

Extrapolated Urban View Synthesis Benchmark

Xiangyu Han^{1,3*} Zhen Jia^{1*} Boyi Li² Yan Wang² Boris Ivanovic² Yurong You²
Lingjie Liu³ Yue Wang^{2,4} Marco Pavone^{2,5} Chen Feng¹ Yiming Li^{1,2†}

¹NYU ²NVIDIA ³UPenn ⁴USC ⁵Stanford

<https://ai4ce.github.io/EUVS-Benchmark>

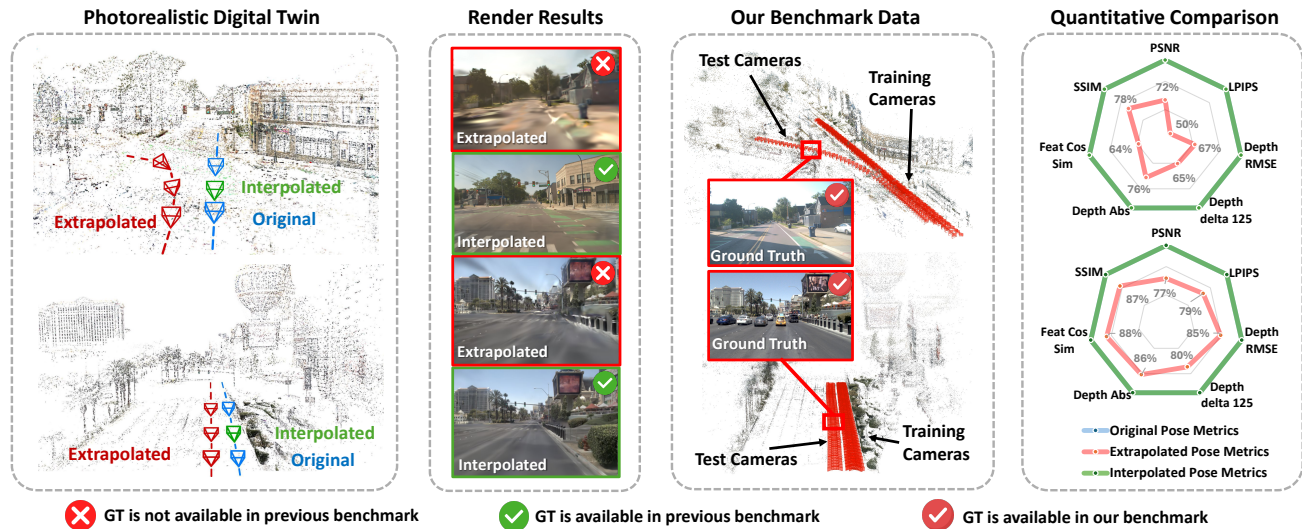


Figure 1. **Our key contributions.** Previous evaluations for urban view synthesis have primarily focused on interpolated poses, as the lack of ground truth data has made it challenging to evaluate extrapolated poses. We address this gap by providing real-world data that enables both quantitative and qualitative evaluations of extrapolated view synthesis in urban scenes. The quantitative results reveal a significant performance drop in 3D Gaussian Splatting [25] when handling extrapolated views, highlighting the need for more robust NVS methods.

Abstract

Photorealistic simulators are essential for the training and evaluation of vision-centric autonomous vehicles (AVs). At their core is Novel View Synthesis (NVS), a crucial capability that generates diverse unseen viewpoints to accommodate the broad and continuous pose distribution of AVs. Recent advances in radiance fields, such as 3D Gaussian Splatting, achieve photorealistic rendering at real-time speeds and have been widely used in modeling large-scale driving scenes. However, their performance is commonly evaluated using an interpolated setup with highly correlated training and test views. In contrast, extrapolation, where test views largely deviate from training views, remains underexplored, limiting progress in generalizable simulation technology. To address this gap, we leverage publicly available AV datasets with multiple traversals,

multiple vehicles, and multiple cameras to build the first *Extrapolated Urban View Synthesis (EUVS) benchmark*. Meanwhile, we conduct both **quantitative** and **qualitative** evaluations of state-of-the-art NVS methods across different evaluation settings. Our results show that current NVS methods are prone to overfitting to training views. Besides, incorporating diffusion priors and improving geometry cannot fundamentally improve NVS under large view changes, highlighting the need for more robust approaches and large-scale training. We have released the data to help advance self-driving and urban robotics simulation technology.

1. Introduction

The development of vision-centric autonomous vehicles (AVs) relies heavily on photorealistic simulators, which provide controlled, reproducible, and scalable environments for training and evaluation of driving models [13, 44, 56]. These simulators enable AVs to learn and adapt to a variety

*Equal contribution.

†Corresponding author.

of real-world scenarios, from crowded urban streets to adverse weather conditions, without the logistical and safety concerns of physical road testing. At the heart of these simulators is the capability for Novel View Synthesis (NVS)—a key technology that generates realistic images of unseen viewpoints, simulating the continuous changes in perspective that occur as AVs navigate through urban environments.

Recent advancements in radiance fields, particularly methods based on 3D Gaussian Splatting [25], have significantly improved the realism and efficiency of NVS. These approaches [8, 53, 64, 66] can produce photorealistic renderings at real-time speeds, making them highly attractive for large-scale driving scene simulation. However, despite their impressive results, the evaluation of NVS methods has predominantly focused on **interpolated** scenarios, where training and test viewpoints are closely related. While interpolation tests are valuable for assessing local consistency, they fall short in addressing the more critical challenge of **extrapolation**—where test viewpoints differ significantly from the training data. As shown in Figure 1, the interpolation test set demonstrates strong performance, with metrics such as PSNR, SSIM, and LPIPS remaining very close to the training set values. In contrast, the extrapolation test set, which includes additional translation and rotation changes relative to the training set, exhibits notable drops in performance. Specifically, the metric decreases relative to the training set are **28%** for PSNR, **22%** for SSIM, and **50%** for LPIPS. These results underscore the urgent need to explore and advance extrapolated view synthesis in complex urban scenes, as real-world driving often involves encountering scenarios with significant viewpoint shifts and diverse spatial transformations that deviate from training distributions. Several recent studies [20, 23] have investigated the generalization capabilities of NVS in 3D Gaussian Splatting. Although they show promising qualitative results, there is **no comprehensive quantitative analysis** due to the absence of standardized datasets. Moreover, their evaluations are primarily limited to specific scenarios or use cases, without investigating varying evaluation settings based on the degree of extrapolation. This gap underscores the urgent need for a benchmark that offers diverse and challenging datasets, enabling a rigorous and systematic assessment of NVS methods.

