

Composing Driving Worlds through Disentangled Control for Adversarial Scenario Generation

Yifan Zhan^{1*} Zhengqing Chen^{2*†} Qingjie Wang^{2*}
 Zhuo He³ Muyao Niu¹ Xiaoyang Guo² Wei Yin²
 Weiqiang Ren² Qian Zhang² Yinqiang Zheng^{1†}

¹The University of Tokyo ²Horizon Robotics ³University of Glasgow
 *Equal contribution †Project lead †Corresponding author




Fig. 1: *CompoSIA* is a powerful simulator for synthesizing rare driving scenes. While existing driving world models focus on generating, we enable explicit control over structure, identity, and ego action.

Abstract. A major challenge in autonomous driving is the “long tail” of safety-critical edge cases, which often emerge from unusual combinations of common traffic elements. Synthesizing these scenarios is crucial, yet current controllable generative models provide incomplete or entangled guidance, preventing the independent manipulation of scene structure, object identity, and ego actions. We introduce *CompoSIA*, a compositional driving video simulator that disentangles these traffic factors, enabling fine-grained control over diverse adversarial driving scenarios. To support controllable identity replacement of scene elements, we propose a noise-level identity injection, allowing pose-agnostic identity generation

across diverse element poses, all from a single reference image. Furthermore, a hierarchical dual-branch action control mechanism is introduced to improve action controllability. Such disentangled control enables adversarial scenario synthesis—systematically combining safe elements into dangerous configurations that entangled generators cannot produce. Extensive comparisons demonstrate superior controllable generation quality over state-of-the-art baselines, with a 17% improvement in FVD for identity editing and reductions of 30% and 47% in rotation and translation errors for action control. Furthermore, downstream stress-testing reveals substantial planner failures: across editing modalities, the average collision rate of 3s increases by 173%.

Keywords: Driving World Models · Diffusion Models · Controllable Generation

1 Introduction

End-to-end autonomous driving systems rely on large-scale video data to learn unified representations for perception, prediction, and planning [10, 42, 44, 60]. However, existing driving benchmarks such as nuScenes [7] and Waymo [69] exhibit long-tail distributions, where rare scenarios often arise from unusual combinations of otherwise common traffic factors. These adversarial samples are severely underrepresented, limiting scenario diversity and hindering the robustness of autonomous systems under complex driving situations.

To enable adversarial construction (*e.g.*, a sudden truck cut-in forcing hard braking), a generative model [57, 62] must support fine-grained control of scene structure, identity, and ego actions. However, as summarized in Tab. 1, existing generative approaches expose substantial limitations in controllability. Video editing methods like DriveEditor [43] struggles to synthesize novel-view videos, while ego-action-based ReCamMaster [3] fails with element-wise location and identity control. More unified driving world models such as MagicDrive-V2 [21] inject multiple conditions through shared pathways, leading to entangled generation and noticeable quality degradation with editing. Overall, prior methods either control only a subset of scene factors, or fail to disentangle multiple signals, making it difficult to intentionally construct valuable adversarial scenarios—like “a puppeteer pulling tangled strings”.

To address these limitations, we argue that disentangled control requires factor-specific signals to be injected at different levels of the diffusion process. Building on this insight, we propose **CompoSIA** (**Composite Structure, Identity, and Action**), a diffusion-based video generation framework with three explicit and complementary conditioning branches for structure, identity, and ego action control. This design enables fine-grained independent and compositional control, allowing *CompoSIA* to precisely steer adversarial scenario generation.

For element-wise control, each scene element is manipulated through explicit structure and identity cues that play fundamentally different roles: structure governs geometric layout and motion trajectories, while identity specifies semantic

Table 1: Comparison of control capabilities across general and driving-specific models. Through disentangled conditioning, *CompoSIA* enables more comprehensive and fine-grained control and editing over adversarial scenario generation.

Type	Method	Control Modes			
		Structure	Element ID	Scene ID	Ego-Action
General Controllable Generation	ReCamMaster [3]				✓
	Wan-Move [14]	✓		✓	✓
	LoRA-Edit [18]		✓	✓	
	TTM [67]	✓	✓	✓	
Driving Model	Vista [23]			✓	✓
	Drive-WM [79]			✓	✓
	Panacea [80]	✓		✓	✓
	GAIA-2 [64]	✓		✓	✓
	DriveEditor [43]	✓	✓		
	MagicDrive-V2 [21]	✓			✓
	<i>CompoSIA</i> (Ours)	✓	✓	✓	✓



appearance. In our formulation, a scene element (*e.g.*, a surrounding vehicle or a traffic cone) is defined by a 3D bounding box sequence and a single reference image. This design is edit-friendly, yet raises a practical challenge: how to robustly inject semantic identity from a single image without over-constraining geometry. To address this, we construct replaced-clean clip pairs and train an identity-guided denoiser with identity injection at multiple noise levels. At sampling time, identity strength is tuned by the injection timestep, enabling controllable identity control across diverse element poses and categories.

In contrast, ego action is naturally represented as a frame-wise continuous camera trajectory that governs global temporal dynamics. To model such motion, we introduce a hierarchical dual-branch action conditioner that captures both local residual changes and global trajectory context. The former introduces adaptive layer normalization (AdaLN) [57] for fast early-stage learning, while the latter applies camera attention with PRoPE [39] to provide accurate long-range guidance. Together, the two branches accelerate convergence and improve the precision and stability of action-conditioned generation.

CompoSIA enables precise independent control and compositional editing of structure, identity, and ego action, turning driving world generation into a controllable simulator rather than a data synthesizer. Our design allows deliberate construction of unusual factor combinations that are difficult to obtain from real-world logs and hard to produce with entangled synthesizer like DriveEditor [43] and MagicDrive-V2 [21]. Leveraging this capability, we design control policies to synthesize precise adversarial corner cases for planner stress-testing, exposing hidden failure modes beyond standard benchmarks. Our contributions are three-fold:

- 1) a noise-level identity cue injection coupled with a tailored training strategy to recover consistent element identity from a single image, improving generalization across diverse element poses, categories, and in-the-wild scenes;

- 2) a hierarchical dual-branch action conditioning scheme that models ego motion through complementary local residual modulation and global P_{RoPE} attention, enabling faster convergence and more precise action control;
- 3) *CompoSIA*, a disentangled driving video simulator that enables independent and compositional control of structure, identity, and action, supporting systematic synthesis of adversarial scenarios for downstream stress-testing.

