
Genesis: Multimodal Driving Scene Generation with Spatio-Temporal and Cross-Modal Consistency

Xiangyu Guo^{1*}, Zhanqian Wu^{2*}, Kaixin Xiong^{2*}, Ziyang Xu¹, Lijun Zhou², Gangwei Xu¹,
Shaoqing Xu², Haiyang Sun^{2†}, Bing Wang², Guang Chen²,
Hangjun Ye², Wenyu Liu¹, Xinggong Wang¹✉

¹Huazhong University of Science and Technology ²Xiaomi EV
<https://xiaomi-research.github.io/genesis/>

Abstract

We present Genesis, a unified world model for joint generation of multi-view driving videos and LiDAR sequences with spatio-temporal and cross-modal consistency. Genesis employs a two-stage architecture that integrates a DiT-based video diffusion model with 3D-VAE encoding, and a BEV-represented LiDAR generator with NeRF-based rendering and adaptive sampling. Both modalities are directly coupled through a shared condition input, enabling coherent evolution across visual and geometric domains. To guide the generation with structured semantics, we introduce DataCrafter, a captioning module built on vision-language models that provides scene-level and instance-level captions. Extensive experiments on the nuScenes benchmark demonstrate that Genesis achieves state-of-the-art performance across video and LiDAR metrics (FVD 16.95, FID 4.24, Chamfer 0.611), and benefits downstream tasks including segmentation and 3D detection, validating the semantic fidelity and practical utility of the synthetic data.

1 Introduction

As autonomous driving systems progress toward higher levels of intelligence, generating diverse, realistic driving scenarios [34, 30, 7, 44] has become essential for improving the robustness and safety of perception and planning modules. While recent advances have explored video generation and LiDAR sequence synthesis independently, achieving holistic multimodal consistency across visual and geometric modalities remains an open challenge, limiting the development of faithful world models that can simulate real-world sensor interactions.

As shown in Fig. 1, existing driving scene generation methods typically focus on generating data in a single modality, usually RGB videos [35, 10, 39, 50, 40, 42, 18, 26] or LiDAR point clouds [54, 53, 2, 3, 21, 22]. While these approaches have significantly advanced the field of driving scene generation, they overlook the synergistic potential of multimodal generation and lack consistency in aligning RGB videos with various sensor data, leading to limitations in real-world applications. Critically, they fall short of constituting a comprehensive world model, as they fail to capture the intricate dependencies between different sensory streams. Many of these methods rely on one-step layout-to-data pipelines conditioned solely on coarse spatial priors, such as BEV maps [36, 10, 43, 46, 39, 50, 40] or 3D boxes [12, 4, 33], which restricts their ability to capture complex scene dynamics and fine-grained semantics. To propel driving scene generation into the multimodal domain, UniScene [20] involves occupancy grids in for multimodal generation, while the acquisition of occupancy labels is

* Equal contribution, the work was done when Xiangyu Guo was an intern of Xiaomi EV.

† Project lead (sunhaiyang1@xiaomi.com). ✉ Corresponding author (xgwang@hust.edu.cn).




Figure 1: **Comparison of Multimodal Scene Generation Pipelines.** (a) *Video-only generation*, (b) *LiDAR-only generation*, (c) *Occupancy-guided dual-branch generation*, (d) **Our Genesis**.

highly costly, which severely limits its application. Moreover, most existing generation approaches generally rely on limited semantic annotations, typically in the form of coarse labels or generic captioning models, without fully leveraging the fine-grained descriptive capabilities of modern vision-language models (VLMs) [5, 23, 1]. This lack of structured semantic grounding restricts the fidelity, controllability, and contextual alignment of generated scenes, further distancing them from a truly grounded and actionable world model.

Therefore, advancing toward a realistic driving world model requires more than unimodal fidelity—it demands tight cross-modal alignment and fine-grained semantic grounding to ensure spatio-temporal coherence. However, existing methods suffer from three core limitations: (1) they often decouple video and LiDAR generation, weakening cross-modal consistency; (2) they rely on intermediate representations like occupancy grids that may cause information loss; (3) they provide limited semantic supervision, typically in the form of coarse layout maps or generic captions, which hinders scene-level controllability and realism.

To address these limitations, we propose **Genesis**, a unified joint generation framework tailored for autonomous driving that synthesizes multi-view RGB videos and LiDAR point clouds in a consistent and semantically grounded manner. Our framework introduces three key innovations:

- **Unified Multimodal Generation Architecture.** Genesis employs a unified pipeline where both video and LiDAR branches use shared conditional inputs, including scene descriptions and scene layouts, etc. To further guarantee the consistency between point clouds and image, BEV features extracted from projected image features are incorporated as conditional inputs into the LiDAR diffusion model. This approach enhances cross-modality consistency without relying on intermediaries like occupancy grids.
- **Structured Semantic Captions via DataCrafter.** To improve semantic controllability, we introduce DataCrafter, a structured captioning module built upon vision-language models. It extracts multi-view, scene-level and instance-level descriptions that are fused into dense, language-guided priors. These captions provide detailed semantic guidance to both video and LiDAR generators, resulting in outputs that are not only realistic but also interpretable and controllable.
- **State-of-the-art performance.** Extensive experiments conducted on the nuScenes benchmark reveal that Genesis achieves the most advanced performance in both video and LiDAR metrics.