To establish a common platform for assessing the robustness of NVS methods, we introduce a comprehensive benchmark for quantitatively and qualitatively evaluating extrapolated novel view synthesis in large-scale urban scenes. Our benchmark leverages publicly available datasets, including NuPlan [4], MARS [30], and Argoverse2 [46], which feature multi-traversal, multi-agent and multi-camera sensory recordings. Multi-traversal data consists of asynchronous traversals of the same location, while multi-agent data refers to data collected simultaneously

from multiple vehicles within the same area. These data provide diverse camera poses within a 3D scene, enabling the training and evaluation of extrapolated view synthesis in outdoor environments. For the experimental setup, we define three evaluation settings: (1) translation only, (2) rotation only, and (3) translation + rotation, as shown in Figure 4. In autonomous driving scenarios, Setting 1 corresponds to maneuvers such as lane changes, Setting 2 involves switching between cameras facing different directions, and Setting 3 addresses complex intersections, such as crossroads with diverse traversal paths. These settings represent common challenges in autonomous driving, and addressing them enables the synthesis of complete scenes from sparse image observations.

We conduct pose estimation and sparse reconstruction using COLMAP [39], which facilitates the initialization of Gaussian Splatting. We then evaluate state-of-the-art Gaussian Splatting-based approaches across each evaluation setting, identifying performance gaps both qualitatively and quantitatively in extrapolated urban view synthesis.

In summary, our main contributions are as follows:

- We initiate the first comprehensive quantitative and qualitative study on the Extrapolated Urban View Synthesis (EUVS) problem, supported by a robust evaluation framework that categorizes evaluation settings (translation-only, rotation-only, and translation + rotation) while assessing performance using diverse metrics, including reconstruction accuracy and visual fidelity.
- We construct a novel dataset by integrating multi-traversal, multi-agent, and multi-camera data from publicly available resources, totaling **90,810** frames across **345** videos. Our dataset effectively addresses the limitations of existing benchmarks, enabling rigorous and robust evaluation for extrapolated urban view synthesis.
- We benchmark state-of-the-art Gaussian Splatting-based and NeRF-based models and analyze key factors that influence the performance of extrapolated NVS, laying a solid foundation for future advancements in this challenging task. Data and code will be released upon acceptance.

2. Related Works

Extrapolated View Synthesis. Extrapolated view synthesis aims to generate novel views beyond observed perspectives, addressing challenges in visual coherence for unseen regions. RapNeRF [60] proposes a random ray-casting policy that enables training on unseen views based on visible ones. Following work [54] enhances this approach by incorporating holistic priors. Besides, some generalizable models [6, 9, 45] have emerged, capable of generating extrapolated novel views from a limited number of input images. While these methods are designed for indoor scenes, several works address extrapolated view synthesis in outdoor driving scenarios, which typically involve forward-facing

cameras and unbounded environments. To tackle the Setting 1 challenge in our benchmark and address the scarcity of lane change data, GGS [20] introduces a novel virtual lane generation module. In parallel, AutoSplat [26] tackles lane change in dynamic scenes by applying geometric and reflected consistency constraints. To address the Setting 1 or 2 challenge, FreeSim [15], VEGS [23], and SGD [58] enhance 3DGS [25] with diffusion priors to improve generalization ability. *Yet existing methods suffer from two major limitations: (1) a lack of real data for quantitative evaluation, which confines them to quantitative analysis, and (2) a narrow focus on a specific setting in our benchmark, preventing a comprehensive and systematic exploration.*

3D Gaussian Splatting. Recent advances in radiance fields, particularly NeRF [34] and 3DGS [25], have garnered significant attention due to their impressive advancements in NVS. NeRF employs an implicit representation through a multi-layer perceptron (MLP). Furthermore, 3DGS explicitly represents scenes using anisotropic 3D Gaussian ellipsoids, enabling high-quality real-time rendering. Several works have addressed issues such as difficulties with reflective surfaces [24], aliasing [57], etc. However, urban scenes introduce unique challenges due to their unbounded and dynamic nature. To address the challenge, several works separate dynamic and static elements in the scene by leveraging a composite dynamic Gaussian graph [53, 66], optical flow prediction [55, 64], etc. PVG [8] presents a unified representation model that simultaneously incorporates both dynamic and static components without relying on priors. To achieve realistic geometry and efficient rendering, 2DGS [22] collapses 3D Gaussians onto 2D planes, while hybrid approaches [41, 48] combine different Gaussians to better capture region-specific features. *In summary, current urban NVS methods primarily focus on effectively handling dynamic elements and enhancing geometry representation, while the challenge of extrapolated view synthesis remains largely underexplored.*

Autonomous Driving Simulators. Current simulators focus on three key challenges: parameter initialization [16, 42], traffic simulation [28, 51, 63], and sensor simulation [1, 13, 18, 27, 56, 62]. Sensor simulation is crucial for generating realistic sensory data that AVs depend on for perception and decision-making. Early sensor simulators [13, 40, 50] provide simulated environments that are valuable for research but lack visual realism. Recent studies have focused on data-driven simulators that extract data from real-world driving logs, creating more realistic and adaptable environments. These methods can be classified into two categories: generation-based and reconstruction-based approaches. The former rely on inputs such as text, video, and other data sources for simulation, supported by world models or prior knowledge [18, 27, 62]. Reconstruction-based simulations leverage real-world data

to ensure both visual fidelity and geometric consistency [47, 49, 56]. UniSim [56] is a pioneering example of this approach, utilizing NeRF-based scene representation to create dynamic scenes with geometric information that are both editable and controllable. *Extrapolated view synthesis is essential for these simulators, as it enables the generation of realistic and consistent views from diverse angles.*

Autonomous Driving Datasets. High-quality datasets play a vital role in advancing autonomous driving research. The KITTI dataset [19], released in 2012, marked a major milestone, significantly accelerating advances in AVs [17, 33, 43]. Since then, many influential autonomous vehicle datasets have been developed to tackle challenges like adverse weather conditions [37], multimodal fusion [3, 4], repeated driving [4, 12], collaborative driving [29, 30, 52], and motion prediction [5, 14, 46], etc. We leverage publicly available datasets with multi-traversal, multi-agent, and multi-camera recordings, enabling a comprehensive and robust evaluation of extrapolated urban view synthesis.

