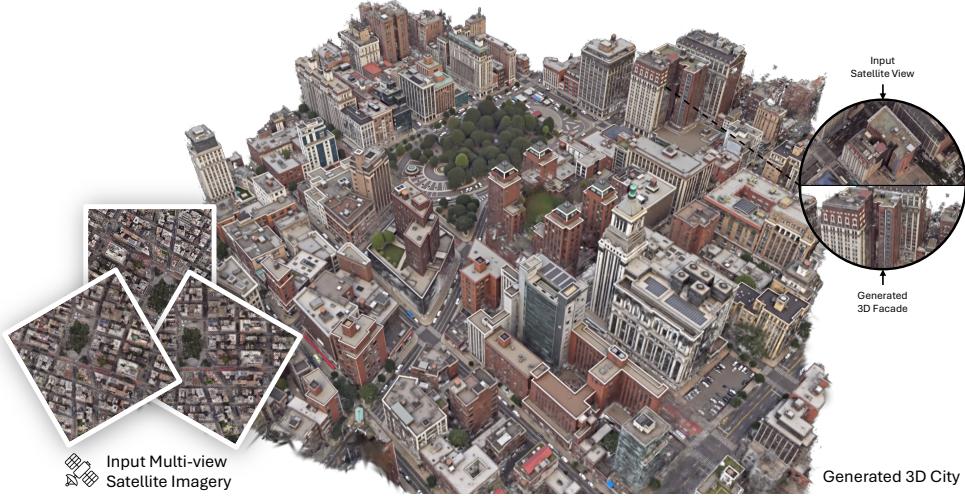


000 SKYFALL-GS: SYNTHESIZING IMMERSIVE 3D URBAN 001 SCENES FROM SATELLITE IMAGERY

002 **Anonymous authors**

003 Paper under double-blind review



004
005 **Figure 1: Our method synthesizes high-quality, immersive 3D urban scenes solely from multi-**
006 **view satellite imagery, enabling realistic drone-view navigation without relying on additional**
007 **3D or street-level training data.** Given multiple satellite images from diverse viewpoints and dates
008 (*left*), our method leverages 3D Gaussian Splatting combined with pre-trained text-to-image diffusion
009 models in an iterative refinement framework to generate realistic 3D block-scale city from limited
010 satellite-view input (*right*). Our method significantly enhances visual fidelity, geometric sharpness,
011 and semantic consistency, enabling real-time immersive exploration.

ABSTRACT

004
005 Synthesizing large-scale, explorable, and geometrically accurate 3D urban scenes is
006 a challenging yet valuable task in providing immersive and embodied applications.
007 The challenges lie in the lack of large-scale and high-quality real-world 3D scans
008 for training generalizable generative models. In this paper, we take an alternative
009 route to create large-scale 3D scenes by synergizing the readily available satellite
010 imagery that supplies realistic coarse geometry and the open-domain diffusion
011 model for creating high-quality close-up appearances. We propose **Skyfall-GS**,
012 a novel hybrid framework that synthesizes immersive city-block scale 3D urban
013 scenes by combining satellite reconstruction with diffusion refinement, eliminat-
014 ing the need for costly 3D annotations, also featuring real-time, immersive 3D
015 exploration. We tailor a curriculum-driven iterative refinement strategy to pro-
016 gressively enhance geometric completeness and photorealistic textures. Extensive
017 experiments demonstrate that Skyfall-GS provides improved cross-view consistent
018 geometry and more realistic textures compared to state-of-the-art approaches.

1 INTRODUCTION

049 Synthetic high-quality, immersive, and semantically plausible 3D urban scenes have crucial appli-
050 cations in gaming, filmmaking, and robotics. The ability to create a large-scale and 3D-grounded
051 environment supports realistic rendering and immersive experience for storytelling, demon-
052 stration, and embodied physics simulation. However, due to limited 3D-informed data, building a generative
053

054 model for realistic and navigable 3D cities remains challenging. It is expensive and labor-intensive
 055 to acquire large-scale 3D and textured reconstructions of cities with detailed geometry, while using
 056 Internet image collections **face** challenges in camera pose registration and excessive data noise
 057 (e.g., transient objects and different times of the day). These constraints set back existing 3D city
 058 generation frameworks from creating realistic and diverse appearances. With this observation, we
 059 propose an alternative route for virtual city creation with a two-stage pipeline: partial and coarse
 060 geometry reconstruction from multi-view satellite imagery, then close-up appearance completion and
 061 hallucination using an open-domain diffusion model.

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

070 Completing and enhancing the geometry and texture in the ground view requires a significant influx
 071 of extra information. In Figure 2(b), we study a few state-of-the-art methods in city generation (Xie
 072 et al., 2024; 2025b). These methods produce oversimplified building geometries and unrealistic
 073 appearances due to strong assumptions, particularly the reliance on semantic maps and height fields as
 074 the sole inputs, and overfitting to small-scale, domain-specific datasets. Such an observation motivates
 075 us to seek help from open-domain foundation vision models as an external information source, which
 076 provides better zero-shot generalization and diversity. Noticing that the ground-view novel-view
 077 renderings from the GS reconstructed scene exhibit noise-like patterns, we treat these renderings as
 078 intermediate results in a denoising diffusion process. Then, we complete the remaining denoising
 079 process to create hallucinated pseudo ground-truth for the GS scene optimization. To stabilize the
 080 convergence, we carefully design a curriculum-based view selection and iterative refinement process,
 081 where the sampled view angles gradually *fall* from the *sky* to the ground over time. Accordingly,
 082 we name our framework **Skyfall-GS**. In Figure 1 and Figure 2, we show that Skyfall-GS yields
 083 significantly enhanced texture with 3D-justified geometry compared to the relevant baselines.

084 Skyfall-GS is a **novel hybrid framework that synthesizes immersive city-block scale 3D urban scenes**
 085 **by combining satellite reconstruction with diffusion refinement, eliminating the need of fixed-domain**
 086 **training on 3D data.** Skyfall-GS operates on readily available satellite imagery as the only input,
 087 then hallucinates realistic aerial-view appearances and maintains a strong satellite-to-ground 3D
 088 consistency. Moreover, Skyfall-GS supports real-time and interactive rendering, as we design our
 089 framework to produce GS results without sophisticated data structures. Through experiments on
 090 diverse environments, we show that Skyfall-GS has better generalization and robustness compared
 091 to state-of-the-art methods. Our ablation shows that each of our designs improves the perceptual
 092 plausibility and semantic consistency. Skyfall-GS paves the way for scalable 3D urban virtual scene
 093 creation, enabling applications in virtual entertainment, simulation, and robotics.

094 In summary, our contributions include:

- 095 • We introduce Skyfall-GS, the first method to synthesize immersive, real-time free-flight navigable
 096 3D urban scenes solely from multi-view satellite imagery using generative refinement.
- 097 • An open-domain refinement approach leveraging pre-trained text-to-image diffusion models without
 098 domain-specific training.
- 099 • A curriculum-learning-based iterative refinement strategy progressively enhances reconstruction
 100 quality from higher to lower viewpoints, significantly improving visual fidelity in occluded areas.

101 2 RELATED WORK

102 **Gaussian Splatting.** 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) offers real-time view syn-
 103 thesis rivaling NeRFs (Mildenhall et al., 2021; Barron et al., 2021; 2022; Müller et al., 2022;
 104 Barron et al., 2023; Martin-Brualla et al., 2021). Mip-Splatting (Yu et al., 2024) fixes scale-
 105 change issues via on-the-fly resizing. Recent advances target satellite and aerial reconstruction:
 106 FusionRF (Sprintson et al., 2024) achieves 17% depth improvement from multispectral acquisitions,
 107 while InstantSplat (Fan et al., 2024) enables 40-second pose-free reconstruction. “In-the-wild” vari-



108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161

Figure 2: **Limitations of existing novel-view synthesis methods from satellite imagery.** (a) Sat-NeRF (Marí et al., 2022) and naive 3DGS (Kerbl et al., 2023) yield blurred or distorted building facades due to insufficient geometric detail and limited parallax from satellite viewpoints. (b) City generation methods (Xie et al., 2024; 2025b) 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.

ants handle appearance and uncertainty (Xu et al., 2024; Sabour et al., 2024; Wang et al., 2024b; Dahmani et al., 2024; Zhang et al., 2024a; Kulhanek et al., 2024; Hou et al., 2025), including SpectroMotion (Fan et al., 2025) for dynamic specular scenes, while large-scene methods use LOD and partitioning (Kerbl et al., 2024; Liu et al., 2025c; 2024; Lin et al., 2024; Turki et al., 2022; Tancik et al., 2022). CAT-3DGS (Zhan et al., 2025) achieves rate-distortion optimization via context-adaptive triplanes. For sparse-view satellite imagery, depth or co-regularization priors guide reconstruction (Li et al., 2024b; Zhang et al., 2024b; Zhu et al., 2023; Niemeyer et al., 2022; Lin et al., 2025), with SparseSat-NeRF (Zhang & Rupnik, 2023) adding dense depth supervision.

Diffusion models for 3D reconstruction and editing. Diffusion models (Rombach et al., 2022; Labs, 2024b) underpin image generation and editing. Early SDS pipelines DreamFusion (Poole et al., 2022) and Magic3D (Lin et al., 2023a) enabled text-to-3D, with ProlificDreamer (Wang et al., 2023) addressing over-smoothing via Variational Score Distillation. DreamGaussian (Tang et al., 2023) achieves 10x speedup via progressive densification, while GaussianDreamer (Yi et al., 2024) bridges 2D and 3D diffusion models. SDEdit (Meng et al., 2022), DDIM inversion (Mokady et al., 2022; Miyake et al., 2024), and FlowEdit (Kulikov et al., 2024) enable fine control. Extensions include sparse-view reconstruction (Wu et al., 2023; Liu et al., 2023b; Chen et al., 2024), with MVDream (Shi et al., 2023) enabling multi-view consistency. For 3D/4D generation (Gao et al., 2024b; Wu et al., 2024b; Melas-Kyriazi et al., 2024; Chung et al., 2023; Liu et al., 2023a) and scene editing (Haque et al., 2023; Wu et al., 2025; Ye et al., 2024b; Fang et al., 2024; Mirzaei et al., 2024; Dihlmann et al., 2024; Weber et al., 2024; Wu et al., 2024a; Wang et al., 2025), SPIn-NeRF (Mirzaei et al., 2023) handles occlusions via perceptual inpainting while CF-NeRF (Shen et al., 2022) provides uncertainty quantification. CorrFill (Liu et al., 2025a) enhances faithfulness via correspondence guidance, while AuraFusion360 (Wu et al., 2025) enables 360° scene inpainting for Gaussian Splatting. Instruct-NeRF2NeRF (Haque et al., 2023) refines NeRF views iteratively with Instruct-Pix2Pix (Brooks et al., 2023) for diffusion-driven 3D editing.

Urban scene modeling. Classic SfM-MVS pipelines extract DSMs from satellite pairs (Schönberger & Frahm, 2016; Zhang et al., 2019; Gao et al., 2023a), with MVS3D (Bosch et al., 2016) benchmarks for evaluation. Neural variants improve geometric fidelity (Derksen & Izzo, 2021; Marí et al., 2022; 2023; Zhou et al., 2024b; Leotta et al., 2019; Liu et al., 2025b; Qu & Deng, 2023; Gao et al., 2024a; Savant Aira et al., 2025; Huang et al., 2025), including Sat-NeRF (Marí et al., 2022), which utilizes NeRF for satellite imagery and SatMVS (Gao et al., 2021; 2023b) with RPC warping, yet both miss occluded facades. Generative synthesis divides into: (i) street-view methods (Li et al., 2024c; 2021; 2024d; Toker et al., 2021; Qian et al., 2023; Shi et al., 2022; Ze et al., 2025; Deng et al., 2024; Xu & Qin, 2025), including GeoDiffusion (Xiong et al., 2024) for mixed-view synthesis, Geospecific View Generation (Xu & Qin, 2024) achieving 10x resolution gains, and SkyDiffusion (Ye et al., 2024a) with Curved-BEV for street-to-satellite mapping, though lacking 3D consistency and temporal coherence; and (ii) full-3D city generation (Lin et al., 2023b; Xie et al., 2024; 2025a;b; Sun et al., 2024; Zhou et al., 2024a; Shang et al., 2024; Li et al., 2024a; Zhang et al., 2024c), with BEVFormer (Li et al., 2022) and MagicDrive (Gao et al., 2023c) using spatiotemporal transformers for view consistency. While Infinicity (Lin et al., 2023b) uses pixel-to-voxel rendering for infinite cities, and CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b) use BEV neural fields or BEV-Point splats for editable scenes, these remain constrained by input representations (semantic maps and height fields) and training distributions, limiting synthesis of realistic textures and complex structures like tunnels, bridges, and multi-level architectures. Our method uses pretrained diffusion priors to recover

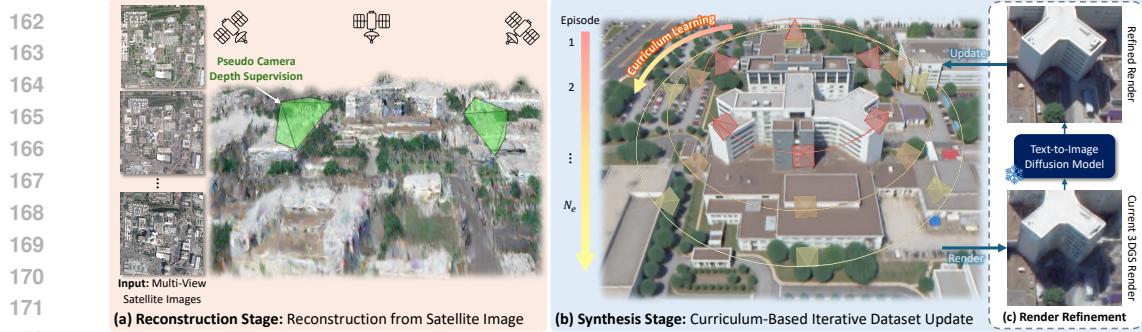


Figure 3: **Overview of the proposed Skyfall-GS pipeline.** Our method synthesizes immersive and free-flight navigable city-block scale 3D scenes solely from multi-view satellite imagery in two stages. (a) In the Reconstruction Stage, we first reconstruct the initial 3D scene using 3DGS, enhanced by pseudo-camera depth supervision to address limited parallax in satellite images. We use an appearance modeling component 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 (Labs, 2024b) with prompt-to-prompt editing (Kulikov et al., 2024). 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.

high-fidelity facades in occluded regions without dataset-specific training, respecting user constraints more faithfully.

