height adjustments. A “max” operation is used to combine the roof \mathbf{x} and tree height map \mathbf{x}_{tree} . Algorithm 3 in supplements has the details.

5.2 Incompleteness

To synthesize incompleteness, we employ multiple Gauss masks with various means and variances to indicate missing points. This approach emulates occlusions caused by more intricate objects with soft (Fig. 2a) or hard (Fig. 2b) boundaries, in the spirit of a Gaussian Mixture Model [66]. The rationale behind using Gauss masks is as follows: When multiple Gaussian masks are randomly positioned on the same side of a building, they can completely block one side of the roof point clouds, akin to the ray-tracing-free simulation of occlusion by [94], where the authors remove points distant from the camera. Conversely, when scattered and smaller Gaussian masks are utilized within the roof structure, they effectively simulate the occlusion caused by small roof features. Lastly, a larger variance of Gaussian simulates softer boundaries. The supplement details efficient mask generation for training in Algorithm 1, and mask creation with a specific portion of missing points in Algorithm 2.

⁵ Although we focus on tree noise for clarity, we also simulate global Gaussian noise and outlier noise in practice.

5.3 Benchmarks

The above height map synthesizing technique allows us to create benchmarks tailored to various research needs. Specifically, we can adjust different levels of sparsity, incompleteness, and the number of trees introduced. This makes it ideal for both training and evaluating machine learning models geared towards 3D reconstruction and other related fields. We have partitioned the dataset into a training split of 10k samples and a test split of 3k samples. Additionally, for tree-height maps, we provide 766 samples for training and 255 for testing.

6 Experiments

Section 6.1 demonstrates the effectiveness of RoofDiffusion and No-FP RoofDiffusion, trained without footprint, by quantitative evaluation on two datasets with ground truth: our PoznanRD and the BuildingNet [73]. Section 6.2 illustrates the ability of No-FP RoofDiffusion to recover footprints from corrupted height maps. Section 6.3 assesses the 3D mesh reconstruction quality for the City3D [29], comparing results with, and without using RoofDiffusion as a pre-processor. Section 6.4 showcases qualitative outcomes through tests on real-world scans in AHN3 [29], Dales3D [82], and USGS 3DEP LiDAR sampled over Cambridge, MA [79], and Wayne County, MI [78]. Section 6.5 discusses the limitations.

The supplement provides ablation of tree augmentation, sampling step analysis, impact of LiDAR scan pattern, additional qualitative results, and implementation details. Additionally, the supplement shows an extension of RoofDiffusion to be comparable to existing unguided depth completion methods on the KITTI dataset [81].

6.1 Quantitative Evaluation of Height Completion

We examine RoofDiffusion in the four relevant scenarios. Footprints (FP) are derived by converting the ground truth height map into a binary mask.

1. (FP) *Both* sparsity *and* incompleteness on PoznanRD in Tab. 1.
2. (FP) *Either* sparsity *or* incompleteness on BuildingNet [73] in Tab. 1.
3. (No-FP) *Both* sparsity *and* incompleteness in the PoznanRD in Tab. 2.
4. (No-FP) *Either* sparsity *or* incompleteness on BuildingNet [73] in Tab. 2.

PoznanRD Dataset. We conducted tests on 1k randomly selected data from test split of PoznanRD, injected with both global and outlier noise to simulate real-world cases. We follow [39, 91, 95] to emulate global noise by incorporating Gaussian noise into all the normalized data points, with $\sigma_{\text{global}} \sim [0, 0.05]$ for each height map. Also, outliers are introduced by randomly assigning nonzero pixels with random values, as suggested in [96], where the probability for occurrence of an outlier is set to 0.01 %. Additionally, we introduced tree noise in 30% of the cases, involving between 1 to 3 trees, following Sec. 5.1.

BuildingNet [73] contains compact and noise-free building meshes spanning various categories. These models are mostly sourced from 3D artists and often

Methods	PoznanRD						BuildingNet									
	s95 i30		s95 i80		s99 i30		s99 i80		s90 i80		s90 i90		s99 i0		s99.75 i0	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Linear	0.236	0.461	0.687	1.037	0.365	0.631	0.868	1.218	0.654	0.997	0.868	1.224	0.297	0.549	0.528	0.833
Spline	0.278	0.508	0.785	1.198	0.391	0.659	0.888	1.260	0.829	1.288	1.033	1.523	0.330	0.586	0.536	0.845
Nearest	0.288	0.541	0.691	1.089	0.424	0.734	0.866	1.271	0.634	1.044	0.856	1.287	0.347	0.662	0.582	0.961
IDW	0.239	0.449	0.648	0.984	0.377	0.619	0.827	1.172	0.573	0.897	0.808	1.159	0.309	0.540	0.537	0.822
P.M. Diff.	0.266	0.473	1.825	2.311	0.523	0.739	3.085	3.548	1.090	1.448	1.775	2.136	0.414	0.706	0.743	0.979
Ours	0.162	0.342	0.430	0.727	0.253	0.463	0.603	0.916	0.508	0.871	0.705	1.071	0.280	0.550	0.447	0.751

Table 1: Evaluation of height map completion on PoznanRD and BuildingNet [73] dataset, measured by Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) in meters. **Bold** represents best outcome.

include additional elements like virtual ground, cars, and trees in the scene. We picked 200 buildings that contain only the building structure, making it easier to infer footprints. We rasterized the meshes in BuildingNet [73] into height maps as ground truths. It is important to note that our tests were conducted using a model trained exclusively on PoznanRD.

(FP) Both sparsity and incompleteness. We evaluated combinations of sparsity at 95% and 99%, along with incompleteness at 30% and 80%, using the method shown in Sec. 5.2. Notably, our sparsity selection is according to USGS 3DEP LiDAR with Quality Level (QL) 2. Note that QL2 only guarantees at least 2 points per square meter and that most of the US territory LiDAR is QL 2 or lower. In fact, it is common to find roof data with extreme sparsity, as shown as yellow bars in Fig. 6. While we noticed several cases whose incompleteness exceeded 95%, the combination of this level of incompleteness with severe sparsity makes the analysis overly difficult and potentially meaningless. Therefore, we set the upper limit for incompleteness at 80%.

In this evaluation, we benchmark against interpolation techniques commonly employed in DEM inpainting, including linear, nearest-neighbor, Spline [32], and IDW [74] interpolation. In particular, we compare with the latest DSM inpainting technique [3] based on Perona-Malik Diffusion (P.M. Diff.) [59]. Table 1 shows that our approach consistently surpasses all these methods in height restoration, regardless of the varying degrees of sparsity and incompleteness.

(FP) Either sparsity or incompleteness. Here, we assess the handling of pure sparsity (s99 i0, s99.75 i0) and incompleteness (s90 i80, s90 i90) independently. We maintain the same noise injection settings as those used for the PoznanRD dataset but exclude tree noise. Table 1 shows RoofDiffusion outperforms most DEM interpolation methods and demonstrates generalizability to unseen BuildingNet [73] datasets. RoofDiffusion demonstrates a stronger performance advantage for the incompleteness task, suggesting particular effectiveness at restoring structural information.

(No-FP) Both sparsity and incompleteness. We compare to the state-of-the-art unguided depth completion methods, pNCNN [15] and CU-Net [85], that we selected for their exclusive use of depth data and hence close relation to height completion tasks. These algorithms were trained on our PoznanRD, using identical data augmentation settings, but no footprint information was used. We

Methods	PoznanRD								BuildingNet							
	s95 i30		s95 i60		s99 i30		s99 i60		s90 i60		s90 i80		s98 i0		s99 i0	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Linear	0.778	1.966	1.330	2.811	1.091	2.542	1.969	3.625	0.568	1.302	1.517	2.484	0.325	0.926	0.426	1.111
pNCNN	1.635	3.016	1.885	3.169	2.012	3.298	2.172	3.378	0.918	1.605	1.185	1.822	0.904	1.632	1.063	1.765
CU-Net	1.246	1.823	1.628	2.244	1.544	2.187	1.923	2.554	0.675	1.280	1.641	2.386	0.323	0.717	0.397	0.829
Ours (No-FP)	0.319	1.232	0.769	2.018	0.722	1.968	1.200	2.600	0.509	1.228	1.501	2.449	0.262	0.803	0.349	0.962

Table 2: Evaluation of height map completion *w/o footprint* on the PoznanRD and BuildingNet [73] datasets.

Methods	s99 i30		s99 i80		Methods	s95 i30	s95 i60	s99 i30	s99 i60
	RMSE	#Face	RMSE	#Face		RMSE	RMSE	RMSE	RMSE
City3D + IDW	0.352	124.40	0.708	105.27	Linear	82.18	68.54	73.81	51.04
City3D + P.M. Diff	0.577	89.04	3.016	97.47	pNCNN	68.97	66.32	64.80	63.94
City3D + Ours	0.244	82.72	0.534	80.12	CU-Net	82.12	73.88	75.92	69.51
					Ours (No-FP)	92.14	81.83	83.59	75.15

Table 3: Evaluation of point cloud pre-processors for City3D [29] on PoznanRD. City3D tested with GT point cloud achieved 0.104 RMSE and 82.68 average faces.