2 Related Work

Driving Scene Video Generation. Driving scene video generation has seen rapid progress, with many methods relying on structured spatial priors for controllability. BEVGen [35] conditions generation on BEV maps to encode road and vehicle layouts, but ignores height information, limiting its 3D representation capacity. BEVControl [46] addresses this by introducing a height-lifting module to partially restore scene geometry. MagicDrive [10] advances 3D-awareness through geometric constraints and cross-view attention, while MagicDriveDiT [9] leverages diffusion transformers for improved temporal fidelity. MagicDrive3D [8] further employs deformable Gaussian splatting for coarse 3D reconstruction. DriveDreamer [39] proposes hybrid Gaussians for temporally consistent synthesis of complex maneuvers. Voxel-based approaches such as WoVoGen [25] and Drive-WM [40]




Figure 2: **Overview of the Genesis architecture for joint video and LiDAR generation.** A dual-branch design processes shared semantic conditions via camera and LiDAR pathways, using STDiT blocks for spatiotemporal generation and a BEV encoder for geometric alignment.

explore 4D voxel diffusion and latent substitution to model spatiotemporal dynamics. Despite progress, most methods remain unimodal and rely on weak or decoupled temporal priors, limiting their ability to generate globally consistent, semantically grounded sequences.

LiDAR Point Cloud Generation. LiDAR point clouds form the perceptual cornerstone of autonomous driving, enabling precise 3D understanding through sparse yet geometrically rich measurements. PointNet [31] and VoxelNet [52] laid the foundation for point- and voxel-based LiDAR processing. Building on these, neural radiance fields (NeRF) [28, 45, 41] incorporate LiDAR priors, improving training efficiency and geometric fidelity. Beyond static scenes, recent work emphasizes capturing spatiotemporal dynamics and generating high-fidelity, temporally consistent LiDAR sequences. Copilot4D [48] models long-range dependencies via LiDAR tokenization and hierarchical Transformers. ViDAR [47] employs video diffusion models to generate temporally consistent LiDAR sequences, while [53, 32] employ denoising and conditional diffusion to enhance geometry and consistency. While recent methods begin to capture temporal dynamics, they typically remain unimodal and underexplore semantic integration. In contrast, our framework jointly generates semantically grounded, spatiotemporally coherent sequences across both video and LiDAR modalities.

3 Method

In this section, we present Genesis, a unified generation framework designed to jointly synthesize multi-view video and LiDAR point cloud data with fine-grained semantic consistency.

Overview. As shown in Fig. 2, **Genesis** is architected as a unified world model with a dual-branch design that jointly synthesizes multimodal driving scenes, including videos and LiDAR point clouds. The camera branch leverages a DiT-based spatiotemporal diffusion backbone with a 3D-VAE encoder for fine-grained visual dynamics (Sec. 3.2), while the LiDAR branch employs a BEV-aware autoencoder with NeRF-style rendering and adaptive sampling for accurate geometric reconstruction (Sec. 3.3). Both branches are jointly conditioned on scene layouts, intermediate video latent features, and structured scene semantics to maintain strong cross-modal consistency. Semantic priors such as captions are derived from our DataCrafter module (Sec. 3.1) and injected as global conditioning to enhance alignment across modalities. More architectural details are provided in the supplementary material (Sec. A).

3.1 DataCrafter Module for Scene-Level Semantics

We propose **DataCrafter**, a structured captioning framework designed for multi-view autonomous driving videos. As illustrated in Fig. 3, given multi-view input $V = \{V_1, \dots, V_K\}$, we segment it




Figure 3: **DataCrafter pipeline for structured multi-view captioning.** Videos are segmented and filtered via a VLM-based quality checker (1–2), then per-view captions are generated and fused into coherent, structured descriptions (3–4). Steps 1–4 are used during training, while only 3–4 are used at inference.

into clips $\mathcal{C} = \{c_1, \dots, c_N\}$ via a scene boundary detector. Each clip is scored by a visual-language model (VLM)-based module:

$$S(c_i) = \lambda_1 Q_{\text{clarity}}(c_i) + \lambda_2 Q_{\text{structure}}(c_i) + \lambda_3 Q_{\text{aesthetics}}(c_i), \quad (1)$$

where Q terms represent VLM-derived subscores and λ_i are fixed weights.

To ensure consistency across overlapping views, we introduce a *multi-view consistency module*. Let $\mathcal{Y}_k$ denote the caption from view V_k . Each caption is encoded via a pre-trained VLM:

$$\hat{\mathcal{Y}} = \mathcal{F}(\{\phi(\mathcal{Y}_k)\}_{k=1}^K), \quad (2)$$

where $\phi(\cdot)$ is a language encoder and $\mathcal{F}(\cdot)$ performs summarization and redundancy removal.

Finally, we produce a structured caption for each clip c_i :

$$\mathcal{Y}_i = \left\{ E_i, \{(o_j, b_j, d_j)\}_{j=1}^{M_i} \right\}, \quad (3)$$

where E_i encodes global scene context (e.g., weather, road, time), and each object is represented by its category o_j , bounding box $b_j = (x_{1j}, y_{1j}, x_{2j}, y_{2j})$, and grounded description d_j .

This hierarchical design enables coherent and detailed scene-level and object-level captioning across time and views. More information about DataCrafter can be found in Sec. A.




Figure 4: **Left:** STDiT-Block-C of Camera branch. **Right:** STDiT-Block-L of Lidar branch.

3.2 Video Generation Model

To ensure coherent multi-view video generation, a DiT-based diffusion backbone is extended with 3D-aware latent encoding and conditioned on scene-level priors—including road topology and linguistic descriptions via attention mechanisms, enabling spatial alignment, temporal consistency, and semantic fidelity.