3. The EUVS Benchmark

3.1. Dataset Curation

We leverage three publicly available, community-verified autonomous driving datasets—nuPlan [4], Argoverse 2 [46], and MARS [30]—to enhance adoption and foster trust through their established reliability. nuPlan [4] provides 1,200 hours of driving data from four cities, serving as the first large-scale planning benchmark. Argoverse 2 [46] supports multimodal perception and forecasting with 1,000 annotated 3D scenarios, 20,000 unlabeled lidar sequences, and 250,000 motion forecasting cases. MARS [30] enables collaborative driving research with multi-agent and multi-traversal scenarios. The original datasets were not designed for evaluating extrapolated view synthesis, which requires our significant labor in data processing—**300+ hours** of manual traversal selection and **800+ hours** of computing to run COLMAP. By integrating these datasets, our benchmark enables the evaluation of view synthesis across diverse and realistic urban environments under varying conditions. [Figure 3](#) illustrates the distribution of the integrated datasets.

3.2. Evaluation Framework

To systematically assess model performance in extrapolated urban view synthesis, our evaluation framework incorporates *three evaluation settings* and *three data configurations*. Data configurations include multi-traversal, multi-agent, and multi-camera, while evaluation settings are categorized into (1) translation only, (2) rotation only, and (3) translation + rotation, as illustrated in [Figure 4](#).

Setting 1. The translation-only experimental setup involves scenarios where the vehicle’s position shifts without any change in orientation. This scenario is commonly observed



Figure 2. **Dataset visualization.** Our dataset features diverse scenes across various locations in different cities, sourced from multiple datasets. Typical driving scenarios include maneuvers such as lane changes, cross intersections, and T-junctions. **Top:** Each column displays images captured at the same location by different agents or traversals. **Bottom:** Each image displays the COLMAP points at a specific location, along with the corresponding camera poses.




Figure 3. **Dataset distribution.** Our dataset comprises **90,810** frames distributed over **104** cases, capturing a diverse array of multi-traversal paths, multi-agent interactions, and multi-camera perspectives across varying evaluation settings.

in lane changes. We use traversals from different lanes in multi-traversal data, focusing on the three front cameras. The data is sourced from nuPlan [4] and Argoverse 2 [46].

Setting 2. The second setting, rotation only, evaluates models on views with significant orientation changes. In vision-centric autonomous vehicles, this corresponds to transitions between cameras capturing different directions. We leverage multi-camera data from nuPlan [4], training on three forward-facing and three rear-facing cameras, while evaluating on two side-facing cameras.

Setting 3. The third setting, combining translation and rotation, includes both positional shifts and orientation changes, posing the greatest challenge for NVS. To address this, we utilize multi-traversal driving data collected from the same location but across different traversal routes. For example, the training and test sets may include routes that approach an intersection from different directions. Typical route combinations feature scenarios such as intersections,

T-junctions, and Y-junctions, as shown in Figure 2, ensuring diversity and comprehensive evaluation. The data for this setting is from MARS [30] and Argoverse 2 [46].

3.3. Algorithm Overview

Vanilla 3D Gaussian Splatting. 3D Gaussian Splatting (3DGS) [25] leverage 3D Gaussians to explicitly represent the scene, which achieves high quality while offering real-time rendering by avoiding unnecessary computation in the empty space. Building on this, 3DGM [31] leverages multi-traversal consensus to differentiate transient and permanent elements, enabling joint 2D segmentation and 3D mapping without using any human supervision.

Planar-based and Geometry Refined Gaussian Splatting. GaussianPro [11] builds on 3DGS [25] by introducing multi-frame geometric optimization, which guides the densification of 3D Gaussians, enhancing scene consistency in complex geometries. It further refines geometry by encour-

