

Sat2Scene: 3D Urban Scene Generation from Satellite Images with Diffusion

Zuoyue Li¹ Zhenqiang Li² Zhaopeng Cui³ Marc Pollefeys^{1,4} Martin R. Oswald^{1,5}
¹ETH Zürich ²The University of Tokyo ³Zhejiang University
⁴Microsoft ⁵University of Amsterdam

Abstract

Directly generating scenes from satellite imagery offers exciting possibilities for integration into applications like games and map services. However, challenges arise from significant view changes and scene scale. Previous efforts mainly focused on image or video generation, lacking exploration into the adaptability of scene generation for arbitrary views. Existing 3D generation works either operate at the object level or are difficult to utilize the geometry obtained from satellite imagery. To overcome these limitations, we propose a novel architecture for direct 3D scene generation by introducing diffusion models into 3D sparse representations and combining them with neural rendering techniques. Specifically, our approach generates texture colors at the point level for a given geometry using a 3D diffusion model first, which is then transformed into a scene representation in a feed-forward manner. The representation can be utilized to render arbitrary views which would excel in both single-frame quality and inter-frame consistency. Experiments in two city-scale datasets show that our model demonstrates proficiency in generating photorealistic street-view image sequences and cross-view urban scenes from satellite imagery.

1. Introduction

The direct generation of scenes from satellite imagery sparks exciting possibilities for seamlessly integrating diverse real-life environments into games, films, and map services. However, this task remains challenging due to the significant view differences between satellite images and street-view images on the ground, coupled with the expansive scale of the scenes under consideration. Prior efforts primarily focused on image or video generation [19, 23, 31, 34], with relatively limited exploration into scene generation. The salient advantage of scenes resides in their adaptability for rendering images from arbitrary views, which is a feature particularly beneficial for the aforementioned practical applications. Conventional cross-view synthesis works [23, 31, 34] tend to generate street-view images from the ge-



Figure 1. **Urban scenes generated by Sat2Scene.** From a single satellite image covering urban streets, Sat2Scene is able to generate videos with photorealistic and consistent textures across different views. Please use **Adobe Reader / KDE Okular** to see **animations**. Use the scroll to drag frames while the mouse is over the video. Leave the current page and back to replay.

ometry estimated based on the satellite height map. Their limitations arise from the individual generation through 2D GANs, leading to the resulting images holding a noticeable lack of inter-frame consistency. Sat2Vid [19] further advances the generation of street-view images with multi-view consistency by explicitly building the 2D-3D correspondence map between pixels in frames and point clouds. Nevertheless, the 3D point cloud is constructed by back-projecting pixels in all 2D frames, restricting its ability to synthesize video along predefined trajectories rather than arbitrary views. Furthermore, memory usage significantly increases with growing trajectory length, imposing constraints on the permissible video length.

In recent years, notable advancements [22, 28, 39] in scene representations have been observed, particularly in the realm of novel view synthesis since the introduction of NeRF [26]. Inspired by this, InfiniCity [21] relies on the pseudo ground truth generated by the 2D network for per-scene training through the optimization of the radiance field. However, it exhibits limitations in flexibility and is contin-

gent on the performance of the pre-trained 2D network.

On the other hand, diffusion models have brought generative models to a new era due to their better performance and generative ability than GANs [10]. In the realm of 3D generation, the majority of these models were introduced for generating 3D objects [27, 41] based on dense voxels. However, large-scale scenes occupy an excessive number of voxels, demanding lots of memory and restricting the scalability of these models for outdoor scene generation.

To address the aforementioned issues, we propose a novel architecture that harnesses the potential generative ability of the diffusion model coupled with the photorealistic rendering effect provided by neural rendering. This architecture directly generates the scene representation tightly associated with the geometry at the point level and produces consistent images from any view by assembling the features in 3D space through a physical-based renderer. Specifically, we accomplish this by implementing a diffusion model based on 3D sparse representations to derive point-wise color through direct denoising on the noised point clouds. We subsequently extract point-anchored features from the initially colorized point cloud and feed them to the neural rendering network with point-wise feature aggregation. This allows us to overcome the resolution limitations of voxel-based representations while ensuring efficient memory usage for handling large-scale outdoor scenes, and to generate images from arbitrary views in the scene while simultaneously maintaining consistency. The proposed architecture surpasses state-of-the-art generation models on the HoliCity [45] dataset in terms of the overall quality of the rendered video and inter-frame consistency. Additionally, the trained model demonstrates superior generalization to another city-scale dataset OmniCity [18] further containing satellite images. Fig. 1 presents rendered videos of our method inferred on two scenes from the satellite images, which demonstrates both photorealism and robust consistency across views.

