

InfiniCube: Unbounded and Controllable Dynamic 3D Driving Scene Generation with World-Guided Video Models

Yifan Lu^{1,2*} Xuanchi Ren^{1,3,4*} Jiawei Yang⁵ Tianchang Shen^{1,3,4} Zhangjie Wu¹
Jun Gao^{1,3,4} Yue Wang^{1,5} Siheng Chen² Mike Chen¹ Sanja Fidler^{1,3,4} Jiahui Huang¹

¹NVIDIA, ²Shanghai Jiao Tong University, ³University of Toronto, ⁴Vector Institute, ⁵University of Southern California

<https://research.nvidia.com/labs/toronto-ai/infinicube/>

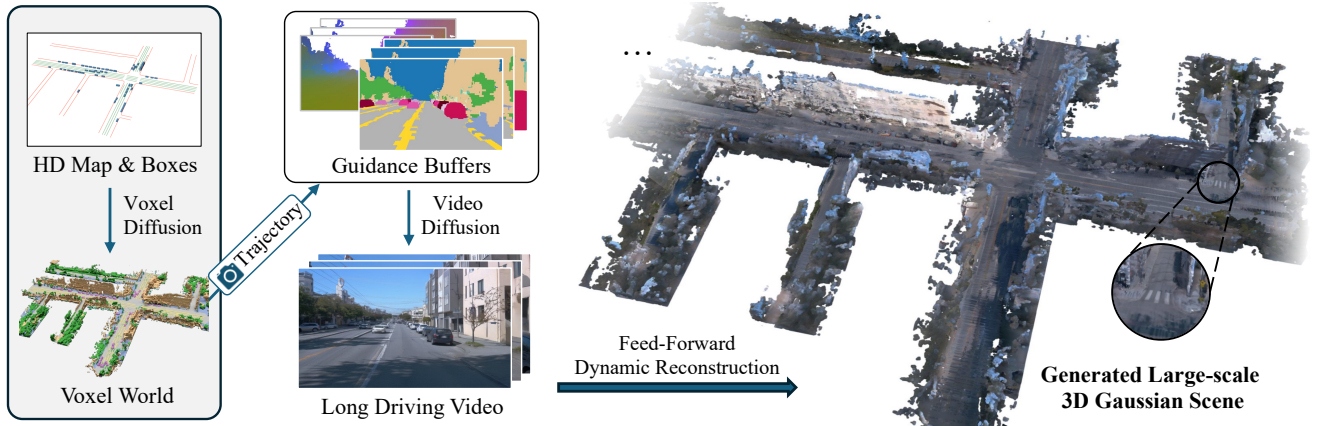


Figure 1. **InfiniCube**. Our model generates large-scale dynamic 3D driving scenes given the HD map, 3D bounding boxes and text controls. Here we show a generated large-scale driving scene spanning $\sim 100\,000\text{ m}^2$ in 3D Gaussians.

Abstract

We present *InfiniCube*, a scalable method for generating unbounded dynamic 3D driving scenes with high fidelity and controllability. Previous methods for scene generation either suffer from limited scales or lack geometric and appearance consistency along generated sequences. In contrast, we leverage the recent advancements in scalable 3D representation and video models to achieve large dynamic scene generation that allows flexible controls through HD maps, vehicle bounding boxes, and text descriptions. First, we construct a map-conditioned sparse-voxel-based 3D generative model to unleash its power for unbounded voxel world generation. Then, we re-purpose a video model and ground it on the voxel world through a set of carefully designed pixel-aligned guidance buffers, synthesizing a consistent appearance. Finally, we propose a fast feed-forward approach that employs both voxel and pixel branches to lift the dynamic videos to dynamic 3D Gaussians with control-

* Equal Contribution.

lable objects. Our method can generate controllable and realistic 3D driving scenes, and extensive experiments validate the effectiveness and superiority of our model design.

1. Introduction

Generating simulatable and controllable 3D scenes is an essential task for a wide spectrum of applications, including mixed reality, robotics, and the training and testing of autonomous vehicles (AV) [25, 33]. In particular, the requirements of AV applications have introduced new challenges for 3D generative models in driving scenarios, posing the following key desiderata: (1) *fidelity and consistency*, to ensure that the generated scenes support photo-realistic rendering while preserving consistent appearance and geometry for reliable and stable physics simulation; (2) *large-scale*, to generate scenes at map-level for traffic simulation; and (3) *controllability*, to allow easy manipulation of the scene layout, appearance, and ego-car behaviors for curating adversarial scenarios.

The investigation into large-scale 3D driving scene gen-

eration that satisfies the above criteria has been an active research area. One type of approaches focuses on the direct learning of 3D priors. After encoding the 3D scene structure into either sparse voxels [42] or Birds-Eye-View (BEV) maps [29, 34], powerful generative models (*e.g.* diffusion models [23]) are employed to model these encodings. While these approaches can generate valid 3D structures, they often fail to capture detailed appearance information. This limitation is partially due to the inductive biases of the 3D backbone networks, but more significantly, it stems from the lack of high-quality and diverse 3D training data containing detailed appearance features.

Alternatively, the recent development of video generative models [4, 6] has shown promising results in generating rich and high-fidelity visual details by pretraining on massive video datasets. Finetuning these models on driving datasets with additional conditions, such as High-Definition (HD) maps or bounding boxes, has demonstrated the capability of generating realistic driving scenes, which can be eventually used as a “world model” to facilitate planning [18, 24]. However, there are two main challenges in this type of methods: First, the generated videos often lack 3D consistency and may even violate physics, leading to disruptive appearance changes along the trajectory. Second, the auto-regressive errors accumulate over time during the generation process, making it hard to simulate long rollouts. Both limitations can be largely attributed to the lack of 3D *grounding* in the generation loop to provide guidance for synthesizing long and 3D-consistent driving scenarios.

In this paper, we identify the key challenges related to the above two families of methods, and present **InfiniCube** (Fig. 1), a scalable method that takes an HD map, vehicle bounding boxes, and a text description as input, and generates *unbounded and controllable dynamic 3D driving scenes* with high fidelity. The key innovation is to build and utilize a semantic 3D voxel world representation as guidance to a video generation model. Specifically, we first train a sparse voxel generation model that can condition on an HD map, and build a pipeline to outpaint the voxels in an unbounded fashion. The generated voxel world is subsequently rendered into a set of guidance buffers to assist in long video generation for appearance synthesis with the video model. For purposes such as location revisiting, visualization, or driving simulation, we further contribute a fast and robust method to lift the video and the voxels to dynamic 3D Gaussians (3DGS) [28] scenes while preserving the controllability of dynamic vehicles.