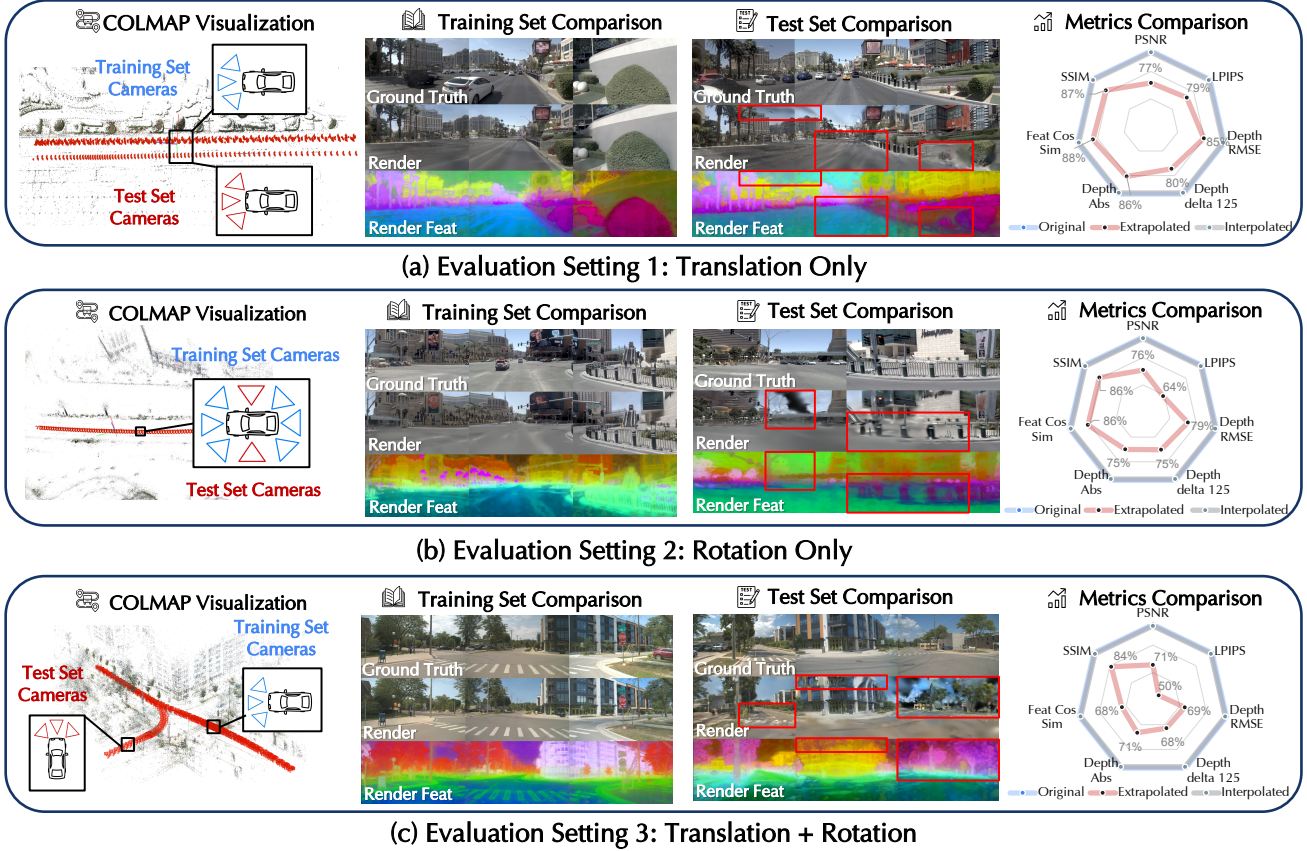


Figure 4. **Qualitative and quantitative results across three evaluation settings.** The performance drop from interpolation to extrapolation is significant in both qualitative and quantitative comparison. Different testing settings have distinct scenario characteristics, enabling the evaluation of a method’s capabilities from various aspects, thus systematically assessing the overall performance of reconstruction algorithms, including geometric accuracy, hallucination ability, view consistency, and depth precision, etc.

aging Gaussian primitives to adopt flat structures. Similarly, 2DGS [22] projects the 3D volume into a set of 2D oriented planar Gaussian disks, enabling high-fidelity surface reconstruction. PGSR [7] introduces an unbiased depth rendering method and integrates single-view geometric, multi-view photometric, and geometric regularization techniques to improve global geometry accuracy.

Gaussian Splatting with Diffusion Priors. VEGS [23] introduces a novel view generalization approach that harnesses pre-extracted surface normals to align 3D Gaussians while generating augmented camera views guided by diffusion priors. These diffusion priors serve a dual purpose: providing denoising loss guidance and supervising the training of augmented cameras. This process effectively mitigates floating artifacts and fragmented geometries, resulting in more accurate and coherent 3D representations.

Feature-Enhanced Gaussian Splatting. Feature 3DGS [65] extends 3D Gaussian Splatting with a Parallel N-dimensional Gaussian Rasterizer, allowing simultaneous rendering of radiance fields and high-dimensional semantic features. By embedding semantic features directly into 3D

Gaussians, the approach enhances optimization, enabling better correspondence with scene semantics and achieving more detailed and accurate spatial representations.

NeRF-based Method. Instant-NGP [35] uses a multiresolution hash encoding to map spatial coordinates into compact latent representations via hash tables. It efficiently encodes high-frequency details by combining trainable feature vectors with interpolation, enabling adaptive and scalable input encodings. Zip-NeRF [2] leverages multisampling with isotropic Gaussians for scale-aware features and introduces a smooth anti-aliasing loss to address z-aliasing. In addition, it incorporates a novel distance normalization technique to better manage close and distant objects, achieving high-quality rendering and fast training.

4. Experiment

4.1. Experiment Setup

Implementation Details. All 3DGS-based methods are initialized only from sparse points obtained by COLMAP [39] and exclude lidar points. We employ Grounded-SAM-2

Table 1. **Quantitative rendering results across three evaluation settings.** *In.* denotes interpolation, while *Ex.* represents extrapolation. Different baselines excel in different settings, reflecting the comprehensiveness and completeness of the evaluation protocol.

