

DriveGen3D: Boosting Feed-Forward Driving Scene Generation with Efficient Video Diffusion

Weijie Wang^{1,2*} Jiagang Zhu^{1,3*} Zeyu Zhang¹ Xiaofeng Wang^{1,4}
 Zheng Zhu¹ Guosheng Zhao^{1,4} Chaojun Ni^{1,5} Haoxiao Wang²
 Guan Huang¹ Xinze Chen¹ Yukun Zhou¹ Wenkang Qin¹
 Duochoa Shi² Haoyun Li^{1,4} Guanghong Jia³ Jiwen Lu³

¹GigaAI ²Zhejiang University ³Tsinghua University
⁴Institute of Automation, Chinese Academy of Sciences ⁵Peking University

*Equal contribution ✉Corresponding author

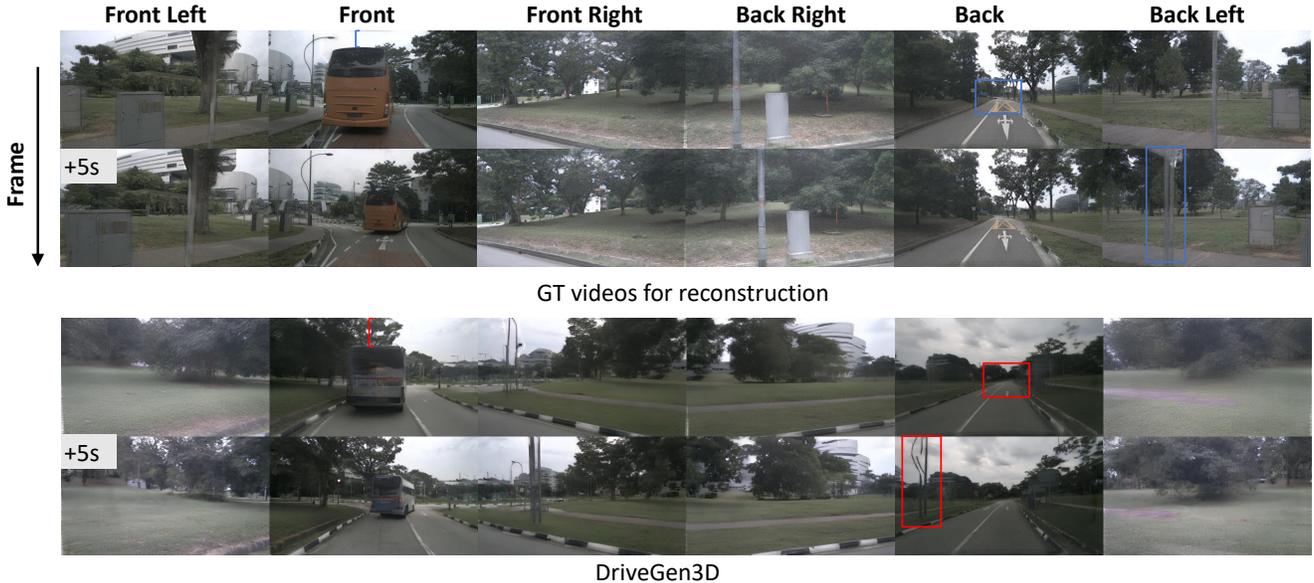


Figure 1. **DriveGen3D visualization.** Comparison of GT videos and generated video for reconstruction.

Abstract

We present *DriveGen3D*, a novel framework for generating high-quality and highly controllable dynamic 3D driving scenes that addresses critical limitations in existing methodologies. Current approaches to driving scene synthesis either suffer from prohibitive computational demands for extended temporal generation, focus exclusively on prolonged video synthesis without 3D representation, or restrict themselves to static single-scene reconstruction. Our work bridges this methodological gap by integrating accelerated long-term video generation with large-scale dynamic scene reconstruction through multimodal conditional con-

trol. *DriveGen3D* introduces a unified pipeline consisting of two specialized components: **FastDrive-DiT**, an efficient video diffusion transformer for high-resolution, temporally coherent video synthesis under text and Bird’s-Eye-View (BEV) layout guidance; and **FastRecon3D**, a feed-forward reconstruction module that rapidly builds 3D Gaussian representations across time, ensuring spatial-temporal consistency. Together, these components enable real-time generation of extended driving videos (up to 424×800 at 12 FPS) and corresponding dynamic 3D scenes, achieving SSIM of **0.811** and PSNR of **22.84** on novel view synthesis, all while maintaining parameter efficiency.

1. Introduction

The synthesis of 3D dynamic driving environments has emerged as a key research frontier in autonomous systems, driven by its wide-ranging applications in simulation, perception, and planning. While recent advances in video generation [1–7] and 3D scene reconstruction [8–16] have made substantial progress, a critical gap remains: the lack of an integrated and efficient framework that unifies long-horizon video synthesis and large-scale 3D scene reconstruction under multimodal control.

Existing methodologies typically address either temporal coherence in video generation or spatial fidelity in scene reconstruction—but not both—often requiring high computational cost or suffering from limited scalability. For example, state-of-the-art diffusion-based models like MagicDriveDiT [2] can produce high-resolution driving sequences but require up to 30 minutes to generate a single 1600×900 video, making them impractical for real-time use.

From the reconstruction perspective, optimization-based approaches [8–10, 17] are similarly time-consuming, often needing 30 minutes per scene. While recent feed-forward methods [11–13, 18–20] have reduced reconstruction time to seconds, they remain limited in scale and rarely integrate with dynamic video generation pipelines.

To bridge this gap, we propose **DriveGen3D**, an efficient and unified framework that integrates two specialized modules: **FastDrive-DiT**, an accelerated video diffusion transformer for high-resolution driving video generation, and **FastRecon3D**, a feed-forward reconstruction pipeline that builds dynamic 3D scenes from multi-view video frames in real time. FastDrive-DiT employs both diffusion step caching and quantized attention to reduce inference time by over $2\times$, while FastRecon3D leverages temporal-aware Gaussian splatting to produce high-fidelity reconstructions with minimal latency. Together, these components enable high-quality video generation and complete 3D reconstruction within 6 minutes, significantly outperforming prior methods in both efficiency and scalability, as shown in Figure 2.

