

Skyfall-GS: Synthesizing Immersive 3D Urban Scenes from Satellite Imagery

Jie-Ying Lee¹, Yi-Ruei Liu², Shr-Ruei Tsai¹, Wei-Cheng Chang¹, Chung-Ho Wu¹, Jiewen Chan¹, Zhenjun Zhao³, Chieh Hubert Lin⁴, and Yu-Lun Liu¹

¹ National Yang Ming Chiao Tung University, ² UIUC

³ University of Zaragoza, ⁴ UC Merced

jayinnn.cs14@nycu.edu.tw, yulunliu@cs.nycu.edu.tw



Fig. 1: Our method synthesizes high-quality, immersive 3D urban scenes solely from multi-view satellite imagery, enabling realistic drone-view navigation without relying on additional 3D or street-level training data. Given multiple satellite images from diverse viewpoints and dates (*left*), our method leverages 3D Gaussian Splatting combined with pre-trained text-to-image diffusion models in an iterative refinement framework to generate realistic 3D block-scale city from limited satellite-view input (*right*). Our method significantly enhances visual fidelity, geometric sharpness, and semantic consistency, enabling real-time immersive exploration.

Abstract. Synthesizing large-scale, explorable, and geometrically accurate 3D urban scenes is a challenging yet valuable task for immersive and embodied applications. The challenge lies in the lack of large-scale and high-quality real-world 3D scans for training generalizable generative models. In this paper, we take an alternative route to create large-scale 3D scenes by leveraging readily available satellite imagery for realistic coarse geometry and open-domain diffusion models for high-quality close-up appearance synthesis. We propose **Skyfall-GS**, a novel hybrid

framework that synthesizes immersive city-block scale 3D urban scenes by combining satellite reconstruction with diffusion refinement, eliminating the need for costly 3D annotations, and also featuring real-time, immersive 3D exploration. We tailor a curriculum-driven iterative refinement strategy to progressively enhance geometric completeness and photorealistic texture. Extensive experiments demonstrate that Skyfall-GS provides improved cross-view consistent geometry and more realistic textures compared to state-of-the-art approaches. Project page: <https://skyfall-gs.jayinnn.dev/>

Keywords: 3D Gaussian Splatting · Satellite Imagery · Scene Synthesis

1 Introduction

High-quality, immersive, and semantically plausible 3D urban scenes are essential for a wide range of applications, including gaming, filmmaking, navigation planning, and robotics. The ability to create a large-scale and 3D-grounded environment supports realistic rendering and immersive experience for storytelling, demonstration, and embodied physics simulation. However, due to limited 3D-informed data, building a generative model for realistic and navigable 3D cities remains challenging. It is expensive and labor-intensive to acquire large-scale 3D and textured reconstructions of cities with detailed geometry, while using Internet image collections faces challenges in camera pose registration and excessive data noise (*e.g.*, transient objects and different times of the day). These constraints prevent existing 3D city generation frameworks from creating realistic and diverse appearances. With this observation, we propose an alternative route for virtual city creation with a two-stage pipeline: partial and coarse geometry reconstruction from multi-view satellite imagery, then close-up appearance completion and synthesis using an open-domain diffusion model.

Satellite imagery offers a compelling alternative due to its extensive geographic coverage, automated collection, and high-resolution capabilities. For instance, Maxar’s WorldView-3 satellite captures approximately 680,000 km² of imagery daily at resolutions up to 31 cm per pixel. Such data encodes a large volume of real-world environment semantics, enabling scalable 3D urban scene creation. However, in Figure 2(a), we show that directly applying 3D reconstruction methods to satellite imagery is insufficient for creating *navigable and immersive* 3D cities. The substantial invisible regions (*e.g.*, building facades) and limited satellite-view parallax create incorrect geometry and artifacts.

Completing and enhancing the geometry and texture in the ground view requires a significant influx of extra information. In Figure 2(b), we study a few state-of-the-art methods in city generation [101, 103]. These methods produce oversimplified building geometries and unrealistic appearances due to strong assumptions, particularly the reliance on semantic maps and height fields as the sole inputs, and overfitting to small-scale, domain-specific datasets. Such an observation motivates us to leverage an open-domain foundation image generation model as an external information source, providing better zero-shot

generalization and diversity. Ground-level novel-view renderings from the GS-reconstructed scene suffer from severe degradation, including floater artifacts and texture smearing, caused by insufficiently constrained Gaussians in regions with limited satellite-view parallax. To address this, we employ an open-domain foundation image generation model to directly refine these degenerate renderings into photorealistic and geometrically consistent outputs, exploiting the model’s rich visual priors to recover plausible appearances. The refined outputs serve as pseudo ground-truth to supervise iterative GS scene optimizations, progressively improving the visual fidelity and geometric consistency across the scene. To stabilize the convergence, we carefully design a curriculum-based view selection and iterative refinement process, where the sampled view angles gradually *fall* from the *sky* to the ground over time. Accordingly, we name our framework **Skyfall-GS**. In Figure 1 and Figure 2, we show that Skyfall-GS substantially enhances texture with 3D-justified geometry compared to the relevant baselines.

Skyfall-GS is a novel hybrid framework that synthesizes immersive 3D urban scenes by combining satellite reconstruction with diffusion refinement, eliminating the need for fixed-domain training on 3D data. Skyfall-GS operates on readily available satellite imagery as the only input, then synthesizes realistic aerial-view appearances and maintains a strong satellite-to-ground 3D consistency. Moreover, Skyfall-GS supports real-time and interactive rendering. Through experiments on diverse environments, we show that Skyfall-GS has better generalization and robustness compared to state-of-the-art methods. Our ablation study shows that each component improves perceptual plausibility and semantic consistency. Skyfall-GS paves the way for scalable 3D urban virtual scene creation, enabling applications in virtual entertainment, simulation, and robotics.

In summary, our contributions include:

- We introduce Skyfall-GS, the first method to synthesize immersive, real-time free-flight navigable 3D urban scenes solely from multi-view satellite imagery, without domain-specific 3D training. Its novelty lies in three integrations tailored to the satellite-to-ground setting:
- A satellite-specific 3DGS reconstruction that pairs entropy-based opacity regularization with pseudo-camera depth supervision; we are the first to demonstrate that these designs tame the limited-parallax, multi-date satellite-imagery setting.
- A curriculum-based iterative dataset update that couples a high-to-low elevation schedule with multi-sample consensus, distilling open-domain diffusion priors into photorealistic, cross-view-consistent facades in occluded regions.

2 Related Work

Gaussian Splatting. 3D Gaussian Splatting (3DGS) [33] offers real-time view synthesis rivaling NeRFs [1–3, 58, 61, 66], with Mip-Splatting [111] addressing scale-change aliasing. In-the-wild variants handle appearance variation and transient objects [10, 14, 26, 35, 76, 92, 104, 113], while scaling to large scenes is addressed via partitioning and LOD [8, 21, 34, 49, 86, 89]. For sparse-view settings, depth and co-regularization priors guide reconstruction [42, 48, 67, 114, 116, 121].



Fig. 2: Limitations of existing novel-view synthesis methods from satellite imagery. (a) Sat-NeRF [56] and naive 3DGS [33] yield blurred or distorted building facades due to insufficient geometric detail and limited parallax from satellite viewpoints. (b) City generation methods [101, 103] produce oversimplified building geometries and unrealistic appearances, primarily due to strong assumptions about the input data, and overfitting to small-scale, domain-specific datasets. In comparison, our method synthesizes more realistic appearances and geometries from aerial views.

