

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 SAT3DGEN: COMPREHENSIVE STREET-LEVEL 3D SCENE GENERATION FROM SINGLE SATELLITE IMAGE

Anonymous authors

Paper under double-blind review

ABSTRACT

Generating a street-level 3D scene from a single satellite image is a crucial yet challenging task. Current methods present a stark trade-off: geometry-colorization models achieve high geometric fidelity but are typically building-focused and lack semantic diversity. In contrast, proxy-based models use feed-forward image-to-3D frameworks to generate holistic scenes by jointly learning geometry and texture, a process that yields rich content but coarse and unstable geometry. We attribute these geometric failures to the extreme viewpoint gap and sparse, inconsistent supervision inherent in satellite-to-street data. We introduce Sat3DGen to address these fundamental challenges, which embodies a geometry-first methodology. This methodology enhances the feed-forward paradigm by integrating novel geometric constraints with a perspective-view training strategy, explicitly countering the primary sources of geometric error. This geometry-centric strategy yields a dramatic leap in both 3D accuracy and photorealism. **For validation, we first constructed a new benchmark by pairing the VIGOR-OOD test set with high-resolution DSM data. On this benchmark, our method improves geometric RMSE from 6.76m to 5.20m.** Crucially, this geometric leap also boosts photorealism, reducing the Fréchet Inception Distance (FID) from ~ 40 to 19 against the leading method, Sat2Density++, despite using no extra tailored image-quality modules. We demonstrate the versatility of our high-quality 3D assets through diverse downstream applications, including semantic-map-to-3D synthesis, multi-camera video generation, large-scale meshing, and unsupervised single-image Digital Surface Model (DSM) estimation.

1 INTRODUCTION

Street-level 3D scenes are useful for mapping, robotics, simulation, and media creation (Workman et al., 2017; Toker et al., 2021; Xie et al., 2024; Zhou et al., 2020; Shi et al., 2022). Ground-level capture is costly and uneven across regions (Anguelov et al., 2010), whereas satellite imagery offers wide coverage, low cost, and frequent updates (Campbell & Wynne, 2011). These characteristics motivate generating street-level 3D from overhead satellite images for large-scale, long-term applications. Our goal is to generate a 3D scene that faithfully preserves the semantics and appearance of an input satellite image and that can be rendered for street-view images and videos.

Existing methods for generating 3D from a single satellite image fall into two categories: *3D geometry colorization* (Hua et al., 2025; Li et al., 2024) and *3D proxy for image rendering* (Qian et al., 2023; 2025). *3D geometry colorization* follows a two-stage pipeline to predict and then texture 3D building geometry. While producing clean building models, these methods fail to capture non-building elements (e.g., zebra crossings, trees), resulting in outputs weakly consistent with the input satellite image (Fig. 1 (a,b)).¹ Extending them beyond buildings would require fine-grained geometry labels for many classes, which are scarce. *3D proxy for image rendering* uses tailored *feed-forward image-to-3D frameworks* (Hong et al., 2024; Xiang et al., 2024; Yu et al., 2021; Zhang et al., 2025) to learn a coarse, differentiable 3D proxy via joint optimization under 2D supervision. These methods are semantically faithful but yield poor geometry: boundaries are degraded, roofs are unrealistic, and floating artifacts are common (Fig. 1 (c)).

¹As of the submission deadline, the official implementations of Sat2Scene and Sat2City had not been fully released. We therefore report the 3D results shown in their papers and project pages.

	Image Renderable	Satellite Faithful	More than Buildings	Supervision Signal	Synthetic or Real	Geometry Quality
Sat2City	✗	✗	✗	3D	Synthetic	✓
Sat2Scene	✓	✗	✗	Imgae	Real	✓
Sat2Density++	✓	✓	✓	Imgae	Real	✗
Ours	✓	✓	✓	Imgae	Real	✓

Comparison of single satellite input to 3D scene methods



High-quality Renderable3D generation from satellite image support multiple applications...

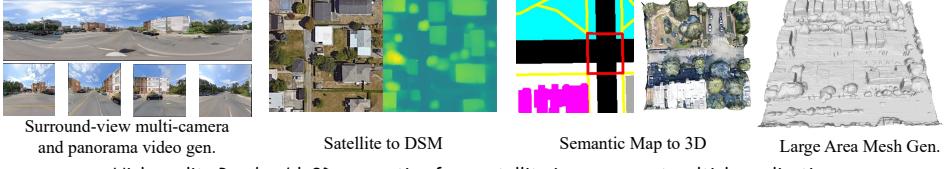


Figure 1: Comparison of 3D scene generation methods (top: attribute table; bottom: visual results). Given an input satellite image, (a) Sat2Scene and (b) Sat2City generate only shells of buildings and roads and miss non-building semantics; (c) Sat2Density++ and (d) Ours are faithful to satellite semantics and appearance, but Sat2Density++ is heavily distorted, whereas our Sat3DGen yields a more structured, higher-quality 3D representation.

Our goal requires preserving the rich semantics of the input satellite image, making the proxy-based paradigm more suitable than geometry colorization. Encouragingly, recent object-level *feed-forward image-to-3D* works (e.g., InstantMesh (Xu et al., 2024), LRM (Hong et al., 2024)) have demonstrated that high-quality 3D can be learned from 2D supervision alone. This suggests that the poor geometry of existing scene-level proxy methods is not a fundamental flaw of the paradigm. Instead, we hypothesize it stems from insufficient geometric constraints to handle the unique challenges of outdoor scenes. Specifically, the supervision from only one satellite patch and a few ground-level panoramas is extremely sparse. This sparsity, coupled with the extreme viewpoint gap, leaves rooftop geometry underconstrained and induces artifacts like holes and floaters on vertical facades. Additionally, a footprint mismatch between the satellite and street views often destabilizes the geometry at scene boundaries.

To solve these specific problems, we propose Sat3DGen, which embodies a holistic, geometry-first methodology. Our strategy is not to invent a new feedforward image-to-3D architecture from scratch, but to elevate a general framework by demonstrating how to effectively solve its core geometric failures. To enforce plausible vertical structures and suppress floating artifacts, we introduce a *Gravity-based Density Variation Loss*. To address boundary errors stemming from the footprint mismatch, a *Spatial Token* regularizes peripheral layouts. To resolve rooftop ambiguity, a *Monocular Relative-Depth Prior* constrains satellite-view depth. Furthermore, to mitigate the issue of sparse supervision, we employ *Perspective View Training*, jointly training on panoramas and their projected views to increase effective viewpoint coverage and photometric consistency.

In evaluation, this emphasis on geometry translates directly to substantial quantitative and perceptual gains. We first validate our geometric improvements against the leading method, Sat2Density++, on a new benchmark we constructed by pairing the VIGOR-OOD test set with 1-meter resolution DSM data. Sat3DGen achieves a geometric RMSE of 5.20m, a significant reduction from Sat2Density++’s 6.76m. Crucially, this leap in 3D accuracy directly fuels a dramatic improvement in photorealism. Even though it includes no components tailored to image quality, our framework reduces the Fréchet Inception Distance (FID) on the VIGOR-OOD unseen-city split from Sat2Density++’s ~ 40 to 19. The resulting assets support diverse downstream applications such as semantic-map-to-3D synthesis, surround-view video from satellite imagery, large-area mesh generation, and single-image Digital Surface Model (DSM) generation without ground-truth depth supervision.

108

2 RELATED WORKS

110 **Feed-Forward Image to 3D Works** has gained popularity for producing high-quality 3D assets.
 111 Recently, large reconstruction models (Hong et al., 2024; Tang et al., 2024; Xu et al., 2024; Xiang
 112 et al., 2024) have focused on generating object-level 3D assets, leveraging larger datasets, more
 113 refined annotations, and more substantial models to improve the quality of the generated assets,
 114 achieving impressive results. However, existing works primarily focus on object-level generation,
 115 presenting additional challenges when applying these models in outdoor scenes. In our work, we
 116 focus on generating high-quality, comprehensive street-level 3D from a single input satellite image,
 117 thereby naturally enhancing the quality of generated videos and supporting various applications.

118 **Single Satellite to Street-view Synthesis.** Early studies generate individual street-view images from
 119 a single satellite patch (Regmi & Borji, 2018; 2019; Toker et al., 2021; Shi et al., 2022; Lu et al.,
 120 2020; Tang et al., 2019), but they do not produce usable 3D or multi-view consistency. Later works
 121 synthesize street-view videos by learning a colored 3D asset from the satellite input (Li et al., 2021;
 122 2024; Qian et al., 2025). Geometry-colorization methods (Li et al., 2021; 2024) often rely on height
 123 maps and vertical-facade assumptions, yielding building-centric scenes and missing non-building
 124 semantics such as roads, crosswalks, and trees. Our work builds on the proxy-based line and focuses
 125 on improving 3D quality under the same single-satellite input setting.

127

3 METHOD

129 As shown in Fig. 2, given a single overhead satellite image I_{sat} and an optional global illumination
 130 feature input f_{ill} used solely to control illumination when rendering street views, our model can
 131 synthesize a renderable 3D scene that (i) preserves the semantics and appearance of I_{sat} , (ii) supports
 132 high-fidelity satellite, perspective street-view, and panoramic rendering under controllable lighting,
 133 and (iii) can be exported as a mesh with Marching Cubes.

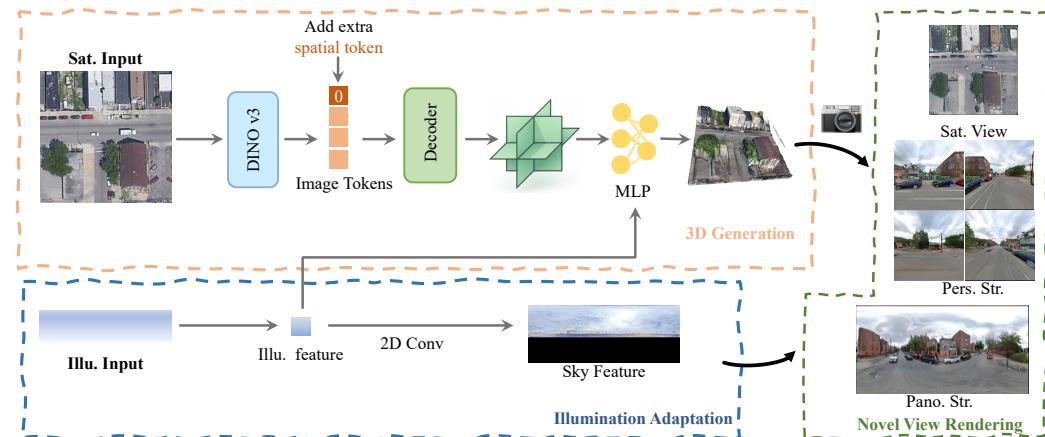
150

Figure 2: Diagram of the proposed Sat3DGen framework.

152 We adopt a feed-forward image-to-3D framework instantiated with a tri-plane NeRF (Chan et al.,
 153 2022) as a baseline. A frozen DINO-v3 encoder (Siméoni et al., 2025) maps I_{sat} to a compact token
 154 grid, which is optionally padded with learnable spatial capacity at the periphery and then decoded
 155 into a high-resolution tri-plane feature field. A lightweight MLP predicts density and color features
 156 from tri-plane features for volumetric rendering. Besides, we follow the illumination-adaptive design
 157 in Qian et al. (2025) to mitigate the sky/illumination mismatch issue.

158 Beyond this backbone, we introduce three novel geometry-oriented components that substantially
 159 enhance performance and depart from prior work (Qian et al., 2025): a gravity-based density variation
 160 loss to favor gravity-aligned structures, a monocular relative-depth prior in satellite view to resolve
 161 rooftop ambiguity, and panoramic-to-perspective supervision to densify viewpoints. The remainder
 of this section details the backbone; the losses and supervision strategy are presented in Section 3.4.

162 3.1 SATELLITE TO 3D GENERATION
163

164 Given a satellite image, our pipeline constructs a radiance field by encoding it into a 2D token
165 grid with a frozen backbone, padding with spatial tokens to expand the effective scene extent, and
166 decoding the tokens into tri-plane features.