Structured Semantic Priors and Cross-View Conditioning. To enable view-consistent and semantically grounded video generation, a structured BEV layout $\mathcal{S}^l = \{\mathcal{L}, \mathcal{H}, \mathcal{B}\}$ is constructed as




Figure 5: Illustration of the scene layouts projection pipeline, which enables view-consistent and semantically grounded video generation.

shown in Fig. 5, where $\mathcal{L}$, $\mathcal{H}$, and $\mathcal{B}$ denote lane segments, human pose keypoints, and 3D vehicle bounding boxes, respectively. Each element in S^l is projected onto the 2D image plane of view v using calibrated intrinsics K_v and extrinsics E_v , resulting in a set of view-specific semantic control maps $\mathcal{M}_v \in \mathbb{R}^{H \times W \times C}$. These projected maps—comprising layout curves, pose skeletons, and instance masks—are subsequently encoded and fed into a *Control-DiT* module.

Specifically, each block of the spatiotemporal diffusion transformer receives guidance from $\mathcal{M}_v$ via cross-attention, allowing the model to integrate structured priors at every denoising timestep. The forward process at timestep t for view v is formulated as:

$$z_v^{(t)} = \text{DiT}_\theta(z_v^{(t-1)}) + \text{CrossAttn}(z_v^{(t-1)}, \mathcal{M}_v), \quad (4)$$

where DiT_θ denotes the base diffusion transformer and CrossAttn represents the conditional modulation induced by $\mathcal{M}_v$.

Pedestrian Pose Enhancement. To enhance dynamic scene semantics, human keypoints are detected using YOLOv8x-Pose and projected into each view, forming semantic channels that augment layout and instance conditions.

Latent Encoding via 3D VAE. To ensure spatiotemporal consistency of control signals, a 3D Variational Autoencoder (VAE) is employed to compress the multi-frame scene layouts S^l and images into a latent representation $z \in \mathbb{R}^{f \times c \times h \times w}$. The decoder reconstructs the RGB images from the denoised latents.

Scene-Level Caption Conditioning. To encode high-level semantics, descriptive captions are generated via the *DataCrafter* module and processed using a pre-trained T5 encoder:

$$e_{\text{cap}} = \mathcal{E}_{\text{text}}(\text{caption}), \quad z_s = \mathcal{E}_{\text{layout}}(s). \quad (5)$$

Both embeddings are used as conditioning inputs via cross-attention in the DiT blocks:

$$\hat{z}_i = \text{CrossAttn}(q = z_i, k = [e_{\text{cap}}, z_s], v = [e_{\text{cap}}, z_s]). \quad (6)$$

Spatiotemporal Control via STDiT-Block-Cam. As shown in Fig. 4, our model integrates a Semantic-Aligned Control Transformer to modulate early-stage diffusion blocks. semantic features are injected via control attention:

$$h_i^{\text{ctrl}} = \text{Attn}(h_i^{\text{base}}, z_s), \quad i = 1, \dots, K. \quad (7)$$

This module works in parallel with the core DiT pipeline, which includes spatial self-attention, cross-view attention, and temporal attention blocks to ensure view consistency and motion smoothness.

3.3 Lidar Generation Model

Lidar generation model is composed of two modules: a Point Cloud AutoEncoder Module for reconstructing point clouds and a SpatioTemporal Diffusion Module for generating Bird’s-Eye-View (BEV) representation.

Point Cloud AutoEncoder. We design a BEV representation autoencoder to compress sparse and unordered point clouds into structured latent embedding. Inspired by [49, 48], we first voxelize the point clouds into BEV grids and then employ a Swin Transformer [24] backbone to downsample the BEV features by a factor of $8 \times$ in spatial aspect and output the latent hidden feature with channel

4, which is further used in the SpatioTemporal Diffusion Module, as shown in Fig. 4. We decode point clouds from this latent embedding via a Swin-based decoder and a NeRF [28]-based rendering module. We adopt spatial skipping algorithm [48] to avoid the accumulated errors introduced by empty grids. The total training losses include vanilla depth L1 loss [15], occupancy loss [27], and surface regularization loss [17]. Furthermore, we propose a simple post-process algorithm to filter those noisy points generated from nerf rendering process. Specifically, we use filter point cloud that falls within grids, the occupancy value of which is smaller than the threshold.

Spatiotemporal Diffusion with ControlNet-Branch. With the latent embedding from the Point Cloud AutoEncoder Module, we design a diffusion network to learn this representation to generate realistic point clouds. Similar to the structure adopted in the video generation, we use a dual DiT-based network with ControlNet assistance. For semantic condition input, we use scene captions and scene layouts to provide structured scene priors. To keep cross modal consistency, we use the camera branch’s RGB video outputs as another condition input. We adopt widely used LSS [29] algorithm to convert the perspective images to BEV feature. For geometric condition input, we encode 3D bounding boxes into embeddings for providing accurate foreground information. The image BEV feature is concatenated with the scene layouts latent feature and then send into the ControlNet branch as input. For a given latent token z_i at block i , the cross-attention operation could be formulated as:

$$\hat{z}_i = \text{CrossAttn}(q = z_i, k = [e_{\text{cap}}, e_{\text{box}}], v = [e_{\text{cap}}, e_{\text{box}}]), \quad (8)$$

where the embeddings e_{cap} and e_{box} are derived from road sketches and 3D bounding boxes respectively. To ensure temporal coherence, the STDiT-Block-L applies multi-head self-attention operation. Given input z' , the tokens are updated as $\bar{z} = \text{MHSA}(z') + z'$. The entire LiDAR diffusion model is trained using a rectified flow schedule [6] to enhance generation quality.