Our key contributions are summarized as follows:

- We introduce **DriveGen3D**, a unified framework for efficient 3D driving scene generation that combines video synthesis and feed-forward 3D reconstruction under multimodal control via text and BEV layout.
- We propose **FastDrive-DiT**, a high-performance video diffusion transformer equipped with step caching and quantized attention, achieving over $2\times$ inference speedup while maintaining visual fidelity. We also develop **FastRecon3D**, a feed-forward 3D scene reconstruction module based on temporal-aware Gaussian splatting, enabling fast and scalable scene modeling

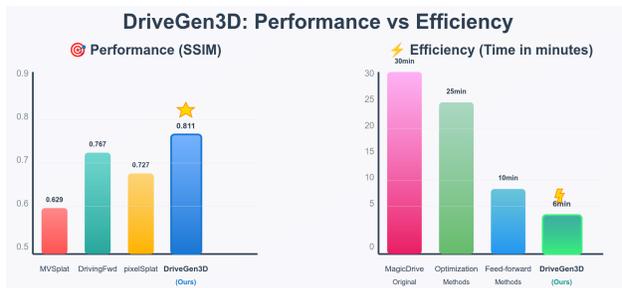


Figure 2. **Performance vs. Efficiency.** DriveGen3D achieves the highest SSIM (0.811) while significantly reducing generation time to 6 minutes—an 80% improvement over optimization-based and diffusion-based baselines—demonstrating superior video quality and real-time capability.

from generated videos.

- We conduct extensive experiments demonstrating that **DriveGen3D** outperforms state-of-the-art baselines in both video quality and 3D reconstruction accuracy. Specifically, our framework achieves a reconstruction SSIM of **0.811** on generated inputs, surpassing the previous best (0.767) while maintaining competitive PSNR. Moreover, DriveGen3D reduces the end-to-end generation time to under **6 minutes**, including video synthesis and 3D scene reconstruction, marking a substantial improvement in efficiency.

2. Related Work

2.1. Video generation for driving scene

Recent advancements in street view generation and autonomous driving scene synthesis have significantly improved the fidelity and controllability of synthetic data. MagicDrive [1] introduces a framework for street view generation with diverse 3D geometry controls, such as camera poses and 3D bounding boxes, enhancing tasks like BEV segmentation and 3D object detection through cross-view consistency. MagicDriveDiT [2] extends this by addressing high-resolution, long video generation for autonomous driving, leveraging flow matching and spatial-temporal conditional encoding to achieve superior scalability and control. InfiniCube [21] focuses on unbounded dynamic 3D driving scene generation, combining scalable 3D representations with video models to ensure geometric and appearance consistency across large-scale scenes. DrivingDiffusion [3] synthesizes multi-view driving videos with precise layout control, ensuring cross-view and cross-frame consistency through a spatial-temporal diffusion framework. UniScene [22] unifies the generation of semantic occupancy, video, and LiDAR data, employing a progressive generation process to reduce complexity and improve downstream task performance. DriveDreamer [4] pioneers real-world-driven

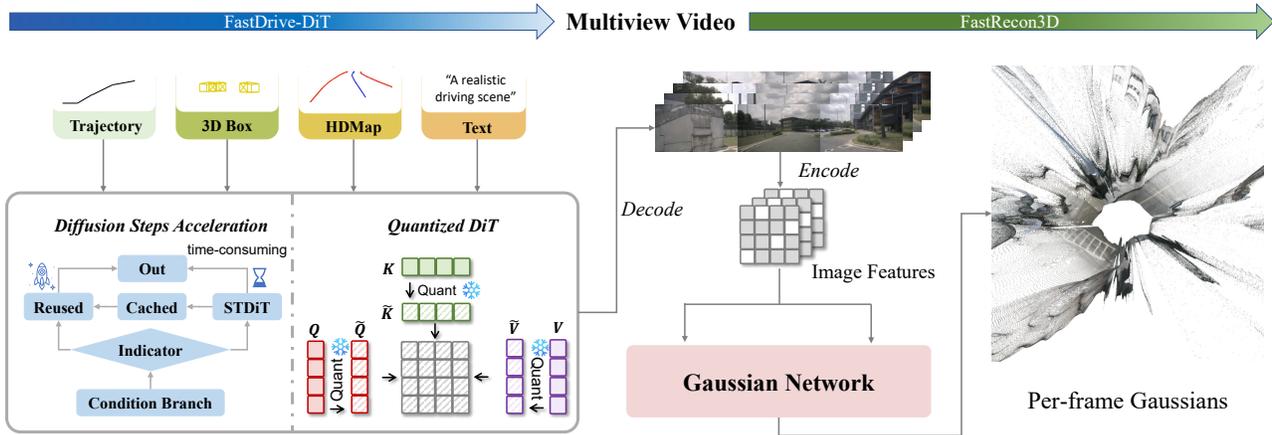


Figure 3. **Overview of DriveGen3D.** (a) Given textual and BEV layout conditions, our model first employs an accelerated Video Diffusion Transformer to synthesize a long driving video. (b) Next, a per-frame 3D Gaussian Splatting representation is utilized to construct entire scene from the generated video frames.

world models, using diffusion models to capture complex driving environments and enhance driving video generation and action prediction. Panacea [23] integrates a novel 3D attention and a two-stage generation pipeline to maintain coherence, supplemented by the ControlNet framework for meticulous control by the Bird’s-Eye-View (BEV) layouts. DriveDreamer-2 [5] integrates LLMs to generate user-defined driving videos, improving temporal and spatial coherence while surpassing state-of-the-art methods in video quality metrics like FID and FVD. Together, these works advance the field of autonomous driving by providing scalable, controllable, and high-fidelity synthetic data generation frameworks.

2.2. Reconstruction for driving scene

The reconstruction of dynamic driving scenes [3, 8, 17, 24–30] has emerged as a critical task in autonomous systems and immersive environment modeling. Contemporary approaches predominantly leverage Gaussian splatting-based representations due to their inherent balance between rendering efficiency and geometric expressiveness. Early methodologies in this domain adopted optimization-based paradigms, exemplified by works such as StreetGaussian [8], DrivingGaussian [10], and HUGS [9]. These frameworks optimize scene representations per instance for specific street segments (typically under 100 meters in scale) through iterative refinement processes spanning approximately 30 minutes.