Satellite and aerial 3D reconstruction. Classical SfM-MVS pipelines extract DSMs from satellite imagery [5, 17, 78, 115], with neural variants improving geometric fidelity [12, 19, 41, 52, 57, 72, 120]. Sat-NeRF [56] and SatMVS [16, 18] apply NeRF and RPC-based warping to satellite imagery respectively, yet neither recovers occluded facades, while FusionRF [84] improves depth via multispectral acquisitions. EOGS [77] adapts 3DGS with affine cameras and shadow mapping for sparse multi-date satellite imagery, and SkySplat [28] outputs Gaussians from satellite imagery in a feed-forward manner. For aerial imagery, AGS [100] introduces Ray-Gaussian Intersection for large-scale surface reconstruction, City-Gaussian [54, 55] achieves city-scale rendering via scene partitioning and LOD, and Horizon-GS [31] unifies aerial-to-ground reconstruction. However, no existing method jointly handles multi-date appearance variations in satellite imagery while synthesizing ground-level perspectives.

Urban scene synthesis. Cross-view synthesis methods [11, 74, 82, 88, 105, 106, 112] generate ground-level images or videos from satellite imagery without explicit 3D representations, while others employ intermediate 3D structures such as point clouds [45], density fields [69, 71], or voxels [43], and recent methods [27, 32, 44, 70, 107] advance toward explicit 3D generation via diffusion- and lifting-based approaches. Layout-conditioned methods synthesize cities from BEV maps: InfiniCity [47] lifts infinite-pixel BEV maps to 3D via octree voxel completion and neural rendering, CityDreamer [101] introduces compositional generative models separating buildings from backgrounds, and GaussianCity [103] extends this paradigm to 3D Gaussian Splatting [102]; LLM-driven procedural approaches [79, 118, 119] further generate layouts from text, OSM, or semantic maps. Although CityX [118] can accept satellite images as layout guidance, none of these methods faithfully synthesize photorealistic ground-level views directly conditioned on the observed satellite texture.

Diffusion models for 3D reconstruction and editing. Diffusion models [38, 75] have emerged as powerful generative priors for image synthesis. Score distillation pipelines [46, 68, 93] lift 2D priors into 3D; DreamGaussian [87] and GaussianDreamer [109] extend this to Gaussian Splatting, while MVDream [81]




Fig. 3: Overview of the proposed Skyfall-GS pipeline. Our method synthesizes immersive and free-flight navigable city-block scale 3D scenes solely from multi-view satellite imagery in two stages. (a) In the Reconstruction Stage, we first reconstruct the initial 3D scene using 3DGS, enhanced by pseudo-camera depth supervision to address limited parallax in satellite images. We integrate an appearance modeling to handle varying illumination conditions across multi-date satellite images. (b) In the Synthesis Stage, we introduce a curriculum-based Iterative Dataset Update (IDU) refinement technique leveraging (c) a pre-trained T2I diffusion model [38] with prompt-to-prompt editing [36]. By iteratively updating training datasets with progressively refined renders, our approach significantly reduces visual artifacts, improving geometric accuracy and texture realism, particularly in previously occluded areas such as building facades.

and sparse-view methods [7, 9, 20, 51, 53, 59, 98, 99] address multi-view consistency. Inversion-based methods [36, 60, 64, 65] and occlusion-aware inpainting works [13, 15, 50, 62, 63, 80, 91, 95–97, 108] enable scene editing and 3D Gaussian inpainting. Instruct-NeRF2NeRF [23] introduced the Iterative Dataset Update (IDU) paradigm, iteratively refining a NeRF via InstructPix2Pix [6] edits; LucidDreamer [9] and WonderWorld [110] adopt IDU for progressive scene expansion via extrapolation, and RealmDreamer [83] generates forward-facing scenes via iterative inpainting and depth diffusion. However, none of these can be naively applied to our setting: inpainting-based approaches treat facade regions as unknown and synthesize them freely, whereas satellite imagery does capture building facades at an oblique angle, providing appearance constraints that must be respected. Ignoring this leads to synthesized facades inconsistent with the satellite observations.

3 Method

Our two-stage pipeline (Figure 3) transforms satellite images into immersive 3D cities. In the Reconstruction Stage (Section 3.1), we fit a 3D Gaussian Splatting model, adding illumination-adaptive appearance modeling and regularizers for sparse, multi-date views. In the Synthesis Stage (Section 3.2), we recover occluded regions, e.g., facades, through curriculum-based Iterative Dataset Update, repeatedly refining renders with text-guided diffusion edits. The loop keeps textures faithful to the satellite input while preserving geometry, yielding complete, navigable urban scenes from satellite data alone.

Preliminary. 3D Gaussian Splatting (3DGS) [33] encodes a scene as Gaussians with center μ_i , covariance Σ_i , opacity α_i , and view-dependent color. Each Gaussian projects to the image plane with covariance: $\Sigma'_i = JW\Sigma_iW^TJ^T$, where W is the viewing transformation and J is the affine-projection Jacobian. Pixels are alpha-composited front-to-back. Parameters are trained with:

$$\mathcal{L}_{\text{color}} = \lambda_{\text{D-SSIM}} \text{DSSIM}(\hat{C}, C) + (1 - \lambda_{\text{D-SSIM}}) \|\hat{C} - C\|_1. \quad (1)$$

3.1 Initial 3DGS Reconstruction from Satellite Imagery

The initial 3DGS reconstruction must faithfully preserve the texture and geometry of satellite imagery to provide a robust foundation for synthesis. We employ appearance modeling to handle variations in multi-date imagery. Since limited satellite parallax creates floating artifacts, we apply regularization techniques to constrain both texture and geometry.

Approximated camera parameters. Satellite imagery typically uses the rational polynomial camera (RPC) model, directly mapping image coordinates to geographic coordinates. To integrate with the 3DGS pipeline, we employ SatelliteSfM [115] to approximate perspective camera parameters (extrinsic and intrinsic) from RPC and generate sparse SfM points as initial 3DGS points.

Appearance modeling. Multi-date satellite imagery exhibits significant appearance variations due to global illumination changes, seasonal factors, and transient objects (Figure 3(a)). Following WildGaussians [35], we use trainable per-image embeddings $\{e_j\}_{j=1}^N$ (with N input images) to handle varying illumination. We also employ trainable per-Gaussian embeddings g_i to capture localized appearance changes, *e.g.* shadow variations. A lightweight MLP f computes affine color transformation parameters (β, γ) as $(\beta, \gamma) = f(e_j, g_i, \bar{c}_i)$, where e_j is the per-image embedding, g_i is the per-Gaussian embedding, and $\bar{c}_i$ denotes the zeroth-order spherical harmonics (SH). Let $\hat{c}_i(\mathbf{r})$ be the i -th Gaussian’s view-dependent color conditioned on the ray direction $\mathbf{r}$. The transformed color $\tilde{c}_i$ is computed as $\tilde{c}_i(\mathbf{r}) = \gamma \cdot \hat{c}_i(\mathbf{r}) + \beta$. To avoid modeling the appearance changes as view-dependent effects, we limit SH coefficients to zero- and first-order terms. At inference, we fuse the learned appearance into a standard 3DGS by selecting a fixed image embedding e^* and evaluating $f(e^*, g_i, \bar{c}_i)$ for every Gaussian to compute a static color. The embeddings and MLP are then discarded, yielding a portable representation compatible with standard 3DGS renderers and enabling real-time rendering at 60 FPS (1920×1080) on a MacBook Pro M4 Pro.

