



OmniNWM: Omniscient Driving Navigation World Models

Bohan Li^{1,2*}, Zhuang Ma^{3*†}, Dalong Du^{3*}, Baorui Peng², Zhujin Liang³, Zhenqiang Liu³,
Chao Ma¹, Yueming Jin⁴, Hao Zhao⁵, Wenjun Zeng², Xin Jin^{2‡}

¹Shanghai Jiao Tong University, ²Eastern Institute of Technology, Ningbo

³PhiGent, ⁴National University of Singapore, ⁵Tsinghua University

<https://github.com/Arlo0o/OmniNWM>

Abstract

Autonomous driving world models are expected to work effectively across three core dimensions: state, action, and reward. Existing models, however, are typically restricted to limited state modalities, short video sequences, imprecise action control, and a lack of reward awareness. In this paper, we introduce OmniNWM, an omniscient panoramic navigation world model that addresses all three dimensions within a unified framework. For **state**, OmniNWM jointly generates panoramic videos of RGB, semantics, metric depth, and 3D occupancy. A flexible forcing strategy enables high-quality long-horizon auto-regressive generation. For **action**, we introduce a normalized panoramic Plücker ray-map representation that encodes input trajectories into pixel-level signals, enabling highly precise and generalizable control over panoramic video generation. Regarding **reward**, we move beyond learning reward functions with external image-based models: instead, we leverage the generated 3D occupancy to directly define rule-based dense rewards for driving compliance and safety. Extensive experiments demonstrate that OmniNWM achieves state-of-the-art performance in video generation, control accuracy, and long-horizon stability, while providing a reliable closed-loop evaluation framework through intrinsic occupancy-grounded rewards.

1. Introduction

Recent advancements in world models have demonstrated significant potential for autonomous driving, facilitating high-fidelity simulation of complex environments and controllable navigation for autonomous agents [5, 23, 45, 78, 79, 105]. These models typically predict plausible future environmental states conditioned on historical observations

*Equal contribution

†Project leader

‡Corresponding

Method	State		Action	Reward
	Modalities	Length		
GenAD [112]	R	6	Waypoint	✗
DrivingGPT [10]	R	60	Waypoint	✗
Drivedreamer [78]	R	32	Velocity, Angle	✗
Drivedreamer-2 [108]	R	26	Velocity, Angle	✗
Panacea [84]	R	8	Waypoint	✗
MagicDrive [16]	R	60	Waypoint	✗
Dist-4D [21]	R, D	17	Waypoint	✗
WoVoGen [56]	R, O	6	Volumetric-based	✗
OccScene [46]	R, O	8	Volumetric-based	✗
UniScene [45]	R, L, O	24	Volumetric-based	✗
Drive-WM [79]	R	8	Waypoint	External Image-based
OmniNWM (Ours)	R, S, D, O	321	Normalized Plücker Ray-map-based	Intrinsic Occupancy-based

Table 1. **Comparison of world model capabilities.** We compare OmniNWM with existing approaches across state, action, and reward dimensions. ‘R’, ‘S’, ‘D’, ‘L’, ‘O’ denote RGB, semantic, depth, LiDAR, and occupancy modalities, respectively.

and actions, then evaluate action feasibility via reward functions [17, 18, 45, 58, 79].

Nevertheless, building robust and generalizable driving world models remains challenging due to the complexity of real-world 3D environments [22, 45, 57, 60, 87, 96, 97, 117]. Specifically, as shown in Table 1, existing methods face three core limitations: **1) State:** Driving world models must integrate diverse data formats over sufficient temporal frames for closed-loop evaluation and downstream tasks [16, 21, 45]. Current approaches predominantly rely on single-modality RGB videos with limited length, which fail to capture the geometric and semantic complexity of large-scale 3D scenes [16–18, 58, 79]. **2) Action:** Precise control for generating panoramic videos remains under-explored. Existing approaches often employ sparse action encodings (e.g., trajectory waypoints) or heavy volumetric conditions that hinder long-term consistency and precise viewpoint manipulation [16, 78, 79, 108]. **3) Reward:** Effective reward functions are essential for evaluating navigation performance, yet few methods have integrated them into a unified




Figure 1. **Versatile capabilities of OmniNWM.** (a) State: Given a conditional image and ego trajectory, OmniNWM jointly generates comprehensive multi-modal outputs, including panoramic RGB, semantic, metric depth, and 3D occupancy videos. (b) Action: Conditioned on different input trajectories, OmniNWM facilitates precise panoramic camera control by converting them into normalized Plücker ray-maps as pixel-level guidance. (c) Reward: OmniNWM enables unbounded long-term navigation through a closed-loop pipeline: the planning trajectory guides the multi-modal generation, while dense rewards are derived from the generated 3D semantic occupancy. Note that all input trajectories are encoded into normalized Plücker ray-maps before guiding the generation process, even when not explicitly visualized.

framework [17, 45, 78, 84, 108]. Although some recent studies propose image-based reward functions, they often rely on external models and fail to provide precise evaluations in driving environments [79]. These limitations collectively limit the robustness and generalization of current methods.

To address these challenges, we propose OmniNWM, an omniscient navigation world model for autonomous driving. OmniNWM unifies multi-modal forecasting over unbounded long-term sequences, precise and generalizable panoramic camera control, and seamless integration with 3D occupancy-grounded dense driving rewards within a single framework. As illustrated in Figure 1 (a), OmniNWM jointly generates panoramic videos of RGB, semantic, metric depth, and 3D semantic occupancy in an autonomous driving scenario. A flexible forcing strategy injects independent multi-level noise during training, enabling robust long-term generation beyond ground-truth sequence lengths (*i.e.*, 321 vs. 241

frames in Figure 1 (c)) and supporting frame- or clip-level auto-regressive inference. Figure 1 (b) illustrates the precise action control for panoramic video generation. In contrast to the sparse representation of trajectory waypoints used in prior work [18, 78, 108], we encode action sequences into normalized panoramic Plücker ray-maps—a unified pixel-level representation that facilitates accurate and generalizable control. Moreover, this normalized panoramic Plücker encoder is lightweight and introduces minimal overhead. Finally, as illustrated in Figure 1 (c), OmniNWM supports closed-loop navigation by leveraging its generated 3D occupancy to provide dense waypoint-wise rewards, while planned trajectories iteratively guide the generation of future multi-modal predictions. Our contributions are summarized as follows:

- We propose OmniNWM, an omniscient navigation world model that supports comprehensive forecasting, precise

action control, and occupancy-grounded dense rewards within a unified framework.

- A normalized panoramic ray-map representation is introduced to enable precise and generalizable camera control.
- Driven by a flexible forcing strategy, OmniNWM achieves long-term auto-regressive generation beyond GT length.
- Leveraging jointly generated 3D occupancy, OmniNWM naturally supports closed-loop navigation and evaluation through intrinsic, occupancy-grounded dense rewards.

Experiments show that OmniNWM achieves state-of-the-art performance in generation quality and control precision. Furthermore, the model exhibits remarkable generalization across diverse and challenging settings, including panoramic generation conditioned on novel trajectories, unbounded generation beyond GT sequence lengths, and zero-shot generation across different datasets and camera configurations.