2 Related Works

2.1 Video-based Driving World Models

World models can be broadly instantiated through three paradigms: explicit 3D/4D reconstruction of dynamic scenes [26, 35, 41, 55, 58, 87, 93, 100], latent-space predictive modeling [2, 4, 54], and end-to-end video generation, with the last providing the most direct formulation for future scene synthesis. Recent progress in video generation has been largely driven by diffusion-based models [57, 62]. Foundational video generation models [6, 24, 36, 51, 63, 71, 72, 86] have demonstrated strong capability in synthesizing high-quality, full-length videos conditioned on text prompts or reference images. Autonomous driving scenarios require leveraging more complex sensor modalities to achieve higher-precision controllability. A common line of work [20–22, 77, 79, 80, 85, 97] leverages geometric constraints, such as scene layouts or object-level spatial representations, to provide explicit structural guidance for generation. Other approaches [23, 27, 31, 49, 64, 94, 95] introduce ego-trajectory signals to better guide ego-vehicle behavior. Semantic-guided driving scene generation remains underexplored, with existing approaches primarily relying on coarse-grained controls such as text prompts [31, 64] or weather tags. In addition, feed-forward 3D methods [12, 13, 15, 19, 70, 81, 101] have been explored for driving scene generation, but they often suffer from limited rendering fidelity from novel views.

2.2 Controllable Video Generation

The injection of conditioning signals significantly improves the controllability of image and video generation models, including structure [52, 82, 83, 99], identity [8, 25, 32, 66, 73, 78, 84], scene flow [56, 65], camera pose [3, 75, 90], 3D point clouds [40, 50, 61, 89], and audio signals [28, 37]. Early UNet-based diffusion models rely on designs such as ControlNet [96] or IP-Adapter [88] for conditioning injection. The rise of diffusion transformers [57] has introduced greater flexibility for conditioning injection. Latent concatenation injects conditions into the data stream at fixed positions. Cross-attention layers treat conditions as external context to provide global control over the generation process. More recently, in-context conditioning [34] has been explored, which places both the input and all conditioning signals into a single token sequence, allowing the model to reason over them jointly within a unified framework.

2.3 Image-based Identity Control and Editing

Image-based generation has been extensively studied. ID-Animator [29] employs a ViT-based face adapter with cross-attention for identity conditioning. SkyReels-A2 [17] further integrates multi-view identity information and prompt concepts into a temporally aware DiT model. ConsisID [91] and MagicMe [53] mitigate identity drift by introducing frequency-aware and dynamic attention, respectively. However, for attention mechanisms, enforcing stronger identity constraints can hinder motion expressiveness, while dominant motion cues may distort identity. TTM [67] performs identity-guided editing at the noise level using reference images, while training-free formulation makes it difficult to achieve identity control under accurate motion. To preserve motion correctness and background consistency, identity guidance in autonomous driving [11, 30, 38, 43, 48, 76, 92] is applied conservatively, operating on pose-aligned single image or image sequences for mild harmonization. Such designs typically require time-consuming pose alignment and sacrifice the ability to generate novel scenes.

3 Methods

This section presents the formulation and architectural design of our controllable driving world simulator, *CompoSIA* (see Fig. 2). We first introduce the Flow Matching [1, 45, 46]-based DiT formulation that serves as the backbone of our framework. Built upon this foundation, we describe how three categories of control signals—3D structure, identity, and ego action—are explicitly decomposed and injected during training under a disentangled yet compositional design principle. Finally, we detail the sampling procedure and show how different combinations of conditions enable structured and controllable video generation.

3.1 Preliminary

For efficient video generation, diffusion models are typically trained and sampled in latent spaces [5, 6, 16, 62, 98], enabled by well-designed compression VAEs. Specifically, a diffusion process that transforms samples ϵ drawn from a prior distribution $p_1 = \mathcal{N}(\mathbf{0}, \mathbf{I})$ into latent data samples $z_{(0)}$ is formulated through a differentiable dynamical equation Θ as

$$dz_{(t)} = v_{\Theta}(z_{(t)}, t) dt, \quad t \in [0, 1]. \quad (1)$$

CompoSIA also adopts a VAE-diffusion framework, incorporating recent flow matching formulations [1, 45, 46] and supervised with a v -prediction loss. Intermediate states are constructed as

$$z_{(t)} = \sigma_t z_{(0)} + (1 - \sigma_t) \epsilon, \quad (2)$$

and the objective is

$$\mathcal{L}_{CFM} = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|v_{\Theta}(z_{(t)}, t) - (z_{(0)} - \epsilon)\|_2^2. \quad (3)$$

where σ denotes a predefined noise schedule.




Fig. 2: Overview of the *CompoSIA* framework. During training (top), structure, identity, and ego-action signals are explicitly decomposed and injected into a Flow Matching-based DiT backbone in a disentangled yet compositional manner. During sampling (bottom), different combinations of conditions enable controllable driving video generation, supporting multiple editing applications.

3.2 Structure Conditioning via Spatiotemporal Layouts

To faithfully represent driving scenes, element-wise annotations of both dynamic and static elements (*e.g.*, vehicles and traffic cones) are essential. Incorporating such structured annotations into the generative model substantially enhances element controllability. Given a video clip $x_{(0)} \in \mathbb{R}^{F \times H \times W \times 3}$, we obtain element-wise dynamic 3D bounding box annotations from video frames or LiDAR sensors. Each element is represented by a bounding box sequence $\mathbf{b} \in \mathbb{R}^{F \times 7}$, encoding its 3D position, size, and orientation. However, these spatial annotations lie in a 3D world coordinate system, which is mismatched with the 2D image domain of latent video generation. To bridge this gap, rather than directly embedding $\mathbf{b}$, we first convert each box into a 3D polyline $\mathbf{p}$ and project it onto the image plane via camera intrinsics K and extrinsics E , following

$$\mathbf{b}_f = \text{Pinhole}(\mathbf{p}, K, E), \quad (4)$$

where $\mathbf{b}_f \in \mathbb{R}^{F \times H \times W \times 3}$ is spatiotemporally aligned with $x_{(0)}$.

The diffusion process is performed in the latent space. Given a VAE encoder $\mathcal{E}$, we obtain $z_{(0)} = \mathcal{E}(x_{(0)}) \in \mathbb{R}^{c \times f \times h \times w}$. This latent is perturbed with noise according to the schedule σ_t (Eq. (2)). It is then partitioned into patches of size P , which are flattened into a sequence of latent tokens $\mathbf{h}_{(t)} \in \mathbb{R}^{f \times s \times d}$ ($s = \frac{h}{P} \times \frac{w}{P}$, and d is the channel dimension of latent tokens).

To incorporate the structure cue $\mathbf{b}_f$, we first encode it with $\mathcal{E}$ to obtain $z_{\mathbf{b}_f} \in \mathbb{R}^{c \times f \times h \times w}$. The latent is further processed by a lightweight convolutional adapter and reshaped into layout tokens $\mathbf{h}_{\mathbf{b}_f} \in \mathbb{R}^{f \times s \times d}$. Structure condition is achieved through a zero-initialized projection module

$$\mathbf{h}_{(t)} \leftarrow \mathbf{h}_{(t)} + f_{\text{zero}}(\mathbf{h}_{\mathbf{b}_f}). \quad (5)$$

3.3 Identity Conditioning via Noise Editing

We aim to control the identity of a specific scene element in a video clip using only a single reference image. This remains challenging for existing methods, as prior video editing works [43, 48, 76] lack precise control over element pose and position, while attention-based conditioning alone provides insufficient identity supervision. To this end, we formulate identity control as a restoration problem in the diffusion process: identity cue is injected at noisy stages, and the model is trained to restore a clean, geometrically consistent sequence.