Method		PSNR \uparrow			SSIM \uparrow			LPIPS \downarrow			Feat Cos Sim \uparrow		
		In.	Ex.	Drop	In.	Ex.	Drop	In.	Ex.	Drop	In.	Ex.	Drop
Setting 1	3DGS [25]	21.36	16.37	23.4%	0.8275	0.7203	13.0%	0.2041	0.2599	27.3%	0.6828	0.6039	11.6%
	3DGM [31]	20.96	16.35	22.0%	0.8293	0.7248	12.6%	0.2003	0.2542	26.9%	0.6802	0.6087	10.5%
	GSPro [11]	21.51	16.39	23.8%	0.8310	0.7189	13.5%	0.1804	0.2450	35.8%	0.7081	0.6130	13.4%
	VEGS [23]	21.26	15.88	25.3%	0.8107	0.7047	13.1%	0.2498	0.3062	22.6%	0.6323	0.5521	12.7%
	PGSR [7]	20.57	16.32	20.7%	0.8104	0.7102	12.4%	0.2262	0.2733	20.8%	0.6515	0.5848	10.2%
	2DGS [22]	20.87	16.30	21.9%	0.8076	0.7103	12.0%	0.2438	0.2890	18.5%	0.6256	0.5644	9.8%
	Feature 3DGS [65]	21.02	16.01	23.8%	0.8096	0.7243	10.5%	0.1876	0.2575	37.3%	0.6958	0.6122	12.0%
	Zip-NeRF [2]	19.68	14.06	28.6%	0.7856	0.6917	12.0%	0.2711	0.3418	26.1%	0.6318	0.5542	12.3%
	Instant-NGP [35]	18.77	12.65	32.6%	0.7631	0.6252	18.1%	0.4874	0.5938	21.8%	0.5465	0.4837	11.5%
AVERAGE		20.67	15.59	24.6%	0.8083	0.7046	12.8%	0.2501	0.3134	25.3%	0.6505	0.5774	11.2%
Setting 2	3DGS [25]	25.75	19.53	24.2%	0.8766	0.7511	14.3%	0.1536	0.2668	73.7%	0.7327	0.6319	13.8%
	3DGM [31]	25.75	18.78	27.1%	0.8786	0.7464	15.0%	0.1556	0.2813	80.8%	0.7278	0.6344	12.8%
	GSPro [11]	26.42	19.39	26.6%	0.8821	0.7470	15.3%	0.1329	0.2246	69.0%	0.7523	0.6487	13.8%
	VEGS [23]	24.54	23.33	4.9%	0.8366	0.7949	5.0%	0.2301	0.2811	22.2%	0.6595	0.6133	7.0%
	PGSR [7]	24.53	18.38	25.1%	0.8612	0.7119	17.3%	0.1555	0.2532	62.8%	0.7200	0.5817	19.2%
	2DGS [22]	25.15	18.83	25.1%	0.8578	0.7204	16.0%	0.1756	0.2917	66.1%	0.6898	0.5785	16.1%
	Feature 3DGS [65]	24.91	19.59	21.4%	0.8800	0.7864	10.6%	0.1377	0.2278	65.4%	0.7427	0.6464	13.0%
	Zip-NeRF [2]	29.06	17.36	40.3%	0.8660	0.6715	22.5%	0.2078	0.3582	72.4%	0.7479	0.5843	21.9%
	Instant-NGP [35]	25.61	17.15	33.0%	0.8596	0.7212	16.1%	0.3340	0.5171	54.8%	0.7254	0.6182	14.8%
AVERAGE		25.75	19.15	25.6%	0.8665	0.7390	14.7%	0.1870	0.3002	62.5%	0.7220	0.6153	14.8%
Setting 3	3DGS [25]	21.22	14.99	29.4%	0.8550	0.7169	16.1%	0.2252	0.4050	79.8%	0.7002	0.4774	31.8%
	3DGM [31]	20.62	14.60	29.2%	0.8543	0.7233	15.3%	0.2254	0.4049	79.6%	0.6874	0.4672	32.0%
	GSPro [11]	21.58	14.82	31.3%	0.8634	0.6996	19.0%	0.2010	0.3877	92.9%	0.7093	0.4541	36.0%
	VEGS [23]	21.13	14.25	32.6%	0.8266	0.6475	21.7%	0.2359	0.4422	87.5%	0.6785	0.4442	34.5%
	PGSR [7]	19.60	14.17	27.7%	0.8238	0.6984	15.2%	0.2934	0.4363	48.7%	0.5867	0.3787	35.5%
	2DGS [22]	17.35	11.36	34.5%	0.7568	0.5447	28.0%	0.4296	0.5459	27.1%	0.3552	0.2327	34.5%
	Feature 3DGS [65]	21.88	14.33	34.5%	0.8643	0.6386	26.1%	0.1400	0.3816	172.6%	0.7411	0.4669	37.0%
	Zip-NeRF [2]	20.61	14.42	30.0%	0.8383	0.6565	21.7%	0.2197	0.4546	106.9%	0.7108	0.3645	48.7%
	Instant-NGP [35]	19.63	14.39	26.7%	0.8179	0.7104	13.1%	0.4956	0.6592	33.0%	0.6083	0.4157	31.7%
AVERAGE		20.40	14.15	30.6%	0.8334	0.6707	19.5%	0.2740	0.4575	70.0%	0.6419	0.4113	35.9%

Table 2. **Quantitative comparison of depth evaluation in extrapolated views in Setting 1.** 3DGM [31] demonstrates superior performance on most evaluation metrics, while VEGS [23] and GSPro [11] excel in SqRel and Delta1, respectively.

Baseline	AbsRel \downarrow	RMSE \downarrow	SqRel \downarrow	Delta1 \uparrow	Delta2 \uparrow	Delta3 \uparrow
3DGS [25]	0.361	14.44	10.41	0.649	0.824	0.895
3DGM [31]	0.301	13.93	8.906	0.651	0.846	0.915
PGSR [7]	0.366	17.57	21.50	0.759	0.834	0.883
GSPro [11]	0.355	19.66	32.01	0.643	0.839	0.909
VEGS [23]	0.368	15.00	8.398	0.441	0.691	0.827

[38] to mask out pixels corresponding to potentially movable objects during both training and evaluation and also exclude them from the initialization of 3DGS.

Evaluation Metrics. We use three widely-used metrics to evaluate visual quality: peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), and learned perceptual image patch similarity (LPIPS) [61]. We also employ DINOv2 [36] feature cosine similarity to evaluate image quality in latent space. For geometry evaluation, we use depth metrics, including RMSE and $\delta_{1.25}$. Additionally, we incorporate the Fréchet Inception Distance (FID) [21] to

Table 3. **Perceptual quality metric of VEGS [23].** Our benchmark also incorporates the Fréchet Inception Distance (FID) to evaluate the inpainting ability of models with generative priors.

	Setting 1	Setting 2	Setting 3	AVERAGE
FID (In.) \downarrow	46.3	33.5	46.1	42.0
FID (Ex.) \downarrow	87.4	67.2	132.4	95.7
FID (Drop)	89%	101%	187%	126.7%

evaluate the inpainting ability of models with strong priors.

4.2. Experimental Results

Table 1 presents the quantitative results across Settings 1-3, while Figure 5 illustrates the qualitative outcomes on the extrapolated test set. Table 3 and Table 2 present the perceptual quality and depth metrics. The results indicate that, while the metrics perform relatively well in the interpolated test set, there is a significant drop in performance in the extrapolated test set across all baselines.