Opacity regularization. We observe that numerous floaters in reconstructed scenes exhibit low opacity. To encourage geometry to adhere closely to actual surfaces, we propose entropy-based opacity regularization:

$$\mathcal{L}_{\text{op}} = - \sum_i [\alpha_i \log(\alpha_i) + (1 - \alpha_i) \log(1 - \alpha_i)] . \quad (2)$$

This regularization promotes binary opacity distributions, allowing low-opacity Gaussians to be more aggressively pruned during densification. Incorporating this term significantly sharpens geometric reconstruction, providing a better foundation for subsequent synthesis.




Fig. 4: The motivation of curriculum strategy. Renderings of the initial 3D reconstruction from varied elevation angles reveal progressive degradation as the viewing angle decreases.



Fig. 5: Render refinement. (a) Original 3DGS render with artifacts and blurry textures; (b) Refined result showing enhanced geometry and texture quality.

Pseudo camera depth supervision. To further reduce floating artifacts, we sample pseudo-cameras positioned closer to the ground during optimization. From these pseudo-cameras, we render RGB images I_{RGB} and corresponding alpha-blended depth maps $\hat{D}_{\text{GS}}$. We then use an off-the-shelf monocular depth estimator, MoGe [90], to predict scale-invariant depths $\hat{D}_{\text{est}}$ from these renders. We use the absolute value of Pearson correlation (PCorr) to supervise the depth:

$$\mathcal{L}_{\text{depth}} = 1 - \|\text{PCorr}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}})\|_1 ; \quad \text{PCorr}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}}) = \frac{\text{Cov}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}})}{\sqrt{\text{Var}(\hat{D}_{\text{GS}})\text{Var}(\hat{D}_{\text{est}})}} . \quad (3)$$

Optimization. Combining all components, the overall loss for the reconstruction stage is defined as:

$$\mathcal{L}_{\text{sat}}(G, C) = \mathcal{L}_{\text{color}} + \lambda_{\text{op}}\mathcal{L}_{\text{op}} + \lambda_{\text{depth}}\mathcal{L}_{\text{depth}} , \quad (4)$$

where G is the 3DGS representation, C is the set of ground-truth satellite images, λ_{op} and λ_{depth} weight opacity regularization and depth supervision relative to the color reconstruction loss.

3.2 Synthesize via Curriculum-Based Iterative Datasets Update

The iterative dataset update (IDU) technique [23, 59] repeatedly executes render-edit-update cycles across multiple episodes to progressively synthesize 3D scenes. Unlike previous methods that sample camera poses from original training views [23] or simple orbits [59], we introduce a curriculum-based refinement schedule over N_e episodes that specifically addresses the geometric and visual limitations inherent to satellite imagery, producing structurally accurate and photorealistic reconstructions of occluded areas.

Curriculum learning strategy. As illustrated in Figure 4, we observe that 3DGS trained from satellite imagery produces higher-quality renders at higher

elevation angles but degenerates at lower elevation angles. Leveraging this insight, we introduce a curriculum-based synthesis strategy, which progressively lowers viewpoints across optimization episodes. Specifically, we define N_p look-at points $\{P_j\}_{j=1}^{N_p}$ uniformly placed throughout the scene and uniformly sample N_v camera positions along orbital trajectories with controlled elevation angles and radii. Our iterative dataset update (IDU) process starts at higher elevations and progressively descends toward lower perspectives. This approach gradually reveals previously occluded regions, improving geometric detail and texture realism, as validated in our ablation studies (Section 4.2).

Render refinement by text-to-image diffusion model. As illustrated in Figure 5(a), renderings from initial 3DGS contain blurry texture and artifacts. To address this, we leverage prompt-to-prompt editing with a pre-trained text-to-image diffusion model to synthesize disocclusion areas, remove artifacts, and enhance geometry. Prompt-to-prompt editing [24] modifies input images, which are described by the source prompt, to align with the target prompt while preserving structural content. Although originally designed for real or synthetically generated images, we demonstrate its effectiveness for refining degraded 3DGS renders produced from satellite-view training, enabling high-quality appearance enhancement without disrupting the scene’s satellite-consistent geometry. Specifically, we use FlowEdit [36] with the FLUX.1 [dev] diffusion model [39]. Our source prompts describe the original degraded features. The target prompts specify the desired high-quality attributes (details in Supplementary). This approach significantly improves rendering quality. It yields sharper geometry, richer textures, and physically coherent shadows (see Figure 5). This strengthens the 3DGS training dataset for more accurate reconstructions.

Multiple diffusion samples. While diffusion models effectively refine individual 3DGS renders, independently applying the diffusion model across viewpoints introduces inconsistencies. Furthermore, 3DGS is well known to suffer from overfitting on single views, as pointed out by CoR-GS [114], causing artifacts when rendering from novel viewpoints.

Ideally, the optimal denoising diffusion process should produce a distribution where all views maintain consistent 3D appearance. However, independent 2D denoising on each view does not preserve 3D consistency, resulting in a denoising trajectory distribution that is a superset of the optimal trajectories. Selecting a single denoising trajectory from this expanded distribution is unlikely to yield the optimal 3D-consistent result, leading to the artifacts observed in Figure 9(c).

To mitigate this, we synthesize N_s independently refined samples per view, effectively sampling multiple trajectories from the denoising distribution. During optimization, the photometric loss $\mathcal{L}_{\text{color}}$ implicitly averages over these N_s samples. Rather than committing to a single potentially suboptimal denoising path, this approach allows the 3DGS optimization to find a consensus representation that balances fidelity to individual samples while promoting geometric coherence across views. Ablation studies (Section 4.2) and Figure 9(c) confirm that this strategy successfully balances detail preservation with structural coherence.




Fig. 6: Visualization of progressive refinement. This figure illustrates the step-by-step evolution of the synthesized 3D scene. Starting from the initial reconstruction state (a), the geometry and texture are progressively refined through successive stages of the iterative process (b-e), culminating in the final high-fidelity result (f).

Iterative dataset update. Our curriculum-based Iterative Dataset Update (IDU), detailed in Algorithm 1, optimizes the 3DGS over N_e episodes. In each episode, we render curriculum-guided views and refine them using FlowEdit [36] with specified prompts and strengths to generate a new training set. As the curriculum descends to lower altitudes, rendering quality steadily improves, particularly in previously occluded regions, as illustrated in Figure 6. We provide detailed parameters in the Supplementary.

Algorithm 1 3DGS Refinement via Iterative Dataset Updates

Input: N_e, N_v, N_s, N_p : Number of episodes, views per point, samples per view, and look-at points.
Input: $\{R_i\}_{i=1}^{N_e}, \{E_i\}_{i=1}^{N_e}$: Radius and elevation sequences; $\{P_j\}_{j=1}^{N_p}$: Target look-at points.
Input: $\Phi = (T_{src}, T_{tgt}, n_{min}, n_{max})$: FlowEdit parameters.
Input: G : Initial 3DGS from satellite-view training.
Output: G' : Refined 3DGS.

```

1:  $G' \leftarrow G$ 
2: for  $i = 1$  to  $N_e$  do
3:    $cam\_views \leftarrow \text{ORBITVIEWS}(\{P\}, R_i, E_i, N_v)$  ▷ Generate  $N_p \times N_v$  views
4:    $render\_views \leftarrow \text{RENDER}(G', cam\_views)$  ▷ Render RGB images
5:    $refine\_views \leftarrow \text{REFINE}(render\_views, \Phi, N_s)$  ▷ Refine renders using FlowEdit
6:    $G' \leftarrow \text{TRAIN}(G', refine\_views)$  ▷ Update 3DGS using refined views
7: end for
8: return  $G'$ 

```

Optimization. For each episode i , we optimize the 3DGS using:

$$\mathcal{L}_{IDU}(G_{i-1}, \tilde{C}_i) = \mathcal{L}_{color} + \lambda_{depth} \mathcal{L}_{depth}, \quad (5)$$

where G_{i-1} denotes the previous episode’s 3DGS model, and $\tilde{C}_i$ are the current refined images. We provide more implementation details in Supplementary.

