

DiST-4D: Disentangled Spatiotemporal Diffusion with Metric Depth for 4D Driving Scene Generation

Jiazhe Guo^{1*} Yikang Ding^{2*†} Xiwu Chen^{3‡} Shuo Chen¹ Bohan Li⁴ Yingshuang Zou¹
 Xiaoyang Lyu⁵ Feiyang Tan³ Xiaojuan Qi⁵ Zhiheng Li¹ Hao Zhao^{1‡}

¹Tsinghua University ²MEGVII ³Mach Drive ⁴Shanghai Jiao Tong University ⁵University of Hong Kong

Abstract

Current generative models struggle to synthesize dynamic 4D driving scenes that simultaneously support temporal extrapolation and spatial novel view synthesis (NVS) without per-scene optimization. A key challenge lies in finding an efficient and generalizable geometric representation that seamlessly connects temporal and spatial synthesis. To address this, we propose DiST-4D, the first disentangled spatiotemporal diffusion framework for 4D driving scene generation, which leverages metric depth as the core geometric representation. DiST-4D decomposes the problem into two diffusion processes: DiST-T, which predicts future metric depth and multi-view RGB sequences directly from past observations, and DiST-S, which enables spatial NVS by training only on existing viewpoints while enforcing cycle consistency. This cycle consistency mechanism introduces a forward-backward rendering constraint, reducing the generalization gap between observed and unseen viewpoints. Metric depth is essential for both accurate reliable forecasting and accurate spatial NVS, as it provides a view-consistent geometric representation that generalizes well to unseen perspectives. Experiments demonstrate that DiST-4D achieves state-of-the-art performance in both temporal prediction and NVS tasks, while also delivering competitive performance in planning-related evaluations. The Project is available at <https://royalmelon0505.github.io/DiST-4D>

1. Introduction

Generative models have recently emerged as a powerful tool for autonomous driving, enabling the creation of large-scale synthetic data to train and evaluate perception, planning, and simulation systems [9, 10, 14, 22, 37, 38, 42]. These methods alleviate the limitations of costly real-world data collection, providing diverse and scalable training scenarios crucial for improving model robustness and generalization. However, existing generative approaches for autonomous driving remain limited in their ability to synthesize realistic,

Method	Temporal Gen.	Spatial NVS	Feed-forward
Vista [11]	✓	✗	✓
UniScene [22]	✓	✗	✓
Drive-WM [38]	✓	✗	✓
DriveDreamer [37]	✓	✗	✓
MagicDrive [10]	✓	✗	✓
MagicDriveDiT [9]	✓	✗	✓
3DGS [18]	✗	✓	✗
PVG [3]	✗	✓	✗
UniSim [49]	✗	✓	✗
EmerNeRF [45]	✗	✓	✗
StreetGaussian [43]	✗	✓	✗
DriveDreamer4D [54]	✗	✓	✗
Stag-1 [32]	✗	✓	✓
FreeVS [34]	✗	✓	✓
STORM [46]	✗	✓	✓
DriveForward [31]	✗	✓	✓
StreetCrafter [44]	✗	✓	✓
DreamDrive [26]	✓	✓	✗
MagicDrive3D [8]	✓	✓	✗
DiST-4D (Ours)	✓	✓	✓

Table 1. **Comparison of different methods.** DiST-4D is the first framework to achieve feed-forward dynamic 4D driving scene generation with both temporal extrapolation and spatial novel view synthesis. This is made possible by its disentangled spatiotemporal diffusion design with metric depth.

controllable open-world environments. A key yet unresolved challenge lies in generating dynamic 4D scenes (3D geometry + temporal evolution) that can support data generation at any novel time and location.

To better understand the limitations of existing approaches, we briefly review current methods for autonomous driving scene generation, as shown in Tab.1. (1) Some methods leverage diffusion models conditioned on BEV layouts or trajectories to predict future scenarios, generating 2D videos [9–11, 22, 25, 26, 37, 42], 3D occupancy maps [22, 33, 40, 55], or LiDAR data [22, 60]. However, these approaches remain tightly coupled with predefined trajectories, preventing them from achieving novel view synthesis (NVS) capabilities. (2) Meanwhile, implicit reconstruction paradigms [3, 4, 43, 45] excel in NVS, but their reliance on per-scene optimization results in high computational costs. (3) Recent feed-forward solutions [31, 32, 34, 46] employ large models to synthesize novel views from captured sensor

*Equal contribution, † Project leader, ‡ Corresponding author

data but lack temporal extrapolation, limiting their ability to predict future scene evolution. (4) Some emerging frameworks attempt to bridge temporal prediction and spatial NVS by reconstructing generated video sequences [8, 26]. However, these hybrid approaches still inherit fundamental limitations from per-scene optimization, restricting their efficiency and scalability. This highlights a critical open challenge: how to achieve simultaneous anytime future forecasting and anywhere NVS without per-scene optimization.

To tackle these challenges, we posit that the key to simultaneously achieving temporal evolution and spatial NVS lies in establishing appropriate 4D representations. Such a representation must satisfy two critical requirements: (1) It should encode geometric information sufficient for NVS at each time step, and (2) It must be learnable by generative models to support temporal prediction. Failing to meet the first requirement would necessitate per-scene reconstruction, making feed-forward generation infeasible. Meanwhile, failing to meet the second criterion would prevent future state generation. Given these considerations, along with the inherent complexity of driving scenes that challenge single-stage 4D representation learning, we propose per-frame metric depth as the representation to enable disentangled spatiotemporal 4D generation.

This choice offers three key advantages: (1) Metric depth simultaneously satisfies the two aforementioned criteria, *i.e.*, providing accurate geometric information while being easily learnable by generative models [7, 17], (2) It exhibits strong generalization for novel view synthesis, as evidenced in [34, 51], and (3) Dense metric depth itself holds significant practical value in autonomous driving applications, such as data annotation and training of perception models.

Building on these insights, we propose DiST-4D, a novel framework that leverages temporally and spatially disentangled modeling to achieve feed-forward future scene generation and spatial NVS for the first time. To enable this paradigm, we first develop a data curation pipeline for generating large-scale multi-view metric depth data. This pipeline integrates visual reconstruction [5], LiDAR point cloud fusion, and semantic segmentation [15] to create dense 3D point cloud prompts, followed by a depth completion network [24] to produce high-precision metric depth maps.

