

From Orbit to Ground: Generative City Photogrammetry from Extreme Off-Nadir Satellite Images

Fei Yu^{1,2,4*}, Yu Liu^{2*}, Luyang Tang², Mingchao Sun², Zengye Ge², Rui Bu³, Yuchao Jin¹,
Haisen Zhao⁴, He Sun¹, Yangyan Li³, Mu Xu^{2†}, Wenzheng Chen^{1†}, Baoquan Chen^{1†}
¹Peking University ²AMAP ³Ant Group ⁴Shandong University




Figure 1. **City-Scale 3D Reconstruction from Satellite Imagery.** We reconstruct a 4 km² real-world urban region from 11 sparse-view satellite images captured from orbit that contain extremely limited parallax. The resulting 3D model, featuring crisp geometry and photorealistic appearance, enables **extreme viewpoint extrapolation**, supporting high-fidelity, close-range rendering from ground-level viewpoints. Please zoom in for details.

Abstract

City-scale 3D reconstruction from satellite imagery presents the challenge of **extreme viewpoint extrapolation**, where our goal is to synthesize ground-level novel views from sparse orbital images with minimal parallax. This requires inferring nearly 90° viewpoint gaps from image sources with severely foreshortened facades and flawed textures, causing state-of-the-art reconstruction engines such as NeRF and 3DGS to fail.

To address this problem, we propose two design choices tailored for city structures and satellite inputs. First, we model city geometry as a 2.5D height map, implemented as a Z-Monotonic signed distance field (SDF) that matches urban building layouts from top-down viewpoints. This stabilizes geometry optimization under sparse, off-nadir satellite views and yields a watertight mesh with crisp roofs and

clean, vertically extruded facades. Second, we paint the mesh appearance from satellite images via differentiable rendering techniques. While the satellite inputs may contain long-range, blurry captures, we further train a generative texture restoration network to enhance the appearance, recovering high-frequency, plausible texture details from degraded inputs.

Our method’s scalability and robustness are demonstrated through extensive experiments on large-scale urban reconstruction. For example in Fig. 1, we reconstruct a 4 km² real-world region from only a few satellite images, achieving state-of-the-art performance in synthesizing photorealistic ground views. The resulting models are not only visually compelling but also serve as high-fidelity, application-ready assets for downstream tasks like urban planning and simulation. Project page can be found at <https://pku-vcl-geometry.github.io/Orbit2Ground/>.

* Joint first authors.

† Corresponding authors.

1. Introduction

The past decade has witnessed remarkable advances in 3D reconstruction. Neural representations, particularly Neural Radiance Fields (NeRFs) and 3D Gaussian Splatting (3DGS) [22, 39], have brought photorealistic novel view synthesis (NVS) closer to reality than ever before. While these methods have fundamentally revolutionized object- and street-level modeling [18, 29, 32, 33], scaling this fidelity to expansive urban environments remains a critical challenge. Reconstructing a large city faces a fundamental data acquisition problem: ground-level capture—whether crowdsourced photographs or cameras on vehicles or drones—is logistically complex and prohibitively expensive for continuous, citywide coverage. In contrast, satellite imagery provides massive, readily available city coverage, positioning it as an economical and scalable data source for city-scale 3D modeling.

However, reconstructing a high-fidelity city model from satellite imagery presents the **extreme viewpoint extrapolation** challenge. The satellite imagery is captured from top-down views with extreme off-nadir angles and minimal parallax, while our goal is to synthesize ground-level novel views, resulting in an almost 90° viewpoint gap between source and target. In satellite capture, as illustrated in Fig. 2, building facades suffer severe foreshortening, and vertical structures lack the viewpoint diversity needed to resolve geometry. Long-range atmospheric distortion and sensor limits further degrade textures. These combined conditions—sparse, minimal-parallax, and blurred inputs—fundamentally violate the dense-parallax and sharp-photometry assumptions required by NeRF and 3DGS, causing these methods to collapse.

To overcome this challenge, we adopt a two-stage design tailored for city photogrammetry from satellite inputs. In the first stage, our focus is on **geometry fidelity**: we leverage the inherent 2.5D structure of urban layouts from top-down views and propose a Z-Monotonic signed distance field (SDF) representation. This strong structural prior enables 2.5D mesh extraction via differentiable iso-surfacing [52] and yields coherent, watertight geometry with precise roofs and vertically extruded facades. Optimizing this SDF against off-nadir inputs preserves as much satellite-derived geometric information as possible, ensuring that even under extreme viewpoint shifts, the reconstructed structure remains faithful.

On this robust geometric foundation, the second stage addresses **appearance fidelity**. Naive back-projection texturing from satellite views produces blurred and distorted facades (Fig. 5). To overcome these artifacts, we adapt the powerful generative prior of the FLUX foundation model [25] into a high-fidelity restoration network. Novel close-range renders—degraded by projection gaps—are fed into this network, which synthesizes photorealistic appear-

Typical Viewpoint Extrapolation in Street View Setting




Figure 2. Unlike dense street views, satellite images are sparse and captured with extreme off-nadir angles. This leads to a severe deficiency in parallax for vertical structures. Yellow points represent 3D locations determined by MVS, whereas satellite images only recover ground and roof surfaces.

ances that drive the final texture optimization.

Our key contribution is a holistic approach that decouples this ill-posed problem into two complementary stages: robust geometry regularization and generative appearance refinement. By first establishing a stable geometric scaffold and then “painting” it with plausible, high-frequency details, our method effectively bridges the satellite-to-ground viewpoint gap where prior methods fail. We demonstrate through extensive experiments that this strategy yields state-of-the-art (SOTA) results, producing scalable and high-fidelity city photogrammetry from sparse satellite imagery.

2. Related Work

2.1. Large-Scale 3D Scene Reconstruction

Large-scale scene reconstruction, such as entire cities, has been a key topic in computer vision and graphics. Pioneering works [4, 15, 50, 51] based on Structure-from-Motion (SfM) and Multi-View Stereo (MVS) enabled city-scale reconstruction from massive, unordered photo collections. Recently, neural scene representations, exemplified by NeRF [39] and 3DGS [22], have been extended to large-scale scenes, achieving photorealistic city-scale reconstruction [18, 33]. Block-NeRF [55], Mega-NeRF [57] and Switch-NeRF [71] divide the scene into blocks, each modeled by a separate NeRF. Methods [46, 58] incorporate multi-modal data such as LiDAR and 2D optical flow to improve the reconstruction quality. Subsequent works [54, 60, 68] combine explicit feature grids with small MLPs to improve efficiency. VastGaussian [29] first demonstrated the viability of 3DGS, tackling challenges like appearance variation across large areas. Subsequent efforts [23, 32, 47] introduce Level-of-Detail (LoD) techniques for efficient multi-scale visualization. Other works [7, 13, 69] have developed distributed optimization frameworks to accelerate the training process and manage heavy memory overhead.

The aforementioned methods adopt the “divide-and-

conquer” strategy, where each block requires sufficient parallax information from massive image collections to achieve high-quality reconstruction. While effective, their dependency on dense, low-altitude aerial images makes data acquisition costly and complex, impeding scalable and economical city-scale 3D modeling. In contrast, satellite imagery, with its inherent advantage of vast coverage, presents an ideal data source for this task. However, such images provide little parallax on vertical structures like building facades, which makes the direct application of existing parallax-dependent methods challenging, thus requiring more suitable representations.

2.2. 3D Reconstruction from Satellite Images

Reconstructing urban scenes from satellite imagery is a long-standing goal in remote sensing and photogrammetry. Approaches in the traditional remote sensing field, while proficient with satellite data, pursue a distinct objective: generating metrically accurate 2.5D Digital Surface Models (DSMs) [70], i.e., height maps. Classic photogrammetric pipelines [10, 45, 65] and modern MVS systems, including recent deep learning-based methods [8, 16, 17, 61], are all optimized for this purpose. These works are fundamentally geared toward geospatial analysis, treating the output as a 2.5D grid of elevation values rather than a true 3D mesh, thus not directly suited for applications requiring visual realism or fine-grained geometric fidelity. Recently, many works have started to improve the visual quality of satellite-based 3D reconstruction by adopting neural rendering techniques. Early efforts [11, 36, 37, 66, 67] adapted NeRF to the satellite domain, while recent approaches [2, 3, 21, 26] leverage 3DGS for its efficiency and real-time rendering. However, as these methods supervise geometry solely through a photometric loss, they face challenges in resolving ambiguities from sparse, top-down satellite images. This leads to geometric inaccuracies that ultimately cap the achievable visual quality.

Unlike the 2.5D DSMs common in remote sensing, we introduce a fully differentiable 2.5D representation based on a Z-Monotonic SDF, which allows for fine-grained mesh extraction. This robust geometric representation serves as a high-fidelity foundation for subsequent appearance optimization, enabling more accurate and visually realistic 3D reconstructions from satellite images.

2.3. Generative Priors for 3D Reconstruction

Generative priors [12, 41, 48, 49], learned from vast natural images, have unlocked a new capability for 3D reconstruction, especially in ill-posed settings like sparse-view or single-view tasks. Seminal works such as DreamFusion [44] and Magic3D [28] pioneered this direction by introducing Score Distillation Sampling (SDS), which “distills” the knowledge from a 2D generative model into a 3D neu-

ral field. Subsequent works demonstrate that a powerful image-conditioned generative prior can be leveraged to infer a complete 3D object [30, 34, 38], or even an entire scene [19, 56, 63]. However, these powerful generative priors inherently favor plausible extrapolation over strict fidelity to the input, making them a double-edged sword for high-fidelity reconstruction tasks.