4 Experiments

Implementation details. The Reconstruction Stage runs for 30,000 iterations with $\lambda_{\text{D-SSIM}} = 0.2$, $\lambda_{\text{op}} = 10$, and $\lambda_{\text{depth}} = 0.5$. The Synthesis Stage comprises $N_e = 5$ episodes of 10,000 iterations each, with $N_v = 6$ cameras and $N_s = 2$ samples per look-at point. Training images are sampled 75% from IDU-refined views and 25% from original satellite views, preserving consistency with the input satellite imagery. All experiments run on a single RTX A6000 (48GB) GPU. Full details are in Supplementary.

Datasets. We evaluate our method on high-resolution RGB satellite imagery from two sources. First, the 2019 IEEE GRSS Data Fusion Contest (DFC2019) [40] featuring WorldView-3 captures of Jacksonville, Florida (2048×2048 pixels, 35 cm/pixel resolution). Camera parameters and sparse points are generated using SatelliteSfM [115]. We evaluate on four standard AOIs: JAX_004, JAX_068, JAX_214, and JAX_260, following Sat-NeRF [56] and EOGS [77]. Second, for geographic diversity, we use the GoogleEarth dataset [101] (training data for CityDreamer [101] and GaussianCity [103]) containing NYC scenes. We use four scenes (004, 010, 219, 336) with training views rendered at 80° elevation to approximate satellite conditions. Google Earth Studio (GES) [22] renders serve as ground truth for both datasets. See Supplementary for more detail.

Baselines. As our method bridges satellite-based 3D reconstruction and city generation, we select baselines from both fields. For *satellite reconstruction*, we compare with Sat-NeRF [56] and EOGS [77] on DFC2019 (they require RPC input unavailable in GoogleEarth), plus Mip-Splatting [111] (with our appearance modeling enabled on DFC2019) and CoR-GS [114] on both datasets. For *city generation*, we compare with CityDreamer [101] and GaussianCity [103] on GoogleEarth dataset (their training dataset).¹ We use official implementations with default settings.

Evaluation metrics. Skyfall-GS primarily aims to synthesize and enhance the invisible or low-coverage regions from the satellite view; therefore, our evaluation focuses on quality and diversity assessment metrics. We report FID_{CLIP} [37] and CMMD [30] which use the CLIP [73] backbone, as InceptionV3 [85] used in classic FID [25] and KID [4] is unsuitable for modern generative tasks. We complement these with user studies for perceptual quality assessment. We also report pixel-aligned metrics (PSNR [29], SSIM [94], LPIPS [117]) as secondary references. However, these metrics are unsuitable for generative tasks, as the synthetic elements in the obscured regions cannot match the invisible ground-truth. Moreover, on DFC2019, systematic illumination and color gaps between WorldView-3 imagery and GES references make pixel-level scores unreliable; we therefore report DFC2019 pixel metrics per scene in the supplementary material.

Table 1: Quantitative comparison on DFC2019 [40]. Distribution metrics computed against GES references; per-scene pixel-level metrics are in the supplementary material. Red, orange, and yellow indicate the best, second, and third best.

Methods	FID _{CLIP} ↓	CMMD ↓
<i>3D Reconstruction</i>		
Sat-NeRF [56]	86.52	4.788
EOGS [77]	87.67	5.291
CoR-GS [114]	84.95	5.692
Mip-Splatting [†] [111]	86.72	5.404
<i>Our Approach</i>		
Ours	27.03	2.110

[†] Enhanced with our appearance modeling.

Table 2: Quantitative comparison on the GoogleEarth dataset [101]. Metrics are computed against GES reference frames. Red, orange, and yellow indicate the best, second best, and third best results, respectively.

Methods	Distribution Metrics		Pixel-level Metrics		
	FID _{CLIP} ↓	CMMD ↓	PSNR ↑	SSIM ↑	LPIPS ↓
<i>City Generation</i>					
CityDreamer [101]	36.66	4.200	12.58	0.267	0.558
GaussianCity [103]	28.76	2.915	13.41	0.291	0.540
<i>3D Reconstruction</i>					
CoR-GS [114]	26.35	3.758	13.35	0.299	0.412
Mip-Splatting [111]	16.09	2.086	14.13	0.302	0.379
<i>Our Approach</i>					
Ours	10.29	1.959	14.42	0.302	0.393

4.1 Comparisons with Baselines

Quantitative comparison. We evaluate against satellite reconstruction and city generation baselines by dividing rendered frames into 144 patches (512×512 pixels). Reference frames are extracted from GES at 17° elevation for DFC2019 (30 frames/AOI, 4,320 total images) and 45° for GoogleEarth (24 frames/scene, 3,456 total images); all methods are then evaluated on matching videos generated with identical camera parameters. As shown in Tables 1 and 2, our method substantially outperforms all baselines on the distribution metrics (FID_{CLIP} and CMMD) that best reflect generative quality, and remains competitive on pixel-level metrics (Mip-Splatting attains a slightly better average LPIPS on GoogleEarth), as further corroborated by the qualitative and user studies below. On DFC2019 pixel metrics (supplementary material), CoR-GS occasionally achieves a higher SSIM, which we attribute to its overly smooth reconstructions inflating SSIM given its known insensitivity to blurring.

Qualitative comparison. Figure 7(a) compares our method on DFC2019 against Sat-NeRF [56], EOGS [77], CoR-GS [114], and Mip-Splatting [111], all of which exhibit significant distortions and blurry textures at lower viewpoints. Figure 7(b) further includes CityDreamer [101] and GaussianCity [103], which over-simplify geometry and miss scene-specific details (*e.g.*, the red pavement in scene 010). Furthermore, our approach yields sharper building contours, higher texture fidelity, and fewer artifacts across both datasets. Notice that Skyfall-GS also recovers the challenging facade details in occluded regions and complex

¹ Many methods lack available code or models: Sat2Scene [44] (inference only), Sat2Vid [45], EO-NeRF [57], Sat-DN [52], SatelliteRF [120], Sat-Mesh [72], CrossViewDiff [43], SkySplat [28], MagicCity [107], Sat3DGen [70], Sat2City [27], and Sat2RealCity [32].




Fig. 7: Qualitative comparison on (a) DFC2019 and (b) GoogleEarth datasets. The leftmost column shows a representative input satellite image. Our method consistently surpasses baselines in geometric accuracy and texture quality at low-altitude novel views. On (a), our synthesis exhibit sharp building contours and detailed rooftops, whereas Sat-NeRF [56] produces severe fog, EOGS [77] yields excessive darkness, CoR-GS [114] suffers from heavy blurring, and Mip-Splatting [111] exhibits prominent floaters and artifacts. On (b), CityDreamer [101] and GaussianCity [103] over-simplify building geometry and miss scene-specific details, CoR-GS [114] and Mip-Splatting [111] produce blurry facades, while our method delivers sharper geometry, richer facade details, and correctly recovers distinctive features such as the red pavement in scene 010.

structures (*e.g.*, vegetation and bridges) that require high precision. Additional qualitative results are provided in Supplementary.