We then design DiST-4D as a dual-branch diffusion architecture, consisting of DiST-T (temporal prediction) and DiST-S (spatial synthesis). Specifically, DiST-T generates future multi-view RGB and metric depth sequences by conditioning on historical observations and control signals (*e.g.*, BEV maps and trajectories). DiST-S synthesizes dense RGB and depth outputs from novel viewpoints, using a diffusion model conditioned on sparse projections from surrounding point clouds. A key challenge in training DiST-S on real-world driving data is the limited trajectory diversity, which hinders generalization to novel out-of-distribution view-

points. To address this, we propose a novel self-supervised cycle consistency (SCC) strategy, inspired by [35, 57, 58], to enhance the robustness of spatial NVS. Specifically, we first randomly generate novel trajectories, then use trained DiST-S to generate the corresponding RGB-D sequences, which will be reprojected back to the original viewpoints to form self-supervised training pairs. This SCC strategy enables self-supervised fine-tuning, improving DiST-S’s adaptability to various trajectories.

Our contributions are summarized as follows:

- We propose DiST-4D, the first model to simultaneously achieve temporal future prediction and spatial novel view synthesis (NVS) without requiring per-scene optimization.
- We introduce metric depth as the geometric bridge between temporal forecasting and spatial NVS, enabling efficient and generalizable 4D scene generation.
- We develop a multi-view metric depth curation pipeline, to provide high-quality supervision for 4D scene generation.
- We propose a self-supervised cycle consistency (SCC) strategy to enhance the generalization ability of DiST-S, improving spatial NVS on out-of-distribution trajectories.
- Extensive experiments demonstrate that DiST-4D outperforms existing approaches in both temporal prediction and spatial novel view synthesis, while maintaining competitive performance in planning-related evaluations.

2. Related Work

2.1. Temporal Driving Video Generation

Recent progress in temporal driving video generation has advanced realistic, controllable autonomous driving simulations [9, 10, 12, 20, 21, 26, 32, 39]. MagicDrive[10] achieves high-fidelity street-view synthesis through tailored encoding and cross-view attention mechanisms for precise 3D geometry control. UniScene[22] unifies multi-modal data generation (semantic occupancy, video, LiDAR) via progressive processes, improving downstream task performance with intermediate representations. Extensions like MagicDriveDiT [9] and MagicDrive3D [8] address scalability challenges using DiT architectures and deformable Gaussian splatting with monocular depth initialization, respectively, enabling high-resolution 3D scene reconstruction. DreamDrive[26] synthesizes 4D spatial-temporal scenes via hybrid Gaussian representations, balancing visual quality and generalizability. Despite these advancements, existing methods struggle to reconcile temporal video generation with spatial novel view synthesis, limiting scalability for autonomous driving applications. To address this, we propose a disentangled spatiotemporal framework that leverages metric depth as a geometric bridge to harmonize spatial and temporal coherence.

2.2. Spatial Novel View Synthesis

Recent advances in spatial novel view synthesis (NVS) [3, 18, 27, 30–32, 34, 43–46, 49, 52, 54] have enhanced the gen-

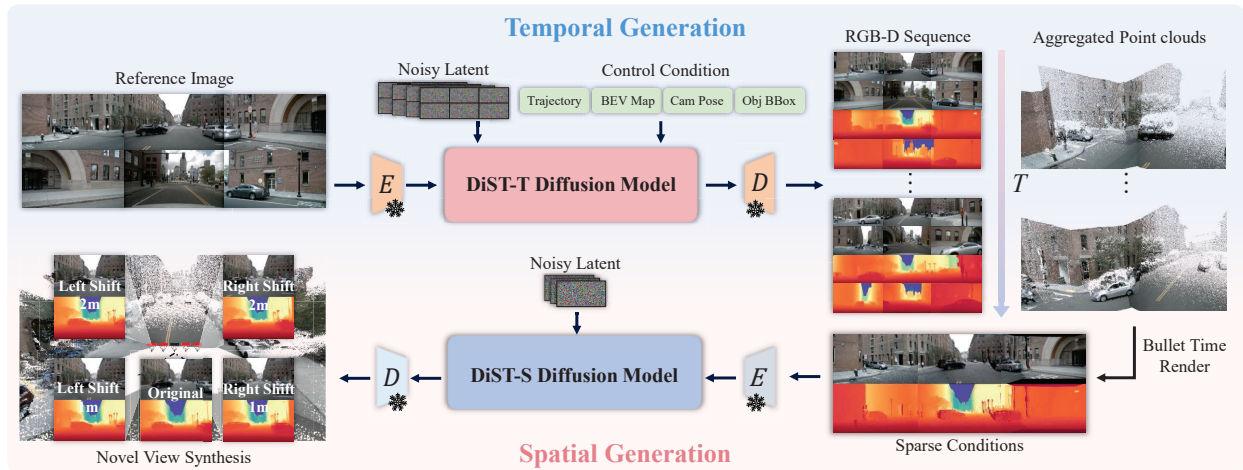


Figure 1. **Overall framework of the proposed DiST-4D.** DiST-4D is a disentangled spatiotemporal diffusion framework for 4D driving scene generation, leveraging metric depth as the core geometric representation to enable both temporal extrapolation and spatial novel view synthesis (NVS). (**Top: Temporal Generation**) DiST-T employs a diffusion model to predict future multi-camera RGB-D sequences from historical multi-camera images and control signals. The generated RGB-D sequences are then aggregated into point clouds, allowing for bullet time rendering. (**Bottom: Spatial Generation**) To enable spatial NVS, DiST-S leverages the predicted RGB-D sequences to generate novel viewpoints by first projecting them into sparse conditions and then refining them into dense RGB-D outputs.

eration of realistic, flexible views for dynamic driving scenes, a critical capability for autonomous driving simulations.

Per-scene Optimizations. Per-scene optimization methods [3, 18, 43, 45, 49, 54] refine scene representations to capture dynamic elements and structural details. For instance, 3D Gaussian Splatting [18] enables real-time radiance field rendering using Gaussian representations, avoiding computations in empty space. Periodic Vibration Gaussian (PVG)[3] extends this framework to dynamic urban scenes by integrating temporal dynamics for unsupervised reconstruction. UniSim [49] models static and dynamic actors via neural feature grids for multi-sensor safety-critical simulations, while EmerNeRF[45] decomposes scenes into static/dynamic fields using self-supervised flow. Street Gaussian [43] combines explicit point clouds and 4D spherical harmonics for efficient rendering.

Feed-forward Approaches. Feed-forward approaches [23, 31, 32, 34, 44, 46] emphasize rapid inference and scalability for real-time applications. STORM [46] leverages transformer-based architectures to infer dynamic 3D scenes from sparse inputs, achieving high-fidelity rendering. FreeVS[34] generates 3D-consistent novel views using pseudo-image priors, bypassing per-scene optimization. Stag-1 [32] decouples spatiotemporal relationships to enable multi-view 4D consistency. In this paper, we propose an innovative feed-forward 4D driving scene paradigm to efficiently achieve unified 4D dynamic driving scene generation without per-scene optimization.