An alternative paradigm, which our work subscribes to, shifts the focus from generation to the refinement of existing 3D appearance [14, 31, 59, 62]. A key challenge in this approach is the inherent stochasticity of generative models, which may produce conflicting details for the same scene. To address this, recent work Skyfall-GS [26] adopts a “generate-and-average” strategy by optimizing over an ensemble of candidates, which is effective but computationally expensive. In contrast, our work introduces a deterministic image restoration process. By fine-tuning a diffusion model to learn a direct mapping from degraded renders to sharp targets, we generate high-quality, holistically-consistent supervision in a single pass.

3. Method

We now describe our method. As illustrated in Fig. 3, our method reconstructs a city in two main stages, each designed to address a core challenge of satellite-based reconstruction. First, to overcome the extreme viewpoint gap inherent to satellite imagery, we introduce a strong geometric prior by representing the city as a Z-Monotonic SDF. This stabilizes optimization from sparse, top-down views and yields a high-fidelity, watertight mesh (Sec. 3.1). Second, to resolve the blur and visual artifacts, we leverage a generative prior by training a large-scale diffusion model to map degraded inputs to sharp, plausible textures (Sec. 3.2). Lastly, we provide implementation details in Sec. 3.3.

3.1. 2.5D Geometry Representation

The primary challenge of satellite imagery for 3D reconstruction lies in its extreme viewpoint constraints. As shown in Fig. 2, standard MVS algorithms struggle to reconstruct vertical facades, yielding point clouds that are dense on rooftops and ground but virtually empty elsewhere. This massive data void on building facades creates a severely ill-posed condition for unconstrained 3D representations like NeRF or 3DGS. As a result, the optimizer is free to produce noisy, collapsed, or “shrink-wrapped” geometry that completely fails to capture the true city structure.

Our key insight is to regularize this ill-posed problem by introducing a strong, urban-specific structural prior: modeling the city as a 2.5D height map. This representation aligns perfectly with both the top-down nature of satellite data and the predominantly vertical geometry of urban architecture. While this design strategically sacrifices the ability to model non-monotonic structures (e.g., bridges), it provides a de-




Figure 3. **The framework of our method.** Our pipeline first reconstructs city geometry, then refines its appearance. **Stage 1 (Geometry):** We optimize a Z-Monotonic SDF against sparse MVS points to extract a high-fidelity, watertight mesh with clean vertical facades. **Stage 2 (Appearance):** Starting with an initial texture (back-projected from source images), we use a restoration network to enhance close-range novel-view renderings, which further serve as sharp, high-fidelity supervision for final texture optimization.

cisive gain in robustness against geometric ambiguity. For most urban scenes, this proves to be a highly effective trade-off, enabling our method to reconstruct high-fidelity geometry where unconstrained approaches inevitably fail.

Z-Monotonic SDF. A 2.5D height map (also called DSM in remote sensing literature) defines the 3D surface as a single-valued height function $z = f(x, y)$ over the 2D ground plane. A naive way to generate this 2.5D height map would be preparing a set of city point cloud $\mathbf{P}$ from the satellite images by applying MVS algorithms [15]. This point cloud $\mathbf{P}$ can be then directly converted into a 2.5D height map, followed by filling a 3D voxel grid and extracting a surface via Marching Cubes [35]. However, this naive conversion introduces significant aliasing or “stair-step” artifacts, primarily due to the noisy and sparse point cloud input, as illustrated in Fig. 4.

To overcome this, we design a novel 2.5D representation using a **Z-Monotonic SDF**, avoiding aliasing and topological holes at building edges. Specifically, we optimize the SDF field $s(\mathbf{x})$ with a special Z-constraint, enforcing its values to be non-decreasing along the vertical Z-axis:

$$\frac{\partial s(x, y, z)}{\partial z} \geq 0 \quad \forall \mathbf{x} = (x, y, z). \quad (1)$$

With this constraint, our geometric representation supports continuous surface extraction for both **continuous surfaces** (e.g., roofs and ground) and **discontinuous surfaces** (e.g., facades). At a continuous surface, the equation $s(x, y, z) = 0$ has a unique solution z , defining the height. Alternatively, at a building edge, the function becomes a vertical plateau where $s(x, y, z) = 0$ for all $z \in [z_{\text{ground}}, z_{\text{roof}}]$, procedurally defining a perfect vertical facade. Thus, the 0-level-set $\{\mathbf{p} \mid s(\mathbf{p}) = 0\}$ implicitly defines this 2.5D surface, trans-

forming the ill-posed facade reconstruction problem into a well-constrained optimization.

We implement the Z-Monotonic SDF via learning monotonic curves on each 2D X-Y plane grid. The SDF value $s(x, y, z)$ can be naturally interpolated. During optimization, we extract a mesh M from the SDF using a differentiable iso-surfacing technique [52]. We then optimize the SDF (parametrized by monotonic curves) by minimizing the Z-axis distance from the MVS points P to the extracted mesh M , together with a Laplacian regularization term and a normal consistency term:


$$\mathcal{L}_{\text{geo}} = \sum_{\mathbf{p} \in P} \min_{\mathbf{m} \in M} \|\mathbf{p}_z - \mathbf{m}_z^*(\mathbf{p})\|_1 + \lambda_{\text{Lap}} \mathcal{L}_{\text{Lap}} + \lambda_{\text{Nrm}} \mathcal{L}_{\text{Nrm}}, \quad (2)$$

where $\mathbf{m}^*(\mathbf{p}) := \underset{\mathbf{m} \in M}{\operatorname{argmin}} \|\mathbf{p}_{xy} - \mathbf{m}_{xy}\|_2$. This optimization process ensures a clean, watertight mesh with accurate vertical facades, as shown in Fig. 4 (c). More details are provided in Appendices B.1 and B.2.

3.2. High-Fidelity Appearance Modeling

With the robust city-scale mesh M established, we proceed to its appearance modeling. A naive approach of optimizing a texture $\mathbf{T}$ by back-projecting the source satellite images I_{sat} is fundamentally limited by the input quality. As demonstrated in Fig. 5, this process “bakes” blurriness, aliasing, and projection artifacts directly into the texture, rendering any synthesized ground-level views unconvincing. To circumvent this, we propose an iterative refinement process that leverages a powerful generative prior to synthesize sharp, plausible textures from these degraded inputs.

Basic Texture Creation. We first compute a preliminary texture $\mathbf{T}_{\text{basic}}$ to provide a coarse but geometrically-aligned



(a) Marching Cubes Low Res. (128) (b) Marching Cubes High Res. (256) (c) Z-Monotonic SDF

Figure 4. **Z-Monotonic SDF vs. Naive Conversion.** (a, b) A naive 2.5D mesh, generated by directly converting sparse MVS points into a voxel grid, suffers from severe “stair-step” artifacts and topological holes. (c) Our Z-Monotonic SDF representation optimizes a continuous field, resulting in a clean, watertight mesh with precise roofs and sharp vertical facades.

starting point. This is achieved by optimizing a texture atlas to match the source satellite images I_{sat} via a differentiable renderer $\mathcal{R}$ [42], minimizing a combination of MSE and SSIM losses:

$$\mathbf{T}_{\text{basic}} = \arg \min_{\mathbf{T}} \sum_{I_i \in I_{\text{sat}}} \lambda_{\text{MSE}} \mathcal{L}_{\text{MSE}} + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}}, \quad (3)$$

where I_i is the i -th satellite image, $\mathbf{C}_i$ is the corresponding camera parameter, $\mathcal{L}_{\text{MSE}} = \|\mathcal{R}(\mathbf{M}, \mathbf{T}, \mathbf{C}_i) - I_i\|_2^2$ and $\mathcal{L}_{\text{SSIM}} = (1 - \text{SSIM}(\mathcal{R}(\mathbf{M}, \mathbf{T}, \mathbf{C}_i), I_i))$. While visually degraded, this texture $\mathbf{T}_{\text{basic}}$ serves as a crucial foundation for the subsequent generative refinement.

Image Restoration Network. The core of our appearance enhancement lies in a powerful generative prior, which we implement by fine-tuning a pre-trained diffusion model. The key is to train this network, denoted D , as a deterministic restorer. It learns to directly map degraded renders (I_{low}) to high-quality targets (I_{high}) from a paired large dataset of diverse 3D urban scenes. This harnesses the rich data distribution of urban appearance learned by the diffusion model while ensuring a stable, repeatable output. The network is optimized by minimizing a combination of a perceptual loss ($\mathcal{L}_{\text{LPIPS}}$) and a fidelity loss (Charbonnier loss $\mathcal{L}_{\text{CHAR}}$ [6]):

$$\mathcal{L}_{\text{restorer}} = \mathcal{L}_{\text{LPIPS}}(\hat{I}, I_{\text{high}}) + \lambda_{\text{CHAR}} \mathcal{L}_{\text{CHAR}}(\hat{I}, I_{\text{high}}). \quad (4)$$

This process yields a robust restorer D that serves as an expert on photorealistic urban appearance, ready to guide the subsequent texture refinement.

Iterative Texture Refinement. The final texture atlas, $\mathbf{T}_{\text{final}}$, is achieved through an iterative refinement process that distills knowledge from our image restoration network D . A single iteration unfolds in two steps: First, in the generation step, we render novel close-range views $I_{\text{low},j} = \mathcal{R}(\mathbf{M}, \mathbf{T}_{\text{cur}}, \mathbf{C}_{\text{novel},j})$ from the current texture atlas (initially $\mathbf{T}_{\text{basic}}$), using simulated UAV paths. These degraded views are then enhanced by D into sharp, high-fidelity pseudo-ground truth targets $\hat{I}_{\text{target},j} = D(I_{\text{low},j})$. Second, in the optimization step, the texture atlas is updated by minimizing a

reconstruction loss against these targets, following the formulation of Eq. (3).