User studies. We conducted two user studies with 44 participants each, evaluating geometric accuracy, spatial alignment, and overall perceptual quality across 4 scenes per study. In each study, participants were presented with side-by-side renderings and asked to select the result that best matched each criterion. In the first study, participants compared Sat-NeRF [56], EOGS [77], CoR-GS [114], Mip-Splatting [111], and our approach; in the second, CityDreamer [101] and GaussianCity [103] replaced Sat-NeRF and EOGS (see Supplementary for full survey details). As shown in Figure 8, our method achieves dominant win rates of $\approx 90\text{--}94\%$ on DFC2019 and $\approx 79\text{--}82\%$ on GoogleEarth across all three criteria, confirming strong and consistent human preference for our approach over all baselines in both *satellite reconstruction* and *city generation* settings.

4.2 Ablation Studies

We conduct ablation studies on the JAX_068 AOI.




Fig. 8: User study results. Our method consistently outperforms Sat-NeRF [56], EOGS [77], CoR-GS [114], Mip-Splatting [111], CityDreamer [101] and GaussianCity [103], achieving particularly high scores in geometric accuracy and overall perceptual quality. (a) details the comparison on the DFC2019 dataset [40], while (b) details the comparison on the GoogleEarth dataset [101].

Table 3: Ablation on the reconstruction stage. Appearance modeling secures convergence. Opacity regularization and depth supervision enhance visual fidelity and geometric accuracy.

Components			Perceptual Metrics		Geometric Metrics	
App. Mod.	Op. Reg.	Depth Sup.	FID _{CLIP} ↓	CMMD ↓	MAE (m) ↓	RMSE (m) ↓
✗	✗	✗	Failed	Failed	Failed	Failed
✓	✗	✗	41.90	2.450	3.542	5.218
✓	✓	✗	39.95	2.395	2.980	4.527
✓	✓	✓	38.01	2.307	2.250	3.483

Table 4: Ablation on the synthesis stage. We evaluate sample counts (N_s), core components, and compare against baselines.

Method Variation	FID _{CLIP} ↓	CMMD ↓	Time (h)
<i>Multiple Samples (N_s)</i>			
$N_s = 1$	34.11	3.189	3.44
Ours ($N_s = 2$)	28.35	2.875	6.37
$N_s = 3$	28.64	2.769	7.19
$N_s = 5$	29.17	2.677	9.80
<i>Component Ablation</i>			
w/o Curriculum (random)	33.79	3.361	-
Reverse Curriculum	53.03	4.170	-
w/ Context-free Pint.	30.78	2.981	-
Replaced w/ SDEdit	64.74	4.138	-

Ablation on the reconstruction stage. We ablate appearance modeling, opacity regularization, and pseudo-camera depth supervision (see Table 3 and Figure 9). For this ablation, we evaluate at higher elevation angles to assess render quality during the IDU process, rather than at the final low-altitude viewpoints. Appearance modeling is crucial for multi-date convergence, opacity regularization removes floating artifacts (Figure 9(a)), and depth supervision sharpens geometry in planar regions (Figure 9(b)). Together, they yield the lowest FID_{CLIP}/CMMD scores. Additionally, we validate geometric accuracy using DFC2019 LiDAR data [40] by unprojecting 3DGS depth renders into DSMs for comparison. Both opacity regularization and pseudo-depth supervision improve geometry, with their combination achieving the lowest MAE/RMSE.

Ablation on the synthesis stage. We isolate multi-sample diffusion and curriculum view progression. $N_s = 2$ achieves optimal visual results (Figure 9(c)), while $N_s = 5$ yields the lowest CMMD but requires $1.5\times$ longer training with marginal gains, so we adopt $N_s = 2$. The “w/o Curriculum” variant in Table 4 replaces our high-to-low schedule with *random* elevation sampling and is markedly worse at restoring occluded geometry (Figure 9(d)); reversing the schedule (low-to-high) is worse still, as it feeds the most degraded low-elevation renders to the diffusion model first, confirming that *both* the curriculum and its high-to-low direction drive the gain. Replacing our refinement with SDEdit [60] under




Fig. 9: Satellite-view training and IDU refinement ablation. (a) Opacity regularization reduces floating artifacts and yields denser reconstructions. (b) Pseudo-camera depth supervision improves geometry in planar, texture-less regions such as rooftops and roads. (c) Multiple diffusion samples per view reduce high-frequency geometric noise and enhance texture consistency, $N_s = 2$ achieves the optimal results. (d) Curriculum learning progressively introduces challenging views, significantly improving geometric coherence in occluded regions over random sampling. (e) A context-free prompt causes only minor degradation in facade details, demonstrating robustness to prompts. Replacing our refinement with SDEdit leads to severe quality degradation, as noising-denoising fails to synthesize details while preserving satellite-defined geometry.

the same schedule causes significant quality degradation, as noising-denoising cannot hallucinate plausible details while preserving satellite-defined geometry (Figure 9(e), Table 4). A context-free prompt yields negligible difference, confirming robustness to prompts, with full prompts provided in Supplementary. The supplement further shows that multi-sample consensus improves cross-view consistency and that the iterative update is essential.

4.3 Performance and Scalability

Training efficiency. We evaluate runtime on the JAX_214 AOI using a single NVIDIA RTX A6000 (48GB). The full pipeline completes in ~ 6 h 45min, split across 1h 35min for reconstruction and 5h 10min for synthesis (per-episode breakdown in Supplementary). This per-scene cost is a one-time, offline expense on par with reconstruction baselines (Sat-NeRF/EO-NeRF report 10–20h; Mip-Splatting 1h 38min on our hardware) and amortizes across unlimited real-time renders, a reasonable trade-off that entirely bypasses time- and labor-intensive physical data collection.

Multi-block scalability via combined imagery. We further combine two adjacent AOIs (JAX_214, JAX_260) overlapping along a highway into a single $\sim 1\text{km} \times 512\text{m}$ dataset (~ 9 h on one RTX A6000). As shown in Figure 10, structures shared across the boundary remain geometrically consistent with no



Fig. 10: Multi-block scalability on combined JAX_214 and JAX_260 AOIs. Skyfall-GS jointly optimizes a seamless $\sim 1\text{km} \times 512\text{m}$ scene from two adjacent satellite tiles; the inset shows a shared building with no stitching artifacts.

stitching artifacts; on a boundary-spanning orbit, the merged model also outperforms either single-AOI model ($\sim 36\text{--}38\%$ lower FID_{CLIP} ; Supplementary), covering boundary regions neither can.

5 Conclusion

We present Skyfall-GS, the first method to synthesize real-time, immersive, and freely navigable 3D urban scenes solely from multi-view satellite imagery, requiring no human intervention or domain-specific 3D training data. By combining 3D Gaussian Splatting with open-domain diffusion priors in a curriculum-based iterative refinement strategy, our method effectively addresses long-standing challenges, including limited parallax, illumination variations, and large-scale occlusions, and consistently outperforms state-of-the-art reconstruction and generation baselines. We hope this work paves the way for scalable, automated 3D urban scene creation, with promising future directions including city-wide scaling and dynamic scene modeling.