3. Methodology

3.1. Overview

As illustrated in Fig. 1, DiST-4D comprises two key components: DiST-T and DiST-S, designed to decouple temporal

prediction from spatial novel view synthesis. Given a historical frame of observations (*i.e.*, multi camera images) and control signals (*i.e.*, trajectories, object BBoxes, and BEV maps), DiST-T leverages a diffusion transformer network to generate future multi camera RGB videos and depth sequences. For spatial NVS, the generated RGB and depth maps are projected onto novel views, forming sparse RGB-D conditions that guide DiST-S in producing dense RGB-D outputs. Since accurate and dense metric depth is crucial, we first introduce the pipeline for obtaining metric depth for model training (Sec. 3.2). We then detail the temporal generation model, DiST-T (Sec. 3.3), followed by the spatial novel view synthesis model, DiST-S (Sec. 3.4).

3.2. Metric Depth Curation

Acquiring metric depth in autonomous driving remains a challenging task. While learning-based models [41, 59] have made remarkable progress in inferring metric depth from RGB images, they still suffer from inherent limitations in precision and detail fidelity. This makes them inadequate for training generative models that demand high-quality supervision. To address this problem, we propose a pseudo-GT curation pipeline to generate high-quality metric depth.

Inspired by the recent generative depth completion method [24], we propose to get dense depth maps conditioned on RGB images and incomplete depth data, which can be obtained from LiDAR point clouds. While LiDAR point clouds offer precise sparse depth measurements, they fall short in reconstructing complete scene geometry, especially in capturing high buildings with upper parts being far from the ground, and distant objects located far away from the ego vehicle. To this end, we additionally combine visual reconstruction results with LiDAR point clouds. As shown in Fig.2, we first aggregate the sparse LiDAR points extracted

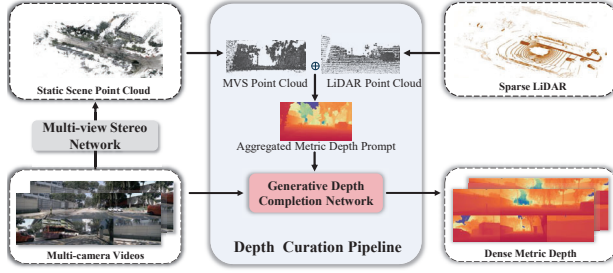


Figure 2. **Illustration of the metric depth curation pipeline.** First, a multi-view stereo network processes multi-camera videos to produce a static scene point cloud. Simultaneously, sparse LiDAR point clouds are collected and fused with the MVS output to obtain an aggregated metric depth prompt. This intermediate depth representation serves as input to a generative depth completion network, which refines and densifies the depth estimates, producing high-fidelity dense metric depth maps.

from adjacent frames, where dynamic points are removed using 3D bounding boxes. Next, we obtain static scene points using an MVS network[5], filtering out dynamic objects and the sky region via 2D semantic segmentation [15]. We then derive the initial metric depth by merging the MVS depth and LiDAR depth, followed by applying a hidden point removal algorithm [16] to eliminate occluded points. Finally, the aggregated metric depth prompt, along with the corresponding RGB images, are fed into a generative depth completion network [24] to produce the final dense depth maps. More details about metric depth curation are provided in supplemental material.

3.3. DiST-T: Temporal RGB-D Generation

Given a frame of historical RGB observations as reference images $\mathbf{I}_{\text{ref}} = \{\mathbf{I}_{\text{ref},c}\}_{c=1}^C$, where C denotes the number of camera views, DiST-T aims to predict the future scenes corresponding to \mathbf{I}_{ref} and control signals. Similar to previous works [9, 10], the control signals are defined as descriptions $\{\mathbf{S}_t\}_{t=1}^T$ of future frames. For each $\mathbf{S}_t = \{\mathbf{M}_t, \mathbf{B}_t, \mathbf{A}_t, \mathbf{P}_t\}$, $\mathbf{M}_t \in \{0, 1\}^{w \times h \times s}$ represents a $w \times h$ road region with s semantic classes. $\mathbf{B}_t = \{(c_i, \mathbf{b}_i)\}_{i=1}^N$ are 3D bounding boxes where each object is described by the corner points of the box $\mathbf{b}_i \in \mathbb{R}^{8 \times 3}$ and class $c_i \in \mathcal{C}$. \mathbf{A}_t represents ego vehicle trajectory, describing the relative transformation from the t -th frame to the first frame of the sequence, and $\mathbf{P}_t \in \mathbb{R}^{c \times 4 \times 4}$ indicates the pose of c cameras. As we treat metric depth as the key to 4D scene generation, DiST-T needs to generate future RGB videos and depth sequences. Let \mathbf{Z}_t^I and \mathbf{Z}_t^D denote the latent feature of RGB image and depth at time t , then DiST-T can be formulated as:

$$\{\mathbf{Z}_t^I, \mathbf{Z}_t^D\} = \mathcal{G}_\theta(\{\mathbf{S}_t, \mathbf{I}_{\text{ref}}\}), \quad (1)$$

where \mathcal{G}_θ indicates DiST-T network with parameters θ .




Figure 3. **Illustration of the diffusion transformer of DiST-T.** STDiT captures multi-camera and multi-frame dependencies for future scene generation. It processes latent features with temporal and spatial STDiT blocks, integrating ego trajectory (A), 3D object bounding boxes (B), camera poses (P) and map sequences (M) as control signals.

Network Architecture. Simultaneously generating RGB-D videos for multi-camera poses significant challenges to GPU memory. To encode the RGB-D videos into highly compressed latent features, we utilize a pre-trained 3D VAE [50] to process RGB and depth videos separately, reducing their size in both temporal and spatial dimensions. For the latent diffusion model, STDiT [56] is employed as the backbone of DiST-T. The down-sample and embedding strategies for B, A and M are adopted from [9]. As illustrated in Fig. 3, the control branch (highlighted in blue) is similar to ControlNet [53]. Map M and reference image \mathbf{I}_{ref} are injected into the control blocks after encoding, while the A, B, and P are provided as input to both the basic block and the control block. Within the STDiT backbone, the multi-view block captures interactions among multiple cameras, whereas the temporal and spatial blocks focus on temporal and spatial interactions within a single camera, respectively.

Training Strategy. Training a multi-modal generative model from scratch results in slow convergence. To speed up training, we first train a low-resolution (224×400) RGB video generation model, then extend it to low-resolution RGB-D generation, and finally refine it for high-resolution (424×800) RGB-D video generation with a mixed-frame-length [9, 17] strategy. Following [56] and pursuing scalability, the model is trained through rectified flow and v-prediction loss [6], more details can be found in the supplementary material.