Recent advancements have shifted toward feed-forward architectures [18–20, 31–41] to enable rapid 3D reconstruction. Methods like PixelSplat [18], MVSplat [19], and DepthSplat [20] employ large pretrained networks to directly infer Gaussian parameters from multi-view in-

puts, reducing reconstruction time from minutes to seconds. Though these approaches demonstrate generalizability across scenes, they often sacrifice reconstruction fidelity in geometrically complex regions or under sparse observational constraints. Parallel innovations address the temporal dimension of driving scenes: InfiniCube [21] extends the Scube [42] framework to 3D street generation, disentangling dynamic vehicles from static backgrounds via hybrid voxel-video control mechanisms. Drive3R [13] adapts the Spann3R [31] architecture for per-frame 3D scene reconstruction through temporal consistency priors. DrivingRecon [12] mimics StreetGaussians’ pipeline but replaces iterative optimization with feed-forward prediction, achieving real-time capability at moderate resolutions. DrivingForward [11] enhances sparse-view reconstruction robustness by jointly learning pose estimation and depth prediction modules within its network architecture. These advancements collectively highlight two persistent limitations:

- Existing feed-forward 3D reconstruction methods operate at constrained spatial resolutions (typically $\leq 512 \times 512$)
- Prior works on generative models for conditional scene synthesis required substantial computational resources and complex framework integration, which hindered their widespread adoption.

3. Method

3.1. Overview

DriveGen3D, as illustrated in Figure 3, is an integrated 3D driving scene generation system composed of two key components: **FastDrive-DiT** for efficient long video generation, and **FastRecon3D** for feed-forward 3D scene recon-

struction. The pipeline begins with FastDrive-DiT, which synthesizes high-resolution, temporally coherent driving videos under conditional guidance. These generated videos are then passed to FastRecon3D, which reconstructs dynamic 3D scenes in a feed-forward manner using temporal-aware Gaussian splatting. Together, these modules enable rapid and scalable 3D scene generation suitable for real-time simulation and autonomous driving applications.

The architecture of our efficient video diffusion module, dubbed FastDrive-DiT, integrates two targeted acceleration strategies to address the high inference cost of driving video generation. *Diffusion steps acceleration*: Given the slow inference speed of the driving video generation diffusion process, we employ TeaCache [43] to expedite the generation. In contrast to the original approach, we utilize TeaCache only for the conditional branch of the video generation model. This modification reduces the calculation cost of TeaCache by half, leading to a speedup of more than $2\times$. *Quantized DiT*: Considering the high cost of the attention part in the video generation model, we use SageAttention [44, 45] for the quantization of Q, K, V during the inference procedure. We conduct inference profiling and apply SageAttention to all the attention blocks, we are able to save an additional 30 seconds of inference time.

The architecture design of FastRecon3D is similar to DrivingForward [11]. The generated image is first processed through an image encoder, after which it is fed into decoders to obtain camera pose estimation and depth information respectively. These derived parameters are subsequently combined with the encoded image features and fed into the Gaussian prediction network, which predicts per-frame Gaussian parameters for dynamic scene representation.

3.2. FastDrive-DiT

Generating 3D driving scenes, especially in the autonomous driving domain, is notoriously time-consuming due to the multi-view nature of the data. The video generation step is particularly costly because of the underlying diffusion process. For instance, MagicDriveDiT can take up to 30 minutes to produce a video of resolution $1600 \times 848 \times 6 \times 233$. To address this inefficiency, we propose **FastDrive-DiT**, an enhanced and lightweight video diffusion model built on MagicDriveDiT with targeted acceleration strategies.

Diffusion steps acceleration. To accelerate the generation of long and high-resolution videos in MagicDriveDiT, we integrate TeaCache [43], a training-free caching approach that enhances the inference speed of diffusion models. TeaCache leverages timestep embeddings to modulate noisy inputs, approximating the fluctuating differences in model outputs across timesteps. By introducing a rescaling strategy, it refines these differences, enabling efficient

output caching with minimal computational overhead. Experiments show that TeaCache achieves up to $4.41\times$ acceleration over Open-Sora-Plan while maintaining visual quality, making it an effective solution for balancing speed and performance in video generation.

In the original TeaCache implementation, coefficients are computed from both the conditional and unconditional branches of the diffusion model. However, in our adaptation for MagicDriveDiT, we calculate coefficients exclusively from the conditional branch. This optimization not only reduces the forwarding time of TeaCache but also ensures no obvious degradation in performance, further enhancing the efficiency of the video generation process.

Quantized DiT. Considering the high cost of the attention part in the video DiT, we use SageAttention [44, 45] for quantization during the inference procedure.

SageAttention [44] is an efficient and accurate quantization method designed to accelerate attention mechanisms in transformers, which are computationally intensive with $O(N^2)$ complexity. By focusing on quantizing attention, it achieves about $2.1\times$ and $2.7\times$ higher OPS compared to FlashAttention2 [46], respectively, while maintaining accuracy across diverse models. SageAttention2 [45] further enhances efficiency by introducing INT4 quantization for Q and K matrices, FP8 for P and V , and precision-enhancing techniques like outlier smoothing and FP32 Matmul buffers. It achieves $3\times$ and $5\times$ higher OPS than FlashAttention2 on RTX4090, with negligible accuracy loss.

We carefully analyze the inference profiling and visualize $Q \times K \times V$ of different attention blocks, pointing out potential improvement techniques. In this paper, applying sageattention to the transformer blocks saves an additional 30 seconds of inference time with nearly no performance degradation. Experiment details are in section 4.2.2.

3.3. FastRecon3D

While the aforementioned methods enable realistic driving scene generation, applications like simulation require complete 3D scene models. To enable rapid novel scene synthesis, we introduce **FastRecon3D**, a feed-forward reconstruction module that avoids costly optimization while preserving quality. Built upon temporal-aware Gaussian Splatting, FastRecon3D reconstructs each frame’s 3D geometry by leveraging both past and future contexts, enabling fast and consistent 3D scene generation.

Per-Frame 3D Representation. Existing methods like Street Gaussians [8] and Infinicube [21] attempt dynamic scene reconstruction through object-scene separation but suffer from two critical limitations: 1) Dynamic elements in static backgrounds (e.g., pedestrians, cyclists) suffer motion blur artifacts 2) Computational constraints prevent large-scale scene reconstruction. Following DrivingForward [11], we propose a temporal-aware Gaussian Splatting formula-

Algorithm 1 Overall pipeline of DriveGen3D.