Limitations. Our fixed heuristic camera trajectory works well in general but may leave blind spots in complex geometries, occasionally causing minor artifacts at extreme street-level views. Strict pedestrian-level synthesis ($< 10^\circ$, facade-close) is beyond our scope, as transient objects across the multi-date captures and GES reference distortion at low elevations preclude faithful close-up reconstruction and reliable evaluation. Finally, synthesizing facades requires off-nadir views; while most satellite imagery is off-nadir (*e.g.*, WorldView spans $\sim \pm 40^\circ$), purely nadir input remains an open challenge.

Acknowledgements

This research was funded by the National Science and Technology Council, Taiwan, under Grants NSTC 112-2222-E-A49-004-MY2 and 113-2628-EA49-023-. The authors are grateful to Google, NVIDIA, and MediaTek Inc. for their generous donations. Yu-Lun Liu acknowledges the Yushan Young Fellow Program by the MOE in Taiwan.

References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields (2021)
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. *CVPR* (2022)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-NeRF: Anti-aliased grid-based neural radiance fields. In: *ICCV* (2023)
4. Binkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying MMD GANs. In: *International Conference on Learning Representations* (2018)
5. Bosch, M., Kurtz, Z., Hagstrom, S., Brown, M.: A multiple view stereo benchmark for satellite imagery. In: *2016 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*. pp. 1–9. *IEEE* (2016)
6. Brooks, T., Holynski, A., Efros, A.A.: InstructPix2Pix: Learning to follow image editing instructions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18392–18402 (2023)
7. Chen, H., Wu, J., Jin, Y., Peng, J., Mao, X., Chi, M., Yao, M., Peng, B., Li, J., Cao, Y.: VI3DRM: Towards meticulous 3D reconstruction from sparse views via photo-realistic novel view synthesis (2024), <https://arxiv.org/abs/2409.08207>
8. Chen, Y., Lee, G.H.: DOGS: Distributed-oriented gaussian splatting for large-scale 3D reconstruction via gaussian consensus. *Advances in Neural Information Processing Systems* **37**, 34487–34512 (2024)
9. Chung, J., Lee, S., Nam, H., Lee, J., Lee, K.M.: LucidDreamer: Domain-free generation of 3D gaussian splatting scenes. *arXiv preprint arXiv:2311.13384* (2023)
10. Dahmani, H., Bennehar, M., Piasco, N., Roldao, L., Tsishkou, D.: SWAG: Splatting in the wild images with appearance-conditioned gaussians (2024), <https://arxiv.org/abs/2403.10427>
11. Deng, B., Tucker, R., Li, Z., Guibas, L., Snavely, N., Wetzstein, G.: Streetscapes: Large-scale consistent street view generation using autoregressive video diffusion. In: *SIGGRAPH 2024 Conference Papers* (2024)
12. Derksen, D., Izzo, D.: Shadow neural radiance fields for multi-view satellite photogrammetry. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 1152–1161 (June 2021)
13. Dihlmann, J.N., Engelhardt, A., Lensch, H.P.: SIGNeRF: Scene integrated generation for neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)
14. Fan, C.D., Chang, C.W., Liu, Y.R., Lee, J.Y., Huang, J.L., Tseng, Y.C., Liu, Y.L.: SpectroMotion: Dynamic 3d reconstruction of specular scenes. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 21328–21338 (2025)
15. Fang, J., Wang, J., Zhang, X., Xie, L., Tian, Q.: GaussianEditor: Editing 3D gaussians delicately with text instructions. In: *CVPR* (2024)
16. Gao, J., Liu, J., Ji, S.: 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*. pp. 6148–6157 (2021)
17. Gao, J., Liu, J., Ji, S.: 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). <https://doi.org/https://doi.org/10.1016/j.isprsjprs.2022.12.012>, <https://www.sciencedirect.com/science/article/pii/S0924271622003276>

18. Gao, J., Liu, J., Ji, S.: 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)
19. Gao, K., Lu, D., He, H., Xu, L., Li, J.: Enhanced 3D urban scene reconstruction and point cloud densification using gaussian splatting and google earth imagery (2024), <https://arxiv.org/abs/2405.11021>
20. Gao, R., Holynski, A., Henzler, P., Brussee, A., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T., Poole, B.: CAT3D: Create anything in 3D with multi-view diffusion models. *Advances in Neural Information Processing Systems* (2024)
21. Gao, Y., Li, H., Chen, J., Zou, Z., Zhong, Z., Zhang, D., Sun, X., Han, J.: CityGS-X: A scalable architecture for efficient and geometrically accurate large-scale scene reconstruction (2025), <https://arxiv.org/abs/2503.23044>
22. Google: Google earth studio. <https://earth.google.com/studio> (2024), accessed: 2025-05-14
23. Haque, A., Tancik, M., Efros, A., Holynski, A., Kanazawa, A.: Instruct-NeRF2NeRF: Editing 3D scenes with instructions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023)
24. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022)
25. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: *Advances in Neural Information Processing Systems*. pp. 6626–6637 (2017)
26. Hou, H.Y., Hsu, C.C., Huang, Y.C., Shen, M.Y., Sun, W.F., Sun, C., Chang, C.C., Liu, Y.L., Lee, C.Y.: 3d gaussian splatting with grouped uncertainty for unconstrained images. In: *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2025)
27. Hua, T., Jiang, L., Chen, Y.C., Zhao, W.: Sat2City: 3D city generation from a single satellite image with cascaded latent diffusion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 27978–27988 (2025)
28. Huang, X., Liu, X., Wan, Y., Zheng, Z., Zhang, B., Xiong, M., Pei, Y., Zhang, Y.: SkySplat: Generalizable 3d gaussian splatting from multi-temporal sparse satellite images. *arXiv preprint arXiv:2508.09479* (2025)
29. Huynh-Thu, Q., Ghanbari, M.: Scope of validity of PSNR in image/video quality assessment. *Electronics Letters* **44**(13), 800–801 (2008)
30. Jayasumana, S., Ramalingam, S., Veit, A., Glasner, D., Chakrabarti, A., Kumar, S.: Rethinking FID: Towards a better evaluation metric for image generation. In: *CVPR* (2024)
31. Jiang, L., Ren, K., Yu, M., Xu, L., Dong, J., Lu, T., Zhao, F., Lin, D., Dai, B.: Horizon-GS: Unified 3D gaussian splatting for large-scale aerial-to-ground scenes. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 26789–26799 (2025)
32. Kang, Y., Wang, X., Wu, Z., Shi, Y., Zhu, H.: Sat2RealCity: Geometry-aware and appearance-controllable 3D urban generation from satellite imagery (2025), <https://arxiv.org/abs/2511.11470>
33. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. *ACM TOG* (2023)
34. Kerbl, B., Meuleman, A., Kopanas, G., Wimmer, M., Lanvin, A., Drettakis, G.: A hierarchical 3D gaussian representation for real-time rendering of very large datasets. *ACM TOG* (2024)