Table 4: Footprint predictions measured by IoU (%).

adhered to the default hyperparameter configurations specified for each method. To ensure a fair comparison, we also trained No-FP RoofDiffusion, a variant of RoofDiffusion that does not require footprints. We tested reconstruction accuracy on PoznanRD (s95 i30, s95 i60, s99 i30, and s99 i60) utilizing MAE and RMSE. The noise injection settings mirror the experiment in Tab. 1. Since predicting both footprint and roof simultaneously is a more difficult problem, the sparsity and incompleteness selection in Tab. 2 is lower than in Tab. 1.

Table 2 shows that RoofDiffusion achieves the most accurate reconstruction while pNCNN [15] fails to accurately predict the height values. CU-Net [85] tends to over-smooth heights. Linear interpolation struggles to recognise noise and incomplete regions.

(No-FP) *Either sparsity or incompleteness.* When tested on the BuildingNet [73] dataset, No-FP RoofDiffusion outperforms linear interpolation, CU-Net [85], and pNCNN [15] in scenarios with dominant sparsity (s98 i0, s99 i0) or incompleteness (s90 i60, s90 i80). While CU-Net’s smoother predictions reduce large errors and improve RMSE, notably at s90 i80, s98 i0, and s99 i0 in Tab. 2, CU-Net does not predict sharp height maps well, as indicated by a higher MAE.

6.2 Footprint Recovery

No-FP RoofDiffusion can help footprint recovery. Given a corrupted height map, we can predict the complete heights and infer a binary footprint by assigning a value of one to non-zero pixel values and zero otherwise. Tab. 4 displays Intersection over Union (IoU) between ground truth and the predicted footprint. RoofDiffusion yields the best the footprint recovery on the PoznanRD dataset compared to linear interpolation, CU-Net [85], and pNCNN [15].

6.3 Enhancement of 3D Reconstruction

This subsection demonstrates that using our model as the point cloud preprocessor can boost the accuracy of 3D building reconstruction. We preprocessed point clouds for City3D [29], a leading algorithm for converting point clouds to compact building meshes, by with our method, IDW [74], and Perona-Malik Diffusion [3]. We report the average RMSE of the distance from each ground truth point to the face of the closest reconstructed mesh⁶. To further demonstrate the advantages of using clean geometric details from RoofDiffusion in City3D [29] reconstructions, we evaluate the minimum number of polygonal faces required for the output while maintaining geometry accuracy. Here, we dissolve all the edges with a dihedral angle of less than 5 degrees. A lower count, while maintaining a similar level of reconstruction accuracy, is preferable as it indicates a more compact representation.

As shown in Tab. 3, City3D achieves the lowest point-to-plane distance when utilizing point clouds processed by RoofDiffusion. When using RoofDiffusion as a preprocessor for City3D, the average minimum face count is nearly identical to feeding ground truth point clouds to City3D.

Figure 8 illustrates that using RoofDiffusion as a preprocessing step significantly enhances the reconstruction quality in City3D on PoznanRD. RoofDiffusion enhances gable details (Fig. 8a), complex geometry recovery (Figs. 8b and 8c), and is robust to tree noise (Fig. 8d).

6.4 Qualitative Evaluation of Height Completion

To illustrate the capability of handling the gap between synthetic and real-world scanning conditions, we evaluated our method on datasets featuring real-world LiDAR scans. Due to the absence of ground truth, we adopt a visual comparison. Importantly, to maintain the real-world properties of the scans, we refrain from injecting any additional noise including global, outlier, and tree noise.

Real-world Noise. We artificially remove points from the rasterized height maps to evaluate the resilience of our model to such corruptions, even in the presence of real-world noise and unseen roof geometries in both AHN3 [29] and Dales3D [82]. The 99% sparsity (Figs. 7a and 7c) and 70% incompleteness (Figs. 7b and 7d) are synthesized. We observe that IDW [74] produces a mosaic effect, particularly in incomplete regions, and fails to capture the true geometry of roofs. Meanwhile, Perona-Malik Diffusion [3] yields smoother outcomes but tends to overlook sharp roof details. In contrast, RoofDiffusion effectively yields clean and faithful geometric reconstructions.

Real-world Incompleteness. We collect point clouds exhibiting incompleteness from real-world scans, sourced from the Wayne County dataset [78]. Notably, the point cloud is rated at QL2 [80], which is the most common quality of LiDAR scans across the United States. Figures 7e and 7f show that

⁶ We tested City3D without preprocessing but only a few meshes are reconstructed due to the severe sparsity.

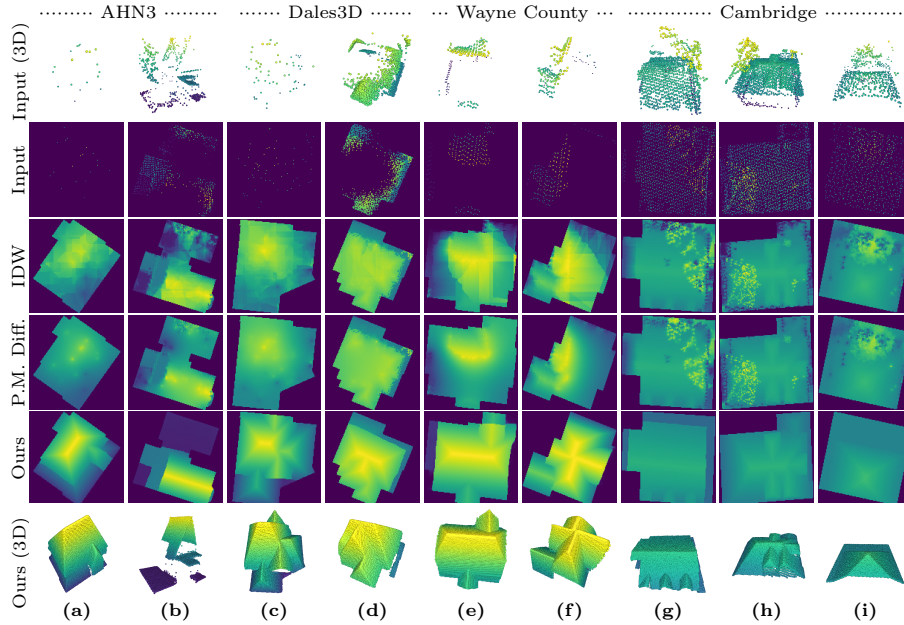


Fig. 7: Evaluation of the completion and denoising of real-world scans.

RoofDiffusion effectively leverages available points in conjunction with footprint information to reconstruct missing parts, even when these parts are distant from the existing points. In contrast, IDW [74] and Perona-Malik Diffusion [3] failed in such scenarios due to a lack of prior information about roof structures. Figures 9a and 9b also shown No-FP RoofDiffusion can recover incompleteness with cleaner and sharper features.

Real-world Tree Noise. We also gather point clouds containing tree noise from the Cambridge dataset [79]. As illustrated in Figs. 7g to 7i, RoofDiffusion demonstrates its efficacy in restoring roof geometry precisely while eliminating the tree points. Moreover, Figs. 9c and 9d highlight that No-FP RoofDiffusion can effectively extract and recover roof geometry from overlapping tree noise.

6.5 Limitations

We selected data containing tree noise from USGS 3DEP LiDAR sampled over Wayne County, MI [78], and Cambridge, MA [79], for analysis. In approximately 36% of these cases, tree noise was mistakenly reconstructed as building structures, see examples in the supplement. This error was due to the resemblance of tree canopies to architectural elements, such as dormers or chimneys, or due to severe occlusion by tree noise obscuring the underlying structures. Footprint-guided RoofDiffusion can encounter challenges due to severe misalignment between the footprint and the height map. This problem also occurs in [29]. Incorporating misalignment into the data augmentation pipeline and relaxing the footprint mask, m , during the denoising process warrants further investigation.




Fig. 8: 3D reconstruction using different height map pre-processors. RoofDiffusion matches the ground truth (GT) meshes generated by feeding GT height maps into City3D [29], except for (b). In (b) RoofDiffusion omits a dormer present in GT due to the lack of data points in that area.




Fig. 9: Evaluation of the completion and denoising in real-world scans without footprints.

Moreover, factors such as floor space, height from the ground, and location are useful indicators of roof geometry but are not yet used in the model. We leave the above as future work.