The **contributions** of this paper are three-fold:

- We present a novel diffusion-based framework Sat2Scene for direct 3D scene generation, which is able to generate 3D urban scenes just from satellite images.
- To ensure consistent image generation from any view, a novel diffusion model with sparse representations is proposed to generate scene features tightly associated with the geometry directly in 3D space. To the best of our knowledge, we are the first to combine diffusion models with 3D sparse representations.
- Our model demonstrates the capability to produce photorealistic street-view image sequences with robust temporal consistency. Superior performance compared to state-of-the-art baselines validates this ability. The model exhibits proficiency when employed for cross-view urban scene generation from satellite images.

2. Related work

Our work is at the cross-section of 3D generative models and neural rendering. In the following, we discuss related works according to their underlying categories.

Diffusion models have gradually become the dominant generative models in recent years since DDPM [16]. These score-based generative models typically hold more stable training, do not require discriminators, and have been proven to achieve better quality than GANs [10]. Advanced techniques such as DDIM [35] and its follow-up [29] accelerate the sampling process while maintaining generation quality or better time scheduling during training. Latent diffusion models (LDMs) [32] further increase the denoising stability by operating in a compressed latent space to reduce the computational cost and memory footprint, making the synthesis of high-resolution images feasible on consumer devices. While image quality is constantly improving, how to utilize diffusion models to ensure multi-view consistency for multiple images generated from the same scene is currently yet less studied, which is crucial for scene or video generation. The concurrent work of this paper, MVDiffusion [36], enables holistic multi-view image generation building by leveraging LDMs and ControlNet [43], however, its cross-frame consistency remains suboptimal.

3D diffusion models, once predominantly concentrated on geometry generation [11, 24], have evolved to encompass texture generation. This expansion achieves better multi-view consistency particularly when coupled with neural differentiable rendering techniques. DiffRF [27] synthesizes a radiance field based on diffusion probabilistic models performed in 3D space. The direct denoising process on the 3D feature space of the radiance field ensures the consistency of the rendered results from different views. Point-UV Diffusion [41] employs a two-level diffusion structure to produce coarse-to-fine texture for mesh. The initial coarse texture is produced on the 3D dense voxel grids of the mesh and is subsequently refined by a secondary diffusion model in 2D space. However, due to the memory-intensive nature of the dense volumetric representation, scalability to large scenes is constrained. NF-LDM [17] solves the issue by compressing the dense scene grid into three disentangled compact representations using an auto-encoder, and training separate diffusion models for each representation. The sampled compact representations are then fed into the decoder section of the auto-encoder to reconstruct the original scene grid, which is subsequently employed for rendering arbitrary views. However, it might not be useful for generating content on a specific geometry, which requires further post-optimization procedures.

3D-aware generative models that do not rely on diffusion models have been extensively researched with various scene representations. PixelNeRF [40] extends NeRF [26] in a feed-forward manner without test-time optimization.

The sampling points in the rendering ray aggregate features from given views, which are used for color and density estimation. However, due to the less accurate point correspondence between views, the rendered results often exhibit blurriness. Alternatively, some methods integrate GAN settings with neural rendering techniques. They commonly employ intermediate explicit representations associated with simplified space structures or sparse grids to enhance memory efficiency and scalability for scenes of varying sizes. For example, GSN [9] generates a 2D floor plan as the intermediate explicit representation for unconditional indoor living scene generation. Similarly, EG3D [3] adopts a tri-plane representation, contributing to both the efficiency and effectiveness of the generation process. CompNVS [20] tries to solve the 3D inpainting or completion task via conditional generation based on NSVF [22] representation using a sparse grid. GANcraft [13] is designed for large-scale scene generation given sparse geometry with semantic information. The approach leverages a pre-trained SPADE model [30] to generate pseudo ground-truth images, which are then used to perform a NeRF-like optimization with a style controller. Notably, GANcraft’s performance is heavily dependent on the image quality of the 2D generator and necessitates per-scene training, which is non-feed-forward.