4 Experiments and Main Results

4.1 Implementation details

Training and evaluation are conducted on the nuScenes [4] dataset, which includes 1,000 urban driving scenes (700 train / 150 val / 150 test). Semantic occupancy labels at 2Hz are interpolated to 12Hz for dense supervision [10, 20]. Multimodal clips are sampled and evaluation follows the standard protocol [10, 20, 11] using 5,369 and 6,019 validation clips. Additional implementation details and training strategy are provided in the supplementary material (Sec. A).

4.2 Qualitative comparison and Versatile Generation Ability

High-quality video samples generated by our method demonstrate strong geometric fidelity and visual coherence across challenging scenes. As shown in Fig. 6 and 10, our model faithfully preserves vehicle shapes, lane structures, and environmental textures. In contrast, MagicDrive suffers from object deformation and layout artifacts, while Panacea exhibits hallucinated content and distorted backgrounds.

Beyond visual realism, the proposed framework enables controllable and diverse scene synthesis. By editing input captions, global scene attributes such as lighting and weather can be modified, as illustrated in Fig. 9. Furthermore, altering the ego-vehicle trajectory in the layout facilitates novel view generation, yielding coherent scene continuations with consistent geometry and appearance across perspectives, as shown in Fig. 11.

4.3 Video Generation Results

Tab. 1 reports the quantitative comparison of video generation performance under various generation modes. We evaluate our model in three configurations: without first-frame conditioning, with first-frame conditioning, and with noisy latent initialization. In all settings, our method consistently achieves superior performance across both FVD and FID metrics.

In the setting without first-frame conditioning, our method achieves an $\text{FVD}_{\text{multi}}$ of 83.10 and a $\text{FID}_{\text{multi}}$ of 14.90, outperforming prior works such as DriveDreamer-2 [50], MagicDrive-V2 [10], and Drive-WM [40]. With first-frame conditioning, our method further improves to 16.95 FVD and 4.24 FID, showing competitive results compared to MiLA [38] while maintaining temporal




Figure 6: **Qualitative comparison of video generation.** From top to bottom: (1) Ground-truth images, (2) Scene Layouts input, (3) Panacea [42], (4) MagicDrive [10], (5) Ours. Panacea suffers from hallucinated textures and geometric misalignment. MagicDrive shows vehicle distortion and broken structures. In contrast, **ours** preserves accurate layout, object shapes, and background integrity.

Table 1: Video Generation Comparison on nuScenes validation set, where green and blue represent the best and the second best values.

Method	Gen. Mode	Multi-view	Video	Sample Num	Frame Num	FVD _{multi} ↓	FID _{multi} ↓
DriveDreamer-2 [50]	w/o first cond	✓	✓	–	–	105.10	25.00
MagicDrive-V2 [9]	w/o first cond	✓	✓	–	16	94.84	20.91
Drive-WM [40]	w/o first cond	✓	✓	–	–	122.70	15.80
Ours	w/o first cond	✓	✓	5369	16	83.10	14.90
MiLA [38]	w/ first cond	✓	✓	5369	16	18.20	3.00
DriveDreamer-2 [50]	w/ first cond	✓	✓	–	–	55.70	11.20
Ours	w/ first cond	✓	✓	5369	16	16.95	4.24
UniScene [20]	w/ noisy latent	✓	✓	6019	–	70.52	6.12
Ours	w/ noisy latent	✓	✓	6019	16	67.87	6.45

consistency and structural fidelity. Under the noisy latent setting, we achieve 67.87 FVD and 6.45 FID on 6,019 samples, surpassing the previous best results reported by UniScene [20].

4.4 LiDAR Generation Results

Tab. 2 reports a quantitative comparison of LiDAR sequence generation performance between prior state-of-the-art methods and our proposed *Genesis* framework. Evaluation is conducted following the HERMES [51] protocol, using Chamfer Distance as the primary metric within a spatial volume of $[-51.2, 51.2]$ meters in the horizontal plane and $[-3, 5]$ meters in height.

Table 2: Lidar Generation Comparison on nuScenes validation set, where green and blue represent the best and the second best values. "gt_img" and "gen_img" indicate using ground-truth or generated images as BEV condition input, respectively.

Method	Condition Type	Chamfer@1s ↓	Chamfer@2s ↓	Chamfer@3s ↓
4D-Occ [19]	gt_img	1.13	1.53	2.11
ViDAR [47]	gt_img	1.12	1.38	1.73
HERMES [51]	gt_img	0.78	0.95	1.17
Ours	gt_img	0.611	0.625	0.633
Ours	gen_img	0.634	0.638	0.641



Figure 7: **Joint generation of LiDAR and multi-view video.** Our method generates spatially aligned LiDAR and camera views conditioned on a shared right turn T-junction layout.



Figure 8: **Comparison with and without DataCrafter.** Our method yields semantically rich, structurally consistent scenes, while the baseline lacks fine-grained details and geometric accuracy.