This creates a powerful feedback loop: the refined texture from the current iteration serves as a higher-quality input for the next generation step, allowing the model to progressively “bootstrap” its way to photorealism [9, 63].

3.3. Implementation Details

Geometry. We implement the Z-Monotonic SDF by parameterizing a field of monotonic curves on a 2D (x, y) grid of resolution 256×256 . Each cell in this grid holds a learnable scalar parameter, h , which dictates the vertical offset of a local basis curve. For any query point $\mathbf{P} = (x, y, z)$, the SDF value is not determined by a single curve, but is synthesized by smoothly interpolating the outputs of multiple basis curves from the local neighborhood on the grid. Specifically, for a set of n neighboring grid locations whose xy -coordinates are $\{(x_j, y_j)\}_{j=1}^n$ around the query’s projection (x, y) , the curve $f(z; x, y)$ is represented by n activation function defined as:

$$f(z; x, y) = \sum_{j=1}^n w_j \cdot \tanh(k \cdot (z - h_j)), \quad (5)$$

where h_j are learnable parameters from the 2D grid, k is a hyperparameter, and w_j is an interpolation weight that depends on the proximity of the query location (x, y) to the neighbor’s location (x_j, y_j) . The grid is optimized using the Adam [24] optimizer with a learning rate of 0.01. The optimization is supervised by the $\mathcal{L}_{\text{geo}}$ loss (Eq. (2)), which balances the Z-axis distance to MVS points and a Laplacian loss $\mathcal{L}_{\text{Lap}}$ and a normal consistency loss $\mathcal{L}_{\text{Nrm}}$ (weighted by λ_{Lap} and λ_{Nrm}).

Appearance. For stability, geometry and texture are optimized separately. The texture is parameterized as a UV atlas, with all rendering performed by EdgeGrad [42]. The preliminary texture $\mathbf{T}_{\text{basic}}$ is optimized for 100 epochs, with $\lambda_{\text{MSE}} = 0.8$, $\lambda_{\text{SSIM}} = 0.2$. Our restoration network D fine-tuned from the FLUX-Schnell [25] on 100,000 aerial image pairs, is trained for 10,000 iterations (batch size 96) with $\lambda_{\text{CHAR}} = 1$ in $\mathcal{L}_{\text{restorer}}$ (Eq. (4)). Finally, the texture is refined over 2 iterations for efficiency. For each iteration, we generate target images of resolution 2048×2048 by sampling novel views from a uniform 150 m grid over the mesh bounding box extended by 100 m (altitude 450 m, pitch 45° , four cardinal orientations) and optimize $\mathbf{T}_{\text{final}}$ for 20 epochs.

Platform. Our method runs on a single NVIDIA A6000 GPU and requires approximately 1.5 hours to process a 1 km^2 urban area. This time comprises 0.5 hours for geometry optimization and 1 hour for appearance refinement. Further details on hyperparameters and network architectures are provided in Appendix B.

Method	MatrixCity					DFC 2019			Google Earth		
	P $\uparrow$	R $\uparrow$	F1 $\uparrow$	CD $\downarrow$	PSNR $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Mip-Splatting	0.282	0.569	0.377	0.161	15.992	10.289	0.346	0.816	12.214	0.245	0.551
2DGS	0.693	0.464	0.556	0.073	13.469	7.366	0.304	0.848	11.022	0.196	0.622
CityGS-X	0.278	0.143	0.189	0.227	13.542	FAIL	FAIL	FAIL	12.674	0.223	0.591
Skyfall-GS	0.352	0.256	0.296	0.359	15.949	12.460	0.330	0.740	12.456	0.233	0.521
Ours	0.673	0.615	0.643	0.036	17.153	13.059	0.358	0.556	12.770	0.253	0.546

Table 1. **Quantitative evaluation of our method with SOTA reconstruction works.** Our method outperforms existing approaches in both geometric accuracy and visual quality for satellite reconstruction. “FAIL” denotes the method fails to converge in experiment, manifested as program crashes.




Figure 5. **Appearance Refinement.** (a) The basic texture T_{basic} , created by naively back-projecting the blurry source satellite images, suffers from low fidelity and “baked-in” artifacts. (b) Our final texture T_{final} , optimized using supervision from the restoration network, recovers sharp, photorealistic, and globally consistent details.

4. Experiments

4.1. Experimental Setup

Datasets. To comprehensively validate the geometric accuracy, visual quality and real-world usage of our method, we evaluate on both synthetic and real-world datasets. Specifically, we select four datasets:

MatrixCity-Satellite. Currently there is no open benchmark for 3D reconstruction from satellite imagery that contains accurate ground-truth point clouds. To evaluate geometry quality, we utilize the synthetic 3D scene in MatrixCity dataset and UE5 engine [27] to synthesize satellite images with ground-truth point clouds, enabling quantitative evaluation of both novel view synthesis quality and geometry accuracy. In particular, we synthesize 50 training images covering an area of 1 km^2 to maximally simulate satellite imagery. For consistent and fair evaluation, we follow the *MatrixCity-Aerial* test protocol, and evaluate geometric accuracy and novel view synthesis on the central $800 \text{ m} \times 800 \text{ m}$ area, thereby avoiding evaluation on regions with degraded reconstruction boundaries.

2019 IEEE GRSS Data Fusion Contest (DFC 2019) dataset is a representative real-world satellite imagery dataset that features in high-quality WorldView-3 satellite images and is widely used in the remote sensing area [2, 36, 67]. Following their protocol, we evaluate on four standard Areas of Interest (AOIs): JAX_004, JAX_068, JAX_214, and JAX_260.

GoogleEarth dataset features an appearance-consistent unification of training satellite views and testing ground-level views of reconstructed New York City by Google Earth. We follow [26] to evaluate visual quality on four AOIs (004, 010, 219, 336).

Urban Scene. For a more rigorous evaluation of scalability and robustness, we curated a challenging real-world test case from satellite imagery of a modern metropolis. This scene is characterized by a large-scale, dense urban environment with numerous high-rise buildings and complex geometric layouts.

More details are provided in Appendix C.1.

Baselines. To comprehensively evaluate our method for city-scale reconstruction, we benchmark it against representative SOTA methods from three key domains. For *geometric accuracy*, we compare against 2DGS [20], which aims to reconstruct a geometrically accurate radiance field. For *visual quality*, we benchmark against Mip-Splatting [64], a leading method for high-quality rendering. For *large-scale urban modeling*, we compare against CityGS-X [18] as a SOTA method for city-scale scene reconstruction. Additionally, we include a comparison with Skyfall-GS [26], as it is a satellite-based reconstruction method that enables low-altitude novel view synthesis. These comparisons allow for a comprehensive evaluation of our method’s capabilities across different aspects of this reconstruction task.

Metrics. We evaluate our method from two perspectives: *geometric accuracy* and *visual quality*. In terms of *geometric accuracy*, we follow [18, 33] to evaluate three metrics Precision (P), Recall (R), F1-Score between sampled and GT point clouds. We also follow [20] to evaluate Chamfer distance (CD) as a supplement. In terms of *visual quality*, we adhere standard practices by measuring PSNR, SSIM and LPIPS to evaluate the rendering quality of novel views. Details about point cloud extraction, test view synthesis, etc., are provided in Appendix C.2 and C.3.

4.2. Experiment Results

Quantitative Results. Quantitative results are provided in Tab. 1. Across all datasets, our method consistently




Figure 6. **Visualization of reconstruction results.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery. “FAIL” denotes the method fails to converge in experiment, manifested as program crashes.

achieves superior geometric accuracy and visual quality over existing approaches.

For *geometric accuracy*, it surpasses all baselines in Recall, F1-Score and Chamfer distance, improving F1-Score by 0.09 and reducing Chamfer distance by 50% relative to the best competitor. These gains stem from our 2.5D Z-Monotonic SDF, which enforces coherent roof-facade geometry under satellite views. Although 2DGS [20] attains slightly higher Precision by extracting detailed rooftops, its lack of meaningful facades produces “shrink-wrapped” structures (Fig. 6), validating our choice of a 2.5D prior.


For *visual quality*, our method establishes a new SOTA, proving particularly dominant on challenging, low-altitude datasets like DFC 2019. While still highly competitive, our lead is less pronounced on the Google Earth benchmark, which we attribute to its higher-altitude test views that limit the visibility of facade texture details.

Qualitative Results. Qualitative results are illustrated in Fig. 6. Our qualitative results visually underscore the superiority of our decoupled geometry and appearance modeling. As shown, our method consistently yields clean ge-

Description	Geometry		NVS
	F1 ↑	CD ↓	PSNR ↑
Geometry			
Naive Marching Cubes 128	0.279	0.0749	16.800
Naive Marching Cubes 256	0.412	0.0757	17.002
w/o Regularization Loss	0.637	0.0364	17.115
Full Modeling	0.643	0.0357	17.153
Appearance			
w/o Image Restoration Network	-	-	17.038
Full Modeling	-	-	17.153

Table 2. **Ablation Studies.** We conduct extensive quantitative comparisons to validate the effectiveness of each modules.

ometry and sharp, photorealistic textures. In contrast, baseline approaches such as Mip-Splatting, 2DGS, and CityGS-X produce fragmented, “floating” geometry because their point-based representations lack a strong structural prior to regularize the ill-posed reconstruction from sparse views. While Skyfall-GS shows some robustness, its results suffer from blurry textures and cross-view inconsistencies, a common pitfall of generative approaches that fail to properly enforce global coherence.