167 **Satellite Encoder and Tokenization.** Following existing object-level feedforward image to 3D
168 works, we use frozen pretrained ViT model as image encoder (Xu et al., 2024; Xiang et al., 2024). In
169 practise, a frozen DINO-v3 ViT encoder (Siméoni et al., 2025) \mathcal{E}_{sat} processes I_{sat} into a 2D token
170 grid:

$$171 \mathbf{F}_{\text{token}} = \mathcal{E}_{\text{sat}}(I_{\text{sat}}) \in \mathbb{R}^{H_t \times W_t \times C}, \quad (1)$$

172 with $H_t = W_t = 16$ and $C = 1024$ in all experiments. This token grid is the minimal scene-level
173 latent that will be lifted into a 3D feature field.

174 **Spatial Tokens.** Street-view supervision often observes buildings and roads extending beyond the
175 satellite crop, which induces boundary artifacts if the 3D field is constrained to the crop footprint. We
176 therefore pad $\mathbf{F}_{\text{token}}$ with a border of N zero-valued spatial tokens on each side:

$$177 \mathbf{F}_{\text{token,pad}} = \text{PAD}_N(\mathbf{F}_{\text{token}}) \in \mathbb{R}^{(H_t+2N) \times (W_t+2N) \times C}, \quad N = 2. \quad (2)$$

179 With $H_t = W_t = 16$, padding yields $\mathbf{F}_{\text{token,pad}} \in \mathbb{R}^{20 \times 20 \times 1024}$. Suppose the original scene cube
180 spans L meters per side (e.g., $L = 50$ m). In that case, padding enlarges the effective cube to
181 $L \cdot \left(1 + \frac{2N}{H_t}\right)$ (e.g., 62.5 m), providing degrees of freedom to accommodate peripheral content while
182 stabilizing interior geometry.

184 **Tokens → Tri-Plane Features.** A lightweight VAE-style decoder (Esser et al., 2021) \mathcal{D} upsamples
185 tokens into a high-resolution tri-plane feature map with an upsampling factor $s = 16$:

$$186 \mathbf{F}_{\text{tri}} = \mathcal{D}(\mathbf{F}_{\text{token,pad}}) \in \mathbb{R}^{\text{res}_{\text{tri}} \times \text{res}_{\text{tri}} \times 96}, \quad (3)$$

187 where $\text{res}_{\text{tri}} = 320$ when padding is used and 256 otherwise. Channels are reshaped into three
188 orthogonal planes (XY, XZ, YZ).

189 **Tri-Plane Sampling.** A 3D query point $\mathbf{x} \in \mathbb{R}^3$ within the normalized scene cube is orthographically
190 projected onto each plane and bilinearly sampled to obtain features $\phi_{XY}(\mathbf{x}), \phi_{XZ}(\mathbf{x}), \phi_{YZ}(\mathbf{x})$. The
191 three plane features are aggregated by elementwise summation to form the fused feature:

$$192 \mathbf{h}(\mathbf{x}) = \phi_{XY}(\mathbf{x}) + \phi_{XZ}(\mathbf{x}) + \phi_{YZ}(\mathbf{x}). \quad (4)$$

193 Then, a shallow MLP predicts density and color:

$$194 \sigma(\mathbf{x}), \mathbf{c}(\mathbf{x}, \mathbf{w}) = \text{MLP}(\mathbf{h}(\mathbf{x}), \mathbf{w}), \quad (5)$$

196 where $\sigma(\mathbf{x})$ denotes the volume density, $\mathbf{c}(\mathbf{x}, \mathbf{w})$ is the radiance color conditioned on an illumination
197 code \mathbf{w} , and $\mathbf{h}(\mathbf{x})$ is the fused tri-plane feature at location \mathbf{x} ; the MLP uses a shared trunk with two
198 output heads for density and color.

199 3.2 ILLUMINATION-ADAPTIVE RENDERING AND SKY GENERATION
200

201 **Global Illumination Code.** Following Sat2Density++, we extract a global illumination feature f_{ill}
202 from a real street-view image I_{ill} in a statistical way (Qian et al., 2025), and then project to a style
203 code \mathbf{w}_{ill} with a light mlp:

$$204 \mathbf{w}_{\text{ill}} = \mathcal{E}_{\text{ill}}(I_{\text{ill}}). \quad (6)$$

205 During training, f_{ill} is extracted from the groundtruth street-view panorama image to mitigate
206 sky/illumination mismatch, and at test time, it enables lighting-controllable rendering.

208 **Sky Region Generation with Spherical Feature Maps.** To natively support perspective view
209 rendering, the sky module must provide consistent appearances for arbitrary viewpoints. We achieve
210 this by modeling the sky as a feature map on the sphere. A lightweight 2D convolutional decoder
211 produces this sky feature map from w_{ill} :

$$212 \mathbf{F}_{\text{sky}} = \mathcal{G}_{\text{sky}}(w_{\text{ill}}) \in \mathbb{R}^{512 \times 512 \times c}, \quad (7)$$

213 where c matches the renderer’s feature channels. For any given ray with normalized direction $\mathbf{d} \in \mathbb{S}^2$,
214 we convert its Cartesian coordinates to spherical angles (θ, ϕ) and bilinearly sample \mathbf{F}_{sky} to obtain the
215 sky color feature $\mathbf{c}_{\text{sky}}(\mathbf{d})$. This design elegantly provides consistent sky features for both panoramic
and perspective-view rendering.

216 3.3 VOLUMETRIC RENDERING AND OUTPUTS
217218 **Ray Marching and Compositing.** For a camera ray $r(t) = \mathbf{o} + t\mathbf{d}$, $t \in [t_n, t_f]$, we sample points
219 $\{\mathbf{x}_k\}$ with step δ_k and compute transmittance $T_k = \exp\left(-\sum_{j < k} \sigma(\mathbf{x}_j) \delta_j\right)$. The rendered color is
220

221
$$\mathbf{C}(r) = \sum_k T_k \left(1 - e^{-\sigma(\mathbf{x}_k) \delta_k}\right) \mathbf{c}(\mathbf{x}_k, w_{\text{ill}}) + T_{\text{out}} \mathbf{c}_{\text{sky}}(\mathbf{d}), \quad (8)$$

222

223 where T_{out} is the remaining transmittance upon exiting the volume. The same renderer supports
224 perspective and spherical cameras; the latter yields full panoramas.
225226 **Renderable Views and Mesh Export.** Our model can render (i) satellite views, (ii) perspective
227 street-view images at arbitrary camera poses, and (iii) panoramic street views. For asset export, we
228 evaluate σ on a dense grid and run marching cubes with a fixed isovalue τ to obtain a watertight
229 mesh. The sky branch is excluded from meshing.
230231 3.4 LOSS FUNCTIONS.
232233 **Gravity-based Density Variation Loss.** Outdoor scenes reconstructed from sparse views often
234 exhibit geometric artifacts like floating debris and hollow grounds. To mitigate these issues, we
235 introduce a regularizer based on a simple design principle: volumetric density should generally be
236 non-increasing with altitude. **The design of this regularizer is motivated by the physical effect of
237 gravity. To translate this concept into the NeRF framework, we leverage the volume density σ .** In
238 NeRF, σ measures light obstruction, making it a natural proxy for physical matter, given that outdoor
239 scenes are predominantly composed of opaque surfaces like terrain, rocks, and tree trunks. Following
240 the intuition that gravity causes matter to accumulate at lower elevations, we establish our principle:
241 σ should generally be non-increasing with altitude. This is consistent with real-world observations;
242 for instance, solid ground and tree trunks are typically found at lower altitudes, while higher altitudes
243 often contain sparser structures like leafy canopies or simply open air. Grounding our regularizer in
244 this physical intuition helps the model learn more plausible geometry.245 Specifically, we sample a 3D point $\mathbf{x} \in \mathbb{R}^3$ and a corresponding point $\mathbf{x}' = \mathbf{x} + \delta\mathbf{z}$ at a slightly higher
246 altitude, where $\delta\mathbf{z}$ is a small displacement vector purely in the upward (anti-gravity) direction. We
247 then penalize cases where the density at the higher point \mathbf{x}' is significantly greater than the density at
248 the lower point \mathbf{x} . This is enforced by minimizing the following loss:

249
$$\mathcal{L}_{\text{grav}} = \mathbb{E}_{\mathbf{x}, \delta\mathbf{z}} [\text{ReLU}(\sigma(\mathbf{x} + \delta\mathbf{z}) - \sigma(\mathbf{x}) - \epsilon)], \quad (9)$$

250 where the slack variable ϵ (set to 1 in our experiments) provides a soft constraint, allowing for
251 legitimate hollow or overhanging structures such as tree canopies, arched roofs, and bridges. This loss
252 effectively suppresses floating artifacts and fills baseless cavities while preserving realistic sparsity
253 under overhangs.254 **Satellite-View Depth Regularization.** Each scene provides one bird’s-eye satellite image and
255 only a few street-view observations; rooftops lack multi-view photometric supervision and tend to
256 be irregular. We therefore impose a relative depth prior in the satellite view using pseudo labels
257 from Depth Anything v2 (Yang et al., 2024). Let D^* be the pseudo relative depth for the satellite
258 camera and \hat{D} the rendered depth from our field. We adopt a scale-and-shift invariant MiDaS-style
259 loss (Ranftl et al., 2022):
260

261
$$\mathcal{L}_{\text{depth}} = \frac{1}{N} \sum_p |s \hat{D}(p) + t - D^*(p)| + \lambda \nabla \frac{1}{N} \sum_p \|\nabla(s \hat{D}(p) + t) - \nabla D^*(p)\|_1, \quad (10)$$

262

263 where (s, t) are optimal scale and shift estimated per image by least squares, N is the number of
264 valid pixels, and ∇ denotes spatial gradients. This encourages consistent depth ordering and smooth
265 rooftops without requiring metric depth.
266267 **Photometric Reconstruction and Adversarial Loss.** We supervise three rendered view types:
268 satellite views, panoramic street views, and perspective crops projected from panoramas. Let \hat{I}_i
269 be a rendered image and I_i^{gt} the corresponding ground truth. The photometric objective combines
per-pixel reconstruction with perceptual similarity, and we add an adversarial term to mitigate blur

270 from pure regression in complex outdoor scenes:
 271

$$272 \quad \mathcal{L}_{\text{RGB}} = \sum_i \left\| \hat{I}_i - I_i^{\text{gt}} \right\|_2^2 + \lambda_{\text{lpips}} \sum_i \mathcal{L}_{\text{LPIPS}}(\hat{I}_i, I_i^{\text{gt}}) + \lambda_{\text{GAN}} \sum_i \mathcal{L}_{\text{GAN}}(\hat{I}_i), \quad (11)$$

273
 274 where $\mathcal{L}_{\text{LPIPS}}$ is the perceptual loss and \mathcal{L}_{GAN} follows the StyleGAN2 hinge objective (Karras et al.,
 275 2020) for realism. In practice, the index i ranges over satellite, panorama, and perspective supervision
 276 views rendered during training.
 277

278 **Sky Losses: Opacity BCE and Masked Sky L1.** To disentangle the sky from the 3D scene and
 279 improve sky quality, we use two complementary losses on panoramic street views.
 280

281 Let $M_{\text{sky}} \in \{0, 1\}^{H \times W}$ be the pseudo binary sky mask of a panorama (1 for sky), which is generated
 282 by the off-the-shelf model (Zhang et al., 2022), and let $T_{\text{out}} \in [0, 1]^{H \times W}$ be the residual transmittance
 283 per pixel from volumetric rendering (interpreted as the fraction attributed to the sky background after
 284 alpha compositing). We apply a binary cross-entropy:
 285

$$\mathcal{L}_{\text{sky-op}} = \mathcal{L}_{\text{BCE}}(T_{\text{out}}, M_{\text{sky}}). \quad (12)$$

286 Denote the rendered panorama \hat{I}_{pano} and the ground-truth panorama $I_{\text{pano}}^{\text{gt}}$. We enforce color fidelity
 287 on sky pixels only:
 288

$$289 \quad \mathcal{L}_{\text{sky-L1}} = \frac{1}{\sum M_{\text{sky}}} \sum_p M_{\text{sky}}(p) \left\| \hat{I}_{\text{pano}}(p) - I_{\text{pano}}^{\text{gt}}(p) \right\|_1. \quad (13)$$