To achieve this, we first construct training pairs using a *select-and-repaint* strategy. For each element in a video clip $x_{(0)}$, we randomly select one frame in which it appears and crop a reference image. Next, for each element, we spatially align its reference image to the sequence of 2D bounding boxes (*i.e.*, enclosing rectangles of projected 3D boxes $\mathbf{b}_f$) across the clip and paste it back frame by frame according to the projected locations, forming the identity cue condition $\mathbf{r}_f \in \mathbb{R}^{F \times H \times W \times 3}$. This step also produces a reference mask $\mathbf{M}_{\mathbf{r}_f} \in \mathbb{R}^{F \times H \times W}$, indicating regions within the 2D bounding box sequences. Now for each element, we construct a training pair consisting of the original natural video and a hard-anchored reference image sequence.

During training, we employ a hard-binding strategy at high-noise stages, replacing the corresponding regions of $x_{(0)}$ with $\mathbf{r}_f$ to enforce identity supervision. Let $z_{\mathbf{r}_f(0)} = \mathcal{E}(\mathbf{r}_f) \in \mathbb{R}^{c \times f \times h \times w}$ and $\mathbf{m}_{\mathbf{r}_f} \in \mathbb{R}^{f \times h \times w}$ denote the encoded identity cue and its downsampled mask. This injection can be formulated as

$$z_{\mathbf{r}_f(t)} = \sigma_t z_{\mathbf{r}_f(0)} + (1 - \sigma_t) \epsilon, \quad (6)$$

$$z_{(t)} \leftarrow \mathbf{m} \odot z_{\mathbf{r}_f(t)} + (1 - \mathbf{m}) \odot z_{(t)}, \quad (7)$$

where

$$\mathbf{m} = \mathbf{m}_{\mathbf{r}_f} \cdot \mathbb{I}(t > T_{\text{id}}). \quad (8)$$

Here, $\mathbb{I}(t > T_{\text{id}})$ indicates that the identity injection is active only at high-noise stages. We avoid hard replacement at low-noise stages, as it would disrupt the model’s generation pathway established during late-stage denoising. This restoration-based formulation enables efficient and stable identity learning during training. As we will detail later, when combined with our sampling strategy, it further achieves dual fidelity in both element identity and pose.

3.4 Hierarchical Dual-Branch Action Conditioning

Ego action is modeled as frame-wise and continuous camera trajectory. In Fig. 3, our action conditioning adopts a hierarchical design with both local and global injection pathways. We observe that local injection facilitates rapid convergence during the early stages of training, while global injection improves the overall accuracy and stability.

Local Residual Modulation. Given a continuous trajectory of shape $F \times 4 \times 4$, we compute the relative transformation $\Delta \mathbf{T}_i = \mathbf{T}_i^{-1} \mathbf{T}_{i+1}$ ($\Delta \mathbf{T}_F = \Delta \mathbf{T}_{F-1}$), from which we extract $\mathbf{a} = (\Delta x, \Delta y, \Delta \text{yaw}) \in \mathbb{R}^{F \times 3}$ to serve as the residual signal. This is injected through a gating mechanism implemented via adaptive layer normalization [57] following $f_{\text{zero}}(\phi(\mathbf{a})) \in \mathbb{R}^{f \times 6 \times d}$, where $\phi(\cdot)$ denotes the sinusoidal frequency embedding and $f_{\text{zero}}(\cdot)$ is a zero-initialized projector. The six channels are split into two groups of three, representing the pre-normalization shift, scale, and post-layer residual gate for the self-attention and feed-forward sublayers, respectively. This local branch enables the model to quickly learn action control during early training.

Global PRoPE Embedding. The local branch inherently loses precise motion information and camera intrinsic cues. To improve the overall accuracy of action control, we embed the global trajectory with Projective Positional Encoding (PRoPE) [39]. Recall the original self attention computation as

$$\text{Attn}_v = \text{Attn}(R^\top \odot Q, R^{-1} \odot K, V), \quad (9)$$

where R is the 3D RoPE [68] and $\odot$ denotes the matrix-vector product. We further compute

$$\text{Attn}_c = D^{\text{proj}} \odot \text{Attn}((D^{\text{proj}})^\top \odot Q_s, (D^{\text{proj}})^{-1} \odot K_s, (D^{\text{proj}})^{-1} \odot V_s). \quad (10)$$

Here, D^{proj} is derived from camera intrinsic and extrinsic parameters.

However, the 3D RoPE design in pre-trained video models and the camera projection design utilize different channel-wise structures. This architectural discrepancy makes it difficult to directly merge relative camera pose embeddings (PRoPE) into the base model without disrupting its generative priors. To address this, we design a lightweight PRoPE attention control branch initialized with zero-convolutions. Furthermore, we observe that camera trajectory signals are inherently low-dimensional, for which a full-dimensional attention space is unnecessary. To minimize computational overhead, we first project the video tokens into a lower-dimensional subspace (e.g., $1/8 \times$) to compute reduced (Q_s, K_s, V_s) . The camera-conditioned attention is efficiently computed within this




Fig. 3: Design of dual-branch action Conditioning. Local residual modulation accelerates early convergence, while global PRoPE embedding improves overall accuracy.




Fig. 4: Identity editing for the first frame. We preserve the background to anchor scene identity while editing the foreground. If the reference area is smaller than the original area, the intermediate region is treated as an inpainting area during denoising.

compact space and subsequently projected back to the original feature dimension, before being safely injected into the main attention branch via the zero-convolution layers. The final attention output is computed as $Attn_v + f_{\text{zero}}(Attn_c)$, where f_{zero} is a zero-initialized projector. This global mechanism enables the model to achieve more precise action control.


3.5 Training–Sampling Design for Condition Decoupling

For the training modalities, we observe that structure cue in Sec. 3.2 can leak ego-action information during training (*e.g.*, when other vehicles gradually move backward, it often indicates that the ego vehicle is moving forward rather than remaining static). To prevent this implicit supervision from biasing the learning process, structure cue is always paired with action conditioning. Further experiments show that structure and identity cues do not conflict with each other and can be jointly trained. In practice, we adopt a training modality ratio of 0.6:0.3:0.1 for [action], [structure, identity, action], and unconditional inputs, respectively. The background region of the first frame is replaced with a clean latent [47], and its denoising timestep is reset to zero to explicitly inject scene identity information. Moreover, we additionally apply localized noise perturbation to the first frame during training to enhance the model’s inpainting capability.