7 Conclusions

We introduced RoofDiffusion, a diffusion model for roof height map repair. RoofDiffusion is capable of repairing extreme sparsity, incompleteness, and noise. Additionally, we unveiled a comprehensive roof dataset containing over 13k complex geometric roof structures, with complete and a clean ground truth mesh and height map. This dataset can serve as a valuable asset for future long-tail research in remote sensing. To approximate real-world conditions, we also introduced methods for synthesizing tree noise and incomplete shapes. By generating intentionally-corrupted height maps from ground truths, these techniques not only mitigate the absence of ground truth in real-world scans but also facilitate data augmentation. They allow for benchmark customization with varying levels of height map corruption. Our experiments demonstrate the robustness of RoofDiffusion across datasets with real [29, 78, 79, 82] and synthetic scans [73], under diverse corruption conditions. RoofDiffusion outperforms both non-learning [3, 32, 74] and learning-based methods [15, 85] and significantly improves the accuracy of the 3D building reconstruction algorithm [29].

Acknowledgement

We thank Oleksandr Moskalenko and Dr. Zhe Jiang for generously providing computational resources on UF HyperGator. We also extend our appreciation to our colleagues at Meta Platforms for their assistance in processing data and offering valuable feedback throughout this project. We acknowledge the map services and data available from the U.S. Geological Survey, National Geospatial Program, and the Spatial Information System of the City of Poznań, Poland. Additionally, we are grateful for the funding provided by the UF Graduate School Preeminence Award and the UF Informatics Institute Fellowship.

References

1. Anciukevičius, T., Xu, Z., Fisher, M., Henderson, P., Bilen, H., Mitra, N.J., Guerrero, P.: RenderDiffusion: Image diffusion for 3D reconstruction, inpainting and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12608–12618 (2023) [2](#)
2. Bagheri, H., Schmitt, M., Zhu, X.: Fusion of multi-sensor-derived heights and osm-derived building footprints for urban 3d reconstruction. ISPRS International Journal of Geo-Information **8**(4), 193 (2019) [3](#)
3. Biasutti, P., Aujol, J.F., Brédif, M., Bugeau, A.: Diffusion and inpainting of reflectance and height LiDAR orthoimages. Computer Vision and Image Understanding **179**, 31–40 (2019) [3](#), [10](#), [12](#), [13](#), [14](#), [31](#), [33](#)
4. Biljecki, F., Ledoux, H., Stoter, J.: An improved LOD specification for 3D building models. Computers, Environment and Urban Systems **59**, 25–37 (2016) [3](#), [7](#), [23](#)
5. Cappelle, C., El Najjar, M.E., Charpillat, F., Pomorski, D.: Virtual 3D city model for navigation in urban areas. Journal of Intelligent & Robotic Systems **66**, 377–399 (2012) [2](#)
6. Chen, M., Hu, Q., Yu, Z., Thomas, H., Feng, A., Hou, Y., McCullough, K., Ren, F., Soibelman, L.: STPLS3D: A large-scale synthetic and real aerial photogrammetry 3d point cloud dataset. In: Proceedings of the British Machine Vision Conference (2022) [4](#), [26](#)
7. Chodosh, N., Wang, C., Lucey, S.: Deep convolutional compressed sensing for LiDAR depth completion. In: Proceedings of the Asian Conference on Computer Vision. pp. 499–513. Springer (2019) [30](#)
8. Conti, A., Poggi, M., Mattoccia, S.: Sparsity agnostic depth completion. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5871–5880 (2023) [31](#)
9. contributors, H.R.I.: 3D City Models of Helsinki. <https://hri.fi/data/dataset/helsingin-3d-kaupunkimalli> (2023), accessed: July 1, 2023 [5](#)
10. contributors, S.C.P.: Three-dimensional model of Poznan. <http://sip.poznan.pl/model3d/> (2023), accessed: July 1, 2023 [5](#), [7](#), [8](#), [23](#)
11. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021) [4](#), [27](#)
12. Dong, G., Chen, F., Ren, P.: Filling SRTM void data via conditional adversarial networks. In: IEEE International Geoscience and Remote Sensing Symposium. pp. 7441–7443. IEEE (2018) [4](#)

13. Dong, G., Huang, W., Smith, W.A., Ren, P.: A shadow constrained conditional generative adversarial net for SRTM data restoration. *Remote Sensing of Environment* **237**, 111602 (2020) [4](#)
14. Dong, Q., Cao, C., Fu, Y.: Incremental transformer structure enhanced image inpainting with masking positional encoding. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11358–11368 (2022) [2](#)
15. Eldesokey, A., Felsberg, M., Holmquist, K., Persson, M.: Uncertainty-aware CNNs for depth completion: Uncertainty from beginning to end. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12014–12023 (2020) [3](#), [4](#), [10](#), [11](#), [14](#), [30](#), [31](#), [37](#)
16. Eldesokey, A., Felsberg, M., Khan, F.S.: Propagating confidences through CNNs for sparse data regression. In: *Proceedings of the British Machine Vision Conference* (2018) [4](#), [30](#)
17. Gao, S., Liu, X., Zeng, B., Xu, S., Li, Y., Luo, X., Liu, J., Zhen, X., Zhang, B.: Implicit diffusion models for continuous super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10021–10030 (2023) [4](#)
18. Gavriil, K., Muntingh, G., Barrowclough, O.J.: Void filling of digital elevation models with deep generative models. *IEEE Geoscience and Remote Sensing Letters* **16**(10), 1645–1649 (2019) [4](#)
19. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014) [4](#)
20. Grohman, G., Kroenung, G., Strebeck, J., et al.: Filling SRTM voids: The Delta surface fill method. *Photogrammetric Engineering and Remote Sensing* **72**(3), 213–216 (2006) [4](#)
21. Gu, S., Zuo, W., Xie, Q., Meng, D., Feng, X., Zhang, L.: Convolutional sparse coding for image super-resolution. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1823–1831 (2015) [2](#)
22. Guizilini, V., Ambrus, R., Burgard, W., Gaidon, A.: Sparse auxiliary networks for unified monocular depth prediction and completion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11078–11088 (2021) [31](#)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016) [27](#)
24. He, Y., Liao, W., Hong, H., Huang, X.: High-precision single building model reconstruction based on the registration between osm and dsm from satellite stereos. *Remote Sensing* **15**(5), 1443 (2023) [3](#)
25. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020) [4](#), [6](#)
26. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research* **23**(1), 2249–2281 (2022) [4](#)
27. Holzmann, T., Maurer, M., Fraundorfer, F., Bischof, H.: Semantically aware urban 3D reconstruction with plane-based regularization. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 468–483 (2018) [2](#)
28. Hu, M., Wang, S., Li, B., Ning, S., Fan, L., Gong, X.: PENet: Towards precise and efficient image guided depth completion. In: *Proceedings of the IEEE Inter-*

- national Conference on Robotics and Automation. pp. 13656–13662. IEEE (2021) [31](#)
29. Huang, J., Stoter, J., Peters, R., Nan, L.: City3D: Large-scale building reconstruction from airborne LiDAR point clouds. *Remote Sensing* **14**(9), 2254 (2022) [2](#), [3](#), [4](#), [9](#), [11](#), [12](#), [13](#), [14](#), [26](#), [29](#), [31](#), [33](#)
 30. Huang, Z., Fan, J., Cheng, S., Yi, S., Wang, X., Li, H.: HMS-Net: Hierarchical multi-scale sparsity-invariant network for sparse depth completion. *IEEE Transactions on Image Processing* **29**, 3429–3441 (2019) [4](#), [30](#)
 31. Jaritz, M., De Charette, R., Wirbel, E., Perrotton, X., Nashashibi, F.: Sparse and dense data with CNNs: Depth completion and semantic segmentation. In: 2018 International Conference on 3D Vision. pp. 52–60. IEEE (2018) [4](#)
 32. Keys, R.: Cubic convolution interpolation for digital image processing. *IEEE transactions on acoustics, speech, and signal processing* **29**(6), 1153–1160 (1981) [3](#), [4](#), [10](#), [14](#), [31](#), [33](#)
 33. Kolbe, T.H., Gröger, G., Plümer, L.: CityGML: Interoperable access to 3D city models. In: *Geo-information for disaster management*, pp. 883–899. Springer (2005) [2](#)
 34. Ku, J., Harakeh, A., Waslander, S.L.: In defense of classical image processing: Fast depth completion on the cpu. In: 2018 15th Conference on Computer and Robot Vision. pp. 16–22. IEEE (2018) [30](#)
 35. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Deep Laplacian pyramid networks for fast and accurate super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 624–632 (2017) [2](#)
 36. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436–444 (2015) [4](#)
 37. Lei, J., Tang, J., Jia, K.: RGBD2: Generative scene synthesis via incremental view inpainting using RGBD diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8422–8434 (2023) [4](#)
 38. Li, B., Xue, K., Liu, B., Lai, Y.K.: BBDM: Image-to-image translation with Brownian bridge diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1952–1961 (2023) [4](#)
 39. Li, L., Song, N., Sun, F., Liu, X., Wang, R., Yao, J., Cao, S.: Point2Roof: End-to-end 3D building roof modeling from airborne LiDAR point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing* **193**, 17–28 (2022) [9](#)
 40. 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 Sensing of Environment* **269**, 112818 (2022) [4](#)
 41. Li, W., Lin, Z., Zhou, K., Qi, L., Wang, Y., Jia, J.: MAT: Mask-aware transformer for large hole image inpainting. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10758–10768 (2022) [2](#)
 42. Lin, L., Liu, Y., Hu, Y., Yan, X., Xie, K., Huang, H.: Capturing, reconstructing, and simulating: the UrbanScene3D dataset. In: *European Conference on Computer Vision*. pp. 93–109. Springer (2022) [4](#), [26](#)
 43. Liu, Q., Tan, Z., Chen, D., Chu, Q., Dai, X., Chen, Y., Liu, M., Yuan, L., Yu, N.: Reduce information loss in transformers for pluralistic image inpainting. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11347–11357 (2022) [2](#)
 44. Liu, T., Cheng, J., Tan, S.: Spectral Bayesian uncertainty for image super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18166–18175 (2023) [2](#)