3 METHOD

Our two-stage pipeline (Figure 3) turns satellite images into immersive 3D cities. Reconstruction Stage (Section 3.1): fit a 3D Gaussian Splatting model, adding illumination-adaptive appearance modeling and regularizers for sparse, multi-date views. Synthesis Stage (Section 3.2): recover occluded regions, e.g., facades, through curriculum 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) (Kerbl et al., 2023) 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 (Zhang et al., 2019) to approximate perspective camera parameters (extrinsic and intrinsic) from RPC and generate sparse SfM points as initial 3DGS points.

Appearance modeling. As highlighted in Section 1, multi-date satellite imagery exhibits significant appearance variations due to global illumination changes, seasonal factors, and transient objects, as illustrated in Figure 3(a). Following WildGaussians (Kulhanek et al., 2024), we use trainable per-image embeddings $\{e_j\}_{j=1}^N$ (with N training images) to handle varying illumination and atmospheric conditions. We also employ trainable per-Gaussian embeddings g_i to capture localized appearance changes like 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 0-th order spherical harmonics (SH). Finally, the transformed color \bar{c}_i

216 is then computed as $\tilde{c}_i(\mathbf{r}) = \gamma \cdot \hat{c}_i(\mathbf{r}) + \beta$, and used in the 3DGS rasterizer. To prevent modeling the
 217 appearance changes as view-dependent effects, we limit SH to zero and first-order terms.
 218

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

$$222 \quad \mathcal{L}_{\text{op}} = - \sum_i \alpha_i \log(\alpha_i) + (1 - \alpha_i) \log(1 - \alpha_i) . \quad (2)$$

$$223$$

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

227 **Pseudo camera depth supervision.** To further reduce floating artifacts, we sample pseudo-cameras
 228 positioned closer to the ground during optimization. From these pseudo-cameras, we render RGB
 229 images I_{RGB} and corresponding alpha-blended depth maps \hat{D}_{GS} . We then use an off-the-shelf
 230 monocular depth estimator, MoGe (Wang et al., 2024a), to predict scale-invariant depths \hat{D}_{est} from
 231 these renders. We use the absolute value of Pearson correlation (PCorr) to supervise the depth:
 232

$$233 \quad \mathcal{L}_{\text{depth}} = \|\text{PCorr}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}})\|_1 ; \quad \text{PCorr}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}}) = \frac{\text{Cov}(\hat{D}_{\text{GS}}, \hat{D}_{\text{est}})}{\sqrt{\text{Var}(\hat{D}_{\text{GS}})\text{Var}(\hat{D}_{\text{est}})}} . \quad (3)$$

$$234$$

$$235$$

236 **Optimization.** Combining all components, the overall loss for the reconstruction stage is defined as:
 237

$$238 \quad \mathcal{L}_{\text{sat}}(G, C) = \mathcal{L}_{\text{color}} + \lambda_{\text{op}} \mathcal{L}_{\text{op}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}} , \quad (4)$$

$$239$$

$$240$$

241 where G is the 3DGS representation, C is the set of ground-truth satellite images, λ_{op} and λ_{depth}
 242 weight opacity regularization and depth supervision relative to the color reconstruction loss.
 243

244 3.2 SYNTHESIZE VIA CURRICULUM-LEARNING BASED ITERATIVE DATASETS UPDATE

245 The iterative dataset update (IDU) technique (Haque et al., 2023; Melas-Kyriazi et al., 2024) re-
 246 peatedly executes render-edit-update cycles across multiple episodes to progressively synthesize 3D
 247 scenes. Unlike previous methods that sample camera poses from original training views (Haque
 248 et al., 2023) or simple orbits (Melas-Kyriazi et al., 2024), we introduce a curriculum-based refine-
 249 ment schedule over N_e episodes that specifically addresses satellite imagery’s geometric and visual
 250 limitations, producing structurally accurate and photorealistic reconstructions of occluded areas.
 251

252 **Curriculum learning strategy.** As illustrated in Figure 4, we observe that 3DGS trained from
 253 satellite imagery produces higher-quality renders at higher elevation angles but degenerates at lower
 254 elevation angles. Leveraging this insight, we introduce a curriculum-based synthesizing strategy,
 255 which progressively lowering viewpoints across optimization episodes. Specifically, we define N_p
 256 look-at points $\{P_i\}_{i=1}^{N_p}$ uniformly placed throughout the scene and uniformly sample N_v camera
 257 positions along orbital trajectories with controlled elevation angles and radii. Our iterative dataset
 258 update (IDU) process starts from higher elevations, progressively moving toward lower perspectives.
 259 This approach gradually reveals previously occluded regions, improving geometric detail and texture
 260 realism, as validated in our ablation studies (Section 4.2).
 261

262 **Render refinement by text-to-image diffusion model.** As illustrated in Figure 5(a), renderings
 263 from initial 3DGS contain blurry texture and artifacts. To address this, we leverage prompt-to-
 264 prompt editing with pre-trained text-to-image diffusion models to synthesize disocclusion areas,
 265 remove artifacts, and enhance geometry. Prompt-to-prompt editing (Hertz et al., 2022) modifies
 266 input images, which are described by the source prompt, to align with the target prompt while
 267 preserving structural content. Although typically used on real or diffusion-generated photos, we
 268 demonstrate its effectiveness for refining degraded satellite-trained 3DGS renders. We employ
 269 FlowEdit (Kulikov et al., 2024) with the pre-trained `FLUX.1 [dev]` diffusion model (Labs, 2024a),
 270 using prompt pairs that transform degenerate renders into high-quality imagery. Our prompts
 271 specifically describe the degraded features in original renders and specify the desired high-quality
 272 attributes in target prompts, see Section A.1 for prompts detail. As illustrated in Figure 5, this
 273 approach significantly improves the visual quality of renders, including sharper geometric details,
 274 enhanced texture richness, and physically coherent shadows, strengthening the 3DGS training dataset
 275 for more accurate reconstructions.
 276

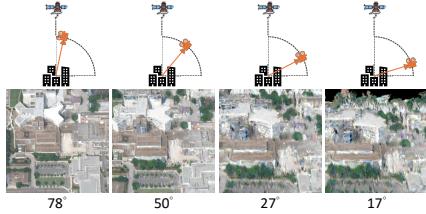


Figure 4: **The motivation of curriculum strategy.** Renderings of the initial 3D reconstruction from varied elevation angles reveal progressive degradation as the viewing angle decreases.



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

Multiple diffusion samples. While diffusion models effectively refine individual 3DGS renders, independently applying them across viewpoints introduces inconsistencies. Furthermore, 3DGS is well known to suffer from overfitting on single views, as pointed out by CoR-GS (Zhang et al., 2024b), causing artifacts when rendering from novel viewpoints.

Ideally, the optimal denoising diffusion process should produce a distribution where all views maintain synchronized 3D appearance. However, independent 2D denoising on each view does not preserve 3D consistency, resulting in a denoising trajectory distribution that is a super-set of the optimal trajectories. Selecting a single denoising trajectory from this expanded distribution has negligible probability of yielding 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.

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 (Kulikov et al., 2024) 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 Section A.1.

Algorithm 1 3DGS Refinement via Iterative Dataset Updates

```

305 Input:  $N_e$ : Number of episodes
306 Input:  $N_v, N_s, N_p$ : Number of views per point, samples per view and look-at points
307 Input:  $\{P_i\}_{i=1}^{N_p}$ : A set of  $N_p$  target look-at points
308 Input:  $\{R_i\}_{i=1}^{N_e}, \{E_i\}_{i=1}^{N_e}$ : Decreasing sequences for radius and elevation with lengths of  $N_e$ 
309 Input:  $T_{\text{src}}, T_{\text{tgt}}, n_{\min}, n_{\max}$ : FlowEdit parameters
310 Input:  $G$ : Initial 3DGS from satellite-view training
311 Output:  $G'$ : Refined 3DGS
312 1:  $G' \leftarrow G$ 
313 2: for  $i = 1$  to  $N_e$  do
314 3:    $\text{radius} \leftarrow R_i$ 
315 4:    $\text{elevation} \leftarrow E_i$ 
316 5:    $\text{cam\_views} \leftarrow \text{ORBITVIEWS}(\{P\}, \text{radius}, \text{elevation}, N_v)$             $\triangleright$  Generate  $N_p \times N_v$  views
317 6:    $\text{render\_views} \leftarrow \text{RENDER}(G', \text{cam\_views})$                                  $\triangleright$  Render RGB images
318 7:    $\text{refine\_views} \leftarrow \text{FLOWEDITREFINE}(\text{render\_views}, T_{\text{src}}, T_{\text{tgt}}, n_{\min}, n_{\max}, N_s)$      $\triangleright$  Refine
319 8:    $G' \leftarrow \text{TRAIN}(G', \text{refine\_views})$                                       $\triangleright$  Update 3DGS using refined views
320 9: end for
321 10: return  $G'$ 

```

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

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

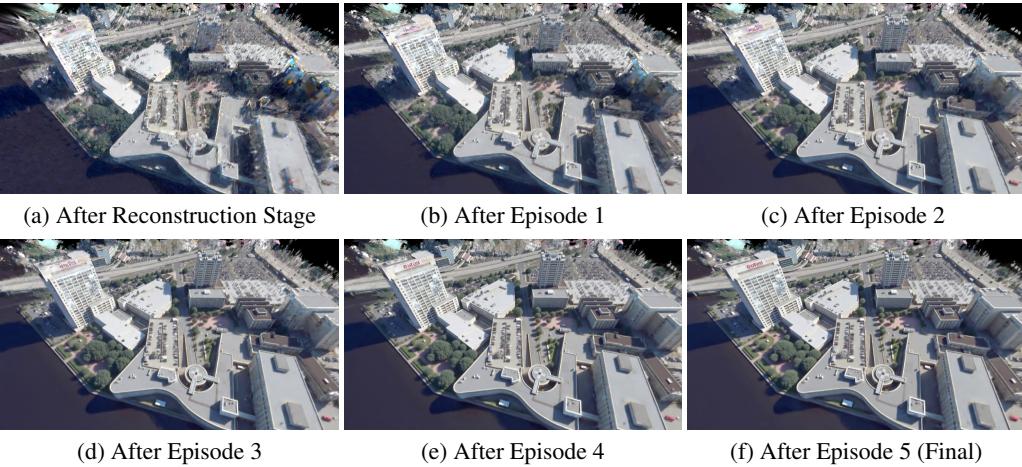


Figure 6: **Visualization of progressive refinement.** This figure illustrates the step-by-step evolution of the synthesized 3D scene. Starting from the initial reconstruction state (a), the geometry and texture are progressively refined through successive stages of the iterative process (b-e), culminating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

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 Section A.1.

4 EXPERIMENTS

Datasets. We evaluate on high-resolution RGB satellite imagery from two sources. First, the 2019 IEEE GRSS Data Fusion Contest (DFC2019) (Le Saux et al., 2019) featuring WorldView-3 captures of Jacksonville, Florida (2048×2048 pixels, 35 cm/pixel resolution). Camera parameters and sparse points were generated using SatelliteSfM (Zhang et al., 2019). We evaluate on four standard AOIs: JAX_004, JAX_068, JAX_214, and JAX_260, following Sat-NeRF (Marí et al., 2022) and EOGS (Savant Aira et al., 2025) protocols. Second, for geographic diversity, we use the GoogleEarth dataset (Xie et al., 2024) (training data for CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b)) containing NYC scenes. We use four scenes (004, 010, 219, 336) with training views rendered at an 80° elevation to approximate satellite conditions. Google Earth Studio (GES) (Google, 2024) renders serve as ground truth for both datasets. See Section A.2 for more detail about datasets.

Baselines. Our method connects satellite-based 3D reconstruction and city generation, requiring baselines from both fields. For *satellite reconstruction*, we compare with Sat-NeRF (Marí et al., 2022) and EOGS (Savant Aira et al., 2025) on DFC2019 (they require RPC input unavailable in GoogleEarth), plus Mip-Splatting (Yu et al., 2024) (enhanced with our appearance modeling) and CoR-GS Zhang et al. (2024b) on both datasets.¹ For *city generation*, we compare with CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b) on GoogleEarth (their training dataset). We use official implementations with default settings. All experiments run on a single RTX A6000 GPU.

Evaluation metrics. We primarily use distribution-based metrics to quantify quality and diversity. We report FID_{CLIP} (Kynkänniemi et al., 2023) and CMMD (Jayasumana et al., 2024) that use the CLIP (Radford et al., 2021) backbone. This is based on their observations that the InceptionV3 (Szegedy et al., 2016) used in the classic FID (Heusel et al., 2017) and KID (Binkowski et al., 2018) is unsuitable for modern generative models. We complement these with user studies for perceptual quality assessment. We also report pixel-aligned metrics (PSNR (Huynh-Thu & Ghanbari, 2008), SSIM (Wang et al., 2004), LPIPS (Zhang et al., 2018)) as secondary references. While generally unsuitable for generative tasks, these metrics are meaningful for the Google Earth dataset, where all images come from the same consistent GES 3D representation, eliminating temporal variations.

¹Many methods lack available code or models: Sat2Scene (Li et al., 2024d), Sat2Vid (Li et al., 2021), EO-NeRF (Marí et al., 2023), Sat-DN (Liu et al., 2025b), SatelliteRF (Zhou et al., 2024b), Sat-Mesh (Qu & Deng, 2023), CrossViewDiff (Li et al., 2024c), SkySplat (Huang et al., 2025), and others.

