

HorizonForge: Driving Scene Editing with Any Trajectories and Any Vehicles

Yifan Wang^{1,2} Francesco Pittaluga¹ Zaid Tasneem¹
Chenyu You² Manmohan Chandraker^{1,3} Ziyu Jiang¹

¹NEC Labs America ²Stony Brook University ³UC San Diego

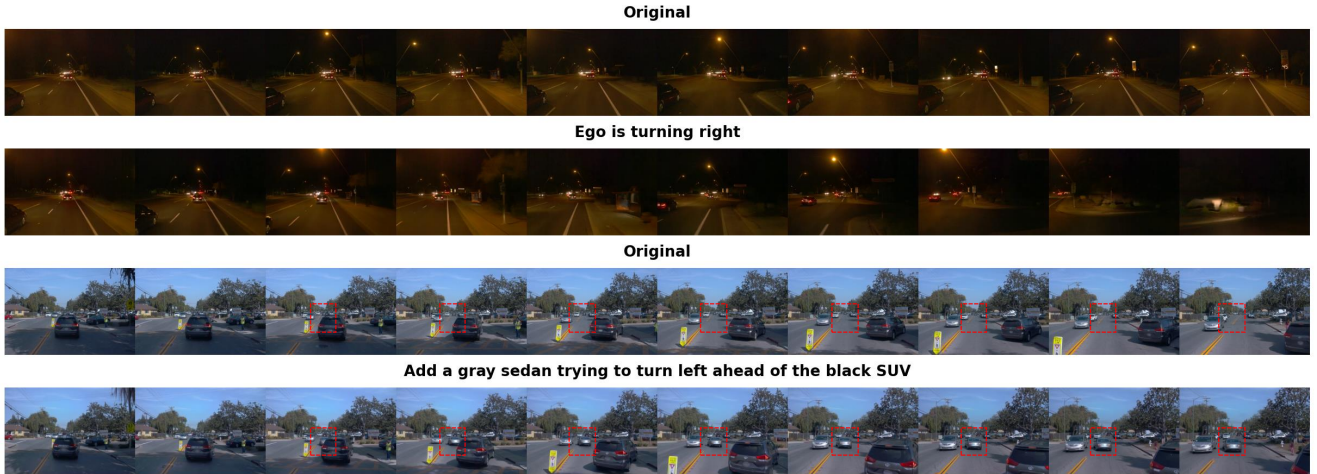


Figure 1. **HorizonForge** is capable of generating high-quality driving scenes in accordance with the provided manipulation instructions. The top two rows of the image depict the transformation of the ego car to the right, while the bottom two rows illustrate the insertion of a gray sedan in front of the selected SUV at the red box location.

Abstract

Controllable driving scene generation is critical for realistic and scalable autonomous driving simulation, yet existing approaches struggle to jointly achieve photorealism and precise control. We introduce **HorizonForge**, a unified framework that reconstructs scenes as editable Gaussian Splats and Meshes, enabling fine-grained 3D manipulation and language-driven vehicle insertion. Edits are rendered through a noise-aware video diffusion process that enforces spatial and temporal consistency, producing diverse scene variations in a single feed-forward pass without per-trajectory optimization. To standardize evaluation, we further propose **HorizonSuite**, a comprehensive benchmark spanning ego- and agent-level editing tasks such as trajectory modifications and object manipulation. Extensive experiments show that Gaussian-Mesh representation delivers substantially higher fidelity than alternative 3D representations, and that temporal priors from video diffusion are essential for coherent synthesis. Combining these findings, **HorizonForge** establishes a simple yet powerful paradigm

for photorealistic, controllable driving simulation, achieving an 83.4% user-preference gain and a 25.19% FID improvement over the second best state-of-the-art method. Project page: <https://horizonforge.github.io/>.

1. Introduction

Evaluating autonomous driving algorithms under long-tail scenarios is essential for ensuring safety during real-world deployment. Although modern autonomous driving systems collect vast amounts of daily driving logs, rare but safety-critical events, such as aggressive lane-switching or sudden braking, remain extremely difficult and expensive to capture at scale. Recent advances in simulation have made it increasingly feasible to synthesize such scenarios by editing existing driving logs. However, achieving both high photorealism and precise controllability remains challenging. Reconstruction-based methods [11, 22] offer explicit geometric fidelity through accurate 3D modeling, yet often fail to generalize to unseen regions [66]. In contrast,

generation-based approaches can hallucinate new content but frequently struggle to preserve scene structure and appearance, limiting their ability to perform fine-grained traffic editing [14]. Although recent hybrid methods [32, 68] have shown promising results, a simple, unified framework coupled with a systematic study of controllable driving editing is still missing.

To bridge this gap, we propose **HorizonForge**, a simple yet unified framework for photorealistic and controllable driving scene generation. Our approach begins by harvesting the input scene into editable 3D representations, such as Gaussian Splats and Meshes, that support flexible manipulation in 3D space. Each edit (e.g., lane changes, cut-ins/outs, or vehicle insertion) is then rendered back into the image domain through a noise-aware video diffusion process, ensuring strong spatial and temporal consistency. Additionally, **HorizonForge** supports language-guided object insertion, enabling the placement of new vehicles described via natural language by generating the corresponding 3D meshes. Once the 3D structure is acquired, diverse controllable variations can be produced in a single feed-forward pass, avoiding the costly per-trajectory optimization required by prior methods [32, 68]. Despite its simplicity, **HorizonForge** achieves state-of-the-art performance across a wide range of editing tasks.

To systematically evaluate controllable editing at scale, we introduce **HorizonSuite**, a comprehensive benchmark for controllable driving simulation. It spans a rich set of challenging tasks involving both ego-vehicle and surrounding-agent manipulation, including lane changes, turns, cut-ins/outs, sudden braking or acceleration, and object insertion or removal. This benchmark establishes a unified protocol for assessing fine-grained controllability and visual fidelity across diverse editing scenarios. Through extensive analysis, we identify two crucial design principles for controllable scene generation:

1. **3D representations matter.** Comparing bounding boxes, colored LiDAR point clouds, and Gaussian Splats (all combined with object meshes), we find that Gaussian Splats encode far richer appearance cues, enabling more accurate edits and substantially improving generation quality.
2. **Temporal priors matter.** Our comparison between image- and video-based paradigms shows that video diffusion models more effectively leverage temporal continuity, leading to significantly more coherent and stable scene synthesis.

In summary, our main contributions are

- **A simple and unified framework (HorizonForge)** for photorealistic and controllable driving scene generation, supporting both 3D editing and language-guided object insertion.
- **A mesh harvesting and insertion pipeline** that enables

insertion of any vehicles with only text prompts.

- **A comprehensive benchmark (HorizonSuite)** enabling standardized evaluation across diverse ego- and agent-level manipulation tasks with fine-grained metrics.
- **Extensive experiments and analysis** showing that HorizonForge outperforms state-of-the-art approaches, delivering an 83.4% user-preference gain and a 25.19% FID improvement over the second-best method, while offering insights into the key design choices for controllable scene generation.

2. Related Works

2.1. Reconstruction-Based Simulation

Reconstruction-based simulation methods focus on faithfully recovering the geometry and appearance of real driving scenes for realistic replay or manipulation. Large-scale benchmarks such as Waymo [44], nuScenes [6], and nuPlan [7] provide the empirical foundation for these efforts. Recent advances in neural rendering [33, 46, 48, 49, 56, 57, 60, 61], particularly 3D Gaussian Splatting [11, 22, 58], have enabled photorealistic multi-view reconstructions in driving scenarios. While these approaches achieve strong visual fidelity and geometric accuracy, their controllability is limited because they primarily replay reconstructed scenes instead of synthesizing novel trajectories or new agents.

2.2. Generation-Based Simulation

In contrast, generation-based methods aim to synthesize controllable driving scenes directly from high-level inputs such as text or semantic maps. Diffusion-based generative models [3, 9, 12, 19–21, 38, 41, 42, 47] have significantly advanced this direction by improving fidelity and temporal coherence, yet most rely on complex multimodal conditioning and 2D cues without explicit physical grounding. As a result, their outputs often exhibit stochastic motion and structural inconsistency, making them unreliable for safety-critical domains such as autonomous driving. Although recent works [2, 4, 31, 67] improve temporal transfer and latent alignment, they still depend heavily on intricate conditioning pipelines and prompt engineering. In this work, we revisit the simplicity of classical video diffusion models, showing that with minimal yet physically grounded conditioning, specifically, trajectory-based control combined with 3D Gaussian and mesh representations, one can achieve both strong generalization and physically consistent controllable generation.