Require: Text prompt \mathcal{T} , BEV layout \mathcal{B}

Ensure: 3D driving scene $\mathcal{G}^{1:T}$

Stage 1: FastDrive-DiT – Efficient Long Video Generation

- 1: Encode \mathcal{T}, \mathcal{B} into conditioning vector \mathcal{C}
- 2: Initialize noise tensor \mathbf{z}_0
- 3: **for** $t = 1$ to T **do**
- 4: $\mathbf{z}_t \leftarrow \text{FastDrive-DiT}(\mathbf{z}_{t-1}, \mathcal{C})$ ▷ Apply diffusion steps with TeaCache
- 5: **end for**
- 6: Generated multi-view video frames $\{I_i^1, \dots, I_i^T\}_{i=1}^{N_{\text{views}}} \leftarrow \text{Decode}(\mathbf{z}_T)$

Stage 2: FastRecon3D – Feed-Forward 3D Scene Reconstruction

- 7: **for** $t = 1$ to T **do**
 - 8: Collect frames $\{I_i^{t-\Delta}, I_i^t, I_i^{t+\Delta}\}_{i=1}^{N_{\text{views}}}$
 - 9: Estimate camera poses: $\{\mathbf{P}_i^t\}_{i=1}^{N_{\text{views}}} \leftarrow \text{PoseNet}(I_i^t)$
 - 10: Estimate depth maps: $\{D_i^t\}_{i=1}^{N_{\text{views}}} \leftarrow \text{DepthNet}(I_i^t)$
 - 11: Predict 3D Gaussian Parameters for each frame: $\mathcal{G}^t \leftarrow \text{GaussianNet}(I_i^t, D_i^t, \mathbf{P}_i^t)$
 - 12: **end for**
 - 13: **return** Complete 3D dynamic scene $\mathcal{G}^{1:T}$
-

tion that reconstructs per-frame 3D Gaussian models while maintaining 3D consistency.

In our approach, each time step is represented by a set of 3D Gaussian primitives:

$$\mathcal{G}_i^t = \{\mathcal{G}_i^t\}_{i=1}^{N_t} = \{(\mu_i, \Sigma_i, \alpha_i, \mathbf{c}_i)\}_{i=1}^{N_t} \quad (1)$$

Recursively reconstruction from videos. We modified DrivingForward [11], a state-of-the-art 3D reconstruction model for autonomous driving, to build the scene.

After the frame generation of the $t + 1$ timestep, we extract all frames from time steps $t - 1$ to $t + 1$ to reconstruct the scene at time t . Formally, for each timestep t , given multi-view images $\{I_i^t\}_{i=1}^{N_t}$ and temporal neighbors $\{I_i^{t\pm\Delta}\}_{i=1}^{N_t}$, our model predicts Gaussian parameters $\mathcal{G}^t = \{\mu^t, \Sigma^t, \alpha^t, \mathbf{c}^t\}$ and save it as a 3D model:

$$\mathcal{G}^t = \mathcal{F}\left(\{I_i^{t-\Delta}, I_i^t, I_i^{t+\Delta}\}_{i=1}^{N_t}\right). \quad (2)$$

By leveraging both past and future context, this recursive reconstruction method effectively captures dynamic scene elements while maintaining high spatial fidelity. As a result, our approach produces complete 3D models in a matter of seconds, meeting the demanding requirements of simulation and other real-time applications without compromising on quality. The overall pipeline of DriveGen3D is outlined in Algorithm 1.

Type	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Static	MVSplat [19]	22.83	0.629	0.317
	pixelSplat [18]	25.00	0.727	0.298
Dynamic	UniPad [47]	16.45	0.375	-
	SelfOcc [48]	18.22	0.464	-
	EmerNeRF [49]	20.95	0.585	-
	DistillNeRF [50]	20.78	0.590	-
	DrivingForward [11] (640p)	21.67	0.727	0.259
	DrivingForward [11] (228p)	21.76	0.767	0.194
	Ours (w/ GT images)	23.71	0.733	0.285
Ours (w/ GEN images)	22.84	0.811	0.332	

Table 1. Comparison of our method against prior feed-forward and optimization-based methods. The last two rows show novel view rendering performance with either GT or generated video input. All metrics are computed at frame t given $t-1$ and $t+1$ as inputs.

4. Experiments

4.1. Experimental Setup

Dataset. The training dataset is obtained from the nuScenes dataset [51]. It consists of 700 training videos and 150 validation videos. For the 3D scene reconstruction model, we split the dataset into 20,000 short sequences for training.

Evaluation Metrics. For video-generation stage, we evaluate both the realism and controllability in street-view video generation. For video generation, we adhere to the benchmarks from [52]. To measure video quality, we use the Fréchet Video Distance (FVD). Regarding controllability, we employ mAP from 3D object detection and mIoU from BEV segmentation. For both of these tasks, we utilize BEVFormer [53], which is a video-based perception model. We generate corresponding videos for the annotations in the validation set. Then, we apply the aforementioned metrics and use perception models pre-trained on real-world data for evaluation.

For 3D scene reconstruction stage, we adopt novel view synthesis (NVS) to assess reconstruction quality, following the evaluation protocol established in Driving Forward [11]. Given sequential video frames from timestamps $t - 1$ and $t + 1$ as input, our model synthesizes the intermediate frame at timestamp t for quantitative comparison against ground truth. We report Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [54] in Table 1. To benchmark performance, we compare our DriveGen3D reconstruction model against the Driving Forward baseline (evaluate on ground truth images) across same validation frames and synthesized video outputs. This dual evaluation strategy disentangles reconstruction capability from generation artifacts, providing comprehensive insights into geometric and appearance recovery accuracy.

Implementation Details. We train the model with a res-

Total	spatial	temporal	cross-view	cross	other
615 s	104 s	82 s	163 s	62 s	204 s

Table 2. Time cost of different attention blocks of MagicDriveDiT during inference.

	FVD ↓	mAP ↑	mIoU ↑	Time cost 17f/233f ↓
MagicDriveDiT	111.58	17.10	21.92	211s/615 s
Ours (w/o Quant)	125.70	16.60	21.27	64 s/309 s
Ours	125.88	16.72	21.24	58 s/278 s

Table 3. Accelerating the inference process of MagicDriveDiT. 17f and 233f denote the frames count of generated videos.

olution of 424×800 . We inference the model on NVIDIA H20 GPUs. When assessing the time cost of our proposed method, the baselines for the video generation model are MagicDriveDiT (17f) and MagicDriveDiT (233f). The 3D scene reconstruction stage is trained on NVIDIA H20 GPUs for 2 days.