45. Liu, X.: Three-dimensional visualized urban landscape planning and design based on virtual reality technology. *IEEE access* **8**, 149510–149521 (2020) [2](#)
46. Lu, K., Barnes, N., Anwar, S., Zheng, L.: From depth what can you see? depth completion via auxiliary image reconstruction. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11306–11315 (2020) [4](#)
47. Luedeling, E., Siebert, S., Buerkert, A.: Filling the voids in the SRTM elevation model — a TIN-based Delta surface approach. *ISPRS Journal of Photogrammetry and Remote Sensing* **62**(4), 283–294 (2007) [4](#)
48. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Re-Paint: Inpainting using denoising diffusion probabilistic models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11461–11471 (2022) [2](#), [4](#)
49. Ma, F., Cavalheiro, G.V., Karaman, S.: Self-supervised sparse-to-dense: Self-supervised depth completion from LiDAR and monocular camera. In: *2019 International Conference on Robotics and Automation (ICRA)*. pp. 3288–3295. *IEEE* (2019) [30](#)
50. McBratney, A., Webster, R.: Choosing functions for semi-variograms of soil properties and fitting them to sampling estimates. *Journal of soil Science* **37**(4), 617–639 (1986) [4](#)
51. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: *Proceedings of the International Conference on Learning Representations* (2021) [4](#)
52. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014) [4](#)
53. Nan, L., Wonka, P.: PolyFit: Polygonal surface reconstruction from point clouds. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2353–2361 (2017) [2](#)
54. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In: *International Conference on Machine Learning*. p. 16784–16804 (2022) [4](#)
55. Ntinas, G.K., Shen, X., Wang, Y., Zhang, G.: Evaluation of CFD turbulence models for simulating external airflow around varied building roof with wind tunnel experiment. In: *Building Simulation*. vol. 11, pp. 115–123. Springer (2018) [2](#)
56. OpenStreetMap contributors: Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org> (2017) [2](#), [3](#)
57. Park, J., Joo, K., Hu, Z., Liu, C.K., So Kweon, I.: Non-local spatial propagation network for depth completion. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII* 16. pp. 120–136. Springer (2020) [31](#)
58. Parmar, G., Kumar Singh, K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation. In: *ACM SIGGRAPH 2023 Conference Proceedings*. pp. 1–11 (2023) [4](#)
59. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on pattern analysis and machine intelligence* **12**(7), 629–639 (1990) [10](#), [33](#)
60. Qian, Y., Zhang, H., Furukawa, Y.: Roof-GAN: Learning to generate roof geometry and relations for residential houses. In: *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition. pp. 2796–2805 (2021) [5](#), [26](#)
61. Qiu, Z., Yue, L., Liu, X.: Void filling of digital elevation models with a terrain texture learning model based on generative adversarial networks. *Remote Sensing* **11**(23), 2829 (2019) [4](#)
 62. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with CLIP latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022) [4](#)
 63. Ran, W., Yuan, W., Shibasaki, R.: Few-shot depth completion using denoising diffusion probabilistic model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6558–6566 (2023) [4](#)
 64. Ren, J., Zhang, B., Wu, B., Huang, J., Fan, L., Ovsjanikov, M., Wonka, P.: Intuitive and efficient roof modeling for reconstruction and synthesis. *ACM Transactions on Graphics* **40** (2021) [5](#), [26](#)
 65. Reuter, H.I., Nelson, A., Jarvis, A.: An evaluation of void-filling interpolation methods for SRTM data. *International Journal of Geographical Information Science* **21**(9), 983–1008 (2007) [4](#)
 66. Reynolds, D.A., et al.: Gaussian mixture models. *Encyclopedia of biometrics* **741**(659-663) (2009) [8](#), [24](#)
 67. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [4](#)
 68. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015) [27](#)
 69. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-Booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023) [4](#)
 70. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–10 (2022) [3](#), [4](#), [7](#), [27](#)
 71. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022) [4](#)
 72. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4713–4726 (2022) [4](#)
 73. Selvaraju, P., Nabail, M., Loizou, M., Maslioukova, M., Averkiou, M., Andreou, A., Chaudhuri, S., Kalogerakis, E.: BuildingNet: Learning to label 3D buildings. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10397–10407 (2021) [5](#), [9](#), [10](#), [11](#), [14](#), [26](#)
 74. Shepard, D.: A two-dimensional interpolation function for irregularly-spaced data. In: Proceedings of the 1968 23rd ACM national conference. pp. 517–524 (1968) [3](#), [10](#), [12](#), [13](#), [14](#), [31](#), [33](#)
 75. Shih, P.T.Y., Huang, C.M.: The building shadow problem of airborne LiDAR. *The Photogrammetric Record* **24**(128), 372–385 (2009) [2](#)

76. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International conference on machine learning. pp. 2256–2265. PMLR (2015) **6**
77. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021) **4**
78. Survey, U.G.: MI Wayne County 2017 A17 345307 LAS. https://rockyweb.usgs.gov/vdelivery/Datasets/Staged/Elevation/LPC/Projects/MI_WayneCounty_2017_A17/ (2017), accessed: July. 1, 2023 **2, 8, 9, 12, 13, 14, 29, 31, 35, 37**
79. Survey, U.G.: MA CentralEastern 2021 B21 19TCG324693 LAS. https://rockyweb.usgs.gov/vdelivery/Datasets/Staged/Elevation/LPC/Projects/MA_CentralEastern_2021_B21/ (2019), accessed: July. 1, 2023 **2, 8, 9, 13, 14, 29, 31, 36, 38**
80. Survey, U.G.: Topographic Data Quality Levels (QLs). <https://www.usgs.gov/3d-elevation-program/topographic-data-quality-levels-qls> (2023), accessed: Oct. 19, 2023 **12**
81. Uhrig, J., Schneider, N., Schneider, L., Franke, U., Brox, T., Geiger, A.: Sparsity invariant CNNs. In: 2017 international conference on 3D Vision (3DV). pp. 11–20. IEEE (2017) **4, 9, 30, 31**
82. Varney, N., Asari, V.K., Graehling, Q.: DALES: A large-scale aerial LiDAR data set for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 186–187 (2020) **9, 12, 14, 29, 31, 34**
83. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017) **27**
84. Wang, R., Huang, S., Yang, H.: Building3D: A urban-scale dataset and benchmarks for learning roof structures from point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 20076–20086 (2023) **4, 26**
85. Wang, Y., Dai, Y., Liu, Q., Yang, P., Sun, J., Li, B.: CU-Net: LiDAR depth-only completion with coupled U-Net. *IEEE Robotics and Automation Letters* **7**(4), 11476–11483 (2022) **3, 4, 10, 11, 14, 30, 31, 37**
86. Wichmann, A., Agoub, A., Kada, M.: ROOFN3D: Deep learning training data for 3D building reconstruction. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **42**, 1191–1198 (2018) **5, 26**
87. Wu, C., Wang, D., Bai, Y., Mao, H., Li, Y., Shen, Q.: HSR-Diff: Hyperspectral image super-resolution via conditional diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7083–7093 (2023) **4**
88. Xie, S., Zhang, Z., Lin, Z., Hinz, T., Zhang, K.: SmartBrush: Text and shape guided object inpainting with diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22428–22437 (2023) **4**
89. Yan, L., Tang, X., Zhang, Y.: High accuracy interpolation of DEM using generative adversarial network. *Remote Sensing* **13**(4), 676 (2021) **4**
90. Yan, W.Y.: Airborne LiDAR data artifacts: What we know thus far. *IEEE Geoscience and Remote Sensing Magazine* (2023) **2**
91. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: PU-Net: Point cloud upsampling network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2790–2799 (2018) **9**