292 **Overall Objective.** The full training objective is a weighted sum of the above terms:
 293

$$294 \quad \mathcal{L}_{\text{total}} = \lambda_{\text{rgb}} \mathcal{L}_{\text{RGB}} + \lambda_{\text{grav}} \mathcal{L}_{\text{grav}} + \lambda_{\text{sky-op}} \mathcal{L}_{\text{sky-op}} + \lambda_{\text{sky-L1}} \mathcal{L}_{\text{sky-L1}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}}, \quad (14)$$

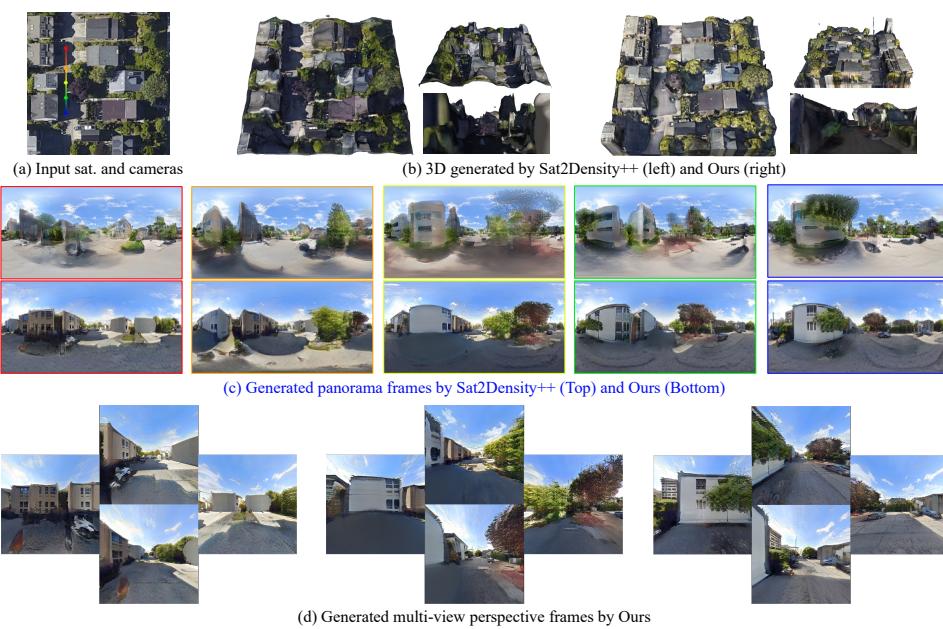
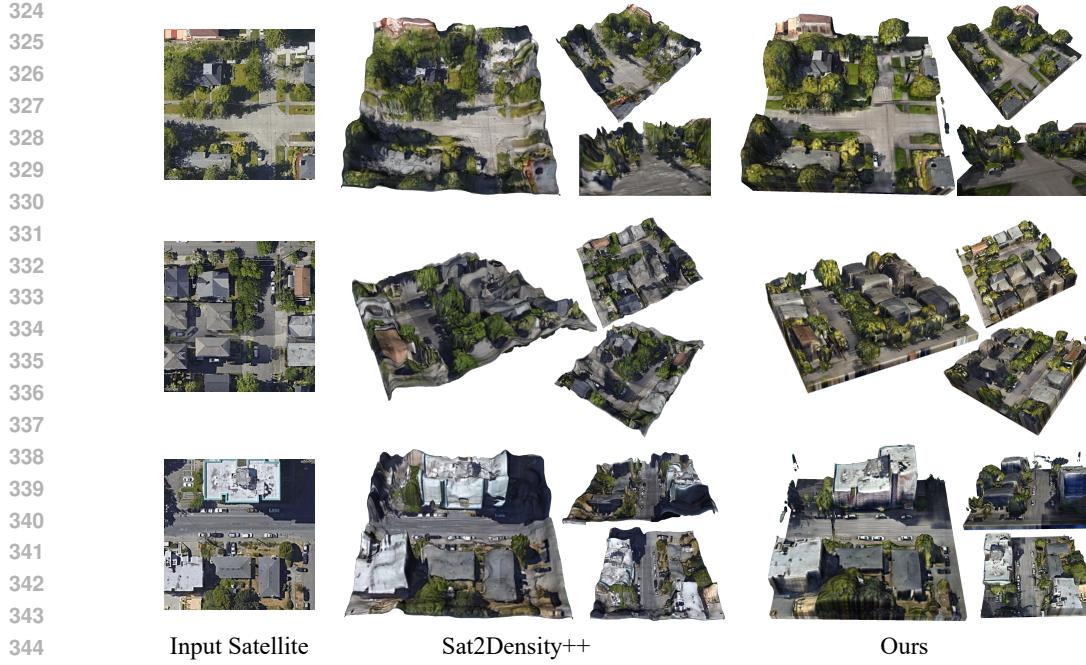
295 where weights λ . are hyperparameters.
 296

297 4 EXPERIMENTS

298 **Datasets and Splits.** We train on GPS-matched satellite–ground image pairs. Training uses three
 299 cities (Chicago, New York, San Francisco) in the VIGOR dataset (Zhu et al., 2021), and out-of-domain
 300 (OOD) testing uses the held-out city Seattle (VIGOR-OOD). VIGOR provides multiple street-view
 301 panoramas per satellite tile together with relative camera poses; the satellite zoom level is fixed at 20,
 302 yielding near-constant ground sampling distance per pixel. In total, we use 78,188 pairs for training
 303 and 11,875 pairs for quantitative evaluation on VIGOR. More details, data preparation, and statistics
 304 are provided in Appendix F.
 305

306 **Implementation Details.** We resize satellite images to 256×256 as input, and the generated triplane
 307 features have dimensions of $320 \times 320 \times 32 \times 3$. For fair comparison, the generated panorama
 308 images are shaped 512×128 , and the perspective images are 256×256 . The training process is
 309 conducted on 8 NVIDIA H20 GPUs with a batch size of 32, comprising 600,000 iterations for the
 310 training phase. More implementation details can be seen in the supplementary materials.
 311

312 **3D Comparision.** We compare our 3D results with Sat2Scene (Li et al., 2024), Sat2City (Hua et al.,
 313 2025) and Sat2Density++ (Qian et al., 2025). The colored meshes are generated by the Marching
 314 Cubes algorithm for Sat2Density++ and ours. Since there are no ground truth 3D assets available
 315 to evaluate the reconstruction quality, we can only perform qualitative comparisons, as shown on
 316 Fig. 1, Fig. 3, Fig. 4 (b), and Fig. 6. We observe consistent improvements in geometric plausibility
 317 and semantic faithfulness across diverse urban layouts. Compared with Sat2Scene and Sat2City,
 318 which mainly texture simplified building blocks and leave non-building regions weakly modeled, our
 319 reconstructions better preserve road markings, crosswalks, medians, tree belts, and sidewalks that are
 320 visible in the satellite input (Fig. 1). Relative to Sat2Density++, although both adopt a *feed-forward*
 321 *image-to-3D framework*, our method jointly integrates several lightweight components to improve
 322 geometry learning at street level under sparse, cross-view supervision. Taken together, these design
 323 choices strengthen scene layout near the satellite patch boundary, bias the volumetric field toward
 324 gravity-aligned structures, and inject rooftop depth cues from the overhead view, while increasing



371 Figure 4: Visual results of generated mesh (b), panorama videos (c), and multi-view perspective
372 video (d) from a single satellite image input and camera trajectories (rainbow line) (a). **The full video**
373 **can be seen in the supplemental materials.**

376 effective viewpoint coverage via panorama-to-perspective supervision. The resulting reconstructions
377 exhibit more coherent ground planes and periphery geometry, with fewer torn edges and warped
378 borders across the tile extent. Rooftops and building bases become geometrically plausible: roofs

378
 379 Table 1: Quantitative results of street-view comparison on the test set of VIGOR-OOD. **Bold** indicates
 380 the best results, while underlined text represents the second-best results.

Method	Realism FID↓	Evaluation KID↓	Semantic DINO↑	Structure SSIM↑	Pixel PSNR↑	Perceptual $P_{\text{alex}} \downarrow$	Similarity $P_{\text{squeeze}} \downarrow$
ControlNet	<u>23.6</u>	/	/	0.34	12.02	0.46	0.34
ControlS2S	28.0	/	/	0.42	<u>13.80</u>	<u>0.38</u>	0.27
Sat2Density	85.6	0.079	0.451	0.32	12.48	0.45	0.37
Sat2Density++	40.8	0.035	0.465	0.34	12.51	0.44	0.34
Canonical Image-to-3D	35.6	<u>0.030</u>	<u>0.479</u>	0.35	12.63	0.42	0.32
Ours	19.2	0.014	0.525	<u>0.37</u>	<u>12.83</u>	0.38	<u>0.30</u>

389 avoid bubbling or sagging, flat roofs remain planar, pitched roofs retain credible tilt, and facades
 390 connect cleanly to the ground (Fig. 1, Fig. 3, and Fig. 6).

391
 392 **Image and Video Comparison.** We provide quantitative and qualitative comparisons. The qualitative
 393 comparison can be seen on Fig. 4, and more video comparisons are provided in the supplementary
 394 ZIP archive. The quantitative comparison is shown on Table 1.

395 Quantitative comparison. We follow prior work (Qian et al., 2025; Ze et al., 2025) for evaluation.
 396 We report Fréchet Inception Distance (FID) (Heusel et al., 2017) and Kernel Inception Distance
 397 (KID) (Bińkowski et al., 2018) to quantify distributional similarity between generated and real image
 398 sets, reflecting realism and coverage. Semantic alignment is assessed using a DINO-based feature
 399 similarity following (Qian et al., 2025). Pixel-level fidelity and structural similarity are evaluated
 400 with PSNR and SSIM, and perceptual similarity is measured with LPIPS (Zhang et al., 2018) using
 401 AlexNet and SqueezeNet backbones, denoted P_{alex} (Krizhevsky et al., 2012) and P_{squeeze} (Iandola
 402 et al., 2016). Given that our task is an input-view conditioned novel view generation problem
 403 with a very large viewpoint gap, pixel-level correspondence to any single real image is inherently
 404 brittle due to occlusions, parallax, and minor pose or scene changes. The primary desiderata are
 405 photorealism and semantic faithfulness rather than exact pixel matching. We therefore treat FID,
 406 KID, and DINO-based semantic similarity as primary metrics, and report PSNR, SSIM, and LPIPS
 407 for completeness.

408 We compare our model with Sat2Density (Qian et al., 2023), Sat2Density++ (Qian et al., 2025), the
 409 diffusion-based image generation model ControlNet (Zhang et al., 2023b), and ControlS2S (Ze et al.,
 410 2025). We also provide results from a *feed-forward image-to-3D* model (denoted as “Canonical
 411 Image-to-3D” in Table 1), which removes the proposed spatial token module, gravity-based density
 412 variation loss, perspective training strategy, $\mathcal{L}_{\text{grav}}$, and $\mathcal{L}_{\text{depth}}$ proposed in our method. For Sat2Density,
 413 Sat2Density++, and our model, we pair each satellite image with a randomly selected global illu-
 414 mination feature input from the training set for a fair comparison. For ControlNet and ControlS2S,
 415 we use the results reported in the ControlS2S paper. It is worth noting that in the ControlS2S work,
 416 they divided the training and testing sets within each city. In contrast, our approach trains on three
 417 cities and conducts out-of-domain testing on an additional city, Seattle. This out-of-domain testing is
 418 significantly more challenging.

419 As shown in Table 1, our method leads in FID, KID, and DINO. Compared with the Canonical
 420 Image-to-3D, we only add our 3D optimization modules, yet the rendered images improve a lot.
 421 When compared with diffusion image generation models, our model also achieves lower FID and
 422 KID. This is because we learn a high-quality, view-consistent 3D representation that handles the large
 423 aerial-to-ground viewpoint change and produces more realistic and semantically correct images.