Setting 1: Translation-only. In Setting 1, training views fully cover test views with moderate translational changes. (1) Results show a consistent drop in performance from

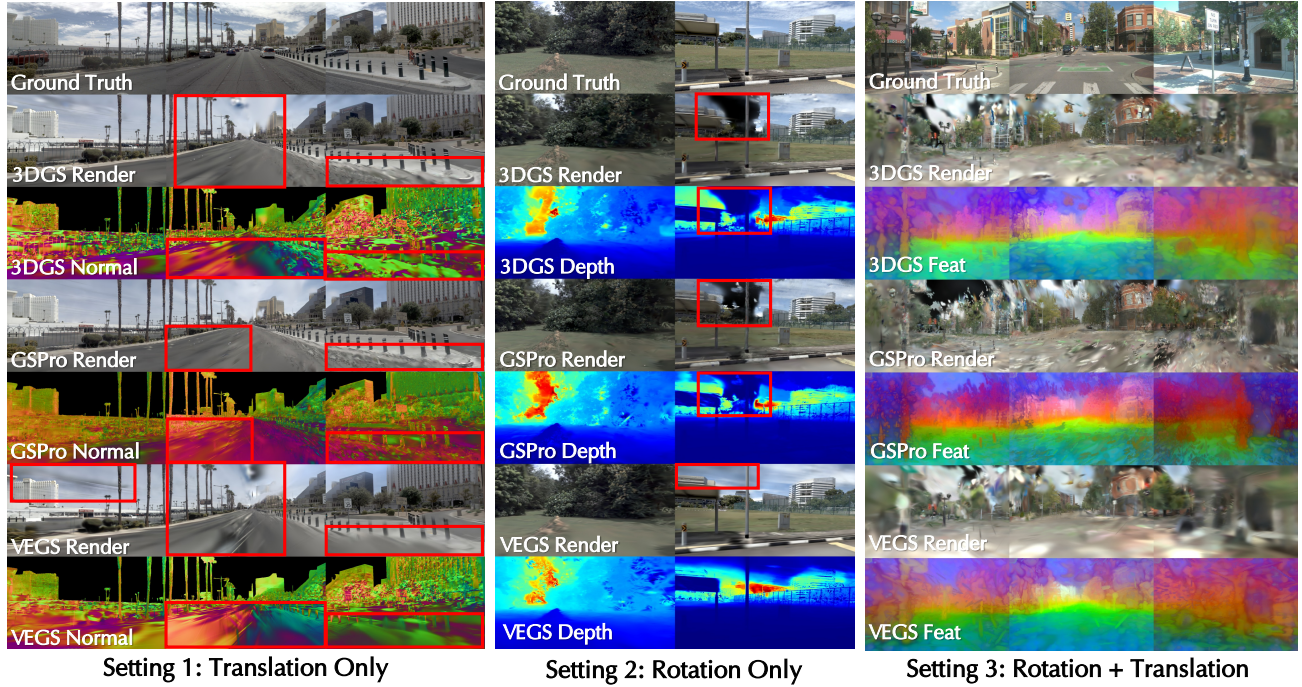


Figure 5. **Qualitative comparison of extrapolated view synthesis across different settings.** For each setting, results from different methods are compared against the ground truth. Red boxes highlight areas where methods are limited in capturing fine details, such as road surfaces, sky regions, or object boundaries, demonstrating the challenges faced by each approach under varying movement complexities.

interpolation to extrapolation across all metrics, highlighting the challenge of generalization to views with variations. Relative drops vary by metric: PSNR falls 23–25% (e.g., GSPRO [11]: 21.51 \rightarrow 16.39), with SSIM and LPIPS showing similar declines. **(2)** On the extrapolative test set, methods perform comparably: GSPRO tops PSNR (16.39) and minimizes LPIPS (0.2450), while 3DGM [31] leads in SSIM (0.7248). Both yield similar feature cosine similarity; VEGS [23] and PGSR [7] lag slightly. Results suggest that although GSPRO marginally outperforms others, differences are small, underscoring the need for improved solutions.

Setting 2: Rotation-only. In Setting 2, training views provide extensive coverage of the surrounding scene. However, most methods show poor generalization, with PSNR dropping by 22.75%. Rotation changes are particularly challenging in textured regions (e.g., trees and intricate details), often resulting in blurring, while regions perpendicular to the vehicle remain difficult to capture. Additionally, distant areas pose reconstruction challenges, frequently leading to missing buildings and sky blackouts. Among the evaluated baselines, VEGS [23] and GSPRO [11] stand out as the best-performing baselines in this setting: VEGS utilizes its diffusion prior for inpainting and refining missing regions, while GSPRO’s robust geometry handling improves generalization.

Setting 3: Translation + Rotation. In Setting 3, the view changes are the largest. **(1)** All methods exhibit notable performance drops from interpolation to extrapolation across metrics. For instance, 3DGS [25] experiences a PSNR


drop of 29.4% (21.22 \rightarrow 14.99), while GSPRO [11] undergoes a similar drop of 31.3% (21.58 \rightarrow 14.82). **(2)** Feature 3DGS [65] and 3DGM [31] stand out as leading methods, excelling in LPIPS (0.3816) and SSIM (0.7233), respectively. However, overall performance remains limited, with PSNR consistently falling below 15, highlighting significant room for improvement in generating high-fidelity outputs for extrapolated views.

5. Discussions

Lighting Inconsistency Handling. Lighting inconsistencies are a widespread challenge in multi-traversal datasets due to varying illumination and weather conditions. To address this, we have taken specific measures. **(1)** Importantly, our multi-traversal dataset is manually curated to ensure that the lighting across images appears consistent to the eye. This step helps reduce the impact of extreme lighting changes and ensures a fairer evaluation of our method under controlled conditions. **(2)** One potential approach is to incorporate an appearance embedding for each image. We experiment with the Gaussian in the Wild (GS-W) baseline [59]. GS-W replaces traditional spherical harmonic-based color modeling with a method that separates intrinsic properties of each Gaussian point from dynamic appearance features of each image. This approach captures the stable, inherent appearance of objects while accommodating dynamic factors like highlights and shadows. Despite these efforts, our results, especially in Setting 3, still exhibit



(a) Mitigating lighting inconsistency by camera embedding.



(b) Qualitative dynamic scenes rendering comparison.

Figure 6. **Qualitative results of dynamic baseline and lighting handling.** We mitigate the lighting inconsistency issue by carefully selecting traversals under similar lighting conditions, and it can be further alleviated by introducing camera embeddings. We provide the quantitative evaluation of dynamic scene reconstruction methods in Setting 2.

Table 4. **Quantitative performance of GS-W [59] in different settings.** After learning the lighting features, the interpolated and extrapolated test metrics show significant improvement compared to other baselines. However, there is still a considerable drop from interpolation to extrapolation.