4.2. Experimental Results

4.2.1 Novel View Synthesis

Table 1 provides a comprehensive comparison of DriveGen3D against both optimization-based dynamic methods and feed-forward reconstruction methods. Notably, when using generated images instead of ground truth, DriveGen3D maintains competitive reconstruction quality with a PSNR of 22.84 and achieves the highest SSIM of 0.811, demonstrating strong temporal coherence and structure preservation in generated scenes. This suggests that despite operating on synthetic inputs, DriveGen3D produces reliable and structurally consistent 3D reconstructions, validating the effectiveness of its end-to-end pipeline.

4.2.2 Experimental Analysis and Ablation Study

Diffusion steps acceleration. Firstly, we visualize the input differences and output differences in consecutive timesteps in Figure 5. It is observed that the default configurations of TeaCache, i.e. *ALL*, exhibits a U-shape for model output, with downward trend initially, nearly constant for the middle and upward until the end. The same is for the *uncondition* branch. The *condition* branch exhibits a slightly different phenomenon, with the start of model output much smaller. We also apply simple polynomial fitting to fit a relationship between model input and output and use these coefficients to predict model output according to the input. As shown in the black lines in Figure 5, the quality of fitting is best for the *condition* branch, while *ALL* and *uncondition* do not fit well for the start. We attribute this to the different input and output relationships in the start of diffusion pro-

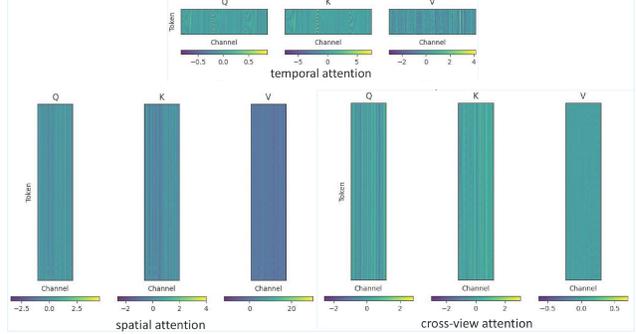


Figure 4. Typical examples of tensors’ data distribution in different attention blocks of MagicDriveDiT.

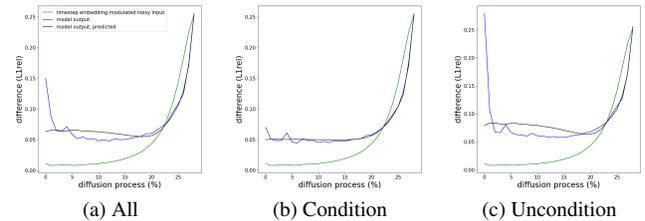


Figure 5. Visualization of input differences and output differences in consecutive timesteps. We plot *all*, *condition* and *uncondition* separately.

cess. So we finally only apply TeaCache to the *condition* branch of MagicDriveDiT. As shown in Table 3, equipping MagicDriveDiT with TeaCache for *condition* branch has a speedup of nearly three times and two times for 17 frames and 233 frames. The perception metrics, mAP and mIoU only shows slight decrease.

Quantized DiT. Secondly, we profile the time cost of MagicDriveDiT. As shown in Table 2, cross-view attention is identified as the most computationally costly process in MagicDriveDiT. We further analyze different attention components in MagicDriveDiT and plot the distribution of Q , K , V , as depicted in Figure 4. Notably, Q and K exhibit a trend similar to that in Figure 4 of SageAttention2. An interesting phenomenon is observed: the numeric range of V follows the order: spatial > temporal > cross-view. Specifically, the cross-view range is 10 times smaller than the spatial range and 5 times smaller than the temporal range. Given the principle of quantization, a smaller range is more conducive to quantization. Therefore, a promising solution is not only to apply SageAttention to cross-view attention but also substitute the high-precision quantization method for V with a lower-resolution one. Concretely, for V , the FP8, E4M3 data type can be replaced with FP8, E5M2 data type and INT8. We leave the latter for future attempt. In this paper, applying sageattention to the transformer blocks saves an additional 30 seconds (233 frames) of inference time after diffusion steps acceleration while maintaining