378
 379 **Table 1: Quantitative comparison of different**
 380 **methods on DFC2019 (Le Saux et al., 2019).**
 381 The results show that our method consistently
 382 achieves the best performance, indicating su-
 383 perior perceptual fidelity compared to all baselines.
 384 Metrics are computed between renders from each
 385 method and reference frames from GES.

Methods	Distribution Metrics		Pixel-level Metrics*		
	FID _{CLIP} ↓	CMMMD ↓	PSNR ↑	SSIM ↑	LPIPS ↓
<i>3D Reconstruction</i>					
Sat-NeRF (Marí et al., 2022)	88.36	4.868	10.05	0.269	0.864
EOGS (Savant Aira et al., 2025)	87.74	5.286	7.26	0.168	0.959
Mip-Splatting (Yu et al., 2024)	87.19	5.405	11.89	0.318	0.819
CoR-GS (Zhang et al., 2024b)	89.03	5.241	11.55	0.350	0.948
<i>Our Approach</i>					
Ours	27.35	2.086	12.38	0.321	0.791

391 4.1 COMPARISONS WITH BASELINES

392 **Quantitative comparison.** We evaluate against both satellite reconstruction and city generation
 393 methods using distribution-based metrics. Evaluation images are created by dividing rendered frames
 394 into 144 patches (512×512 pixels). For comparison in the DFC2019 dataset, we render GES
 395 reference videos at 17° elevation, extracting 30 frames per AOI (4,320 images total). For comparison
 396 in the GoogleEarth dataset, we use 45° elevation with 24 frames per scene (3,456 images total).
 397 We generate matching videos from all methods using identical camera parameters. Our method
 398 consistently outperforms all baselines across all metrics on both the DFC2019 and Google Earth
 399 datasets (Tables 1 and 2), demonstrating effective reconstruction across diverse urban environments.

400 **Qualitative comparison.** Figure 7(a) presents comparisons on the DFC2019 dataset against Sat-
 401 NeRF (Marí et al., 2022), EOGS (Savant Aira et al., 2025), and Mip-Splatting (Yu et al., 2024).
 402 All baselines exhibit significant distortions and blurry textures at lower viewpoints, while our
 403 baseline without IDU improves geometry but still shows floating artifacts and lacks facade detail.
 404 Our full approach achieves superior image quality. Figure 7(b) compares our approach on the
 405 GoogleEarth dataset against CityDreamer (Xie et al., 2024), GaussianCity (Xie et al., 2025b), and
 406 CoR-GS (Zhang et al., 2024b). While CityDreamer and GaussianCity generate plausible scenes,
 407 they produce oversimplified geometry and inaccurate textures, missing distinctive features such as
 408 the red pavement in scene 010 that our method correctly synthesizes. In contrast, our complete
 409 method achieves sharper building contours, enhanced texture fidelity, and reduced artifacts across
 410 both comparison scenarios. Notably, our approach successfully synthesizes plausible details for
 411 building facades occluded in the input satellite imagery and accurately reconstructs complex features
 412 including vegetation and multi-level architectures with finer surface details that better match the
 413 reference images. The visual quality approaches GES reference renders despite using only satellite
 414 imagery without ground-level data. Additional qualitative results are presented in Section A.2.

415 **User studies.** We conducted two comparative evaluations with 89 participants each: first, participants
 416 assessed the satellite input, GES reference video, Sat-NeRF, EOGS, CoR-GS, and our approach;
 417 second, participants compared the satellite input, GES reference video, CityDreamer, GaussianCity,
 418 CoR-GS, and our approach. Both studies evaluated geometric accuracy, spatial alignment, and overall
 419 quality, with full survey details in Section A.2. On the DFC2019 dataset, our method achieved
 420 dominant winrates of $\approx 97\% / 97\% / 97\%$ vs. Sat-NeRF's $\approx 3\% / 3\% / 3\%$, while EOGS and CoR-GS
 421 achieved $0\% / 0\% / 0\%$. On the GoogleEarth dataset, our approach maintained a clear advantage with
 422 $\approx 90\% / 90\% / 92\%$ winrates vs. CityDreamer's $\approx 4\% / 3\% / 3\%$, GaussianCity's $\approx 3\% / 3\% / 3\%$, and CoR-
 423 GS's $\approx 3\% / 4\% / 2\%$. These results consistently validate that our approach significantly outperforms all
 424 baselines under human perception across geometric accuracy, spatial alignment, and overall quality.

425 **Rendering efficiency.** Our method achieves 11 FPS on the modest NVIDIA T4 GPU, significantly
 426 outperforming CityDreamer's 0.18 FPS despite running on the far more powerful NVIDIA A100,
 427 which offers 5x the CUDA cores and 10x the memory bandwidth. GaussianCity reaches comparable
 428 speeds (10.72 FPS) but requires the high-end A100. Furthermore, our fused representation enables
 429 real-time rendering at 40 FPS on consumer hardware (MacBook Air M2), demonstrating that our
 430 method enables high-quality 3D urban navigation without specialized computing resources.

431 4.2 ABLATION STUDIES

We conduct ablation studies on the JAX_068 AOI.

Table 2: Quantitative comparison of different methods on GoogleEarth dataset (Xie et al., 2024). The results show that our approach consistently achieves the best performance, indicating superior perceptual fidelity compared to all baselines. Metrics are computed between renders from each method and reference frames from GES.

Methods	Distribution Metrics		Pixel-level Metrics		
	FID _{CLIP} ↓	CMMMD ↓	PSNR ↑	SSIM ↑	LPIPS ↓
<i>City Generation</i>					
CityDreamer (Xie et al., 2024)	36.52	4.152	12.58	0.267	0.558
GaussianCity (Xie et al., 2025b)	28.73	2.917	13.41	0.291	0.541
<i>3D Reconstruction</i>					
CoR-GS (Zhang et al., 2024b)	27.32	3.752	12.85	0.291	0.455
<i>Our Approach</i>					
Ours	9.91	2.009	14.28	0.298	0.394



Figure 7: **Qualitative comparison on (a) DFC2019 and (b) GoogleEarth datasets.** The leftmost column shows one representative example of the input satellite images. Our method outperforms all baselines in geometric accuracy and texture quality in low-altitude novel views, demonstrating enhanced building geometry, detailed facades, and reduced floating artifacts. Notably, our approach correctly preserves distinctive features such as the red pavement in scene 010 that competing methods miss. Unlike CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b), our method operates directly on satellite imagery without requiring pixel-aligned semantic maps or height-fields, enabling synthesis of complex geometric structures that more closely match GES references.

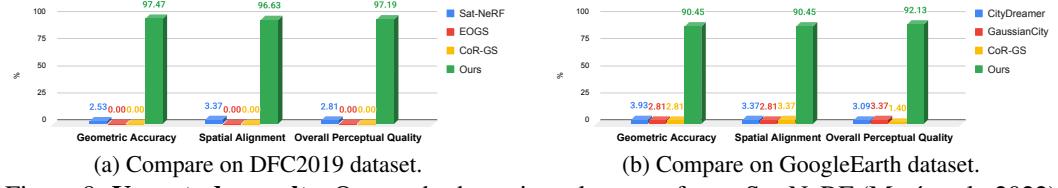


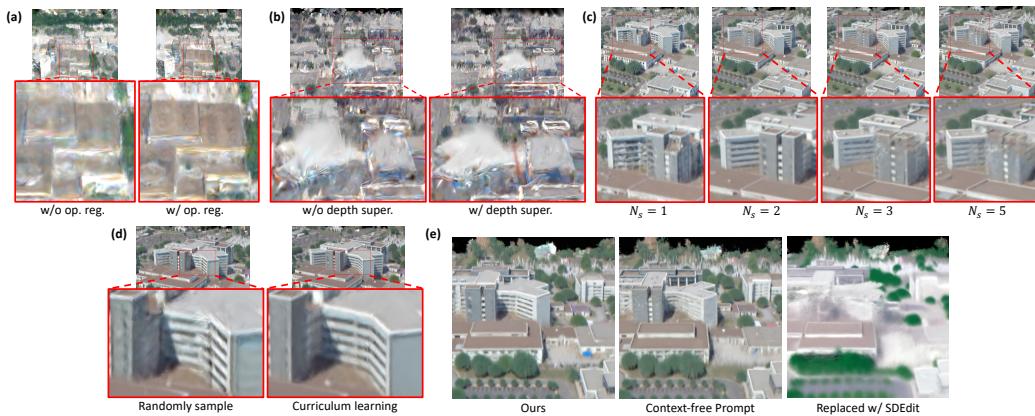
Figure 8: **User study results.** Our method consistently outperforms Sat-NeRF (Marí et al., 2022), EOGS (Savant Aira et al., 2025), CoR-GS (Zhang et al., 2024b), CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b), achieving particularly high scores in geometric accuracy and overall perceptual quality. (a) details the comparison on the DFC2019 dataset (Le Saux et al., 2019), while subfigure (b) details the comparison on the GoogleEarth dataset (Xie et al., 2024).

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 the quality of renders during the IDU process, rather than testing the final low-angle performance. Appearance modeling is crucial for multi-date convergence, opacity regularization removes floating artifacts (Figure 9(a)), and depth supervision flattens planar regions (Figure 9(b)). Together, they yield the lowest FID_{CLIP}/CMMD scores. Furthermore, we validate geometric accuracy using LiDAR data from the DFC2019 dataset (Le Saux et al., 2019). To quantify this, we unproject 3DGS depth renders into point clouds and rasterize them into Digital Surface Models (DSMs) for comparison. Our results show that both opacity regularization and pseudo-depth supervision improve geometric accuracy, with their combination achieving the lowest MAE/RMSE.

Ablation on the synthesis stage. We isolate two key factors: multi-sample diffusion and curriculum view progression. As Figure 9(c) shows, $N_s = 2$ achieves the optimal visual results. Although $N_s = 5$ yields the lowest CMMD, it requires a $1.5\times$ increase in training time with marginal returns in quality; thus, we adopt $N_s = 2$ for all experiments. Additionally, Figure 9(d) highlights that employing a curriculum strategy (vs. random views) effectively restores geometry in occluded areas, a benefit confirmed by Table 4. We further benchmark our refinement module against the SDEdit (Meng et al., 2022) baseline. As evident in Figure 9(e), SDEdit causes significant degradation, primarily due to its inability to hallucinate details while maintaining the structural integrity defined by the satellite imagery. Finally, we evaluate prompt sensitivity by utilizing generic context-free prompts. The negligible visual difference in Figure 9(e) confirms that our method is robust to prompt

486
487
488
489
490
Table 3: **Ablation on the reconstruction stage.** Appearance
491 modeling secures convergence. Opacity regularization
492 and depth supervision enhance visual fidelity and geo-
493 metric accuracy.
494

Components			Perceptual Metrics		Geometric Metrics	
App. Mod.	Op. Reg.	Depth Sup.	FID _{CLIP} ↓	CMMD ↓	MAE (m) ↓	RMSE (m) ↓
✗	✗	✗	Failed	Failed	Failed	Failed
✓	✗	✗	41.90	2.45	3.542	5.218
✓	✓	✗	39.95	2.40	2.980	4.527
✓	✓	✓	38.01	2.31	2.250	3.483



511
512
513
514
515
516
517
518
519
520
521
Figure 9: **Satellite-view training and IDU refinement ablation.** (a) Opacity regularization reduces
522 floating artifacts and yields denser reconstructions. (b) Pseudo-camera depth supervision improves
523 geometry in planar, texture-less areas like rooftops and roads. (c) Using multiple diffusion samples
524 per view enhances texture consistency and reduces high-frequency geometric noise, $N_s = 2$ achieves
525 the optimal visual results. (d) Curriculum learning progressively introduces challenging views,
526 significantly improving geometric coherence in previously occluded regions compared to random
527 sampling. (e) Refinement analysis: Using a generic context-free prompt results in a minor degradation
528 of facade details but maintains structure, demonstrating robustness. In contrast, replacing our
529 refinement method with SDEdit leads to a severe drop in quality, as standard noising-denoising
530 struggles to hallucinate details while preserving the underlying geometry defined by the satellite
531 imagery.

532 engineering and driven primarily by the diffusion model’s internal priors. Please refer to Table 10 for
533 the specific text prompts.

5 CONCLUSION

534 Skyfall-GS synthesizes real-time, immersive 3D urban scenes from multi-view satellite imagery,
535 using 3D Gaussian Splatting and text-to-image diffusion models in a curriculum-based iterative
536 refinement approach. Our method surpasses existing methods like Sat-NeRF, , EOGS, CityDreamer,
537 and GaussianCity, effectively addressing challenges such as limited parallax, illumination variations,
538 and occlusions. Future work includes scaling to larger environments and dynamic scenes.

539 **Limitations.** Our method requires significant computational resources, primarily due to the refine-
540 ment process. The fixed heuristic camera trajectory creates blind spots in complex urban geo-
541 metries, particularly in heavily occluded regions and scene boundaries. This results in artifacts and
542 over-smoothed textures at extreme street-level perspectives. Additionally, our hybrid reconstruc-
543 tion-generation framework requires off-nadir satellite views. It cannot synthesize facades from purely
544 top-down (nadir) imagery.

545
546
547
548
549
550
Table 4: **Ablation on the synthesis stage.** We evaluate sample counts (N_s),
551 core components, and compare against baselines.