424 Qualitative Comparison. We provide a qualitative comparison in Fig. 4 (c), and more video compar-
 425 isons can be seen in the zip supplemental materials. We can render videos from a given satellite image
 426 and any street-view camera trajectory, as illustrated in Fig. 4 (c). We compared the panorama video
 427 results generated by Sat2Density++ and our model. Consistent with the conclusions drawn from
 428 the generated 3D assets, our model enhances the generated videos primarily by reducing artifacts
 429 and producing smoother edges around buildings and scene boundaries through the generation of
 430 higher-quality 3D representations.

431 **Geometric Comparison** To quantitatively evaluate the geometric accuracy of the generated 3D
 432 models, we compare the predicted satellite-view depth against the ground truth. The detailed

432
 433 Table 2: Ablation results on the VIGOR-OOD test set. The first row removes all proposed key
 434 components: \mathcal{L}_{dep} , $\mathcal{L}_{\text{grav}}$, Spatial Tokens, and perspective training. The next three rows ablate each
 435 component individually under the same setting (without perspective training). “Base (Full model w/o
 436 perspective training)” enables \mathcal{L}_{dep} , $\mathcal{L}_{\text{grav}}$, and Spatial Tokens but omits perspective training. “Full
 437 model” enables all components, including perspective training.

	DINOv3	\mathcal{L}_{dep}	$\mathcal{L}_{\text{grav}}$	Sp. Tok.	Per. Train	$\text{FID} \downarrow$	$\text{KID}_{\times 100} \downarrow$	$\text{RMSE} \downarrow$
Canonical Image-to-3D	✓					35.6	30.1	6.21
Base w/o \mathcal{L}_{dep}	✓			✓	✓	23.7	18.4	5.75
Base w/o $\mathcal{L}_{\text{grav}}$	✓		✓		✓	25.9	19.0	5.21
Base w/o Spatial Tokens	✓		✓	✓		24.8	18.1	5.64
Base (Full w/o per. training)	✓		✓	✓	✓	21.6	16.2	5.23
Full model	✓	✓	✓	✓	✓	19.2	13.6	5.20

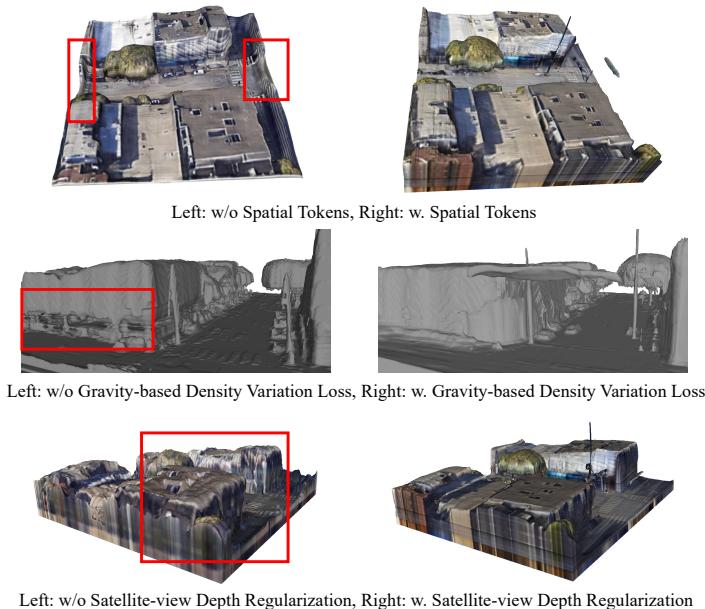


Figure 5: Qualitative ablation on key modules.

468 procedure for preparing the ground truth DSM pairs for our VIGOR-OOD test set is provided in the
 469 Appendix Section D. As shown in Table 3, we adopt the standard evaluation metrics from satellite
 470 stereo literature (Gao et al., 2021), including Mean Absolute Error (MAE), Root Mean Square Error
 471 (RMSE), and the percentage of pixels with errors below 2.5m and 7.5m.

472 The results clearly demonstrate the effectiveness of our proposed methodology. The state-of-the-art
 473 method, Sat2Density++, achieves an RMSE of 6.76m. Our controlled baseline, Canonical Image-to-
 474 3D, which benefits from a stronger backbone, already improves upon this with an RMSE of 6.21m.
 475 Our full model, however, significantly outperforms both, establishing a new state-of-the-art with an
 476 RMSE of 5.20m and a MAE of 3.47m. Notably, our model reconstructs 62.69% of the surface with an
 477 error of less than 2.5m, a substantial improvement over Sat2Density++ (49.69%). This underscores
 478 the superior capability of our method in generating geometrically precise urban scenes. In addition to
 479 these quantitative results, we provide extensive qualitative comparisons of the rendered satellite-view
 480 depth and the groundtruth DSM in the Appendix Fig. 8, which visually corroborate our model’s
 481 superior geometric fidelity.

482
 483 **Ablation Study.** We conduct a comprehensive ablation study to validate our design choices, with
 484 quantitative results in Table 2, Table 3 and qualitative visualizations in Fig. 5. Starting from our
 485 Canonical Image-to-3D baseline (FID 35.6, RMSE 6.21m), integrating our core geometric priors
 $(\mathcal{L}_{\text{dep}}, \mathcal{L}_{\text{grav}}, \text{Spatial Tokens})$ synergistically boosts performance to an FID of 21.6. Among them,

486

487

Table 3: Quantitative comparison for predicted DSM.

488

489

	MAE \downarrow	RMSE \downarrow	$< 2.5 \text{ m} \uparrow$	$< 7.5 \text{ m} \uparrow$
Sat2Density++	4.72	6.76	49.69	83.65
Canonical Image-to-3D	4.23	6.21	52.73	84.54
Base w/o $\mathcal{L}_{\text{grav}}$	3.53	5.21	61.17	88.94
Base w/o \mathcal{L}_{dep}	3.82	5.75	59.88	87.04
Base w/o Spatial Tokens	3.87	5.64	57.10	86.46
Base (Full w/o per. training)	3.52	5.23	61.97	88.52
Ours (Full model)	3.47	5.20	62.69	88.68

490

491

492

493

494

495

496

497

498

499 the gravity-based loss ($\mathcal{L}_{\text{grav}}$) is most critical for photorealism, as its removal causes the largest500 FID degradation (to 25.9). Conversely, removing \mathcal{L}_{dep} or Spatial Tokens leads to more significant

501 geometric errors (RMSE increases to 5.75m and 5.64m, respectively). Finally, adding perspective

502 training achieves our best results across both photorealism (FID 19.2) and geometric accuracy (RMSE

503 5.20m).

504

505 These quantitative gains are visually corroborated in Fig. 5. As intended, Spatial Tokens regularize

506 boundaries, $\mathcal{L}_{\text{grav}}$ yields straighter facades and reduces floaters, and \mathcal{L}_{dep} corrects rooftop geometry.

507 The consistent improvements across metrics and visuals confirm that each component is essential,

508 targeting distinct aspects of the final high-fidelity reconstruction.

509

510

511

512

513

514

515

516

517

Applications. We show a few application examples in Fig. 1. Additional downstream results and implementation details are provided in Appendix Section B, including satellite-to-DSM (metric depth) conversion without ground-truth depth data supervision (Fig. 9), semantic map to 3D reconstruction (Fig. 11), large-area 3D mesh generation from a single large satellite patch (Fig. 10), and surround-view multi-camera video synthesis from satellite imagery (Section B.2).

5 CONCLUSION

518

519

520

521

522

523

524

525

526

527

ETHICS STATEMENT

528

529

530

531

Our work builds on a publicly available dataset (VIGOR (Zhu et al., 2021), which consists of satellite images and street-view panoramas collected from widely accessible map platforms. We do not gather or release any personally identifiable information (PII). All data sources follow the original dataset licenses and terms of use, and no effort is made to identify individuals or private properties beyond what is already visible in the released benchmarks. Potential misuse of our approach should be considered: while our method advances urban-scale 3D scene reconstruction for beneficial applications such as autonomous driving simulation, AR/VR urban planning, and geographic visualization, it could also be applied to large-scale surveillance if deployed irresponsibly. To mitigate such risks, we emphasize that our framework is intended solely for academic research and positive societal use cases, and we release neither additional sensitive data nor pre-trained models tied to private regions.

540 REPRODUCIBILITY STATEMENT
541

542 We aim to ensure the reproducibility of our results. All implementation is based on standard deep
543 learning frameworks (PyTorch (Paszke et al., 2019)), and our model architecture, training strategies,
544 and evaluation protocols are fully described in Sec. 4 and Appendix F. We specify dataset splits,
545 preprocessing steps, and evaluation metrics in detail, and we adopt widely used benchmarks (VIGOR)
546 to facilitate fair comparisons. Hyperparameters, batch sizes, number of iterations, and hardware
547 configurations are reported in the Implementation Details paragraph. We will release the training
548 scripts, configuration files, and inference code, together with instructions for data preparation and
549 evaluation, upon publication. This enables other researchers to reproduce our quantitative scores and
550 qualitative visualizations, and to extend our method to new geographic regions or related cross-view
551 generation tasks.

552 REFERENCES
553

554 Dragomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stéphane Lafon, Richard Lyon,
555 Abhijit Ogale, Luc Vincent, and Josh Weaver. Google street view: Capturing the world at street
556 level. *Computer*, 43(6):32–38, 2010.

557 Mikolaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD
558 GANs. In *ICLR*, 2018.

559 James B Campbell and Randolph H Wynne. *Introduction to remote sensing*. Guilford press, 2011.

560 Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio
561 Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein.
562 Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022.

563 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image
564 synthesis. In *CVPR*, pp. 12873–12883, 2021.

565 Jian Gao, Jin Liu, and Shunping Ji. Rational polynomial camera model warping for deep learning
566 based satellite multi-view stereo matching. In *Proceedings of the IEEE/CVF International
567 Conference on Computer Vision (ICCV)*, pp. 6148–6157, October 2021.

568 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans
569 trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, pp.
570 6629–6640, 2017.

571 Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli,
572 Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3d. In *ICLR*, 2024.
573 URL <https://openreview.net/forum?id=s11U8vvsFF>.

574 Tongyan Hua, Lutao Jiang, Ying-Cong Chen, and Wufan Zhao. Sat2city: 3d city generation from a
575 single satellite image with cascaded latent diffusion. In *ICCV*, 2025.

576 Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt
577 Keutzer. SqueezeNet: Alexnet-level accuracy with 50x fewer parameters and $\frac{1}{10}$mb model size,
578 2016. URL <https://arxiv.org/abs/1602.07360>.

579 Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing
580 and improving the image quality of stylegan. In *CVPR*, pp. 8107–8116, 2020. doi: 10.1109/
581 CVPR42600.2020.00813.

582 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional
583 neural networks. In *NeurIPS*, pp. 1106–1114, 2012.

584 Zuoyue Li, Zhenqiang Li, Zhaopeng Cui, Rongjun Qin, Marc Pollefeys, and Martin R. Oswald.
585 Sat2vid: Street-view panoramic video synthesis from a single satellite image. In *ICCV*, pp.
586 12416–12425, 2021. doi: 10.1109/ICCV48922.2021.01221.

587 Zuoyue Li, Zhenqiang Li, Zhaopeng Cui, Marc Pollefeys, and Martin R. Oswald. Sat2scene: 3d
588 urban scene generation from satellite images with diffusion. In *CVPR*, pp. 7141–7150, June 2024.

594 Xiaohu Lu, Zuoyue Li, Zhaopeng Cui, Martin R. Oswald, Marc Pollefeys, and Rongjun Qin.
 595 Geometry-aware satellite-to-ground image synthesis for urban areas. In *CVPR*, pp. 856–864,
 596 2020.

597 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 598 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
 599 high-performance deep learning library. *Advances in neural information processing systems*, 32,
 600 2019.

601 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 602 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
 603 synthesis. In *ICLR*, 2023.

604 Ming Qian, Jincheng Xiong, Gui-Song Xia, and Nan Xue. Sat2density: Faithful density learning
 605 from satellite-ground image pairs. In *ICCV*, pp. 3660–3669, 2023. doi: 10.1109/ICCV51070.2023.
 606 00341.

607 Ming Qian, Bin Tan, Qiuyu Wang, Xianwei Zheng, Hanjiang Xiong, Gui-Song Xia, Yujun Shen, and
 608 Nan Xue. Seeing through satellite images at street views, 2025. URL <https://arxiv.org/abs/2505.17001>.