35. Kulhanek, J., Peng, S., Kukulova, Z., Pollefeys, M., Sattler, T.: WildGaussians: 3D gaussian splatting in the wild. *NeurIPS* (2024)
36. Kulikov, V., Kleiner, M., Huberman-Spiegelglas, I., Michaeli, T.: FlowEdit: Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629* (2024)
37. Kynkäänniemi, T., Karras, T., Aittala, M., Aila, T., Lehtinen, J.: The role of ImageNet classes in fréchet inception distance. In: *Proc. ICLR* (2023)
38. Labs, B.F.: FLUX. <https://github.com/black-forest-labs/flux> (2024), accessed: 2025-02-28
39. Labs, B.F.: Official weights of FLUX.1 dev. <https://huggingface.co/black-forest-labs/FLUX.1-dev> (2024), accessed: 2025-02-28
40. Le Saux, B., Yokoya, N., Hänsch, R., Brown, M.: Data fusion contest 2019 (DFC2019) (2019). <https://doi.org/10.21227/c6tm-vw12>, <https://dx.doi.org/10.21227/c6tm-vw12>
41. Leotta, M.J., Long, C., Jacquet, B., Zins, M., Lipsa, D., Shan, J., Xu, B., Li, Z., Zhang, X., Chang, S.F., Purri, M., Xue, J., Dana, K.: Urban semantic 3D reconstruction from multiview satellite imagery. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (June 2019)
42. Li, J., Zhang, J., Bai, X., Zheng, J., Ning, X., Zhou, J., Gu, L.: DNGaussian: Optimizing sparse-view 3D gaussian radiance fields with global-local depth normalization. *arXiv preprint arXiv:2403.06912* (2024)
43. Li, W., He, J., Ye, J., Zhong, H., Zheng, Z., Huang, Z., Lin, D., He, C.: CrossViewDiff: A cross-view diffusion model for satellite-to-street view synthesis (2024), <https://arxiv.org/abs/2408.14765>
44. Li, Z., Li, Z., Cui, Z., Pollefeys, M., Oswald, M.R.: Sat2Scene: 3D urban scene generation from satellite images with diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7141–7150 (June 2024)
45. Li, Z., Li, Z., Cui, Z., Qin, R., Pollefeys, M., Oswald, M.R.: Sat2Vid: Street-view panoramic video synthesis from a single satellite image (2021), <https://arxiv.org/abs/2012.06628>
46. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3D: High-resolution text-to-3D content creation. In: *CVPR* (2023)
47. Lin, C.H., Lee, H.Y., Menapace, W., Chai, M., Siarohin, A., Yang, M.H., Tulyakov, S.: InfiniCity: Infinite-scale city synthesis. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023)
48. Lin, C.Y., Wu, C.H., Yeh, C.H., Yen, S.H., Sun, C., Liu, Y.L.: FrugalNeRF: Fast convergence for extreme few-shot novel view synthesis without learned priors. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 11227–11238 (2025)
49. Lin, J., Li, Z., Tang, X., Liu, J., Liu, S., Liu, J., Lu, Y., Wu, X., Xu, S., Yan, Y., Yang, W.: VastGaussian: Vast 3D gaussians for large scene reconstruction (2024), <https://arxiv.org/abs/2402.17427>
50. Liu, K.H., Yang, C.K., Chen, M.H., Liu, Y.L., Lin, Y.Y.: Corrfill: Enhancing faithfulness in reference-based inpainting with correspondence guidance in diffusion models. In: *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 1618–1627. IEEE (2025)
51. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3D object. In: *CVPR* (2023)

52. Liu, T., Zhao, S., Jiang, W., Guo, B.: Sat-DN: Implicit surface reconstruction from multi-view satellite images with depth and normal supervision (2025), <https://arxiv.org/abs/2502.08352>
53. Liu, X., Chen, J., Kao, S.h., Tai, Y.W., Tang, C.K.: Deceptive-NeRF/3DGS: Diffusion-generated pseudo-observations for high-quality sparse-view reconstruction. arXiv preprint arXiv:2305.15171 (2023)
54. Liu, Y., Luo, C., Fan, L., Wang, N., Peng, J., Zhang, Z.: CityGaussian: Real-time high-quality large-scale scene rendering with gaussians. In: European Conference on Computer Vision. pp. 265–282. Springer (2025)
55. Liu, Y., Luo, C., Mao, Z., Peng, J., Zhang, Z.: CityGaussianV2: Efficient and geometrically accurate reconstruction for large-scale scenes (2024), <https://arxiv.org/abs/2411.00771>
56. Marí, R., Facciolo, G., Ehret, T.: Sat-NeRF: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using RPC cameras. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 1310–1320 (2022)
57. Marí, R., Facciolo, G., Ehret, T.: Multi-date earth observation NeRF: The detail is in the shadows. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 2034–2044 (June 2023)
58. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: NeRF in the wild: Neural radiance fields for unconstrained photo collections. In: CVPR (2021)
59. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Neverova, N., Vedaldi, A., Gafni, O., Kokkinos, F.: IM-3D: Iterative multiview diffusion and reconstruction for high-quality 3D generation. International Conference on Machine Learning (2024)
60. 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: International Conference on Learning Representations (2022)
61. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM (2021)
62. Mirzaei, A., Aumentado-Armstrong, T., Brubaker, M.A., Kelly, J., Levinshtein, A., Derpanis, K.G., Gilitschenski, I.: Watch your steps: Local image and scene editing by text instructions. In: ECCV (2024)
63. Mirzaei, A., Aumentado-Armstrong, T., Derpanis, K.G., Kelly, J., Brubaker, M.A., Gilitschenski, I., Levinshtein, A.: Spin-NeRF: Multiview segmentation and perceptual inpainting with neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20669–20679 (2023)
64. Miyake, D., Iohara, A., Saito, Y., Tanaka, T.: Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models (2024), <https://arxiv.org/abs/2305.16807>
65. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. arXiv preprint arXiv:2211.09794 (2022)
66. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans. Graph. **41**(4), 102:1–102:15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://doi.org/10.1145/3528223.3530127>

67. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: RegNeRF: Regularizing neural radiance fields for view synthesis from sparse inputs. In: CVPR (2022)
68. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3D using 2D diffusion. arXiv preprint arXiv:2209.14988 (2022)
69. Qian, M., Tan, B., Wang, Q., Zheng, X., Xiong, H., Xia, G.S., Shen, Y., Xue, N.: Seeing through satellite images at street views. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–18 (2026). <https://doi.org/10.1109/TPAMI.2026.3652860>
70. Qian, M., Xia, Z., Liu, C., Ma, S., Wang, W., Ke, Z., Tan, B., Zhang, H., Xia, G.S.: Sat3DGen: Comprehensive street-level 3D scene generation from single satellite image. In: International Conference on Learning Representations (ICLR) (2026)
71. Qian, M., Xiong, J., Xia, G.S., Xue, N.: Sat2density: Faithful density learning from satellite-ground image pairs. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
72. Qu, Y., Deng, F.: Sat-Mesh: Learning neural implicit surfaces for multi-view satellite reconstruction. *Remote Sensing* **15**(17) (2023). <https://doi.org/10.3390/rs15174297>, <https://www.mdpi.com/2072-4292/15/17/4297>
73. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proceedings of the 38th International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
74. Regmi, K., Borji, A.: Cross-view image synthesis using conditional gans (2018), <https://arxiv.org/abs/1803.03396>
75. 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 (CVPR). pp. 10684–10695 (June 2022)
76. Sabour, S., Goli, L., Kopanas, G., Matthews, M., Lagun, D., Guibas, L., Jacobson, A., Fleet, D.J., Tagliasacchi, A.: SpotLessSplats: Ignoring distractors in 3D gaussian splatting. arXiv preprint arXiv:2406.20055 (2024)
77. Savant Aira, L., Facciolo, G., Ehret, T.: Gaussian splatting for efficient satellite image photogrammetry. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2025)
78. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
79. Shang, Y., Lin, Y., Zheng, Y., Fan, H., Ding, J., Feng, J., Chen, J., Tian, L., Li, Y.: UrbanWorld: An urban world model for 3D city generation. arXiv preprint arXiv:2407.11965 (2024)
80. Shen, J., Agudo, A., Moreno-Noguer, F., Ruiz, A.: Conditional-flow NeRF: Accurate 3d modelling with reliable uncertainty quantification. In: European Conference on Computer Vision. pp. 540–557. Springer (2022)
81. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: MVDream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023)
82. Shi, Y., Campbell, D., Yu, X., Li, H.: Geometry-guided street-view panorama synthesis from satellite imagery. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(12), 10009–10022 (2022)
83. Shriram, J., Trevithick, A., Liu, L., Ramamoorthi, R.: RealmDreamer: Text-driven 3D scene generation with inpainting and depth diffusion. In: International Conference on 3D Vision (3DV) (2025)