2.3. Hybrid Simulation

Hybrid simulation methods seek to combine the geometric accuracy of reconstruction with the flexibility of generative modeling by integrating 3D structural priors into diffusion

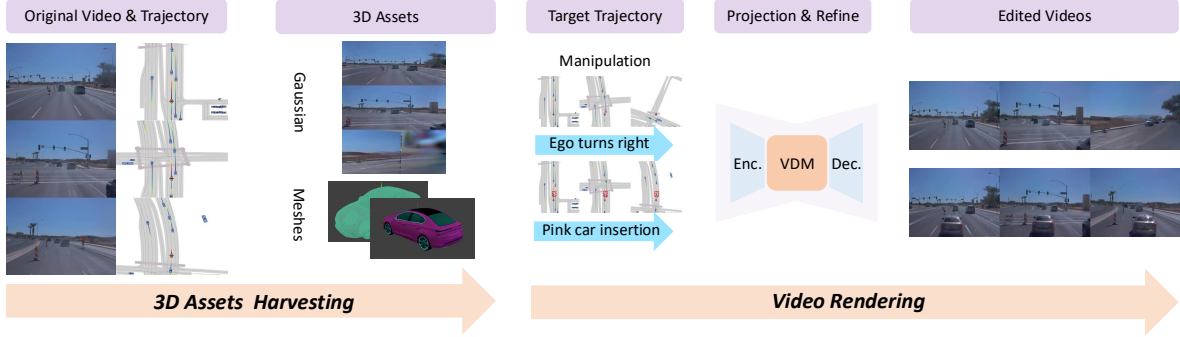


Figure 2. Overview of the **HorizonForge** framework. With original video and trajectory, we will firstly extract corresponding 3D assets according to the manipulated novel trajectories, then feed them into our rendering model for final generation results.

or transformer frameworks [5, 10, 13, 15, 17, 24, 25, 27, 28, 30, 32, 35, 45, 52–55, 59, 63, 64, 69]. While these approaches improve realism and controllability, they often rely on complex architectures and per-trajectory optimization [32, 68], leading to high computational cost and limited scalability. In contrast, **HorizonForge** adopts a simple yet effective paradigm: it reconstructs the scene into editable 3D Gaussian splats and meshes, and employs a standard video diffusion model to render temporally coherent results directly from trajectory-based control. This streamlined design eliminates redundant conditioning and heavy optimization while achieving superior photorealism and multi-agent controllability. Together with **HorizonSuite**, it offers a unified and scalable foundation for physically grounded driving simulation.

3. HorizonForge

3.1. Overview

As illustrated in Figure 2, **HorizonForge** is a unified framework for controllable driving scene generation. Given a user-specified trajectory set $\mathcal{T} = \{\tau_i\}_{i=1}^N$ that jointly specifies the motion of all vehicles in the scene—including both the ego vehicle and surrounding agents—our system synthesizes a photorealistic and physically consistent driving video through two major stages:

1. **3D Assets Harvesting.** We first reconstruct the input video into a collection of editable 3D assets that can be flexibly manipulated under arbitrary motions in 3D space. To preserve both geometric structure and appearance fidelity, we extract high-quality 3D Gaussian Splat representations [11]. Through a systematic comparison with alternative modalities, including colored point clouds [59] and 3D bounding boxes, we find that Gaussian Splats consistently offer superior controllability and rendering quality, owing to their substantially richer appearance detail (Section 4.2). In addition, we harvest 3D meshes from the input video to support text-driven novel object insertion. Specifically, the rendering model is trained with these meshes to enable realistic,

geometry-aligned reconstruction of newly inserted vehicles. Combined with a text-conditioned mesh generation pipeline (Section 3.3), this allows **HorizonForge** to introduce new vehicles described purely in natural language while maintaining scene structure and compatibility with downstream rendering.

2. **Video Rendering:** After repositioning all harvested 3D assets according to the target trajectories \mathcal{T} , the edited 3D scene is rendered into 2D view sequences. These rasterized frames are subsequently refined by a Video Diffusion Model (VDM), which enforces spatial-temporal consistency and synthesizes high-fidelity driving videos aligned with the specified multi-agent motions. This stage yields photorealistic and temporally stable outputs that faithfully preserve the scene geometry and adhere to the dynamics dictated by \mathcal{T} . In Section 4.2, we systematically studied the image and video diffusion model and reveal that video diffusion model provides significant boost on the temporal consistency and helps the model to better resolving the artifacts in the rasterized frames.

This paradigm enables **HorizonForge** to seamlessly integrate explicit multi-agent control with implicit generative modeling, allowing flexible scene manipulation in which the trajectories of all agents, including the ego vehicle, collectively define the resulting dynamics. The design maintains both structural consistency and photorealistic fidelity. Moreover, once the 3D scene representations are harvested, each edited variation can be rendered efficiently with a single feed-forward pass.

In the following sections, we first introduce the necessary preliminaries in Section 3.2. We then describe the 3D Asset Harvesting and Video Rendering components in Section 3.3 and Section 3.4, respectively. Finally, we present the proposed benchmark **HorizonSuite** in Section 3.5.

3.2. Preliminary

3.2.1. 3D Representation Generation

Our framework leverages explicit 3D representations as conditions to achieve high-fidelity and robust control. We

provide a brief overview of 3D Gaussian Splatting.

3D Gaussian Splatting (3DGS). 3DGS [22] is a high-fidelity, explicit 3D scene representation that enables real-time, photorealistic rendering. A scene is represented by a set of N anisotropic 3D Gaussians. Each Gaussian G_i is defined by following four components, Position (Mean): $\mu_i \in \mathbb{R}^3$; Covariance: $\Sigma_i \in \mathbb{R}^{3 \times 3}$, typically parameterized by a scale vector $s_i \in \mathbb{R}^3$ and a rotation quaternion $q_i \in \mathbb{R}^4$; Opacity: $o_i \in \mathbb{R}$ and Color: A set of Spherical Harmonic (SH) coefficients c_i to model view-dependent color.

To render an image from a camera pose, these 3D Gaussians are projected onto the 2D image plane. The final pixel color C is computed by alpha-blending the sorted 2D Gaussians that overlap the pixel in a front-to-back order: $C = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$, where α_i is the computed opacity of the i -th Gaussian in 2D. 3DGS is known for its ability to capture fine appearance details but can produce noisy artifacts under large, novel view shifts.

3.2.2. Video Diffusion Models

Video Diffusion Models [20] consist of two main components: a forward process and a learned reverse process.

Forward Process. The forward process q gradually adds noise to a clean video x_0 over T discrete timesteps. This is a fixed Markov process defined as $q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1})$, where $q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I})$, and β_t is a predefined noise schedule. And we can sample x_t at any arbitrary timestep t in a closed form: $q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\alpha_t}x_0, (1 - \alpha_t)\mathbf{I})$, where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. As $t \rightarrow T$, x_T approaches a pure Gaussian noise $\mathcal{N}(0, \mathbf{I})$.

Reverse Process. The reverse process p_θ is a neural network trained to denoise x_t back to x_{t-1} for all $t \in [1, T]$. This process is parameterized as a Gaussian: $p_\theta(x_{t-1}|x_t, c) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t, c), \Sigma_\theta(x_t, t, c))$, where c represents the conditioning information. In practice, the network ϵ_θ is trained to predict the noise ϵ added to x_0 to create x_t . During inference, a video x_0 is generated by iteratively sampling $x_{t-1} \sim p_\theta(x_{t-1}|x_t, c)$ starting from pure noise $x_T \sim \mathcal{N}(0, \mathbf{I})$. This iterative process is the source of the train-inference discrepancy that our work addresses.