609 René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust
 610 monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE TPAMI*, 44
 611 (3), 2022.

612 Krishna Regmi and Ali Borji. Cross-view image synthesis using conditional gans. In *CVPR*, June
 613 2018.

614 Krishna Regmi and Ali Borji. Cross-view image synthesis using geometry-guided conditional gans.
 615 *Comput. Vis. Image Underst.*, 187, 2019.

616 Yujiao Shi, Dylan Campbell, Xin Yu, and Hongdong Li. Geometry-guided street-view panorama
 617 synthesis from satellite imagery. *IEEE TPAMI*, 44(12):10009–10022, 2022.

618 Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose,
 619 Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel
 620 Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana,
 621 Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé
 622 Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025. URL <https://arxiv.org/abs/2508.10104>.

623 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 624 Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021.
 625 URL <https://openreview.net/forum?id=PxTIG12RRHS>.

626 Hao Tang, Dan Xu, Nicu Sebe, Yanzhi Wang, Jason J. Corso, and Yan Yan. Multi-channel attention
 627 selection GAN with cascaded semantic guidance for cross-view image translation. In *CVPR*, pp.
 628 2417–2426, 2019.

629 Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large
 630 multi-view gaussian model for high-resolution 3d content creation. In *ECCV*, pp. 1–18, 2024.

631 A. Toker, Q. Zhou, M. Maximov, and L. Leal-Taixe. Coming down to earth: Satellite-to-street view
 632 synthesis for geo-localization. In *CVPR*, pp. 6484–6493, 2021. doi: 10.1109/CVPR46437.2021.
 633 00642.

634 Scott Workman, Menghua Zhai, David J Crandall, and Nathan Jacobs. A unified model for near and
 635 remote sensing. In *ICCV*, pp. 2688–2697, 2017.

636 Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin
 637 Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv*
 638 preprint [arXiv:2412.01506](https://arxiv.org/abs/2412.01506), 2024.

648 Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. CityDreamer: Compositional generative
 649 model of unbounded 3D cities. In *CVPR*, 2024.
 650

651 Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh:
 652 Efficient 3d mesh generation from a single image with sparse-view large reconstruction models.
 653 *arXiv preprint arXiv:2404.07191*, 2024.

654 Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang
 655 Zhao. Depth anything v2. *arXiv:2406.09414*, 2024.
 656

657 Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural Radiance Fields
 658 from One or Few Images . In *CVPR*, pp. 4576–4585, June 2021. doi: 10.1109/CVPR46437.2021.
 659 00455.
 660

661 Xianghui Ze, Zhenbo Song, Qiwei Wang, Jianfeng Lu, and Yujiao Shi. Controllable satellite-to-
 662 street-view synthesis with precise pose alignment and zero-shot environmental control. In *ICLR*,
 663 2025. URL <https://openreview.net/forum?id=f92M45YRfh>.

664 Menghua Zhai, Zachary Bessinger, Scott Workman, and Nathan Jacobs. Predicting ground-level
 665 scene layout from aerial imagery. In *CVPR*, pp. 4132–4140, 2017.
 666

667 Jiahui Zhang, Yuelei Li, Anpei Chen, Muyu Xu, Kunhao Liu, Jianyuan Wang, Xiao-Xiao Long,
 668 Hanxue Liang, Zexiang Xu, Hao Su, et al. Advances in feed-forward 3d reconstruction and view
 669 synthesis: A survey. *arXiv preprint arXiv:2507.14501*, 2025.
 670

671 Jiaming Zhang, Kailun Yang, Chaoxiang Ma, Simon Reiß, Kunyu Peng, and Rainer Stiefelhagen.
 672 Bending reality: Distortion-aware transformers for adapting to panoramic semantic segmentation.
 673 In *CVPR*, pp. 16896–16906, 2022.

674 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 675 diffusion models. In *ICCV*, 2023a.

676 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 677 diffusion models. In *ICCV*, 2023b.

678 R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep
 679 features as a perceptual metric. In *CVPR*, pp. 586–595, 2018. doi: 10.1109/CVPR.2018.00068.
 680

681 Yichao Zhou, Jingwei Huang, Xili Dai, Linjie Luo, Zhili Chen, and Yi Ma. Holarity: A city-scale
 682 data platform for learning holistic 3d structures. *CoRR*, abs/2008.03286, 2020.

683 Sijie Zhu, Taojiannan Yang, and Chen Chen. Vigor: Cross-view image geo-localization beyond
 684 one-to-one retrieval. In *CVPR*, pp. 3640–3649, 2021.
 685

686 This appendix provides more details and results of our Sat3DGen.
 687

688 LLM USAGE

689 In this section, we clarify the role of large language models (LLMs) in preparing this work. The
 690 model was used exclusively for language polishing, such as refining grammar, style, and readability,
 691 without contributing to the research design, analysis, or conclusions.
 692

693 A MORE 3D AND VIDEO RESULTS

694 A.1 MESH RESULTS COMPARED TO SAT2DENSITY++

695 We provide additional mesh results in Fig. 6, with all satellite images sourced from the VIGOR-OOD
 696 test set.
 697

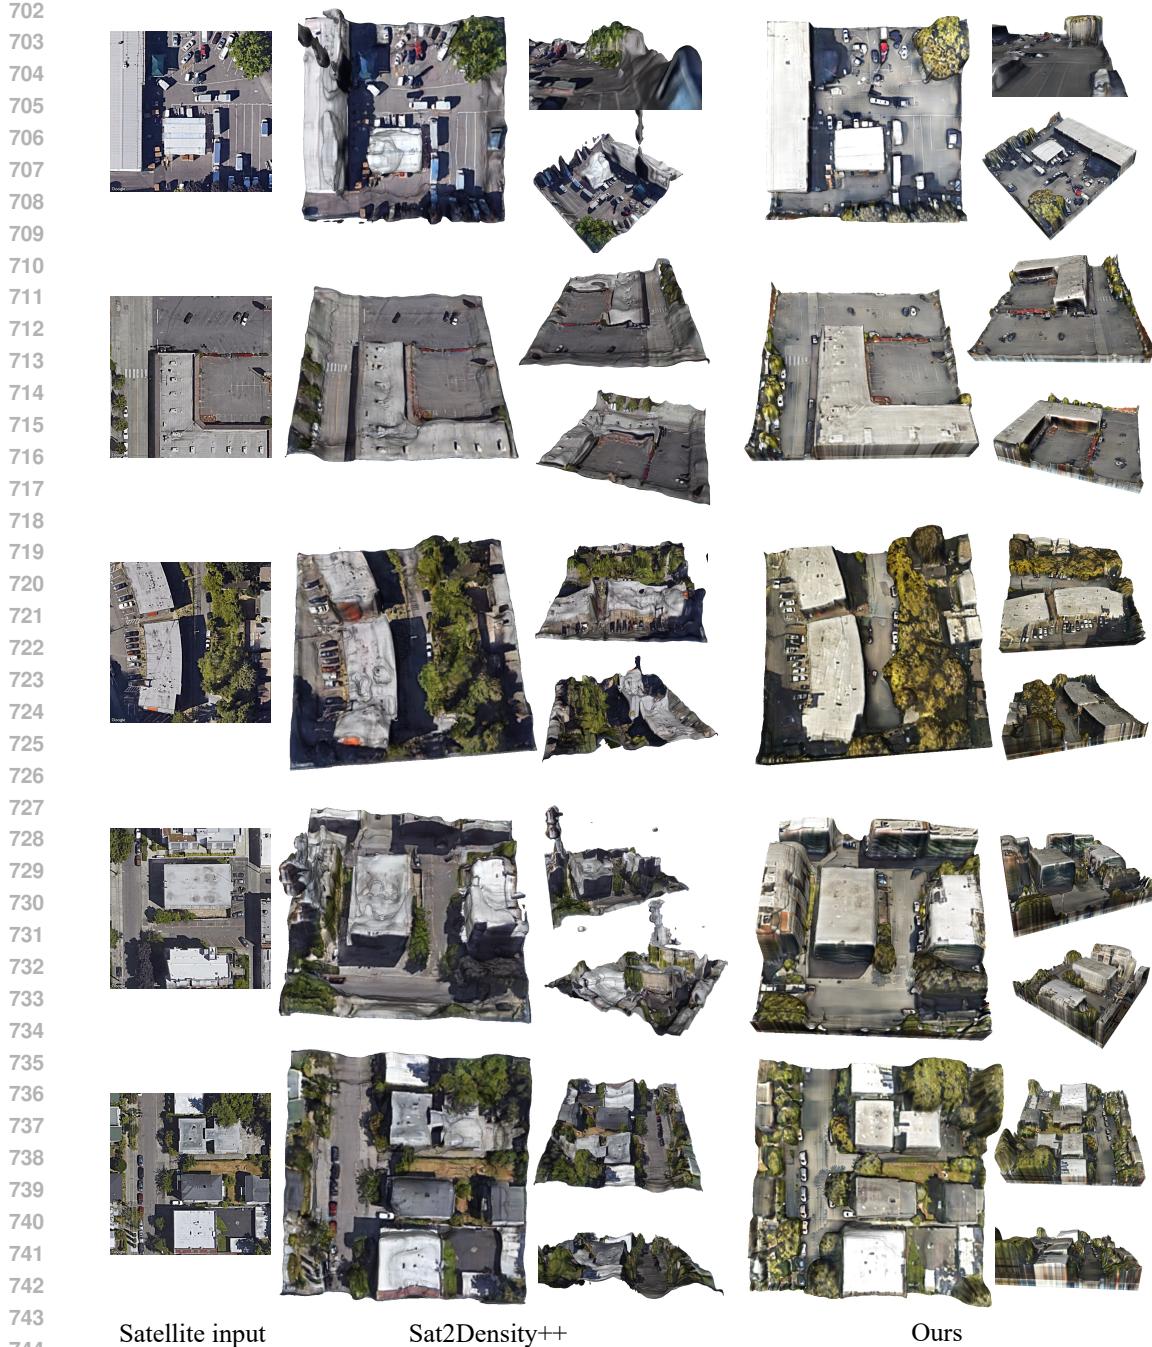


Figure 6: Comparison of generation 3D assets between Sat2Density++(Qian et al., 2025) and Ours.

A.2 MESH RESULTS COMPARED TO CANONICAL IMAGE-TO-3D

As shown in Fig. 7, the Canonical Image-to-3D baseline produces significantly inferior geometry compared to our full model. Specifically, its ground surfaces are noisy and uneven, and it fails to capture distinct shapes for trees or produce flat building rooftops. The mesh boundaries also exhibit irregular, spiky extrusions. These geometric flaws result in a much lower rate of watertight meshes. In contrast, our model consistently generates smoother surfaces, more plausible object structures,