3.4. DiST-S: Spatial RGB-D NVS

Task Formulation. The DiST-S aims at synthesizing spatial novel view RGB images \mathbf{I}_{tgt} and the corresponding depth maps \mathbf{D}_{tgt} in a feed-forward manner. Unlike FreeVS [34], which uses sparse LiDAR data, and ViewCrafter [51], which relies on relative depth projections, DiST-S utilizes metric depth projections as inputs, including projected RGB and depth. As depicted in Fig. 1, given an arbitrary future moment t , we convert the generated multi-view RGB images \mathbf{I}_{gen} and dense depth maps \mathbf{D}_{gen} into a per-frame point cloud using the known camera intrinsics and extrinsics. After that,

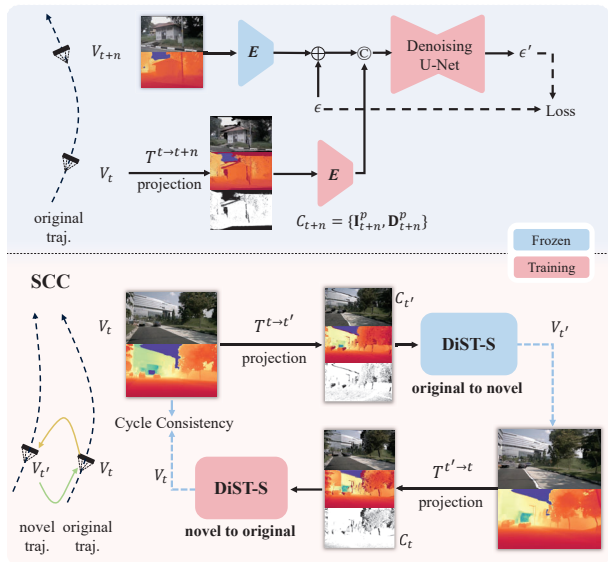


Figure 4. (Top) In the first stage, DiST-S is trained on the original trajectory using n frame projection (Bottom) In the second stage, self-supervised cycle consistency (SCC) is introduced, where novel trajectories are generated, and DiST-S learns to project between original and novel viewpoints.

we project point cloud at t into novel viewpoint to get condition results \mathbf{I}_{cond} and \mathbf{D}_{cond} . Let $\mathbf{Z}_{\text{cond}}^I$ and $\mathbf{Z}_{\text{cond}}^D$ denote the latent feature of \mathbf{I}_{cond} and \mathbf{D}_{cond} from the frozen VAE decoder, then DiST-S can be formulated as:

$$\{\mathbf{Z}_{\text{tgt}}^I, \mathbf{Z}_{\text{tgt}}^D\} = \mathcal{F}_\theta(\{\mathbf{Z}_{\text{cond}}^I, \mathbf{Z}_{\text{cond}}^D\}), \quad (2)$$

where \mathcal{F}_θ indicates DiST-S network with parameters θ .

Network Architecture. As shown in Fig. 4, DiST-S consists of a denoising UNet [1], a frozen VAE encoder, and a trainable condition encoder. The frozen encoder compresses the dense target RGB images and depth maps into latent features, while the trainable encoder is used to extract latent features of the sparse condition maps.

Since the VAE encoder is designed for 3-channel RGB input, we replicate the single-channel depth map into three channels to simulate an RGB image. During training, the RGB latent $\mathbf{Z}_{\text{tgt}}^I$ and depth latent $\mathbf{Z}_{\text{tgt}}^D$ are concatenated along the channel dimension. For the condition inputs, we combine the single-channel valid mask with the channel-wise replicated two-channel depth input to create the final depth condition \mathbf{D}_{cond} . Both \mathbf{D}_{cond} and the image condition \mathbf{I}_{cond} are then encoded into $\mathbf{Z}_{\text{cond}}^D$ and $\mathbf{Z}_{\text{cond}}^I$ through a condition encoder, which is trained from scratch. Finally, $\mathbf{Z}_{\text{cond}}^I$ and $\mathbf{Z}_{\text{cond}}^D$ are concatenated with $\mathbf{Z}_{\text{tgt}}^I$ and $\mathbf{Z}_{\text{tgt}}^D$ along the channel dimension to form the denoising inputs.

As our denoising UNet is based on the video diffusion model [1], we need to modify its original input and output channels to accommodate multi-modal generation. Follow-

ing the approach in [17], we increase the input channels from 8 to 16 by replicating the input layer’s weight tensor and scaling its values by half. Similarly, we modify the output channels from 4 to 8 using the same weight adjustment.

Training Strategy and Cycle Consistency. The training phase of DiST-S includes two stages. In the first stage, we use the training data from the original trajectories to train the denoising UNet and condition encoder. To obtain condition inputs, we use projection matrix $T^{t \rightarrow t+n}$ to project the surround point cloud at time t to time $t+n$ to serve as conditions $C_{t+n} = \{\mathbf{I}_{t+n}^p, \mathbf{D}_{t+n}^p\}$, with the n is set to ± 4 . While $n = \pm 4$ and voxel-down-sampling for surround point cloud is utilized as data augmentation. By incorporating data from different cameras, DiST-S is capable of learning camera motions under various pose changes, where the data from the forward-facing camera enables the model to learn how to move forward and backward, while the data from the side-facing camera teaches the model to move laterally.

However, relying solely on real trajectory data for NVS remains insufficient to address the full range of diverse and unseen trajectories. In the second training stage, inspired by [35, 57, 58], we introduce a Self-supervised Cycle Consistency (SCC) strategy to enhance the model’s robustness in synthesizing novel views under extreme viewpoint variations. Specifically, we first randomly generate novel trajectories $\mathbf{A}_{\text{novel}}$ with the corresponding input conditions based on the original trajectories \mathbf{A}_{ori} and surrounding point clouds. For viewpoint V_t at t frame in \mathbf{A}_{ori} , we randomly shift $\tau \in [-3m, +3m]$ laterally to get the novel viewpoint V'_t in $\mathbf{A}_{\text{novel}}$. Then, we use the trained DiST-S from the first training stage to synthesize dense RGB-D outputs in $\mathbf{A}_{\text{novel}}$, which will be reprojected to \mathbf{A}_{ori} to produce condition inputs. These condition maps together with the original dense RGB-D maps in \mathbf{A}_{ori} form the additional training pairs in the second training stage. The SCC strategy allows the DiST-S to encounter a diverse range of new trajectories during training, significantly enhancing its generalization and robustness to novel perspectives. Following [1] and pursuing stability, the training objective in both training stages is EDM-styled ϵ -prediction, more details about training can be found in the supplementary material.

4. Experiments

Our experiments are conducted on nuScenes dataset [2]. We interpolate the 2Hz keyframe annotation in a higher frame rate of 12Hz, similar to previous studies [10, 36]. We adhere to the official splits, using 700 videos for training and 150 for validation. Further details about model setup are provided in the supplementary material.