3.3. 3D Assets Harvesting

The rapid progress of 3D reconstruction has made it possible to lift 2D video frames back into 3D space, enabling game-engine-style manipulation while preserving real-world photorealism. Our framework begins by harvesting diverse 3D assets that support fine-grained scene editing with both geometric and appearance consistency. To systematically compare different 3D representations, we extract Gaussian splats, colored LiDAR point clouds, and 3D bounding boxes. Gaussian splats offer a dense, view-consistent radiance representation; colored LiDAR point clouds provide strong novel-view synthesis performance




Figure 3. A demonstration of 3D Meshes Harvesting pipeline.

as demonstrated in StreetCrafter [59]; and 3D bounding boxes remain a common control signal in generation methods [13, 40].

We obtain high-quality Gaussian splats using OmniRe [11], optimized across multiple nodes for large-scale driving logs. Following StreetCrafter [59], we harvest colored LiDAR point clouds by projecting 3D points into the image plane for color assignment. 3D bounding boxes are extracted directly from the Waymo dataset [44]. Our study in Section 4.2 shows that Gaussian splats preserve fine appearance details significantly better than sparse 3D primitives such as bounding boxes or colored point clouds.

However, although Gaussian splats excel at editing existing vehicles, inserting new assets remains challenging. To address this limitation, we incorporate mesh representations, which can be generated from text using off-the-shelf foundation models such as Hunyuan3D [70]. We introduce two complementary pipelines to enable mesh-based insertion: *3D Mesh Harvesting* and *Inference-Time Insertion*. The first pipeline extracts meshes from video and reinserts them during training to reduce the photorealism gap between meshes and real imagery. The second pipeline converts arbitrary text descriptions into insertable 3D assets for downstream controllable editing tasks.

3D Meshes Harvesting. As shown in Figure 3, for mesh harvesting, we extract vehicle meshes from the Waymo dataset [44] using ground-truth 3D bounding boxes and LiDAR points. For each object, we first identify its best observation frame f^* , defined as the frame containing the largest number of associated LiDAR points, which typically provides minimal occlusion and the closest viewpoint. We then apply SAM [23] to segment and crop the corresponding image, producing C_{f^*} . To handle occlusions, we perform Amodal inpainting using Pix2Gestalt [34]. The inpainted crop is then passed to Hunyuan3D [70] to obtain a raw mesh:

$$M_{\text{raw}} = \mathcal{H}(C_{f^*}). \quad (1)$$

Meshes produced by Hunyuan3D do not preserve the vehicle’s original pose or metric dimensions. Since accurate geometry is required for re-insertion during rendering, we need to align each generated mesh with its ground-truth 3D bounding box. While vehicles typically have a dominant length dimension, and Hunyuan3D outputs a consistent upright orientation, reducing the candidate poses to two. We then further use a GPT-based heading reasoning module [1] to select the correct orientation.

Given the correct rotation, we assume the mesh center

coincides with the center of the ground-truth box and optimize only the scale s . We jointly minimize a depth discrepancy term and an IoU-based alignment term:

$$\text{Score} = \|D(M, s) - D_{\text{gt}}\| - \lambda \text{IoU}(B(M, s), B_{\text{gt}}), \quad (2)$$

$$s^* = \arg \min_s \text{Score}, \quad (3)$$

where $D(\cdot)$ denotes rendered depth and $B(\cdot)$ denotes the projected bounding box. We solve for s^* using a grid search initialized around the ground-truth box dimensions.

Inference-Time Insertion. At inference, to enable inserting any vehicles with a text description P , we start from leveraging GPT [1] to produce a reference image I_{ref} and feed it into Hunyuan3D [70] to obtain M_{gen} . Afterwards, the rotation and scale would be determined with a Vision Language Model reasoning process leveraging GPT [1]. This guarantees that all inserted or modified vehicles follow their physically valid paths without spatial conflict.

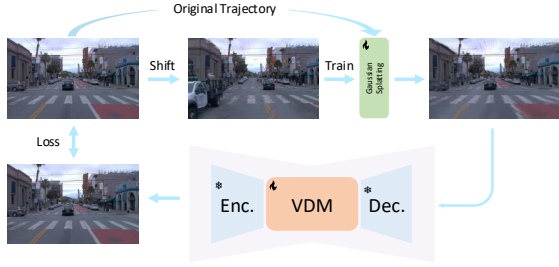


Figure 4. A demonstration of constructing data pairs for Gaussian Splats.

3.4. Video Rendering

With the harvested 3D assets, we can reposition and manipulate any agent in the scene according to a target trajectory. However, rendering these 3D assets back into 2D videos inevitably introduces artifacts, as perfect 3D reconstruction is difficult, if not impossible, to obtain. To bridge this gap, we introduce a video rendering module that maps imperfect 3D renderings into photorealistic video sequences. This section first describes how we construct training pairs for each representation, followed by the diffusion-based training objective.

Constructing data pairs for Gaussian Splats. Artifacts in Gaussian splats commonly arise when rendering from novel viewpoints. A robust rendering module must therefore learn to correct such distortions and synthesize high-quality, temporally consistent video. Motivated by prior work on reconstruction-based denoising [32, 55], we adopt a cycle-reconstruction strategy similar to [55] as shown in Figure 4.

Given a Gaussian field G_1 reconstructed from a real scene under trajectories \mathcal{T} , we perturb the trajectories to obtain $\tilde{\mathcal{T}}$:

$$\tilde{\mathcal{T}} = \mathcal{T} + \Delta\mathcal{T}, \quad \Delta\mathcal{T} = \{\delta_i\}_{i=1}^N, \quad (4)$$

where each δ_i applies horizontal or vertical shifts or a change in heading angle. We then render perturbed frames \tilde{v} from G_1 under $\tilde{\mathcal{T}}$, train a new Gaussian field G_2 on \tilde{v} , and render G_2 back under the original \mathcal{T} to form clean-noisy training pairs:

$$G_2 = \mathcal{T}_{\text{recon}}(\tilde{v}), \quad \hat{v} = \mathcal{R}(G_2, \mathcal{T}). \quad (5)$$

This process yields paired samples (\tilde{v}, v) that enable the rendering module to learn to remove Gaussian-splat artifacts and restore photorealistic appearance.

Constructing data pairs for Meshes. Meshes often exhibit a substantial photorealism gap compared to real driving videos. To train a rendering model that can compensate for this discrepancy, we construct paired training examples by randomly replacing Gaussian-splat assets with harvested meshes during rendering. For each scene, we synthesize two versions: one with mesh-insertion probability $p = 0.5$ and one with $p = 0.0$ (mimicking scenes without object insertion). We additionally randomize Gaussian-splat lighting conditions to enhance robustness to illumination differences between mesh renderings and real scenes.

For rendering compatibility across representations, Gaussian RGB and depth maps are rendered using OmniRe [11], while mesh RGB and depth are generated using PyVista [43]. The two modalities are then merged via depth composition. Example data pairs are shown in Figure 5.



Figure 5. A demonstration of mesh-Gaussian training data pairs. The top frames are the original Gaussian Splats and the bottom ones are Gaussian with mesh vehicle replacements

Diffusion Objective. For each training pair, we render the conditioning video v_c from Gaussian splats or meshes under synchronized camera poses and train a video diffusion model to map these imperfect renderings to photorealistic outputs. Given a clean target video x_0 and a timestep $t \sim \mathcal{U}\{1, \dots, T\}$, we sample noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ and construct the noisy input

$$x_t = \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon. \quad (6)$$

The model predicts $\epsilon_\theta(x_t, t, v_c)$ and is trained with the standard denoising objective

$$\mathcal{L}_{\text{vdm}} = \mathbb{E}_{t, \epsilon} \left[\|\epsilon - \epsilon_\theta(x_t, t, v_c)\|_2^2 \right]. \quad (7)$$

This formulation works effectively when paired with explicit 3D conditioning, resulting in stable and coherent multi-agent video synthesis. In practice, we fine-tune the TrajectoryCrafter pretrained model [65] using the CogVideoX backbone [62] for video diffusion.

We additionally compare this video-based rendering module with an image diffusion approach following [36, 55]. As shown in Section 4.2, our experiments show that video diffusion yields substantially better results, benefiting from richer temporal priors and improved cross-frame consistency.

3.5. HorizonSuite

To quantitatively evaluate controllable multi-agent scene generation, we introduce **HorizonSuite**, a comprehensive benchmark that assesses both ego-vehicle and surrounding-vehicle manipulations under a unified trajectory set \mathcal{T} .