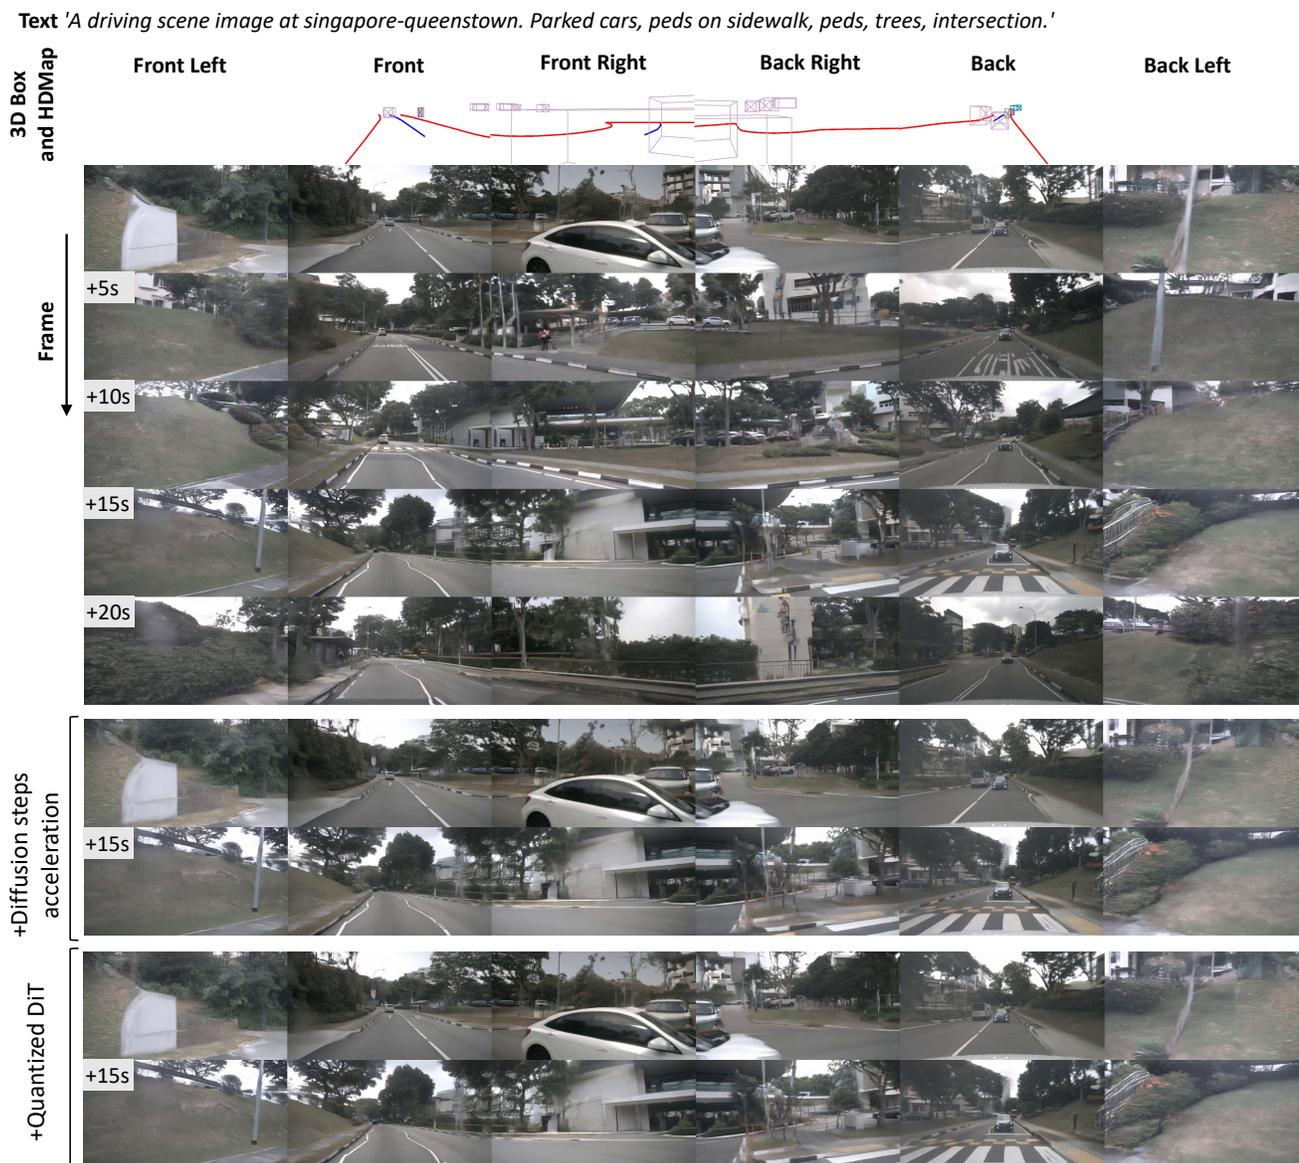


Figure 6. Comparison of video generation for MagicDriveDiT, Diffusion steps acceleration and Quantized DiT.

performance.

4.2.3 Visualization and Comparison

Visualization. In Figure 6, we compare the generated videos of baseline model, baseline model with diffusion steps accerelation and further quantization with SageAttention. With the proposed techniques, no obvious change is observed. It can still generate video longer than 15 s with a much faster speed. This highlights the efficiency of *DriveGen3D*.

Full pipeline results. In Figure 7, we show a typical output result of *DriveGen3D*. *DriveGen3D* can generate and

reconstruct videos for more than 20S, 12FPS.

Comparison with raw videos. By default, our reconstruction stage takes the generated videos as input. We compare with *DrivingForward* to validate the difference of raw videos and generated videos. As evidenced in Table 1, *DriveGen3D* maintains fidelity in video generation that is essentially comparable to existing methodologies with GT images, exhibiting only a marginal degradation. Qualitative analysis is shown in Figure 1. In the front view image of reconstruction videos, the spatial location of generated bus is different from gt videos, we check the 3D Box layout and find the generation videos doesn't match well. In the view of back left, the generation images is blurry, underscoring

Text 'A driving scene image at singapore-onenorth. Parking lot, barrier, exit parking lot.'