Method Variation	FID _{CLIP} ↓	CMMD ↓	Time (h)
<i>Multiple Samples (N_s)</i>			
$N_s = 1$	34.11	3.19	3.44
Ours ($N_s = 2$)	28.35	2.88	6.37
$N_s = 3$	28.64	2.77	7.19
$N_s = 5$	29.17	2.68	9.80
<i>Component Ablation</i>			
w/o Curriculum	33.79	3.36	-
w/ Context-free Pmt.	30.78	2.98	-
Replaced w/ SDEdit	64.74	4.14	-

540
541
ETHICS STATEMENT542
543
544
545
546
547
This work included a small-scale user study where anonymous participants were asked to compare
our results with baselines through an online survey. No personally identifiable information was
collected, and all responses were stored anonymously. Participation was entirely voluntary, and no
risks were posed to participants. The study did not require institutional review board (IRB) approval
under our institution’s policies, as it involved only anonymous survey responses with minimal risk.548
549
REPRODUCIBILITY STATEMENT550
551
552
553
554
555
556
557
We have taken several steps to ensure the reproducibility of our work. Implementation details of
our reconstruction and synthesis pipeline are provided in Section 3, including the architecture, loss
functions, and optimization objectives. All hyperparameters, training schedules, and regularization
terms are described in Section 3 and Section A.1. Details of datasets, splits, and evaluation protocols
are described in Section 4 and Section A.2, with clear references to the publicly available DFC2019
dataset and GoogleEarth dataset. Details of the user study are described in Section A.2. We will
release our source code upon acceptance to further support transparency and reproducibility.558
559
REFERENCES560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and
Pratul P. Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance
fields, 2021.
Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF
360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.
Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-NeRF:
Anti-aliased grid-based neural radiance fields. In *ICCV*, 2023.
Mikołaj Binkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD
GANs. In *International Conference on Learning Representations*, 2018.
Marc Bosch, Zachary Kurtz, Shea Hagstrom, and Myron Brown. A multiple view stereo benchmark
for satellite imagery. In *2016 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pp.
1–9. IEEE, 2016.
Tim Brooks, Aleksander Holynski, and Alexei A Efros. InstructPix2Pix: Learning to follow image
editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
Recognition*, pp. 18392–18402, 2023.
Hao Chen, Jiafu Wu, Ying Jin, Jinlong Peng, Xiaofeng Mao, Mingmin Chi, Mufeng Yao, Bo Peng,
Jian Li, and Yun Cao. VI3DRM: Towards meticulous 3D reconstruction from sparse views via
photo-realistic novel view synthesis, 2024. URL <https://arxiv.org/abs/2409.08207>.
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.
Hiba Dahmani, Moussab Bennehar, Nathan Piasco, Luis Roldao, and Dzmitry Tsishkou. SWAG:
Splatting in the wild images with appearance-conditioned gaussians, 2024. URL <https://arxiv.org/abs/2403.10427>.
Boyang Deng, Richard Tucker, Zhengqi Li, Leonidas Guibas, Noah Snavely, and Gordon Wetzstein.
Streetscapes: Large-scale consistent street view generation using autoregressive video diffusion.
In *SIGGRAPH 2024 Conference Papers*, 2024.
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 (CVPR
Workshops*, pp. 1152–1161, June 2021.

594 Jan-Niklas Dihlmann, Andreas Engelhardt, and Hendrik P.A. Lensch. SIGNeRF: Scene integrated
 595 generation for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer*
 596 *Vision and Pattern Recognition (CVPR)*, 2024.

597

598 Cheng-De Fan, Chen-Wei Chang, Yi-Ruei Liu, Jie-Ying Lee, Jiun-Long Huang, Yu-Chee Tseng, and
 599 Yu-Lun Liu. Spectromotion: Dynamic 3d reconstruction of specular scenes. In *Proceedings of the*
 600 *Computer Vision and Pattern Recognition Conference*, pp. 21328–21338, 2025.

601 Zhiwen Fan, Wenyang Cong, Kairun Wen, Kevin Wang, Jian Zhang, Xinghao Ding, Danfei Xu, Boris
 602 Ivanovic, Marco Pavone, Georgios Pavlakos, et al. InstantSplat: Unbounded sparse-view pose-free
 603 gaussian splatting in 40 seconds. *arXiv preprint arXiv:2403.20309*, 2(3):4, 2024.

604

605 Jiemin Fang, Junjie Wang, Xiaopeng Zhang, Lingxi Xie, and Qi Tian. GaussianEditor: Editing 3D
 606 gaussians delicately with text instructions. In *CVPR*, 2024.

607

608 Jian Gao, Jin Liu, and Shunping Ji. Rational polynomial camera model warping for deep learning-
 609 based satellite multi-view stereo matching. In *Proceedings of the IEEE/CVF international*
 610 *conference on computer vision*, pp. 6148–6157, 2021.

611

612 Jian Gao, Jin Liu, and Shunping Ji. A general deep learning based framework for 3D recon-
 613 struction from multi-view stereo satellite images. *ISPRS Journal of Photogrammetry and Re-
 614 mote Sensing*, 195:446–461, 2023a. ISSN 0924-2716. doi: <https://doi.org/10.1016/j.isprsjprs.2022.12.012>. URL <https://www.sciencedirect.com/science/article/pii/S0924271622003276>.

615

616 Jian Gao, Jin Liu, and Shunping Ji. A general deep learning based framework for 3d reconstruc-
 617 tion from multi-view stereo satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing*,
 618 195:446–461, 2023b.

619

620 Kyle Gao, Dening Lu, Hongjie He, Linlin Xu, and Jonathan Li. Enhanced 3D urban scene recon-
 621 struction and point cloud densification using gaussian splatting and google earth imagery, 2024a.
 622 URL <https://arxiv.org/abs/2405.11021>.

623

624 Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P.
 625 Srinivasan, Jonathan T. Barron, and Ben Poole. CAT3D: Create anything in 3D with multi-view
 626 diffusion models. *Advances in Neural Information Processing Systems*, 2024b.

627

628 Ruiyan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, and Qiang Xu. Mag-
 629 icdrive: Street view generation with diverse 3d geometry control. *arXiv preprint arXiv:2310.02601*,
 630 2023c.

631

632 Google. Google earth studio. <https://earth.google.com/studio>, 2024. Accessed:
 633 2025-05-14.

634

635 Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-
 636 NeRF2NeRF: Editing 3D scenes with instructions. In *Proceedings of the IEEE/CVF International*
 637 *Conference on Computer Vision*, 2023.

638

639 Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-
 640 to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.

641

642 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 643 GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances*
 644 *in Neural Information Processing Systems*, pp. 6626–6637, 2017.

645

646 Hao-Yu Hou, Chia-Chi Hsu, Yu-Chen Huang, Mu-Yi Shen, Wei-Fang Sun, Cheng Sun, Chia-
 647 Che Chang, Yu-Lun Liu, and Chun-Yi Lee. 3d gaussian splatting with grouped uncertainty for
 648 unconstrained images. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech*
 649 *and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025.

650

651 Xuejun Huang, Xinyi Liu, Yi Wan, Zhi Zheng, Bin Zhang, Mingtao Xiong, Yingying Pei, and
 652 Yongjun Zhang. SkySplat: Generalizable 3d gaussian splatting from multi-temporal sparse satellite
 653 images. *arXiv preprint arXiv:2508.09479*, 2025.

648 Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of PSNR in image/video quality
 649 assessment. *Electronics Letters*, 44(13):800–801, 2008.

650

651 Sadeep Jayasumana, Sri Kumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and
 652 Sanjiv Kumar. Rethinking FID: Towards a better evaluation metric for image generation. In *CVPR*,
 653 2024.

654 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting
 655 for real-time radiance field rendering. *ACM TOG*, 2023.

656 Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and
 657 George Drettakis. A hierarchical 3D gaussian representation for real-time rendering of very large
 658 datasets. *ACM TOG*, 2024.

659

660 Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. WildGaus-
 661 sians: 3D gaussian splatting in the wild. *NeurIPS*, 2024.

662 Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. FlowEdit:
 663 Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629*,
 664 2024.

665

666 Tuomas Kynkänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of
 667 ImageNet classes in fréchet inception distance. In *Proc. ICLR*, 2023.

668 Black Forest Labs. Official weights of FLUX.1 dev. <https://huggingface.co/black-forest-labs/FLUX.1-dev>, 2024a. Accessed: 2025-02-28.

669

670 Black Forest Labs. FLUX. <https://github.com/black-forest-labs/flux>, 2024b.
 671 Accessed: 2025-02-28.

672

673 Bertrand Le Saux, Naoto Yokoya, Ronny Hänsch, and Myron Brown. Data fusion contest 2019
 674 (DFC2019), 2019. URL <https://dx.doi.org/10.21227/c6tm-vw12>.

675

676 Matthew J. Leotta, Chengjiang Long, Bastien Jacquet, Matthieu Zins, Dan Lipsa, Jie Shan, Bo Xu,
 677 Zhixin Li, Xu Zhang, Shih-Fu Chang, Matthew Purri, Jia Xue, and Kristin Dana. Urban semantic
 678 3D reconstruction from multiview satellite imagery. In *Proceedings of the IEEE/CVF Conference
 679 on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.

680

681 Haoran Li, Haolin Shi, Wenli Zhang, Wenjun Wu, Yong Liao, Lin Wang, Lik-hang Lee, and
 682 Peng Yuan Zhou. DreamScene: 3D gaussian-based text-to-3D scene generation via formation
 683 pattern sampling. In *European Conference on Computer Vision*, pp. 214–230. Springer, 2024a.

684

685 Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. DNGaussian:
 686 Optimizing sparse-view 3D gaussian radiance fields with global-local depth normalization. *arXiv
 687 preprint arXiv:2403.06912*, 2024b.

688

689 Weijia Li, Jun He, Junyan Ye, Huaping Zhong, Zhimeng Zheng, Zilong Huang, Dahua Lin, and
 690 Conghui He. CrossViewDiff: A cross-view diffusion model for satellite-to-street view synthesis,
 691 2024c. URL <https://arxiv.org/abs/2408.14765>.

692

693 Z Li, W Wang, H Li, E Xie, C Sima, T Lu, Q Yu, and J Dai. Bevformer: Learning bird’s-eye-view
 694 representation from multi-camera images via spatiotemporal transformers. arxiv 2022. In *ECCV*,
 695 2022.

696

697 Zuoyue Li, Zhenqiang Li, Zhaopeng Cui, Rongjun Qin, Marc Pollefeys, and Martin R. Oswald.
 698 Sat2Vid: Street-view panoramic video synthesis from a single satellite image, 2021. URL <https://arxiv.org/abs/2012.06628>.

699

700 Zuoyue Li, Zhenqiang Li, Zhaopeng Cui, Marc Pollefeys, and Martin R. Oswald. Sat2Scene: 3D
 701 urban scene generation from satellite images with diffusion. In *Proceedings of the IEEE/CVF
 702 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7141–7150, June 2024d.

703

704 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten
 705 Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3D: High-resolution text-to-3D content
 706 creation. In *CVPR*, 2023a.

702 Chieh Hubert Lin, Hsin-Ying Lee, Willi Menapace, Menglei Chai, Aliaksandr Siarohin, Ming-Hsuan
 703 Yang, and Sergey Tulyakov. InfiniCity: Infinite-scale city synthesis. In *Proceedings of the*
 704 *IEEE/CVF International Conference on Computer Vision*, 2023b.

705

706 Chin-Yang Lin, Chung-Ho Wu, Chang-Han Yeh, Shih-Han Yen, Cheng Sun, and Yu-Lun Liu.
 707 Frugalnerf: Fast convergence for extreme few-shot novel view synthesis without learned priors. In
 708 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 11227–11238, 2025.

709

710 Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu,
 711 Songcen Xu, Youliang Yan, and Wenming Yang. VastGaussian: Vast 3D gaussians for large scene
 712 reconstruction, 2024. URL <https://arxiv.org/abs/2402.17427>.

713

714 Kuan-Hung Liu, Cheng-Kun Yang, Min-Hung Chen, Yu-Lun Liu, and Yen-Yu Lin. Corrfill: En-
 715 hancing faithfulness in reference-based inpainting with correspondence guidance in diffusion
 716 models. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp.
 717 1618–1627. IEEE, 2025a.

718

719 Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick.
 720 Zero-1-to-3: Zero-shot one image to 3D object. In *CVPR*, 2023a.

721

722 Tianle Liu, Shuangming Zhao, Wanshou Jiang, and Bingxuan Guo. Sat-DN: Implicit surface
 723 reconstruction from multi-view satellite images with depth and normal supervision, 2025b. URL
 724 <https://arxiv.org/abs/2502.08352>.

725

726 Xinhang Liu, Jiaben Chen, Shiu-hong Kao, Yu-Wing Tai, and Chi-Keung Tang. Deceptive-
 727 NeRF/3DGS: Diffusion-generated pseudo-observations for high-quality sparse-view reconstruction.
 728 *arXiv preprint arXiv:2305.15171*, 2023b.

729

730 Yang Liu, Chuanchen Luo, Zhongkai Mao, Junran Peng, and Zhaoxiang Zhang. CityGaussianV2:
 731 Efficient and geometrically accurate reconstruction for large-scale scenes, 2024. URL <https://arxiv.org/abs/2411.00771>.

732

733 Yang Liu, Chuanchen Luo, Lue Fan, Naiyan Wang, Junran Peng, and Zhaoxiang Zhang. CityGaussian:
 734 Real-time high-quality large-scale scene rendering with gaussians. In *European Conference on*
 735 *Computer Vision*, pp. 265–282. Springer, 2025c.

736

737 Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Sat-NeRF: Learning multi-view satellite pho-
 738 togrammetry with transient objects and shadow modeling using RPC cameras. In *2022 IEEE/CVF*
 739 *Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 1310–1320,
 740 2022.

741

742 Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Multi-date earth observation NeRF: The detail
 743 is in the shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 744 *Recognition (CVPR) Workshops*, pp. 2034–2044, June 2023.

745

746 Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy,
 747 and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo
 748 collections. In *CVPR*, 2021.

749

750 Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, Natalia Neverova, Andrea Vedaldi, Oran Gafni,
 751 and Filippos Kokkinos. IM-3D: Iterative multiview diffusion and reconstruction for high-quality
 752 3D generation. *International Conference on Machine Learning*, 2024.

753

754 Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon.
 755 SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International*
 756 *Conference on Learning Representations*, 2022.

757

758 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 759 Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications*
 760 *of the ACM*, 2021.

756 Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G Derpanis, Jonathan Kelly, Marcus A
 757 Brubaker, Igor Gilitschenski, and Alex Levinstein. Spin-NeRF: Multiview segmentation and
 758 perceptual inpainting with neural radiance fields. In *Proceedings of the IEEE/CVF Conference on*
 759 *Computer Vision and Pattern Recognition*, pp. 20669–20679, 2023.

760 Ashkan Mirzaei, Tristan Aumentado-Armstrong, Marcus A. Brubaker, Jonathan Kelly, Alex Levin-
 761 shtein, Konstantinos G. Derpanis, and Igor Gilitschenski. Watch your steps: Local image and
 762 scene editing by text instructions. In *ECCV*, 2024.