Satellite-to-ground synthesis is a very typical application of cross-view generation, which focuses on synthesizing images from a distinctly different perspective of a provided top-view satellite image. S2G [23] is the first work in this area, considering geometry consistency in satellite-to-ground synthesis. The approach predicts a density voxel grid from the satellite height map and transforms it to ground depth and semantics panorama, followed by a 2D generative model to obtain the street view. Its follow-up Sat2Vid [19] further addresses the temporal inconsistency challenge when generalizing S2G to video synthesis, via maintaining a correspondence map between 3D points and 2D frame pixels. The generator performs directly in 3D space, facilitating the natural maintenance of multi-view consistency in the rendered frames, showing better plausibility than 2D video-to-video translation advances Vid2Vid [38] and WC-Vid2Vid [25]. Nevertheless, the point cloud is pre-defined following the rendering poses hence the size of the point cloud would explode with the length of the video sequence increasing. Recently, InfiniCity [21] implements the satellite-to-ground generation on a city-wide scale. However, since it adopts GANcraft [13] as the generative model, the method still requires per-scene training and test-time optimization. Sat2Density [31] places emphasis on top-view density prediction without depth supervision. It achieves this by leveraging the view relationship between satellite and ground, and integrating neural rendering techniques to enhance synthesis quality. But due to this, the resulting videos lack good temporal consistency.

3. Method

Our method disentangles the entire scene into two components: the foreground, comprising buildings and roads, and the background sky. We represent the foreground as a point cloud, denoted as (\mathbf{P}, \mathbf{F}) , where $\mathbf{P} \in \mathbb{R}^{N \times 3}$ is the world coordinates of points in 3D space, $\mathbf{F} \in \mathbb{R}^{N \times D}$ represents per-point-associated feature map with D channels, and N is the number of points. The background is modeled as a dome at an infinite distance and represented by a 2D panorama image $\mathbf{B} \in \mathbb{R}^{H_B \times W_B \times 3}$ with height H_B and width W_B . Our pipeline is specifically designed to generate these representations for synthesizing texture information of a given input geometry \mathbf{P} . As illustrated in the Fig. 2, the proposed pipeline mainly comprises three steps. The generation phase initiates color scenes by colorizing the foreground point clouds with **a diffusion model in a 3D sparse setting** and synthesizing the background skies using a 2D diffusion model (Sec. 3.1). Then the foreground features are extracted from the initial color scene with a 3D encoder in a feed-forward manner. Finally, the neural renderer produces the color frames for the given poses (Sec. 3.2).

3.1. Generation

The generation phase creates initial textures, producing a per-point color map $\mathbf{C} \in [0, 1]^{N \times 3}$ on the normalized RGB space for the fixed geometry \mathbf{P} . We have generalized diffusion model techniques [16, 29, 35] to the 3D space with a sparse setting. Specifically, the 3D denoising network \mathcal{D} utilizes 3D sparse convolutional layers, which operate exclusively on the occupied areas (building facades and road surfaces). The loss function of \mathcal{D} is defined as

$$L_F = \mathbb{E}_{\mathbf{C}, \epsilon, t} [\|(\epsilon - \mathcal{D}(\mathbf{P}, \mathbf{C}_t, t)) \otimes \mathbf{m}\|_F^2], \quad (1)$$

where t is uniformly sampled from $\{1, \dots, T\}$ with T the number of denoising steps, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, \mathbf{C}_t denotes the noisy version of the color map \mathbf{C} at time step t , $\mathbf{m} \in [0, 1]^N$ is a mask indicating the point confidence with \otimes the broadcasted Hadamard multiplication. The background generation follows the SOTA diffusion model [32], denoting \mathcal{D}_B , resulting in a panoramic sky image $\mathbf{B} \in [0, 1]^{H_B \times W_B \times 3}$.

3.2. Rendering

The feature extraction step serves as a bridge between the generation and rendering stages by extracting 3D scene features for neural rendering. Instead of fitting a radiance field during test-time optimization, the scene representation (\mathbf{P}, \mathbf{F}) is obtained in a feed-forward way. A 3D point cloud encoder \mathcal{E} is employed to extract the features $\mathbf{F} = \mathcal{E}(\mathbf{P}, \mathbf{C})$. The features are individually anchored to the points \mathbf{P} , representing the local scene appearance and geometry information. The rendering stage utilizes the physically-based volume rendering techniques [26, 39] to synthesize frames




Figure 2. **Pipeline overview of our method.** Three steps compose the full pipeline to generate the scene representation and render street views based on satellite-inferred geometries. The **generation step** initiates colors for the foreground point cloud by using a 3D diffusion model with sparse convolutions, as well as synthesizing the background panorama with a 2D diffusion model. The scene features tightly anchored with the point cloud are extracted at the **feature extraction step**. The final **rendering step** produces images from arbitrary views through neural rendering.