4.1. Temporal Generation

Evaluation Metric. We evaluate the quality of the generated RGB videos using the FID [28] for images and FVD for videos. The generated RGB videos have a resolution of

Method	Multi-view	Video	Depth	FID ↓	FVD ↓
BEVGen [29]	✓	✗	✗	25.54	-
BEVControl [47]	✓	✗	✗	24.85	-
DriveGAN [19]	✗	✓	✗	73.40	502.30
DriveDreamer [37]	✗	✓	✗	52.60	452.00
Vista [11]	✗	✓	✗	6.90	89.40
WoVoGen [25]	✓	✓	✗	27.60	417.70
Panacea [42]	✓	✓	✗	16.96	139.00
MagicDrive [10]	✓	✓	✗	16.20	217.94
MagicDriveDiT [9]	✓	✓	✗	20.91	94.84
Drive-WM [38]	✓	✓	✗	15.80	122.70
Vista* [11]	✓	✓	✗	13.97	112.65
UniScene [22]	✓	✓	✗	6.45	71.94
DiST-T (Ours-D)	✓	✓	✓	7.40	<u>25.55</u>
DiST-T (Ours)	✓	✓	✓	<u>6.83</u>	22.67

Table 2. **Quantitative results of RGB Video Generation.** Evaluation on the NuScenes validation set with FID and FVD. Unlike baselines lacking depth modeling, DiST-T generates multi-view RGB videos with depth, achieving state-of-the-art performance. The results of the multi-view variant of Vista* [11] are reported in [22].

Method	Detection ↑		BEV Segmentation ↑				L2 ↓		
	NDS	mAP	Lan.	Dri.	Div.	Cro.	1.0	2.0	3.0
Ori. GT	49.85	37.98	31.31	69.14	25.93	14.36	0.51	0.98	1.65
MD [10]	28.36	12.92	21.95	51.46	17.10	5.25	0.57	1.14	1.95
DA [48]	30.03	16.06	26.14	59.37	20.79	8.92	0.56	1.10	1.89
Ours	32.44	15.63	26.80	60.32	21.69	10.99	0.56	1.11	1.91

Table 3. **Quantitative evaluation of video generation for perception and planning.** We use UniAD [13] to assess detection, BEV segmentation, and L2 planning errors on generated videos compared to original nuScenes data [2]. Our method outperforms baselines in perception and segmentation while maintaining competitive planning accuracy, showing the effectiveness of DiST-4D for various autonomous driving tasks.

$424 \times 800 \times 6$ and a total of 17 frames. Since DiST-T conditions RGB-D video generation on a reference image, the first frame of the reference image is excluded from evaluations. For depth sequence generation, as no prior method jointly generates RGB-D videos in driving scenarios, we compare the depth generation performance with SOTA surround depth estimation methods [41, 59] using common metrics including Absolute Relative Error (Abs.Rel.), Root Mean Squared Error (RMSE) and Threshold Accuracy (δ) as doing in [41].

RGB Video Generation Quality. The quantitative comparison of video generation is illustrated in Tab. 2. DiST-T achieves a remarkable improvement in the FVD metric compared to previous models, indicating high fidelity in the generated videos. In terms of FID, DiST-T achieves 6.83, which is close to the previous SOTA model [22]. We also modify DiST-T to take multi-modal (RGB+depth) reference images as input for comparison, denoted as Ours-D, where the reference depth is obtained from our depth curation pipeline. Tab. 2 indicates that incorporating reference depth has certain

Method	Abs. Rel. ↓	RMSE ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$
<i>Single Frame LiDAR GT</i>				
SD [41]	0.27 / 0.28	7.29 / 7.47	0.68 / 0.66	0.86 / 0.84
M ² Depth [59]	0.24 / 0.26	6.82 / 6.90	0.72 / 0.73	0.88 / 0.87
DiST-T (Ours-D)	0.21 / 0.27	6.79 / 7.06	0.77 / 0.73	0.89 / 0.88
DiST-T (Ours)	0.25 / 0.39	7.39 / 7.75	0.70 / 0.58	0.87 / 0.81
<i>Multiple Frame LiDAR GT</i>				
SD [41]	0.27 / 0.28	6.50 / 6.59	0.67 / 0.63	0.87 / 0.85
M ² Depth [59]	0.25 / 0.26	6.02 / 6.16	0.72 / 0.72	0.89 / 0.88
DiST-T (Ours-D)	0.20 / 0.25	5.33 / 5.58	0.79 / 0.75	0.92 / 0.91
DiST-T (Ours)	0.24 / 0.35	6.24 / 6.89	0.71 / 0.61	0.88 / 0.83

Table 4. **Quantitative evaluation of generated depth.** We compare DiST-T with state-of-the-art surround depth estimation methods [41, 59] on the nuScenes dataset [2]. Results are reported *with / without* median scaling for both single-frame and multi-frame LiDAR ground truth.

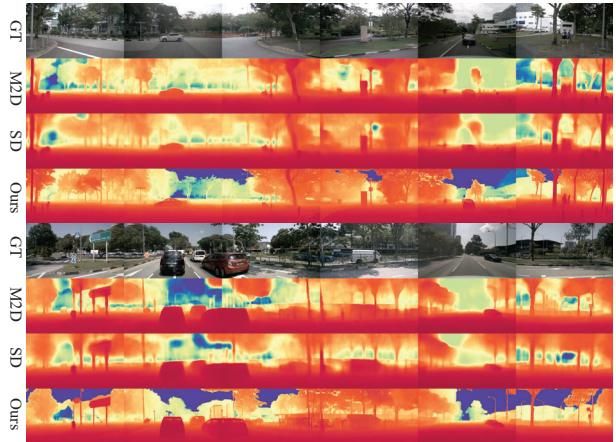


Figure 5. **Visualization results of generated surround depth.** Compared to M²Depth [59] (M2D) and SurroundDepth [41] (SD), DiST-T produces more fine-grained depth with enhanced details.

side effects on RGB forecasting.

Downstream Evaluation. In addition to evaluating the visual quality of the generated images and videos, we assess their effectiveness in downstream tasks, such as perception and open-loop planning using UniAD [13]. The quantitative results in Tab. 3 show that the high-quality videos produced by DiST-T achieve performance comparable to the original ground truth data, highlighting the potential of synthetic data for downstream applications.

Depth Generation Quality. For depth generation, we report both the results of single-frame LiDAR GT [41] and our multi-frame LiDAR GT, which are reported *with / without* median scaling. It is worth noting that our model is different from depth estimation models [41, 59], which take GT image as input to predict the corresponding depth, our method simultaneously generates future RGB frames and their corresponding depth sequences with only one frame of reference images. As shown in Tab. 4, the generated future




Figure 6. **Visualization results of DiST-T.** DiST-T produces high-fidelity multi-view RGB videos (left) and fine-grained depth sequences (right) over time steps.

depths (Ours) achieve comparable performance to [41, 59]. When the reference depths are available for the modified DiST-T (Ours-D), our model surpasses existing models on both single-frame and multi-frame depth GT.