2. Related Work

World Models for Autonomous Driving. Recent progress in driving world models has advanced visual realism and controllability [8, 16, 18, 19, 34, 36, 42, 44, 45, 49, 58, 64, 66, 76, 79, 90, 98–100, 109, 116]. Early efforts like DriveDreamer [78] leverage multi-stage pipelines to incorporate traffic rules. MagicDrive [17] introduces cross-view attention for street-view synthesis, while Vista [18] improves generation via latent replacement strategies. UniScene [45] proposes a hierarchical occupancy-centric framework. Other works include ReSim [93] for simulation data augmentation, HUGSIM [116] for 3D Gaussian Splatting-based rendering, Drive-WM [79] for multi-view forecasting, and GEM [23] for egocentric prediction. However, existing models typically lack unified multi-modal generation (*e.g.*, semantic map, 3D occupancy) [17, 18, 78], and integrated reward modeling for closed-loop evaluation [16, 45, 108]. Our work addresses these limitations with long-horizon, multi-modal generation, precise control, and intrinsic reward modeling.

Camera-controlled Video Generation. Camera control has become a critical technique for guiding generative models in dynamic 3D environments [4, 24, 25, 29, 39, 48, 81, 95]. CameraCtrl [24] introduces a plug-and-play module to parameterize camera trajectories. MotionCtrl [81] decouples camera and object motion for independent control, while CamCo [89] introduces epipolar constraints to improve geometric consistency. VD3D [1] embeds Plücker coordinates into transformer architectures for view synthesis, and ReCamMaster [4] focuses on re-rendering existing videos from novel camera paths. However, these methods focus on short-range, monocular sequences, limiting their applicability to large-scale multi-view environments in autonomous driving [1, 4, 24, 25, 81, 89]. OmniNWM extends camera controllability to panoramic, long-horizon settings through a normalized Plücker ray-map representation.

3. Methodology

At its core, OmniNWM unifies the *state–action–reward* triad within a single framework, enabling comprehensive state generation, precise action control, and intrinsic dense rewards. As illustrated in Figure 2, OmniNWM supports multi-modal generation (Section 3.1), producing pixel-aligned panoramic videos (Section 3.1.1), as well as 3D semantic occupancy for intrinsic dense rewards (Section 3.1.2) and future planning trajectories (Section 3.1.3). Section 3.2 introduces the normalized panoramic Plücker ray-map camera control scheme, while Section 3.3 describes the flexible forcing strategy for long-term robust generation.

3.1. Comprehensive Generation within OmniNWM

The overview of the model architecture is illustrated in Figure 2. We leverage a pretrained 3D VAE [38] to encode input video frames into spatiotemporal latents. A panoramic diffusion transformer (PDiT) with cross-view attention layers [46, 85] is leveraged to denoise the resulting latents. The denoised latents are decoded into pixel-aligned RGB, semantic, and depth outputs, which are subsequently used to produce 3D semantic occupancy and future planning trajectories.

3.1.1. Pixel-level Aligned Multi-modal Generation

Joint generation of pixel-aligned multi-modal outputs has shown great potential to enhance the representational capacity of generative models and downstream task performance [45, 54]. In OmniNWM, we initialize modality-specific noise latents for RGB, semantics, and depth, then concatenate them channel-wise for joint denoising using the PDiT. A shared decoder produces spatially aligned RGB, semantic, and depth video outputs. To ensure cross-modal consistency during training, semantic maps are colorized before VAE encoding and converted back to discrete labels via nearest-neighbor matching after decoding. The details of data curation on semantic and depth labels are provided in the supplementary material. This joint generation scheme ensures strong pixel-level alignment across modalities, which is essential for generating high-quality 3D occupancy, as detailed in the following subsection.

3.1.2. 3D Occupancy Generation for Dense Rewards

As shown in Figure 3, the occupancy module generates a 3D voxel volume $\hat{\mathbf{V}}$ from panoramic RGB, depth, and semantic inputs. Following previous work [46], a UNet based on a pre-trained EfficientNet-B7 [70] extracts image features F_i . Depth and semantic maps are processed via down-sampling and convolutional layers to produce depth features F_d and semantic features F_s , respectively. SE3D blocks [30, 44] then perform adaptive aggregation between F_i and F_d/F_s . Finally, the 3D semantic voxel volume is computed by:

$$\hat{\mathbf{V}} = \text{Adap}_d(F_i, F_d) \otimes \text{Adap}_s(F_i, F_s), \quad (1)$$




Figure 2. **Architecture overview of OmniNWM.** Our model takes as input a historical trajectory, a reference image, and flexible multi-level noise applied view-wise and frame-wise. It simultaneously forecasts panoramic RGB, semantic, and metric depth videos (Section 3.1.1), along with 3D occupancy for intrinsic dense rewards (Section 3.1.2) and future planning trajectories (Section 3.1.3). Note that input trajectories guide the generation of the current video clip, while future trajectories are forecasted for the next clip, forming a closed-loop system. The input trajectories with camera poses are encoded into normalized panoramic Plücker ray-maps, providing a unified pixel-level description and enabling precise action control (Section 3.2). A flexible forcing strategy injects independent multi-level noise at the view and frame levels to support robust and flexible auto-regressive generation (Section 3.3).




Figure 3. **Occupancy generation.** The generated panoramic RGB images, semantics, and depth are used to produce a 3D voxel volume. This design enables lightweight occupancy prediction compared with prior volumetric-based methods [45, 111], while naturally supporting the integration of occupancy-grounded rewards.

where $\otimes$ denotes the outer product. Adap_d and Adap_s are adaptive aggregation with SE3D for depth and semantic contexts. The resulting volume feature $\hat{V}$ is passed to a generation head with up-sampling and a softmax layer to produce the final 3D semantic occupancy. Unlike volumetric-based models [45, 111] that directly synthesize full occupancy grids at high computational cost, our design computes occupancy from lightweight pixel-aligned features, which scales to long sequences and, crucially, enables intrinsic occupancy-grounded dense driving rewards.

Occupancy-grounded Rewards. The generated occupancy enables waypoint-wise reward computation following the class taxonomy in the NuScenes-Occupancy [77] dataset. The occupancy-grounded reward function $\hat{R}$ evaluates the safety, efficiency, and compliance of driving behaviors with detailed 3D environmental understanding as follows:

$$\hat{R} = 1 + (R_{\text{col}} + R_{\text{bd}} + R_{\text{vel}})/N_{\text{reward}}. \quad (2)$$

(1) The Collision Reward R_{col} penalizes collisions between ego vehicle and obstacles, scaled by speed v to reflect risks:

$$R_{\text{col}} = -\alpha_{\text{col}} \cdot \mathbb{I}_{\text{col}} \cdot |v|, \quad (3)$$

where $\mathbb{I}_{\text{col}}$ is a binary indicator function that equals 1 if a collision occurs and 0 otherwise. This is derived by

querying the 3D occupancy for obstacle classes (e.g., ‘car’, ‘pedestrian’, ‘barrier’), and α_{col} is the balancing coefficient. (2) The Boundary Reward R_{bd} penalizes off-road behavior by detecting whether the ego vehicle leaves the drivable region (identified by the ‘drivable surface’ class):

$$R_{\text{bd}} = -\alpha_{\text{bd}} \cdot \mathbb{I}_{\text{non-drivable}}, \quad (4)$$

where $\mathbb{I}_{\text{non-drivable}}$ is a binary indicator function that equals 1 when the vehicle is in a non-drivable area and 0 otherwise.