with multi-view consistency. Suppose we have sampled M points along a single rendering ray with origin \mathbf{o} and direction \mathbf{d} . \mathbf{x}_i denotes the 3D location of the i -th point with $i \in \{1, \dots, M\}$ from near to far. At each location \mathbf{x}_i , we retrieve K nearest neighboring neural points around the point within a certain radius R and obtain the set of indices $\mathbf{K}_i = \text{KNN}(\mathbf{x}_i, \mathbf{P}, K, R)$. The feature \mathbf{f}_i for the location \mathbf{x}_i is aggregated by a linear combination of the features anchored to the points with indices in \mathbf{K}_i , *i.e.*, $\mathbf{f}_i = \frac{\mathbf{F}^\top \mathbf{w}_i}{\|\mathbf{w}_i\|_1}$, where each element of the weight vector \mathbf{w}_i is defined as $w_{ik} = \frac{1}{\|\mathbf{x}_i - \mathbf{p}_k\|_2}, \forall k \in \mathbf{K}_i$ and $w_{ik} = 0, \forall k \notin \mathbf{K}_i$. A network \mathcal{M} is used to infer the color and the density of each sampling point $(c_i, \sigma_i) = \mathcal{M}(\mathbf{f}_i, \mathbf{d})$. The background color \mathbf{c}_0 of the rendering ray is obtained by querying the 2D panoramic map \mathbf{B} with bilinear interpolation using the latitude and longitude coordinates computed from the ray direction \mathbf{d} , *i.e.*, $\mathbf{c}_0 = \text{Interp}(\mathbf{B}, \mathbf{d})$. The final predicted color $\hat{\mathbf{c}}$ and depth \hat{z} of the pixel are blended along the ray using volume rendering as

$$\begin{bmatrix} \hat{\mathbf{c}} \\ \hat{z} \end{bmatrix} = \sum_{i=1}^m \tau_{i-1} (1 - \exp(-\sigma_i \delta_i)) \begin{bmatrix} \mathbf{c}_i \\ z_i \end{bmatrix} + \tau_m \begin{bmatrix} \mathbf{c}_0 \\ 0 \end{bmatrix}, \quad (2)$$

where $\tau_0 = 1$, $\tau_i = \prod_{j=1}^i \exp(-\sigma_j \delta_j)$ denotes the accumulated transmittance along the ray, δ_i represents the distance between adjacent sampled points, and z_i is the Z coordinate of the sampling point in the camera coordinate system. The loss function L_d of the ray with direction \mathbf{d} is defined as

$$L_d = \|\hat{\mathbf{c}} - \mathbf{c}\|_1 + \beta |\hat{z} - z|, \quad (3)$$

where β is a coefficient controlling the importance of depth signal during training, \mathbf{c} and z are the ground-truth color and depth value of the pixel, respectively.

3.3. Implementation details

Point resampling. We notice that the training of the 3D sparse diffusion model is notably affected by the structure of the point cloud, particularly the balance in point density. The denoising process may fail when the point cloud lacks a balanced density in space, as observed in cases where the point cloud is derived from a pinhole view. This observation emphasizes the importance of maintaining a well-distributed point cloud for effective training and denoising performance. Therefore, the point cloud input to our pipeline is evenly sampled from the given geometry, ensuring a balanced density for improved training and generation results. The resampling scheme adopts the Poisson disk method [42] on the geometry surfaces, where each point has approximately the same distance to the neighboring points. Fig. 3 (a) and (b) illustrate an example of the point cloud before and after resampling. (a) is fused from multiple pinhole views thus nearby points are dense and distant points are sparse, and (b) is the resampling result where the point cloud is uniformly distributed. The colors of resampled points are interpolated and their confidence is computed based on their distance to the originally existing points. The confidence map \mathbf{m} is used for computing the denoising loss in Eq. (1). For experiments using different point cloud structures, please refer to Sec. 4.4.

Training. The entire pipeline is implemented and trained using PyTorch and Minkowski Engine [6]. The foreground denoising network \mathcal{D} is a time-conditional UNet [33], while the feature extractor \mathcal{E} is a vanilla UNet. Both networks adopt the provided default network architectures with 3D sparse voxel convolutions. In the 3D diffusion model, we use $T=1000$, and for the rendering phase, the hyper-

parameters are set as follows: $D=32$, $K=12$, $R=3m$, and $\beta=0.5$. The framework runs on a single NVIDIA Tesla A100 GPU with a memory of 40GB. The generation and rendering phases are trained separately. The training of the 3D diffusion model takes approximately 5 days, while the training of the neural renderer takes around 2 days, both starting from scratch. The network voxelization size is set to 6.67cm and 5cm in the two phases respectively in order to fit the GPU memory.