We provide a schematic comparison of our method with similar solutions in Tab. 1, InfiniCube not only enjoys long video generation with high-quality but also scales to large 3D dynamic scenes up to $\sim 100\,000\text{ m}^2$ ¹, providing support for various applications based upon the generated 3D scene.

¹Around $300\text{ m} \times 400\text{ m}$.

	Output Type		Detailed Geometry	Driving Length [†]
	Video	3D Rep.		
Vista [18]	✓	✗	✗	15 s
MagicDrive3D [16]	✓	3DGS	✓	6 s
InfiniCity [34]	✗	Voxels	✗	N/A
WoVoGen [38]	✓	Voxels	✓	0.4 s
InfiniCube (Ours)	✓	Voxels + 3DGS	✓	20 s [‡]

[†]: Measuring the video model generation length at 10Hz.

[‡]: Our method allows unlimited free driving in our 3DGS scene.

Table 1. **High-level comparison with existing solutions.** Our method outputs both video and a renderable 3D representation (‘3D Rep.’) with detailed geometry and a longer driving length.

2. Related Work

3D Generation. Training 3D generative models for object-level shapes has witnessed significant progress in recent years, with either directly learning the 3D shape distribution [15, 61, 74], or first generating multi-view images and then lifting them to 3D [17, 47, 62]. Extending these techniques to larger-scale scenes is however non-trivial due to its high complexity. Some works directly learn the 3D scene distribution [29, 37, 42, 71, 77] with carefully curated ground-truth data, while some others take the combinatorial [14, 32] or hierarchical [34, 60] approach of generating or retrieving objects and then arranging them using guided scene layouts. Notably, a recently emerged trend is to reconstruct scenes generated by video generative models [69, 70] thanks to its rich appearance and flexibility.

Controllable Video Generation. The prevalence of video generation models fueled by either diffusion models [4, 5, 8] or autoregressive models [30, 41, 63] has brought up the need for more fine-grained control over the generated content, including camera trajectories [3, 31], object motions [56], and scene structures [21, 35]. In the more specific domain of driving video generation, several works have conditioned the video generation on the HD maps and the car bounding boxes [18, 24, 55, 57, 67], and can generate a local occupancy map alongside [19, 38], turning them into a “world model” that facilitate planning applications. Comparably, our method enables consistent large-scale scene generation deeply grounded in 3D, unlocking a longer simulation capability.

Driving Scene Reconstruction. Driving scene reconstruction plays a critical role in creating realistic simulations for autonomous vehicle training and testing. Existing methods, such as those based on neural radiance fields (NeRFs) [20, 50, 59, 65] or 3DGS [10, 13, 64, 76], achieve impressive visual fidelity but suffer from long training times, limiting their application at scale. A concurrent trend is the development of large reconstruction models that leverage data priors for improved generalizability and fast inference [2, 43, 72]; however, these works almost exclusively

focus on scene reconstruction. In contrast, our InfiniCube also tackles large dynamic driving scene *generation* with high fidelity and object-level controllability, unmatched by all previous approaches, to the best of our knowledge.

3. Preliminaries

Our method is based heavily on the following key concepts:

Latent Diffusion Models (LDM). Latent diffusion models [44] are a class of generative models that learn the distribution from the latent \mathbf{X} of the data \mathbf{D} . It is used together with an auto-encoder (or a tokenizer) that maps the data into the latent space $\mathbf{X} = \mathcal{E}(\mathbf{D})$. A diffusion model [23] starts with a noise vector $\mathcal{N}(0, \mathbf{I})$ and iteratively denoises it to generate a sample \mathbf{X} , potentially guided by a condition \mathbf{C} . The decoder is eventually applied to the denoised latent to generate the output data $\hat{\mathbf{D}} = \mathcal{D}(\mathbf{X})$. LDMs have been shown to be effective and efficient in modeling many data modalities, including images [44], videos [4], and 3D [74].

Sparse Voxel LDM. As the 3D equivalent of pixels, voxels are uniform and structured and are hence suitable in deep learning applications. In practice, only voxels intersecting actual surfaces need to be stored, forming a *sparse* voxel grid. In InfiniCube, we consider the sparse voxel grid that stores a *semantic label* in each of its voxels, represented as \mathbf{D}^{vx} . The work of XCube [42] that we base on provides an efficient way of encoding sparse voxels \mathbf{D}^{vx} into a dense latent feature cube $\mathbf{X}^{\text{vx}} = \mathcal{E}^{\text{vx}}(\mathbf{D}^{\text{vx}}) \in \mathbb{R}^{N^3 \times C}$ (where N is the edge length), and decoding it back with high fidelity. Crucially, XCube [42] *trains an LDM of sparse voxels*. Readers are encouraged to refer to the original paper for more details.

4. Method

InfiniCube aims to generate large-scale dynamic 3D scenes guided by the input HD maps, vehicle bounding boxes and text prompts. As shown in Fig. 2, we first generate a large-scale semantic voxel world of the target scene (§ 4.1) based on given conditions. Such a world representation is then used to render several *guidance buffers* at the given vehicle trajectories to support long-range video generation (§ 4.2). Finally, we take both the voxels and the synthesized videos to reconstruct a dynamic scene with the 3DGS representation (§ 4.3). We will detail each of the components in the following subsections.

4.1. Unbounded Voxel World Generation

This step takes the HD map and the vehicle bounding boxes as input and synthesizes a corresponding 3D voxel world² with semantic labels. We opt to re-use the sparse voxel LDM from XCube [42] (Sec. 3) for this task.

²The term ‘voxel world’ simply refers to a large-scale sparse voxel grid.