(3) The Velocity Reward R_{vel} encourages maintaining a target speed of the ego vehicle to promote traffic efficiency:

$$R_{\text{vel}} = -\alpha_{\text{vel}} \cdot \tanh(|v - v_{\text{target}}|) \cdot \mathbb{I}_v, \quad (5)$$

where $v > 0$, and $\mathbb{I}_v$ is a binary indicator function that equals 0 when $v \in [v_{\text{min}}, v_{\text{max}}]$ and 1 otherwise.

By transforming 3D occupancy into dense waypoint-level feedback, the proposed reward function enables fine-grained evaluation of policy behavior under diverse scenarios. In a critical oncoming-truck scenario (Figure 4), the reward function effectively distinguishes driving quality: a high-speed collision (left) yields the lowest reward with severe R_{col} penalty; an insufficient avoidance (middle) reduces reward; and a successful evasion (right) achieves the highest reward with minimal control penalty R_{col} and velocity penalty R_{vel} , confirming discriminative capability. More evaluation results can be found in the supplementary material.

3.1.3. Trajectory Planning for Closed-loop Navigation

As illustrated in Figure 1 (c), OmniNWM functions as an environment simulator to enable closed-loop navigation with planning agents. The planning trajectories generated from the model’s own outputs guide future state generation when encoded into Plücker ray-maps.

To fully leverage OmniNWM’s multi-modal outputs, we developed a semantic-geometric reasoning vision-language-action (VLA) model based on Qwen-2.5-VL [3], termed




Figure 4. **Average rewards of different trajectories**, computed using the proposed 3D occupancy-grounded reward function. The rewards effectively evaluate the feasibility of planning trajectories in the presence of obstacles (e.g., an oncoming truck).

OmniNWM-VLA, as the default planning agent. We introduce a plug-and-play Tri-Modal Mamba-based Interpreter (Tri-MIDI) to integrate semantic, depth, and RGB contexts before tokenization and input into OmniNWM-VLA. To adapt planning trajectories suitable as Plücker ray-map inputs, we extend the output channels of OmniNWM-VLA to predict not only the x and y coordinates of waypoints, but also their heading angles. In contrast to prior methods [11, 103], which produce key frames at 2HZ, our method operates at 12HZ on the NuScenes dataset. Note that for the initial video clip, the trajectory is produced by first generating multi-modal results (*i.e.*, RGB, depth, and semantics) using a static trajectory, which are then fed into OmniNWM-VLA to enable multi-modal reasoning for the planner.

Moreover, OmniNWM supports straightforward integration of various planning agents for closed-loop evaluation. Additional experimental comparisons are provided in Section 4.2, while further architectural and implementation details of OmniNWM-VLA are available in the supplementary.

3.2. Control with Panoramic Normalized Ray-map

Existing driving world models face two key limitations in action control: (1) sparse vectorized representations (e.g., waypoints) hinder accurate multi-view depiction [18, 78, 108], and (2) limited trajectory diversity in datasets restricts model generalization (Figure 5 (c)). To overcome these issues, we propose converting trajectories into panoramic normalized Plücker ray-maps, providing a unified pixel-level representation that enables precise camera control (Figure 1 (b)) and supports zero-shot generalization across datasets and camera configurations (Figure 10).




Figure 5. **Panoramic normalized ray-map encoding**. (a) Plücker ray-maps derived from panoramic camera poses are normalized in both scale and pose. (b) The ray-map normalization process constructs different trajectories within a unified 3D Plücker space. (c) Compared to the original NuScenes dataset, our strategy significantly enhances the diversity of trajectory distributions.

3.2.1. Efficient Plücker Ray-map Encoding

As illustrated in Figure 2, input trajectories are encoded into ray-maps via a parameter-free normalized Plücker encoder and injected into PDiT blocks. Specifically, the normalized panoramic ray-maps are first downsampled in spatial and temporal dimensions to align with diffusion latents, and patchified into Plücker embedding tokens. These tokens are concatenated with the diffusion latent tokens (obtained via patchification of the diffusion latents) before being fed into the 3D full-attention layers within the PDiT blocks.

Following [24], Plücker ray-maps are constructed from camera intrinsics K and extrinsics $E = [R; t]$, where R is the rotation matrix and t the translation vector. For a pixel (u, v) , its Plücker embedding is defined as:

$$p_{u,v} = (o \times d_{u,v}, d_{u,v}) \in \mathbb{R}^6, \quad (6)$$

where $o = t$ is the camera center in world coordinates, and $d_{u,v}$ is the direction vector from o to the pixel, computed as:

$$d_{u,v} = RK^{-1}[u, v, 1]^T. \quad (7)$$

The embedding of i^{th} frame and j^{th} view is represented as $P_{i,j} \in \mathbb{R}^{6 \times h \times w}$, while a panoramic sequence of n frames and v views yields $P \in \mathbb{R}^{n \times 6 \times v \times h \times w}$.

3.2.2. Scale and Pose Invariant Normalization

As shown in Figure 5 (a), we apply scale and pose normalization to each frame $P_j \in \mathbb{R}^{6 \times v \times h \times w}$ of the panoramic ray-map to obtain a unified Plücker representation. Our normalization method decouples camera parameters: a shared

intrinsic matrix from a reference camera (defaulting to the front view) ensures scale consistency, while individual extrinsics retain distinct viewpoints. These parameters are transformed into a unified coordinate system. For each pixel (u, v) in the k^{th} camera, the direction vector is computed using the reference intrinsics K_0 and the camera’s rotation matrix R_k :

$$\hat{d}_{u,v}^{(k)} = R_k K_0^{-1} [u, v, 1]^T. \quad (8)$$

This projects all rays into a consistent metric scale. For pose invariance, both the camera center and direction vector are transformed into the reference camera’s coordinate system. The camera center becomes:

$$o^{(k \rightarrow 0)} = R_0 R_k^{-1} \hat{t}^{(k \rightarrow 0)}, \quad (9)$$

and the direction vector is rotated accordingly:

$$\hat{d}_{u,v}^{(k \rightarrow 0)} = R_0 R_k^{-1} \hat{d}_{u,v}^{(k)}. \quad (10)$$

The normalized Plücker coordinate is then given by:

$$\hat{p}_{u,v}^{(k)} = \left(o^{(k \rightarrow 0)} \times \hat{d}_{u,v}^{(k \rightarrow 0)}, \hat{d}_{u,v}^{(k \rightarrow 0)} \right). \quad (11)$$

This yields *scale- and pose-normalized* ray-maps $\hat{P}_j$ that are geometrically consistent across all views. This process constructs trajectories from different camera views within a unified 3D Plücker space (Figure 5 (b)), significantly enriching the diversity of trajectory distributions (Figure 5 (c)).

3.3. Long-term Generation with Flexible Forcing

To improve generation flexibility and robustness in long-term forecasting, we propose a flexible forcing strategy with multi-level noise injection. This approach corrupts preceding frames along multiple independent dimensions, enabling robust denoising under partially corrupted history. This is critical for mitigating error accumulation in extended autoregressive generation [14, 46].