(a) w/o Reg. Loss (b) Full Model (c) w/o Img. Res. (d) Full Model

Figure 7. **Qualitative results on ablation study.** Please refer to Fig. 4 for Marching Cubes results. The absence of Z-Monotonic SDF, regularization loss and image restoration network consistently results in significant artifacts in geometry and texture.

We further demonstrate our method’s scalability and robustness on our challenging “Urban Scene” test case. As seen in Fig. 6, our approach demonstrates remarkable resilience in this demanding scenario. The strong geometric prior proves essential, maintaining structural integrity across the vast and complex scene where other methods might fail. On this stable geometric scaffold, our generative refinement then synthesizes sharp, view-consistent details, preserving distinctive facades and vegetation across the entire district. This validates that our two-stage design is not only effective but also highly scalable and robust to the complexities of large-scale urban modeling.

More experiment results are provided in Appendix D.

4.3. Ablation Study

We conduct ablation studies on the MatrixCity dataset to validate our key design choices. The results, summarized in Tab. 2, highlight the contribution of each component to both geometric and appearance quality. More results are provided in Appendix D.6.

Geometry Our Z-Monotonic SDF representation effectively avoids aliasing and other artifacts in surface extraction, ensuring the quality and precision of the reconstructed mesh. Replacing it (optimized with [52]) with a naive baseline that directly extracts a mesh from a voxel grid via Marching Cubes [35] results in Precision and PSNR drops, regardless of the resolution setting. Our regularization losses $\mathcal{L}_{\text{Lap}}$ and $\mathcal{L}_{\text{Nrm}}$ are designed to encourage smoother and more plausible surfaces. Removing these losses results in a degradation of geometric accuracy, adversely affecting visual quality. Moreover, the geometry becomes much noisier when the regularization loss is removed (Fig. 7 (a, b)).

Appearance Disabling our image restoration network also results in a PSNR drop. Moreover, removing the restoration stage results in degraded textures with pronounced artifacts (Fig. 7 (c, d)), highlighting its critical role in bridging the satellite-to-ground viewpoint gap.

4.4. Applications

A key application of our method is the creation and maintenance of city-scale Digital Twins from satellite imagery. Our high-fidelity geometry and elevation reconstructions



Figure 8. **Visualization of snow simulation in Urban Scene.** Our mesh output enables wide downstream applications.

provide the foundation for large-scale urban simulations such as snow modeling, as illustrated in Fig. 8. Furthermore, our method’s versatility extends to aerial view reconstruction, as demonstrated in Appendix D.7.

5. Discussion

Conclusions In this paper, we propose a novel framework for high-quality, city-scale 3D reconstruction from challenging satellite imagery. Our approach successfully addresses the critical limitations of traditional methods, which often fail under the narrow parallax and poor image quality inherent in satellite data. The core of our solution lies in two tailored design choices: (1) a Z-Monotonic SDF representation that robustly models urban geometry as a 2.5D height map, producing clean, watertight meshes with crisp roofs and vertically extruded facades; and (2) a large-scale texture restoration network that effectively recovers high-fidelity appearance details from blurry inputs. Through extensive experiments, we have demonstrated that our method achieves state-of-the-art performance in both geometric accuracy and texture fidelity.

Limitations Our method’s core strengths also define its primary limitations, suggesting two main avenues for future work. First, our reliance on a strong 2.5D geometric prior, while crucial for stabilizing reconstruction from sparse views, imposes fundamental constraints. It inherently precludes modeling of non-monotonic structures (*e.g.*, bridges) and makes the final quality contingent on a topologically sound initial mesh, as our pipeline is designed for refinement, not holistic correction. A failure case is illustrated in Appendix D.9. Future work could explore hybrid representations that integrate local, full 3D models for complex areas. Second, our generative appearance refinement prioritizes visual plausibility over strict factual accuracy. This can lead to the replacement of unique, real-world features with generic textures. Similarly, it synthesizes a single, canonical appearance, rather than explicitly modeling temporal variations present in multi-date satellite imagery. A promising direction is to enhance factual fidelity by incorporating multi-modal data, such as street-view imagery, or by developing time-conditioned appearance models.

References

- [1] Lewis A G Stuart, Andrew Morton, Ian Stavness, and Michael P Pound. 3dgs-to-pc: 3d gaussian splatting to dense point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3730–3739, 2025. 5
- [2] Luca Savant Aira, Gabriele Facciolo, and Thibaud Ehret. Gaussian splatting for efficient satellite image photogrammetry. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5959–5969, 2025. 3, 6, 7, 22
- [3] Nan Bai, Anran Yang, Hao Chen, and Chun Du. Satgs: Remote sensing novel view synthesis using multi-temporal satellite images with appearance-adaptive 3dgs. *Remote Sensing*, 17(9):1609, 2025. 3, 6
- [4] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009. 2
- [5] Chenjie Cao, Xinlin Ren, and Yanwei Fu. Mvsformer++: Revealing the devil in transformer’s details for multi-view stereo. *arXiv preprint arXiv:2401.11673*, 2024. 1
- [6] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st international conference on image processing*, pages 168–172. IEEE, 1994. 5
- [7] Yu Chen and Gim Hee Lee. Dogs: Distributed-oriented gaussian splatting for large-scale 3d reconstruction via gaussian consensus. *Advances in Neural Information Processing Systems*, 37:34487–34512, 2024. 2
- [8] Ziyang Chen, Wenting Li, Zhongwei Cui, and Yongjun Zhang. Surface depth estimation from multi-view stereo satellite images with distribution contrast network. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024. 3
- [9] Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*, 2023. 5
- [10] Carlo De Franchis, Enric Meinhardt-Llopis, Julien Michel, Jean-Michel Morel, and Gabriele Facciolo. An automatic and modular stereo pipeline for pushbroom images. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2014. 3
- [11] Dawa Derksen and Dario Izzo. Shadow neural radiance fields for multi-view satellite photogrammetry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1152–1161, 2021. 3, 6
- [12] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 3
- [13] Guofeng Feng, Siyan Chen, Rong Fu, Zimu Liao, Yi Wang, Tao Liu, Boni Hu, Linning Xu, Zhilin Pei, Hengjie Li, et al. Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26652–26662, 2025. 2
- [14] Tobias Fischer, Samuel Rota Bulò, Yung-Hsu Yang, Nikhil Keetha, Lorenzo Porzi, Norman Müller, Katja Schwarz, Jonathon Luiten, Marc Pollefeys, and Peter Kotschieder. Flowr: Flowing from sparse to dense 3d reconstructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 27702–27712, 2025. 3
- [15] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009. 2, 4
- [16] Jian Gao, Jin Liu, and Shunping Ji. Rational polynomial camera model warping for deep learning based satellite multi-view stereo matching. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6148–6157, 2021. 3
- [17] Jian Gao, Jin Liu, and Shunping Ji. A general deep learning based framework for 3d reconstruction from multi-view stereo satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 195:446–461, 2023. 3
- [18] Yuanyuan Gao, Hao Li, Jiaqi Chen, Zhengyu Zou, Zhihang Zhong, Dingwen Zhang, Xiao Sun, and Junwei Han. Citygs-x: A scalable architecture for efficient and geometrically accurate large-scale scene reconstruction. *arXiv preprint arXiv:2503.23044*, 2025. 2, 6, 5
- [19] Tongyan Hua, Luta Jiang, Ying-Cong Chen, and Wufan Zhao. Sat2city: 3d city generation from a single satellite image with cascaded latent diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 27978–27988, 2025. 3
- [20] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024. 6, 7, 5
- [21] Xuejun Huang, Xinyi Liu, Yi Wan, Zhi Zheng, Bin Zhang, Mingtao Xiong, Yingying Pei, and Yongjun Zhang. Skysplat: Generalizable 3d gaussian splatting from multi-temporal sparse satellite images. *arXiv preprint arXiv:2508.09479*, 2025. 3
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2
- [23] Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics (TOG)*, 43(4):1–15, 2024. 2
- [24] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [25] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2, 5, 3, 20
- [26] Jie-Ying Lee, Yi-Ruei Liu, Shr-Ruei Tsai, Wei-Cheng Chang, Chung-Ho Wu, Jiewen Chan, Zhenjun Zhao, Chieh Hubert Lin, and Yu-Lun Liu. Skyfall-gs: Synthesiz-

- ing immersive 3d urban scenes from satellite imagery. *arXiv preprint arXiv:2510.15869*, 2025. 3, 6, 4, 5, 7, 14, 15, 16
- [27] Yixuan Li, Lihan Jiang, Linning Xu, Yuanbo Xiangli, Zhenzhi Wang, Dahua Lin, and Bo Dai. Matrixcity: A large-scale city dataset for city-scale neural rendering and beyond. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3205–3215, 2023. 6, 3
- [28] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 300–309, 2023. 3
- [29] Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, et al. Vastgaussian: Vast 3d gaussians for large scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5166–5175, 2024. 2
- [30] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 3
- [31] Xi Liu, Chaoyi Zhou, and Siyu Huang. 3dgs-enhancer: Enhancing unbounded 3d gaussian splatting with view-consistent 2d diffusion priors. *Advances in Neural Information Processing Systems*, 37:133305–133327, 2024. 3
- [32] Yang Liu, Chuanchen Luo, Lue Fan, Naiyan Wang, Junran Peng, and Zhaoxiang Zhang. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. In *European Conference on Computer Vision*, pages 265–282. Springer, 2024. 2
- [33] Yang Liu, Chuanchen Luo, Zhongkai Mao, Junran Peng, and Zhaoxiang Zhang. Citygaussianv2: Efficient and geometrically accurate reconstruction for large-scale scenes. *arXiv preprint arXiv:2411.00771*, 2024. 2, 6, 5
- [34] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9970–9980, 2024. 3
- [35] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 4, 8
- [36] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Sat-nerf: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using rpc cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1311–1321, 2022. 3, 6
- [37] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Multi-date earth observation nerf: The detail is in the shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2035–2045, 2023. 3, 6
- [38] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8446–8455, 2023. 3
- [39] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [40] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1
- [41] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 3
- [42] Stanislav Pidhorskyi, Tomas Simon, Gabriel Schwartz, He Wen, Yaser Sheikh, and Jason Saragih. Rasterized edge gradients: Handling discontinuities differentiably. In *European Conference on Computer Vision*, pages 335–352. Springer, 2024. 5
- [43] Pix4D. Example datasets - pix4dmatric. <https://support.pix4d.com/hc/en-us/articles/360048957691>, 2025. Accessed: 2025-11. 6
- [44] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3
- [45] Rongjun Qin. Rpc stereo processor (rsp)—a software package for digital surface model and orthophoto generation from satellite stereo imagery. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3: 77–82, 2016. 3
- [46] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12932–12942, 2022. 2
- [47] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898*, 2024. 2
- [48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [49] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 3
- [50] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE con-*