Genesis consistently outperforms existing approaches across both short- and long-term horizons. At 1 second, it achieves a Chamfer Distance of 0.611, surpassing the previous best (0.78 by HERMES) by 21%. At 3 seconds, the advantage widens to a 45% relative reduction (from 1.17 to 0.633), underscoring the model’s ability to maintain geometric fidelity over extended temporal predictions. Notably, the performance remains stable when replacing ground-truth images with generated images as conditional input, indicating strong robustness and effective cross-modal generalization. These results highlight the efficacy of our structured semantic conditioning and hierarchical generation design, establishing *Genesis* as a new state-of-the-art for LiDAR sequence synthesis, especially in long-horizon prediction scenarios where geometric consistency is critical.

4.5 Joint Generation Results

The proposed framework enables coherent joint generation of LiDAR and multi-view videos with consistent semantic and spatial alignment. As illustrated in Fig. 7, 12 and 13, our model maintains cross-modal correspondence even in complex scenes and over long temporal horizons (Fig. 14, 15). Dynamic objects and road structures are faithfully preserved across modalities, demonstrating the model’s ability to synthesize structured scene dynamics with high temporal and geometric consistency.

4.6 Ablation Studies

We conduct an ablation study to assess the effectiveness of the proposed *DataCrafter* and *PoseNet* modules, which introduce structured semantic captions into the video generation pipeline. As shown in Tab. 3, removing *DataCrafter* causes a notable degradation in generation quality, with FVD_{multi} increasing from 85.91 to 117.49 and FID_{multi} rising from 15.20 to 22.32. Including the *PoseNet* module further improves performance, reducing FVD_{multi} to 83.10 and FID_{multi} to 14.90. These results underscore the importance of semantic guidance for maintaining temporal consistency and structural fidelity. As illustrated in Fig. 8, caption-conditioned generation produces sharper structures and more coherent layouts, whereas the ablated variant exhibits missing objects and visual distortions. In addition, *PoseNet* contributes fine-grained dynamic cues by localizing human anatomical keypoints.

Another ablation study on LiDAR generation is presented in Tab. 4 to assess the contributions of BEV latent features and first-frame conditioning. Removing first-frame input leads to a consistent performance drop across all prediction horizons (e.g., Chamfer@1s increases from 0.634 to 0.668), highlighting its importance for temporal guidance. Further excluding BEV latent features causes additional degradation, indicating their role in preserving spatial structure. These results demonstrate that both components are critical for achieving geometrically consistent long-range LiDAR synthesis.

Table 3: Ablation in the video generation model.

Method	Gen. Mode	Sample Num	Frame Num	FVD _{multi} ↓	FID _{multi} ↓
baseline	w/o first cond	5369	16	117.49	22.32
w/ DataCrafter	w/o first cond	5369	16	85.91	15.20
w/ DataCrafter and PoseNet	w/o first cond	5369	16	83.10	14.90

Table 4: Ablation in the Lidar generation model.

Method	Chamfer@1s ↓	Chamfer@2s ↓	Chamfer@3s ↓
w/o Img BEV Latent + w/o Ref. Frame	0.661	0.669	0.673
w/ Img BEV Latent + w/o Ref. Frame	0.668	0.672	0.677
w/ Img BEV Latent + w/ Ref. Frame	0.634	0.638	0.641

4.7 Domain gap and Downstream task Utility

Our evaluation approach examines generated videos from two perspectives: the degree of reality measured by domain map metrics, and the functional utility demonstrated in downstream training applications such as 3D object detection tasks.

Table 5: Domain gap on bev segmentation.

Method	mIoU↑	mAP↑
MagicDrive [10]	18.34	11.86
MagicDrive3D [8]	18.27	12.05
MagicDriveDiT [9]	20.40	18.17
DiVE [18]	35.96	24.55
Cogen [16]	37.80	27.88
Ours	38.01	27.90

Table 6: Effect of Multimodal Data Generation on 3D Object Detection. Inputs of the methods in the table are both camera and lidar modalities.

Method	mAP↑	NDS↑
Baseline	66.87	69.65
Ours(+cam_gen)	67.09 (+0.22)	70.12 (+0.47)
Ours(+lidar_gen)	67.69 (+0.82)	70.58 (+0.93)
Ours(+cam&lidar_gen)	67.78 (+0.91)	71.13 (+1.48)

As shown in Table 5, we generate videos according to conditions from validation set to metric the domain gap following [9]. Our method achieves the best mIoU (38.01) and mAP (27.90) on bev segmentation, surpassing prior methods like DiVE [18] and Cogen [16], demonstrating strong semantic alignment and visual fidelity. As shown in Tab.6, we evaluate the effectiveness of our generative data on the BEVFusion[22] framework for 3D object detection. Our approach yields consistent improvements across all settings, increasing the mAP from 66.87 to 67.78 and NDS from 69.65 to 71.13. Notably, joint generation of both camera and LiDAR modalities achieves the highest gains (+0.91 mAP / +1.48 NDS), demonstrating the complementary benefits of multimodal generation. These results validate the utility of high-quality synthetic data in enhancing downstream perception tasks, especially in data-scarce or long-tail scenarios.

5 Conclusion

We propose **Genesis**, a unified world model for joint multi-view video and LiDAR point cloud generation in autonomous driving. Integrating structured semantic priors via the *DataCrafter* module and ensuring cross-modal alignment through a shared conditioning pipeline, it bridges visual and geometric modalities to enable high-fidelity, spatiotemporally coherent, and semantically consistent multimodal sequence synthesis. Experiments on nuScenes confirm Genesis achieves SOTA performance in both modalities, validating its robustness and scalability.

Limitations and future work. Despite Genesis’s strong performance, several limitations remain. Training demands significant computational resources, motivating future efforts in model compression or efficient architectures. The framework is currently evaluated on nuScenes, and generalizing to other domains—such as robotics or aerial perception—requires further adaptation. Addressing these challenges is essential for broader deployment in safety-critical systems.