The qualitative results of the RGB-D sequence generated by DiST-T are shown in Fig. 6. Even in complex driving scenarios, DiST-T can generate high-quality RGB videos along with fine-grained depth maps. Additional visualizations of DiST-T are provided in the supplementary materials.

4.2. Spatial Novel View Synthesis

Evaluation Metric. Following the evaluation methodology introduced in [34], we primarily assess the performance of DiST-S on novel trajectories with FID and FVD. We apply different lateral offsets $\tau \in \{\pm 1m, \pm 2m, \pm 4m\}$ to the camera viewpoints and calculate the FID and FVD metrics between the synthesized RGB images of the novel trajectory and the ground truth images of the original trajectory.

Benchmark Evaluation. In Tab. 5, we compare the NVS results of DiST-S with existing methods under shifted novel views. The quantitative results demonstrate that our method achieves substantial improvements in both FID and FVD metrics. When using conditions generated from real images as inputs, DiST-S achieves several times better performance in these metrics, highlighting its strong capability in NVS tasks based on real data. As illustrated in Fig. 7, the proposed DiST-S significantly outperforms previous reconstruction-

Method	Shift $\pm 1m$		Shift $\pm 2m$		Shift $\pm 4m$	
	FID \downarrow	FVD \downarrow	FID \downarrow	FVD \downarrow	FID \downarrow	FVD \downarrow
PVG [3]	48.15	246.74	60.44	356.23	84.50	501.16
EmerNeRF [45]	37.57	171.47	52.03	294.55	76.11	497.85
StreetGaussian [43]	32.12	153.45	43.24	256.91	67.44	429.98
OmniRe [4]	31.48	152.01	43.31	254.52	67.36	428.20
FreeVS* [34]	51.26	431.99	62.04	497.37	77.14	556.14
DiST-4D (Ours)	20.64	130.98	25.08	156.60	33.56	189.04
DiST-4D (+SCC)	16.40	112.86	19.50	124.97	25.16	146.56
DiST-S (Ours)	12.00	55.80	15.72	89.54	22.54	137.03
DiST-S (+SCC)	10.12	45.14	12.97	68.80	17.57	105.29

Table 5. **Quantitative comparison of novel view synthesis.** We evaluate FID and FVD across different viewpoint shifts ($\pm 1m$, $\pm 2m$, $\pm 4m$). FreeVS* denotes we retrain FreeVS [34] using the official code on the nuScenes dataset [2]. DiST-4D uses generated RGB-D videos from DiST-T, while DiST-S leverages real videos with preprocessed metric depth. DiST-S (+SCC) achieves the best performance, demonstrating the effectiveness of self-supervised cycle consistency in improving novel view synthesis quality.

	Depth	Valid Mask	Data Aug.	FID-1m \downarrow	FID-2m \downarrow
(a)	-	-	✓	26.33	33.54
(b)	✓	-	✓	30.31	32.80
(c)	✓	✓	-	26.19	29.69
DiST-S	✓	✓	✓	25.51	27.75

Table 6. **Ablation study of DiST-S.** We analyze the impact of depth modality, valid mask usage, and data augmentation on synthesis quality, reporting FID at 1m and 2m shifts. Results are evaluated on the front camera of a selected nuScenes validation subset [2].

based models in synthesizing novel spatial viewpoints for both RGB and depth, which is nearly free from image degradation and artifacts. Even when using conditions from generated videos of DiST-T (denoted as DiST-4D), our model achieves significantly lower FID and FVD scores than other methods.

When inferring on H20 GPU, DiST-S synthesizes novel views at 1.56s/frame(25 denoising steps). While slower than GS-based methods (which require 1h to reconstruct each 15s video), DiST-S eliminates optimization overhead during inference, offering a promising alternative for NVS.

4.3. Ablation Study

We conduct ablation studies for DiST-S in Tab. 6, reporting FID and FVD results for novel trajectory synthesis with lane shifting. First, we modify DiST-S into an RGB-only spatial NVS model, denoted as (a). By comparing (a) and naive multimodal DiST-S (b), we find that multi-modal shows a slight decrease in FID. We observe that *Data Augmentation* with $n = \pm 4$ frame projection and point cloud downsampling outperforms (c) in generation quality by 2.5%. The results in (b) demonstrate that incorporating the *Valid Mask* significantly enhances DiST-S’s capability, reducing the FID score by 15.8%.