Building Map Conditions. From our input, we first build a condition volume $\mathbf{C}^{\text{vx}} \in \mathbb{R}^{N^3 \times S}$ that shares the same structure as \mathbf{X}^{vx} . It contains the following three components. (i) **HD Map Condition $\mathbf{C}_{\text{HD}}^{\text{vx}}$:** Our input HD Map contains two sets of 3D polylines, road *edges* defining the road boundary and road *lines* that separate the lanes. We rasterize the polylines into two separate channels of the condition $\mathbf{C}_{\text{HD}} \in \mathbb{R}^{N^3 \times 2}$, with the voxel value set to 1 if any part of the voxel intersects with the polyline, and 0 otherwise. (ii) **Road Surface Condition $\mathbf{C}_{\text{Road}}^{\text{vx}}$:** We empirically find it hard for the model to generate accurate drivable road regions from $\mathbf{C}_{\text{HD}}^{\text{vx}}$ alone due to its value sparsity. We hence add another condition $\mathbf{C}_{\text{Road}}^{\text{vx}} \in \mathbb{R}^{N^3 \times 1}$ that delineates the voxelized road *surface* as an extra signal to inform the model of the drivable area. To identify the road surface, we fit a 3D plane to the provided road edges and remove its non-road parts, which are determined by finding all connected components (via their BEV projections) that do not contain the road lines. This procedure is applied in a block-wise manner to handle possible variations in road elevation. (iii) **Bounding Box Condition $\mathbf{C}_{\text{Box}}^{\text{vx}}$:** Bounding boxes contain detailed information about the vehicles’ poses. However, naively voxelizing bounding box occupancies can lead to information lost due to the coarse latent voxel size (e.g. 1.6 m). We hence encode the vehicles’ heading angles α as two-channel vectors $[\sin \alpha, \cos \alpha]$, and set the feature in $\mathbf{C}_{\text{Box}}^{\text{vx}} \in \mathbb{R}^{N^3 \times 2}$ to this encoding if more than half of the corresponding voxel is occupied by the bounding box. The full condition volume is a concatenation of the three $\mathbf{C}^{\text{vx}} = \{\mathbf{C}_{\text{HD}}^{\text{vx}}, \mathbf{C}_{\text{Road}}^{\text{vx}}, \mathbf{C}_{\text{Box}}^{\text{vx}}\}$ with $S = 4$, as visualized in Fig. 3.

Scene Chunk Generation. With the above \mathbf{C}^{vx} , we apply the diffusion sampling procedure in [42] to generate our desired semantic voxel grid representation \mathbf{D}^{vx} . In a nutshell, the sampling process starts with a Gaussian white noise $\mathcal{N}(0, \mathbf{I})$ in shape $\mathbb{R}^{N^3 \times C}$ which is subsequently concatenated with the conditions \mathbf{C}^{vx} in the channel dimension. A 3D U-Net will then iteratively denoise the white noise to obtain the latent \mathbf{X}^{vx} , which could be decoded to the desired semantic sparse voxel grid. More details about the sampling procedure can be found in the Supplement.

However, a single pass of the LDM sampling can only generate one chunk of the scene. We hence propose the following strategy to *outpaint* the grid to a large voxel world.

Unbounded Scene Outpainting. Here we use a strategy similar to Repaint [39], a training-free outpainting technique that is commonly used in diffusion models, to iteratively extend the scene in a seamless manner. Specifically, during the generation of a new chunk, we ensure its sufficient overlap with the existing part of the scene, and take the latent $\mathbf{X}_{\text{exist}}^{\text{vx}}$ from the overlapping area. During the diffusion process, we keep $\mathbf{X}_{\text{exist}}^{\text{vx}}$ fixed and only update the current la-

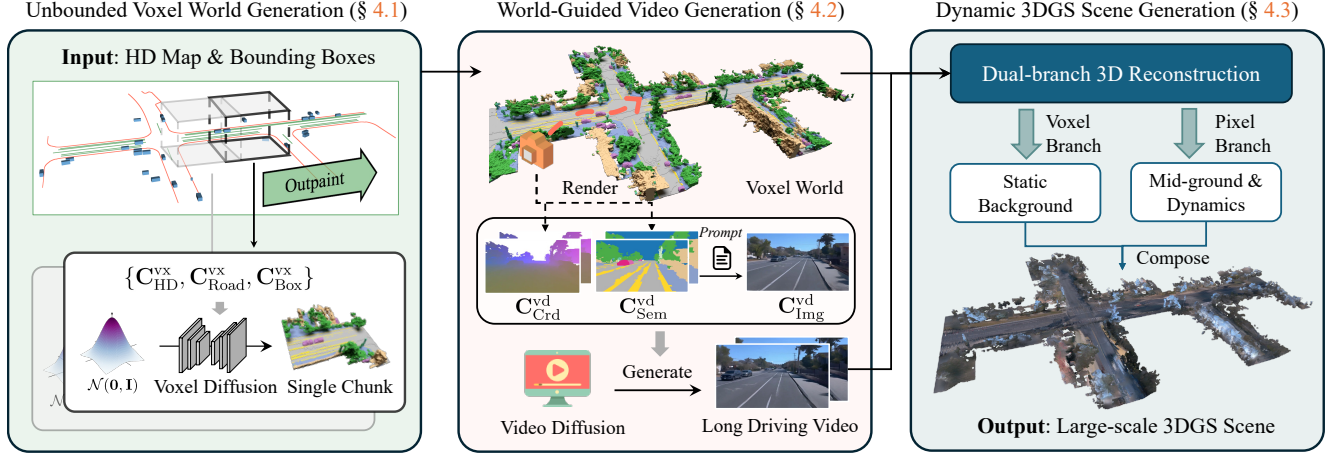


Figure 2. **Pipeline.** Conditioned on HD maps and bounding boxes, we first generate a 3D voxel world representation. We then render the voxel world into several guidance buffers to boost video generation. The generated video and voxel world are jointly fed into a feed-forward dynamic reconstruction module to obtain the final 3DGS representation.

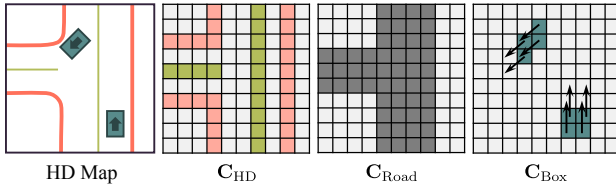


Figure 3. **Conditions for Voxel World Generation.** We illustrate the conditions in 2D for clarity but the actual conditions are in 3D. We use the latent $\mathbf{X}_{\text{new}}^{\text{vx}}$ for the newly generated part, as follows:

$$\mathbf{X}_{\text{new}}^{\text{vx}} \leftarrow (1 - \mathbf{M}) \odot \hat{\mathbf{X}}_{\text{new}}^{\text{vx}} + \mathbf{M} \odot \hat{\mathbf{X}}_{\text{exist}}^{\text{vx}} \quad (1)$$

where \mathbf{M} is the overlapping mask and \odot is the element-wise product. While $\hat{\mathbf{X}}_{\text{new}}$ is sampled from the learned diffusion posterior, $\hat{\mathbf{X}}_{\text{exist}}$ is sampled from the noised version of the fixed $\mathbf{X}_{\text{exist}}$. With this, we can generate a large-scale scene that is consistent across different chunks.