During sampling, independent control of each condition becomes possible. The structure condition can be edited separately (*e.g.*, modifying the motion or presence of other elements). When editing action alone, the structure condition must be explicitly re-projected to ensure that the motion of other elements remains unchanged. The reference image for identity editing can be drawn from the dataset or from in-the-wild images, and may correspond to vehicles or other traffic elements. We now provide a detailed description of the first-frame processing (Fig. 4) in our identity editing pipeline. The background region is explicitly preserved as a clean latent to anchor scene identity. For element identity, when the reference image is smaller than the original edited region, the area between the foreground and background is treated as an inpainting region during denoising. Such inpainting configurations are explicitly exposed during training, enabling robust identity reconstruction under partial spatial replacement. See Sec. 4 for more visualizations.

Table 2: Quantitative comparison of different methods. We evaluate multi-view scene following, identity control, and action control tasks.

Type	Method	FVD↓	VBench Score (%)↑
6v-Scene Following	MagicDrive-V2 [21]	152.80	77.23
	<i>CompoSIA</i> (Ours)	133.66	81.05
Identity Control	TTM [67]	231.17	75.16
	LoRA-Edit [18]	161.32	79.83
	DriveEditor [43]	179.57	79.13
	<i>CompoSIA</i> (Ours)	149.15	80.30
Action Control	ReCamMaster [3]	190.52	74.29
	Vista [23]	171.49	75.35
	MagicDrive-V2 [21]	279.61	73.44
	<i>CompoSIA</i> (Ours)	137.21	80.79

**Fig. 5: Comparison of generation quality.** *CompoSIA* renders vehicles and pedestrians with more structurally coherent geometry, while maintaining stronger cross-view and temporal consistency.

4 Experiments

4.1 Experimental Setups

Implementation Details. Our *CompoSIA* is initialized from Wan2.1-T2V-1.3B [72]. Considering the fast temporal dynamics in driving scenarios, we fine-tune the Wan VAE without temporal downsampling, setting the temporal stride to 1 instead of the original $4\times$ compression. For the DiT module, we additionally introduce a small projector to inject structure condition. Similar projectors are also added for achieving the low-dimensional PRoPE attention and the residual modulation for ego-action conditioning. We train *CompoSIA* on 16 NVIDIA 80GB A100 GPUs, with learning rate 2×10^{-4} for the action projector and 1×10^{-5} for all other components. The weight decay of 5×10^{-2} is applied to all parameters. For training of identity injection, we replace latents with identity image only at high noise levels ($t > 0.2$ in Eq. (1)), to avoid disrupting the denoising path at low noise levels for high-quality video generation.

Dataset. We train *CompoSIA* on nuScenes [7] and our self-collected driving data. For nuScenes, we use 700 multi-view videos of 20 seconds for training and



Fig. 6: Comparison on identity control. *CompoSIA* enables precise identity injection across diverse element poses, categories, and in-the-wild reference images.

150 for validation. To cover a wider range of element types, more diverse scenes, and more complex high-speed driving cases, we further incorporate 100 hours of internal, multi-view autonomous driving data for training. A mixed-resolution training strategy is applied using both $33 \times 256 \times 512$ and $33 \times 480 \times 960$ clips to facilitate faster convergence of the condition modeling. All data are trained at 10 Hz.

4.2 Comparisons of Generation Quality

For video generation, we follow benchmarks from [9], utilizing FVD for evaluating of video quality. We further report VBench Score [33] (Quality Score only) for a more comprehensive and multi-dimensional evaluation. The results are presented in Tab. 2 and Fig. 5. *CompoSIA* renders vehicles and pedestrians with higher fidelity, while maintaining cross-view and temporal consistency.

4.3 Structure and Identity Control

For identity control shown in Fig. 6, the measured video quality is reported in Tab. 2. We apply TTM [67] as a training-free sampling strategy on a model trained on our driving data, which fails to generalize to reference images captured from novel element poses. LoRA-Edit [18] achieves high fidelity via first-frame editing but cannot precisely control the insertion location of the generated element. DriveEditor [43] injects reference images through a conditional branch, yet the transferred identity deviates noticeably from the original. In contrast, *CompoSIA* supports editing diverse elements with controllable positions and orientations, across heterogeneous categories, while maintaining strong identity fidelity to reference images.




Fig. 7: Comparison on action control. *CompoSIA* faithfully executes diverse action commands with precise frame-wise controllability.

4.4 Action Control

In Tab. 3, we evaluate the action control by estimating trajectories from the generated videos using VGGT [74], aligning the predicted trajectories with the ground-truth trajectories. We report

Table 3: Evaluation of action control. TransErr is $1000\times$.

Method	Action Following		Action Editing	
	RotErr ↓	TransErr ↓	RotErr ↓	TransErr ↓
ReCamMaster [3]	1.12	20.35	2.17	25.45
Vista [23]	0.81	14.25	2.33	28.12
MagicDrive-V2 [21]	0.76	13.66	2.21	22.86
<i>CompoSIA</i> (Ours)	0.55	7.37	1.54	12.15

$$\text{RotErr} = \sum_{i=1}^n \arccos \frac{\text{tr}(\mathbf{R}_{gen}^i \mathbf{R}_{gt}^{i\top}) - 1}{2}, \quad (11)$$

$$\text{TransErr} = \sum_{i=1}^n \|\mathbf{T}_{gt}^i - \mathbf{T}_{gen}^i\|_2^2.$$

where $\mathbf{R}_{gen}^i$ and $\mathbf{R}_{gt}^i$ denote the estimated and ground-truth rotation matrices, and $\mathbf{T}_{gen}^i$ and $\mathbf{T}_{gt}^i$ denote the corresponding translation vectors, all at the i -th frame. We compare MagicDrive-V2 [21] (trajectory embedding), Vista [23] (trajectory cross-attention), and ReCamMaster [3] (3D attention) under the same first-frame image condition. For each method, we modify the action to perform different maneuvers, including left lane change, right lane change, stop, and accelerating straight. As shown in Fig. 7, only our approach accurately follows the frame-wise action sequence.

4.5 Ablation Study

Ablating Structure Control. To demonstrate the role of structure control, we disable structure conditioning while retaining action guidance during sampling




Fig. 8: Ablation studies on structure, action, and identity injection. (a) Structure ablation. (b) Action ablation. (c) Effect of T_{id} . (d) Identity injection improves cross-frame consistency under complex illumination.

in Fig. 8 (a). As a result, surrounding vehicles fail to maintain consistent motion and spatial alignment.

Ablating Action Control. To verify that action control is effective rather than ego motion being leaked from structure cues, we remove the action branch during sampling. In Fig. 8 (b), the structure conditioning are kept identical to the ground truth while no action control is provided. Without which the ego motion becomes unstable. We further ablate different action injection strategies (r.m. for residual modulation and p.a. for PRoPE attention) in Tab. 4 to study their impact on controllability, showing that our hierarchical dual-branch design achieves the best performance.