Unlike prior benchmarks that primarily focus on ego lane-changing performance [32, 59], **HorizonSuite** targets realistic editing requirements encountered in real-world driving scenarios. We define a diverse suite of manipulation tasks for both ego and non-ego agents, including speed changes, lane changes, direction changes, insertions, and removals. For each task category, we begin by describing the desired trajectory modification in natural language, then apply the language-driven trajectory editing model LangTraj [8] to produce corresponding trajectory edits within the scenes sampled from Waymo evaluation dataset [44]. Afterward, we manually select 109 challenging, high-quality edited trajectories spanning all manipulation types, forming a rigorous testbed for controllable trajectory editing. Figure 6 illustrates an example of ego-vehicle direction change. Details of **HorizonSuite** are provided in Appendix 7.




Figure 6. A demonstration for trajectory modification in our proposed **HorizonSuite**.

To rigorously assess video quality, temporal consistency, and alignment with both the source videos and input trajectories, we employ the following five metrics:

- **FID & FVD** [18, 50]: Fréchet Inception Distance and Fréchet Video Distance evaluate spatial fidelity and temporal realism of generated videos relative to the source sequences.
- **Vehicle Identity Matching Score (VIMS)**: Based on the intuition that vehicles occupying similar positions in the ego coordinate frame should retain consistent appearance across the source and edited videos, VIMS jointly measures trajectory adherence and visual identity preservation. For each non-occluded vehicle in each frame, we locate its source counterpart with a similar ego-frame pose

and compute their CLIP similarity. The final score averages these similarities across all vehicles and frames. A formal definition is provided in Appendix 7.1.1.

- **Background Alignment Score (BAS)**: BAS quantifies static-scene preservation by computing CLIP similarity [39] between the background regions of the edited frames and its closest source frames under ego coordinates after masking out all dynamic objects. A formal definition is provided in Appendix 7.1.2.
- **Operation Success Rate (OSR)**: For insertion and removal tasks, we measure semantic correctness using GPT-based reasoning [1]. The model assigns a score from 1 to 10 reflecting whether the manipulation is correctly realized and visually realistic. A formal definition is provided in Appendix 7.1.3.

Together, **HorizonForge** and **HorizonSuite** establish a unified foundation for multi-agent, physically consistent, and photorealistic driving-scene editing and evaluation.

4. Experiments

In this section, we conduct a series of experiments to validate the effectiveness of our proposed **HorizonForge** comparing with multiple baselines and settings on **HorizonSuite**. We finetuned our model for 60k steps based on CogVideoX [62] backbone, utilizing AdamW [29] as optimizer. For more implementation details, please refer to Appendix 6. In Section 4.1, we evaluate its task performances comparing against state-of-the-art and representative rendering methods to demonstrate its efficacy. In Section 4.2, we present systematic ablation studies to explore the impact of different modalities on autonomous driving scene rendering tasks and the importance of temporal consistency for such tasks. Finally, Section 4.3 further conducts a comprehensive user study to demonstrate that the quality of videos generated by **HorizonForge** are aligned with human preference. As for downstream task, please refer to Appendix 8.

4.1. Quality Comparisons

We compare **HorizonForge** with three representative baselines covering different paradigms of controllable scene generation: StreetCrafter [59], Difx3D [55] and OmniRe [11]. All methods are evaluated on **HorizonSuite** under the same settings, covering both ego-vehicle and multi-agent manipulation tasks such as speed change, lane change, direction change, insertion, and removal. As shown in Table 1 and Figure 7, **HorizonForge** consistently outperforms all baselines across almost all metrics and achieves the best overall FID and FVD scores. Compared with baselines, our Gaussian-mesh representation knowledge demonstrates superior generalizability and produces finer appearance details and preserves stable geometry under strong ego and agent motion. These results showcase that

(a) Ego Car Manipulation												
Method	Speed Change				Lane Change				Direction Change			
	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑
StreetCrafter [59]	150.11	1232.13	0.9623	0.8754	134.65	1347.02	0.9568	0.8612	199.28	1846.96	0.9573	0.8497
Difix3D [55]	146.11	1019.25	<u>0.9643</u>	0.8357	135.80	1049.53	<u>0.9586</u>	0.8191	203.11	1416.44	0.9590	0.8052
OmniRe [11]	<u>46.89</u>	548.22	0.9619	<u>0.9121</u>	<u>54.02</u>	667.70	0.9541	<u>0.9102</u>	80.60	<u>1237.23</u>	0.9517	<u>0.8719</u>
Ours	39.92	<u>583.78</u>	0.9657	0.9326	50.33	<u>729.76</u>	0.9615	0.9303	<u>144.83</u>	1175.52	<u>0.9585</u>	0.8947

(b) Other Vehicle Manipulation I												
Method	Speed Change				Lane Change				Direction Change			
	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑
StreetCrafter [59]	102.90	1181.27	0.9674	0.8887	163.21	1467.54	0.9600	0.8818	140.32	1310.09	0.9683	0.8820
Difix3D [55]	128.27	1038.93	<u>0.9716</u>	0.8500	149.11	1358.60	<u>0.9624</u>	0.8449	133.91	994.03	<u>0.9708</u>	0.8387
OmniRe [11]	<u>38.83</u>	<u>527.26</u>	0.9711	<u>0.9355</u>	<u>48.05</u>	<u>580.17</u>	0.9619	<u>0.9238</u>	<u>70.35</u>	<u>598.01</u>	0.9707	<u>0.9159</u>
Ours	38.62	499.43	0.9770	0.9512	44.34	539.40	0.9633	0.9391	54.54	508.53	0.9732	0.9333

(c) Other Vehicle Manipulation II & Overall Assessment												
Method	Insertion					Removal					Overall	
	FID↓	FVD↓	VIMS↑	BAS↑	OSR↑	FID↓	FVD↓	VIMS↑	BAS↑	OSR↑	FID↓	FVD↓
StreetCrafter [59]	267.76	2558.27	0.9652	0.8276	4.62	92.19	1209.77	0.9785	0.8918	6.51	91.16	1245.96
Difix3D [55]	267.53	2122.44	0.9664	0.7824	<u>4.71</u>	122.80	786.16	<u>0.9823</u>	0.8526	<u>6.89</u>	80.84	991.23
OmniRe [11]	<u>182.29</u>	<u>1636.95</u>	0.9754	<u>0.8732</u>	4.23	<u>33.97</u>	<u>390.97</u>	0.9712	<u>0.9363</u>	6.42	<u>44.37</u>	<u>546.00</u>
Ours	117.46	1142.09	0.9784	0.8913	5.86	29.01	380.71	0.9884	0.9579	7.00	33.19	536.49

Table 1. Quantitative comparison on **HorizonSuite** with different methods. Metrics include FID↓, FVD↓, Vehicle Identity Matching Score (VIMS)↑, Background Alignment Score (BAS)↑ and Operation Success Rate (OSR)↑. Lower is better for ↓ and higher is better for ↑. The best results are highlighted in **bold**, and the second best results are underlined.




Figure 7. Qualitative Comparison among different methods. From top to bottom is Original scene, StreetCrafter [59], Difix3D [55], OmniRe [11] and **Ours**. It’s a operation of turning ego direction to left.

HorizonForge, integrating editable Gaussian splats and meshes with a video diffusion backbone, provides a simple yet powerful paradigm for photorealistic, physically reliable, and controllable driving scene generation.

4.2. Ablation Study

We conduct detailed ablations on **HorizonForge** using **HorizonSuite** to analyze the contribution of different conditioning modalities and model designs. Specifically, we compare (1) different 3D conditioning inputs for the diffusion model and (2) the effectiveness of using a video diffusion model (VDM) versus an image diffusion model (Image DM). The quantitative results are reported in Table 2.