4.2. World-Guided Video Generation

Our video generation model is based on Stable Video Diffusion (SVD) [4], a state-of-the-art LDM for video generation. Under the hood, the video \mathbf{D}^{vd} is represented as a volume of shape $\mathbb{R}^{H \times W \times T \times 3}$, where H, W, T are the height, width, and number of frames, respectively. The latent of the video modeled by the LDM is $\mathbf{X}^{\text{vd}} = \mathcal{E}^{\text{vd}}(\mathbf{D}^{\text{vd}}) \in \mathbb{R}^{h \times w \times T \times 4}$, downsampling the spatial resolution with $h = \frac{H}{8}$ and $w = \frac{W}{8}$. Similarly to the voxel generation model, the condition to the model $\mathbf{C}^{\text{vd}} \in \mathbb{R}^{h \times w \times T \times M}$ shares the same size as \mathbf{X}^{vd} except for the last channel dimension M . The official model of SVD is trained on large-scale Internet videos, and supports conditioning on the first video frame with $\mathbf{C}_{\text{Img}}^{\text{vd}} \in \mathbb{R}^{h \times w \times T \times 4}$, derived by repeating the frame T times and then encoding it with the SVD encoder $\mathcal{E}^{\text{vd}}(\cdot)$. One can then auto-regressively generate videos with lengths longer than T by conditioning on the last generated frame.

However, generating long and consistent driving videos is challenging since the model has to maintain an internal *implicit* ‘memory’ of the surrounding world. We hence propose to use the generated semantic voxel world from § 4.1 to assist the video model with an explicit condition of the world geometry and the camera trajectory, implemented as video-space *guidance buffers*.

Guidance Buffers. Our guidance buffers are renderings of the 3D world into the video frames, using the desired camera parameters. They are composed of the following components. (i) The **Semantic Buffer** $\mathbf{C}_{\text{Sem}}^{\text{vd}}$ is the rendering of the voxel world’s semantic labels. We use a manually-designed fixed discrete color palette to map the semantic labels to RGB values. The color palette is chosen to add distinction between different semantic categories and fits the value range of the pretrained encoder $\mathcal{E}^{\text{vd}}(\cdot)$. To differentiate between the vehicle instances that belong to the same semantic category, we assign different saturation levels to their rendered colors. This enables the video model to maintain the distinct appearance of individual cars when there are multiple cars overlapping each other in the buffer. (ii) The **Coordinate Buffer** $\mathbf{C}_{\text{Crd}}^{\text{vd}}$ contains the 3D scene coordinates of the first voxel that each pixel ray hits (*cf.* [52, 54]). For a T -frame video model, we accumulate³ all the 3D points in world space from these frames and normalize them to the range of $[-1, 1]$ to be accepted by $\mathcal{E}^{\text{vd}}(\cdot)$. For the same locations in the 3D scene across different frames, although they may be projected to different *pixel* coordinates due to ego and object motion, their pixel values remain consistent. This helps establish scene correspondences across frames and is useful when the semantic buffer exhibits a repeating pattern. A visualization of the guidance buffers can be found in Fig. 2. Together the channel size of the video

³Moving vehicles are transformed according to their bounding boxes.



Figure 4. **Illustration of the pixel branch supervision.** The mid-ground area and (masked) voxel depth.

model condition $C^{vd} = \{C_{\text{Img}}^{vd}, C_{\text{Sem}}^{vd}, C_{\text{Crd}}^{vd}\}$ is $M = 12$.

Our guidance buffers can be built very efficiently by rendering the voxel world. Compared to other conditioning strategies such as pose embedding [16] or displacements [18], our method is agnostic to the metric scale of the trajectory while achieving precise controls.

Adding Text Prompts. Our video model necessitates the first frame as a condition. To achieve this, we train a ControlNet [73] based on FLUX [1] to generate the initial frame using semantic buffers as control images. This approach allows us to incorporate text descriptions for scene generation while enhancing the quality and diversity of the video content by leveraging a pre-trained image diffusion model.

4.3. Dynamic 3DGS Scene Generation

While the above world-guided video model can already provide a photo-realistic and consistent appearance, in some applications (*e.g.*, real-time visualization and simulation) where a 3D structure is still needed, we present a novel feed-forward method that is deeply integrated into our pipeline to reconstruct a dynamic 3DGS [28] scene.

State-of-the-art feed-forward scene reconstruction models that leverage the 3DGS representation infer Gaussian attributes either in the *voxel space* [43] or in the *pixel space* [72]. While the former typically outputs a better 3D geometry around the camera, the latter could better capture contents in the *mid-ground* areas and the per-frame movement for the dynamic objects. Here we define the *mid-ground* as the pixel regions that (1) have no overlap with the projected voxels, and (2) do not belong to the sky (visualized in Fig. 4). We hence propose a new *dual-branch* reconstruction method that combines the above two strategies for our large-scale dynamic 3DGS scene generation.

Voxel Branch. One branch of our model takes the voxel world and the generated posed video frames as input, and outputs a set of 3D Gaussians for each voxel. Similar to the appearance reconstruction branch in SCube [43], we unproject the features of the images to the voxel world and then apply a 3D sparse convolution U-Net architecture to transform the features into the per-voxel Gaussian attributes. We mask the dynamic objects out from the image features and only use the static background voxels in this branch.

Pixel Branch. Our pixel branch employs a 2D UNet [45] backbone to convert input images into per-pixel 3D Gaus-

sians, similar to the representation in GS-LRM [72]. The pixel-aligned formulation simplifies the 3D-lifting task to a depth estimation problem. Here, we further propose to fully utilize the rendered voxel depth Z to enhance the depth prediction capability and generalizability in this branch. Specifically, We incorporate Z as both an input and a supervisory signal: During training, we use a randomly masked version of Z (denoted as \hat{Z}) to simulate the region that is not grounded by voxels, and supervise the predicted depth with the full Z — see the illustration in Fig. 4. This helps our pixel branch predict reasonable mid-ground depth at inference time. We also supplement the network with the ViT backbone features F_{DAV2} from a state-of-the-art depth estimation model Depth Anything V2 [66]. In summary, the network takes the input image, \hat{Z} , and F_{DAV2} as input, and outputs the per-pixel 3DGS attributes (color, rotation, depth, scale, *etc.*).