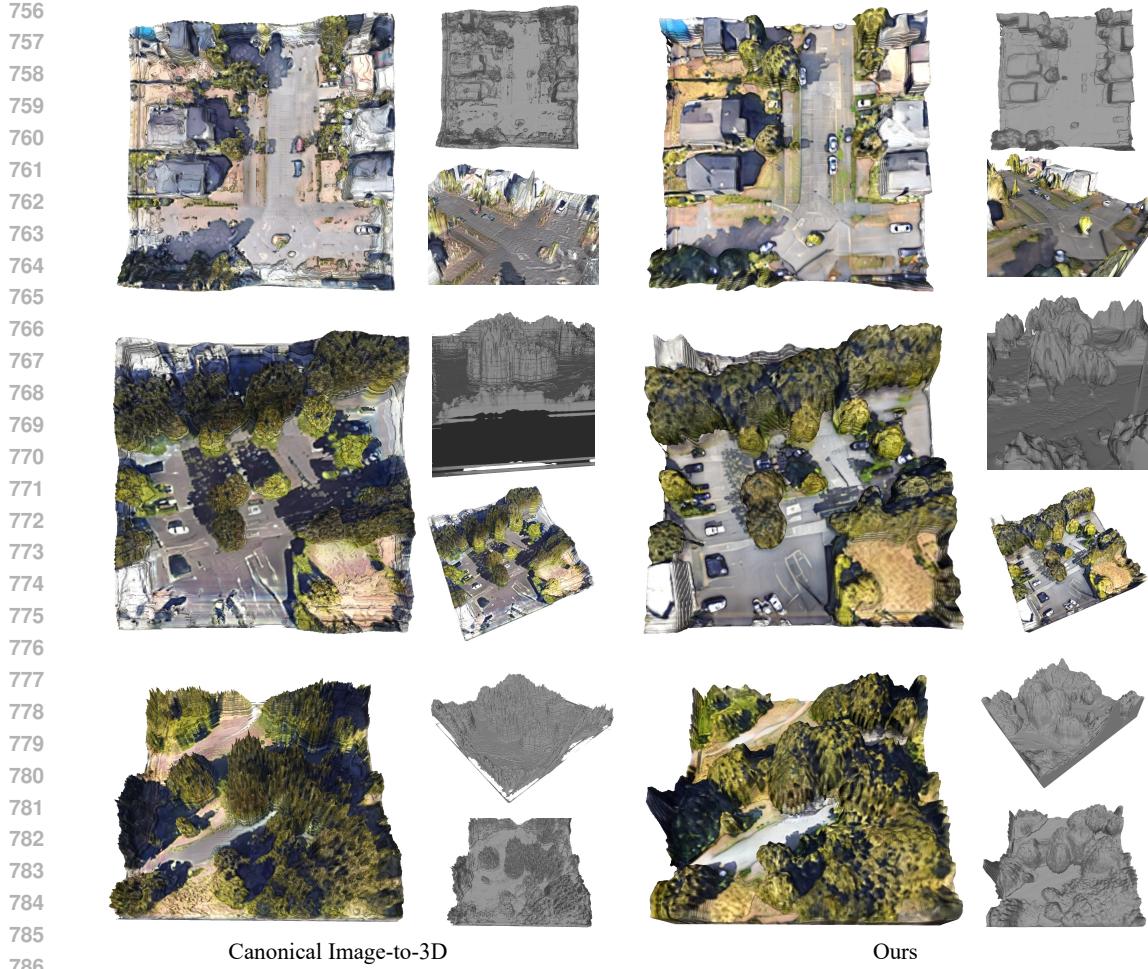


Figure 7: Comparison of generation 3D assets between the baseline Canonical image-to-3D model and Ours.

and cleaner boundaries. This visual evidence strongly corroborates the quantitative results in Table 2, where the baseline’s poor geometric fidelity is reflected in its high RMSE score.

A.3 PANORAMA VIDEO RESULTS.

In the supplementary ZIP archive, the ‘video_result’ folder contains results from 16 sets of VIGOR-OOD test data. Each video is named according to its latitude and longitude, allowing you to view the latest satellite images by entering these coordinates into Google Maps². In each video, the top-left corner displays the satellite image and camera trajectory, the top-right section presents the results of the Sat2Density++ (Qian et al., 2025), and the adjacent section showcases our results. From the presented videos, it is clear that our method, by producing enhanced 3D assets, achieves superior panorama video generation—yielding more regular building shapes, fewer floating artifacts, and improved generation of ground vehicles, road signs, and trees.

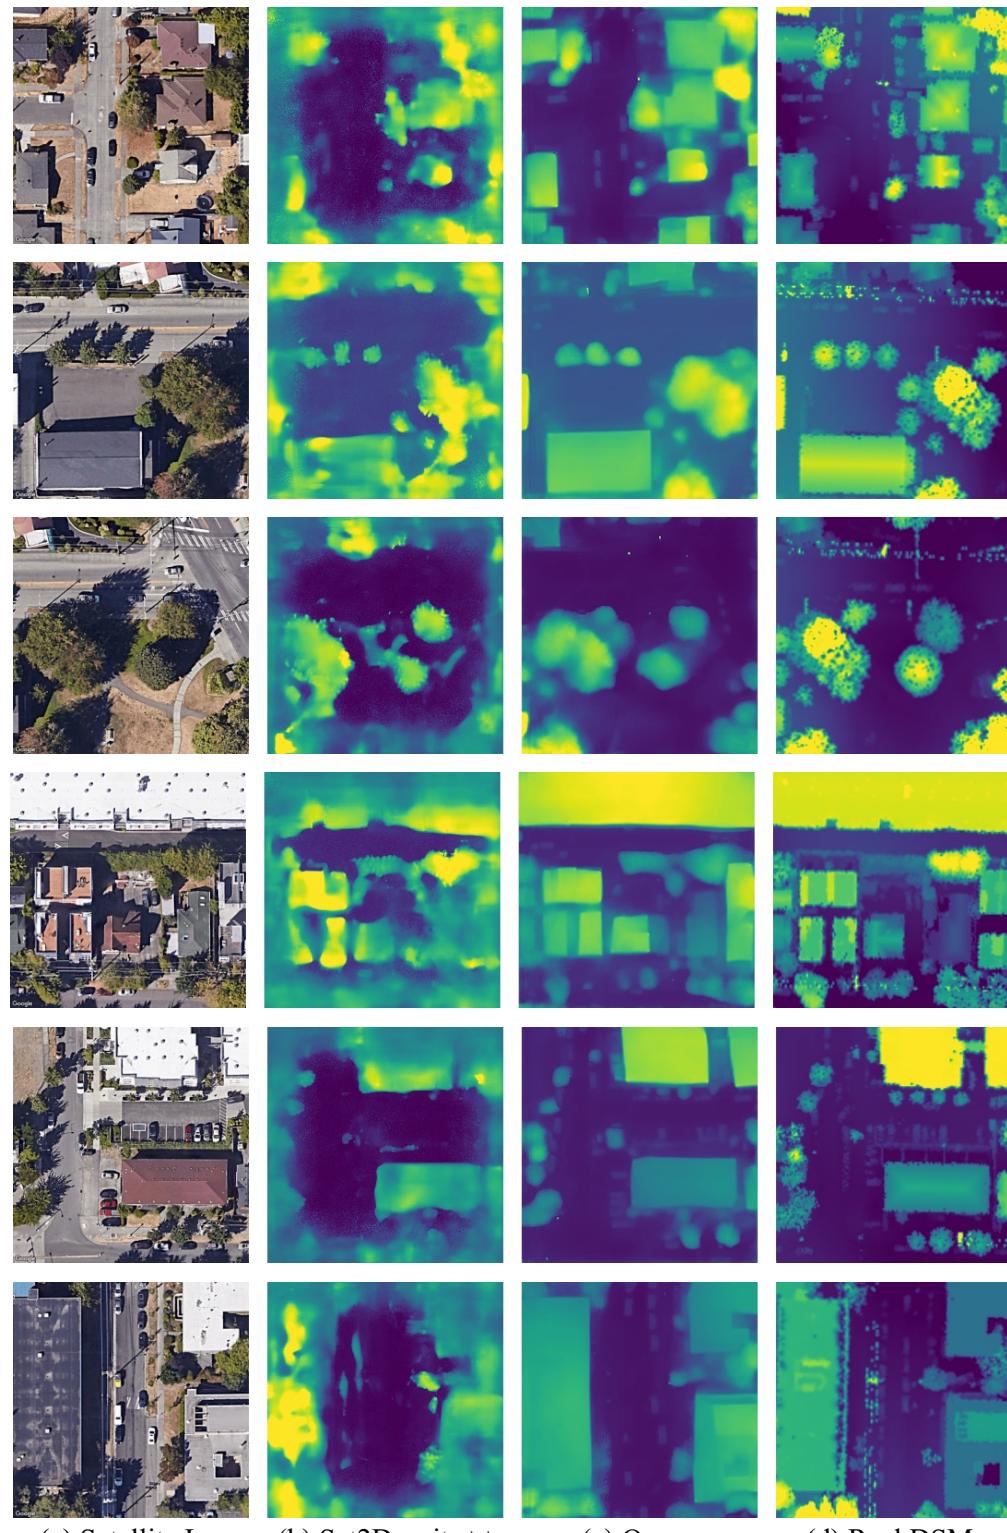
A.4 SATELLITE-VIEW DEPTH COMPARISON

We present a qualitative comparison of the generated depth maps in Fig. 8. It is important to note the inherent challenge of temporal misalignment between the satellite images and the ground truth DSM,

²<https://www.google.com/maps/>

810

811



860

861

862

863

Figure 8: Comparison of predicted satellite-view height map.

864 which was collected over a year (Section D). This leads to discrepancies from transient objects (e.g.,
 865 vehicles) and noise within the DSM itself, making a perfect reconstruction unattainable.
 866

867 Despite these challenges, a clear difference emerges. Sat2Density++ (b) yields overly-smoothed
 868 results with indistinct building edges and noisy ground planes. In contrast, our model (c) generates
 869 significantly sharper and more geometrically plausible structures, featuring flat rooftops and clean
 870 ground surfaces. These visual improvements directly corroborate our superior quantitative metrics
 871 reported in the main paper, demonstrating a better capability to reconstruct high-fidelity geometry
 872 from a single image.
 873

874 B MORE APPLICATIONS RESULTS.

876 B.1 SINGLE SATELLITE IMAGE TO DSM (METRIC DEPTH)

877 AS shown in Fig. 9, our model can render satellite-view metric depth from the learned NeRF-based
 878 3D representation, even though no metric-depth annotations are used during training.
 879

881 B.2 SURROUND-VIEW MULTI-CAMERA VIDEO GENERATION FROM A SINGLE SATELLITE 882 IMAGE.

883 We provide some surround-view multi-camera video results in Fig. 4 and the supplementary ZIP
 884 archive, with four fixed 120-degree FOV perspectives shown at the bottom of each video. Our
 885 algorithm leverages NeRF-based 3D representations, allowing the generation of perspective im-
 886 ages/videos with varying FOV and image sizes. The videos demonstrate that the quality of our
 887 generated perspective images is nearly on par with panorama images, underscoring the superiority of
 888 our model design.
 889

890 To the best of our knowledge, our approach is the first to generate diverse content in multi-view
 891 perspective videos from a single satellite image without requiring video data or 3D geometry as
 892 training input. The most relevant existing method for generating perspective videos from a single
 893 satellite image is Sat2Scene (Li et al., 2024). However, this method is limited to synthesizing
 894 buildings and ground surfaces, primarily due to its strong reliance on building height maps converted
 895 into point clouds. Moreover, as their code and dataset are not fully open-sourced, a direct comparison
 896 with our results is not feasible. In summary, we demonstrate the capability of our algorithm to
 897 generate multi-view perspective videos, underscoring its potential applications in vehicle driving
 898 simulation.
 899

900 B.3 LARGE-SCALE MESH GENERATION

901 The process for generating a big mesh from a large satellite image is as follows: We first download
 902 an extensive satellite image from online platforms and then resize it to ensure the per-pixel resolution
 903 remains consistent with the pixel resolution of our training data. We then perform inference using a
 904 sliding window approach, processing each patch. Each patch resolution is 256x256 with a step set to
 905 128. For the density at the edges, we simply average the values to obtain the final result. In our demo,
 906 as shown in Fig. 10, the input for the large-scale satellite image is at zoom level 19, covering a spatial
 907 area of approximately 150 m × 150 m. In theory, we can download larger remote sensing images to
 908 generate results over even broader areas. However, at a fixed zoom level, the Google Maps Static API
 909 returns at most 640×640 pixels per request.
 910

911 Our demo results clearly demonstrate that the proposed algorithm produces smooth and coherent
 912 ground surfaces, seamlessly integrating multiple patches. The transitions between buildings show
 913 minimal discontinuities, highlighting the effectiveness of our method and its strong potential for
 914 large-scale 3D content generation.
 915

916 B.4 SEMANTIC MAPS TO 3D ASSETS.

917 Generating 3D scenes from 2D semantic maps is a highly effective application. We can utilize open
 918 street map data to obtain ground semantic maps or directly create semantic maps through drawing.
 919