Formally, during training, we independently inject noise into each frame i and view j . Let $\mathbf{x}^{(i,j)}$ denote the latent representation for frame i and view j . We corrupt it as:

$$\tilde{\mathbf{x}}^{(i,j)} = \mathbf{x}^{(i,j)} + \alpha^{(i)} \cdot \epsilon_{\text{frame}} + \beta^{(j)} \cdot \epsilon_{\text{view}}, \quad (12)$$

where $\epsilon_{\text{frame}} \sim \mathcal{N}(0, I)$ and $\epsilon_{\text{view}} \sim \mathcal{N}(0, I)$ are independent noise vectors, and $\alpha^{(i)}, \beta^{(j)}$ are scaling factors. This per-view/frame noise scheme improves handling of spatio-temporal distortions compared to previous uniform noise application [18, 45].

As shown in Figure 6 (b), this formulation enables two distinct autoregressive inference modes:

(1) Frame-level autoregression generates a single future frame conditioned on multiple preceding frames:

$$\hat{\mathbf{x}}^{(t+1)} = f_{\theta} \left(\tilde{\mathbf{x}}^{(t-K)}, \tilde{\mathbf{x}}^{(t-K+1)}, \dots, \tilde{\mathbf{x}}^{(t)} \right), \quad (13)$$




Figure 6. **Flexible forcing strategy.** (a) During training, independent multi-level noise is injected along view-wise and frame-wise dimensions. (b) During inference, flexible and robust generation is enabled via either frame-level auto-regression (many-to-one) or clip-level auto-regression (one-to-many or many-to-many).

where f_{θ} is the diffusion model. K is the input context window. $\tilde{\mathbf{x}}^{(t)}$ is the potentially corrupted latent at time t . This mode is particularly suitable for frame-wise trajectory planning simulation (Figure 4).

(2) Clip-level auto-regression generates multiple future frames from fewer inputs:

$$\hat{\mathbf{x}}^{(t+1)}, \dots, \hat{\mathbf{x}}^{(t+L)} = f_{\theta} \left(\tilde{\mathbf{x}}^{(t-M)}, \dots, \tilde{\mathbf{x}}^{(t)} \right), \quad (14)$$

where L is the output horizon and M is the limited input context window. This approach is efficient for long-horizon generation while maintaining temporal coherence. During inference, historical frames can be partially noised to support both generation modes. For instance, in clip-level auto-regression, a single historical frame can be used as conditioning to generate multiple future frames initially, while multiple historical frames can be utilized in later stages to better exploit dynamic priors. Ablation studies in Figure 9 further validate the effectiveness of our approach.

4. Experiments

This section presents evaluation results of our approach against existing approaches across multiple applications.

4.1. Experimental Setup

We train our model on the NuScenes [6] and NuScenes-Occupancy [77] datasets. Our Panoramic Diffusion Transformer (PDiT) comprises 11.22B parameters, including an 11B Diffusion Transformer backbone consistent with prior works [40, 115] and 0.22B new parameters for cross-view attention layers. Training is conducted on 48 NVIDIA A800 GPUs with a batch size of 48. We employ the AdamW optimizer [55] with a learning rate of 1e-4. Further details regarding training objectives, data processing, experimental implementations of the occupancy generation module, and the OmniNWM-VLA are provided in the supplementary.

Method	Volumetric Condition-Free	Multi-view	Video	FID ↓	FVD ↓
DriveGAN [37]	✓	✗	✓	73.40	502.30
BEVControl [94]	✓	✓	✗	24.85	-
DriveDreamer [78]	✓	✗	✓	14.90	340.80
BEVGen [69]	✓	✓	✗	25.54	-
DrivingGPT [10]	✓	✗	✓	12.78	142.61
Vista [18]	✓	✗	✓	6.90	89.40
DrivingWorld [31]	✓	✗	✓	7.40	90.90
Epona [105]	✓	✗	✓	7.50	82.80
WoVoGen [56]	✗	✓	✓	27.60	417.70
X-Scene [97]	✗	✓	✗	11.29	-
OccScene [46]	✗	✓	✓	11.87	-
UniScene [45]	✗	✓	✓	6.45	71.94
DiST-4D [21]	✗	✓	✓	7.40	25.55
MagicDrive [17]	✓	✓	✓	16.20	-
Panacea [84]	✓	✓	✓	16.96	139.00
Drive-WM [79]	✓	✓	✓	15.80	122.70
GenAD [112]	✓	✓	✓	15.40	184.00
DriveDreamer-2 [108]	✓	✓	✓	11.20	55.70
MagicDrive-V2 [16]	✓	✓	✓	20.91	94.84
OmniNWM	✓	✓	✓	5.45	23.63

Table 2. **Quantitative evaluation of RGB video generation** on the NuScenes validation set. Our proposed method outperforms previous SOTA methods without heavy volumetric-based conditions.

Category	Method	Abs. Rel. ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$
Discriminative	SurroundDepth [82]	0.28	0.66	0.84
	M ² Depth [119]	0.26	0.73	0.87
Generative	Dist-4D [21]	0.39	0.58	0.81
	OmniNWM	0.23	0.81	0.93

Table 3. **Quantitative evaluation of generated panoramic depth** on the NuScenes validation set.

4.2. Main Results

Video Generation Quality. We evaluate RGB video generation quality using Frechet Inception Distance (FID) [26] and Frechet Video Distance (FVD) [73]. As shown in Table 2, our method achieves state-of-the-art performance of 5.45 FID and 23.63 FVD without requiring the heavy volumetric-based conditions (e.g., semantic occupancy [45, 46, 56], aggregated point clouds [21]), surpassing all previous approaches. These results demonstrate the effectiveness of our method for high-fidelity driving video generation.

We further evaluate the generated panoramic depth maps using standard metrics [53, 119]: absolute relative error (Abs. Rel.), and accuracy thresholds (δ), with LiDAR-projected ground truth. As shown in Table 3, our method achieves superior performance on Abs. Rel. and δ metrics. Our method significantly outperforms the generative method of Dist-4D [21], and also surpasses the discriminative methods [82, 119], whose generalization ability is limited.

Occupancy Prediction Quality. We evaluate occupancy prediction performance using IoU and mIoU metrics as in previous works [45, 46]. AICNet* [47] and 3DSketch* [9] utilize both camera-based RGB images and LiDAR-projected depth maps as inputs, following [77]. As shown in Table 4, our method significantly outperforms all comparison approaches, including LiDAR-based methods [63, 77, 92] that benefit

Category	Method	Input	IoU ↑	mIoU ↑
Discriminative	AICNet* [47]	Camera&Depth	23.8	10.6
	3DSketch* [9]	Camera&Depth	25.6	10.7
	LMSCNet [63]	LiDAR	27.3	11.5
	JS3C-Net [92]	LiDAR	30.2	12.5
	L-CONet [77]	LiDAR	30.9	15.8
	MonoScene [7]	Camera	18.4	6.9
	TPVFormer [32]	Camera	15.3	7.8
	C-CONet [77]	Camera	20.1	12.8
	HTCL [42]	Camera	21.4	14.1
	SparseOcc [71]	Camera	21.8	14.1
Generative	OccScene [46]	Camera	-	12.2
	OccGen [74]	Camera	23.4	14.5
	OmniNWM	Camera	33.3	19.8

Table 4. **Quantitative results of generated semantic occupancy** on the NuScenes-Occupancy validation set. Our method surpasses previous SOTA methods with a generative paradigm.




Figure 7. (a) **Closed-loop evaluation of trajectory planning**, reporting pass/fail counts, Scenario Pass Rate (SPR), and average rewards. (b) **Distribution of average rewards** across different methods on the NuScenes validation set.

from precise 3D geometric measurements, despite using only RGB images as inputs. These results demonstrate the effectiveness of our joint generation paradigm.