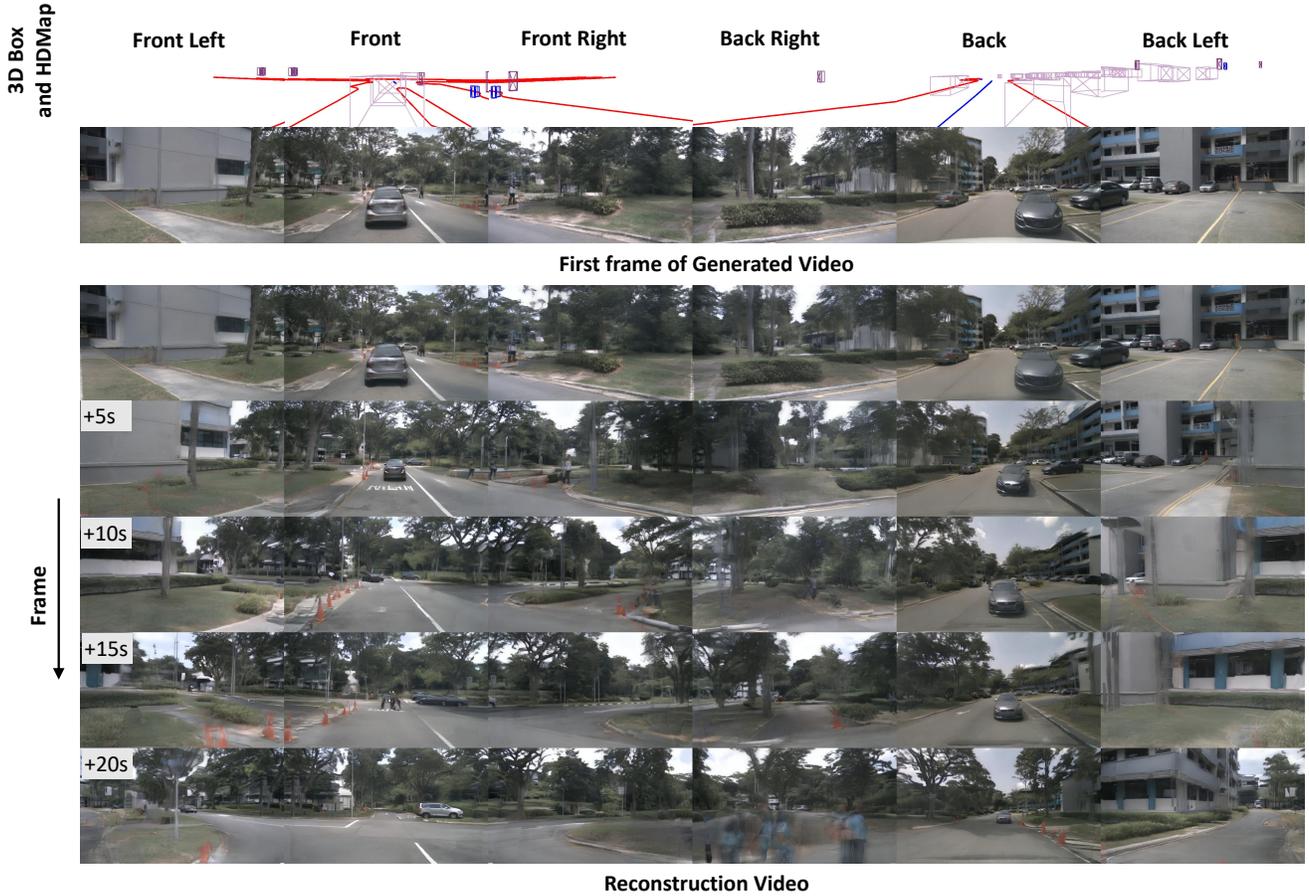


Figure 7. Visualization of multiview reconstruction video.

the needs of optimizing video generation model for reconstruction. Moreover, both GT videos and generation videos display drawbacks of flickering pole, pointing out the flaw of second stage reconstruction model.

5. Limitation and Future Work

While DriveGen3D demonstrates strong performance in efficient long-horizon video generation and rapid 3D scene reconstruction, several limitations remain. First, the reconstruction process relies on per-frame Gaussian representations without explicitly modeling object-level dynamics, which may limit fidelity in highly interactive or densely populated scenes. Second, although our feed-forward reconstruction achieves significant speedup, it may still suffer from minor temporal inconsistencies in low-texture regions or occluded areas.

In future work, we plan to explore semantic-aware 3D scene reconstruction that disentangles dynamic objects from static backgrounds, enabling object-level controllability and persistent identity tracking. Additionally, we aim to

extend DriveGen3D toward full-scene temporal modeling by integrating recurrent or transformer-based consistency modules. Finally, enhancing reconstruction robustness under noisy video inputs or adverse weather conditions remains an important direction to improve real-world deployment readiness.

6. Conclusion

We present DriveGen3D, an efficient framework for synthesizing high-resolution, long-duration driving videos from textual descriptions and BEV layouts, and generating high-quality large-scale dynamic scenes. Our approach marks a significant advancement in world modeling by bypassing conventional voxel-based generation paradigms [21, 39] through a novel integration of longitudinal video generation and scene reconstruction modules. This architecture enables faithful reproduction of real-world driving scenarios and thus paves the way for novel applications in autonomous vehicle simulation and dynamic world modeling.

References

- [1] 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. 2
- [2] Ruiyuan Gao, Kai Chen, Bo Xiao, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrivedit: High-resolution long video generation for autonomous driving with adaptive control. *arXiv preprint arXiv:2411.13807*, 2024. 2
- [3] Xiaofan Li, Yifu Zhang, and Xiaoqing Ye. Drivingdiffusion: Layout-guided multi-view driving scenarios video generation with latent diffusion model. In *European Conference on Computer Vision*, pages 469–485. Springer, 2024. 2, 3
- [4] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiayang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. *arXiv preprint arXiv:2309.09777*, 2023. 2
- [5] 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. *arXiv preprint arXiv:2403.06845*, 2024. 2, 3
- [6] 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. In *Advances in Neural Information Processing Systems*, 2024. 2
- [7] Zhuoran Yang, Xi Guo, Chenjing Ding, Chiyu Wang, and Wei Wu. Physical informed driving world model, 2024. 2
- [8] 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, 4
- [9] 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, June 2024. 2, 3
- [10] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21634–21643, 2024. 2, 3
- [11] Qijian Tian, Xin Tan, Yuan Xie, and Lizhuang Ma. Drivingforward: Feed-forward 3d gaussian splatting for driving scene reconstruction from flexible surround-view input. *arXiv preprint arXiv:2409.12753*, 2024. 2, 3, 4, 5
- [12] Hao Lu, Tianshuo Xu, Wenzhao Zheng, Yunpeng Zhang, Wei Zhan, Dalong Du, Masayoshi Tomizuka, Kurt Keutzer, and Yingcong Chen. Drivingrecon: Large 4d gaussian reconstruction model for autonomous driving. *arXiv preprint arXiv:2412.09043*, 2024. 2, 3
- [13] Xin Fei, Wenzhao Zheng, Yueqi Duan, Wei Zhan, Masayoshi Tomizuka, Kurt Keutzer, and Jiwen Lu. Driv3r: Learning dense 4d reconstruction for autonomous driving. *arXiv preprint arXiv:2412.06777*, 2024. 2, 3
- [14] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-irm: Large reconstruction model for 3d gaussian splatting, 2024. 2
- [15] Junge Zhang, Qihang Zhang, Li Zhang, Ramana Rao Kompella, Gaowen Liu, and Bolei Zhou. Urban scene diffusion through semantic occupancy map, 2024. 2
- [16] Chaojun Ni, Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Wenkang Qin, Guan Huang, Chen Liu, Yuyin Chen, Yida Wang, Xueyang Zhang, et al. Recondreamer: Crafting world models for driving scene reconstruction via online restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1559–1569, 2025. 2
- [17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 2023. 2, 3
- [18] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 2, 3, 5
- [19] 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, 3, 5
- [20] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Depthsplat: Connecting gaussian splatting and depth. *arXiv preprint arXiv:2410.13862*, 2024. 2, 3
- [21] 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. *arXiv preprint arXiv:2412.03934*, 2024. 2, 3, 4, 8
- [22] Bohan Li, Jiazhe Guo, Hongsu Liu, Yingshuang Zou, Yikang Ding, Xiwu Chen, Hu Zhu, Feiyang Tan, Chi Zhang, Tiancai Wang, et al. Uniscene: Unified occupancy-centric driving scene generation. *arXiv preprint arXiv:2412.05435*, 2024. 2
- [23] Yuqing Wen, Yucheng Zhao, Yingfei Liu, Fan Jia, Yanhui Wang, Chong Luo, Chi Zhang, Tiancai Wang, Xiaoyan Sun, and Xiangyu Zhang. Panacea: Panoramic and controllable video generation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6902–6912, 2024. 3
- [24] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [25] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron,