Sky Modeling. Modeling the sky region at infinite depth is challenging since most of it is not visible from the images. We hence adopt an implicit sky representation from STORM [2] that is highly generalizable to unseen regions. Specifically, we use a light-weight encoder to summarize a single sky feature vector $c \in \mathbb{R}^{192}$ from the images, and use AdaGN [27] to modulate a Multi-Layer Perceptron (MLP) that takes a viewing angle and outputs RGB colors. More network details can be found in the Supplement.

Supervision. We train two branches separately using photometric loss. For the pixel branch, we additionally add the aforementioned depth loss.

Inference with Dynamic Objects. During inference, the voxel branch is applied only to the static part of the scene (including static vehicles). The pixel branch is applied iteratively for every frame, but we only keep the 3DGS corresponding to the pixels of mid-ground regions and dynamic objects. Notably, for the dynamic vehicles, we extract the 3DGS belonging to each *individual object* using the segmentation from the Semantic Buffer in § 4.2, transform and aggregate them using the vehicle trajectory to composite an *amodal* geometry. The motions of the objects are fully controllable by simply altering their trajectories.

5. Experiments

5.1. Data Processing

Our model is trained on Waymo Open Dataset [49], which provides LiDAR data, images, and accurate annotations of HD maps and vehicle bounding boxes. To extract the ground-truth scene geometry to supervise the semantic voxel generation, we follow the approach of [43] by combining accumulated LiDAR points and the dense geometry obtained from the multi-view stereo pipeline of COLMAP [46]. The geometry of the dynamic cars is accumulated from the LiDAR points in their canonical space

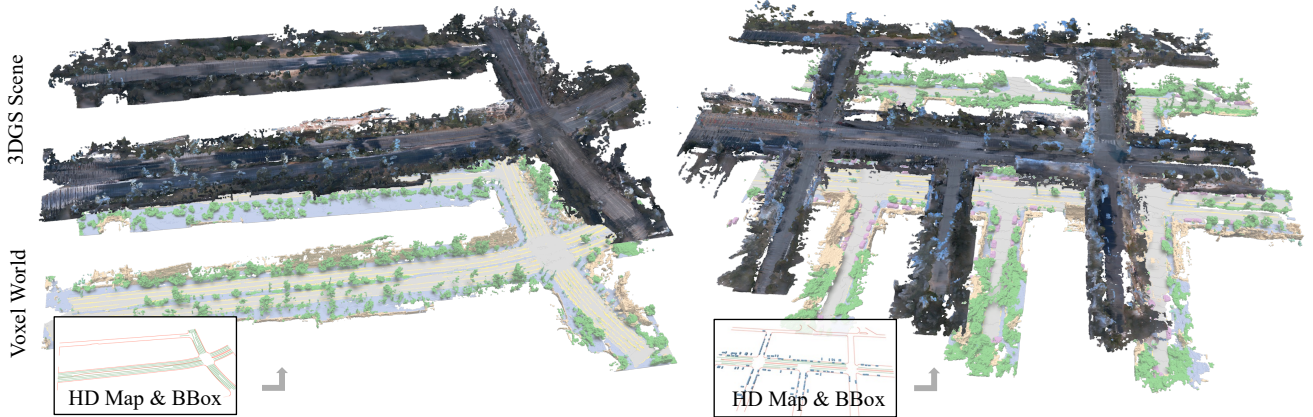


Figure 5. **Full pipeline results.** The generated 3DGS scene and voxel world adhere to the input HD map and bounding boxes and extend over hundreds of meters with fine details. The text prompt is "Daytime, Sunny with a bright blue sky".

defined by their bounding box trajectories. To enable text prompt conditioning, we annotate the video frames with textual descriptions. We employ Llama-3.2-90B-Vision-Instruct [12] to summarize the weather and time of day based on a stitched thumbnail created from 8 timestamps across 3 different views.

For each data sample used in training the voxel generation stage, we crop and voxelize the extracted geometry into a local chunk of $51.2\text{ m} \times 51.2\text{ m}$ with a voxel size of 0.2 m , centered around a randomly sampled ego-vehicle pose. We remove low-quality voxel grids and sequences with mostly static ego trajectories, resulting in 618 sequences for training and 90 sequences for evaluation.

5.2. Implementation Details

For both the voxel world (§ 4.1) and the 3DGS scene (§ 4.3) generation stages where a 3D sparse network backbone is needed, we take heavy use of the *fVDB* [58] framework, and design our 3D auto-encoder and diffusion backbone similar to [42, 43]. Please refer to the Supplement for network architecture details. For the video generation stage (§ 4.2), our SVD [4] backbone is implemented with the *diffusers* [51] library. We finetune the pretrained model with a resolution of 576×1024 together with our designed conditions. We add Gaussian noise augmentation to perturb the conditioning latent features to improve the model robustness. During inference, we set the classifier-free guidance [22] weight to 3.0 and use a denoising step of 25. The voxel generation stage is trained for 48 GPU days, the video generation stage is trained for 192 GPU days, and both the voxel and pixel branches of the scene generation stage are trained for 32 GPU days, all using NVIDIA A100 GPUs.

5.3. Large-scale Dynamic Scene Generation

We visualize the generated scenes from our full pipeline in Fig. 5 and Fig. 1. Furthermore, Fig. 6 shows a close-up view of the dynamic objects. For some parts of the



Figure 6. **Dynamic 3DGS visualization.** We generate dynamic 3DGS with full controllability of dynamic objects across frames.

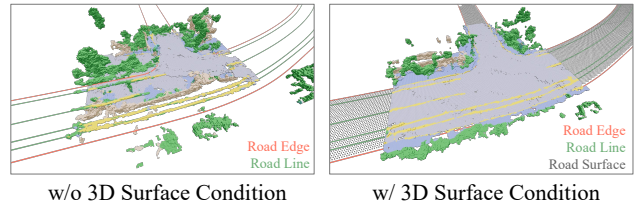


Figure 7. **Impact of 3D Road Surface Condition C_{Road} .** Without C_{Road} , the model fails to generate the ground accurately.

scenes without any coverage of the ego-car trajectories, we use a customized strategy to generate the trajectories for the guidance buffers, whose details are deferred to the Supplement. Given just the HD map with 3D bounding boxes, our method can generate a complete scene with a high-fidelity appearance and controllable dynamic actors. The scene is also rich in details and accurate in geometry, allowing for large-scale bird-eye visualizations that no prior work could generate. These results are made possible by the synergy among the proposed components (§§ 4.1 to 4.3). In the following sections, we will analyze their importance in detail.