Trajectory Planning Evaluation. Figure 7 (a) presents closed-loop evaluation results for trajectory planning, showing Scenario Pass Rate (SPR) [67] and average rewards calculated across 150 scenes from the NuScenes validation set. Impromptu-VLA [11] achieves higher performance compared to Qwen-2.5-VL [3], with results comparable to our OmniNWM-VLA. Figure 7 (b) illustrates the distribution of average rewards that are evaluated using our occupancy-grounded driving rewards for different VLA planning baselines. The horizontal axis represents the percentage of the number of evaluation scenes. The results confirm that our reward function effectively distinguishes between the planning trajectories generated by these VLA planners.

Camera Control Accuracy. We evaluate camera control accuracy on the NuScenes validation set using Rotation Error (RotErr) and Translation Error (TransErr) metrics following previous works [24, 25]. Camera pose sequences are extracted using VGGT [75], with similarity scores calculated via Similarity Transformation (Sim(3)) [28]. As shown in Table 5 and Figure 8, our method with the panoramic normalized ray-map encoding significantly outperforms UniScene [45], yielding lower error distributions across

Method	Norm	RotErr (10^{-2}) ↓	TransErr (10^{-2}) ↓
GT Video	-	1.36	1.75
UniScene [45]	-	1.98	12.74
OmniNWM	✗	2.11	14.19
OmniNWM	✓	1.42	5.14

Table 5. **Quantitative results of camera control accuracy** on the NuScenes validation set. VGGT [75] is leveraged for camera pose extraction, and Similarity Transformation (Sim(3)) [28] is utilized for similarity scoring. Our method outperforms UniScene [45] obviously with ray-map normalization (‘Norm’).




Figure 8. **Distribution of camera control performance** with the rotation and translation errors on the NuScenes validation set.

RGB	Semantics	Depth	IoU ↑	mIoU ↑
✓	✓		28.9	17.1
✓		✓	31.5	16.8
✓	✓	✓	33.3	19.8

Table 6. **Ablation study on the occupancy prediction module.** The jointly generated semantics and depth effectively benefit occupancy prediction in contextual and geometric modeling.

Method	FVD ₁₇	FVD ₃₃	FVD ₆₅	FVD ₁₂₉	FVD ₂₀₁
w/. flexible forcing	23.63	24.14	24.72	25.05	25.22
w/o. flexible forcing	26.79	59.64	102.79	249.74	386.72

Table 7. **Ablation study on the flexible forcing strategy** on different video lengths. Our strategy yields high-quality results with long-term generation sequences.

the validation set. The RotErr achieved by our approach closely approaches ground-truth performance, while the higher TransErr reflects natural variability in driving distances across different scenarios. More evaluation results are provided in the supplementary material.

Zero-shot Generalization. Figure 10 presents zero-shot generalization results. We evaluate OmniNWM without fine-tuning across different datasets (e.g., nuplan or in-house collected datasets) and camera view configurations (e.g., 3 or 6 camera views). Leveraging the panoramic normalized ray-map encoding strategy with normalized Plücker ray-maps, our method demonstrates strong generalization capabilities across these diverse settings.




Figure 9. **Visualization on the effect of flexible forcing strategy.** Our approach effectively mitigates image degradation across multiple panoramic frames for long-term generation.




Figure 10. **Zero-shot generalization** across different new datasets and camera configurations without any fine-tuning.

4.3. Ablation Study

Occupancy Prediction. Table 6 presents ablation results for the occupancy prediction module. The baseline configuration predicts semantic occupancy from generated RGB, semantic, and depth maps within OmniNWM. The results indicate that the jointly generated semantic and depth maps yield noticeable improvements, enhancing the mIoU by 3.0 and 2.7, respectively. These results validate that the complementary benefits of the joint generated semantics and depth in contextual and geometric modeling.

Flexible Forcing. To validate the effectiveness of the flexible forcing strategy, we conduct ablation studies as shown in Table 7 and Figure 9. The results demonstrate that the proposed strategy effectively mitigates degradation for high-quality long-term generation, particularly in structural and textural details.

5. Conclusion and Future Work

OmniNWM establishes a new paradigm for driving world models by unifying state, action, and reward in a single omniscient framework. It jointly generates long-horizon, pixel-aligned panoramic videos of RGB, semantics, depth, and 3D occupancy (state); enables precise, generalizable camera control via normalized Plücker ray-maps (action); and introduces occupancy-grounded dense rewards for closed-loop evaluation (reward). This integration delivers state-of-the-art performance in generation fidelity, control accuracy, and zero-shot generalization, making OmniNWM a powerful and holistic framework for autonomous driving development. **Limitations and Future Work.** Despite its capabilities, the model remains computationally demanding, motivating future research into efficient deployment strategies.

References

- [1] Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, et al. Vd3d: Taming large video diffusion transformers for 3d camera control. *arXiv preprint arXiv:2407.12781*, 2024. 3
- [2] 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. 14, 15, 16
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. 4, 7, 15
- [4] Jianhong Bai, Menghan Xia, Xiao Fu, Xintao Wang, Lianrui Mu, Jinwen Cao, Zuoze Liu, Haoji Hu, Xiang Bai, Pengfei Wan, et al. Recammaster: Camera-controlled generative rendering from a single video. *arXiv preprint arXiv:2503.11647*, 2025. 3
- [5] Amir Bar, Gaoyue Zhou, Danny Tran, Trevor Darrell, and Yann LeCun. Navigation world models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15791–15801, 2025. 1
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 6, 15
- [7] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *CVPR*, 2022. 7, 14, 16
- [8] Rui Chen, Zehuan Wu, Yichen Liu, Yuxin Guo, Jingcheng Ni, Haifeng Xia, and Siyu Xia. Unimlv: Unified framework for multi-view long video generation with comprehensive control capabilities for autonomous driving, 2025. 3
- [9] Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, and Hongsheng Li. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *CVPR*, 2020. 7
- [10] Yuntao Chen, Yuqi Wang, and Zhaoxiang Zhang. Driving-gpt: Unifying driving world modeling and planning with multi-modal autoregressive transformers. *arXiv preprint arXiv:2412.18607*, 2024. 1, 7
- [11] Haohan Chi, Huan-ang Gao, Ziming Liu, Jianing Liu, Chenyu Liu, Jinwei Li, Kaisen Yang, Yangcheng Yu, Zeda Wang, Wenyi Li, et al. Impromptu vla: Open weights and open data for driving vision-language-action models. *arXiv preprint arXiv:2505.23757*, 2025. 5, 7, 14
- [12] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 15
- [13] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds. In *ISVC*, 2020. 14
- [14] Jiawei Du, Yidi Jiang, Vincent YF Tan, Joey Tianyi Zhou, and Haizhou Li. Minimizing the accumulated trajectory error to improve dataset distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3749–3758, 2023. 6
- [15] Hao Gao, Shaoyu Chen, Bo Jiang, Bencheng Liao, Yiang Shi, Xiaoyang Guo, Yuechuan Pu, Haoran Yin, Xiangyu Li, Xinbang Zhang, et al. Rad: Training an end-to-end driving policy via large-scale 3dgs-based reinforcement learning. *arXiv preprint arXiv:2502.13144*, 2025. 14
- [16] Ruiyuan Gao, Kai Chen, Bo Xiao, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrive-v2: High-resolution long video generation for autonomous driving with adaptive control. *arXiv preprint arXiv:2411.13807*, 2024. 1, 3, 7
- [17] 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. In *ICLR*, 2024. 1, 2, 3, 7
- [18] Shenyan 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. *Advances in Neural Information Processing Systems*, 37:91560–91596, 2025. 1, 2, 3, 5, 6, 7
- [19] Junhao Ge, Zuhong Liu, Longteng Fan, Yifan Jiang, Jiaqi Su, Yiming Li, Zhejun Zhang, and Siheng Chen. Unraveling the effects of synthetic data on end-to-end autonomous driving. *arXiv preprint arXiv:2503.18108*, 2025. 3
- [20] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 14
- [21] Jiazhe Guo, Yikang Ding, Xiwu Chen, Shuo Chen, Bohan Li, Yingshuang Zou, Xiaoyang Lyu, Feiyang Tan, Xiaojuan Qi, Zhiheng Li, et al. Dist-4d: Disentangled spatiotemporal diffusion with metric depth for 4d driving scene generation. *arXiv preprint arXiv:2503.15208*, 2025. 1, 7, 15
- [22] Xiangyu Guo, Zhanqian Wu, Kaixin Xiong, Ziyang Xu, Li-jun Zhou, Gangwei Xu, Shaoqing Xu, Haiyang Sun, Bing