- ference on computer vision and pattern recognition, pages 4104–4113, 2016. 2
- [51] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European conference on computer vision*, pages 501–518. Springer, 2016. 2
- [52] Tianchang Shen, Jacob Munkberg, Jon Hasselgren, Kangxue Yin, Zian Wang, Wenzheng Chen, Zan Gojcic, Sanja Fidler, Nicholas Sharp, and Jun Gao. Flexible isosurface extraction for gradient-based mesh optimization. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023. 2, 4, 8, 1, 3
- [53] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 1
- [54] Kaiwen Song, Xiaoyi Zeng, Chenqu Ren, and Juyong Zhang. City-on-web: Real-time neural rendering of large-scale scenes on the web. In *European Conference on Computer Vision*, pages 385–402. Springer, 2024. 2
- [55] Matthew Tancik, Vincent Casser, Xincheng Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8248–8258, 2022. 2
- [56] HunyuanWorld Team, Zhenwei Wang, Yuhao Liu, Junta Wu, Zixiao Gu, Haoyuan Wang, Xuhui Zuo, Tianyu Huang, Wenhuan Li, Sheng Zhang, et al. Hunyuanworld 1.0: Generating immersive, explorable, and interactive 3d worlds from words or pixels. *arXiv preprint arXiv:2507.21809*, 2025. 3
- [57] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12922–12931, 2022. 2
- [58] Haithem Turki, Jason Y Zhang, Francesco Ferroni, and Deva Ramanan. Suds: Scalable urban dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12375–12385, 2023. 2
- [59] Jay Zhangjie Wu, Yuxuan Zhang, Haithem Turki, Xuanchi Ren, Jun Gao, Mike Zheng Shou, Sanja Fidler, Zan Gojcic, and Huan Ling. Diffr3d+: Improving 3d reconstructions with single-step diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26024–26035, 2025. 3
- [60] Lining Xu, Yuanbo Xiangli, Sida Peng, Xingang Pan, Nanxuan Zhao, Christian Theobalt, Bo Dai, and Dahua Lin. Grid-guided neural radiance fields for large urban scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8296–8306, 2023. 2
- [61] Shuting Yang, Hao Chen, Fachuan He, Wen Chen, Ting Chen, and Jianjun He. A learning-based dual-scale enhanced confidence for dsm fusion in 3d reconstruction of multi-view satellite images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2025. 3
- [62] Xingyilang Yin, Qi Zhang, Jiahao Chang, Ying Feng, Qingnan Fan, Xi Yang, Chi-Man Pun, Huaqi Zhang, and Xiaodong Cun. Gsfixer: Improving 3d gaussian splatting with reference-guided video diffusion priors. *arXiv preprint arXiv:2508.09667*, 2025. 3
- [63] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5916–5926, 2025. 3, 5
- [64] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19447–19456, 2024. 6, 5
- [65] Kai Zhang, Noah Snively, and Jin Sun. Leveraging vision reconstruction pipelines for satellite imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3, 1, 4
- [66] Lulin Zhang and Ewelina Rupnik. Sparsesat-nerf: Dense depth supervised neural radiance fields for sparse satellite images. *arXiv preprint arXiv:2309.00277*, 2023. 3, 6
- [67] Tongtong Zhang, Yu Zhou, Yuanxiang Li, and Xian Wei. Satensorf: Fast satellite tensorial radiance field for multirate satellite imagery of large size. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–15, 2024. 3, 6, 7, 22
- [68] Yuqi Zhang, Guanying Chen, and Shuguang Cui. Efficient large-scale scene representation with a hybrid of high-resolution grid and plane features. *Pattern Recognition*, 158: 111001, 2025. 2
- [69] Hexu Zhao, Haoyang Weng, Daohan Lu, Ang Li, Jinyang Li, Aurojit Panda, and Saining Xie. On scaling up 3d gaussian splatting training. In *European Conference on Computer Vision*, pages 14–36. Springer, 2024. 2
- [70] Li Zhao, Haiyan Wang, Yi Zhu, and Mei Song. A review of 3d reconstruction from high-resolution urban satellite images. *International Journal of Remote Sensing*, 44(2):713–748, 2023. 3
- [71] MI Zhenxing and Dan Xu. Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In *The Eleventh International Conference on Learning Representations*, 2022. 2

From Orbit to Ground: Generative City Photogrammetry from Extreme Off-Nadir Satellite Images

Supplementary Material

A. Appendix Overview

In this appendix, we provide additional details and results to complement the main paper. Specifically, Sec. B presents extended implementation details of our geometry and appearance pipelines, including the acquisition of satellite-based point clouds, the optimization of our Z-Monotonic SDF, and details for appearance modeling. Sec. C further describes the experimental settings, including dataset configurations, baseline setups, and evaluation metrics. Finally, Sec. D reports more extensive quantitative and qualitative results, ablation studies, and an additional application example that validate the scalability and robustness of our method.

B. Implementation Details

In this section, we introduce additional implementation details in our framework.

B.1. Point Clouds from Multi-View Stereo

In our framework, we first extract rooftop point clouds from MVS. In implementation, we utilize MVSFormer++ and Colmap to extract point clouds. We now describe the details.

To adapt MVSFormer++ [5] for satellite-based MVS, we constructed a new large-scale dataset comprising 2,600 unique satellite scenes with corresponding ground truth depth maps. This dataset was used to train the network specifically for satellite imagery. To enhance point cloud quality, we upgraded the feature extraction backbone in MVSFormer++ by replacing the pre-trained DINOv2 [40] with the more powerful DINOv3 [53]. During inference on real-world multi-view satellite images with their corresponding RPC camera models, we first employ SatelliteSfM [65] to obtain approximate pinhole camera projections, which then serve as the input for our trained MVSFormer++.

B.2. Z-Monotonic SDF Optimization

Our geometry stage reconstructs a high-fidelity, watertight mesh from an MVS point cloud P . We detail the optimization process below, which is applied to 4 spatial divisions of the full scene. The final city model is produced by merging the resulting meshes from each part.

Z-Monotonic SDF Representation For each of the 4 scene divisions, we implement the Z-Monotonic SDF by parameterizing a field of monotonic curves on a 2D (x, y) grid. This is implemented as a learnable 2D parameter grid of resolution 256×256 . When querying the SDF value $s(\mathbf{x})$ at a 3D point $\mathbf{x} = (x, y, z)$, we compute the function from Eq. 5. This is achieved by:

- Defining a 2D neighborhood of size 3×3 centered at the query’s (x, y) location.
- Retrieving the height offset parameters $\{h_j\}$ for this neighborhood by bilinearly sampling from the learnable parameter grid.
- Calculating the weights $\{w_j\}$ using a softmax of the inverse xy -plane distance to these neighbors.
- Computing the final weighted sum $f(z; x, y) = \sum_{j=1}^n w_j \cdot \tanh(k \cdot (z - h_j))$, where k is a fixed hyperparameter controlling the curve sharpness, and we set it to 80.

This formulation ensures the SDF is continuous, differentiable, and inherently Z-monotonic.

Differentiable Optimization We optimize the learnable parameters of the 2D grid for each scene part. We use the Adam optimizer with a learning rate of 0.01 and optimize the grid for 2000 steps. At each optimization step, we extract a mesh M from the current SDF using FlexiCubes [52]. The FlexiCubes module operates on a grid of resolution 128^3 , enabling gradients to flow from the mesh-based loss back to the 2D grid parameters. The optimization is supervised by the $\mathcal{L}_{\text{geo}}$ loss function, as defined in Eq. 2:

$$\mathcal{L}_{\text{geo}} = \sum_{\mathbf{p} \in P} \min_{\mathbf{m} \in M} \|\mathbf{p}_z - \mathbf{m}^*(\mathbf{p})_z\|_1 + \lambda_{\text{Lap}} \mathcal{L}_{\text{Lap}} + \lambda_{\text{Nrm}} \mathcal{L}_{\text{Nrm}}.$$

We detail the definition for each loss term below.