Additionally, we perform an ablation study on the pro-

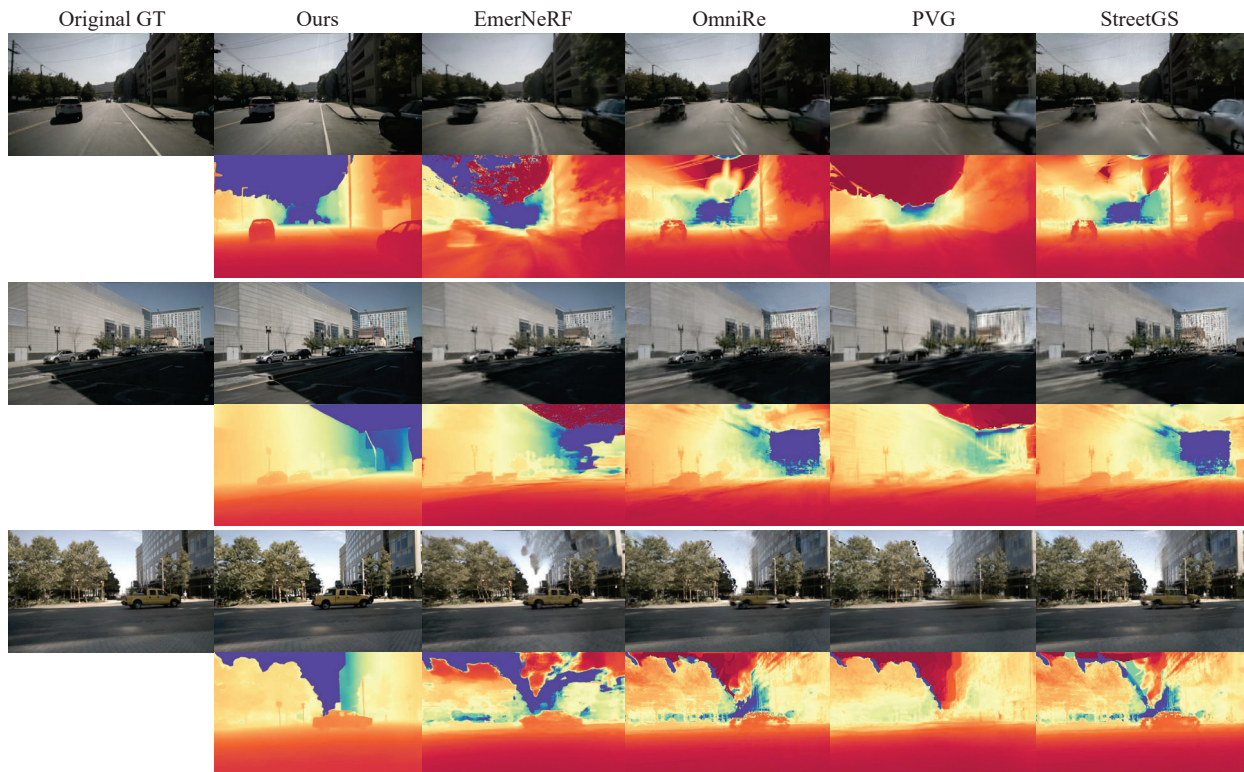


Figure 7. **Qualitative comparison on RGB-D spatial NVS.** Comparison between DiST-S and state-of-the-art reconstruction-based methods [3, 4, 43, 45] on spatial NVS. DiST-S preserves fine-grained scene geometry and reduces artifacts in shifted viewpoints.

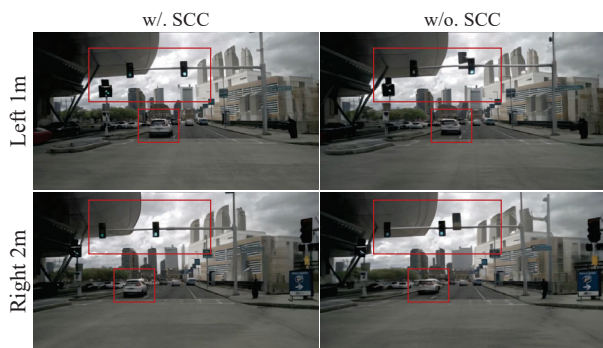


Figure 8. **Qualitative ablation study on the SCC strategy.** We compare novel view synthesis results with and without self-supervised cycle consistency (SCC) under different viewpoint shifts (left 1m and right 2m). Red boxes highlight improved object fidelity and scene consistency when SCC is applied, particularly in fine-grained structures like traffic lights and vehicles.

posed self-supervised cycle consistency (SCC) strategy. As shown in Fig. 8 and Tab. 5, SCC leads to substantial improvements in dynamic vehicle representation and fine object details, significantly boosting the quality of novel view synthesis, with the FID metric also reduced by 20%.

5. Conclusion and Future Work

In this paper, we introduce DiST-4D, a novel framework for feed-forward 4D driving scene generation. With a disentangled spatiotemporal paradigm that bridges temporal forecasting and spatial novel view synthesis with dense metric depth, DiST-4D is capable of generating RGB-D 4D scenarios without per-scene optimization. Extensive experiments show the DiST-4D achieves state-of-the-art performance in both temporal and spatial generation.

Limitations and future work DiST-4D leverages metric depth to connect future temporal predictions with novel spatial viewpoint synthesis. However, its performance is still influenced by the quality of the pseudo metric depth ground truth, especially when it comes to distant buildings. This underscores the need for further research on acquiring high-quality metric depth in outdoor environments. Additionally, expanding DiST-4D’s application to areas such as embodied intelligence and robotics offers a promising avenue for future development.

Acknowledgements

This work is supported by National Key Research and Development Program of China under Grant 2024YFB2505105.

References

- [1] 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. 5
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020. 5, 6, 7
- [3] Yurui Chen, Chun Gu, Junzhe Jiang, Xiatian Zhu, and Li Zhang. Periodic vibration gaussian: Dynamic urban scene reconstruction and real-time rendering. *arXiv preprint arXiv:2311.18561*, 2023. 1, 2, 3, 7, 8
- [4] 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. 1, 7, 8
- [5] Yikang Ding, Wentao Yuan, Qingtian Zhu, Haotian Zhang, Xiangyue Liu, Yuanjiang Wang, and Xiao Liu. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *CVPR*, 2022. 2, 4
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 4
- [7] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision*, pages 241–258. Springer, 2024. 2
- [8] 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. 1, 2
- [9] Ruiyuan Gao, Kai Chen, Bo Xiao, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrivedit: High-resolution long video generation for autonomous driving with adaptive control. *arXiv preprint arXiv:2411.13807*, 2024. 1, 2, 4, 6
- [10] Ruiyuan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, and Qiang Xu. Magicdrive: Street view generation with diverse 3d geometry control. In *ICLR*, 2024. 1, 2, 4, 5, 6
- [11] 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. 1, 6
- [12] 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. *Advances in Neural Information Processing Systems*, 37:91560–91596, 2025. 2
- [13] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *CVPR*, 2023. 6
- [14] Binyuan Huang, Yuqing Wen, Yucheng Zhao, Yaosi Hu, Yingfei Liu, Fan Jia, Weixin Mao, Tiancai Wang, Chi Zhang, Chang Wen Chen, et al. Subjectdrive: Scaling generative data in autonomous driving via subject control. *arXiv preprint arXiv:2403.19438*, 2024. 1
- [15] Jitesh Jain, Jiachen Li, MangTik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. OneFormer: One Transformer to Rule Universal Image Segmentation. 2023. 2, 4
- [16] Sagi Katz, Ayellet Tal, and Ronen Basri. Direct visibility of point sets. In *ACM SIGGRAPH 2007 papers*, pages 24–es. 2007. 4
- [17] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 5
- [18] 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. 1, 2, 3
- [19] Seung Wook Kim, Jonah Philion, Antonio Torralba, and Sanja Fidler. Drivegan: Towards a controllable high-quality neural simulation. In *CVPR*, 2021. 6
- [20] Bohan Li, Yasheng Sun, Zhujin Liang, Dalong Du, Zhuanghui Zhang, Xiaofeng Wang, Yunnan Wang, Xin Jin, and Wenjun Zeng. Bridging stereo geometry and bev representation with reliable mutual interaction for semantic scene completion. *arXiv preprint arXiv:2303.13959*, 2023. 2
- [21] Bohan Li, Jiajun Deng, Wenyao Zhang, Zhujin Liang, Dalong Du, Xin Jin, and Wenjun Zeng. Hierarchical temporal context learning for camera-based semantic scene completion. In *European Conference on Computer Vision*, pages 131–148. Springer, 2024. 2
- [22] Bohan Li, Jiazhe Guo, Hongsi Liu, Yingshuang Zou, Yikang Ding, Xiwu Chen, Hu Zhu, Feiyang Tan, Chi Zhang, Tiancai Wang, et al. Uniscene: Unified occupancy-centric driving scene generation. *arXiv preprint arXiv:2412.05435*, 2024. 1, 2, 6
- [23] Hanxue Liang, Jiawei Ren, Ashkan Mirzaei, Antonio Torralba, Ziwei Liu, Igor Gilitschenski, Sanja Fidler, Cengiz Oztireli, Huan Ling, Zan Gojcic, et al. Feed-forward bullet-time reconstruction of dynamic scenes from monocular videos. *arXiv preprint arXiv:2412.03526*, 2024. 3
- [24] Zhiheng Liu, Ka Leong Cheng, Qiuyu Wang, Shuzhe Wang, Hao Ouyang, Bin Tan, Kai Zhu, Yujun Shen, Qifeng Chen, and Ping Luo. Depthlab: From partial to complete. *arXiv preprint arXiv:2412.18153*, 2024. 2, 3, 4
- [25] Jiachen Lu, Ze Huang, Jiahui Zhang, Zeyu Yang, and Li Zhang. Wovogen: World volume-aware diffusion for controllable multi-camera driving scene generation. *arXiv preprint arXiv:2312.02934*, 2023. 1, 6
- [26] Jiageng Mao, Boyi Li, Boris Ivanovic, Yuxiao Chen, Yan Wang, Yurong You, Chaowei Xiao, Danfei Xu, Marco Pavone, and Yue Wang. Dreamdrive: Generative 4d scene modeling from street view images. *arXiv preprint arXiv:2501.00601*, 2024. 1, 2