Design of T_{id} for Identity Control. We further ablate the effect of different T_{id} , the stopping step of identity injection, which directly influences the balance between editing freedom and identity preservation. As can be found in Fig. 8 (c), earlier stopping ($T_{id} = 0.6$) increases generation freedom but reduces similarity to the reference image, whereas later stopping ($T_{id} = 0.2$) enforces stronger identity preservation but increasingly anchors the generation to the reference image. Nevertheless, since identity cues are injected during training across a wide range of noise levels ($t > 0.2$), the stopping step remains robust for different elements. Therefore, no per-case tuning is required, and we adopt $T_{id} = 0.4$ as a practical trade-off.

Effect of Identity Cue on ID Preservation. Beyond its role in controllable editing, identity cue injection also improves ID stability during generation. As shown in Fig. 8 (d), in scenarios with complex illumination transitions (*e.g.*, driving through tunnels), relying solely on the first-frame reference leads to inconsistent scene element appearances across frames. In contrast, injecting identity cues throughout the denoising process substantially reduces identity drift. This noise-level identity conditioning offers a principled and temporally coherent mechanism for stabilizing identity without sacrificing generative flexibility.

Table 4: Ablation of action conditioning. TransErr is 1000 $\times$.

Condition	Action Following	
	RotErr $\downarrow$	TransErr $\downarrow$
w/o r.m.	0.71	10.68
w/o p.a.	0.62	11.24
Full	0.55	7.37

Table 5: End-to-end motion planning performance on nuScenes [7].

Type	L2(m) ↓				Collision Rate(%) ↓			
	1s	2s	3s	Avg.	1s	2s	3s	Avg.
Following GT	0.68	1.33	2.24	1.42	0.04	0.24	0.76	0.35
Following Generation	0.72	1.41	2.83	<u>1.65</u>	0.08	0.36	1.32	<u>0.59</u>
Editing Structure	—	—	—	—	0.72	2.68	5.28	2.89 ↑390%
Editing Identity	0.89	1.77	3.91	2.19 ↑33%	0.12	0.48	1.64	0.75 ↑27%
Editing Action	0.91	1.83	4.22	2.32 ↑41%	0.16	0.76	2.64	1.19 ↑102%

4.6 Evaluation of Planning Robustness

We conduct planning robustness evaluation on nuScenes [7]. Specifically, we synthesize multiple groups of data. The first group strictly follows the original structure, identity and action, serving to isolate the impact of potential video quality degradation on planner decisions. Other groups edit both surrounding agents and ego actions, together with identity replacement, to construct more challenging scenarios. These are designed to demonstrate that our synthesized data represent long-tail and adversarial cases for planning. We adopt Epona [95] as the end-to-end planner and evaluate all groups in an open-loop setting, reporting L2 distance and collision rate based on the predicted trajectories. Tab. 5 shows noticeable metric degradation relative to the generated baseline, indicating that the introduced edits increase task complexity and challenge the planner’s generalization capability. Detailed editing modalities can be found in the supplementary materials.

5 Limitations

Despite its effectiveness, *CompoSIA* has some limitations. First, our training data primarily consists of driving scenes, which restricts identity editing generalization to total out-of-distribution element categories (*e.g.*, animals). Although our identity formulation does not explicitly constrain category type, broader generalization would benefit from scaling training data to include more diverse video sources. Second, our current identity editing pipeline requires specifying an approximate 3D bounding box size (*i.e.*, height, width, length) for the reference image during replacement or insertion. At present, this estimation is facilitated by a large language model (*e.g.*, Gemini [59]), but the process remains semi-automatic. Future work may explore more user-friendly editing interfaces with automatic spatial grounding and interactive refinement.

6 Conclusion

In this paper, we present *CompoSIA*, a compositional driving world simulator to generate adversarial Scenario through unified and disentangled control. By introducing factor-specific conditioning branches, our framework enables both independent and joint manipulation of scene layout, object identity, and ego actions.

Extensive experiments demonstrate that *CompoSIA* consistently surpasses prior driving world models, driving-specific editing approaches, and general controllable generation frameworks in both generation fidelity and condition alignment. Building upon these capabilities, we construct controllable adversarial scenarios for systematic planner stress-testing, uncovering robustness limitations that remain hidden under conventional benchmarks. We believe *CompoSIA* advances controllable driving world modeling and provides a principled foundation for reliable scenario synthesis and robustness auditing in autonomous driving systems.

Acknowledgements

This work was supported in part by JSPS KAKENHI Grant Number 24KK0209 and by the Forest Digital Twin Project under the Partnership Agreement for Social Value Creation between UTokyo and SMBC Group. We also thank the BOOST NAIS Program for supporting this research.

References

1. Albergo, M.S., Vanden-Eijnden, E.: Building normalizing flows with stochastic interpolants. arXiv preprint arXiv:2209.15571 (2022) [5](#)
2. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., Ballas, N.: Self-supervised learning from images with a joint-embedding predictive architecture. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15619–15629 (2023) [4](#)
3. Bai, J., Xia, M., Fu, X., Wang, X., Mu, L., Cao, J., Liu, Z., Hu, H., Bai, X., Wan, P., et al.: Recammaster: Camera-controlled generative rendering from a single video. arXiv preprint arXiv:2503.11647 (2025) [2](#), [3](#), [4](#), [10](#), [12](#)
4. Balestrieri, R., LeCun, Y.: Lejepa: Provable and scalable self-supervised learning without the heuristics. arXiv preprint arXiv:2511.08544 (2025) [4](#)
5. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023) [5](#)
6. Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., et al.: Video generation models as world simulators. OpenAI Blog **1**(8), 1 (2024) [4](#), [5](#)
7. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020) [2](#), [10](#), [14](#)
8. Chen, B., Zhao, M., Sun, H., Chen, L., Wang, X., Du, K., Wu, X.: Xverse: Consistent multi-subject control of identity and semantic attributes via dit modulation. arXiv preprint arXiv:2506.21416 (2025) [4](#)
9. Chen, K., Gao, R., Hong, L., Xu, H., Jia, X., Caesar, H., Dai, D., Liu, B., Tsishkou, D., Xu, S., et al.: Eccv 2024 w-coda: 1st workshop on multimodal perception and comprehension of corner cases in autonomous driving. arXiv preprint arXiv:2507.01735 (2025) [11](#)