For the first aspect, we compare our Gaussian–mesh conditioning with alternative 3D modalities, including colored point clouds [59] and bounding boxes (BBoxes) [44]. For

fair comparison, we combine all of them with meshes. The results show that Gaussian representation provides competitive geometric and photometric knowledge for high-quality, controllable scene generation. While colored point clouds capture accurate spatial positions and exhibit robustness to trajectory perturbations, their sparsity and lack of fine appearance cues limit their ability to guide diffusion synthesis at high fidelity. BBoxes, on the other hand, offer precise localization and motion information but fail to encode structural or visual context, resulting in unstable textures and poor alignment with the original scene. In contrast, Gaussian splats inherently encode both continuous density and color information, providing rich mid-level priors that effectively balance spatial precision and appearance detail, leading to the overall best performances across all manip-

(a) Ego Car Manipulation												
Method	Speed Change				Lane Change				Direction Change			
	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑
Point Cloud	62.16	869.06	0.9633	0.9051	81.65	877.87	0.9599	0.8946	152.10	1427.25	0.9593	0.8751
BBox	108.54	2105.98	0.9491	0.7986	125.19	1677.12	0.9469	0.7744	228.06	2662.92	0.9475	0.7663
Image DM	82.19	920.29	0.9655	0.8353	89.77	967.54	0.9607	0.8188	114.11	1180.18	0.9608	0.8063
Ours	39.92	583.78	0.9657	0.9326	50.33	729.76	0.9615	0.9303	144.83	1175.52	0.9595	0.8947

(b) Other Vehicle Manipulation I												
Method	Speed Change				Lane Change				Direction Change			
	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑	FID↓	FVD↓	VIMS↑	BAS↑
Point Cloud	69.47	844.01	0.9714	0.9077	74.80	905.17	0.9598	0.9023	84.89	840.43	0.9699	0.8975
BBox	100.61	1428.47	0.9584	0.8156	122.76	1994.33	0.9501	0.8001	124.87	1499.41	0.9587	0.8304
Image DM	88.99	870.43	0.9733	0.8416	89.44	1271.34	0.9636	0.8470	103.89	821.33	0.9720	0.8386
Ours	38.62	499.43	0.9770	0.9512	44.34	539.40	0.9633	0.9391	54.54	508.53	0.9732	0.9333

(c) Other Vehicle Manipulation II & Overall Assessment												
Method	Insertion					Removal					Overall	
	FID↓	FVD↓	VIMS↑	BAS↑	OSR↑	FID↓	FVD↓	VIMS↑	BAS↑	OSR↑	FID↓	FVD↓
Point Cloud	148.92	1665.92	0.9713	0.8633	6.07	56.55	763.86	0.9807	0.9139	6.68	54.14	813.67
BBox	166.21	2645.56	0.9621	0.7707	4.28	87.12	1351.49	0.9672	0.8203	6.89	81.74	1521.07
Image DM	176.96	1615.10	0.9758	0.8995	4.00	74.29	703.47	0.9842	0.8454	7.00	75.83	837.99
Ours	117.46	1142.09	0.9784	0.8913	5.86	29.01	380.71	0.9884	0.9579	7.00	33.19	536.49

Table 2. Ablation study on **HorizonSuite** with different settings. Metrics include FID↓, FVD↓, Vehicle Identity Matching Score (VIMS)↑, Background Alignment Score (BAS)↑ and Operation Success Rate (OSR)↑. Lower is better for ↓ and higher is better for ↑. The best results are highlighted in **bold**, and the second best results are underlined.

ulation categories. For the second aspect, we compare our full video diffusion model with a single-frame Image DM. We utilize Difix3D [55] as backbone and finetune it under the same conditioning setup with our method. Although the Image DM baseline achieves competitive fidelity, it lacks temporal awareness and thus struggles to maintain cross-frame consistency. This leads to noticeable temporal flickering. In contrast, our **HorizonForge** explicitly models temporal correlations, preserving coherent motion across frames and maintaining consistent visual features during complex multi-agent edits. These findings highlight that temporal priors are indispensable for realistic controllable video generation, and that combining Gaussian-mesh conditioning with a temporal diffusion backbone yields the most stable and photorealistic results.

4.3. User Study

Method	Wins	Total	Win Rate
StreetCrafter [59]	2	501	0.40%
Difix3D [55]	4	501	0.80%
OmniRe [11]	39	501	7.78%
Ours	456	501	91.02%

Table 3. Comparison of win rates among different methods.

To further assess the perceptual quality and controllability of generated videos, we conducted a comprehensive user study comparing under the same comparison setting

shown in Sec. 4.1. The study was designed to evaluate how well each method produces photorealistic and trajectory-consistent driving videos under identical editing tasks. Each participant was shown 20 video samples randomly drawn from **HorizonSuite**. For the design details, please refer to Appendix 9. As shown in Table 3, **HorizonForge** overwhelmingly outperformed other methods. These results strongly confirm that users consistently perceive **HorizonForge**’s outputs as more realistic, visually stable, and physically consistent with the intended trajectories, validating both its photorealism and controllability advantages.

5. Conclusion

This paper introduced **HorizonForge**, a simple yet effective framework for controllable driving scene generation that combines editable 3D Gaussian splats and meshes with a standard video diffusion model, enabling photorealistic and physically grounded multi-agent editing through trajectory-based control. **HorizonForge** supports fine-grained manipulation of both ego and surrounding vehicles and consistently outperforms reconstruction-based, generation-based, and hybrid baselines across all tasks in our newly proposed benchmark **HorizonSuite**. Our systematic study further highlights the effectiveness of Gaussian conditioning and the importance of temporal diffusion for stable, coherent scene synthesis. Together, **HorizonForge** and **HorizonSuite** provide a unified, scalable, and physically reliable foundation for future research in controllable driving simulation.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 4, 5, 6, 2
- [2] Hassan Abu Alhaija, Jose Alvarez, Maciej Bala, Tiffany Cai, Tianshi Cao, Liz Cha, Joshua Chen, Mike Chen, Francesco Ferroni, Sanja Fidler, et al. Cosmos-transfer1: Conditional world generation with adaptive multimodal control. *arXiv preprint arXiv:2503.14492*, 2025. 2
- [3] 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
- [4] 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 (CVPR)*, 2023. 2
- [5] Daniel Bogdoll, Yitian Yang, Tim Joseph, Melih Yazgan, and J Marius Zollner. Muvo: A multimodal generative world model for autonomous driving with geometric representations. In *2025 IEEE Intelligent Vehicles Symposium (IV)*, pages 2243–2250. IEEE, 2025. 3
- [6] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 2
- [7] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. *arXiv preprint arXiv:2106.11810*, 2021. 2
- [8] Wei-Jer Chang, Wei Zhan, Masayoshi Tomizuka, Manmohan Chandraker, and Francesco Pittaluga. Langtraj: Diffusion model and dataset for language-conditioned trajectory simulation. *arXiv preprint arXiv:2504.11521*, 2025. 6
- [9] 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 (CVPR)*, pages 7310–7320, 2024. 2
- [10] Jiacheng Chen, Ziyu Jiang, Mingfu Liang, Bingbing Zhuang, Jong-Chyi Su, Sparsh Garg, Ying Wu, and Manmohan Chandraker. Autoscape: Geometry-consistent long-horizon scene generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25700–25711, 2025. 3
- [11] 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, 2, 3, 4, 5, 6, 7, 8
- [12] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [13] 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. *arXiv preprint arXiv:2310.02601*, 2023. 3, 4
- [14] 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
- [15] 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. 3
- [16] Sylvain Gugger, Lysandre Debut, Thomas Wolf, Philipp Schmid, Zachary Mueller, Sourab Mangrulkar, Marc Sun, and Benjamin Bossan. Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>, 2022. 1
- [17] Mariam Hassan, Sebastian Stapf, Ahmad Rahimi, Pedro M. B. Rezende, Yasaman Haghighi, David Bruggemann, Isinsu Katircioglu, Lin Zhang, Xiaoran Chen, Suman Saha, et al. GEM: A generalizable ego-vision multimodal world model for fine-grained ego-motion, object dynamics, and scene composition control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 2020. 2
- [20] Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 4
- [21] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2312.08270*, 2023. 2
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42 (4), 2023. 1, 2, 4

- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 4
- [24] 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. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 11971–11981, 2025. 3
- [25] Jiusi Li, Jackson Jiang, Jinyu Miao, Miao Long, Tuopu Wen, Peijin Jia, Shengxiang Liu, Chunlei Yu, Maolin Liu, Yuzhan Cai, et al. Realistic and controllable 3d gaussian-guided object editing for driving video generation. *arXiv preprint arXiv:2508.20471*, 2025. 3
- [26] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: learning bird’s-eye-view representation from lidar-camera via spatiotemporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):2020–2036, 2024. 3
- [27] Yiyuan Liang, Zhiying Yan, Liqun Chen, Jiahuan Zhou, Luxin Yan, Sheng Zhong, and Xu Zou. Driveeditor: A unified 3d information-guided framework for controllable object editing in driving scenes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5164–5172, 2025. 3
- [28] Buyu Liu, Kai Wang, Yansong Liu, Jun Bao, Tingting Han, and Jun Yu. Mvpbev: Multi-view perspective image generation from bev with test-time controllability and generalizability. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 8393–8401, 2024. 3
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [30] Yifan Lu, Xuanchi Ren, Jiawei Yang, Tianchang Shen, Zhangjie Wu, Jun Gao, Yue Wang, Siheng Chen, Mike Chen, Sanja Fidler, et al. Infinicube: Unbounded and controllable dynamic 3d driving scene generation with world-guided video models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 27272–27283, 2025. 3
- [31] Zeyu Lu, Zidong Wang, Di Huang, Chengyue Wu, Xihui Liu, Wanli Ouyang, and Lei Bai. Fit: Flexible vision transformer for diffusion model. *CoRR*, abs/2405.06270, 2024. 2
- [32] Chaojun Ni, Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Wenkang Qin, Guan Huang, Chen Liu, Yuyin Chen, Yida Wang, Xueyang Zhang, Yifei Zhan, Kun Zhan, Peng Jia, Xianpeng Lang, Xingang Wang, and Wenjun Mei. Recondreamer: Crafting world models for driving scene reconstruction via online restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1559–1569, 2025. 2, 3, 5, 6, 4
- [33] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2856–2865, 2021. 2
- [34] Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3931–3940. IEEE Computer Society, 2024. 4
- [35] Sung-Yeon Park, Adam Lee, Juanwu Lu, Can Cui, Luyang Jiang, Rohit Gupta, Kyungtae Han, Ahmadreza Moradipari, and Ziran Wang. Simsplat: Predictive driving scene editing with language-aligned 4d gaussian splatting. *arXiv preprint arXiv:2510.02469*, 2025. 3
- [36] Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. One-step image translation with text-to-image models. *arXiv preprint arXiv:2403.12036*, 2024. 6
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 1
- [38] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 2
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 6, 1, 2
- [40] Xuanchi Ren, Yifan Lu, Tianshi Cao, Ruiyuan Gao, Shengyu Huang, Amirmojtaba Sabour, Tianchang Shen, Tobias Pfaff, Jay Zhangjie Wu, Runjian Chen, et al. Cosmos-drive-dreams: Scalable synthetic driving data generation with world foundation models. *arXiv preprint arXiv:2506.09042*, 2025. 4
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2022. 2
- [42] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [43] Cory Sullivan and Andrew Kaszynski. Pyvista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (vtk). *Journal of Open Source Software*, 4(37):1450, 2019. 5
- [44] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijay Patnaik, Philip Tsui, Jiyang 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 (CVPR)*, pages 2446–2454, 2020. 2, 4, 6, 7, 3

- [45] Shanlin Sun, Gabriel De Araujo, Jiaqi Xu, Shenghan Zhou, Hanwen Zhang, Ziheng Huang, Chenyu You, and Xiaohui Xie. Coma: Compositional human motion generation with multi-modal agents. *arXiv preprint arXiv:2412.07320*, 2024. 3
- [46] Shanlin Sun, Bingbing Zhuang, Ziyu Jiang, Buyu Liu, Xiaohui Xie, and Manmohan Chandraker. Lidarf: Delving into lidar for neural radiance field on street scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19563–19572, 2024. 2
- [47] Shanlin Sun, Yifan Wang, Hanwen Zhang, Yifeng Xiong, Qin Ren, Ruogu Fang, Xiaohui Xie, and Chenyu You. Ouroboros: Single-step diffusion models for cycle-consistent forward and inverse rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10386–10397, 2025. 2
- [48] Matthew Tancik, Vincent Casser, Xincheng Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretschmar. 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
- [49] 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
- [50] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 6
- [51] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 1
- [52] 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. 3
- [53] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jagang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-drive world models for autonomous driving. In *European conference on computer vision*, pages 55–72. Springer, 2024.
- [54] Yuxi Wei, Zi Wang, Yifan Lu, Chenxin Xu, Changxing Liu, Hao Zhao, Siheng Chen, and Yanfeng Wang. Editable scene simulation for autonomous driving via collaborative llm-agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15077–15087, 2024. 3
- [55] Jay Zhangjie Wu, Yuxuan Zhang, Haithem Turki, Xuanchi Ren, Jun Gao, Mike Zheng Shou, Sanja Fidler, Zan Gojic, and Huan Ling. Diffrax3D+: Improving 3d reconstructions with single-step diffusion models. *arXiv preprint arXiv:2503.01774*, 2025. 3, 5, 6, 7, 8
- [56] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuntao Chen, Runyi Yang, Yuxin Huang, Xiaoyu Ye, Zike Yan, Yongliang Shi, Yiyi Liao, and Hao Zhao. Mars: An instance-aware, modular and realistic simulator for autonomous driving. In *CICAI*, 2023. 2
- [57] Ziyang Xie, Junge Zhang, Wenye Li, Feihu Zhang, and Li Zhang. S-nerf: Neural radiance fields for street views. In *International Conference on Learning Representations*, 2023. 2
- [58] 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. 2
- [59] Yunzhi Yan, Zhen Xu, Haotong Lin, Haian Jin, Haoyu Guo, Yida Wang, Kun Zhan, Xianpeng Lang, Hujun Bao, Xiaowei Zhou, and Sida Peng. Streetcrafter: Street view synthesis with controllable video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 822–832, 2025. 3, 4, 6, 7, 8
- [60] Jiawei Yang, Boris Ivanovic, Or Litany, Xinshuo Weng, Seung Wook Kim, Boyi Li, Tong Che, Danfei Xu, Sanja Fidler, Marco Pavone, and Yue Wang. Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. *arXiv preprint arXiv:2311.02077*, 2023. 2
- [61] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1389–1399, 2023. 2
- [62] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazhen 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. 5, 6, 1
- [63] Chenyu You, Yifei Mint, Weicheng Dai, Jasjeet S Sekhon, Lawrence Staib, and James S Duncan. Calibrating multi-modal representations: A pursuit of group robustness without annotations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26140–26150, 2024. 3
- [64] Chenyu You, Haocheng Dai, Yifei Min, Jasjeet S Sekhon, Sarang Joshi, and James S Duncan. Uncovering memorization effect in the presence of spurious correlations. *Nature Communications*, 16(1):5424, 2025. 3
- [65] Mark Yu, Wenbo Hu, Jinbo Xing, and Ying Shan. Trajectorycrafter: Redirecting camera trajectory for monocular videos via diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 100–111, 2025. 5, 1
- [66] Jian Zhang, Yuanqing Zhang, Huan Fu, Xiaowei Zhou, Bowen Cai, Jinchi Huang, Rongfei Jia, Binqiang Zhao, and Xing Tang. Ray priors through reprojection: Improving neural radiance fields for novel view extrapolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18376–18386, 2022. 1
- [67] Shiwei Zhang, Jiayu Wang, Yingya Zhang, Kang Zhao, Hangjie Yuan, Zhiwu Qin, Xiang Wang, Deli Zhao, and Jin-

gren Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. *CoRR*, abs/2308.06792, 2023. [2](#)

- [68] Guosheng Zhao, Chaojun Ni, Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Boyuan Wang, Youyi Zhang, Wenjun Mei, and Xingang Wang. Drivedreamer4d: World models are effective data machines for 4d driving scene representation. *arXiv preprint arXiv:2410.13571*, 2024. [2](#), [3](#), [4](#)
- [69] Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Xinze Chen, Guan Huang, Xiaoyi Bao, and Xingang Wang. Drivedreamer-2: Llm-enhanced world models for diverse driving video generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10412–10420, 2025. [3](#)
- [70] Zibo Zhao, Zeqiang Lai, Qingxiang Lin, Yunfei Zhao, Haolin Liu, Shuhui Yang, Yifei Feng, Mingxin Yang, Sheng Zhang, Xianghui Yang, et al. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202*, 2025. [4](#), [5](#)

HorizonForge: Driving Scene Editing with Any Trajectories and Any Vehicles

Supplementary Material

6. Implementation Detail

We fine-tune our Video Diffusion Model (VDM) using the 5B CogVideoX backbone [62], initialized from the TrajectoryCrafter checkpoint [65]. All experiments are implemented with `pytorch` [37], `accelerate` [16] and `diffuser` [51]. Four H100 (80GB) GPUs are employed for training. It uses a per-GPU batch size of 1 with gradient accumulation of 1, yielding an effective batch size of 4. We adopt `bfloat16` mixed-precision training to reduce memory consumption while maintaining numerical stability.