- [27] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6121–6132, 2025. 2
- [28] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, 2020. Version 0.3.0. 5
- [29] Alexander Swerdlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *IEEE Robotics and Automation Letters*, 2024. 6
- [30] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 2
- [31] Qijian Tian, Xin Tan, Yuan Xie, and Lizhuang Ma. Drivingforward: Feed-forward 3d gaussian splatting for driving scene reconstruction from flexible surround-view input. *arXiv preprint arXiv:2409.12753*, 2024. 1, 3
- [32] Lening Wang, Wenzhao Zheng, Dalong Du, Yunpeng Zhang, Yilong Ren, Han Jiang, Zhiyong Cui, Haiyang Yu, Jie Zhou, Jiwen Lu, et al. Stag-1: Towards realistic 4d driving simulation with video generation model. *arXiv preprint arXiv:2412.05280*, 2024. 1, 2, 3
- [33] Lening Wang, Wenzhao Zheng, Yilong Ren, Han Jiang, Zhiyong Cui, Haiyang Yu, and Jiwen Lu. Occsora: 4d occupancy generation models as world simulators for autonomous driving. *arXiv preprint arXiv:2405.20337*, 2024. 1
- [34] Qitai Wang, Lue Fan, Yuqi Wang, Yuntao Chen, and Zhaoxiang Zhang. Freevs: Generative view synthesis on free driving trajectory. *arXiv preprint arXiv:2410.18079*, 2024. 1, 2, 3, 4, 7
- [35] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2566–2576, 2019. 2, 5
- [36] Xiaofeng Wang, Zheng Zhu, Yunpeng Zhang, Guan Huang, Yun Ye, Wenbo Xu, Ziwei Chen, and Xingang Wang. Are we ready for vision-centric driving streaming perception? the asap benchmark. *arXiv preprint arXiv:2212.08914*, 2022. 5
- [37] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. *ECCV*, 2024. 1, 6
- [38] 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 *CVPR*, 2024. 1, 6
- [39] 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
- [40] Julong Wei, Shanshuai Yuan, Pengfei Li, Qingda Hu, Zhongxue Gan, and Wenchao Ding. Occllama: An occupancy-language-action generative world model for autonomous driving. *arXiv preprint arXiv:2409.03272*, 2024. 1
- [41] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Yongming Rao, Guan Huang, Jiwen Lu, and Jie Zhou. Surround-depth: Entangling surrounding views for self-supervised multi-camera depth estimation. In *Conference on robot learning*, pages 539–549. PMLR, 2023. 3, 6, 7
- [42] 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, 2023. 1, 6
- [43] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians: Modeling dynamic urban scenes with gaussian splatting. In *European Conference on Computer Vision*, pages 156–173. Springer, 2024. 1, 2, 3, 7, 8
- [44] Yunzhi Yan, Zhen Xu, Haotong Lin, Haian Jin, Haoyu Guo, Yida Wang, Kun Zhan, Xianpeng Lang, Hujun Bao, Xiaowei Zhou, et al. Streetcrafter: Street view synthesis with controllable video diffusion models. *arXiv preprint arXiv:2412.13188*, 2024. 1, 3
- [45] 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. 1, 3, 7, 8
- [46] Jiawei Yang, Jiahui Huang, Yuxiao Chen, Yan Wang, Boyi Li, Yurong You, Apoorva Sharma, Maximilian Igl, Peter Karkus, Danfei Xu, et al. Storm: Spatio-temporal reconstruction model for large-scale outdoor scenes. *arXiv preprint arXiv:2501.00602*, 2024. 1, 2, 3
- [47] Kairui Yang, Enhui Ma, Jibin Peng, Qing Guo, Di Lin, and Kaicheng Yu. Bevcontrol: Accurately controlling street-view elements with multi-perspective consistency via bev sketch layout. *arXiv preprint arXiv:2308.01661*, 2023. 6
- [48] 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. 6
- [49] Ze Yang, Yun Chen, Jingkan Wang, Sivabalan Manivasagam, Wei-Chiu Ma, Anqi Joyce Yang, and Raquel Urtasun. Unisim: A neural closed-loop sensor simulator. In *CVPR*, 2023. 1, 2, 3
- [50] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 4
- [51] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024. 2, 4
- [52] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d

- gaussian splatting for unconstrained image collections. In *European Conference on Computer Vision*, pages 341–359. Springer, 2024. [2](#)
- [53] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. [4](#)
- [54] Guosheng Zhao, Chaojun Ni, Xiaofeng Wang, Zheng Zhu, Xueyang Zhang, Yida Wang, Guan Huang, Xinze Chen, Boyuan Wang, Youyi Zhang, et al. Drivedreamer4d: World models are effective data machines for 4d driving scene representation. *arXiv preprint arXiv:2410.13571*, 2024. [1](#), [2](#), [3](#)
- [55] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. *arXiv preprint arXiv:2311.16038*, 2023. [1](#)
- [56] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. [4](#)
- [57] Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, and Alexei A Efros. Learning dense correspondence via 3d-guided cycle consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 117–126, 2016. [2](#), [5](#)
- [58] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#), [5](#)
- [59] Yingshuang Zou, Yikang Ding, Xi Qiu, Haoqian Wang, and Haotian Zhang. M2depth: Self-supervised two-frame multi-camera metric depth estimation. In *European Conference on Computer Vision*, pages 269–285. Springer, 2024. [3](#), [6](#), [7](#)
- [60] Vlas Zyrianov, Henry Che, Zhijian Liu, and Shenlong Wang. Lidardm: Generative lidar simulation in a generated world. *arXiv preprint arXiv:2404.02903*, 2024. [1](#)