Figure 3. **Visualization of the point resampling scheme.** Yellow / purple color means high / low point confidence.

4. Experiments

4.1. Configuration

Datasets. We conduct our experiments on two datasets, HoliCity [45] and OmniCity [18]. HoliCity is a city-scale 3D dataset with rich annotation, containing 6k+ real-world urban scenes in central London. Each scene has a corresponding high-resolution panorama image, a depth (distance) map, and 8 posed RGB-D pinhole views. Since the CAD model is not publicly available, we obtain the scene geometry by lifting the pixels in the panorama image using the distance map. A scene mesh is built using the pixel connectivity in the panorama image and the point cloud is sampled on the surface using the resampling scheme described in Sec. 3.3. The number of points N is around one million for each scene, with a density of 400/m². The geometry for dynamic objects in the scene such as cars and pedestrians is unavailable in the dataset, thus we used the segmentation method Mask2Former [5] with ViT-Adapter [4] backbone pre-trained on the Cityscapes dataset [7] to obtain the semantics of the panorama image and further set the confidence of points with mismatched semantics and geometry to zero. Fig. 3 (c) gives a visualization example of the confidence map of the points. OmniCity is another city-scale dataset captured in New York City, further containing satellite images associated with ground-view panoramas. However, its street-level geometry is not available, making it difficult to compare with the baselines, thus we show qualitative results for demo purposes in Sec. 4.5.

Metrics. We evaluate our model and the baselines using the combined validation and test splits of the HoliCity dataset, containing 624 scenes and around 5k single images in total. For quantitative evaluation, we adopt commonly used

FID [15] and KID [1] to measure the generated image quality by comparing the distribution similarity between sets of synthesized **single** frames and real images. We evaluate the temporal consistency across video frames via lower-level metrics PSNR, SSIM, and LPIPS [44], which are only computed in the overlapping areas between neighboring frames. Also, considering that none of the above metrics can simultaneously measure generation quality and inter-frame consistency, we further adopt FVD and KVD [37], which are the video version of FID & KID, being able to focus more on the overall quality. They use video-classification network I3D [2] as the backbone, where the spatiotemporal features are utilized and the temporal relationship can be well modeled. To overcome the insufficient number of test scenes to compute the metrics, during calculation, we pass augmented videos with each frame as the starting frame circularly, as the 8 poses of each scene in the dataset can form a loop. This enhancement ensures that the resulting number of videos is still enough to compute those two metrics. Additionally, we incorporated a user study of 32 people to evaluate the quality of generated videos, consisting of 30 multiple-choice questions. For each question, the participant was asked to select a single generated video from three options (baselines and ours) that they considered more visually plausible.

Baselines. We compare our method to three baselines.

- *MVDiffusion* [36] is a **concurrent** advancement in multi-view generative models, conditioned on depths and camera poses. The architecture is based on LDMs [32] and ControlNet [43], together with the proposed correspondence-aware attention (CAA) module designed to enhance multi-view consistency. We train the model on the HoliCity dataset by following the proposed two-stage process, *i.e.*, training the single-view depth-conditioned diffusion model backbone first and fixing it, followed by training the ControlNet part with the CAA module for multiple views. We use their default hyper-parameters during training but use a uniform text prompt, “a street with buildings”.

- *Sat2Vid* [19] was proposed to generate street panorama videos from the geometry inferred from a satellite image. To ensure consistency across the resulting frames, it creates an explicit 3D point cloud representation for the scene and maintains its correspondence with every frame pixel. The texture information is generated directly in 3D space by a GAN method in a 3D sparse setting. In our experiment, we exclude their satellite stage but mainly train their generative parts on the HoliCity dataset, using the lifted **pinhole**-view point cloud and established 3D-2D correspondence map.

- *InfiniCity* [21] is a recent advancement for the city-scale generation based on GANcraft [13]. Since their code is not publicly available but the model is trained with the HoliCity [45] dataset, we directly take the FID and KID scores from their paper for reference only.




Figure 4. **Qualitative baseline comparison on the HoliCity [45] dataset.** Our method produces higher-quality video with better temporal consistency compared with the baselines. Please use **Adobe Reader / KDE Okular** to see **animations**.