10. Chen, S., Jiang, B., Gao, H., Liao, B., Xu, Q., Zhang, Q., Huang, C., Liu, W., Wang, X.: Vadv2: End-to-end vectorized autonomous driving via probabilistic planning. arXiv preprint arXiv:2402.13243 (2024) [2](#)
11. Chen, Z., Mei, R., Guo, X., Wang, Q., Hu, Y., Yin, W., Ren, W., Zhang, Q.: Synthdrive: Scalable real2sim2real sensor simulation pipeline for high-fidelity asset generation and driving data synthesis. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 912–919. IEEE (2025) [5](#)
12. Cheng, C., Chen, X., Xie, T., Yin, W., Ren, W., Zhang, Q., Guo, X., Wang, H.: Longstream: Long-sequence streaming autoregressive visual geometry. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 272–283 (2026) [4](#)
13. Cheng, C., Tao, P., Yao, N., Ding, G., Chen, X., Du, Y., Guo, X., Yin, W., Ren, W., Zhang, Q., et al.: Horizonstream: Long-horizon attention for streaming 3d reconstruction. arXiv preprint arXiv:2605.23889 (2026) [4](#)
14. Chu, R., He, Y., Chen, Z., Zhang, S., Xu, X., Xia, B., Wang, D., Yi, H., Liu, X., Zhao, H., et al.: Wan-move: Motion-controllable video generation via latent trajectory guidance. arXiv preprint arXiv:2512.08765 (2025) [3](#)
15. Deng, J., Li, H., Xie, T., Ren, W., Zhang, Q., Tan, P., Guo, X.: Sail-recon: Large sfm by augmenting scene regression with localization. In: 2026 International Conference on 3D Vision (3DV). pp. 309–329. IEEE (2026) [4](#)
16. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024) [5](#)
17. Fei, Z., Li, D., Qiu, D., Wang, J., Dou, Y., Wang, R., Xu, J., Fan, M., Chen, G., Li, Y., et al.: Skyreels-a2: Compose anything in video diffusion transformers. arXiv preprint arXiv:2504.02436 (2025) [5](#)
18. Gao, C., Ding, L., Cai, X., Huang, Z., Wang, Z., Xue, T.: Lora-edit: Controllable first-frame-guided video editing via mask-aware lora fine-tuning. arXiv preprint arXiv:2506.10082 (2025) [3](#), [10](#), [11](#)
19. Gao, H., Chen, S., Jiang, B., Liao, B., Shi, Y., Guo, X., Pu, Y., Li, X., Liu, W., Zhang, Q., et al.: Rad: Training an end-to-end driving policy via large-scale 3dgs-based reinforcement learning. Advances in Neural Information Processing Systems **38**, 32551–32576 (2026) [4](#)
20. Gao, R., Chen, K., Li, Z., Hong, L., Li, Z., Xu, Q.: Magicdrive3d: Controllable 3d generation for any-view rendering in street scenes. arXiv preprint arXiv:2405.14475 (2024) [4](#)
21. Gao, R., Chen, K., Xiao, B., Hong, L., Li, Z., Xu, Q.: Magicdrive-v2: High-resolution long video generation for autonomous driving with adaptive control. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 28135–28144 (2025) [2](#), [3](#), [4](#), [10](#), [12](#)
22. Gao, R., Chen, K., Xie, E., Hong, L., Li, Z., Yeung, D.Y., Xu, Q.: Magicdrive: Street view generation with diverse 3d geometry control. arXiv preprint arXiv:2310.02601 (2023) [4](#)
23. Gao, S., Yang, J., Chen, L., Chitta, K., Qiu, Y., Geiger, A., Zhang, J., Li, H.: Vista: A generalizable driving world model with high fidelity and versatile controllability. Advances in Neural Information Processing Systems **37**, 91560–91596 (2024) [3](#), [4](#), [10](#), [12](#)
24. Gao, Y., Guo, H., Hoang, T., Huang, W., Jiang, L., Kong, F., Li, H., Li, J., Li, L., Li, X., et al.: Seedance 1.0: Exploring the boundaries of video generation models. arXiv preprint arXiv:2506.09113 (2025) [4](#)

25. Gong, Y., Pang, Y., Cun, X., Xia, M., He, Y., Chen, H., Wang, L., Zhang, Y., Wang, X., Shan, Y., et al.: Talecrafter: Interactive story visualization with multiple characters. arXiv preprint arXiv:2305.18247 (2023) [4](#)
26. Gong, Y., Liu, Y., Zhan, Y., Niu, M., Li, X., Liao, Y., Chen, J., Gao, Y., Chen, J., Chen, M., et al.: Visionary: The world model carrier built on webgpu-powered gaussian splatting platform. arXiv preprint arXiv:2512.08478 (2025) [4](#)
27. Guo, X., Ding, C., Dou, H., Zhang, X., Tang, W., Wu, W.: Infinitydrive: Breaking time limits in driving world models. arXiv preprint arXiv:2412.01522 (2024) [4](#)
28. Haji-Ali, M., Menapace, W., Siarohin, A., Skorokhodov, I., Canberk, A., Lee, K.S., Ordonez, V., Tulyakov, S.: Av-link: Temporally-aligned diffusion features for cross-modal audio-video generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19373–19385 (2025) [4](#)
29. He, X., Liu, Q., Qian, S., Wang, X., Hu, T., Cao, K., Yan, K., Zhang, J.: Id-animator: Zero-shot identity-preserving human video generation. arXiv preprint arXiv:2404.15275 (2024) [5](#)
30. He, Y., Pittaluga, F., Jiang, Z., Zwicker, M., Chandraker, M., Tasneem, Z.: Lang-drivectrl: Natural language controllable driving scene editing with multi-modal agents. arXiv preprint arXiv:2512.17445 (2025) [5](#)
31. Hu, A., Russell, L., Yeo, H., Murez, Z., Fedoseev, G., Kendall, A., Shotton, J., Corrado, G.: Gaia-1: A generative world model for autonomous driving. arXiv preprint arXiv:2309.17080 (2023) [4](#)
32. Hu, L.: Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8153–8163 (2024) [4](#)
33. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., et al.: Vbench: Comprehensive benchmark suite for video generative models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21807–21818 (2024) [11](#)
34. Ju, X., Ye, W., Liu, Q., Wang, Q., Wang, X., Wan, P., Zhang, D., Gai, K., Xu, Q.: Fulldit: Multi-task video generative foundation model with full attention. arXiv preprint arXiv:2503.19907 (2025) [4](#)
35. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139–1 (2023) [4](#)
36. Kong, W., Tian, Q., Zhang, Z., Min, R., Dai, Z., Zhou, J., Xiong, J., Li, X., Wu, B., Zhang, J., et al.: Hunyuanvideo: A systematic framework for large video generative models. arXiv preprint arXiv:2412.03603 (2024) [4](#)
37. Lee, S.H., Oh, G., Byeon, W., Kim, C., Ryoo, W.J., Yoon, S.H., Cho, H., Bae, J., Kim, J., Kim, S.: Sound-guided semantic video generation. In: European Conference on Computer Vision. pp. 34–50. Springer (2022) [4](#)
38. Li, J., Jiang, J., Miao, J., Long, M., Wen, T., Jia, P., Liu, S., Yu, C., Liu, M., Cai, Y., et al.: Realistic and controllable 3d gaussian-guided object editing for driving video generation. arXiv preprint arXiv:2508.20471 (2025) [5](#)
39. Li, R., Yi, B., Liu, J., Gao, H., Ma, Y., Kanazawa, A.: Cameras as relative positional encoding. arXiv preprint arXiv:2507.10496 (2025) [3](#), [8](#)
40. Li, R., Torr, P., Vedaldi, A., Jakab, T.: Vmem: Consistent interactive video scene generation with surfel-indexed view memory. arXiv preprint arXiv:2506.18903 (2025) [4](#)
41. Li, T., Slavcheva, M., Zollhoefer, M., Green, S., Lassner, C., Kim, C., Schmidt, T., Lovegrove, S., Goesele, M., Newcombe, R., et al.: Neural 3d video synthesis from multi-view video. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5521–5531 (2022) [4](#)