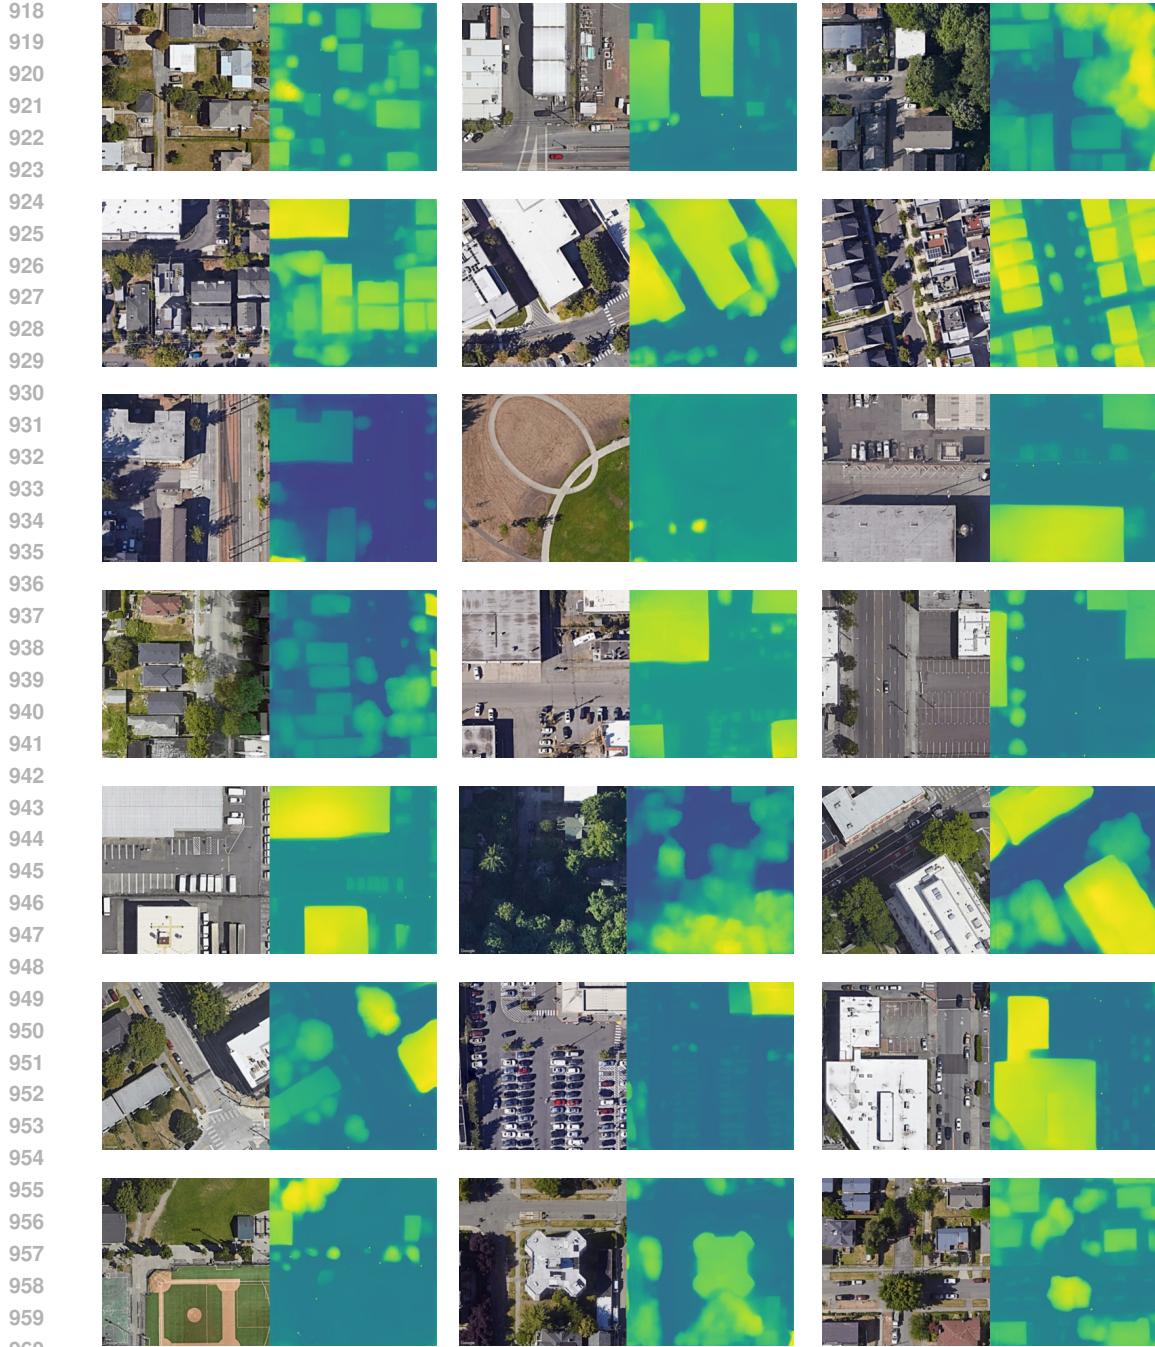


Figure 9: Visual results of our model generated DSM (metric depth) from the monocular satellite image, which is rendered from the satellite view, with no metric depth data for training.

In our work, we collected multiple semantic map-satellite image pairs using OpenStreetMap³ and Google Satellite Maps⁴ to train an additional model for transforming color semantic maps into satellite images. This model is based on diffusion (Song et al., 2021), composed of ControlNet (Zhang et al., 2023a) and SDXL (Podell et al., 2023).

³<https://www.openstreetmap.org/>

⁴<https://mapsplatform.google.com/>

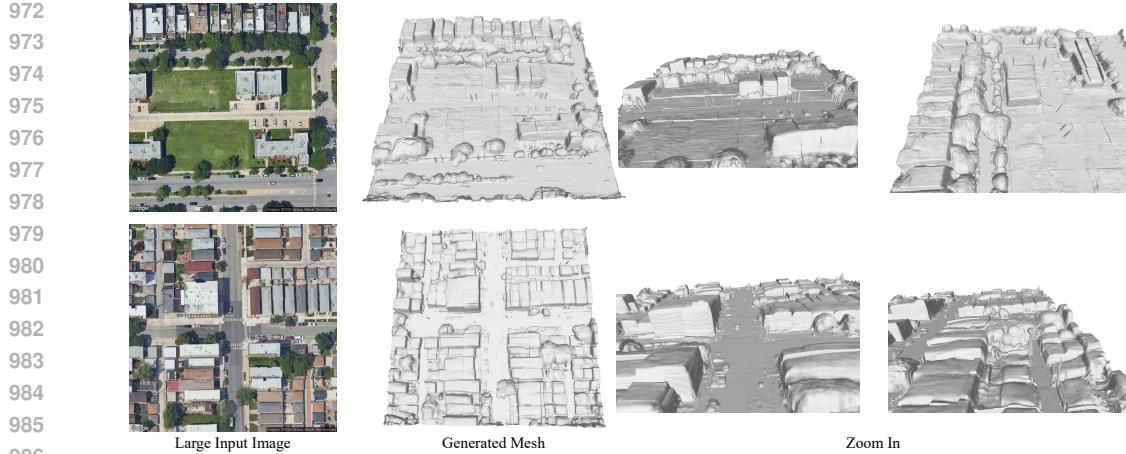


Figure 10: Given a large satellite image, our model can generate mesh with sliding window inference mode.

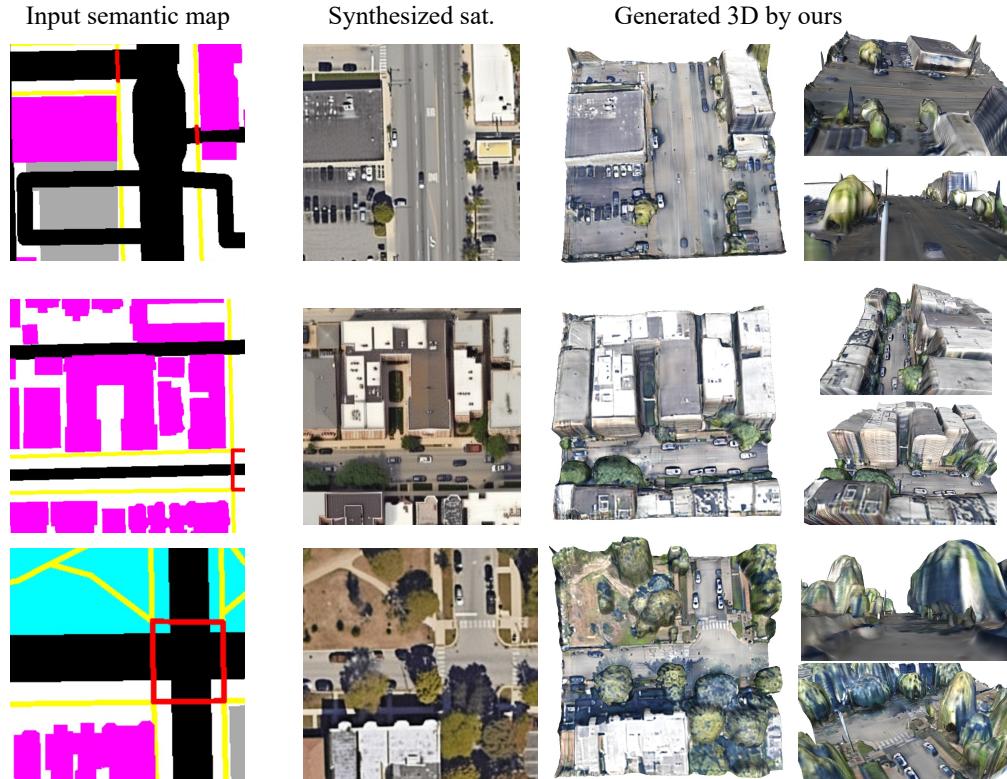


Figure 11: Given a colored semantic map, our model can generate 3D mesh through a pipeline that first converts the semantic map to a satellite map and then transforms the satellite map into 3D assets.

As shown in Fig. 11, given a colored semantic map, the diffusion model first generates a satellite map, followed by Sat3DGen generating 3D assets based on the satellite image input. The results indicate that the generated 3D assets maintain spatial consistency with the semantic positions in the input semantic map. This application of generating 3D assets from semantic maps can be advantageous in spatial planning and game modeling, underscoring the significance of our work.

1026 **C IMPLEMENTATION DETAILS**
10271028 **Perspective sampling.** During training, we randomly obtain perspective images from panorama
1029 street view images. We set the pitch range to [-30, 30] and fix the roll at 0. The yaw is selected within
1030 the range [-79, 179], and the field of view (FOV) is randomly chosen from [90, 105, 120]; the render
1031 size is 256×256 .
10321033 **Volume Rendering.** The triplane dimension is set to 32, and there are 96 samples per ray.
10341035 **Other Hyperparameters.** In our reconstruction of the panorama and satellite view perspectives,
1036 as well as in the sections of GAN loss and opacity loss, we maintain the same weight settings as
1037 Sat2Density++. For the newly introduced *Gravity-based Density Variation Loss*, we assign a weight
1038 of 3.5. As for the Satellite View Depth Regularization based on Midas loss, we set the weight to 0.1,
1039 acknowledging that the pseudo-labels for relative depth predicted by existing models may not be
1040 entirely accurate, hence opting for a smaller weight to mitigate potential adverse effects on the model.
10411042 Regarding the perspective view image reconstruction loss, we assign a weight that is half of the
1043 panorama view reconstruction loss. Conversely, the weight for the perspective view GAN loss is
1044 kept consistent with the panorama view GAN loss weight. This approach is due to the significant
1045 detail present in perspective images and the substantial difference between input satellite views and
1046 output perspective views. Given that the satellite input offers limited effective information for such
1047 detailed perspective views, we aim to prioritize realism and image quality in generating perspective
1048 images (as constrained by GAN loss) rather than achieving visual consistency with the ground truth
1049 images (which is the target of the reconstruction loss constraint). The full code will be released after
1050 acceptance.
10511052 **D GROUND TRUTH DSM PREPARATION FOR VIGOR-OOD**
10531054 Quantitative evaluation of 3D geometry requires accurate ground truth Digital Surface Models
1055 (DSMs). However, high-quality, publicly available LiDAR-derived DSM data is scarce and typically
1056 limited to a few cities. We were fortunate that our out-of-distribution (OOD) test set, VIGOR-OOD,
1057 is based on the city of Seattle, for which we were able to obtain a corresponding high-precision DSM
1058 dataset. This section details the entire pipeline for processing and aligning this raw data to create the
1059 ground truth for our geometric evaluation.
10601061 **D.1 DATA SOURCE AND JUSTIFICATION**
10621063 The ground truth DSMs were derived from the **King County West 2021** dataset, which is publicly
1064 available through the Washington State Department of Natural Resources (DNR) Lidar Portal⁵. As
1065 shown in the data's official report, this consists of Quality Level 1 (QL1) LiDAR data collected in
1066 the spring and summer of 2021. We specifically chose this data as its acquisition timeline closely
1067 matches the period when the VIGOR dataset (Zhu et al., 2021) was being created.
10681069 We selected six large GeoTIFF tiles from this collection that collectively cover the geographical
1070 extent of the VIGOR-OOD test images in Seattle, as visualized in Fig. 12. The official metadata
1071 confirms the high quality of this source data, reporting positional errors of less than 5.6 cm with 95%
1072 confidence, which provides a reliable basis for our geometric evaluation.
10731074 **D.2 METHODOLOGY FOR SATELLITE-DSM ALIGNMENT**
10751076 A significant technical challenge lies in aligning the raw DSM GeoTIFF files with the individual
1077 satellite images from the VIGOR-OOD test set. The datasets use different Coordinate Reference
1078 Systems (CRS), resolutions, and are not spatially aligned. We developed a robust processing pipeline
1079 to address this, which is outlined in Algorithm 1.
10801081 The key to our approach is leveraging the metadata provided by the VIGOR dataset, which, fortunately,
1082 includes the precise WGS84 latitude/longitude coordinates and the Google Maps zoom level for each
10835^{https://lidarportal.dnr.wa.gov/}