- Wang, Guang Chen, et al. Genesis: Multimodal driving scene generation with spatio-temporal and cross-modal consistency. *arXiv preprint arXiv:2506.07497*, 2025. 1
- [23] Mariam Hassan, Sebastian Stapf, Ahmad Rahimi, Pedro Rezende, Yasaman Haghighi, David Brüggemann, Isinsu Katircioglu, Lin Zhang, Xiaoran Chen, Suman Saha, et al. Gem: A generalizable ego-vision multimodal world model for fine-grained ego-motion, object dynamics, and scene composition control. In *CVPR*, pages 22404–22415, 2025. 1, 3
- [24] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024. 3, 5, 7
- [25] Hao He, Ceyuan Yang, Shanchuan Lin, Yinghao Xu, Meng Wei, Liangke Gui, Qi Zhao, Gordon Wetzstein, Lu Jiang, and Hongsheng Li. Cameractrl ii: Dynamic scene exploration via camera-controlled video diffusion models. *arXiv preprint arXiv:2503.10592*, 2025. 3, 7
- [26] 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. *NeurIPS*, 2017. 7
- [27] Yuanduo Hong, Huihui Pan, Weichao Sun, and Yisong Jia. Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes. *arXiv preprint arXiv:2101.06085*, 2021. 15
- [28] Berthold KP Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the optical society of America A*, 4(4):629–642, 1987. 7, 8
- [29] Chen Hou, Guoqiang Wei, Yan Zeng, and Zhibo Chen. Training-free camera control for video generation. *arXiv preprint arXiv:2406.10126*, 2024. 3
- [30] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 3
- [31] Xiaotao Hu, Wei Yin, Mingkai Jia, Junyuan Deng, Xiaoyang Guo, Qian Zhang, Xiaoxiao Long, and Ping Tan. Driving-world: Constructing world model for autonomous driving via video gpt. *arXiv preprint arXiv:2412.19505*, 2024. 7
- [32] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *CVPR*, 2023. 7, 14
- [33] Jyh-Jing Hwang, Runsheng Xu, Hubert Lin, Wei-Chih Hung, Jingwei Ji, Kristy Choi, Di Huang, Tong He, Paul Covington, Benjamin Sapp, et al. Emma: End-to-end multimodal model for autonomous driving. *arXiv preprint arXiv:2410.23262*, 2024. 14
- [34] Hanwen Jiang, Hao Tan, Peng Wang, Haian Jin, Yue Zhao, Sai Bi, Kai Zhang, Fujun Luan, Kalyan Sunkavalli, Qixing Huang, et al. Rayzer: A self-supervised large view synthesis model. *arXiv preprint arXiv:2505.00702*, 2025. 3
- [35] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. TSIT: A simple and versatile framework for image-to-image translation. In *ECCV*, 2020. 15
- [36] Haian Jin, Hanwen Jiang, Hao Tan, Kai Zhang, Sai Bi, Tianyuan Zhang, Fujun Luan, Noah Snively, and Zexiang Xu. Lvsm: A large view synthesis model with minimal 3d inductive bias. *arXiv preprint arXiv:2410.17242*, 2024. 3
- [37] Seung Wook Kim, Jonah Philion, Antonio Torralba, and Sanja Fidler. Drivegan: Towards a controllable high-quality neural simulation. In *CVPR*, 2021. 7
- [38] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuo Zhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 3
- [39] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas J Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *Advances in Neural Information Processing Systems*, 37:16240–16271, 2024. 3
- [40] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 6
- [41] Bohan Li, Yasheng Sun, Xin Jin, Wenjun Zeng, Zheng Zhu, Xiaofeng Wang, Yunpeng Zhang, James Okae, Hang Xiao, and Dalong Du. Stereoscene: Bev-assisted stereo matching empowers 3d semantic scene completion. *arXiv preprint arXiv:2303.13959*, 2023. 14, 16
- [42] Bohan Li, Jiajun Deng, Wenyao Zhang, Zhujin Liang, Dalong Du, Xin Jin, and Wenjun Zeng. Hierarchical temporal context learning for camera-based semantic scene completion. In *European Conference on Computer Vision*, pages 131–148. Springer, 2024. 3, 7
- [43] Bohan Li, Xin Jin, Jiajun Deng, Yasheng Sun, Xiaofeng Wang, and Wenjun Zeng. Hierarchical context alignment with disentangled geometric and temporal modeling for semantic occupancy prediction. *arXiv preprint arXiv:2412.08243*, 2024. 14
- [44] Bohan Li, Yasheng Sun, Zhujin Liang, Dalong Du, Zhuanghui Zhang, Xiaofeng Wang, Yunnan Wang, Xin Jin, and Wenjun Zeng. Bridging stereo geometry and bev representation with reliable mutual interaction for semantic scene completion. In *IJCAI*, 2024. 3
- [45] Bohan Li, Jiazhe Guo, Hongsi Liu, Yingshuang Zou, Yikang Ding, Xiwu Chen, Hu Zhu, Feiyang Tan, Chi Zhang, Tiancai Wang, et al. Uniscene: Unified occupancy-centric driving scene generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 11971–11981, 2025. 1, 2, 3, 4, 6, 7, 8
- [46] Bohan Li, Xin Jin, Jianan Wang, Yukai Shi, Yasheng Sun, Xiaofeng Wang, Zhuang Ma, Baao Xie, Chao Ma, Xiaokang Yang, et al. Occscene: Semantic occupancy-based cross-task mutual learning for 3d scene generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 1, 3, 6, 7, 14
- [47] Jie Li, Kai Han, Peng Wang, Yu Liu, and Xia Yuan. Anisotropic convolutional networks for 3d semantic scene completion. In *CVPR*, 2020. 7
- [48] Ruilong Li, Brent Yi, Junchen Liu, Hang Gao, Yi Ma, and Angjoo Kanazawa. Cameras as relative positional encoding. *arXiv preprint arXiv:2507.10496*, 2025. 3
- [49] Samuel Li, Pujith Kachana, Prajwal Chidananda, Saurabh Nair, Yasutaka Furukawa, and Matthew Brown. Rig3r: Rig-