5.4. Main Components Analysis

5.4.1 Voxel World Generation

We perform an ablation study for this component to verify the design of our HD map conditions. Specifically, in Fig. 7 we compare our generation result with the one trained without the *Road Surface* condition $C_{\text{Road}}^{\text{vx}}$. One can clearly see the benefits of the condition by helping the network better disambiguate and localize the actual drivable regions.

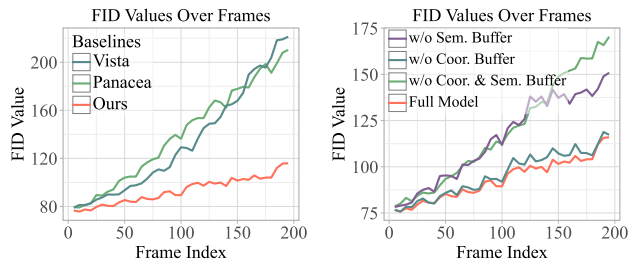


Figure 8. **Video Model Comparison.** Our model can generate high-quality videos of 200 frames without severe degradation by conditioning on guidance buffers. These buffers also enable the motions in the videos to be consistent with the scale of the physical world.

5.4.2 World-Guided Video Generation

To evaluate the performance of our world-guided video generation, we compare it with two baselines: Panacea [57] and Vista [18]. Both of them are specifically designed for driving video generation (in other words, ‘world models’). The original Panacea model is trained based on Stable Diffusion 1.5 — for a fair comparison, we re-implement their method with the Stable Video Diffusion backbone same as ours, while adopting their map conditioning strategies. Since Vista does not support analytic ego trajectory input, we only use the first frame as its condition without specifying any additional controls. We use the ground-truth first frames of the 90 test sequences from the Waymo Open Dataset and apply the video generation models on top of them. We compute the Fréchet Inception Distance (FID) over the generated video frames at different frame indexes to measure the video quality over time.

Shown in Fig. 9a and visualized in Fig. 8, our model maintains a lower FID score and a better visual quality over a long time horizon. This is in contrast to our baselines where we observe significant quality degradation after ~ 100 frames. The result highlights the effectiveness of our world-guided video generation strategy and the guidance buffer design, being able to reduce the auto-regressive



(a) Baseline Comparison on FID. (b) Ablation of Guidance Buffers.

Figure 9. **Comparison of long video generation quality** based on the first frame from the Waymo Dataset. FID: lower is better.

	InfiniCube (Ours)			Panacea [57]		
Frame Index	40	80	120	40	80	120
Positive Rate \uparrow	84.6%	83.9%	84.8%	76.8%	54.0%	53.4%

Table 2. **Human evaluation of HD map alignment.** We highlight the **best** and the **second**.

errors typically persistent in long video generation tasks. In Fig. 9b, we show a detailed ablation study on the effects of different guidance buffers: While the semantic buffer plays the most crucial role in maintaining video quality, the coordinate buffer helps to resolve detailed ambiguities of motion-induced scene changes.

	Novel View ($T + 5$)			Novel View ($T + 10$)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
PixelNeRF [68]	15.21	0.52	0.64	14.61	0.49	0.66
PixelSplat [7]	20.11	0.70	0.60	18.77	0.66	0.62
DUS $\text{\textcircled{3}}\text{R}$ [53]	17.08	0.62	0.56	16.08	0.58	0.60
MVSplat [9]	20.14	0.71	0.48	18.78	0.69	0.52
MVSGaussian [36]	16.49	0.70	0.60	16.42	0.60	0.59
SCube [43]	19.90	0.72	0.47	18.78	0.70	0.49
InfiniCube (Ours)	20.80	0.73	0.42	19.93	0.72	0.45

Table 3. **Quantitative comparisons of novel view rendering.** Metrics are computed at frames $T + 5$ and $T + 10$ given frame T as input. We highlight the **best**, **second best** and **third best**.

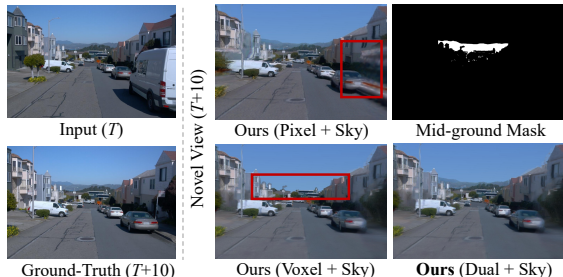


Figure 10. **Novel view rendering** of different branches given input images and voxels from Waymo Dataset. Dual inference eliminates the artifacts in either single branch, as shown in the **red box**.

To assess the *controllability* of our model, we conducted a user study to evaluate the alignment of our generation w.r.t. the HD map condition. We extract the 40th, 80th, and 120th frames from the video generated by our method and Panacea (sharing the same first-frame generated by ControlNet), and put the map projection beside to ask users if the image aligns with the map. A total of 180 image samples are extracted for each frame index, and we report the positive response rate comparison in Tab. 2. Our method is more preferential than Panacea in all cases, especially for bigger frame indices, further demonstrating the effectiveness of our world-guided model design.

5.4.3 3DGS Scene Generation

Our 3DGS reconstruction stage works in synergy with the previous two stages to generate high-quality dynamic 3DGS scenes. To showcase the advantage of our dual-branch reconstruction, we follow the setting from SCube [43] to evaluate the reconstruction quality by synthesizing novel views at frame $T + 5$ and $T + 10$ given the input views from frame T with 3 front views. We show the Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [75] results in Tab. 3. Our method outperforms the baselines in all metrics; by introducing the pixel branch, we gain improvement over SCube [43]. We further show in Fig. 10 a qualitative analysis of the renderings coming from different branches of our model. While each branch has

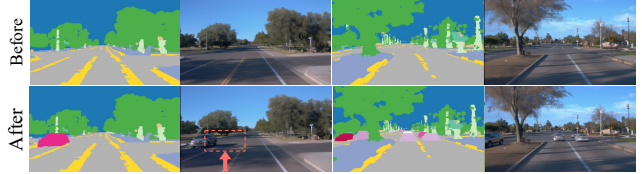


Figure 11. **Object insertion by the video model.** We observe realistic shadows cast by the inserted objects (indicated by **arrow**).