764 Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast
 765 image inversion for editing with text-guided diffusion models, 2024. URL <https://arxiv.org/abs/2305.16807>.

767 Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for
 768 editing real images using guided diffusion models. *arXiv preprint arXiv:2211.09794*, 2022.

770 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives
 771 with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022.
 772 doi: 10.1145/3528223.3530127. URL <https://doi.org/10.1145/3528223.3530127>.

773 Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and
 774 Noha Radwan. RegNeRF: Regularizing neural radiance fields for view synthesis from sparse
 775 inputs. In *CVPR*, 2022.

776 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D
 777 diffusion. *arXiv preprint arXiv:2209.14988*, 2022.

779 Ming Qian, Jincheng Xiong, Gui-Song Xia, and Nan Xue. Sat2Density: Faithful density learning
 780 from satellite-ground image pairs. In *ICCV*, 2023.

782 Yingjie Qu and Fei Deng. Sat-Mesh: Learning neural implicit surfaces for multi-view satellite
 783 reconstruction. *Remote Sensing*, 15(17), 2023. ISSN 2072-4292. doi: 10.3390/rs15174297. URL
 784 <https://www.mdpi.com/2072-4292/15/17/4297>.

785 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 786 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 787 models from natural language supervision. In *Proceedings of the 38th International Conference on*
 788 *Machine Learning*, pp. 8748–8763. PMLR, 2021.

789 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 790 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Confer-
 791 ence on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

793 Sara Sabour, Lily Goli, George Kopanas, Mark Matthews, Dmitry Lagun, Leonidas Guibas, Alec
 794 Jacobson, David J. Fleet, and Andrea Tagliasacchi. SpotLessSplats: Ignoring distractors in 3D
 795 gaussian splatting. *arXiv preprint arXiv:2406.20055*, 2024.

796 Luca Savant Aira, Gabriele Facciolo, and Thibaud Ehret. Gaussian splatting for efficient satellite
 797 image photogrammetry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 798 *Pattern Recognition*, 2025.

800 Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference
 801 on Computer Vision and Pattern Recognition (CVPR)*, 2016.

802 Yu Shang, Yuming Lin, Yu Zheng, Hangyu Fan, Jingtao Ding, Jie Feng, Jiansheng Chen, Li Tian,
 803 and Yong Li. UrbanWorld: An urban world model for 3D city generation. *arXiv preprint
 804 arXiv:2407.11965*, 2024. URL <https://arxiv.org/abs/2407.11965>.

805 Jianxiong Shen, Antonio Agudo, Francesc Moreno-Noguer, and Adria Ruiz. Conditional-flow
 806 NeRF: Accurate 3d modelling with reliable uncertainty quantification. In *European Conference on*
 807 *Computer Vision*, pp. 540–557. Springer, 2022.

809 Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. MVDream: Multi-view
 diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023.

810 Yujiao Shi, Dylan Campbell, Xin Yu, and Hongdong Li. Geometry-guided street-view panorama
 811 synthesis from satellite imagery. In *IEEE TPAMI*, 2022.

812

813 Michael Sprintson, Rama Chellappa, and Cheng Peng. FusionRF: High-fidelity satellite neural
 814 radiance fields from multispectral and panchromatic acquisitions. *arXiv preprint arXiv:2409.15132*,
 815 2024.

816 Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang.
 817 DimensionX: Create any 3D and 4D scenes from a single image with controllable video diffusion.
 818 *arXiv preprint arXiv:2411.04928*, 2024.

819

820 Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking
 821 the inception architecture for computer vision. In *CVPR*, 2016.

822

823 Matthew Tancik, Vincent Casser, Xincheng Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan,
 824 Jonathan T Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view
 825 synthesis. In *CVPR*, 2022.

826 Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative
 827 gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023.

828 Aysim Toker, Qunjie Zhou, Maxim Maximov, and Laura Leal-Taixe. Coming down to earth: Satellite-
 829 to-street view synthesis for geo-localization. In *Proceedings of the IEEE/CVF Conference on
 830 Computer Vision and Pattern Recognition (CVPR)*, pp. 6488–6497, June 2021.

831 Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-NeRF: Scalable construction
 832 of large-scale NeRFs for virtual fly-throughs. In *CVPR*, 2022.

833

834 Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang.
 835 MoGe: Unlocking accurate monocular geometry estimation for open-domain images with optimal
 836 training supervision, 2024a. URL <https://arxiv.org/abs/2410.19115>.

837 Yuxuan Wang, Xuanyu Yi, Zike Wu, Na Zhao, Long Chen, and Hanwang Zhang. View-consistent
 838 3D editing with gaussian splatting, 2025. URL <https://arxiv.org/abs/2403.11868>.

839

840 Yuze Wang, Junyi Wang, and Yue Qi. WE-GS: An in-the-wild efficient 3D gaussian representation for
 841 unconstrained photo collections, 2024b. URL <https://arxiv.org/abs/2406.02407>.

842

843 Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Pro-
 844 lificDreamer: High-fidelity and diverse text-to-3d generation with variational score distillation.
 845 *Advances in neural information processing systems*, 36:8406–8441, 2023.

846

847 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: From
 848 error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612,
 849 2004.

850

851 Ethan Weber, Aleksander Holynski, Varun Jampani, Saurabh Saxena, Noah Snavely, Abhishek Kar,
 852 and Angjoo Kanazawa. NeRFiller: Completing scenes via generative 3D inpainting. In *CVPR*,
 853 2024.

854

855 Chung-Ho Wu, Yang-Jung Chen, Ying-Huan Chen, Jie-Ying Lee, Bo-Hsu Ke, Chun-Wei Tuan Mu,
 856 Yi-Chuan Huang, Chin-Yang Lin, Min-Hung Chen, Yen-Yu Lin, and Yu-Lun Liu. AuraFusion360:
 857 Augmented unseen region alignment for reference-based 360° unbounded scene inpainting, 2025.
 858 URL <https://arxiv.org/abs/2502.05176>.

859

860 Jing Wu, Jia-Wang Bian, Xinghui Li, Guangrun Wang, Ian Reid, Philip Torr, and Victor Prisacariu.
 861 GaussCtrl: Multi-view consistent text-driven 3D gaussian splatting editing. *ECCV*, 2024a.

862

863 Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P.
 864 Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. ReconFusion:
 865 3D reconstruction with diffusion priors. *arXiv preprint arXiv:2312.02981*, 2023.

866

867 Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T. Barron, and
 868 Aleksander Holynski. CAT4D: Create anything in 4D with multi-view video diffusion models,
 869 2024b. URL <https://arxiv.org/abs/2411.18613>.

864 Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. CityDreamer: Compositional generative
 865 model of unbounded 3D cities. In *CVPR*, 2024.

866

867 Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. CityDreamer4D: Compositional generative
 868 model of unbounded 4D cities. *arXiv preprint arXiv:2501.08983*, 2025a.

869 Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. Generative gaussian splatting for
 870 unbounded 3D city generation. In *CVPR*, 2025b.

871

872 Zhexiao Xiong, Xin Xing, Scott Workman, Subash Khanal, and Nathan Jacobs. Mixed-view panorama
 873 synthesis using geospatially guided diffusion. *Transactions on Machine Learning Research*, 2024.

874 Jiacong Xu, Yiqun Mei, and Vishal M. Patel. Wild-GS: Real-time novel view synthesis from
 875 unconstrained photo collections, 2024. URL <https://arxiv.org/abs/2406.10373>.

876

877 Ningli Xu and Rongjun Qin. Geospecific view generation geometry-context aware high-resolution
 878 ground view inference from satellite views. In *European Conference on Computer Vision*, pp.
 879 349–366. Springer, 2024.

880 Ningli Xu and Rongjun Qin. Satellite to GroundScape—large-scale consistent ground view generation
 881 from satellite views. *arXiv preprint arXiv:2504.15786*, 2025.

882 Junyan Ye, Jun He, Weijia Li, Zhutao Lv, Jinhua Yu, Haote Yang, and Conghui He. Skydiffusion:
 883 Street-to-satellite image synthesis with diffusion models and bev paradigm. *arXiv e-prints*, pp.
 884 arXiv–2408, 2024a.

885

886 Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit
 887 anything in 3D scenes. In *ECCV*, 2024b.

888 Taoran Yi, Jiemin Fang, Junjie Wang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian,
 889 and Xinggang Wang. GaussianDreamer: Fast generation from text to 3d gaussians by bridging 2d
 890 and 3d diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 891 *Pattern Recognition*, pp. 6796–6807, 2024.

892

893 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-Splatting: Alias-free
 894 3D gaussian splatting. In *CVPR*, 2024.

895

896 Xianghui Ze, Zhenbo Song, Qiwei Wang, Jianfeng Lu, and Yujiao Shi. Controllable satellite-to-
 897 street-view synthesis with precise pose alignment and zero-shot environmental control, 2025. URL
 898 <https://arxiv.org/abs/2502.03498>.

899

900 Yu-Ting Zhan, Cheng-Yuan Ho, Hebi Yang, Yi-Hsin Chen, Jui Chiu Chiang, Yu-Lun Liu, and
 901 Wen-Hsiao Peng. Cat-3dgs: A context-adaptive triplane approach to rate-distortion-optimized
 902 3dgs compression. *arXiv preprint arXiv:2503.00357*, 2025.

903

904 Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang.
 Gaussian in the wild: 3D gaussian splatting for unconstrained image collections. *arXiv preprint*
 905 *arXiv:2403.15704*, 2024a.

906

907 Jiawei Zhang, Jiahe Li, Xiaohan Yu, Lei Huang, Lin Gu, Jin Zheng, and Xiao Bai. CoR-GS:
 908 Sparse-view 3d gaussian splatting via co-regularization. In *ECCV*, 2024b.

909

910 Kai Zhang, Jin Sun, and Noah Snavely. Leveraging vision reconstruction pipelines for satellite
 911 imagery. In *IEEE International Conference on Computer Vision Workshops*, 2019.

912

913 Lulin Zhang and Ewelina Rupnik. Sparsesat-NeRF: Dense depth supervised neural radiance fields
 914 for sparse satellite images. *arXiv preprint arXiv:2309.00277*, 2023.

915

916 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 917 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on*
 918 *Computer Vision and Pattern Recognition*, pp. 586–595, 2018.

919

920 Shougao Zhang, Mengqi Zhou, Yuxi Wang, Chuanchen Luo, Rongyu Wang, Yiwei Li, Zhaoxiang
 921 Zhang, and Junran Peng. CityX: Controllable procedural content generation for unbounded 3D
 922 cities, 2024c. URL <https://arxiv.org/abs/2407.17572>.

918 Mengqi Zhou, Yuxi Wang, Jun Hou, Shougao Zhang, Yiwei Li, Chuanchen Luo, Junran Peng, and
919 Zhaoxiang Zhang. SceneX: Procedural controllable large-scale scene generation via large-language
920 models, 2024a. URL <https://arxiv.org/abs/2403.15698>.

921
922 Xin Zhou, Yang Wang, Daoyu Lin, Zehao Cao, Biqing Li, and Junyi Liu. SatelliteRF: Accelerating
923 3D reconstruction in multi-view satellite images with efficient neural radiance fields. *Applied*
924 *Sciences*, 14(7), 2024b. ISSN 2076-3417. doi: 10.3390/app14072729. URL <https://www.mdpi.com/2076-3417/14/7/2729>.

925
926 Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. FSGS: Real-time few-shot view
927 synthesis using gaussian splatting, 2023.

928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

972 **A APPENDIX**
973974 This supplementary material provides additional details that complement our main paper. We include:
975

- 976 **Implementation Details:** This section details the pseudo-camera depth supervision strategy,
977 3DGS reconstruction parameters for satellite imagery, and the FlowEdit-based refinement
978 process. **We also provide a detailed breakdown of training time and memory consumption.**
979 **Furthermore, we include a discussion on the validity of the RPC to perspective camera**
980 **model approximation, including quantitative error analysis.**
- 981 **Experiments Detail:** We provide dataset details, including training image counts and geo-
982 graphical coordinates for each Area of Interest (AOI), alongside the user study methodology
983 and evaluation protocol.
- 984 **Additional Qualitative Results:** We present extended visual comparisons with state-of-
985 the-art methods and results on four additional AOIs of the DFC2019 dataset. **Furthermore,**
986 **we provide visualizations of renders conditioned on varying per-image embeddings e_j to**
987 **illustrate temporal stability.**
- 988 **Additional Experiments & Results:** We encompass a comprehensive set of new experi-
989 ments, including: (i) synthesis results for complex, irregular geometries (e.g., castles and
990 cathedrals) to demonstrate the framework’s robustness; (ii) a sensitivity analysis of refine-
991 ment text prompts; (iii) an episode-vs-coverage analysis to quantify the effectiveness of the
992 curriculum strategy; and (iv) synthesized results with different random seed.

993 Additionally, we provide an interactive HTML visualization (available in the folder, `main.html`)
994 that allows readers to explore our video results and compare reconstructions across different viewing
995 conditions and scenes. This visualization enables direct comparison of our method’s geometric
996 accuracy, spatial alignment, and overall perceptual quality against baseline approaches and Google
997 Earth Studio reference video.

998 We also provide example datasets via Zenodo, which can be accessed at this URL. However, due to
999 storage limitations, we only provide training data for an AOI as an example. We plan to release the
1000 complete dataset upon acceptance.

1002 **A.1 IMPLEMENTATION DETAILS**
1003

1004 **Codebase.** Our method extends the Mip-Splatting (Yu et al., 2024) codebase with custom modules
1005 for satellite imagery processing and our curriculum-based IDU refinement pipeline.