- and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 3
- [26] Ziyang Xie, Junge Zhang, Wenye Li, Feihu Zhang, and Li Zhang. S-nerf: Neural radiance fields for street views, 2023. 3
- [27] 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
- [28] Qiang Zhang, Seung-Hwan Baek, Szymon Rusinkiewicz, and Felix Heide. Differentiable point-based radiance fields for efficient view synthesis, 2023. 3
- [29] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting, 2024. 3
- [30] Ze Yang, Yun Chen, Jingkang Wang, Sivabalan Manivasagam, Wei-Chiu Ma, Anqi Joyce Yang, and Raquel Urtasun. Unisim: A neural closed-loop sensor simulator, 2023. 3
- [31] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. *arXiv preprint arXiv:2408.16061*, 2024. 3
- [32] Weijie Wang, Donny Y Chen, Zeyu Zhang, Duochoa Shi, Akide Liu, and Bohan Zhuang. Zpressor: Bottleneck-aware compression for scalable feed-forward 3dgs. *arXiv preprint arXiv:2505.23734*, 2025. 3
- [33] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5294–5306, 2025. 3
- [34] Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. Long-irm: Long-sequence large reconstruction model for wide-coverage gaussian splats. *arXiv preprint arXiv:2410.12781*, 2024. 3
- [35] Chuanrui Zhang, Yingshuang Zou, Zhuoling Li, Minmin Yi, and Haoqian Wang. Transplat: Generalizable 3d gaussian splatting from sparse multi-view images with transformers. In *AAAI Conf. Artif. Intell.*, volume 39, pages 9869–9877, 2025. 3
- [36] Yunsong Wang, Tianxin Huang, Hanlin Chen, and Gim Hee Lee. Freesplat: Generalizable 3d gaussian splatting towards free view synthesis of indoor scenes. *Adv. Neural Inform. Process. Syst.*, 37:107326–107349, 2024. 3
- [37] Botao Ye, Sifei Liu, Haofei Xu, Xueting Li, Marc Pollefeys, Ming-Hsuan Yang, and Songyou Peng. No pose, no problem: Surprisingly simple 3d gaussian splats from sparse unposed images. *arXiv preprint arXiv:2410.24207*, 2024. 3
- [38] Duochoa Shi, Weijie Wang, Donny Y. Chen, Zeyu Zhang, Jiawang Bian, Bohan Zhuang, and Chunhua Shen. Revisiting depth representations for feed-forward 3d gaussian splatting. *arXiv preprint arXiv:2506.05327*, 2025. 3
- [39] Weijie Wang, Yeqing Chen, Zeyu Zhang, Hengyu Liu, Haoxiao Wang, Zhiyuan Feng, Wenkang Qin, Zheng Zhu, Donny Y. Chen, and Bohan Zhuang. Volsplat: Rethinking feed-forward 3d gaussian splatting with voxel-aligned prediction. *arXiv preprint arXiv:2509.19297*, 2025. 3, 8
- [40] Chaojun Ni, Xiaofeng Wang, Zheng Zhu, Weijie Wang, Haoyun Li, Guosheng Zhao, Jie Li, Wenkang Qin, Guan Huang, and Wenjun Mei. Wonderturbo: Generating interactive 3d world in 0.72 seconds. *arXiv preprint arXiv:2504.02261*, 2025. 3
- [41] Haofei Xu, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Resplat: Learning recurrent gaussian splats. *arXiv preprint arXiv:2510.08575*, 2025. 3
- [42] Xuanchi Ren, Yifan Lu, Hanxue Liang, Jay Zhangjie Wu, Huan Ling, Mike Chen, Francis Fidler, Sanja and Williams, and Jiahui Huang. Scube: Instant large-scale scene reconstruction using voxsplats. In *Adv. Neural Inform. Process. Syst.*, 2024. 3
- [43] Feng Liu, Shiwei Zhang, Xiaofeng Wang, Yujie Wei, Haonan Qiu, Yuzhong Zhao, Yingya Zhang, Qixiang Ye, and Fang Wan. Timestep embedding tells: It’s time to cache for video diffusion model. *arXiv preprint arXiv:2411.19108*, 2024. 4
- [44] Jintao Zhang, Jia Wei, Pengle Zhang, Jun Zhu, and Jianfei Chen. Sageattention: Accurate 8-bit attention for plug-and-play inference acceleration. In *Int. Conf. Learn. Represent.*, 2025. 4
- [45] Jintao Zhang, Haofeng Huang, Pengle Zhang, Jia Wei, Jun Zhu, and Jianfei Chen. Sageattention2: Efficient attention with thorough outlier smoothing and per-thread int4 quantization, 2024. 4
- [46] Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *Int. Conf. Learn. Represent.*, 2024. 4
- [47] Honghui Yang, Sha Zhang, Di Huang, Xiaoyang Wu, Haoyi Zhu, Tong He, Shixiang Tang, Hengshuang Zhao, Qibo Qiu, Binbin Lin, et al. Unipad: A universal pre-training paradigm for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15238–15250, 2024. 5
- [48] Yuanhui Huang, Wenzhao Zheng, Borui Zhang, Jie Zhou, and Jiwen Lu. Selfocc: Self-supervised vision-based 3d occupancy prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19946–19956, 2024. 5
- [49] 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. 5
- [50] Letian Wang, Seung Wook Kim, Jiawei Yang, Cunjun Yu, Boris Ivanovic, Steven Waslander, Yue Wang, Sanja Fidler, Marco Pavone, and Peter Karkus. Distillnerf: Perceiving 3d scenes from single-glance images by distilling neural fields and foundation model features. *Advances in Neural Information Processing Systems*, 37:62334–62361, 2024. 5

- [51] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2020. 5
- [52] Zhiying Du and Zhen Xing. Challenge report: Track 2 of multimodal perception and comprehension of corner cases in autonomous driving. In *ECCV 2024 Workshop on Multimodal Perception and Comprehension of Corner Cases in Autonomous Driving*, 2024. 5
- [53] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022. 5
- [54] 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. 5