Height Supervision ($\sum_{\mathbf{p} \in P} \min_{\mathbf{m} \in M} \|\mathbf{p}_z - \mathbf{m}^*(\mathbf{p})_z\|_1$) We supervise the Z-Monotonic SDF optimization using height-based supervision. To ensure both efficiency and numerical stability, we avoid computing per-point nearest neighbors on the mesh at each iteration. Instead, we project the MVS point cloud onto a fixed 2D grid and compare it with a rasterized height map rendered from the current mesh. Given the normalized MVS point cloud $P = \{\mathbf{p}_i\}_{i=1}^{N_P}$ in $[-1, 1]^3$, we first convert it into a 2D height map on a regular grid of resolution $R \times R$. For each point $\mathbf{p}_i = (x_i, y_i, z_i)$, we compute its grid indices

$$u_i = \left\lfloor \frac{x_i + 1}{2} R \right\rfloor, \quad v_i = \left\lfloor \frac{y_i + 1}{2} R \right\rfloor, \quad (\text{A1})$$

and accumulate the maximum height per grid cell:

$$H_{\text{target}}(u, v) = \max\{z_i \mid (u_i, v_i) = (u, v)\}, \quad (\text{A2})$$

for all (u, v) that receive at least one point. Cells that do not receive any point are marked as invalid and excluded from the loss via a binary mask.

In parallel, at each optimization step we render a height map H_{pred} from the current mesh M by rasterizing it onto the same $R \times R$ grid. We first transform the mesh vertices from normalized coordinates $[-1, 1]^3$ to image coordinates:

$$\tilde{\mathbf{v}} = \left(\frac{x+1}{2} R, \frac{y+1}{2} R, z \right), \quad (\text{A3})$$

and then use a standard differentiable rasterizer to project the mesh triangles onto the (u, v) plane. For each pixel (u, v) , we compute the interpolated height z via barycentric interpolation of the triangle vertices that cover that pixel. This yields a dense height map $H_{\text{pred}}(u, v)$ that is fully differentiable with respect to the mesh vertices, and thus with respect to the underlying SDF parameters.

We then define the height-map loss as an L_1 difference between the two height maps, computed only on valid pixels where MVS observations exist:

$$\mathcal{L}_H = \frac{1}{|\Omega|} \sum_{(u,v) \in \Omega} |H_{\text{pred}}(u, v) - H_{\text{target}}(u, v)|, \quad (\text{A4})$$

where Ω denotes the set of pixels with valid MVS-derived heights. In implementation, we set $R = 1024$.

Laplacian Regularization ($\mathcal{L}_{\text{Lap}}$) We employ the standard laplacian loss. This loss encourages a smooth mesh surface by penalizing the L_2 distance between each vertex $\mathbf{v}_i$ and the uniform average of its 1-ring neighboring vertices $\mathcal{N}(i)$. The loss is defined as:

$$\mathcal{L}_{\text{Lap}} = \frac{1}{|V|} \sum_{\mathbf{v}_i \in V} \left\| \mathbf{v}_i - \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{v}_j \right\|_2^2, \quad (\text{A5})$$

where V is the set of all vertices in the extracted mesh M . This helps to reduce high-frequency ‘‘bumpy’’ artifacts from the MVS point cloud.

In implementation, we set $\lambda_{\text{Lap}} = 0.5$.

Normal Consistency ($\mathcal{L}_{\text{Nrm}}$) The normal consistency term $\mathcal{L}_{\text{Nrm}}$ is implemented as a Total Variation (TV) loss on the *rendered normal map* $\mathbf{N}$. This encourages smooth changes in surface orientation, which is particularly important for flat facades and ground planes.

At each optimization step, we render a normal map $\mathbf{N}$ from the current mesh M . The loss is defined as the sum of the mean L_2 norms of the finite differences of the normalized normal vectors $\hat{\mathbf{n}}$ between adjacent pixels, computed in the horizontal (x) and vertical (y) directions:

$$\mathcal{L}_{\text{Nrm}} = \mathbb{E}_{u,v} [\|\hat{\mathbf{n}}_{u,v} - \hat{\mathbf{n}}_{u+1,v}\|_2] + \mathbb{E}_{u,v} [\|\hat{\mathbf{n}}_{u,v} - \hat{\mathbf{n}}_{u,v+1}\|_2], \quad (\text{A6})$$

where $\hat{\mathbf{n}}_{u,v}$ is the 3D unit normal vector at pixel (u, v) in the rendered map $\mathbf{N}$. This loss effectively penalizes sharp, local changes in surface normals, promoting piecewise-planar structures.

In implementation, we set $\lambda_{\text{Nrm}} = 0.01$.

FlexiCubes [52] Regularization Following [52], we also add a regularization term $\lambda_{\text{reg}} \mathcal{L}_{\text{reg}}$ in the geometric optimization, where $\lambda_{\text{reg}} = 0.1$, and $\mathcal{L}_{\text{reg}}$ follows the same definition in [52].

Final Mesh Generation After the optimization converges for all 4 divisions, we extract the final mesh for each part. The vertices are un-normalized to their original world coordinates. These 4 meshes are then merged to produce the single city model. Finally, to prepare the mesh for the appearance stage, we compute a single UV atlas for the entire geometry with a resolution of 8192×8192 .

B.3. Appearance Modeling

The core of our appearance enhancement is the image restoration network D , a generative model fine-tuned to deterministically map degraded, low-quality renders to sharp, photorealistic targets.

Network Architecture We construct our restorer D by adapting the pre-trained FLUX-Schnell [25], a state-of-the-art diffusion transformer based on Rectified Flow (RF). While standard RF formulations model image synthesis as an Ordinary Differential Equation (ODE) initiated from Gaussian noise, this stochastic initialization introduces variance that is detrimental to our texture optimization. Our iterative framework strictly requires a stable, one-to-one mapping between the degraded render and the enhanced output to ensure convergence.

To enforce determinism, we reformulate the generation task as a direct, single-step restoration. Instead of sampling a random noise vector, we encode the degraded render I_{low} into a latent code $z_{\text{low}} = E(I_{\text{low}})$ using the pre-trained VAE encoder. This latent code serves as the deterministic boundary condition for the ODE trajectory. Through fine-tuning, the network learns to map the corrupted latent z_{low} directly to its clean counterpart z_{high} in a single function evaluation. This establishes a deterministic mapping $D : z_{\text{low}} \rightarrow z_{\text{high}}$, providing the consistent supervisory signals essential for the stability of our optimization.

Fine-tuning Dataset Construction To train D effectively, we constructed a specialized dataset of 100,000 paired images. This dataset was generated from our extensive internal collection of high-quality 3D urban assets, which is completely disjoint from any of our test scenes to prevent data leakage. From each 3D asset, we rendered multiple sets of paired images $(I_{\text{low}}, I_{\text{high}})$ from low-altitude (e.g., 400–600 m height) and oblique-angle (e.g., 45° – 60° pitch). Each pair consists of a high-resolution, photorealistic rendering I_{high} serving as the ground-truth target, and a corresponding degraded version I_{low} rendered from the identical viewpoint but using a reduced Level-of-Detail (LoD). This degradation strategy naturally suppresses fine-grained textures and high-frequency details, thereby accurately simulating the coarse inputs encountered during the iterative refinement stage.

C. Experimental Settings

C.1. Datasets

In the main paper, we evaluate our method on both synthetic and real-world datasets. We now provide the detailed configurations for each benchmark.

MatrixCity-Satellite We build a synthetic satellite benchmark on top of the MatrixCity-Satellite [27] dataset using the UE5 engine, following the overall protocol described in the main paper. We now detail our capture settings.

For training views, we adopt a clean and stable environment configuration with: no fog, no dynamic weather, and no traffic or moving objects. We place virtual cameras at a fixed altitude of 2000 m above the ground (the maximal value, to the best of knowledge, that yields stable rendering in UE5), with a near-nadir viewing direction. Concretely, we set the yaw angle to 89° relative to the ground plane, a resolution of 2560×1440 , and a field-of-view (FOV) of 22.42° such that the resulting ground sampling distance (GSD) is 0.31 m/px, matching WorldView-3 satellite imagery.

We derive the field-of-view based on the desired ground sampling distance and image resolution. Let H denote the camera altitude above the ground, W the image width in pixels, and θ the horizontal FOV. The ground footprint width L covered by the image is related to the FOV by

$$L = 2H \tan\left(\frac{\theta}{2}\right). \quad (\text{A7})$$

For a target ground sampling distance GSD, the footprint width is also given by

$$L = W \cdot \text{GSD}. \quad (\text{A8})$$

Equating the two expressions for L yields

$$2H \tan\left(\frac{\theta}{2}\right) = W \cdot \text{GSD}, \quad (\text{A9})$$

from which the FOV can be solved as

$$\theta = 2 \arctan\left(\frac{W \cdot \text{GSD}}{2H}\right). \quad (\text{A10})$$

Substituting our settings $H = 2000$ m, $W = 2560$, and $\text{GSD} = 0.31$ m/px, we obtain

$$\theta = 2 \arctan\left(\frac{2560 \times 0.31}{2 \times 2000}\right) \approx 22.42^\circ, \quad (\text{A11})$$

which is the FOV used in all our synthetic satellite captures.

To ensure sufficient overlap between neighboring satellite views, following practical experience, we enforce approximately 60% horizontal overlap between adjacent captures. Given the horizontal footprint width $L \approx 792.7$ m (derived from the FOV and altitude described above), this corresponds to a target stride of

$$s = (1 - 0.6) L \approx 0.4 \times 792.7 \approx 317.1 \text{ m}. \quad (\text{A12})$$

In practice, we set the sampling distance to $s = 317.44$ m, which yields a horizontal overlap very close to 60% between neighboring views.