42. Li, Z., Li, K., Wang, S., Lan, S., Yu, Z., Ji, Y., Li, Z., Zhu, Z., Kautz, J., Wu, Z., et al.: Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation. arXiv preprint arXiv:2406.06978 (2024) [2](#)
43. Liang, Y., Yan, Z., Chen, L., Zhou, J., Yan, L., Zhong, S., Zou, X.: Driveeditor: A unified 3d information-guided framework for controllable object editing in driving scenes. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 5164–5172 (2025) [2](#), [3](#), [5](#), [7](#), [10](#), [11](#)
44. Liao, B., Chen, S., Yin, H., Jiang, B., Wang, C., Yan, S., Zhang, X., Li, X., Zhang, Y., Zhang, Q., et al.: Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12037–12047 (2025) [2](#)
45. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022) [5](#)
46. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022) [5](#)
47. Liu, Y., Ren, Y., Artola, A., Hu, Y., Cun, X., Zhao, X., Zhao, A., Chan, R.H., Zhang, S., Liu, R., et al.: Pusa v1. 0: Surpassing wan-i2v with 500 dollors training cost by vectorized timestep adaptation. arXiv preprint arXiv:2507.16116 (2025) [9](#)
48. Ljungbergh, W., Taveira, B., Zheng, W., Tonderski, A., Peng, C., Kahl, F., Petersson, C., Felsberg, M., Keutzer, K., Tomizuka, M., et al.: R3d2: Realistic 3d asset insertion via diffusion for autonomous driving simulation. arXiv preprint arXiv:2506.07826 (2025) [5](#), [7](#)
49. Lu, J., Huang, Z., Yang, Z., Zhang, J., Zhang, L.: Wovogen: World volume-aware diffusion for controllable multi-camera driving scene generation. In: European Conference on Computer Vision. pp. 329–345. Springer (2024) [4](#)
50. Ma, B., Gao, H., Deng, H., Luo, Z., Huang, T., Tang, L., Wang, X.: You see it, you got it: Learning 3d creation on pose-free videos at scale. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 2016–2029 (2025) [4](#)
51. Ma, G., Huang, H., Yan, K., Chen, L., Duan, N., Yin, S., Wan, C., Ming, R., Song, X., Chen, X., et al.: Step-video-t2v technical report: The practice, challenges, and future of video foundation model. arXiv preprint arXiv:2502.10248 (2025) [4](#)
52. Ma, Y., He, Y., Cun, X., Wang, X., Chen, S., Li, X., Chen, Q.: Follow your pose: Pose-guided text-to-video generation using pose-free videos. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 4117–4125 (2024) [4](#)
53. Ma, Z., Zhou, D., Wang, X.S., Yeh, C.H., Li, X., Yang, H., Dong, Z., Keutzer, K., Feng, J.: Magic-me: Identity-specific video customized diffusion. In: European Conference on Computer Vision. pp. 19–37. Springer (2024) [5](#)
54. Maes, L., Lidec, Q.L., Scieur, D., LeCun, Y., Balestriero, R.: Leworldmodel: Stable end-to-end joint-embedding predictive architecture from pixels. arXiv preprint arXiv:2603.19312 (2026) [4](#)
55. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021) [4](#)
56. Niu, M., Cun, X., Wang, X., Zhang, Y., Shan, Y., Zheng, Y.: Mofa-video: Control-lable image animation via generative motion field adaptations in frozen image-to-video diffusion model. In: European Conference on Computer Vision. pp. 111–128. Springer (2024) [4](#)

57. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4195–4205 (2023) [2](#), [3](#), [4](#), [8](#)
58. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10318–10327 (2021) [4](#)
59. Reid, M.e.a., DeepMind, G.: Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context (2024) [14](#)
60. Ren, A.Z., Lidard, J., Ankile, L.L., Simeonov, A., Agrawal, P., Majumdar, A., Burchfiel, B., Dai, H., Simchowitz, M.: Diffusion policy policy optimization. arXiv preprint arXiv:2409.00588 (2024) [2](#)
61. Ren, X., Shen, T., Huang, J., Ling, H., Lu, Y., Nimier-David, M., Müller, T., Keller, A., Fidler, S., Gao, J.: Gen3c: 3d-informed world-consistent video generation with precise camera control. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6121–6132 (2025) [4](#)
62. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [2](#), [4](#), [5](#)
63. RunwayML, G.: Introducing gen-3 alpha: a new frontier for video generation (2024) [4](#)
64. Russell, L., Hu, A., Bertoni, L., Fedoseev, G., Shotton, J., Arani, E., Corrado, G.: Gaia-2: A controllable multi-view generative world model for autonomous driving. arXiv preprint arXiv:2503.20523 (2025) [3](#), [4](#)
65. Shi, X., Huang, Z., Wang, F.Y., Bian, W., Li, D., Zhang, Y., Zhang, M., Cheung, K.C., See, S., Qin, H., et al.: Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024) [4](#)
66. Shi, X., Li, B., Han, X., Cai, Z., Yang, L., Lin, D., Wang, Q.: Consistcompose: Unified multimodal layout control for image composition. arXiv preprint arXiv:2511.18333 (2025) [4](#)
67. Singer, A., Rotstein, N., Mann, A., Kimmel, R., Litany, O.: Time-to-move: Training-free motion controlled video generation via dual-clock denoising. arXiv preprint arXiv:2511.08633 (2025) [3](#), [5](#), [10](#), [11](#)
68. Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., Liu, Y.: Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **568**, 127063 (2024) [8](#)
69. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020) [2](#)
70. Tao, P., Cheng, C., Du, Y., Song, C., Chen, Z., Guo, X., Yin, W., Ren, W., Zhang, Q., Cui, H., et al.: Anchor3r: Streaming 3d reconstruction with transient anchors for long-horizon visual mapping. arXiv preprint arXiv:2606.05035 (2026) [4](#)
71. Team, M.L., Cai, X., Huang, Q., Kang, Z., Li, H., Liang, S., Ma, L., Ren, S., Wei, X., Xie, R., et al.: Longcat-video technical report. arXiv preprint arXiv:2510.22200 (2025) [4](#)
72. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025) [4](#), [10](#)