1006 **Pseudo camera depth supervision.** We sample cameras with varied azimuths and decreasing
1007 elevations, using random per-image embeddings. MoGe (Wang et al., 2024a) provides scale-invariant
1008 depth estimation. We sample 24 views every 10 iterations, with look-at points (x, y, z) , where
1009 $x, y \sim \mathcal{N}(0, 128)$ and $z = 0$, as illustrated in Figure 10. Camera azimuths are uniformly sampled
1010 between 0 and 2π , while elevation angles and radii linearly decrease from 80° to 45° and 300 to 250
1011 units, respectively. Rendered RGB images (I_{RGB}) are 1024×1024 pixels. We illustrate the 3DGS
1012 rendered RGB image I_{RGB} , scale-invariant depth D_{est} estimated by MoGe (Wang et al., 2024a) and
1013 depth from 3DGS D_{GS} in Figure 11.

1014 **3DGS reconstruction from satellite imagery.** Our satellite-view optimization process runs for
1015 30,000 iterations, with densification enabled between iterations 1,000 and 21,000. We modify several
1016 key parameters in the standard 3DGS implementation to address satellite imagery’s unique challenges.
1017 First, to prevent undesirable Gaussian elongation artifacts common with overhead views, we reduce
1018 the scaling learning rate from 0.005 to 0.001. Second, we address sparsity issues of Gaussian points
1019 in close-up renderings by lowering the densification gradient threshold from 0.002 to 0.001, ensuring
1020 sufficient detail when viewed from ground level. Furthermore, we implement pruning of Gaussians
1021 with maximum covariance exceeding 20 to eliminate floating artifacts. The loss function weights are
1022 set to $\lambda_{\text{D-SSIM}} = 0.2$, $\lambda_{\text{op}} = 10$, and $\lambda_{\text{depth}} = 0.5$ for optimal reconstruction quality. For appearance
1023 modeling, we adopt the architecture from WildGaussians (Kulhanek et al., 2024), implementing
1024 an appearance MLP with 2 hidden layers (128 neurons each) and ReLU activation functions. The
1025 per-image and per-Gaussian embedding dimensions are set to 32 and 24 respectively, with learning
rates of 0.001, 0.005, and 0.0005 for per-image embeddings e_j , per-Gaussian embeddings g_i , and the

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043

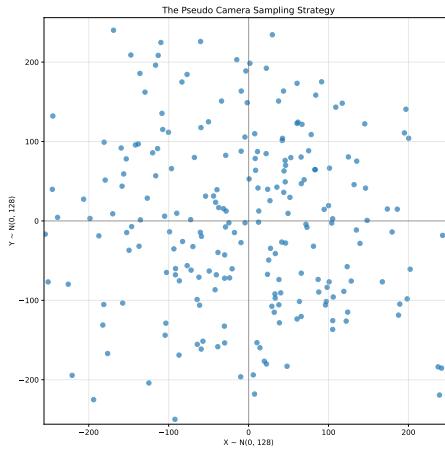
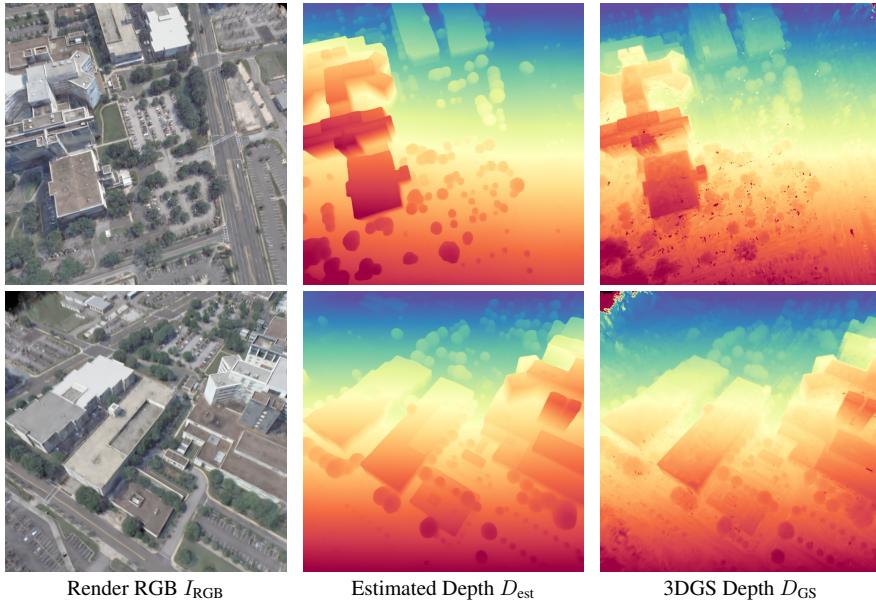


Figure 10: **The sampling strategy of pseudo camera.** In this example, we sample 240 points using the strategy.

1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065



1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

Figure 11: **Pseudo-cam Depth Supervision.** We use MoGe (Li et al., 2024c) to estimate the scale-invariant depth D_{est} from the rendered RGB image I_{RGB} . The rightmost figures show the rasterized depth D_{GS} from 3DGS.

appearance MLP f , respectively. The complete satellite-view training requires approximately 1 hour on a single NVIDIA RTX A6000 GPU.

FlowEdit-based refinement. We set FlowEdit noise parameters $n_{\min} = 4$ and $n_{\max} = 10$ to balance artifact removal with detail preservation. Our source prompt (“*Satellite image of an urban area with modern and older buildings, roads, green spaces. Some areas appear distorted, with blurring and warping artifacts.*”) characterizes initial renders, while the target prompt (“*Clear satellite image of an urban area with sharp buildings, smooth edges, natural lighting, and well-defined textures.*”) guides refinement. These parameters were determined through experimentation, with lower noise values preserving more original structure but removing fewer artifacts, and higher values creating

1080 more significant changes but potentially altering underlying geometry. All other FlowEdit parameters
 1081 use default values.
 1082

1083 **Curriculum-based refinement details.** Our IDU process comprises $N_e = 5$ episodes of 10,000
 1084 iterations each, with densification through iteration 9,000. At the start of IDU, we randomly select
 1085 and fix a single per-image appearance embedding e_j . Opacity regularization is disabled during IDU,
 1086 as our curriculum naturally mitigates floating artifacts through multi-view consistency, enabling
 1087 Gaussians to retain variable opacities beneficial for semi-transparent structures (Kerbl et al., 2023).
 1088 For DFC2019 (Le Saux et al., 2019) dataset, we utilize $N_p = 9$ look-at points in a 3×3 grid (512
 1089 units wide, centered at origin), with $N_v = 6$ cameras per point and $N_s = 2$ samples per view.
 1090 Camera elevations decrease from 85° to 45° and radii from 300 to 250 units across episodes. For
 1091 GoogleEarth (Xie et al., 2024) dataset, we utilize $N_p = 16$ look-at point at origin, with $N_v = 6$
 1092 cameras per point and $N_s = 2$ samples per view. Camera elevations decrease from 85° to 45° and
 1093 radius is fixed 600-unit across episodes. All training images are rendered at 2048×2048 resolution.
 1094 Our training strategy samples 75% from refined images and 25% from original satellite images, this
 1095 sampling strategy makes sure that the final 3DGS scene faithfully follows the semantic and layout
 1096 in the input satellite imagery. The complete synthesizing stage requires approximately 6 hours on a
 1097 single NVIDIA RTX A6000 GPU.

1098 **Detail of training time.** All time measurements were conducted on the JAX_214 AOI using a
 1099 single NVIDIA RTX A6000 (48GB) GPU. The total training time increases from approximately
 1100 1 hour 35 minutes for the baseline reconstruction (30K iterations) to 6 hours 45 minutes for the
 1101 full pipeline. The majority of this additional cost is attributed to the Curriculum-based Iterative
 1102 Dataset Update (IDU) process (5 episodes), which accounts for approximately 5 hours and 10 minutes
 1103 combined. Specifically, a single IDU episode requires roughly 1 hour, where the computational load
 1104 is split almost evenly between render refinement (~ 30 min) and 3DGS reconstruction update (~ 32
 1105 min), while the initial rendering step is negligible (~ 4 s). While this results in a total training time
 1106 increase of approximately $4.3 \times$, we view this as a justifiable offline investment to bypass physical
 1107 data collection limitations.

1108 **Detail of memory consumption.** We distinguish between peak memory and final memory. The
 1109 peak memory usage reaches 46 GB during the synthesis stage, driven by the overhead of loading the
 1110 diffusion model (FLUX.1) and temporary densification of Gaussians. However, the final training
 1111 memory footprint is significantly lower (28.04 GB) as our method actively prunes redundant and
 1112 low-opacity points. In terms of scene complexity, the refinement process densifies the scene by
 1113 approximately 27%, increasing the Gaussian count from ~ 1.65 million (reconstruction stage) to
 1114 ~ 2.1 million, specifically targeting the vertical facade geometry missing in the initial satellite
 1115 reconstruction.

1116 **Validity of RPC to perspective approximation.** We adopt the methodology proposed in Satel-
 1117 liteSfM (Zhang et al., 2019) to approximate the satellite linear pushbroom sensor as a perspective
 1118 camera. This approximation relies on the “weak perspective” assumption, which holds valid when
 1119 the satellite altitude (Z) is significantly larger than the depth variation within the scene (ΔZ), i.e.,
 1120 $Z \gg \Delta Z$. Given that satellites orbit at distances of hundreds of kilometers while terrestrial depth
 1121 variations are limited to a few hundred meters, the ratio $\Delta Z/Z$ remains negligible, allowing the
 1122 geometry to converge to a perspective model. The approximation is achieved by generating a dense
 1123 grid of 3D-2D correspondences using the rigorous RPC model and solving for a projection
 1124 matrix P via the Direct Linear Transformation (DLT) method, which is subsequently decomposed
 1125 ($P = K[R|t]$) to recover camera parameters. Quantitative evaluations demonstrate that this process
 1126 introduces negligible error: the average maximum forward projection error against the rigorous RPC
 1127 model is only **0.126 pixels**, and the difference in triangulated 3D points is typically less than **5 cm**.
 1128 Furthermore, this initialization allows Bundle Adjustment to achieve sub-pixel accuracy, with median
 1129 reprojection errors recorded at **0.864 pixels**, confirming the suitability of this approximation for
 1130 high-fidelity 3D reconstruction.

1130 A.2 MAIN PAPER EXPERIMENTS DETAIL & RESULTS

1131 **DFC2019 (Le Saux et al., 2019) dataset details.** The number of training images and geographical
 1132 coordinates for each AOI is provided in Table 5. We also include four additional AOIs from

1134 **Table 5: Number of training images and geographical coordinate per Area of Interest (AOI).**
 1135 These AOIs correspond to standard evaluation scenarios established by previous works, ensuring
 1136 consistent and fair comparisons with existing baselines (e.g., Sat-NeRF (Marí et al., 2022)).

AOI	JAX_004	JAX_068	JAX_214	JAX_260
# of training image	9	17	21	15
Geographical coordinate	81.70643°W, 30.35782°N	81.66375°W, 30.34880°N	81.66353°W, 30.31646°N	81.66350°W, 30.31184°N

1140 **Table 6: Number of training images and geographical coordinates for additional AOIs.** We
 1141 selected 4 additional AOIs with distinct characteristics: JAX_164 features a city hall building,
 1142 JAX_175 contains an American football stadium, while the remaining two AOIs present other notable
 1143 urban structures.

AOI	JAX_164	JAX_168	JAX_175	JAX_264
# of training image	20	21	21	21
Geographical coordinate	81.66362°W, 30.33032°N	81.65297°W, 30.33037°N	81.63696°W, 30.32583°N	81.65285°W, 30.31189°N

1144 Jacksonville to demonstrate our method’s robustness across varying scene characteristics. The
 1145 number of training images and geographical coordinates for these additional AOIs is provided in
 1151 Table 6. These additional AOIs feature distinct characteristics: one contains a city hall building
 1152 (JAX_164), another includes an American football stadium (JAX_175), while the remaining two
 1153 exhibit other notable urban features (JAX_168 and JAX_264).

1154 **GoogleEarth (Xie et al., 2024) dataset details.** The GoogleEarth dataset, introduced by City-
 1155 Dreamer (Xie et al., 2024), contains semantic maps, height fields and renders from Google Earth
 1156 Studio (Google, 2024) of New York City. This dataset is used to train the generative model in
 1157 CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b). We pick four AOIs which
 1158 contain diverse city elements, including complex architectures (004), squares (010), resident area
 1159 (219) and riverside (336). However, original GES renders provided in GoogleEarth dataset are
 1160 rendered from a lower elevation angle, which is not similar to satellite imagery. Therefore, for each
 1161 AOI, we render 60 images from GES using an orbit trajectory with 80° of elevation angle and 2219 of
 1162 radius. These new renders serve as the input of our methods. The AOI ID, geographical coordinates,
 1163 and the number of input images are detailed in Table 7.

1164 **User study details.** We asked participants three specific questions and instructed them to select one
 1165 video that best addressed each question:

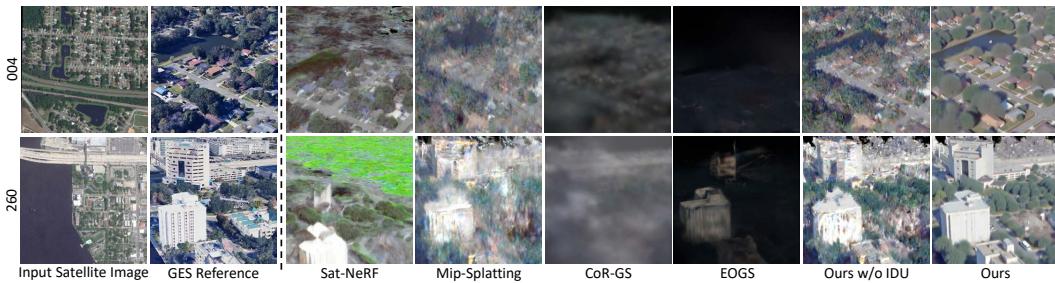
1. **Geometric Accuracy:** "Which video’s 3D structures (buildings, terrain, objects) more
 accurately represent the real-world geometry when compared to the ground truth video?"
2. **Spatial Alignment:** "Which video’s layout and positioning of elements better matches the
 satellite imagery reference?"
3. **Overall Perceptual Quality:** "Considering all aspects (geometry, textures, lighting, consis-
 tency), which video presents a more convincing and high-quality 3D representation of the
 scene?"