- aware conditioning for learned 3d reconstruction. *arXiv preprint arXiv:2506.02265*, 2025. 3
- [50] Shuai Liu, Quanmin Liang, Zefeng Li, Boyang Li, and Kai Huang. Gaussianfusion: Gaussian-based multi-sensor fusion for end-to-end autonomous driving. *arXiv preprint arXiv:2506.00034*, 2025. 14
- [51] Yang Liu, Ming Ma, Xiaomin Yu, Pengxiang Ding, Han Zhao, Mingyang Sun, Siteng Huang, and Donglin Wang. Ssr: Enhancing depth perception in vision-language models via rationale-guided spatial reasoning. *arXiv preprint arXiv:2505.12448*, 2025. 14, 15
- [52] Zhiheng Liu, Ka Leong Cheng, Qiuyu Wang, Shuzhe Wang, Hao Ouyang, Bin Tan, Kai Zhu, Yujun Shen, Qifeng Chen, and Ping Luo. Depthlab: From partial to complete. *arXiv preprint arXiv:2412.18153*, 2024. 15
- [53] Xiaoxiao Long, Lingjie Liu, Wei Li, Christian Theobalt, and Wenping Wang. Multi-view depth estimation using epipolar spatio-temporal networks. In *CVPR*, pages 8258–8267, 2021. 7
- [54] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9970–9980, 2024. 3
- [55] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6, 16
- [56] Jiachen Lu, Ze Huang, Jiahui Zhang, Zeyu Yang, and Li Zhang. Wovogen: World volume-aware diffusion for controllable multi-camera driving scene generation. *arXiv preprint arXiv:2312.02934*, 2023. 1, 7
- [57] Yifan Lu, Xuanchi Ren, Jiawei Yang, Tianchang Shen, Zhangjie Wu, Jun Gao, Yue Wang, Siheng Chen, Mike Chen, Sanja Fidler, and Jiahui Huang. Infinicube: Unbounded and controllable dynamic 3d driving scene generation with world-guided video models, 2025. 1
- [58] Jiageng Mao, Boyi Li, Boris Ivanovic, Yuxiao Chen, Yan Wang, Yurong You, Chaowei Xiao, Danfei Xu, Marco Pavone, and Yue Wang. Dreamdrive: Generative 4d scene modeling from street view images. *arXiv preprint arXiv:2501.00601*, 2024. 1, 3
- [59] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *IROS*. IEEE, 2019. 14
- [60] Arian Mousakhan, Sudhanshu Mittal, Silvio Galesso, Karim Farid, and Thomas Brox. Orbis: Overcoming challenges of long-horizon prediction in driving world models. *arXiv preprint arXiv:2507.13162*, 2025. 1
- [61] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kontschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 4990–4999, 2017. 15
- [62] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 15
- [63] Luis Roldao, Raoul de Charette, and Anne Verroust-Blondet. Lmscnet: Lightweight multiscale 3d semantic completion. In *3DV*, 2020. 7
- [64] Jan Philipp Schneider, Pratik Singh Bisht, Ilya Chugunov, Andreas Kolb, Michael Moeller, and Felix Heide. Neural atlas graphs for dynamic scene decomposition and editing. *arXiv preprint arXiv:2509.16336*, 2025. 3
- [65] Chen Shi, Shaoshuai Shi, Kehua Sheng, Bo Zhang, and Li Jiang. Drivex: Omni scene modeling for learning generalizable world knowledge in autonomous driving. *arXiv preprint arXiv:2505.19239*, 2025. 14
- [66] Rui Song, Chenwei Liang, Yan Xia, Walter Zimmer, Hu Cao, Holger Caesar, Andreas Festag, and Alois Knoll. Coda4dgs: Dynamic gaussian splatting with context and deformation awareness for autonomous driving. *arXiv preprint arXiv:2503.06744*, 2025. 3
- [67] Jian Sun, He Zhang, Huajun Zhou, Rongjie Yu, and Ye Tian. Scenario-based test automation for highly automated vehicles: A review and paving the way for systematic safety assurance. *IEEE transactions on intelligent transportation systems*, 23(9):14088–14103, 2021. 7
- [68] 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. 15
- [69] Alexander Szwedlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *IEEE Robotics and Automation Letters*, 2024. 7
- [70] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 3
- [71] Pin Tang, Zhongdao Wang, Guoqing Wang, Jilai Zheng, Xiangxuan Ren, Bailan Feng, and Chao Ma. Sparseocc: Rethinking sparse latent representation for vision-based semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15035–15044, 2024. 7
- [72] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Yang Wang, Zhiyong Zhao, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*, 2024. 14
- [73] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 7
- [74] Guoqing Wang, Zhongdao Wang, Pin Tang, Jilai Zheng, Xiangxuan Ren, Bailan Feng, and Chao Ma. Occgen: Generative multi-modal 3d occupancy prediction for autonomous driving. *arXiv preprint arXiv:2404.15014*, 2024. 7
- [75] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. 7, 8