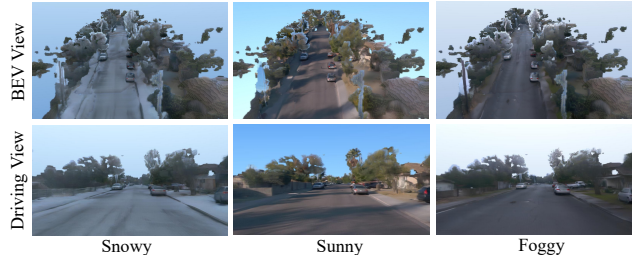


Figure 12. **Weather control for scene generation.** We show 3 scenes generated with different text prompts (listed at the bottom).

its own artifacts, our dual-branch inference can effectively eliminate them and generate high-quality novel views.

5.5. Applications

With a 3DGS scene and a corresponding voxel world, our method naturally supports applications such as novel view synthesis or collision simulation. Meanwhile, the coherent design of our pipeline also enables more advanced applications as shown below.

Vehicle Insertion. New vehicles (assuming known trajectories) can be inserted into the scene by simply placing a voxelized car CAD model into the 3D voxel world and re-running the subsequent steps. To maintain the video appearance of the scene before and after the insertion, we keep the first-frame condition unchanged and only update the guidance buffer. Two insertion examples are shown in Fig. 11.

Weather Control. Our method allows users to generate scenes with different weather conditions by just altering the text prompts. We show 3 scenes with different weather conditions sharing the same underlying voxel world in Fig. 12, where the 3D Gaussians have different appearances.

6. Discussion

Limitations. While our method could generate 3D driving scenes with rich *appearance* details thanks to the power of the world-guided video model, the diversity of the *geometry* is still bounded by the voxel generation stage where ground-truth 3D training data is hard to acquire. The full pipeline also consists of multiple stages involving multiple time-consuming diffusion sampling steps and it is hard to recover from the potential failure of intermediate stages.

Conclusion. In this work, we present InfiniCube, a novel

method for generating large-scale and high-quality dynamic 3D driving scenes. Our method is deeply rooted in the synergy among our voxel world generation model, the world-guided video model, and the dynamic 3DGS generation model. Together we can generate realistic 3D scenes with rich appearance details and full controllability. Future works include scaling up the training data with more diverse driving scenarios and speeding up the entire generation process with end-to-end models.

References

- [1] Announcing black forest labs. <https://blackforestlabs.ai/announcing-black-forest-labs/>. Accessed: 2024-08-01. 5, 14
- [2] Anonymous. STORM: Spatio-temporal reconstruction model for large-scale outdoor scenes. In *Submitted to The Thirteenth International Conference on Learning Representations*, 2024. under review. 2, 5
- [3] Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, et al. Vd3d: Taming large video diffusion transformers for 3d camera control. *arXiv preprint arXiv:2407.12781*, 2024. 2
- [4] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 2, 3, 4, 6
- [5] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023. 2
- [6] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. 2024. URL <https://openai.com/research/video-generation-models-as-world-simulators>, 3, 2024. 2
- [7] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. Pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, 2024. 8
- [8] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7310–7320, 2024. 2
- [9] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *ECCV*, 2024. 8
- [10] Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo de Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, et al. Omnire: Omni urban scene reconstruction. *arXiv preprint arXiv:2408.16760*, 2024. 2
- [11] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 15
- [12] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 6
- [13] Tobias Fischer, Jonas Kulhanek, Samuel Rota Bulò, Lorenzo Porzi, Marc Pollefeys, and Peter Kotschieder. Dynamic 3d gaussian fields for urban areas. *arXiv preprint arXiv:2406.03175*, 2024. 2
- [14] Gege Gao, Weiyang Liu, Anpei Chen, Andreas Geiger, and Bernhard Schölkopf. Graphdreamer: Compositional 3d scene synthesis from scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21295–21304, 2024. 2
- [15] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems*, 35:31841–31854, 2022. 2
- [16] Ruiyuan Gao, Kai Chen, Zhihao Li, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrive3d: Controllable 3d generation for any-view rendering in street scenes. *arXiv preprint arXiv:2405.14475*, 2024. 2, 5
- [17] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul Srinivasan, Jonathan T Barron, and Ben Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv preprint arXiv:2405.10314*, 2024. 2
- [18] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. Vista: A generalizable driving world model with high fidelity and versatile controllability. *arXiv preprint arXiv:2405.17398*, 2024. 2, 5, 7
- [19] Songen Gu, Wei Yin, Bu Jin, Xiaoyang Guo, Junming Wang, Haodong Li, Qian Zhang, and Xiaoxiao Long. Dome: Taming diffusion model into high-fidelity controllable occupancy world model. *arXiv preprint arXiv:2410.10429*, 2024. 2
- [20] Jianfei Guo, Nianchen Deng, Xinyang Li, Yeqi Bai, Botian Shi, Chiyu Wang, Chenjing Ding, Dongliang Wang, and Yikang Li. Streetsurf: Extending multi-view implicit surface reconstruction to street views. *arXiv preprint arXiv:2306.04988*, 2023. 2
- [21] Yuwei Guo, Ceyuan Yang, Anyi Rao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Sparsectrl: Adding sparse controls to text-to-video diffusion models. In *European Conference on Computer Vision*, pages 330–348. Springer, 2025. 2
- [22] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022. 6, 13

- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3
- [24] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023. 2
- [25] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhui Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. 1
- [26] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. 14
- [27] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 5, 15
- [28] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2, 5
- [29] 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*, pages 8496–8506, 2023. 2
- [30] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vignesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023. 2
- [31] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *arXiv preprint arXiv:2405.17414*, 2024. 2
- [32] Haoran Li, Haolin Shi, Wenli Zhang, Wenjun Wu, Yong Liao, Lin Wang, Lik-hang Lee, and Pengyuan Zhou. Dreamscene: 3d gaussian-based text-to-3d scene generation via formation pattern sampling. *arXiv preprint arXiv:2404.03575*, 2024. 2
- [33] Zhenxin Li, Kailin Li, Shihao Wang, Shiyi Lan, Zhiding Yu, Yishen Ji, Zhiqi Li, Ziyue Zhu, Jan Kautz, Zuxuan Wu, et al. Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation. *arXiv preprint arXiv:2406.06978*, 2024. 1
- [34] 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*, pages 22808–22818, 2023. 2
- [35] Fangfu Liu, Wenqiang Sun, Hanyang Wang, Yikai Wang, Haowen Sun, Junliang Ye, Jun Zhang, and Yueqi Duan. Reconx: Reconstruct any scene from sparse views with video diffusion model. *arXiv preprint arXiv:2408.16767*, 2024. 2
- [36] Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao, Wei Li, and Ziwei Liu. Mvs gaussian: Fast generalizable gaussian splatting reconstruction from multi-view stereo. *arXiv preprint arXiv:2405.12218*, 2, 2024. 8
- [37] Yuheng Liu, Xinke Li, Xueting Li, Lu Qi, Chongshou Li, and Ming-Hsuan Yang. Pyramid diffusion for fine 3d large scene generation. *arXiv preprint arXiv:2311.12085*, 2023. 2
- [38] Jiachen Lu, Ze Huang, Zeyu Yang, Jiahui Zhang, and Li Zhang. Wovogen: World volume-aware diffusion for controllable multi-camera driving scene generation. In *European Conference on Computer Vision*, pages 329–345. Springer, 2025. 2
- [39] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022. 3
- [40] Nerfstudio Project. Viser: A library for interactive 3d visualization in python. <https://github.com/nerfstudio-project/viser>, 2024. 16
- [41] Xuanchi Ren and Xiaolong Wang. Look outside the room: Synthesizing a consistent long-term 3d scene video from a single image. In *CVPR*, 2022. 2
- [42] Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, Ken Museth, Sanja Fidler, and Francis Williams. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4209–4219, 2024. 2, 3, 6, 13
- [43] Xuanchi Ren, Yifan Lu, Hanxue Liang, Zhangjie Wu, Huan Ling, Mike Chen, Sanja Fidler, Francis Williams, and Jiahui Huang. Scube: Instant large-scale scene reconstruction using voxplats. *arXiv preprint arXiv:2410.20030*, 2024. 2, 5, 6, 8, 14, 15
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [45] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 5, 15
- [46] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 5
- [47] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 2
- [48] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 13
- [49] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 5, 15

- [50] Adam Tonderski, Carl Lindström, Georg Hess, William Ljungbergh, Lennart Svensson, and Christoffer Petersson. Neurad: Neural rendering for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14895–14904, 2024. 2
- [51] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models, 2024. 6, 14
- [52] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019. 4
- [53] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. *arXiv preprint arXiv:2312.14132*, 2023. 8
- [54] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 4
- [55] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024. 2
- [56] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2
- [57] Yuqing Wen, Yucheng Zhao, Yingfei Liu, Fan Jia, Yanhui Wang, Chong Luo, Chi Zhang, Tiancai Wang, Xiaoyan Sun, and Xiangyu Zhang. Panacea: Panoramic and controllable video generation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6902–6912, 2024. 2, 7
- [58] Francis Williams, Jiahui Huang, Jonathan Swartz, Gergely Klar, Vijay Thakkar, Matthew Cong, Xuanchi Ren, Ruilong Li, Clement Fuji-Tsang, Sanja Fidler, et al. fvd: A deep-learning framework for sparse, large scale, and high performance spatial intelligence. *ACM Transactions on Graphics (TOG)*, 43(4):1–15, 2024. 6
- [59] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuantao Chen, Runyi Yang, et al. Mars: An instance-aware, modular and realistic simulator for autonomous driving. In *CAAI International Conference on Artificial Intelligence*, pages 3–15. Springer, 2023. 2
- [60] Haozhe Xie, Zhaoxi Chen, Fangzhou Hong, and Ziwei Liu. Citydreamer: Compositional generative model of unbounded 3d cities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9666–9675, 2024. 2
- [61] Kevin Xie, Jonathan Lorraine, Tianshi Cao, Jun Gao, James Lucas, Antonio Torralba, Sanja Fidler, and Xiaohui Zeng. Latte3d: Large-scale amortized text-to-enhanced3d synthesis. *arXiv preprint arXiv:2403.15385*, 2024. 2
- [62] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arXiv preprint arXiv:2403.14621*, 2024. 2
- [63] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2
- [64] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339*, 2024. 2
- [65] Jiawei Yang, Boris Ivanovic, Or Litany, Xinshuo Weng, Seung Wook Kim, Boyi Li, Tong Che, Danfei Xu, Sanja Fidler, Marco Pavone, et al. Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. *arXiv preprint arXiv:2311.02077*, 2023. 2
- [66] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 5, 15
- [67] Xuemeng Yang, Licheng Wen, Yukai Ma, Jianbiao Mei, Xin Li, Tiantian Wei, Wenjie Lei, Daocheng Fu, Pinlong Cai, Min Dou, et al. Drivearena: A closed-loop generative simulation platform for autonomous driving. *arXiv preprint arXiv:2408.00415*, 2024. 2
- [68] 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*, pages 4578–4587, 2021. 8
- [69] Heng Yu, Chaoyang Wang, Peiye Zhuang, Willi Menapace, Aliaksandr Siarohin, Junli Cao, Laszlo A Jeni, Sergey Tulyakov, and Hsin-Ying Lee. 4real: Towards photorealistic 4d scene generation via video diffusion models. *arXiv preprint arXiv:2406.07472*, 2024. 2
- [70] Haiyu Zhang, Xinyuan Chen, Yaohui Wang, Xihui Liu, Yunhong Wang, and Yu Qiao. 4diffusion: Multi-view video diffusion model for 4d generation. *arXiv preprint arXiv:2405.20674*, 2024. 2
- [71] Junge Zhang, Qihang Zhang, Li Zhang, Ramana Rao Kompella, Gaowen Liu, and Bolei Zhou. Urban scene diffusion through semantic occupancy map. *arXiv preprint arXiv:2403.11697*, 2024. 2
- [72] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-irm: Large reconstruction model for 3d gaussian splatting. In *European Conference on Computer Vision*, pages 1–19. Springer, 2025. 2, 5, 15
- [73] 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*, pages 3836–3847, 2023. 5, 14
- [74] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu.

- Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024. [2](#), [3](#)
- [75] 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. [8](#)
- [76] Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21634–21643, 2024. [2](#)
- [77] Vlas Zyrianov, Henry Che, Zhijian Liu, and Shenlong Wang. Lidardm: Generative lidar simulation in a generated world. *arXiv preprint arXiv:2404.02903*, 2024. [2](#)