Acknowledgement: This work was partially supported by the National Natural Science Foundation of China (No. 62276108).

References

- [1] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- [2] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1090–1099, 2022.
- [3] Julie Stephany Berrio, Mao Shan, Stewart Worrall, and Eduardo Nebot. Camera-lidar integration: Probabilistic sensor fusion for semantic mapping. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):7637–7652, 2021.
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024.
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- [7] 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.
- [8] Ruiyuan Gao, Kai Chen, Zhihao Li, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magic-drive3d: Controllable 3d generation for any-view rendering in street scenes. *arXiv preprint arXiv:2405.14475*, 2024.
- [9] Ruiyuan Gao, Kai Chen, Bo Xiao, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrivedit: High-resolution long video generation for autonomous driving with adaptive control. *arXiv preprint arXiv:2411.13807*, 2024.
- [10] Ruiyuan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, and Qiang Xu. Magicdrive: Street view generation with diverse 3d geometry control. *arXiv preprint arXiv:2310.02601*, 2023.
- [11] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. Vista: A generalizable driving world model with high fidelity and versatile controllability. *arXiv preprint arXiv:2405.17398*, 2024.
- [12] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013.
- [13] 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.
- [14] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022.
- [15] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *2019 IEEE winter conference on applications of computer vision (WACV)*, pages 1043–1051. IEEE, 2019.

- [16] Yishen Ji, Ziyue Zhu, Zhenxin Zhu, Kaixin Xiong, Ming Lu, Zhiqi Li, Lijun Zhou, Haiyang Sun, Bing Wang, and Tong Lu. Cogen: 3d consistent video generation via adaptive conditioning for autonomous driving. *arXiv preprint arXiv:2503.22231*, 2025.
- [17] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6001–6010, 2020.
- [18] Junpeng Jiang, Gangyi Hong, Lijun Zhou, Enhui Ma, Hengtong Hu, Xia Zhou, Jie Xiang, Fan Liu, Kaicheng Yu, Haiyang Sun, et al. Dive: Dit-based video generation with enhanced control. *arXiv preprint arXiv:2409.01595*, 2024.
- [19] Tarasha Khurana, Peiyun Hu, David Held, and Deva Ramanan. Point cloud forecasting as a proxy for 4d occupancy forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1116–1124, 2023.
- [20] Bohan Li, Jiazhe Guo, Hongsi Liu, Yingshuang Zou, Yikang Ding, Xiwu Chen, Hu Zhu, Feiyang Tan, Chi Zhang, Tiancai Wang, et al. Uniscene: Unified occupancy-centric driving scene generation. *arXiv preprint arXiv:2412.05435*, 2024.
- [21] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17182–17191, 2022.
- [22] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. *Advances in Neural Information Processing Systems*, 35:10421–10434, 2022.
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [25] Jiachen Lu, Ze Huang, Zeyu Yang, Jiahui Zhang, and Li Zhang. Wovogen: World volume-aware diffusion for controllable multi-camera driving scene generation. In *European Conference on Computer Vision*, pages 329–345. Springer, 2024.
- [26] Enhui Ma, Lijun Zhou, Tao Tang, Zhan Zhang, Dong Han, Junpeng Jiang, Kun Zhan, Peng Jia, Xianpeng Lang, Haiyang Sun, et al. Unleashing generalization of end-to-end autonomous driving with controllable long video generation. *arXiv preprint arXiv:2406.01349*, 2024.
- [27] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019.
- [28] 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.
- [29] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 194–210. Springer, 2020.
- [30] Ethan Pronovost, Meghana Reddy Ganesina, Nouredin Hendy, Zeyu Wang, Andres Morales, Kai Wang, and Nick Roy. Scenario diffusion: Controllable driving scenario generation with diffusion. *Advances in Neural Information Processing Systems*, 36:68873–68894, 2023.
- [31] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

- [32] Haoxi Ran, Vitor Guizilini, and Yue Wang. Towards realistic scene generation with lidar diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14738–14748, 2024.
- [33] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020.
- [34] Shuo Sun, Zekai Gu, Tianchen Sun, Jiawei Sun, Chengran Yuan, Yuhang Han, Dongen Li, and Marcelo H Ang. Drivesceneneg: Generating diverse and realistic driving scenarios from scratch. *IEEE Robotics and Automation Letters*, 2024.
- [35] Alexander Swerdlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *IEEE Robotics and Automation Letters*, 2024.
- [36] Alexander Swerdlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *IEEE Robotics and Automation Letters*, 2024.
- [37] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. 2019.
- [38] Haiguang Wang, Daqi Liu, Hongwei Xie, Haisong Liu, Enhui Ma, Kaicheng Yu, Limin Wang, and Bing Wang. Mila: Multi-view intensive-fidelity long-term video generation world model for autonomous driving. *arXiv preprint arXiv:2503.15875*, 2025.
- [39] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. Drive-dreamer: Towards real-world-drive world models for autonomous driving. In *European Conference on Computer Vision*, pages 55–72. Springer, 2024.
- [40] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024.
- [41] Pengcheng Wei, Li Yan, Hong Xie, Dashi Qiu, Changcheng Qiu, Hao Wu, Yinghao Zhao, Xiao Hu, and Ming Huang. Lidenerf: Neural radiance field reconstruction with depth prior provided by lidar point cloud. *ISPRS Journal of Photogrammetry and Remote Sensing*, 208:296–307, 2024.
- [42] Yuqing Wen, Yucheng Zhao, Yingfei Liu, Fan Jia, Yanhui Wang, Chong Luo, Chi Zhang, Tiancai Wang, Xiaoyan Sun, and Xiangyu Zhang. Panacea: Panoramic and controllable video generation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6902–6912, 2024.
- [43] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023.
- [44] Zhanqian Wu, Kaichong Liang, Dengcheng Liu, and Zhiguo Zhao. Driver lane change intention recognition based on attention enhanced residual-mbi-ilstm network. *IEEE Access*, 10:58050–58061, 2022.
- [45] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5438–5448, 2022.
- [46] Kairui Yang, Enhui Ma, Jibin Peng, Qing Guo, Di Lin, and Kaicheng Yu. Bevcontrol: Accurately controlling street-view elements with multi-perspective consistency via bev sketch layout. *arXiv preprint arXiv:2308.01661*, 2023.