Method / Metric	FVD↓	KVD _{×100} ↓	FID↓	KID _{×100} ↓	PSNR↑	SSIM↑	LPIPS↓	User study
Sat2Vid [19]	37.06	4.03 \pm 0.05	137.84	13.76 \pm 0.10	25.25	0.741	0.252	2.92%
InfiniCity [21]	-	-	108.47	8.40 \pm 0.10	-	-	-	-
MVDiffusion [36]	22.79	2.36 \pm 0.03	50.78	4.14\pm0.07	17.56	0.593	0.259	15.62%
Ours	20.30	1.90\pm0.03	71.98	5.91 \pm 0.06	31.54	0.956	0.237	81.46%

Table 1. **Quantitative baseline comparison.** Sat2Scene outperforms the baselines in most of the metrics including the overall video quality and the temporal consistency measures.

4.2. Quantitative comparison

The quantitative evaluation results can be found in Tab. 1. Our method outperforms the baselines in most cases except FID and KID where we are second only to MVDiffusion [36]. Meanwhile, we hold better FVD and KVD scores, which are more comprehensive metrics for the overall video quality comparison. This is very well aligned with the user study, where the videos generated by our method are more recognized and liked by the users. Since MVDiffusion [36] is a 2D diffusion model built on top of powerful LDMs and ControlNet with billions of parameters and using a pre-trained model, the generated (single) images are therefore reasonably photorealistic. This is reflected in their FID and KID scores that emphasize more on the reality of iso-

lated individual frames. However, the video realism of image sequence not only depends on the quality of individual frames but also on the temporal consistency between them. Despite the proposed CAA module helping to maintain temporal consistency, it may not generate frames as consistent as rendering from a 3D representation, resulting in inferior FVD and KVD scores than our method. This indicates that our generated frames hold very well consistency, which can be also verified in the cross-frame consistency measures PSNR, SSIM, and LPIPS, where we outperform all the baselines. It is noteworthy that Sat2Vid [19]’s 3D-2D point-to-pixel correspondence is strictly maintained, the inferior score in consistency measures may be because their 2D upsampling module breaks the original consistency.

4.3. Qualitative comparison

Fig. 4 includes the visualization results for six samples generated by our method and two baselines. For MVDiffusion [36] and our method, we generate 48 frames for every scene with the interpolated poses between the 8 poses given by the dataset. Since the memory needed by Sat2Vid [19] inflates extremely with a lot of frame generation, we keep its frame number to 8 and decrease its video FPS for synchronized visualization with other 48-frame videos. Both Sat2Vid and our method show good consistency between frames since the texture features are tiled on the scene geometry. Instead, while MVDiffusion connects frame texture features via geometric correspondences and limits the similarity, it may still struggle to guarantee complete consistency between frames. This challenge could become particularly severe for 2D methods when dealing with a smaller overlap ratio in the training data. For instance, the average foreground overlap between two neighboring “key” frames is $\sim 42\%$ in HoliCity [45], as opposed to $\sim 65\%$ in ScanNet [8] used in MVDiffusion, which could be the primary reason that it presents an inferior consistency compared with the results reported in its original paper. When considering photorealism, our method and MVDiffusion generate more plausible textures for the building facades than Sat2Vid. MVDiffusion can output more diversified building textures than ours. We infer this is because the LDMs [32] embedded in MVDiffusion have a higher ability to memorize and generate texture patterns with its large number of pre-trained parameters, compared with the diffusion model used by our method. In short, ours achieves a commendable equilibrium between the photorealism and consistency of synthesized results, showcasing a balance superior to Sat2Vid and MVDiffusion.

4.4. Ablation study

We further conduct ablative experiments to validate the effectiveness of the three components utilized in our method.

- **w/o pt-rsmp**: Do not use point resampling for balancing the point density of point clouds in the generation phase as mentioned in Sec. 3.3, use the pinhole views back-projection; The 3D diffusion model is directly trained on this structure and no neural rendering in this variant as we can directly take the colors from the denoised point cloud;

Variant / Metric	FID ↓	KID $\times 100$ ↓	Dep. RMSE
w/o pt-rsmp	131.38	12.66 ± 0.12	-
w/o pt-aggr	85.58	7.79 ± 0.08	3.22
w/o dep-sup	80.40	7.22 ± 0.08	3.44
Ours	71.98	5.91± 0.06	3.07

Table 2. **Ablative evaluation of our method.** The result shows the effectiveness of each component.