1080 satellite image. This allows us to first estimate the geographic bounding box of each image. We
 1081 then reproject the corresponding section of the high-resolution DSM onto the exact pixel grid of the
 1082 satellite image, ensuring perfect alignment. The main steps are detailed below.
 1083
 1084

Algorithm 1 DSM Ground Truth Preparation Pipeline

1: **Input:** Set of VIGOR-OOD satellite images I , set of raw DSM GeoTIFF tiles T .
 2: **Output:** A ground truth DSM array D_i for each valid input image I_i .
 3: Initialize an in-memory spatial index `dsm_index` for all tiles in T for fast lookups.
 4: **for** each satellite image I_i in I **do**
 5: {Step 1: Estimate satellite image’s geographic footprint}
 6: Parse center latitude/longitude and zoom level from I_i ’s filename.
 7: Estimate the WGS84 bounding box `target_bounds` for I_i .
 8: {Step 2: Find and Reproject Overlapping DSM Data}
 9: Find all candidate tiles `candidate_tiles` from `dsm_index` that spatially overlap with
 10: `target_bounds`.
 11: **if** `candidate_tiles` is empty **then**
 12: **continue** {No DSM coverage for this image}
 13: **end if**
 14: Initialize `best_dsm` to ‘null’ and `max_coverage` to 0.
 15: **for** each tile T_j in `candidate_tiles` **do**
 16: Create an empty destination array `temp_dsm` with the same dimensions as I_i .
 17: Reproject the data from T_j onto `temp_dsm` using bilinear resampling.
 18: **if** coverage of `temp_dsm` > `max_coverage` **then**
 19: `best_dsm` \leftarrow `temp_dsm`, `max_coverage` \leftarrow coverage.
 20: **end if**
 21: **end for**
 22: {Step 3: Post-processing and Quality Control}
 23: **if** `best_dsm` is not ‘null’ **then**
 24: Convert elevation values in `best_dsm` from feet to meters.
 25: Calculate the percentage of NaN pixels `nan_percent` in `best_dsm`.
 26: **if** `nan_percent` \leq 5.0% **then**
 27: Save `best_dsm` as the final ground truth D_i .
 28: **else**
 29: Discard this pair due to insufficient DSM coverage.
 30: **end if**
 31: **end for**

1116
 1117
 1118 This automated pipeline ensures that for every test image, we generate an aligned, unit-corrected, and
 1119 quality-controlled DSM that can be used for direct, pixel-wise comparison in our geometric accuracy
 1120 evaluation. Fig. 8 shows visualizations of several paired satellite images and their corresponding
 1121 DSM data.
 1122
 1123

E MORE ABLATIONS

E.1 VARIATION REGULARIZATION ABLATION.

1124 As shown in Table 4, removing variation regularization gives the worst results (FID 25.90, KID 19.0).
 1125 Replacing it with a TV loss slightly lowers FID (24.83) but hurts KID (20.1), indicating oversmoothing
 1126 and weak structural guidance. Our $\mathcal{L}_{\text{grav}}$ outperforms both, with $\epsilon = 1.0$ yielding the best metrics
 1127 (FID 21.60, KID 16.2). Setting $\epsilon = 0$ degrades performance because it fails to accommodate genuine
 1128 voids (e.g., gaps beneath tree canopies outside the trunk), leading to over-penalization and smoothing.
 1129
 1130
 1131
 1132
 1133

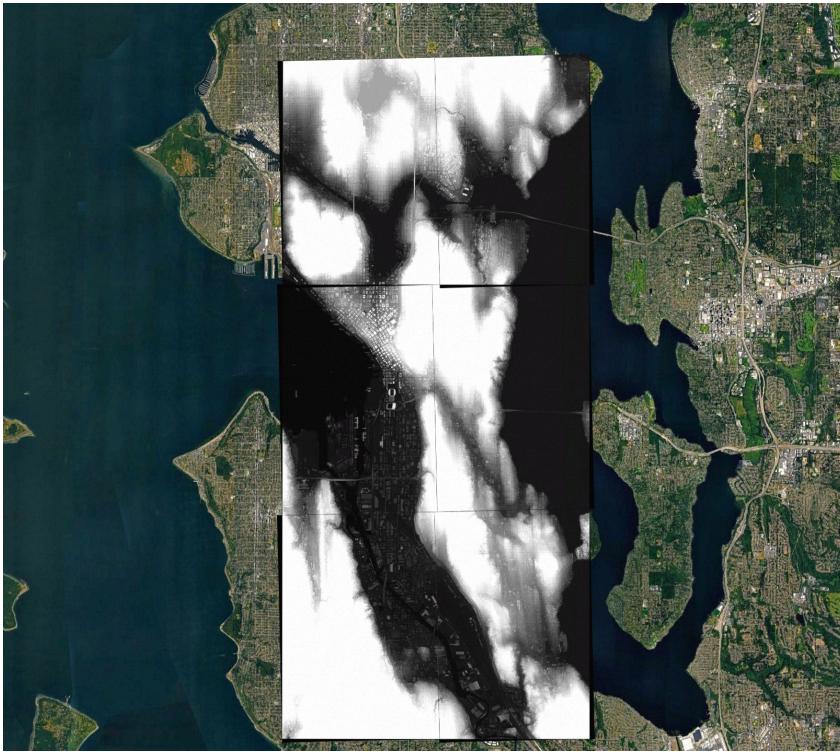


Figure 12: The collected DSM data in Seattle City.

Table 4: Ablation on variation regularization. “No variation loss” means starting from the base (Full model w/o perspective training) and removing $\mathcal{L}_{\text{grav}}$. “w. TV loss” replaces $\mathcal{L}_{\text{grav}}$ with total variation (TV) loss Chan et al. (2022). The remaining rows use $\mathcal{L}_{\text{grav}}$ with different ϵ ; our choice $\epsilon=1.0$ (Ours) performs best overall. Lower is better.

Method	FID \downarrow	KID $\times 100\downarrow$
No variation loss	25.90	19.0
w. TV loss	24.83	20.1
$\mathcal{L}_{\text{grav}} (\epsilon=0)$	24.52	18.7
$\mathcal{L}_{\text{grav}} (\epsilon=0.01)$	22.61	17.2
$\mathcal{L}_{\text{grav}} (\epsilon=0.1)$	22.63	16.2
$\mathcal{L}_{\text{grav}} (\epsilon=0.5)$	21.94	17.9
$\mathcal{L}_{\text{grav}} (\epsilon=1.0)$	21.60	16.2
$\mathcal{L}_{\text{grav}} (\epsilon=5.0)$	21.74	18.5
$\mathcal{L}_{\text{grav}} (\epsilon=10.0)$	21.66	17.5

F DATASET PREPARATION AND DETAILS

VIGOR. VIGOR (Zhu et al., 2021) contains four cities and, for each satellite image, multiple associated street-view panoramas with known relative poses. The satellite zoom level is 20. We train in Chicago, New York, and San Francisco, and use Seattle as an unseen-city OOD test set. We use 78,188 satellite–panorama pairs for training and 11,875 pairs for metric evaluation. In principle, additional cities could be incorporated to further scale training.

Discussion on other Datasets. CVACT (Shi et al., 2022) and CVUSA (Zhai et al., 2017) are widely used for satellite-to-street image synthesis and localization. However, they are unsuitable for our 3D generation task due to two main reasons.

1188 First, their data structure provides **insufficient geometric supervision**. Each satellite tile in these
 1189 datasets is paired with a single ground-level image at the tile center, without multi-view coverage. This
 1190 data design is incompatible with our method, which relies on panoramas captured at spatially distinct
 1191 positions to learn consistent street-level 3D. With only a centered street view, the supervision is too
 1192 sparse to learn reliable geometry, rendering our optimization components ineffective. Consequently,
 1193 we do not use CVACT or CVUSA for training. In practice, VIGOR-style datasets can be constructed
 1194 for any geographic region where street-level panoramic imagery is available (e.g., via Google Maps)
 1195 by pairing a single satellite image with N panoramic views.

1196 Second, these datasets are unsuitable for a fair out-of-distribution (OOD) evaluation. The satellite
 1197 images in CVACT and CVUSA were captured by different satellites and at different zoom levels
 1198 than those in VIGOR. This results in significant appearance shifts and differences in pixel-level
 1199 spatial resolution. Evaluating our VIGOR-trained model on these datasets would conflate the test of
 1200 geometric generalization with a test of resilience to data source domain shift, which is beyond the
 1201 scope of this work. We test on an unseen city from the same data source.

1202 **Our Strategy for Generalization Evaluation.** We explicitly test generalization using the VIGOR-
 1203 OOD split, by training on three cities and testing on the unseen city of Seattle. This setup introduces
 1204 a significant domain gap in terms of urban layouts and architectural styles, providing a strong signal
 1205 of our model’s robustness. Furthermore, the notion that VIGOR is purely ”urban” is a misconception.
 1206 Like other datasets derived from vehicle-based captures, it covers a wide range of environments,
 1207 including many less-dense, suburban areas, as shown in our qualitative results (e.g., Figure 7). The
 1208 primary generalization challenge is not urban vs. suburban density, but rather the architectural and
 1209 environmental shifts between disparate geographic regions.

1210 G LIMITATIONS

1211 Our work faces challenges rooted in both data availability and model assumptions.

1212 **Pose Inaccuracy.** A primary limitation is the lack of precise pose data. We treat satellite images as
 1213 ideal orthogonal projections, which is not always the case. Moreover, the panoramas lack authentic
 1214 intrinsic/extrinsic parameters; we only have GPS data. Our model assumes panoramas are captured
 1215 perpendicular to the ground, ignoring roll angles from terrain or road banking. Accessing or predicting
 1216 accurate poses is a valuable direction for future work.

1217 **Geometric and Terrain Assumptions.** Our model’s performance is also constrained by its underlying
 1218 assumptions. By its generative nature, it struggles with **atypical architectures** that are rare in the
 1219 training data, as we lack explicit 3D ground-truth shapes for supervision. Furthermore, our framework
 1220 assumes a **locally flat ground plane** and does not model significant terrain variations like hills, which
 1221 are difficult to infer from sparse imagery alone. Future work could address this by incorporating
 1222 multi-modal data, such as terrain maps.

1223 **Evaluation Metrics.** Finally, while metrics like multi-view photometric consistency or temporal
 1224 flicker are powerful, they are not applicable to the VIGOR dataset. VIGOR’s sparse, non-sequential
 1225 collection of still images does not support the evaluation of temporal stability or dense view consis-
 1226 tency.