- [47] Zetong Yang, Li Chen, Yanan Sun, and Hongyang Li. Visual point cloud forecasting enables scalable autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14673–14684, 2024.
- [48] Lunjun Zhang, Yuwen Xiong, Ze Yang, Sergio Casas, Rui Hu, and Raquel Urtasun. Copilot4d: Learning unsupervised world models for autonomous driving via discrete diffusion. *arXiv preprint arXiv:2311.01017*, 2023.
- [49] Yumeng Zhang, Shi Gong, Kaixin Xiong, Xiaoqing Ye, Xiao Tan, Fan Wang, Jizhou Huang, Hua Wu, and Haifeng Wang. Bevworlde: A multimodal world model for autonomous driving via unified bev latent space. *arXiv preprint arXiv:2407.05679*, 2024.
- [50] Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Xinze Chen, Guan Huang, Xiaoyi Bao, and Xingang Wang. Drivedreamer-2: Llm-enhanced world models for diverse driving video generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 10412–10420, 2025.
- [51] Xin Zhou, Dingkan Liang, Sifan Tu, Xiwu Chen, Yikang Ding, Dingyuan Zhang, Feiyang Tan, Hengshuang Zhao, and Xiang Bai. Hermes: A unified self-driving world model for simultaneous 3d scene understanding and generation. *arXiv preprint arXiv:2501.14729*, 2025.
- [52] Yin Zhou and Oncel Tuzel. Voxnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018.
- [53] Vlas Zyrianov, Henry Che, Zhijian Liu, and Shenlong Wang. Lidardm: Generative lidar simulation in a generated world. *arXiv preprint arXiv:2404.02903*, 2024.
- [54] Vlas Zyrianov, Xiyue Zhu, and Shenlong Wang. Learning to generate realistic lidar point clouds. In *European Conference on Computer Vision*, pages 17–35. Springer, 2022.

A Technical Appendices and Supplementary Material

(a) DataCrafter Setup

To enable structured semantic supervision during training, the *DataCrafter* module was developed based on the Qwen2-VL 7B vision-language model. To enhance its understanding of driving scenes and its ability to assess sample quality, the base model was fine-tuned using Low-Rank Adaptation (LoRA) on two domain-specific datasets: *OmniDrive* and *CoDALM*. During inference, the fine-tuned model evaluated multi-view clips along three predefined dimensions: *visual quality*—including clarity and aesthetics—and *semantic descriptiveness*. Only samples that met a predefined quality threshold were retained for training, ensuring the exclusion of visually poor or semantically inconsistent sequences. The selected high-quality images were then forwarded to the captioning module. The captioning procedure is detailed in the following table.

You are provided with a set of descriptions of images captured by a vehicle's multiple cameras. Please carefully analyze these descriptions and answer the following questions. Your answer must strictly adhere to the format provided below, including all start and end tags for each field. All responses must be in English.

Time: Select one: "Daytime", "Night", "Indoor", "No visible sign"

Weather: Select one from "Sunny", "Cloudy", "Overcast", "Rain", "Snow" or "Night with no visible sign"

Road Type: Select one from "Highway", "Urban Road", "Rural Road", "Tunnel", "Bridge" or "No visible sign" **Road Surface:** Select one from "Asphalt", "Concrete", "Gravel", "resin (Indoor)" **Lane:** Select one from "No visible sign", "Single Lane", "Dual Lane", "Multi-Lane", "Other(should explain in Details)" and describe specifically what the lanes are.

Environment Type: Select one from the following options: "Highway", "Roundabout", "Intersection", "Ramp", "Tunnel", "Parking Lot", "Urban Road", "Rural Road", "Bridge" or "Other(should explain in Details)"

Surroundings Details: Provide a general description of the environment (e.g., a busy street, a quiet neighborhood).

Traffic: Details: Describe the vehicles present in the scene. Include their spatial relationship to the ego vehicle—for example. Categorized by direction: Front, Left, Right, Behind. Each line follows the format: <position> <object> <action>.

(b) Model Setup

The video generation pipeline is constructed upon a DiT-based architecture. The backbone is initialized with pretrained weights from MagicDrive [9], while the variational autoencoder (VAE) is initialized using the publicly available encoder from CogVideo-XL [14]. In contrast, the LiDAR generation branch is trained entirely from scratch. Both its diffusion backbone and 3D autoencoder are designed and optimized specifically for sparse geometry modeling, enabling independent learning of modality-specific representations without reliance on pretrained vision priors.