- **w/o pt-aggr**: Do not use point aggregation in the rendering phase, use voxel-based interpolation from the output of the feature extractor;
- **w/o dep-sup**: Do depth supervision in training the neural renderer, loss only calculated from color difference.

We remove each operation individually from our full model training procedure and present the quantitative evaluation results in Tab. 2 and the visualization results in Fig. 5. Since the inter-frame consistency performance of each method is very similar, we only compare the quality measures FID and KID, and further calculate the RMSE value of the depth map (except w/o pt-rsmp because it always gets ground-truth depth). It can be observed that removing any of these oper-




Figure 5. **Qualitative ablation study.** We present exemplary qualitative results for various ablations of our method. The rendered images visibly contain more details and the depths are recovered better with our full method. **The second line of each example shows the depth in pseudo colors, except the bottom left ones which are GT images.**

ations results in a decline in image quality, as evidenced by both metrics and visualization. In particular, the omission of the point resampling scheme leads to a dramatic decrease in resulting image quality. This underscores the critical importance of maintaining an even density in the point cloud for the diffusion model’s denoising process to effectively generate meaningful textures. Without the use of point aggregation, the features are organized in a slightly lower resolution and thus the representations are more coarse, leading to a decline in the rendered image quality and increased blurriness. As shown in Fig. 5, the results without depth supervision consistently exhibit holes in the rendered frames, particularly for buildings near the camera. While more refined depths for patterns on the buildings, such as windows and doors, can be generated, their geometries are not well-connected. This results in the ray hitting directly into the sky in these open areas and generating holes in the images.

4.5. Generalization

In order to demonstrate the usability of our pipeline on the satellite-to-ground cross-view generation task, and highlight its generalization capabilities, we conducted additional

experiments on the OmniCity [18] dataset. Specifically, we utilize the pre-trained monocular depth estimation network DRON [12] and the instance segmentation model Mask R-CNN [14] as provided by [18]. Exemplary prediction results of the two networks are shown in the left column of Fig. 6. The scene geometry is inferred from the combined results of the two models: the top-view height map and building footprints. For each scene, we create ground-view and bird-view navigation trajectories along the roads, containing around 500 frames. We then generate videos along the trajectories using the model trained on the HoliCity dataset. We compare the generation qualitatively on selected scenes with MVDiffusion [36] and show the videos in the right columns of Fig. 6. Our method consistently produces frames with superior consistency across different positions and view angles. In contrast, MVDiffusion generates images with varied textures even for the same position under different view angles. Moreover, in the bird-view generation, the results produced by MVDiffusion are still biased towards street-view images. This is probably because the depth image used as input to the model follows the bird-view characteristics, which is out-of-distribution with the street-view training set used to train MVDiffusion, thus failing to show good generalization ability. Since our model operates directly in 3D space, the bird-view images can be still rendered even if these views are not fed to the network during the training process.

5. Conclusion

In this paper, we explored direct scene generation from satellite imagery. Our novel architecture brings diffusion models into 3D sparse representations and integrates them with neural rendering techniques. This approach involves generating texture colors at the point level using the proposed 3D diffusion model, subsequently converting it into a scene representation without test-time optimization. The resulting representation allows the rendering of arbitrary views with excellence in both single-frame quality and inter-frame consistency. Experimental results demonstrate the superiority of our model in generating photorealistic street-view videos. To the best of our knowledge, we presented the first work that utilizes the diffusion model in a 3D sparse setting for cross-view urban scene generation from satellite imagery. We regard our work as a step forward in overcoming the challenges associated with urban scene generation, offering a promising avenue for future developments in realistic and adaptable visual content synthesis.

Acknowledgement. Zuoyue Li was supported by the Swiss Data Science Center (SDSC) fellowship program. Zhaopeng Cui was affiliated with the State Key Lab of CAD & CG at Zhejiang University. Martin R. Oswald was supported by a FIFA research grant.




Figure 6. **Street-view videos generated on the OmniCity [18] dataset.** Ours shows better generalization ability, whereas MVDiffusion [36] generates bird-view videos as from ground views. Please use **Adobe Reader / KDE Okular** to see **animations**.