We perform finetuning in two stages. In the first stage, we train for 40K iterations on clips of resolution $33 \times 384 \times 576$ (frames \times height \times width), requiring approximately three days. In the second stage, we train for an additional 20K iterations on longer clips of resolution $81 \times 384 \times 576$, which takes another three days.

Optimization. We use the AdamW optimizer with a learning rate of 2×10^{-5} , weight decay of 1×10^{-4} , and $\epsilon = 10^{-8}$. The learning rate follows a constant schedule with linear warmup over the first 100 steps.

7. Benchmark Design Details

In this section, we will introduce the benchmark design details. Table. 4 showcases the statistic of tasks within the benchmark.

Object Type	Manipulation Type	# Samples
Ego Vehicle	Direction Change	10
	Lane Change	15
	Speed Change	14
Other Vehicles	Direction Change	8
	Lane Change	13
	Speed Change	15
	Insertion	14
	Removal	20

Table 4. Statistics of manipulation categories in **HorizonSuite**. The benchmark with in total 109 samples covers two object types (Ego vehicle and Other vehicles) and five manipulation types, providing diverse scene editing scenarios for controllable driving generation evaluation.

7.1. Evaluation metrics formulation

In this section, we formally describe the three task-specific controllability metrics used in **HorizonSuite**: Vehicle Identity Matching Score (VIMS), Background Alignment Score (BAS), and Operation Success Rate (OSR).

7.1.1. Vehicle Identity Matching Score (VIMS)

VIMS measures how well the generated video preserves the identity of each vehicle instance along its trajectory, by aligning frames in 3D space and computing CLIP similarity [39] on instance-masked regions.

3D Frame Correspondence. For a given vehicle instance j , the modified trajectory provides a sequence of vehicle-to-world transforms

$$T_{t,j}^{\text{mod}} \in SE(3), \quad t = 1, \dots, T,$$

and ego (camera) poses

$$E_t^{\text{mod}} \in SE(3).$$

We first express the vehicle pose in the ego coordinate system by composing the transforms:

$$P_{t,j}^{\text{mod}} = (E_t^{\text{mod}})^{-1} T_{t,j}^{\text{mod}} = \begin{bmatrix} R_{t,j}^{\text{mod}} & p_{t,j}^{\text{mod}} \\ \mathbf{0}^\top & 1 \end{bmatrix}, \quad (8)$$

where $R_{t,j}^{\text{mod}} \in SO(3)$ and $p_{t,j}^{\text{mod}} \in \mathbb{R}^3$ denote, respectively, the rotation and position of vehicle j in ego coordinates at time t .

Similarly, from the original trajectory we obtain vehicle-to-world transforms $T_{k,j}^{\text{gt}} \in SE(3)$ and ego poses $E_k^{\text{gt}} \in SE(3)$ for $k = 1, \dots, K$, and we convert them to ego coordinates as

$$P_{k,j}^{\text{gt}} = (E_k^{\text{gt}})^{-1} T_{k,j}^{\text{gt}} = \begin{bmatrix} R_{k,j}^{\text{gt}} & p_{k,j}^{\text{gt}} \\ \mathbf{0}^\top & 1 \end{bmatrix}, \quad (9)$$

with $R_{k,j}^{\text{gt}} \in SO(3)$ and $p_{k,j}^{\text{gt}} \in \mathbb{R}^3$.

Given these ego-frame poses, we define the distance between the modified pose at time t and the ground-truth pose at time k as

$$D_{t,k,j} = \left\| p_{t,j}^{\text{mod}} - p_{k,j}^{\text{gt}} \right\|_2 + 0.1 \left\| R_{t,j}^{\text{mod}} - R_{k,j}^{\text{gt}} \right\|_F. \quad (10)$$

To avoid spurious matches due to heavy occlusions, we restrict the matching to frames in which vehicle j is visible. We mark vehicle j as occluded in frame k if more than 80% of its projected 2D bounding box is covered by other vehicles; otherwise it is treated as non-occluded. Let $\mathcal{K}_{t,j} \subseteq \{1, \dots, K\}$ denote the set of ground-truth frames in which vehicle j is non-occluded. The best-matching ground-truth frame for (t, j) is then

$$k^*(t, j) = \arg \min_{k \in \mathcal{K}_{t,j}} D_{t,k,j}. \quad (11)$$

Instance-Masked CLIP Similarity. Let I_t^{gen} and $I_{k^*(t,j)}^{\text{gt}}$ denote the generated and ground-truth frames, respectively.

For instance j at time t , we use its binary masks $M_{t,j}^{\text{gen}}$ and $M_{k^*(t),j}^{\text{gt}}$ to extract instance-only regions:

$$\tilde{I}_{t,j}^{\text{gen}} = I_t^{\text{gen}} \odot M_{t,j}^{\text{gen}}, \quad (12)$$

$$\tilde{I}_{t,j}^{\text{gt}} = I_{k^*(t),j}^{\text{gt}} \odot M_{k^*(t),j}^{\text{gt}}, \quad (13)$$

where \odot denotes element-wise masking.

We then compute CLIP image embeddings $f(\cdot)$ and cosine similarity:

$$s_{t,j} = \frac{\langle f(\tilde{I}_{t,j}^{\text{gen}}), f(\tilde{I}_{t,j}^{\text{gt}}) \rangle}{\|f(\tilde{I}_{t,j}^{\text{gen}})\|_2 \|f(\tilde{I}_{t,j}^{\text{gt}})\|_2}. \quad (14)$$

For a single scene, we aggregate over all valid (t, j) pairs:

$$\text{VIMS}_{\text{scene}} = \frac{1}{N_{\text{inst}}} \sum_j \frac{1}{T_j} \sum_{t \in \mathcal{T}_j} s_{t,j}, \quad (15)$$

where \mathcal{T}_j is the set of frames where instance j is visible and N_{inst} is the number of evaluated instances. In the benchmark, we report the average across scenes:

$$\text{VIMS} = \frac{1}{N_{\text{scene}}} \sum_{n=1}^{N_{\text{scene}}} \text{VIMS}_{\text{scene},n}. \quad (16)$$

7.1.2. Background Alignment Score (BAS)

BAS evaluates how well the static background is preserved after editing, by explicitly masking out all dynamic instances using masks and computing CLIP similarity [39] on the remaining background regions.

Ego-Pose-Based Frame Matching. For background evaluation, we match frames based on ego poses. Let

$$E_t^{\text{mod}}, E_k^{\text{gt}} \in SE(3)$$

be the ego poses at modified frame t and ground-truth frame k , respectively. We define a combined distance

$$D(t, k) = \|x_t^{\text{mod}} - x_k^{\text{gt}}\|_2 + 0.1 \|R_t^{\text{mod}} - R_k^{\text{gt}}\|_F, \quad (17)$$

where x is the translation component and R is the 3×3 rotation matrix extracted from E , and $\|\cdot\|_F$ denotes the Frobenius norm. The best-matching ground-truth frame for background at time t is

$$k^*(t) = \arg \min_k D(t, k). \quad (18)$$

Background Masking. Let M_t^{mod} and $M_{k^*(t)}^{\text{gt}}$ be the binary foreground masks for the generated and ground-truth videos, respectively. We invert them to obtain background masks:

$$BM_t^{\text{mod}} = 1 - M_t^{\text{mod}}, \quad BM_{k^*(t)}^{\text{gt}} = 1 - M_{k^*(t)}^{\text{gt}}. \quad (19)$$

We then extract background-only regions:

$$\tilde{I}_t^{\text{gen,bg}} = I_t^{\text{gen}} \odot BM_t^{\text{mod}}, \quad (20)$$

$$\tilde{I}_t^{\text{gt,bg}} = I_{k^*(t)}^{\text{gt}} \odot BM_{k^*(t)}^{\text{gt}}. \quad (21)$$

CLIP Similarity and Aggregation. We compute CLIP embeddings and cosine similarity:

$$s_t^{\text{bg}} = \frac{\langle f(\tilde{I}_t^{\text{gen,bg}}), f(\tilde{I}_t^{\text{gt,bg}}) \rangle}{\|f(\tilde{I}_t^{\text{gen,bg}})\|_2 \|f(\tilde{I}_t^{\text{gt,bg}})\|_2}. \quad (22)$$

For a single scene:

$$\text{BAS}_{\text{scene}} = \frac{1}{T} \sum_{t=1}^T s_t^{\text{bg}}, \quad (23)$$

and the benchmark reports the average across all scenes.