(c) Training Strategy & Training Setup

A three-stage curriculum is adopted. (1) In the first stage, image-level generation is trained at 512×768 resolution to warm-start spatial representation learning. (2) The second stage focuses on video synthesis, where a two-phase training protocol is employed. Multi-resolution pretraining is first conducted by gradually increasing input resolution from 144p (144×256) to 900p (900×1600), paired with clip lengths ranging from 128 frames at lower resolutions to 6 frames at higher ones. This is followed by adapter-based fine-tuning at a fixed resolution of 360p (360×640) and 16 frames, where lightweight spatiotemporal adapters are inserted into DiT blocks. (3) The third stage performs joint training of video and LiDAR generation with shared conditioning signals, enabling cross-modal temporal alignment. For inference, video frames are generated at 900p with the first frame observed as input, consistent with prior work [8].

Training is performed with PyTorch using 64 NVIDIA H20 GPUs and mixed-precision acceleration. Stages 1, 2, and 3 are trained for 300, 800, and 200 epochs respectively. The optimizer is AdamW

with a weight decay of 0.01. A cosine annealing learning rate schedule is adopted with linear warm-up over the first 10% of steps. Learning rates are set to 2×10^{-4} for the image and video stages, and 1×10^{-5} for the joint stage. A global batch size of 1024 is used, distributed evenly across GPUs.

(d) Evaluation Metrics

To comprehensively evaluate the quality of multimodal generation, we adopt a suite of metrics tailored to each modality. For video synthesis, we report Fréchet Inception Distance (FID) and Fréchet Video Distance (FVD), which assess spatial and spatiotemporal realism, respectively, in accordance with prior works [13, 37]. For LiDAR sequence generation, we employ Chamfer Distance at 1s, 2s, and 3s horizons. These metrics quantify both geometric fidelity and distributional alignment.

For evaluating downstream 3D perception utility, we adopt mean Average Precision (mAP) and NuScenes Detection Score (NDS) on generated sequences using a BEVFormer & BEVFusion [22] detector, following the protocol in [16]. These metrics provide a practical proxy for assessing the semantic controllability and structural consistency of generated data.

(e) Supplementary visualization and quantitative analysis results



Figure 9: **Controllable generation across time-of-day.** By altering scene-level conditions, our method produces consistent multi-view videos aligned with the same underlying map and object layout, while adapting appearance to represent daytime and nighttime settings.



Figure 10: **Qualitative comparison of video generation quality.** Our method (second column) preserves accurate layout, object shapes, and background integrity. MagicDrive (third column) shows vehicle distortion and broken structures. Panacea (fourth column) often suffers from hallucinated textures and geometric misalignment.




Figure 11: **Trajectory-conditioned novel view synthesis.** Given a ground-truth trajectory (middle), we modify the layout (top) by shifting the ego path 4 meters right (bottom). Our model generates plausible and consistent scenes across all views under these layout changes.



Figure 12: **Joint generation of LiDAR and multi-view video.** Our method generates spatially aligned LiDAR and camera views conditioned on a shared straight street layout.



Figure 13: **Joint generation of LiDAR and multi-view video.** Our method generates spatially aligned LiDAR and camera views conditioned on a Busy junction layout.




Figure 14: Long-term multi-view video generation over 3 seconds in an urban driving scene, conditioned on the straight street layout.




Figure 15: Long-term multi-view video generation over 3 seconds in an urban driving scene, conditioned on the busy junction layout.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly state the main contributions of the paper, including the proposed algorithm, its theoretical properties, and empirical results. These claims are consistently supported by the theoretical analysis and experiments presented in the main body of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Despite Genesis's strong performance in joint video and LiDAR generation, several limitations remain. Training is computationally intensive, suggesting future work on efficient architectures or model compression. Generalization beyond nuScenes, such as to embodied agents or robotics, requires further adaptation. Real-time and interactive generation also remains an open challenge.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not contain formal theoretical results or proofs. The focus is on the design, implementation, and empirical evaluation of the proposed method.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper fully discloses all information necessary to reproduce the main experimental results. Detailed training configurations, including hyperparameters, model architecture, optimization settings, and evaluation protocols, are provided in the appendix to ensure reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The datasets used in our experiments are publicly available, while the code is not yet released. Detailed experimental procedures are provided to ensure the verifiability of our findings.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).

- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper provides comprehensive details on the experimental setup, including dataset splits, hyperparameter settings and optimizer type.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We report the generated metrics for the joint generation task and provide visualizations of the generated LiDAR point clouds and video images to demonstrate the effectiveness of our method. Nevertheless, the current experiments lack error bars and formal statistical significance analysis.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).

- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We report the number of GPUs used during training, each equipped with 96GB of memory.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics and confirm our research fully complies with its principles. No ethical issues arise.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This work primarily focuses on algorithm/model improvements and technical methods. There are currently no direct or significant societal impacts associated with the research; therefore, no discussion on positive or negative societal impacts is included.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not involve the release of models or datasets that pose a high risk of misuse. Therefore, no specific safeguards are necessary.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use publicly available datasets and pretrained models, all of which are properly cited in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [No]

Justification: The new assets introduced in the paper (such as the self-developed training framework and private dataset) have not yet provided public documentation. The main reasons include that the models involve commercial privacy protection.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing experiments or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve research with human subjects or crowdsourcing, and therefore no IRB or equivalent ethical approval is required.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method of this research does not involve any important, original, or non-standard usage of large language models (LLMs). Any use of LLMs was limited to writing assistance and does not impact the scientific contributions of the paper.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.