84. Sprintson, M., Chellappa, R., Peng, C.: FusionRF: High-fidelity satellite neural radiance fields from multispectral and panchromatic acquisitions. arXiv preprint arXiv:2409.15132 (2024)
85. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)
86. Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-NeRF: Scalable large scene neural view synthesis. In: CVPR (2022)
87. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: DreamGaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023)
88. Toker, A., Zhou, Q., Maximov, M., Leal-Taixé, L.: Coming down to earth: Satellite-to-street view synthesis for geo-localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6488–6497 (2021)
89. Turki, H., Ramanan, D., Satyanarayanan, M.: Mega-NeRF: Scalable construction of large-scale NeRFs for virtual fly-throughs. In: CVPR (2022)
90. Wang, R., Xu, S., Dai, C., Xiang, J., Deng, Y., Tong, X., Yang, J.: MoGe: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision (2024), <https://arxiv.org/abs/2410.19115>
91. Wang, Y., Yi, X., Wu, Z., Zhao, N., Chen, L., Zhang, H.: View-consistent 3D editing with gaussian splatting (2025), <https://arxiv.org/abs/2403.11868>
92. Wang, Y., Wang, J., Qi, Y.: WE-GS: An in-the-wild efficient 3D gaussian representation for unconstrained photo collections (2024), <https://arxiv.org/abs/2406.02407>
93. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: ProlificDreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in neural information processing systems* **36**, 8406–8441 (2023)
94. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
95. Weber, E., Holynski, A., Jampani, V., Saxena, S., Snavely, N., Kar, A., Kanazawa, A.: NeRFiller: Completing scenes via generative 3D inpainting. In: CVPR (2024)
96. Wu, C.H., Chen, Y.J., Chen, Y.H., Lee, J.Y., Ke, B.H., Mu, C.W.T., Huang, Y.C., Lin, C.Y., Chen, M.H., Lin, Y.Y., Liu, Y.L.: AuraFusion360: Augmented unseen region alignment for reference-based 360° unbounded scene inpainting (2025), <https://arxiv.org/abs/2502.05176>
97. Wu, J., Bian, J.W., Li, X., Wang, G., Reid, I., Torr, P., Prisacariu, V.: GaussCtrl: Multi-view consistent text-driven 3D gaussian splatting editing. ECCV (2024)
98. Wu, R., Gao, R., Poole, B., Trevithick, A., Zheng, C., Barron, J.T., Holynski, A.: CAT4D: Create anything in 4D with multi-view video diffusion models (2024), <https://arxiv.org/abs/2411.18613>
99. Wu, R., Mildenhall, B., Henzler, P., Park, K., Gao, R., Watson, D., Srinivasan, P.P., Verbin, D., Barron, J.T., Poole, B., Holynski, A.: ReconFusion: 3D reconstruction with diffusion priors. arXiv preprint arXiv:2312.02981 (2023)
100. Wu, Y., Liu, J., Ji, S.: 3D gaussian splatting for large-scale surface reconstruction from aerial images (2024), <https://arxiv.org/abs/2409.00381>
101. Xie, H., Chen, Z., Hong, F., Liu, Z.: CityDreamer: Compositional generative model of unbounded 3D cities. In: CVPR (2024)
102. Xie, H., Chen, Z., Hong, F., Liu, Z.: CityDreamer4D: Compositional generative model of unbounded 4D cities. arXiv preprint arXiv:2501.08983 (2025)

103. Xie, H., Chen, Z., Hong, F., Liu, Z.: Generative gaussian splatting for unbounded 3D city generation. In: CVPR (2025)
104. Xu, J., Mei, Y., Patel, V.M.: Wild-GS: Real-time novel view synthesis from unconstrained photo collections (2024), <https://arxiv.org/abs/2406.10373>
105. Xu, N., Qin, R.: Geospecific view generation geometry-context aware high-resolution ground view inference from satellite views. In: European Conference on Computer Vision. pp. 349–366. Springer (2024)
106. Xu, N., Qin, R.: Satellite to groundscene-large-scale consistent ground view generation from satellite views. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6068–6077 (2025)
107. Yao, X., Wang, X., Wu, H., Ping, C., Zhang, D., Xiong, H.: MagicCity: Geometry-aware 3d city generation from satellite imagery with multi-view consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 25325–25334 (2025)
108. Ye, M., Danelljan, M., Yu, F., Ke, L.: Gaussian grouping: Segment and edit anything in 3D scenes. In: ECCV (2024)
109. Yi, T., Fang, J., Wang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: GaussianDreamer: Fast generation from text to 3d gaussians by bridging 2d and 3d diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6796–6807 (2024)
110. Yu, H.X., Duan, H., Herrmann, C., Freeman, W.T., Wu, J.: WonderWorld: Interactive 3D scene generation from a single image. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5916–5926 (2025)
111. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-Splatting: Alias-free 3D gaussian splatting. In: CVPR (2024)
112. Ze, X., Song, Z., Wang, Q., Lu, J., Shi, Y.: Controllable satellite-to-street-view synthesis with precise pose alignment and zero-shot environmental control (2025), <https://arxiv.org/abs/2502.03498>
113. Zhang, D., Wang, C., Wang, W., Li, P., Qin, M., Wang, H.: Gaussian in the wild: 3D gaussian splatting for unconstrained image collections. arXiv preprint arXiv:2403.15704 (2024)
114. Zhang, J., Li, J., Yu, X., Huang, L., Gu, L., Zheng, J., Bai, X.: CoR-GS: Sparse-view 3d gaussian splatting via co-regularization. In: ECCV (2024)
115. Zhang, K., Sun, J., Snavely, N.: Leveraging vision reconstruction pipelines for satellite imagery. In: IEEE International Conference on Computer Vision Workshops (2019)
116. Zhang, L., Rupnik, E.: Sparsesat-NeRF: Dense depth supervised neural radiance fields for sparse satellite images. arXiv preprint arXiv:2309.00277 (2023)
117. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 586–595 (2018)
118. Zhang, S., Zhou, M., Wang, Y., Luo, C., Wang, R., Li, Y., Zhang, Z., Peng, J.: Cityx: Controllable procedural content generation for unbounded 3d cities (2024), <https://arxiv.org/abs/2407.17572>
119. Zhou, M., Wang, Y., Hou, J., Zhang, S., Li, Y., Luo, C., Peng, J., Zhang, Z.: SceneX: Procedural controllable large-scale scene generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 10806–10814 (2025)
120. Zhou, X., Wang, Y., Lin, D., Cao, Z., Li, B., Liu, J.: SatelliteRF: Accelerating 3D reconstruction in multi-view satellite images with efficient neural radiance fields. Applied Sciences **14**(7) (2024). <https://doi.org/10.3390/app14072729>, <https://www.mdpi.com/2076-3417/14/7/2729>

121. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: FSGS: Real-time few-shot view synthesis using gaussian splatting (2023)