Setting	PSNR \uparrow		SSIM \uparrow		LPIPS \downarrow		Feat Cos Sim \uparrow	
	In.	Ex.	In.	Ex.	In.	Ex.	In.	Ex.
Setting 1	28.15	20.22	0.89	0.78	0.15	0.23	0.71	0.64
Setting 2	30.10	21.21	0.91	0.82	0.13	0.20	0.76	0.67
Setting 3	28.62	19.36	0.87	0.73	0.15	0.31	0.74	0.50

a significant performance drop, as shown in Table 4 and Figure 6a, underscoring the limitations of Gaussian-based models in handling extrapolated, unseen scenarios.

Dynamic Scenes. In addition to the static scene reconstruction methods that we evaluated in Table 1, some dynamic scene reconstruction techniques have recently emerged. We conduct evaluation for Setting 2 to assess these approaches, e.g., OmniRe [10]. It organizes rigid-deformable nodes and background nodes to capture dynamic scene structures and employs SMPL [32] for non-rigid object modeling. As illustrated in Figure 6b, the rendering results reveal a significant performance gap between the training and extrapolated test cameras. In the extrapolated test views, objects such as trees and stakes lose texture and geometric details, resulting in noticeably blurry outputs, whereas the training views maintain high fidelity. As shown in Table 5, the reconstruction metrics indicate an average drop of 25% when transitioning from interpolated to extrapolated settings. The results highlight the challenges of extrapolated view synthesis in dynamic scenes and the need for further research.

6. Conclusions, Limitations and Future Work

Conclusions. We introduce the first benchmark enabling quantitative evaluation of extrapolated view synthesis, advancing photorealistic simulation for self-driving

Table 5. **Quantitative performance of OmniRe [10] in Setting 2.** The experiment uses a single traversal with dynamic objects, showing a noticeable drop from interpolation to extrapolation.

Setting	PSNR \uparrow		SSIM \uparrow		LPIPS \downarrow		Feat Cos Sim \uparrow	
	In.	Ex.	In.	Ex.	In.	Ex.	In.	Ex.
Setting 2	19.78	15.32	0.65	0.45	0.38	0.53	0.73	0.58

and robotics. Our benchmark integrates real-world multi-traversal, multi-agent, and multi-camera data, categorizes scenes into different evaluation settings, and evaluates state-of-the-art NVS models. Experimental results reveal that while some methods address specific challenges, current models demonstrate limited generalization, with significant overfitting to training views and suboptimal performance in extrapolated view synthesis. To support further research, we will release the dataset and benchmark, addressing the long-standing data scarcity and providing evaluation protocols. We believe the EUVS benchmark will catalyze meaningful advancements in self-driving and robotics innovation.

Limitations and Future Work. Our benchmark does have limitations. First, although we provide the ground truth to enable quantitative evaluation for static scenes in all settings and the foreground objects evaluation in Setting 2, we lack the foreground evaluation in Settings 1 and 3. Future work will aim to expand our work to include the dynamic object evaluation in Settings 1 and 3 using multi-agent data. Secondly, although we carefully manually selected and ensured that the test trajectory viewpoints are well covered by training trajectories, a more in-depth evaluation of how observed and unseen regions differ is left for future work.

Acknowledgment. This work was supported in part through NSF grants 2238968 and 2121391, and the NYU IT High Performance Computing resources, services, and staff expertise. Yiming Li is supported by NVIDIA Graduate Fellowship (2024-2025).

References

- [1] Alexander Amini, Tsun-Hsuan Wang, Igor Gilitschenski, Wilko Schwarting, Zhijian Liu, Song Han, Sertac Karaman, and Daniela Rus. Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2419–2426. IEEE, 2022. 3
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19697–19705, 2023. 5, 6
- [3] 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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 3
- [4] 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, 3, 4
- [5] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8748–8757, 2019. 3
- [6] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19457–19467, 2024. 2
- [7] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *arXiv preprint arXiv:2406.06521*, 2024. 5, 6, 7
- [8] 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. 2, 3
- [9] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*, 2024. 2
- [10] Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo de Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gojcic, Sanja Fidler, Marco Pavone, et al. Omnire: Omni urban scene reconstruction. *arXiv preprint arXiv:2408.16760*, 2024. 8
- [11] Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. Gaussian-pro: 3d gaussian splatting with progressive propagation. In *Forty-first International Conference on Machine Learning*, 2024. 4, 6, 7
- [12] Carlos A Diaz-Ruiz, Youya Xia, Yurong You, Jose Nino, Junan Chen, Josephine Monica, Xiangyu Chen, Katie Luo, Yan Wang, Marc Emond, et al. Ithaca365: Dataset and driving perception under repeated and challenging weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21383–21392, 2022. 3
- [13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 1, 3
- [14] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9710–9719, 2021. 3
- [15] Lue Fan, Hao Zhang, Qitai Wang, Hongsheng Li, and Zhaoxiang Zhang. Freesim: Toward free-viewpoint camera simulation in driving scenes, 2024. 3
- [16] Lan Feng, Quanyi Li, Zhenghao Peng, Shuhan Tan, and Bolei Zhou. Trafficgen: Learning to generate diverse and realistic traffic scenarios. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3567–3575. IEEE, 2023. 3
- [17] Jannik Fritsch, Tobias Kuehnl, and Andreas Geiger. A new performance measure and evaluation benchmark for road detection algorithms. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pages 1693–1700. IEEE, 2013. 3
- [18] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. Vista: A generalizable driving world model with high fidelity and versatile controllability. *arXiv preprint arXiv:2405.17398*, 2024. 3
- [19] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 3
- [20] Huasong Han, Kaixuan Zhou, Xiaoxiao Long, Yusen Wang, and Chunxia Xiao. Ggs: Generalizable gaussian splatting for lane switching in autonomous driving. *arXiv preprint arXiv:2409.02382*, 2024. 2, 3
- [21] 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
- [22] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 3, 5, 6
- [23] Sungwon Hwang, Min-Jung Kim, Taewoong Kang, Jayeon Kang, and Jaegul Choo. Vegs: View extrapolation of urban scenes in 3d gaussian splatting using learned priors. *arXiv preprint arXiv:2407.02945*, 2024. 2, 3, 5, 6, 7