We then translate the camera centers on a regular grid with this fixed stride along the x - and y -axes, covering the central $[-500, 500] \text{ m} \times [-500, 500] \text{ m}$ area of the scene. At each grid location, we place two virtual cameras with different off-nadir viewing directions to simulate multi-directional satellite passes. Concretely, we use two sets of rotations:

$$(\text{pitch}, \text{yaw}, \text{roll}) = (0^\circ, 89^\circ, 0^\circ) \quad \text{and} \quad (0^\circ, 89^\circ, 90^\circ), \quad (\text{A13})$$

corresponding to south-looking and west-looking near-nadir views, respectively. This configuration results in a sparse multi-view satellite observation pattern that closely approximates real-world acquisition geometries.

For ground-truth point clouds, we use the point clouds provided by MatrixCity-Satellite and select the central $[-400, 400] \text{ m}^2$ region for evaluation. For test views, we follow the default *MatrixCity-Aerial* evaluation protocol and extend it to two altitudes to assess multi-scale performance. We place cameras at heights of 200 m and 500 m, with a resolution of 1920×1080 , FOV of 45° , and a pitch angle of 45° toward the ground. Camera centers are sampled on a uniform grid over the central $[-200, 200] \text{ m}^2$ region with a fixed sampling interval (45.01 m) along both axes. This yields 72 test views that cover the evaluation area with sufficient angular diversity while avoiding boundary regions where the reconstruction quality naturally degrades. These views are used for novel view synthesis evaluation in the main paper.

DFC 2019 For the DFC 2019 benchmark, we follow the overall evaluation protocol of Skyfall-GS [26] in terms of AOI selection and input modality. We obtain approximate pinhole camera intrinsics and extrinsics from the provided RPC parameters using SatelliteSfM [65], and use these cameras both for all competing methods and for rendering our results.

To obtain ground-level reference views, we follow Skyfall-GS [26] and use video sequences from Google Earth Studio (GES) along a low-altitude orbital trajectory around each AOI as test views. In practice, directly using the camera trajectories of [26] may lead to invalid pixels near the image boundaries (e.g., black borders). To ensure a fair and robust evaluation, we additionally apply a simple binary mask that removes boundary pixels before computing the image-based metrics.

GoogleEarth For the GoogleEarth benchmark, we also follow the evaluation protocol of Skyfall-GS [26] in terms of AOI selection, input data, and camera trajectories. For each AOI, we use the training and test views from Google Earth Studio, matching the satellite-like input configuration of [26]. Similar to the DFC 2019 case, these renders may contain invalid pixels near the image borders. We therefore apply the same binary boundary mask as in the DFC 2019 evaluation.

In addition, we empirically observe a systematic misalignment in the NYC_004 scene. The camera parameters provided in [26] do not accurately reproduce the released GES images, and this discrepancy cannot be resolved by simple refinement. To avoid introducing bias in the quantitative comparison, we exclude NYC_004 from our metric reporting and only include NYC_010, NYC_219, and NYC_336 in the main paper.

C.2. Baselines

We compare our method against four representative baselines, as described in the main paper: Mip-Splatting [64], 2DGS [20], CityGS-X [18], and Skyfall-GS [26].

For [26], we use the official implementation and default hyperparameters recommended by the authors for satellite-based reconstruction. Specifically, we evaluate the MatrixCity-Satellite dataset with the recommended DFC 2019 hyperparameters, and to match the scale of DFC 2019 scenes, we uniformly scale up the MatrixCity-Satellite scene by a factor of $50\times$ before running Skyfall-GS.

For the other baselines, directly adopting their default hyperparameters (which are typically designed for dense aerial or street-view imagery) leads to numerical instability or divergence when applied to our extreme off-nadir satellite setting. To ensure a fair and stable comparison, we carefully re-tune a small subset of hyperparameters while keeping the overall model architectures unchanged. In particular, for [18, 20, 64], we reduce the initial position learning rate from 1.6×10^{-4} to 1.6×10^{-5} , decrease the densification ratio from 0.01 to 0.001, slightly increase the densification gradient threshold from 2.0×10^{-4} to 4.0×10^{-4} , and disable frequent opacity resets. In addition, we extend the far plane from 100.0 to 1.0×10^8 and forward the principal point (c_x, c_y) estimated by COLMAP into the projection matrix, so that these baselines can handle our kilometer-scale, satellite-like scenes without modifying their core models.

C.3. Metrics

For PSNR and SSIM, we follow standard image quality evaluation protocols. For LPIPS, we adopt the official implementation with the AlexNet backbone, as commonly done in prior work.

For geometric evaluation, we report the Chamfer distance (CD) and the Precision/Recall/F1 metrics between the reconstructed and ground-truth point clouds, following [18, 20, 33]. The Precision/Recall/F1 metrics are adopted in [18, 33]. For the distance threshold d_r for aerial reconstructions, it is typically set to 0.006 in normalized scene units, which corresponds to an effective ground sampling distance (GSD) of about 5 cm/px in aerial settings. In our satellite setting, however, the GSD is approximately 31 cm/px, *i.e.*, about six times coarser than in the aerial case. To maintain a comparable geometric tolerance in physical units, we therefore scale the threshold proportionally and set $d_r = 0.036$ which yields a consistent relative matching criterion under our satellite imaging configuration.

For point cloud extraction, we follow the native pipelines of each method whenever possible: for [18, 20] we directly use their official mesh export routines, while for [26, 64] we convert the optimized 3D Gaussian representations into meshes using the SOTA method [1]. We then uniformly sample points on all resulting meshes following the protocol of [18] to obtain the predicted point clouds used in our geometric evaluation.

D. More Results

D.1. Quantitative Comparison

Tabs. A1 and A2 present per-scene quantitative comparisons on the DFC 2019 and GoogleEarth datasets, respectively. Across all AOIs, our method consistently matches or surpasses state-of-the-art baselines in PSNR and SSIM, while achieving competitive or better LPIPS. Notably, the “Ours w/o Image Restoration Network” variant already improves over most baselines, and further gains are obtained by our full appearance modeling pipeline, highlighting the contribution of our deterministic restoration network.

D.2. Qualitative Comparison with Ground-Truth Views

Figs. A1 to A8 provide per-scene qualitative comparisons between our method, competing baselines, and the ground-truth test views. Our reconstructions preserve fine-scale structural and texture details, such as building facades and roof patterns, which are often blurred, distorted, or missing in alternative methods.

D.3. Qualitative Analysis at Varying Camera Altitudes

To further investigate how our reconstructed city models behave under varying observation altitudes, we conduct a controlled rendering study with fixed camera intrinsics and identical view geometry. Specifically, we fix the camera’s horizontal position, look-at point, and intrinsics ($f_x, f_y = 1500$), while systematically varying the camera height at [50, 200, 400, 600, 800] meters.

Fig. A9 presents a qualitative comparison between our method and Skyfall-GS [26] in real-world scenes. Our model exhibits stronger robustness, consistently delivering superior results across different observation altitudes. As the camera height decreases, the performance gap becomes increasingly evident: our method preserves markedly sharper facade textures

and more refined geometric details. While our reconstructions may exhibit minor seams due to the joint optimization over multiple satellite views, Skyfall-GS suffers from pronounced quality degradation at lower altitudes, including widespread blurring and noticeable “floating object” artifacts.

D.4. Close-Up Views

The extreme viewpoint gap between orbital inputs and low-altitude target views poses a particular challenge for satellite-based reconstruction. To further evaluate our method under such challenging conditions, we provide additional close-up comparisons that focus on fine-scale geometric and texture details.

Fig. A10 shows close-up novel views on the MatrixCity-Satellite benchmark, comparing our results against Skyfall-GS [26] and the ground truth. Our method produces sharper facades and more detailed roof textures, with significantly fewer ghosting and aliasing artifacts.

Similarly, Figs. A11 and A12 present close-up comparisons on DFC 2019 and GoogleEarth, respectively. Across all real-world scenes, our reconstructions retain fine-grained details such as window patterns and vertical edges better than Skyfall-GS, validating the effectiveness of our two-stage geometry and appearance modeling for close-up observations.

D.5. Additional Real Urban Scenes

Beyond the benchmark datasets used in the main paper, we further validate the generalization ability of our method on two additional real-world urban scenes. These scenes are of the same type and scale as the *Urban Scene* in the main paper, but are completely disjoint in terms of geographic location and appearance.

Figs. A13 and A14 show bird’s-eye overviews of two scenes reconstructed by our method. The results demonstrate that, without dataset-specific tuning, our approach can robustly handle diverse urban layouts, ranging from dense high-rise clusters to wide road networks and large open areas, while maintaining globally coherent geometry and visually plausible textures across each entire scene.

D.6. Ablation Study

We provide additional ablation results to complement the analysis in the main paper.

Fig. A15 compares different geometric representations and meshing strategies on the MatrixCity-Satellite scene. From left to right, we visualize the ground truth, a naive 2.5D mesh obtained via Marching Cubes at low and high resolutions (MC 128/256), a full 3D SDF optimized with FlexiCubes, a variant trained with a Chamfer distance-based loss only, and our full Z-Monotonic SDF model. Higher-resolution Marching Cubes still suffer from stair-step artifacts, and full 3D SDFs completely fail due to topological inconsistencies. Our Z-Monotonic SDF yields cleaner roofs and vertically extruded facades, leading to the most faithful reconstruction.

Fig. A16 examines the impact of different appearance modeling choices. We show the ground truth, a variant excluding our image restoration network (“w/o Image Restoration”), a variant that directly uses the off-the-shelf FLUX-Kontext model for enhancement, and our full appearance pipeline. Without explicit restoration, textures remain blurry and exhibit baked-in projection artifacts. Using FLUX-Kontext improves sharpness but introduces view-inconsistent hallucinations. In contrast, our fine-tuned, deterministic restorer produces sharp and globally consistent textures across views, which is critical for stable texture optimization.