1176 For the user study on DFC2019 dataset, each participant viewed videos from Sat-NeRF (Marí et al.,
 1177 2022), our method without IDU, and our complete method, alongside Google Earth Studio reference
 1178 footage and the original satellite imagery. For the user study on the GoogleEarth dataset, each
 1179 participant viewed videos from CityDreamer (Xie et al., 2024), GaussianCity (Xie et al., 2025b) and
 1180 our complete method, alongside Google Earth Studio reference footage and the reference satellite
 1181 imagery.

1182 **Comparison details.** For quantitative comparisons with Sat-NeRF (Marí et al., 2022), Mip-
 1183 Splatting (Yu et al., 2024) and our method without IDU refinement, we used consistent camera
 1184 parameters across all methods: 17° elevation angle, 328-unit radius, and 20° field of view, with
 1185 cameras targeting the AOI’s origin. For comparisons with CityDreamer (Xie et al., 2024) and Gaus-
 1186 sianCity (Xie et al., 2025b), we use 45° elevation angle, 1067-unit radius, and 20° field of view,
 1187 with cameras also targeting the AOI’s origin. These parameters were selected to ensure equitable
 1188 comparison with similar scene coverage across methods.

1188 Table 7: **Number of training images and geographical coordinate per Area of Interest (AOI).** We
 1189 pick 4 AOIs from the GoogleEarth (Xie et al., 2024) dataset, ensuring fair comparisons with existing
 1190 baselines (e.g., CityDreamer (Xie et al., 2024) and GaussianCity (Xie et al., 2025b))

1191	AOI	4WorldFinancialCtr (004)	10UnionSquareE#5P (010)	219E12thSt (219)	336AlbanySt (336)
1192	# of training image	60	60	60	60
1193	Geographical coordinate	74.01587°W, 40.71473°N	73.98975°W, 40.73482°N	73.98690°W, 40.73187°N	74.01753°W, 40.71020°N



1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619
 1620
 1621
 1622
 1623
 1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639
 1640
 1641
 1642
 1643
 1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673
 1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727
 1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781
 1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835
 1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889
 1890
 1891
 1892
 1893
 1894
 1895
 1896
 1897
 1898
 1899
 1900
 1901
 1902
 1903
 1904
 1905
 1906
 1907
 1908
 1909
 1910
 1911
 1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943
 1944
 1945
 1946
 1947
 1948
 1949
 1950
 1951
 1952
 1953
 1954
 1955
 1956
 1957
 1958
 1959
 1960
 1961
 1962
 1963
 1964
 1965
 1966
 1967
 1968
 1969
 1970
 1971
 1972
 1973
 1974
 1975
 1976
 1977
 1978
 1979
 1980
 1981
 1982
 1983
 1984
 1985
 1986
 1987
 1988
 1989
 1990
 1991
 1992
 1993
 1994
 1995
 1996
 1997
 1998
 1999
 2000
 2001
 2002
 2003
 2004
 2005
 2006
 2007
 2008
 2009
 2010
 2011
 2012
 2013
 2014
 2015
 2016
 2017
 2018
 2019
 2020
 2021
 2022
 2023
 2024
 2025
 2026
 2027
 2028
 2029
 2030
 2031
 2032
 2033
 2034
 2035
 2036
 2037
 2038
 2039
 2040
 2041
 2042
 2043
 2044
 2045
 2046
 2047
 2048
 2049
 2050
 2051
 2052
 2053
 2054
 2055
 2056
 2057
 2058
 2059
 2060
 2061
 2062
 2063
 2064
 2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075
 2076
 2077
 2078
 2079
 2080
 2081
 2082
 2083
 2084
 2085
 2086
 2087
 2088
 2089
 2090
 2091
 2092
 2093
 2094
 2095
 2096
 2097
 2098
 2099
 2100
 2101
 2102
 2103
 2104
 2105
 2106
 2107
 2108
 2109
 2110
 2111
 2112
 2113
 2114
 2115
 2116
 2117
 2118
 2119
 2120
 2121
 2122
 2123
 2124
 2125
 2126
 2127
 2128
 2129
 2130
 2131
 2132
 2133
 2134
 2135
 2136
 2137
 2138
 2139
 2140
 2141
 2142
 2143
 2144
 2145
 2146
 2147
 2148
 2149
 2150
 2151
 2152
 2153
 2154
 2155
 2156
 2157
 2158
 2159
 2160
 2161
 2162
 2163
 2164
 2165
 2166
 2167
 2168
 2169
 2170
 2171
 2172
 2173
 2174
 2175
 2176
 2177
 2178
 2179
 2180
 2181
 2182
 2183
 2184
 2185
 2186
 2187
 2188
 2189
 2190
 2191
 2192
 2193
 2194
 2195
 2196
 2197
 2198
 2199
 2200
 2201
 2202
 2203
 2204
 2205
 2206
 2207
 2208
 2209
 2210
 2211
 2212
 2213
 2214
 2215
 2216
 2217
 2218
 2219
 2220
 2221
 2222
 2223
 2224
 2225
 2226
 2227
 2228
 2229
 2230
 2231
 2232
 2233
 2234
 2235
 2236
 2237
 2238
 2239
 2240
 2241
 2242
 2243
 2244
 2245
 2246
 2247
 2248
 2249
 2250
 2251
 2252
 2253
 2254
 2255
 2256
 2257
 2258
 2259
 2260
 2261
 2262
 2263
 2264
 2265
 2266
 2267
 2268
 2269
 2270
 2271
 2272
 2273
 2274
 2275
 2276
 2277
 2278
 2279
 2280
 2281
 2282
 2283
 2284
 2285
 2286
 2287
 2288
 2289
 2290
 2291
 2292
 2293
 2294
 2295
 2296
 2297
 2298
 2299
 2300
 2301
 2302
 2303
 2304
 2305
 2306
 2307
 2308
 2309
 2310
 2311
 2312
 2313
 2314
 2315
 2316
 2317
 2318
 2319
 2320
 2321
 2322
 2323
 2324
 2325
 2326
 2327
 2328
 2329
 2330
 2331
 2332
 2333
 2334
 2335
 2336
 2337
 2338
 2339
 2340
 2341
 2342
 2343
 2344
 2345
 2346
 2347
 2348
 2349
 2350
 2351
 2352
 2353
 2354
 2355
 2356
 2357
 2358
 2359
 2360
 2361
 2362
 2363
 2364
 2365
 2366
 2367
 2368
 2369
 2370
 2371
 2372
 2373
 2374
 2375
 2376
 2377
 2378
 2379
 2380
 2381
 2382
 2383
 2384
 2385
 2386
 2387
 2388
 2389
 2390
 2391
 2392
 2393
 2394
 2395
 2396
 2397
 2398
 2399
 2400
 2401
 2402
 2403
 2404
 2405
 2406
 2407
 2408
 2409
 2410
 2411
 2412
 2413
 2414
 2415
 2416
 2417
 2418
 2419
 2420
 2421
 2422
 2423
 2424
 2425
 2426
 2427
 2428
 2429
 2430
 2431
 2432
 2433
 2434
 2435
 2436
 2437
 2438
 2439
 2440
 2441
 2442
 2443
 2444
 2445
 2446
 2447
 2448
 2449
 2450
 2451
 2452
 2453
 2454
 2455
 2456
 2457
 2458
 2459
 2460
 2461
 2462
 2463
 2464
 2465
 2466
 2467
 2468
 2469
 2470
 2471
 2472
 2473
 2474
 2475
 2476
 2477
 2478
 2479
 2480
 2481
 2482
 2483
 2484
 2485
 2486
 2487
 2488
 2489
 2490
 2491
 2492
 2493
 2494
 2495
 2496
 2497
 2498
 2499
 2500
 2501
 2502
 2503
 2504
 2505
 2506
 2507
 2508
 2509
 2510
 2511
 2512
 2513
 2514
 2515
 25

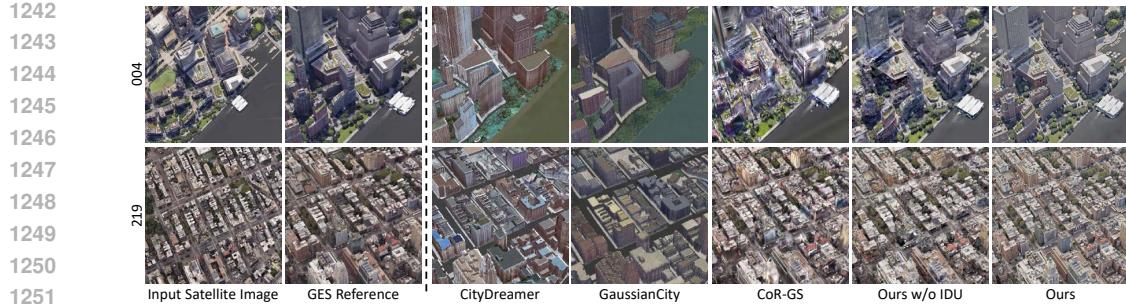


Figure 13: **Additional qualitative comparison on the GoogleEarth dataset with CityDreamer (Xie et al., 2024), GaussianCity (Xie et al., 2025b), and CoR-GS (Zhang et al., 2024b).** Our method is able to synthesize texture and geometry that is closer to the reference GES render.

Table 8: **Quantitative comparison on each AOI of DFC2019 (Le Saux et al., 2019).** Our method consistently outperforms baseline methods on distribution metrics and most pixel-level metrics, indicating superior image synthesis quality. Metrics are computed between renders from each method and reference frames from GES.

Scene	Methods	Distribution Metrics		Pixel-level Metrics*		
		$\text{FID}_{\text{CLIP}} \downarrow$	$\text{CMMMD} \downarrow$	$\text{PSNR} \uparrow$	$\text{SSIM} \uparrow$	$\text{LPIPS} \downarrow$
JAX_004	Sat-NeRF	79.97	3.838	11.95	0.2290	0.8700
	EOGS	107.23	5.913	8.22	0.1271	1.0174
	Mip-Splatting	85.33	4.986	13.06	0.2412	0.8157
	CoR-GS	91.01	5.131	11.25	0.2554	0.9793
	Ours	24.45	1.474	12.90	0.2446	0.846
JAX_068	Sat-NeRF	93.70	5.376	9.86	0.2607	0.8414
	EOGS	85.57	5.516	6.39	0.1593	0.9953
	Mip-Splatting	92.95	6.163	11.64	0.2900	0.8444
	CoR-GS	90.34	5.864	11.77	0.3230	1.0073
	Ours	28.35	2.845	11.79	0.2931	0.8210
JAX_214	Sat-NeRF	90.76	5.376	8.97	0.2684	0.8394
	EOGS	71.02	4.342	7.40	0.2293	0.8883
	Mip-Splatting	82.04	5.088	11.23	0.3844	0.8048
	CoR-GS	86.33	5.258	11.66	0.4074	0.9079
	Ours	26.69	1.964	12.24	0.3881	0.7420
JAX_260	Sat-NeRF	89.00	4.881	9.43	0.3172	0.9068
	EOGS	87.15	5.372	7.04	0.1574	0.9342
	Mip-Splatting	88.42	5.385	11.61	0.3579	0.8130
	CoR-GS	88.44	4.710	11.50	0.4162	0.8977
	Ours	29.83	2.076	12.59	0.3574	0.7540

significant architectures. As shown in Figure 20, we present synthesis results for **Neuschwanstein Castle** and **Wells Cathedral**. These scenes pose significant challenges due to their intricate non-Manhattan geometries, including sharp spires, varying elevations, and gothic architectural details. Despite these complexities, our method successfully disentangles the underlying geometry from the satellite input and hallucinates plausible high-frequency details for facades that are heavily occluded in the nadir views. This confirms that our hybrid reconstruction-generation approach is not limited to simple urban prisms but extends effectively to complex, free-form structures.

Synthesis of bridges. In addition to dense building clusters, we evaluate our method’s performance on scenes with complex topological structures, such as bridges. Figure 16 illustrates renders of bridges in JAX_068, JAX_214 and JAX_175, a typically difficult case for standard photogrammetry due to the thin structural components. Our method successfully recovers the connectivity of the bridge span while synthesizing realistic water textures. The diffusion-based refinement effectively

1296 Table 9: **Quantitative comparison with CityDreamer (Xie et al., 2024), GaussianCity (Xie et al.,**

1297 2025b), CoR-GS (Zhang et al., 2024b) on each AOI of the GoogleEarth dataset (Xie et al., 2024).

1298 The results show that our approach consistently achieves the best performance, indicating superior

1299 geometric and perceptual fidelity compared to all baselines. Metrics are computed between renders

1300 from each method and reference frames from GES.

1301	1302	1303	Scene	Methods	Distribution Metrics		Pixel-level Metrics		
					FID _{CLIP} ↓	CMMD ↓	PSNR ↑	SSIM ↑	LPIPS ↓
1304	004		CityDreamer	39.88	3.869	13.06	0.3519	0.5643	
1305			GaussianCity	28.71	2.710	14.00	0.3786	0.5656	
1306			CoR-GS	33.69	4.203	11.55	0.3440	0.6120	
1307			Ours	10.43	2.491	15.09	0.3793	0.3978	
1308	010		CityDreamer	34.29	4.270	12.24	0.1387	0.5544	
1309			GaussianCity	29.67	2.850	12.90	0.1661	0.5335	
1310			CoR-GS	29.75	3.672	12.90	0.1807	0.4209	
1311			Ours	11.03	1.631	13.58	0.1769	0.4073	
1312	219		CityDreamer	42.38	4.372	11.63	0.1344	0.5471	
1313			GaussianCity	32.83	2.883	12.37	0.1676	0.5254	
1314			CoR-GS	29.55	3.958	12.64	0.1792	0.3974	
1315			Ours	7.83	2.635	13.12	0.1699	0.3975	
1316	336		CityDreamer	29.53	4.097	13.39	0.4431	0.5654	
1317			GaussianCity	23.72	3.224	14.36	0.4533	0.5382	
1318			CoR-GS	16.29	3.173	14.29	0.4592	0.3879	
1319			Ours	10.36	1.279	15.32	0.4662	0.3719	

1320
1321
1322 regularizes the geometry, preventing the characteristic "melting" artifacts often observed in thin
1323 structures when using satellite-only reconstruction.