92. Yu, M., Lafarge, F.: Finding good configurations of planar primitives in unorganized point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6367–6376 (2022) [2](#)
93. Yu, Q., Chu, L., Wu, Q., Pei, L.: Grayscale and normal guided depth completion with a low-cost LiDAR. In: 2021 IEEE International Conference on Image Processing (ICIP). pp. 979–983. IEEE (2021) [4](#)
94. Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J., Zhou, J.: PoinTr: Diverse point cloud completion with geometry-aware transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12498–12507 (2021) [8](#)
95. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: PCN: Point completion network. In: 2018 international conference on 3D vision (3DV). pp. 728–737. IEEE (2018) [9](#)
96. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: LION: Latent point diffusion models for 3D shape generation. Advances in Neural Information Processing Systems (2022) [9](#)
97. Zhang, C., Shi, S., Ge, Y., Liu, H., Cui, W.: DEM void filling based on context attention generation model. ISPRS International Journal of Geo-Information **9**(12), 734 (2020) [4](#)
98. Zhang, K., Gool, L.V., Timofte, R.: Deep unfolding network for image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3217–3226 (2020) [2](#)
99. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023) [4](#)
100. Zhao, Y., Bai, L., Zhang, Z., Huang, X.: A surface geometry model for lidar depth completion. IEEE Robotics and Automation Letters **6**(3), 4457–4464 (2021) [30](#)
101. Zhao, Y., Elhousni, M., Zhang, Z., Huang, X.: Distance transform pooling neural network for lidar depth completion. IEEE Transactions on Neural Networks and Learning Systems (2021) [30](#)
102. Zhou, G., Song, B., Liang, P., Xu, J., Yue, T.: Voids filling of DEM with multiattention generative adversarial network model. Remote Sensing **14**(5), 1206 (2022) [4](#)
103. Zhu, D., Cheng, X., Zhang, F., Yao, X., Gao, Y., Liu, Y.: Spatial interpolation using conditional generative adversarial neural networks. International Journal of Geographical Information Science **34**(4), 735–758 (2020) [4](#)

Supplementary Material for RoofDiffusion

Kyle Shih-Huang Lo¹, Jörg Peters¹, and Eric Spellman²

¹ University of Florida, Gainesville FL 32611, USA

² Meta Platforms, Inc.

kyleshihuanglo@ufl.edu, jorg@cise.ufl.edu, espellman@meta.com

This supplement contains:

- Selection of Hyper-parameters (Sec. 1)
 - Roof Height Normalization (Sec. 1.1)
 - Notations (Sec. 1.2)
- PoznanRD Dataset (Sec. 2)
 - Data Balancing (Sec. 2.1)
 - Synthesizing Incompleteness (Sec. 2.2)
 - Synthesizing Tree Noise (Sec. 2.3)
 - Comparison to Existing Datasets (Sec. 2.4)
- Experiments (Sec. 3)
 - Implementation Details (Sec. 3.1)
 - Ablation of Tree Noise Augmentation (Sec. 3.2)
 - Sampling Step Analysis (Sec. 3.3)
 - LiDAR Scan Patterns (Sec. 3.4)
 - Variety of the Roofs (Sec. 3.5)
 - How Prone is the Model to Hallucination based on Footprints? (Sec. 3.6)
 - Failure Cases (Sec. 3.7)
 - Depth Completion on the KITTI Dataset (Sec. 3.8)
 - Additional Qualitative Experimental Results (Sec. 3.9)

1 Selection of Hyper-parameters

1.1 Roof Height Normalization

Our choice of 10-meters as the upper bound on the difference between maximum minus minimum roof height, for mapping into the range of $[-1, 1]$, is justified by the distribution of roof height differences in the PoznanRD, see Fig. 1: 99% of the height differences are less or equal to 10 meters, marked as the black dash line. For the buildings with height differences exceeding 10 meters, we scale them to 10 meters and record the scaling factor for recovering the true height.




Fig. 1: Distribution of height differences and roof counts.

1.2 Notations

As in the main document, in this supplement, “s” represents Sparsity (%), the ratio of randomly removed to total pixels in the footprint. Furthermore, “i” denotes the Incompleteness ratio (%), the proportion of pixels removed due to incompleteness in the footprint. Lastly, “t” signifies the Tree count used to synthesize tree noise.

2 PoznanRD Dataset

2.1 Data Balancing

We started with collecting 16k compact and high-detail LoD 2.2 [4] roof meshes from the city of Poznan in Poland [10]. To match our focus on complex roof structures, we rebalanced the dataset by excluding 3k flat roofs. This resulted in our PoznanRD comprising 13k buildings. The original longitude, latitude, height, and scaling for each roof are recorded in a Comma-Separated Values (CSV) file for easy access.

Algorithm 1 Incompleteness Mask (Training)

```

1: set  $\mathbf{m}_{\text{incomplete}}$  to ZEROS( $H, W$ )
2: for  $g = 1 \dots G$  do                                ▷ Index of Gauss distribution.
3:   for  $i = 1 \dots H$  do                                ▷ Image height.
4:     for  $j = 1 \dots W$  do                                ▷ Image width.
5:        $p \sim \mathcal{U}(0, 1)$ 
6:       if  $p_{\text{Gauss},g}(i, j) > p$  then
7:          $\mathbf{m}_{\text{incomplete}}(i, j) = 1$ 
8:       end if
9:     end for
10:  end for
11: end for
12:  $\mathbf{m}_{\text{incomplete}} = \mathbf{m} \odot \mathbf{m}_{\text{incomplete}}$   ▷ Set outside footprint to zero.

```

Algorithm 2 Incompleteness Mask (Benchmark)

```

1: set  $\mathbf{m}_{\text{incomplete}}$  to ZEROS( $H, W$ )
2: while SUM(M) <  $N_{\text{thres}}$  do
3:    $(i, j) \sim p_{\text{GMM}}(i, j)$ 
4:   if  $\mathbf{m}(i, j) == 1$  then                                ▷ Should be within footprint.
5:      $\mathbf{m}_{\text{incomplete}}(i, j) = 1$ 
6:   end if
7: end while

```

2.2 Synthesizing Incompleteness

We show the detail of generating a binary mask for synthesizing incompleteness, denoted as $\mathbf{m}_{\text{incomplete}}$. In this mask, a value of 1 indicates the pixels that will be removed, thereby yielding incompleteness, otherwise, 0. First, we employ a Gaussian Mixture Model (GMM) [66] to establish a probability density function,

$$p_{\text{GMM}}(i, j) := \sum_{g=1}^G \frac{1}{G} \underbrace{e^{-\frac{\|(i,j) - \boldsymbol{\mu}_g\|^2}{2\sigma_g^2}}}_{p_{\text{Gauss},g}}, \quad (1)$$

where i and j are the pixel coordinate in x and y direction of right-handed coordinate system, respectively, and

$$\mu_{g,x} \sim \mathcal{U}(1, H) \quad \mu_{g,y} \sim \mathcal{U}(1, W), \quad (2)$$

$$\sigma_{g,x} \sim \mathcal{U}(\sigma_{\min}, \sigma_{\max}), \quad \sigma_{g,y} \sim \mathcal{U}(\sigma_{\min}, \sigma_{\max}), \quad (3)$$

$$\boldsymbol{\mu}_g = (\mu_{g,x}, \mu_{g,y}), \quad (4)$$

$$\boldsymbol{\sigma}_g = (\sigma_{g,x}, \sigma_{g,y}), \quad (5)$$




Fig. 2: Examples of tree height map in PoznanRDdataset. (The examples are color mapped and resized for clear visualization. The height increases as the color becomes more yellow.)

Algorithm 3 Tree Noise Injection

```

1: Given  $\mathbf{x}_{\text{gt}} \sim q(\mathbf{x}_{\text{gt}})$  and  $\mathbf{x}_{\text{tree}} \sim q(\mathbf{x}_{\text{tree}})$ 
2:  $N_{\text{tree}} \sim \mathcal{U}_{\mathbb{Z}}(N_{\text{tree}}^{\min}, N_{\text{tree}}^{\max})$ 
3: for  $k = 1 \dots N_{\text{tree}}$  do
4:    $d \sim \mathcal{U}(1, 360)$  ▷ Tree rotation degree.
5:    $s_{xy} \sim \mathcal{U}(s_{xy}^{\min}, s_{xy}^{\max})$  ▷ Tree coverage area scaler.
6:    $s_z \sim \mathcal{U}(s_z^{\min}, s_z^{\max})$  ▷ Tree height scaler.
7:    $\mathbf{x}_{\text{tree}} = \text{ROTATE}(\mathbf{x}_{\text{tree}}, d)$ 
8:    $\mathbf{x}_{\text{tree}} = \text{RESIZE}(\mathbf{x}_{\text{tree}}, s_{xy})$ 
9:    $\mathbf{x}_{\text{tree}} = s_z \mathbf{x}_{\text{tree}}$ 
10:  repeat
11:     $(i, j) \sim \{(i, j) \mid \mathbf{m}(i, j) = 0\}$  ▷ Sample coordinates outside of footprint.
12:    set  $\mathbf{x}'_{\text{tree}}$  to ZEROS( $H, W$ )
13:    merge  $\mathbf{x}_{\text{tree}}$  into  $\mathbf{x}'_{\text{tree}}$  at center  $(i, j)$ 
14:     $\mathbf{x}, c_{\text{replaced}} = \text{MAX}(\mathbf{x}_{\text{gt}}, \mathbf{x}'_{\text{tree}})$  ▷ Get occluded height map and pixel count.
15:    until  $c_{\text{replaced}} > 0$  ▷ Enforce tree occlusion.
16:  end for

```

and G denotes the number of Gaussian distributions utilized in the construction of the GMM. Furthermore, $\mu_{g,x}$ and $\mu_{g,y}$ represent the mean in x and y direction, respectively, for g -th Gaussian distribution in the model. Similarly, $\sigma_{g,x}$ and $\sigma_{g,y}$ represent the standard deviation in x and y direction, respectively, for g -th Gaussian distribution. Also, σ_{\min} , σ_{\max} are hyper-parameters. H and W are the height and width of the image. Note that a greater value of G can yield a more intricate GMM, which, in turn, can synthesize a more complex shape of incompleteness.