D.7. Applications

We illustrate an additional application of our framework beyond satellite-to-ground reconstruction. Fig. A17 shows an example where our method is applied to reconstruct an aerial-view scene from a set of oblique aerial images [43]. The top row visualizes a subset of the input views, while the bottom row shows renderings from novel viewpoints using our reconstructed mesh and textures. The results demonstrate that our 2.5D geometry prior and generative appearance refinement are also applicable to aerial photogrammetry, yielding high-quality assets.

D.8. Comparison with Remote Sensing Methods

In the remote sensing community, several methods have also investigated 3D reconstruction from satellite imagery [2, 3, 11, 36, 37, 66, 67]. These methods are primarily designed for accurate digital surface model (DSM) estimation and elevation recovery over large areas, and thus mainly focus on metric height accuracy. In contrast, our work targets high-fidelity reconstruction of the *full 3D scene appearance*, including building facades and complex urban details, for photorealistic novel view synthesis from ground and near-ground viewpoints. This difference in objective naturally leads to complementary strengths:

DSM-oriented methods excel at producing geographically accurate surface models, whereas our method emphasizes visually coherent, view-consistent 3D assets suitable for immersive rendering and city-scale visual applications.

Following Skyfall-GS [26], we compare our method against two representative methods Sat-NeRF [67] and EOGS [2] on the DFC 2019 dataset in terms of novel view synthesis quality. As shown in Tab. A3, our method achieves superior performance across all metrics, indicating that our reconstruction better preserves the visual fidelity of the 3D scene from challenging novel viewpoints.

We further provide qualitative comparisons in Fig. A18. As illustrated, Sat-NeRF and EoGS are able to recover plausible elevation and coarse structure, but their rendered views tend to exhibit oversmoothed textures and limited facade details when viewed from oblique or near-ground viewpoints. In contrast, our method produces sharper, more realistic facades and richer high-frequency details. These results underscore the complementary nature of our approach with respect to DSM-oriented remote sensing methods.

D.9. Failure Cases

As discussed in the Limitations section, our reliance on a 2.5D Z-monotonic geometric prior prevents our method from faithfully modeling non-monotonic structures such as bridges, overpasses, and other multi-layer configurations. In these cases, the Z-monotonic assumption forces the geometry to collapse along the vertical axis, leading to missing underpasses. Fig. A19 shows two representative examples from the *MatrixCity-Satellite* dataset, where the ground-truth geometry exhibits multiple vertical layers, while our reconstruction either merges them into a single surface or produces incomplete structures. We believe that future work could further alleviate these issues, for example by augmenting the 2.5D scaffold with local full-3D representations or topology-aware modules that handle multi-layer and non-monotonic structures, thereby making the framework more flexible in such challenging cases.

Scene	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
JAX_004	Mip-Splatting	13.124	0.262	0.868
	2DGS	6.958	0.209	0.906
	CityGS-X	FAIL	FAIL	FAIL
	Skyfall-GS	12.955	0.248	0.790
	Ours w/o Image Restoration Network	13.524	0.267	0.664
	Ours	13.857	0.294	0.606
JAX_068	Mip-Splatting	7.950	0.314	0.834
	2DGS	8.079	0.300	0.839
	CityGS-X	FAIL	FAIL	FAIL
	Skyfall-GS	11.855	0.300	0.762
	Ours w/o Image Restoration Network	12.824	0.348	0.576
	Ours	12.935	0.341	0.554
JAX_214	Mip-Splatting	9.003	0.404	0.766
	2DGS	6.794	0.357	0.813
	CityGS-X	FAIL	FAIL	FAIL
	Skyfall-GS	12.318	0.398	0.696
	Ours w/o Image Restoration Network	12.566	0.386	0.582
	Ours	12.467	0.397	0.539
JAX_260	Mip-Splatting	11.078	0.402	0.795
	2DGS	7.633	0.349	0.832
	CityGS-X	FAIL	FAIL	FAIL
	Skyfall-GS	12.713	0.371	0.715
	Ours w/o Image Restoration Network	12.793	0.373	0.576
	Ours	12.977	0.400	0.526

Table A1. Quantitative evaluation of our method compared to prior works on every scene of **DFC 2019** datasets.

Scene	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NYC_336	Mip-Splatting	13.389	0.453	0.597
	2DGS	11.014	0.311	0.681
	CityGS-X	14.143	0.393	0.661
	Skyfall-GS	13.497	0.420	0.507
	Ours w/o Image Restoration Network	13.512	0.434	0.491
	Ours	13.894	0.436	0.496
NYC_010	Mip-Splatting	12.022	0.148	0.533
	2DGS	10.899	0.151	0.642
	CityGS-X	12.151	0.140	0.544
	Skyfall-GS	12.184	0.142	0.531
	Ours w/o Image Restoration Network	12.007	0.164	0.538
	Ours	12.513	0.162	0.577
NYC_219	Mip-Splatting	11.230	0.135	0.524
	2DGS	11.151	0.129	0.542
	CityGS-X	11.727	0.135	0.569
	Skyfall-GS	11.686	0.135	0.526
	Ours w/o Image Restoration Network	11.474	0.161	0.534
	Ours	11.901	0.160	0.566

Table A2. Quantitative evaluation of our method compared to prior works on every scene of **GoogleEarth** datasets.




Figure A1. **Results of the MatrixCity-Satellite scene.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery.




Figure A2. **Results of the JAX_004 scene in the DFC 2019 dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery. Results of CityGS-X are removed since the method crashes while recovering this scene.




Figure A3. **Results of the JAX_068 scene in the DFC 2019 dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery. Results of CityGS-X are removed since the method crashes while recovering this scene.




Figure A4. **Results of the JAX_214 scene in the DFC 2019 dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery. Results of CityGS-X are removed since the method crashes while recovering this scene.



Figure A5. **Results of the JAX_260 scene in the DFC 2019 dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery. Results of CityGS-X are removed since the method crashes while recovering this scene.




Figure A6. **Results of the NYC_010 scene in the GoogleEarth dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery.



Figure A7. **Results of the NYC_219 scene in the GoogleEarth dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery.




Figure A8. **Results of the NYC_336 scene in the GoogleEarth dataset.** Compared to baselines, our method successfully achieves high-quality city reconstruction from satellite imagery.





Figure A9. **Qualitative comparison between our method and Skyfall-GS across different camera altitudes.** While our reconstructions may exhibit minor seam artifacts, Skyfall-GS even suffers from substantial quality degradation at lower altitudes.



G.T.

Skyfall-GS

Ours

Figure A10. **Close-Up views of reconstruction results of the MatrixCity-Satellite scene.** Our method retains small-scale details such as window grids, facade lines, and roof textures significantly better than Skyfall-GS [26], evidencing the benefit of our deterministic diffusion-based texture refinement.




Figure A11. **Close-Up views of reconstruction results of the DFC 2019 datasets.** Our reconstructions preserve fine facade structures and building edges under low-altitude viewpoints, whereas Skyfall-GS [26] often produces blurred or smeared textures, confirming our improved near-view fidelity on real satellite data.



Skyfall-GS

Ours

Skyfall-GS

Ours

Figure A12. **Close-Up views of reconstruction results of the Google Earth datasets.** Our reconstructions preserve fine facade structures and building edges under low-altitude viewpoints, whereas Skyfall-GS [26] produces smeared textures and even broken parts, confirming our improved near-view fidelity.




Figure A13. **Bird's-Eye view of a Real Urban Scene Reconstruction.** Our pipeline generalizes to new urban areas without dataset-specific tuning while preserving overall layout, building geometry, and texture plausibility.



Figure A14. **Bird's-Eye view of a Real Urban Scene Reconstruction.** Our pipeline generalizes to new urban areas without dataset-specific tuning while preserving overall layout, building geometry, and texture plausibility.





Figure A15. **Visualization of ablation results on geometry.** We compare our full geometry pipeline against variants using low-/high-resolution Marching Cubes, a generic 3D SDF with FlexiCubes, and a Chamfer-distance-only supervision. Only our Z-Monotonic SDF with height-map + regularized training recovers clean roofs, vertical facades, and watertight, artifact-free structures.



G.T.

w/o Image Restoration

w/ Flux-Kontext

Ours

Figure A16. **Visualization of ablation results on appearance modeling.** Disabling image restoration or directly plugging in a generic FLUX-Kontext [25] model produces blurry and inconsistent textures, whereas our fine-tuned, deterministic restorer yields both sharp and view-consistent appearances that best match the ground truth.



Input



Result

Figure A17. **Overview of aerial-view reconstruction results of our method.** We demonstrate that the same pipeline can be applied to an aerial-view reconstruction task: given a set of oblique aerial images (top), our method reconstructs a geometrically accurate and visually realistic scene (bottom), illustrating its applicability beyond satellite imagery to general photogrammetric settings.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Sat-NeRF [67]	10.220	0.278	0.816
EOGS [2]	7.338	0.181	0.931
Ours	13.059	0.358	0.556

Table A3. **Quantitative comparison with remote sensing methods on DFC 2019.** Our method achieves higher PSNR and SSIM and lower LPIPS for novel view synthesis compared to Sat-NeRF and EOGS, highlighting its focus on reconstructing visually faithful 3D scene appearance.




Figure A18. **Qualitative comparison with remote sensing methods on DFC 2019.** Compared to Sat-NeRF and EOGS, our method produces sharper facades and more realistic textures from oblique and near-ground viewpoints, reflecting our focus on reconstructing the visual appearance of the full 3D scene.




Figure A19. **Failure case on non-monotonic structures.** Non-monotonic scenes from the *MatrixCity-Satellite* dataset where our 2.5D prior fails to capture the multi-layer geometry. We believe that future work could further alleviate these issues.