- [24] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5322–5332, 2024. 3
- [25] 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, 3, 4, 6, 7
- [26] Mustafa Khan, Hamidreza Fazlali, Dhruv Sharma, Tongtong Cao, Dongfeng Bai, Yuan Ren, and Bingbing Liu. Autosplat: Constrained gaussian splatting for autonomous driving scene reconstruction. *arXiv preprint arXiv:2407.02598*, 2024. 3
- [27] Seung Wook Kim, Jonah Philion, Antonio Torralba, and Sanja Fidler. Drivegan: Towards a controllable high-quality neural simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5820–5829, 2021. 3
- [28] Quanyi Li, Zhenghao Mark Peng, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and Bolei Zhou. Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling. *Advances in neural information processing systems*, 36, 2024. 3
- [29] Yiming Li, Dekun Ma, Ziyang An, Zixun Wang, Yiqi Zhong, Siheng Chen, and Chen Feng. V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving. *IEEE Robotics and Automation Letters*, 7(4): 10914–10921, 2022. 3
- [30] Yiming Li, Zhiheng Li, Nuo Chen, Moonjun Gong, Zonglin Lyu, Zehong Wang, Peili Jiang, and Chen Feng. Multiagent multitraversal multimodal self-driving: Open mars dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22041–22051, 2024. 2, 3, 4
- [31] Yiming Li, Zehong Wang, Yue Wang, Zhiding Yu, Zan Gojcic, Marco Pavone, Chen Feng, and Jose M. Alvarez. Memorize what matters: Emergent scene decomposition from multitraverse. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 4, 6, 7
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 8
- [33] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 3
- [34] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 5, 6
- [36] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 6
- [37] Matthew Pitropov, Danson Evan Garcia, Jason Rebello, Michael Smart, Carlos Wang, Krzysztof Czarnecki, and Steven Waslander. Canadian adverse driving conditions dataset. *The International Journal of Robotics Research*, 40(4-5):681–690, 2021. 3
- [38] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos, 2024. 6
- [39] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5
- [40] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics: Results of the 11th International Conference*, pages 621–635. Springer, 2018. 3
- [41] Xi Shi, Lingli Chen, Peng Wei, Xi Wu, Tian Jiang, Yonggang Luo, and Lecheng Xie. Dhgs: Decoupled hybrid gaussian splatting for driving scene. *arXiv preprint arXiv:2407.16600*, 2024. 3
- [42] Shuhan Tan, Kelvin Wong, Shenlong Wang, Sivabalan Manivasagam, Mengye Ren, and Raquel Urtasun. Scenegen: Learning to generate realistic traffic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 892–901, 2021. 3
- [43] Izzeddin Teeti, Valentina Musat, Salman Khan, Alexander Rast, Fabio Cuzzolin, and Andrew Bradley. Vision in adverse weather: Augmentation using cyclecars with various object detectors for robust perception in autonomous racing. *arXiv preprint arXiv:2201.03246*, v3, 2023. 3
- [44] 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. 1
- [45] Christopher Wewer, Kevin Raj, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. *arXiv preprint arXiv:2403.16292*, 2024. 2
- [46] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint arXiv:2301.00493*, 2023. 2, 3, 4
- [47] Chenming Wu, Jiadai Sun, Zhelun Shen, and Liangjun Zhang. Mapnerf: Incorporating map priors into neural ra-

- diance fields for driving view simulation. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7082–7088. IEEE, 2023. 3
- [48] Ke Wu, Kaizhao Zhang, Zhiwei Zhang, Shanshuai Yuan, Muer Tie, Julong Wei, Zijun Xu, Jieru Zhao, Zhongxue Gan, and Wenchao Ding. Hgs-mapping: Online dense mapping using hybrid gaussian representation in urban scenes. *arXiv preprint arXiv:2403.20159*, 2024. 3
- [49] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuantao Chen, Runyi Yang, et al. Mars: An instance-aware, modular and realistic simulator for autonomous driving. In *CAAI International Conference on Artificial Intelligence*, pages 3–15. Springer, 2023. 3
- [50] Bernhard Wymann, Eric Espié, Christophe Guionneau, Christos Dimitrakakis, Rémi Coulom, and Andrew Sumner. Torcs, the open racing car simulator. *Software available at <http://torcs.sourceforge.net>*, 4(6):2, 2000. 3
- [51] Danfei Xu, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Bits: Bi-level imitation for traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2929–2936. IEEE, 2023. 3
- [52] Runsheng Xu, Xin Xia, Jinlong Li, Hanzhao Li, Shuo Zhang, Zhengzhong Tu, Zonglin Meng, Hao Xiang, Xiaoyu Dong, Rui Song, et al. V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13712–13722, 2023. 3
- [53] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339*, 2024. 2, 3
- [54] Chen Yang, Peihao Li, Zanwei Zhou, Shanxin Yuan, Bingbing Liu, Xiaokang Yang, Weichao Qiu, and Wei Shen. Nerfvs: Neural radiance fields for free view synthesis via geometry scaffolds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16549–16558, 2023. 2
- [55] 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. 3
- [56] 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. 1, 3
- [57] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. 3
- [58] Zhongrui Yu, Haoran Wang, Jinze Yang, Hanzhang Wang, Zeke Xie, Yunfeng Cai, Jiale Cao, Zhong Ji, and Mingming Sun. Sgd: Street view synthesis with gaussian splatting and diffusion prior. *arXiv preprint arXiv:2403.20079*, 2024. 3
- [59] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. *arXiv preprint arXiv:2403.15704*, 2024. 7, 8
- [60] 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. 2
- [61] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [62] 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. 3
- [63] Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen, Sushant Veer, Tong Che, Baishakhi Ray, and Marco Pavone. Guided conditional diffusion for controllable traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3560–3566. IEEE, 2023. 3
- [64] Hongyu Zhou, Jiahao Shao, Lu Xu, Dongfeng Bai, Weichao Qiu, Bingbing Liu, Yue Wang, Andreas Geiger, and Yiyi Liao. Hugs: Holistic urban 3d scene understanding via gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21336–21345, 2024. 2, 3
- [65] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. 5, 6, 7
- [66] Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21634–21643, 2024. 2, 3