During training, there is no requirement to generate an incomplete mask with a precise number of missing points. Therefore, for more efficient mask generation within the dataloader, we propose a method outlined in Algorithm 1.

For benchmarking, we can sample an incompleteness mask, $\mathbf{m}_{\text{incomplete}}$, with a specific number of pixels to be removed, N_{thres} , using Algorithm 2.

2.3 Synthesizing Tree Noise

Figure 2 shows examples of tree height maps in our dataset. Algorithm 3 outlines injecting tree noise into ground truth height maps, \mathbf{x}_{gt} , to generate a corrupted height map, \mathbf{x} .

Dataset	Roof Variety	Point Cloud	Mesh Property	Roof Only	#Roof (k)
RoofGAN [60]	Limit	-	Compact (Noise Free)	✓	0.5
Ren <i>et al.</i> [64]	Limit	-	Compact (Noise Free)	✓	3
RoofN3D [86]	Limit	Real-world Scan	Compact (Noise Free)	✓	118
UrbanScene3D [42]	High	Real-world Scan	Dense (Reconstructed)		-
STPLS3D [6]	High	Real-world Scan	Dense (Reconstructed)		-
City3D [29]	High	Real-world Scan	Coarse (Reconstructed)		20
Building3D [84]	High	Real-world Scan	Coarse (Reconstructed)	✓	160
BuildingNet [73]	High	Sample from Mesh	Compact (Noise Free)		2
PoznanRD (Ours)	High	Sample from Mesh	Compact (Noise Free)	✓	13

Table 1: Comparison of Building Datasets.

2.4 Comparison to Existing Datasets

Table 1 provides the comparison between our PoznanRD and the existing building or roof datasets. RoofGAN [60], RoofN3D [86], and the dataset [64] provide compact, noise-free roof meshes but offer only a limited range of roof types. Specifically, RoofGAN [60] encompasses 0.5k roofs, constructed with 2 to 5 hip roof primitives. The dataset [64] presents 3k more complex meshes compared to RoofGAN [60], yet it primarily features buildings constructed with hip-based primitives. Moreover, most roofs in [64] lack detailed structures like dormers, which are central to our focus. RoofN3D [86] is confined to pyramid, saddleback, and two-sided hip roofs.

UrbanScene3D [42] and STPLS3D [6] offer a diverse range of roof types but provide ground truth meshes with scan noise and dense triangles, reconstructed from large-coverage real-world scanning. City3D [29] and Building3D [84] provide datasets featuring coarser ground truths, utilizing plane partition-based reconstruction algorithms and artist-assisted refinement methods, respectively. These datasets can contribute to research in compact mesh reconstruction. However, the meshes, are reconstructed from real-world scans and may contain scan noise and algorithmic errors, which do not align with our requirement for clean, error-free data.

BuildingNet [73] provides a wide variety of compact meshes but of only 2k buildings. Furthermore, direct usage is often impractical due to the inclusion of non-relevant elements like humans, virtual ground, trees, cars, and landscaping. Although BuildingNet [73] provides the classifying labels, we have encountered several misclassifications that can adversely affect the accuracy of roof-only ground truth extraction.

3 Experiments

3.1 Implementation Details

Network Architecture. We followed Palette [70] by employing a U-Net [68] architecture with an attention mechanism [83] in its deeper layers to construct our conditional diffusion models. Specifically, the input size is $2 \times 128 \times 128$, one channel for the corrupted height map, \mathbf{x} , and another for the estimated height map at step t , $\hat{\mathbf{x}}_t$. The output is the predicted noise in $\hat{\mathbf{x}}_t$ and is of size $1 \times 128 \times 128$. Our network contains four down-sampling modules and four up-sampling modules. These modules operate at resolutions of 128×128 , 64×64 , 32×32 , and 16×16 , with channel dimensions set to 64, 128, 256, and 512, respectively. Each module includes two residual blocks [11, 23] and the attention mechanism [83] is integrated into the modules when the resolution reaches 32×32 and 16×16 .

Data Augmentation. We augment each height map by rotating 90, 180, and 270 degrees. Outlier noise occurs with a probability of 0.01%. Global noise is synthesized by sampling σ_{global} from a uniform distribution $\mathcal{U}(0, 0.05)$. For each pixel, a Gaussian distribution is constructed with the height as the mean and σ_{global} as the variance; subsequently, the height value is re-sampled using this distribution. We also augment the data with varying sparsity including 99, 98, 90, 80, 50, and 25%.

For synthesizing incompleteness, we follow Sec. 2.2 and use Algorithm 1 with five Gaussian distributions, $G = 5$, to synthesize various types of incompleteness. For each Gauss distribution, the minimum variance, σ_{min} , and maximum variance, σ_{max} , are set to 0 and 0.3, respectively.

The tree noise is injected according to Algorithm 3 using a 30% probability. We set the minimum and maximum tree count, $N_{\text{tree}}^{\text{min}}$ and $N_{\text{tree}}^{\text{max}}$, to 1 and 3, respectively. Furthermore, the tree coverage scaling parameters, $s_{\text{xy}}^{\text{min}}$ and $s_{\text{xy}}^{\text{max}}$, are set to 0.5 and 2.0. Lastly, s_z^{min} and s_z^{max} are set to 2.0 and 4.0 for height scaling.

Training is conducted on 8 NVIDIA Ampere A100 GPUs with a batch size of 512 for 260 epochs. This took approximately 2.5 days. The settings for variance scheduling and the use of exponential moving averages are adopted from [70]. The learning rate is set at 7×10^{-5} , and a warm-up learning rate is employed for the first 10k steps, starting with a factor of 0.2.

Training No-FP RoofDiffusion. During data augmentation, we retain the entire tree shape and do not remove tree noise outside the footprint. While RoofDiffusion uses a footprint image, \mathbf{m} , that is 0 for inside pixels and 1 outside, \mathbf{m} is an array of 1s when training No-FP RoofDiffusion. This guarantees that the model possesses no prior footprint information.

Methods	Tree Counts			
	0	1	3	5
w/o tree aug.	0.203	0.363	0.523	0.768
w/ tree aug.	0.208	0.278	0.356	0.504

Table 2: Evaluation of tree noise resilience in training with tree augmentation (RMSE in meters).

Test Set	Steps					
	60	125	250	500	1000	2000
Easy (s95 i30)	0.444	0.343	0.339	0.337	0.338	0.344
Hard (s99 i80)	1.123	0.940	0.919	0.893	0.899	0.922

Table 3: Impact of inference steps on height completion quality (RMSE in meters). **Bold** denotes the most efficient in terms of achieving acceptable quality with the fewest steps.

3.2 Ablation of Tree Noise Augmentation

We investigated the impact of tree augmentation during model training. We constructed multiple test sets featuring varying numbers of tree intrusions per building. Specifically, tree count specifies the number of trees that will 100% appear in each building. Table 2 compares height map restoration quality for models trained with and without tree augmentation. The results clearly indicate that tree augmentation during training effectively enhances the resilience of the model to tree noise.

3.3 Sampling Step Analysis

We analyze the relationship between height completion quality and the number of inference steps. The model was trained using a fixed 2k steps, and we tested it by changing the number of inference steps, as detailed in Tab. 3. For easier datasets with 95% sparsity and 30% incompleteness, reducing the inference steps to 250 resulted in only a minor loss of quality. However, for harder datasets, 500 inference steps were needed for satisfactory results. We observed that for more difficult tasks more steps are needed. Also, we attempted to train the model directly with fewer steps such as 1k. This led to failure to converge.




Fig. 3: RoofDiffusion reconstruction of height maps corrupted by scan line strip patterns.




Fig. 4: Samples in the PoznanRD (Poznan Roof Dataset).

3.4 LiDAR Scan Patterns

Our observations indicate that simulating sparse LiDAR points through random point removal yields results without noticeable gaps, closely resembling real-world scans. We found that most real-world point clouds [29, 78, 79, 82] appear similar to those obtained via random sampling. While a few examples exhibit line strip patterns, these do not hinder the ability of RoofDiffusion to restore roofs. Figure 3 showcases the corrupted height maps featuring line strip patterns, and RoofDiffusion effectively reconstructs complete, noise-free height maps from these.