73. Wang, H., Ouyang, H., Wang, Q., Yu, Y., Meng, Y., Wang, W., Cheng, K.L., Ma, S., Bai, Q., Li, Y., et al.: The world is your canvas: Painting promptable events with reference images, trajectories, and text. *arXiv preprint arXiv:2512.16924* (2025) [4](#)
74. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 5294–5306 (2025) [12](#)
75. Wang, Q., Luo, Y., Shi, X., Jia, X., Lu, H., Xue, T., Wang, X., Wan, P., Zhang, D., Gai, K.: Cinemaster: A 3d-aware and controllable framework for cinematic text-to-video generation. In: *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*. pp. 1–10 (2025) [4](#)
76. Wang, S., Sun, H., Wang, B., Ye, H., Yu, X.: Mirage: One-step video diffusion for photorealistic and coherent asset editing in driving scenes. *arXiv preprint arXiv:2512.24227* (2025) [5](#), [7](#)
77. Wang, X., Zhu, Z., Huang, G., Chen, X., Zhu, J., Lu, J.: Drivedreamer: Towards real-world-drive world models for autonomous driving. In: *European conference on computer vision*. pp. 55–72. Springer (2024) [4](#)
78. Wang, X., Fu, S., Huang, Q., He, W., Jiang, H.: Ms-diffusion: Multi-subject zero-shot image personalization with layout guidance. *arXiv preprint arXiv:2406.07209* (2024) [4](#)
79. Wang, Y., He, J., Fan, L., Li, H., Chen, Y., Zhang, Z.: 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*. pp. 14749–14759 (2024) [3](#), [4](#)
80. Wen, Y., Zhao, Y., Liu, Y., Jia, F., Wang, Y., Luo, C., Zhang, C., Wang, T., Sun, X., Zhang, X.: Panacea: Panoramic and controllable video generation for autonomous driving. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6902–6912 (2024) [3](#), [4](#)
81. Xie, T., Yang, P., Jin, Y., Cai, Y., Yin, W., Ren, W., Zhang, Q., Hua, W., Peng, S., Guo, X., et al.: Scal3r: Scalable test-time training for large-scale 3d reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 21760–21771 (2026) [4](#)
82. Xing, J., Liu, H., Xia, M., Zhang, Y., Wang, X., Shan, Y., Wong, T.T.: Toon-crafter: Generative cartoon interpolation. *ACM Transactions on Graphics (TOG)* **43**(6), 1–11 (2024) [4](#)
83. Xing, J., Xia, M., Liu, Y., Zhang, Y., Zhang, Y., He, Y., Liu, H., Chen, H., Cun, X., Wang, X., et al.: Make-your-video: Customized video generation using textual and structural guidance. *IEEE Transactions on Visualization and Computer Graphics* **31**(2), 1526–1541 (2024) [4](#)
84. Xu, R., Zhou, D., Ma, F., Yang, Y.: Contextgen: Contextual layout anchoring for identity-consistent multi-instance generation. *arXiv preprint arXiv:2510.11000* (2025) [4](#)
85. Yang, K., Ma, E., Peng, J., Guo, Q., Lin, D., Yu, K.: Bevcontrol: Accurately controlling street-view elements with multi-perspective consistency via bev sketch layout. *arXiv preprint arXiv:2308.01661* (2023) [4](#)
86. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072* (2024) [4](#)

87. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20331–20341 (2024) 4
88. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. arXiv preprint arXiv:2308.06721 (2023) 4
89. Yu, H.X., Duan, H., Hur, J., Sargent, K., Rubinstein, M., Freeman, W.T., Cole, F., Sun, D., Snavely, N., Wu, J., et al.: Wonderjourney: Going from anywhere to everywhere. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6658–6667 (2024) 4
90. Yu, W., Xing, J., Yuan, L., Hu, W., Li, X., Huang, Z., Gao, X., Wong, T.T., Shan, Y., Tian, Y.: Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. arXiv preprint arXiv:2409.02048 (2024) 4
91. Yuan, S., Huang, J., He, X., Ge, Y., Shi, Y., Chen, L., Luo, J., Yuan, L.: Identity-preserving text-to-video generation by frequency decomposition. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12978–12988 (2025) 5
92. Zeng, K., Wu, Z., Xiong, K., Wei, X., Guo, X., Zhu, Z., Ho, K., Zhou, L., Zeng, B., Lu, M., et al.: Rethinking driving world model as synthetic data generator for perception tasks. arXiv preprint arXiv:2510.19195 (2025) 5
93. Zhan, Y., Li, Z., Niu, M., Zhong, Z., Nobuhara, S., Nishino, K., Zheng, Y.: Kfd-nerf: Rethinking dynamic nerf with kalman filter. In: European Conference on Computer Vision. pp. 1–18. Springer (2024) 4
94. Zhang, C., Zhan, Y., Wang, Q., Ouyang, Z., Li, Y., Yang, Z., Guo, X., Ren, W., Zhang, Q., Dong, Z., et al.: Horizontdrive: Self-corrective autoregressive world model for long-horizon driving simulation. arXiv preprint arXiv:2605.11596 (2026) 4
95. Zhang, K., Tang, Z., Hu, X., Pan, X., Guo, X., Liu, Y., Huang, J., Yuan, L., Zhang, Q., Long, X.X., et al.: Epona: Autoregressive diffusion world model for autonomous driving. arXiv preprint arXiv:2506.24113 (2025) 4, 14
96. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3836–3847 (2023) 4
97. Zhao, G., Wang, X., Zhu, Z., Chen, X., Huang, G., Bao, X., Wang, X.: Drivedreamer-2: Llm-enhanced world models for diverse driving video generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 10412–10420 (2025) 4
98. Zheng, Z., Peng, X., Yang, T., Shen, C., Li, S., Liu, H., Zhou, Y., Li, T., You, Y.: Open-sora: Democratizing efficient video production for all, 2024. URL <https://github.com/hpcaitech/Open-Sora> (2024) 5
99. Zhou, D., Xie, J., Yang, Z., Yang, Y.: 3dis: Depth-driven decoupled instance synthesis for text-to-image generation. arXiv preprint arXiv:2410.12669 (2024) 4
100. Zhu, Q., Wei, Z., Zheng, Z., Zhan, Y., Yao, Z., Zhang, J., Wu, K., Zheng, Y.: Rpbg: Towards robust neural point-based graphics in the wild. In: European Conference on Computer Vision. pp. 389–406. Springer (2024) 4
101. Zhu, R., Guo, J., Guo, X., Liu, Z., Hui, K.H., Yin, W., Chen, K., Chen, W., Ren, W., Liu, Y., et al.: Eps3d: End-to-end feed-forward 3d panoptic segmentation. arXiv preprint arXiv:2606.08980 (2026) 4