7.1.3. Operation Success Rate (OSR)

OSR is used for *insertion* and *removal* tasks, and measures whether the intended operation has been successfully carried out, using GPT-4o [1] as a learned judge. Our evaluation code follows a two-step process: (1) preparing a small set of annotated frames, and (2) querying GPT-4o for a scalar score in [1, 10].

Frame Sampling and Bounding-Box Annotation. For a given insertion or removal scene, we uniformly sample a small set of frame indices

$$\mathcal{S} = \{t_1, \dots, t_K\}, \quad (24)$$

typically $K = 5$.

- **Insertion:** We locate the inserted object's instance ID and load its binary mask $M_{t_k}^{\text{ins}}$ on each sampled frame. We then draw a *green* bounding box around the support of this mask on $I_{t_k}^{\text{gen}}$.
- **Removal:** We locate the removed agent ID and corresponding original instance, load the mask $M_{t_k}^{\text{rem}}$ for the original video, and draw a *red* bounding box on the edited frame indicating where the object used to be.

The resulting annotated frames $\{\hat{I}_{t_k}\}_{k=1}^K$ are sent to GPT-4o together with a task-specific textual instruction.

GPT-4o Prompting. We use a concise, rating-oriented prompt.

Insertion prompt:

You are evaluating a generated driving video for object insertion quality. The goal was to insert: *<vehicle description from metadata>*.

The inserted object is marked with a **green** bounding box in the frames.

Please rate the generated video on a scale of 1–10 based on:

1. Realism: Does the inserted object (in green box) look realistic and well-integrated?
2. Consistency: Is the object consistent across frames?
3. Alignment: Does it match the text description?
4. Physical plausibility: Does it follow realistic motion and positioning?
5. Visual quality: Are there any artifacts or inconsistencies in the boxed region?

Provide ONLY a single number from 1 to 10 as your response.

Removal prompt:

You are evaluating a generated driving video for object removal quality. The goal was to remove: *<vehicle description from metadata>*.

The inserted object is marked with a **red** bounding box in the frames.

Please rate the generated video on a scale of 1–10 based on:

1. Completeness: Is the object fully removed from the red box area?
2. Naturalness: Does the background look natural where the object was (in red box)?
3. Consistency: Is the removal consistent across frames?
4. Artifacts: Are there visible artifacts or inconsistencies in the boxed region?
5. Inpainting quality: How well is the removed region filled in?

Provide ONLY a single number from 1 to 10 as your response.

GPT-4o returns a string which is parsed as a scalar score

$$g_i \in [1, 10] \quad (25)$$

for scene i .

Aggregation. We define OSR as the average GPT score over all scenes:

$$\text{OSR} = \frac{1}{N_{\text{ops}}} \sum_{i=1}^{N_{\text{ops}}} g_i, \quad (26)$$

where N_{ops} is the number of evaluated insertion or removal scenes. This gives a human-aligned, continuous measure of operation success on a 1–10 scale.

8. Downstream Tasks

To evaluate the utility of our generated data for downstream perception, we conduct a 3D object detection study using BEVFormer [26] on the Waymo Open Dataset [44]. Following the protocol of ChatSim [54], we mix the real

training set with simulated data and compare performance before and after mixing. We start from 8000 real training frames and generate 8000 corresponding simulated versions. For each sequence, we produced two edited variants with our framework: (i) **Insertion (Ins)**: randomly inserting four novel vehicles as replacements for existing ones; (ii) **Removal (Rem)**: randomly removing vehicles from the scene. Each operation contributes an additional 8000 edited frames.

We evaluate BEVFormer on the full Waymo validation set using mean Average Precision (AP) at IoU thresholds of 0.3, 0.5, and 0.7.

Training Data	AP@0.3	AP@0.5	AP@0.7
Org	0.1211	0.0533	0.0103
Org + Ins	0.1343	0.0635	0.0125
Org + Rem	0.1359	0.0613	0.0134

Table 5. Downstream 3D detection results of BEVFormer [26] on the Waymo validation set under a front-camera-only configuration. **Org** denotes the original 8000 real frames. **Ins** and **Rem** denote additional frames generated using our method with random vehicle insertion and removal, respectively.

As shown in Table 5, augmenting the training set with our edited scenes leads to *substantial* relative gains across all IoU thresholds. Adding insertion-based edits (**Org + Ins**) improves AP from 0.1211 to 0.1343 at IoU 0.3 (+10.9% relative), from 0.0533 to 0.0635 at IoU 0.5 (+19.1%), and from 0.0103 to 0.0125 at IoU 0.7 (+21.4%). Removal-based edits (**Org + Rem**) provide similarly strong or even larger gains, boosting AP to 0.1359, 0.0613, and 0.0134 at IoU 0.3/0.5/0.7, corresponding to +12.2%, +15.0%, and +30.1% relative improvements over the **Org** baseline. These large relative gains demonstrate that our generated data is not only visually realistic but also highly informative, and that controllable scene editing with our framework can effectively enrich training distributions for autonomous driving perception.

9. User Study Design

In this section we will introduce the design of our user study. For every sample, participants were provided with: (1) the original video and trajectory visualization before editing, (2) the target trajectory visualization specifying the desired vehicle behavior and (3) four edited results from different methods, denoted as methods A, B, C, and D. The participant’s task was to select the one video that best aligns with the target trajectory and maintains the highest visual realism. To ensure fairness and eliminate potential bias or cherry-picking, the sampling and ordering were fully randomized: each participant received a unique questionnaire

where both the 20 selected samples and the order of methods (A–D) were shuffled independently. This guarantees that no two participants saw identical question sets, and that every video sample was equally likely to appear across users.

10. Fixed-Offset Shifting Performance

We further evaluate fixed-offset shifting, as shown in Table 6, where our method notably outperforms ReconDreamer [32] and DriveDreamer4D [68] in NTA-IoU, NTL-IoU and FID, proving the superior positional control accuracy and robustness of the proposed method.

Method	NTA-IoU (\uparrow)		NTL-IoU (\uparrow)		FID (\downarrow)	
	Lane Shift @ 3m	Lane Shift @ 6m	Lane Shift @ 3m	Lane Shift @ 6m	Lane Shift @ 3m	Lane Shift @ 6m
DriveDreamer4D	0.340	0.121	51.32	49.28	129.05	210.37
ReconDreamer	0.539	0.467	54.58	52.58	93.56	149.19
Ours	0.552	0.523	56.73	54.01	49.47	53.61

Table 6. Comparison on large-offset manipulation.

11. Challenging Inputs Performance

As shown in Figure 8, our method demonstrates robust performance under challenging inputs: correcting vehicle artifacts caused by partial occlusion in original views, addressing severe sky artifacts, and naturally rendering dynamic headlight illumination during nighttime turns—proving our capability to handle poor-quality inputs, heavy occlusion, or dynamic background.



Figure 8. Visualization under challenging conditions.

12. OSR Reliability

We validated OSR reliability through a human evaluation using the identical 1-10 rubric. As shown in Table 7, this alignment confirms OSR serves as a reliable metric.

Metric	Ours		Difix3D		OmniRe		StreetCrafter	
	Ins.	Rem.	Ins.	Rem.	Ins.	Rem.	Ins.	Rem.
OSR	5.86	7.00	4.71	6.89	4.23	6.42	4.62	6.51
Human	7.50	7.25	3.79	4.40	3.58	4.75	3.12	3.25

Table 7. Comparison of OSR and Human Scores.

13. More visualization results

For more visualization results, please check the webpage.