3.5 Roof Variety

Figure 4 illustrates the variety of roof types in the PoznanRD dataset. Constructions of valid roofs of different types and with a variety of features have been shown for RoofDiffusion and No-FP RoofDiffusion: gables (Fig. 10e), hips (Figs. 8a and 8h), shads (Figs. 8e and 10d), leans (Figs. 7b and 7i), and flats (Fig. 10a).

3.6 How Prone is the Model to Hallucination based on Footprints?

With only the footprint available, and no height data provided, we conducted tests using a height map set to all zeros, and varied the footprints. The output was predominantly zeros or values less than one meter. This suggests that the model uses footprints to enhance height prediction but does not by itself cause the diffusion model to hallucinate roof shapes.

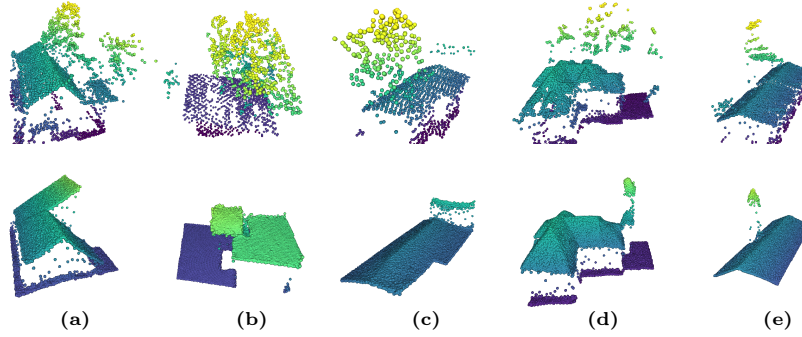


Fig. 5: Cases of tree noise misinterpreted as roof structures.

Methods	iRMSE	iMAE	RMSE	MAE
SparseConvfs [81]	4.94	1.78	1601.33	481.27
IP_Basic [34]	3.78	1.29	1288.46	302.60
ADNN [7]	59.39	3.19	1325.37	439.48
Nconv_CNN [16]	4.67	1.52	1268.22	360.28
S2D (depth-only) [49]	3.21	1.35	954.36	288.64
HMS-Net [30]	2.93	1.14	937.27	258.48
pNCNN [15]	3.37	1.05	960.05	251.77
Physical_Surface [100]	3.76	1.21	1239.84	298.30
DTP [101]	2.94	1.07	937.27	247.81
CU-Net [85]	2.69	1.04	917.76	244.36
Ours (KITTI)	5.69	2.66	1641.41	491.61
Ours + M.O. (KITTI)	4.84	1.84	1631.98	448.97

Table 4: Evaluation of depth completion methods on KITTI dataset [81]. M.O. stands for applying mean shift offset.

3.7 Failure Cases

Figure 5 provides some examples of failure cases where RoofDiffusion misclassifies tree noise as roof structures. Figures 5a and 5b illustrate erroneous reconstructions of unlikely slope and flat planes, respectively, caused by extensive occlusion from tree points. Figure 5c incorrectly identifies tree points as thin wall structures. Figures 5d and 5e mistakenly interpret tree noise as non-existent chimney structures.

3.8 Depth Completion on the KITTI Dataset

We extend RoofDiffusion to address the unguided depth completion task on the KITTI dataset [81]. Given the original KITTI depth map resolution of 352×1216 , training with such size is impractical due to the excessive processing time required by the diffusion model. Therefore, we downsampled the depth maps to

a resolution of 176×608 and trained our model using randomly cropped sections of size 176×304 . We trained on the official training split with 86k samples. We conducted tests on the 1k validation depth maps [81], where the predicted results with the size of 176×608 are upsampled back to 352×1216 for evaluation.

We compared to the results of state-of-the-art and representative unguided depth completion methods, reported in [85]. Table 4 combines these results with those newly reported by our method. We observed a mean shift phenomenon in our method. To address this, we aligned the predicted depth map with the input sparse depth map by offsetting it with the difference between their respective means. This adjustment resulted in a significant performance boost.

Nevertheless, our results do not improve on prior methods, likely due to the differences in KITTI depth completion v.s. roof completion. Available pixels in the KITTI dataset [81] are more uniformly distributed, and regional incompleteness size is smaller compared to roof completion tasks. Consequently, KITTI depth completion resembles an interpolation task more than inpainting for roof completion.

From this we conclude: diffusion models are more adept at tasks requiring the inpainting of large missing regions, while for smaller interpolation tasks, their performance is on par with other methods. Our conclusion is supported by two observations. First, in our PoznanRD benchmark (see Table 2 in the main paper), the performance margin of RoofDiffusion over existing methods [3, 32, 74] is greater in tasks of restoring incompleteness compared to sparsity. Secondly, SpAgNet [8], a diffusion model-based depth completion method, does not outperform existing methods [15, 22, 28, 57] when completing sparsely scanned data with 64 lines. However, as the scan lines become sparser, 32, 16, 8, and 4 lines, SpAgNet [8] outperforms others.

3.9 Additional Qualitative Experimental Results

RoofDiffusion.

- Figure 6 showcases examples of applying various height map pre-processors for 3D reconstruction algorithm, City3D [29], on our PoznanRD.
- Figure 7 illustrates the evaluation on AHN3 dataset [29].
- Figure 8 shows the evaluation on Dales3D dataset [82].
- Figure 9 provides the results on USGS 3DEP LiDAR sampled over Wayne County, MI [78].
- Figure 10 provides the results on USGS 3DEP LiDAR sampled over Cambridge, MA [79].

No-FP RoofDiffusion.³

- Figure 11 provides the results on USGS 3DEP LiDAR sampled over Wayne County, MI [78].
- Figure 12 provides the results on USGS 3DEP LiDAR sampled over Cambridge, MA [79].

³ Since the point clouds in AHN3 [29] and Dales3D [82] are already cropped by footprints, we do not include these two datasets in our testing.




Fig. 6: 3D reconstruction using different height map pre-processors. RoofDiffusion closely matches the ground truth (GT).

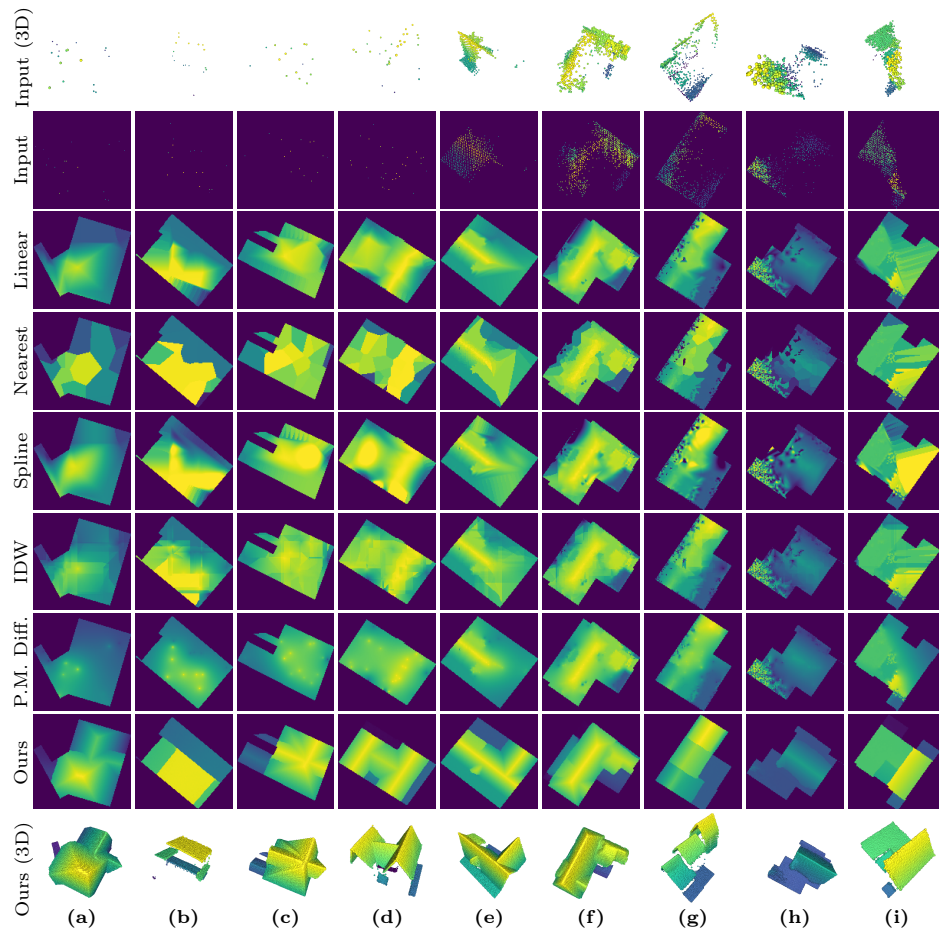


Fig. 7: Evaluation of completion and denoising on the AHN3 dataset [29]. Compared with Linear, Nearest, Spline [32], Inverse Distance Weighting (IDW) [74], and Perona-Malik Diffusion (P.M. Diff.) [3, 59] interpolation methods.