References

- [1] Mikolaj Binkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD gans. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. 5
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [3] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16123–16133, 2022. 3
- [4] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. In *The Eleventh International Conference on Learning Representations*, 2023. 5
- [5] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 5
- [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019. 4
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. 7
- [9] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W. Taylor, and Joshua M. Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14313, 2021. 3
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, pages 8780–8794. Curran Associates, Inc., 2021. 2
- [11] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14300–14310, 2023. 2
- [12] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 8
- [13] Zekun Hao, Arun Mallya, Serge Belongie, and Ming-Yu Liu. GANcraft: Unsupervised 3D Neural Rendering of Minecraft Worlds. In *ICCV*, 2021. 3, 5
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 8
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. 5
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851. Curran Associates, Inc., 2020. 2, 3
- [17] Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchical latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8496–8506, 2023. 2
- [18] Weijia Li, Yawen Lai, Linning Xu, Yuanbo Xiangli, Jinhua Yu, Conghui He, Gui-Song Xia, and Dahua Lin. Omnicity: Omnipotent city understanding with multi-level and multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17397–17407, 2023. 2, 5, 8
- [19] Zuoyue Li, Zhenqiang Li, Zhaopeng Cui, Rongjun Qin, Marc Pollefeys, and Martin R. Oswald. Sat2vid: Street-view panoramic video synthesis from a single satellite image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12436–12445, 2021. 1, 3, 5, 6, 7
- [20] Zuoyue Li, Tianxing Fan, Zhenqiang Li, Zhaopeng Cui, Yoichi Sato, Marc Pollefeys, and Martin R. Oswald. Compnvs: Novel view synthesis with scene completion. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part I*, pages 447–463. Springer, 2022. 3
- [21] Chieh Hubert Lin, Hsin-Ying Lee, Willi Menapace, Menglei Chai, Aliaksandr Siarohin, Ming-Hsuan Yang, and Sergey Tulyakov. Infinicity: Infinite-scale city synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22808–22818, 2023. 1, 3, 5, 6
- [22] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Thirty-four Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 1, 3
- [23] Xiaohu Lu, Zuoyue Li, Zhaopeng Cui, Martin R. Oswald, Marc Pollefeys, and Rongjun Qin. Geometry-aware satellite-to-ground image synthesis for urban areas. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [1](#), [3](#)
- [24] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)
- [25] Arun Mallya, Ting-Chun Wang, Karan Sapra, and Ming-Yu Liu. World-consistent video-to-video synthesis. In *Proceedings of the European Conference on Computer Vision*, 2020. [3](#)
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *The European Conference on Computer Vision (ECCV)*, 2020. [1](#), [2](#), [3](#)
- [27] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, and Matthias Nießner. DiffRF: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4328–4338, 2023. [2](#)
- [28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Trans. Graph. (SIGGRAPH)*, 41(4):102:1–102:15, 2022. [1](#)
- [29] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [2](#), [3](#)
- [30] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [3](#)
- [31] Ming Qian, Jincheng Xiong, Gui-Song Xia, and Nan Xue. Sat2density: Faithful density learning from satellite-ground image pairs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3683–3692, 2023. [1](#), [3](#)
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. [2](#), [3](#), [5](#), [7](#)
- [33] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241. Springer, 2015. [4](#)
- [34] Yujiao Shi, Dylan Campbell, Xin Yu, and Hongdong Li. Geometry-guided street-view panorama synthesis from satellite imagery. In *TPAMI*, 2022. [1](#)
- [35] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021. [2](#), [3](#)
- [36] Shitao Tang, Fuayng Zhang, Jiacheng Chen, Peng Wang, and Furukawa Yasutaka. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023. [2](#), [5](#), [6](#), [7](#), [8](#)
- [37] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. FVD: A new metric for video generation. In *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019*, 2019. [5](#)
- [38] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. [3](#)
- [39] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixian Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5438–5448, 2022. [1](#), [3](#)
- [40] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4578–4587, 2021. [2](#)
- [41] Xin Yu, Peng Dai, Wenbo Li, Lan Ma, Zhengzhe Liu, and Xiaojuan Qi. Texture generation on 3d meshes with point-uv diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4206–4216, 2023. [2](#)
- [42] Cem Yuksel. Sample Elimination for Generating Poisson Disk Sample Sets. *Computer Graphics Forum*, 2015. [4](#)
- [43] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3836–3847, 2023. [2](#), [5](#)
- [44] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [5](#)
- [45] Yichao Zhou, Jingwei Huang, Xili Dai, Linjie Luo, Zhili Chen, and Yi Ma. HoliCity: A city-scale data platform for learning holistic 3D structures. *arXiv*, 2020. arXiv:2008.03286 [cs.CV]. [2](#), [5](#), [6](#), [7](#)