1324
1325 **Visualizing transient object handling via per-image embeddings.** A key challenge in multi-date
1326 satellite reconstruction is the handling of dynamic elements, such as moving vehicles and pedestrians,
1327 which can introduce ghosting artifacts. Our approach addresses this by learning per-image appearance
1328 embeddings e_j that capture photometric variations specific to each capture date. As visualized in
1329 Figure 14, rendering the same viewpoint across 20 distinct appearance embeddings reveals that
1330 transient objects exhibit significant variability, appearing clearly in some embeddings while fading
1331 or vanishing in others. This qualitative evidence suggests that our appearance modeling effectively
1332 acts as a "sink" for transient data that does not align with the static 3D geometry. By absorbing
1333 these inconsistencies into the appearance code rather than the geometric parameters, the optimization
1334 naturally disentangles transient elements from the underlying static structure, ensuring a clean and
1335 consistent geometric reconstruction.

1336
1337 **Episode-vs-coverage analysis of curriculum strategy.** To quantify the effectiveness of the IDU
1338 module in revealing occluded regions, we present an Episode-vs-Coverage analysis (Figure 18). Since
1339 ground truth 3D geometry is unavailable for these satellite scenes, we use the final converged 3DGS
1340 model as a proxy for the total scene surface. We compute the cumulative coverage by optimizing
1341 a visibility attribute for every Gaussian point against the camera poses utilized in each episode.
1342 As shown in the figure, the coverage ratio steadily increases from ~ 0.50 in Episode 1 to ~ 0.75 in
1343 Episode 5. This consistent gain confirms that our curriculum strategy, which progressively lowers
1344 camera elevation from 85° to 45° , successfully reveals and reconstructs vertical facade geometry
1345 that was initially occluded in the top-down satellite views. However, we acknowledge a limitation in
1346 this metric: because it calculates coverage based on reconstructed points, it cannot account for "true
1347 holes" (surface areas that were never generated at all because they were completely occluded from
1348 all sampled views). Future work could address this by dynamically sampling IDU cameras to target
1349 specific geometric uncertainties or detected holes.

1348
1349 **Stochastic appearance diversity.** To demonstrate the generative capacity of our hybrid framework,
1350 we evaluate the stochastic diversity of the synthesized textures in Figure 22. By varying the random
1351 seed during the diffusion refinement stage while maintaining the same geometric initialization, our

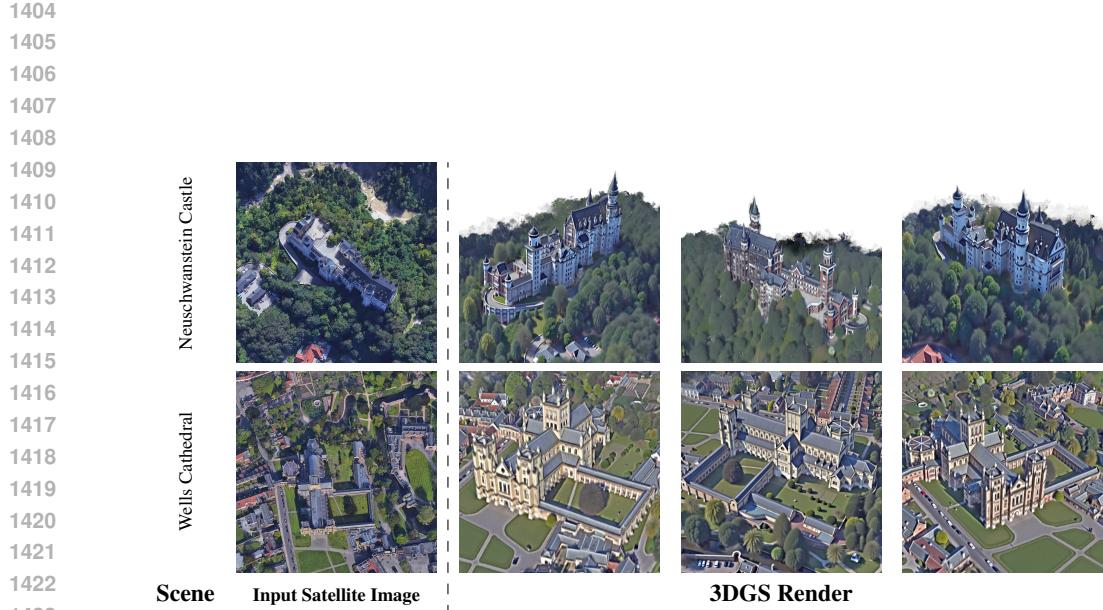


Figure 14: **Visualizing transient object handling via per-image embeddings.** We render the same viewpoint using 20 different learned appearance embeddings (Emb. 00–19). Observe that transient objects, such as the vehicles on the road, exhibit varying degrees of visibility across different embeddings (e.g., clearly visible in some, faded or absent in others), while the static building geometry remains consistent. This qualitatively demonstrates that our per-image appearance modeling effectively disentangles transient elements from the underlying static 3D structure, preventing dynamic artifacts from corrupting the geometric reconstruction.

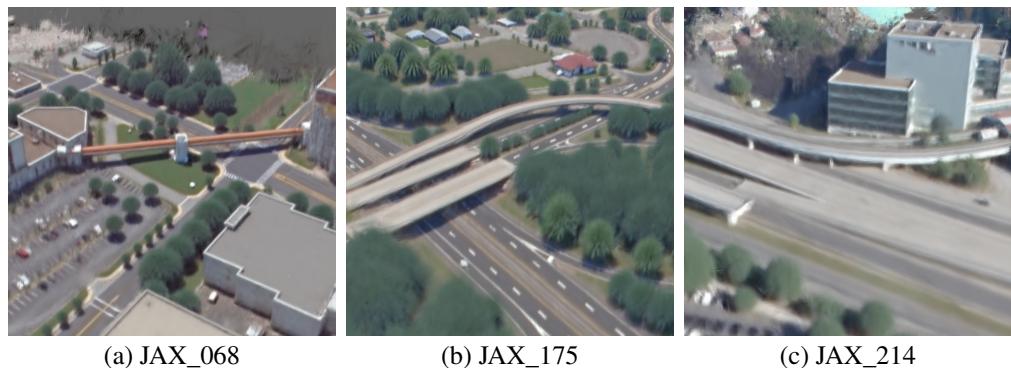
method produces diverse yet plausible surface details for identical underlying structures, 5643]. As illustrated in the figure, detailed features such as the text on the red building signage vary distinctively (e.g., “Outeil” vs. “CUTAN”). Crucially, the macroscopic building footprint remains geometrically fixed, confirming that our framework successfully disentangles the reconstruction of physical geometry (grounded in satellite constraints) from the generative synthesis of high frequency appearance.

A.4 LLM USAGE DISCLOSURE

Large language models (LLMs) were used to assist in improving the clarity and conciseness of the writing and in searching for related work. All technical ideas, algorithm designs, experiments, and analysis were conceived, implemented, and validated by the authors. The authors have carefully verified all content and take full responsibility for the correctness and integrity of this paper.



1424 **Figure 15: Qualitative results on complex geometries.** Visualization of satellite image inputs
1425 and corresponding rendered frames. We demonstrate the model’s capability on irregular historical
1426 architectures, including Neuschwanstein Castle and Wells Cathedral, showing the synthesis of
1427 complex geometry.

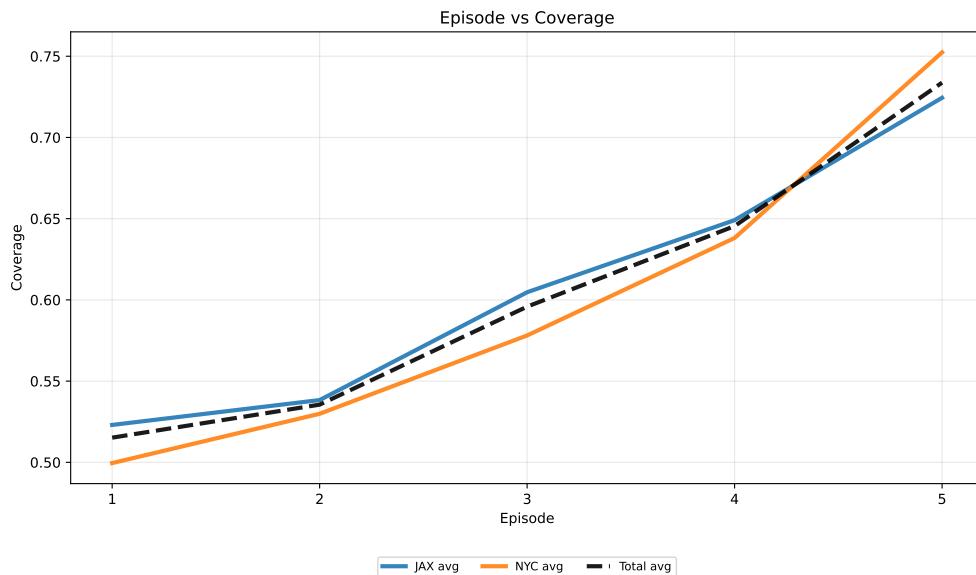


1450 **Figure 16: Qualitative results for bridges.** We present the render results for bridges appears in
1451 JAX_068, JAX_214 and JAX_175, demonstrating the method’s ability to handle complex topological
1452 structures and water surfaces that are typically challenging for standard reconstruction pipelines.
1453

Figure 17: **Refine renders with different prompt strategies.**

1512
 1513
 1514 **Table 10: List of text prompts used in sensitivity analysis.** We evaluate six different prompting
 1515 strategies to test the robustness of our method.
 1516

Strategy	Source Prompt (P_{src})	Target Prompt (P_{tar})
Baseline	Satellite image of an urban area with modern and older buildings, roads, green spaces. Some areas appear distorted, with blurring and warping artifacts.	Clear satellite image of an urban area with sharp buildings, smooth edges, natural lighting, and well-defined textures.
Vague Source	A blurry satellite image of an urban area.	Clear satellite image of an urban area with sharp buildings, smooth edges, natural lighting, and well-defined textures.
Vague Target	Satellite image of an urban area with modern and older buildings, roads, green spaces. Some areas appear distorted, with blurring and warping artifacts.	A clear satellite image of an urban area.
Focus Geometry	Satellite image of an urban area with modern and older buildings, roads, green spaces. Some areas appear distorted, with blurring and warping artifacts.	Clear satellite image of an urban area with geometrically precise buildings, flat rooftops, straight edges, and well-defined roads.
Focus Texture	Satellite image of an urban area with modern and older buildings, roads, green spaces. Some areas appear distorted, with blurring and warping artifacts.	Clear satellite image of an urban area with realistic, high-resolution textures, detailed facades, clear vegetation, and natural lighting.
Context Free	distorted, blurring, warping artifacts	clear, sharp, smooth edges, natural lighting, well-defined textures



1561 **Figure 18: Episode-vs-Coverage analysis.** The plot illustrates the cumulative surface coverage ratio
 1562 increasing across refinement episodes. The curriculum-based strategy effectively exposes occluded
 1563 regions, particularly vertical facades, as the camera elevation descends.
 1564
 1565



Figure 19: **Qualitative results across primary scenes.** Visualization of satellite image inputs and corresponding rendered frames for our four main AOIs.

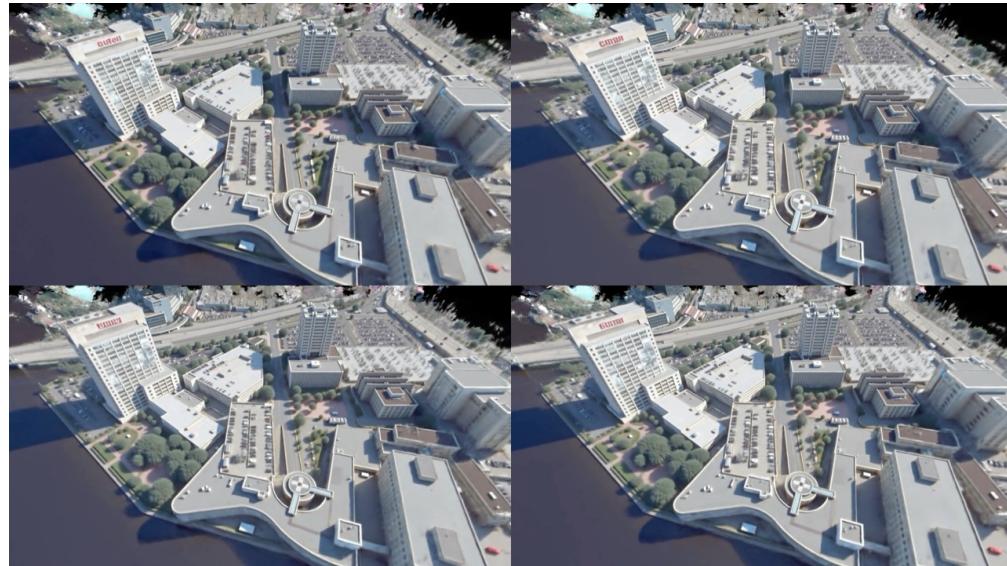


Figure 20: **Qualitative results across additional scenes.** Visualization of satellite image inputs and corresponding rendered frames for four additional AOIs with distinctive characteristics: JAX_164 features a city hall building, JAX_175 contains an American football stadium, while JAX_168 and JAX_264 present other notable urban structures.



Figure 21: **Visualization of multi-date satellite imagery of the DFC2019 dataset.** Note the substantial shifts in appearance, including changes in illumination, cloud cover, and surface characteristics, which introduce challenges for consistent 3D reconstruction.

1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742



1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781

Figure 22: Demonstration of stochastic appearance diversity while preserving geometric consistency. Our method generates diverse plausible textures for identical underlying geometry across different random seeds. Notice how the red signage text on the building facade varies distinctively (e.g., “Outeil” vs. “CUTAN”) while the building’s structural footprint remains fixed, confirming that our framework successfully disentangles geometric reconstruction (grounded in satellite data) from generative appearance synthesis (variable via diffusion).