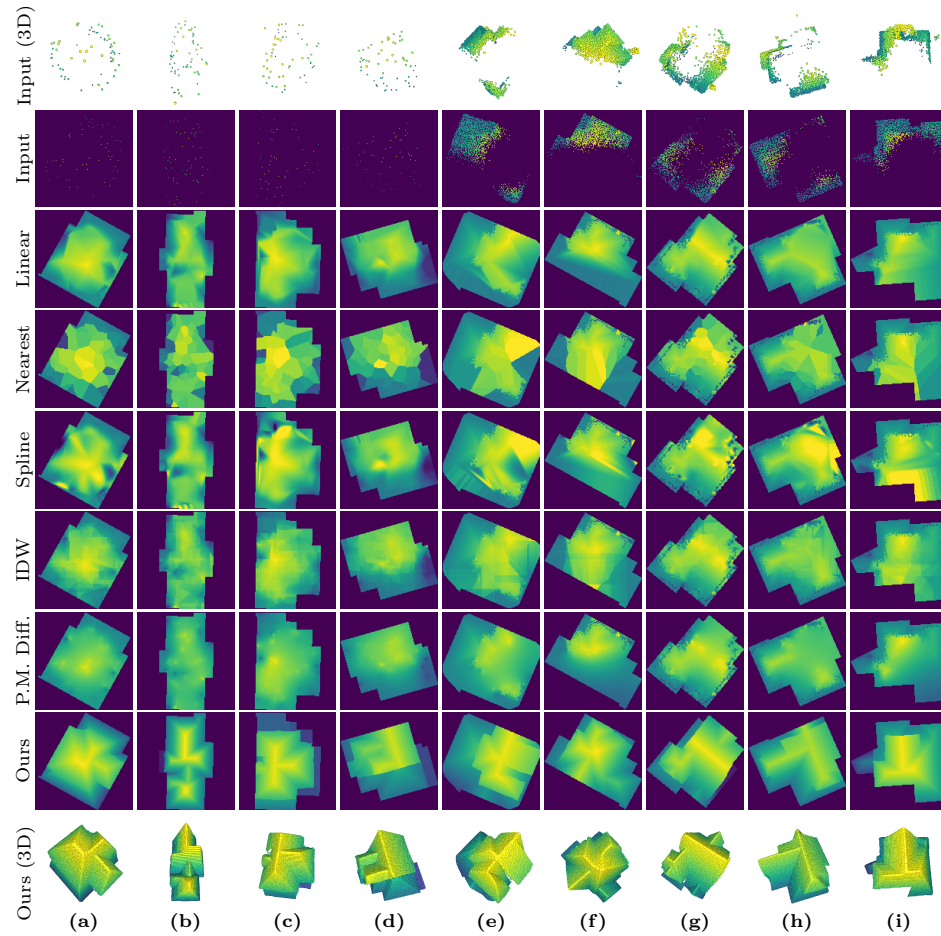


Fig. 8: Evaluation of the completion and denoising on the Dales3D dataset [82].




Fig. 9: Evaluation of the completion and denoising on USGS 3DEP LiDAR data sampled over Wayne County, MI [78].

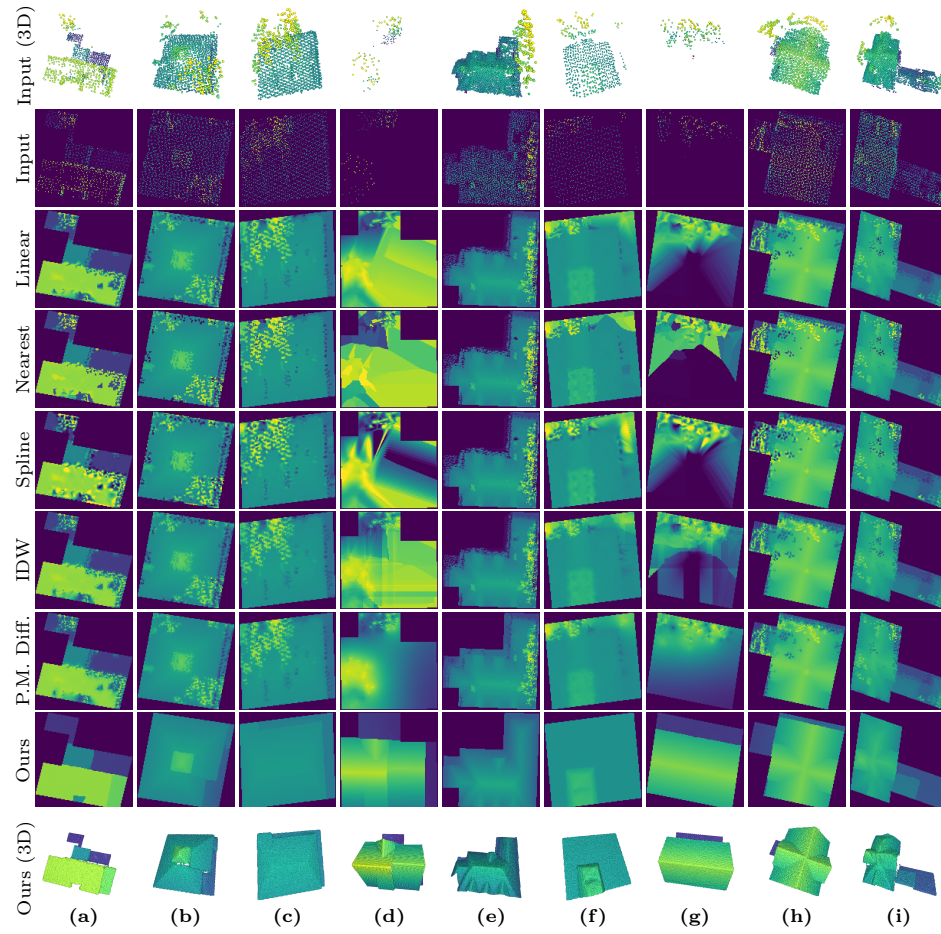


Fig. 10: Evaluation of the completion and denoising on USGS 3DEP LiDAR data sampled over Cambridge, MA [79].

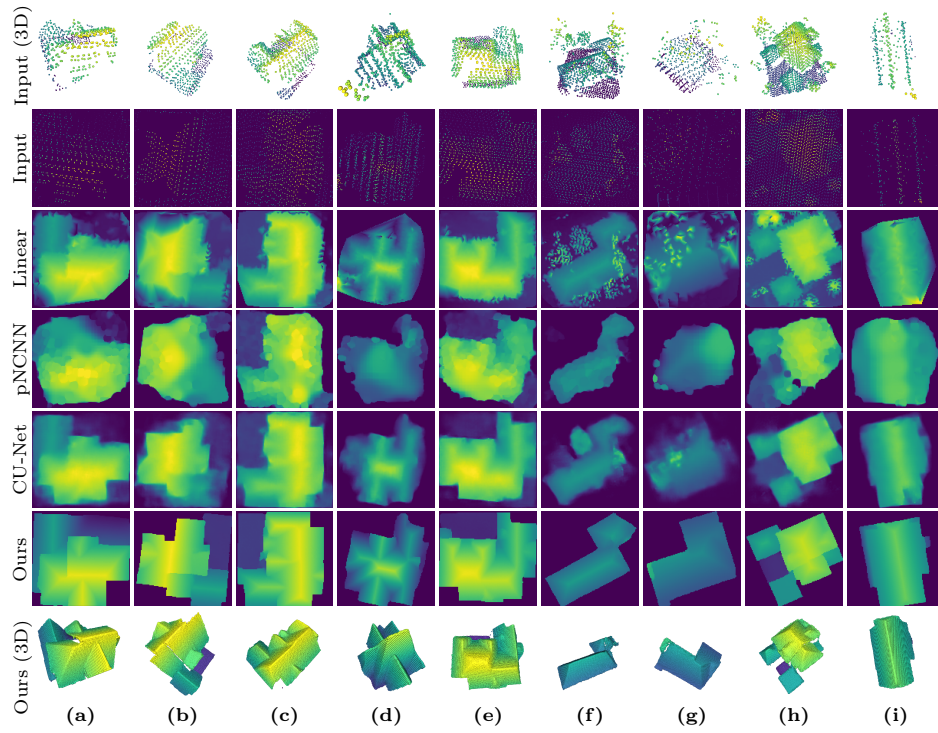


Fig. 11: Evaluation of the completion and denoising for No-FP RoofDiffusion on USGS 3DEP LiDAR data sampled over Wayne County, MI [78]. Comparison with linear, pNCNN [15], and CU-Net [85].




Fig. 12: Evaluation of the completion and denoising for No-FP RoofDiffusion on USGS 3DEP LiDAR data sampled over Cambridge, MA [79].