- [76] Lening Wang, Wenzhao Zheng, Dalong Du, Yunpeng Zhang, Yilong Ren, Han Jiang, Zhiyong Cui, Haiyang Yu, Jie Zhou, Jiwen Lu, et al. Stag-1: Towards realistic 4d driving simulation with video generation model. *arXiv preprint arXiv:2412.05280*, 2024. 3
- [77] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Jiwen Lu, and Xingang Wang. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. *ICCV*, 2023. 4, 6, 7
- [78] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. *ECCV*, 2024. 1, 2, 3, 5, 7
- [79] 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 *CVPR*, 2024. 1, 2, 3, 7
- [80] Yuping Wang, Xiangyu Huang, Xiaokang Sun, Mingxuan Yan, Shuo Xing, Zhengzhong Tu, and Jiachen Li. Uniocc: A unified benchmark for occupancy forecasting and prediction in autonomous driving. *arXiv preprint arXiv:2503.24381*, 2025. 14
- [81] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 3
- [82] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Yongming Rao, Guan Huang, Jiwen Lu, and Jie Zhou. Surround-depth: Entangling surrounding views for self-supervised multi-camera depth estimation. In *Conference on robot learning*, pages 539–549. PMLR, 2023. 7
- [83] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *ICCV*, 2023. 14
- [84] 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 *CVPR*, 2024. 1, 2, 7
- [85] 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 *ICCV*, 2023. 3
- [86] Zhaoyang Xia, Youquan Liu, Xin Li, Xinge Zhu, Yuexin Ma, Yikang Li, Yuenan Hou, and Yu Qiao. Scpnnet: Semantic scene completion on point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17642–17651, 2023. 14
- [87] Baihui Xiao, Chengjian Feng, Zhijian Huang, Yujie Zhong, Lin Ma, et al. Robotron-sim: Improving real-world driving via simulated hard-case. *arXiv preprint arXiv:2508.04642*, 2025. 1
- [88] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34: 12077–12090, 2021. 15
- [89] Dejia Xu, Weili Nie, Chao Liu, Sifei Liu, Jan Kautz, Zhangyang Wang, and Arash Vahdat. Camco: Camera-controllable 3d-consistent image-to-video generation. *arXiv preprint arXiv:2406.02509*, 2024. 3
- [90] Jiawei Xu, Kai Deng, Zexin Fan, Shenlong Wang, Jin Xie, and Jian Yang. Ad-gs: Object-aware b-spline gaussian splatting for self-supervised autonomous driving. *arXiv preprint arXiv:2507.12137*, 2025. 3
- [91] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 14
- [92] Xu Yan, Jiantao Gao, Jie Li, Ruimao Zhang, Zhen Li, Rui Huang, and Shuguang Cui. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In *AAAI*, 2021. 7
- [93] Jiazhi Yang, Kashyap Chitta, Shenyuan Gao, Long Chen, Yuqian Shao, Xiaosong Jia, Hongyang Li, Andreas Geiger, Xiangyu Yue, and Li Chen. Resim: Reliable world simulation for autonomous driving. *arXiv preprint arXiv:2506.09981*, 2025. 3
- [94] 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. 7
- [95] Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. Direct-a-video: Customized video generation with user-directed camera movement and object motion. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–12, 2024. 3
- [96] Xueming Yang, Licheng Wen, Yukai Ma, Jianbiao Mei, Xin Li, Tiantian Wei, Wenjie Lei, Daocheng Fu, Pinlong Cai, Min Dou, et al. Drivearena: A closed-loop generative simulation platform for autonomous driving. *arXiv preprint arXiv:2408.00415*, 2024. 1
- [97] Yu Yang, Alan Liang, Jianbiao Mei, Yukai Ma, Yong Liu, and Gim Hee Lee. X-scene: Large-scale driving scene generation with high fidelity and flexible controllability. *arXiv preprint arXiv:2506.13558*, 2025. 1, 7
- [98] Zhuoran Yang, Xi Guo, Chenjing Ding, Chiyu Wang, Wei Wu, and Yanyong Zhang. Instadrive: Instance-aware driving world models for realistic and consistent video generation. In *ICCV*, 2025. 3
- [99] Zeyu Yang, Zijie Pan, Yuankun Yang, Xiatian Zhu, and Li Zhang. Driving view synthesis on free-form trajectories with generative prior, 2025.
- [100] Xin Ye, Burhaneddin Yaman, Sheng Cheng, Feng Tao, Abhirup Mallik, and Liu Ren. Bevdiffuser: Plug-and-play diffusion model for bev denoising with ground-truth guidance. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1495–1504, 2025. 3
- [101] Senthil Yogamani, Ciarán Hughes, Jonathan Horgan, Ganesh Sistu, Pádraig Varley, Derek O’Dea, Michal Uricár, Stefan

- Milz, Martin Simon, Karl Amende, et al. Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9308–9318, 2019. [15](#)
- [102] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. [15](#)
- [103] Shuang Zeng, Xinyuan Chang, Mengwei Xie, Xinran Liu, Yifan Bai, Zheng Pan, Mu Xu, and Xing Wei. Futuresight-drive: Thinking visually with spatio-temporal cot for autonomous driving. *arXiv preprint arXiv:2505.17685*, 2025. [5](#), [14](#)
- [104] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023. [15](#)
- [105] Kaiwen Zhang, Zhenyu Tang, Xiaotao Hu, Xingang Pan, Xiaoyang Guo, Yuan Liu, Jingwei Huang, Li Yuan, Qian Zhang, Xiao-Xiao Long, et al. Epona: Autoregressive diffusion world model for autonomous driving. *arXiv preprint arXiv:2506.24113*, 2025. [1](#), [7](#)
- [106] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *CVPR*, 2020. [14](#)
- [107] Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. *ICCV*, 2023. [14](#)
- [108] Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Xinze Chen, Guan Huang, Xiaoyi Bao, and Xingang Wang. Drivedreamer-2: Llm-enhanced world models for diverse driving video generation. *arXiv preprint arXiv:2403.06845*, 2024. [1](#), [2](#), [3](#), [5](#), [7](#)
- [109] Guosheng Zhao, Xiaofeng Wang, Chaojun Ni, Zheng Zhu, Wenkang Qin, Guan Huang, and Xingang Wang. Recon-dreamer++: Harmonizing generative and reconstructive models for driving scene representation. *arXiv preprint arXiv:2503.18438*, 2025. [3](#)
- [110] Rui Zhao, Yuze Fan, Ziguo Chen, Fei Gao, and Zhenhai Gao. Diffe2e: Rethinking end-to-end driving with a hybrid action diffusion and supervised policy. *arXiv preprint arXiv:2505.19516*, 2025. [14](#)
- [111] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. *arXiv preprint arXiv:2311.16038*, 2023. [4](#)
- [112] Wenzhao Zheng, Ruiqi Song, Xianda Guo, and Long Chen. Genad: Generative end-to-end autonomous driving. *arXiv preprint arXiv:2402.11502*, 2024. [1](#), [7](#)
- [113] Wenzhao Zheng, Zetian Xia, Yuanhui Huang, Sicheng Zuo, Jie Zhou, and Jiwen Lu. Doe-1: Closed-loop autonomous driving with large world model. *arXiv preprint arXiv:2412.09627*, 2024. [14](#)
- [114] Yupeng Zheng, Pengxuan Yang, Zebin Xing, Qichao Zhang, Yuhang Zheng, Yinfeng Gao, Pengfei Li, Teng Zhang, Zhongpu Xia, Peng Jia, et al. World4drive: End-to-end autonomous driving via intention-aware physical latent world model. *arXiv preprint arXiv:2507.00603*, 2025. [14](#)
- [115] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. [6](#), [15](#), [16](#)
- [116] Hongyu Zhou, Longzhong Lin, Jiabao Wang, Yichong Lu, Dongfeng Bai, Bingbing Liu, Yue Wang, Andreas Geiger, and Yiyi Liao. Hugsim: A real-time, photo-realistic and closed-loop simulator for autonomous driving. *arXiv preprint arXiv:2412.01718*, 2024. [3](#)
- [117] 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. [1](#)
- [118] Zewei Zhou, Tianhui Cai, Seth Z Zhao, Yun Zhang, Zhiyu Huang, Bolei Zhou, and Jiaqi Ma. Autovla: A vision-language-action model for end-to-end autonomous driving with adaptive reasoning and reinforcement fine-tuning. *arXiv preprint arXiv:2506.13757*, 2025. [14](#)
- [119] Yingshuang Zou, Yikang Ding, Xi Qiu, Haoqian Wang, and Haotian Zhang. M2depth: Self-supervised two-frame multi-camera metric depth estimation. In *European Conference on Computer Vision*, pages 269–285. Springer, 2024. [7](#)

Supplementary Material for OmniNWM

A. More Related Works and Discussions

A.1. Semantic Occupancy Prediction

Semantic Occupancy Prediction (SOP), also referred to as Semantic Scene Completion (SSC), is a holistic 3D perception task that jointly infers the semantic labels and geometric completeness of a scene from partial observations [7, 13, 32, 41, 43, 46, 80, 86, 106, 107]. Existing methods can be broadly categorized into LiDAR-based and camera-based approaches. LiDAR-based methods benefit from precise 3D measurements and have advanced significantly through architectures designed to handle sparse and irregular point clouds. For instance, PolarNet [106] leverages a polar BEV representation to address uneven point distribution, enabling efficient real-time inference. SalsaNext [13] improves segmentation via residual dilated convolutions and Bayesian uncertainty estimation. RangeNet++ [59] uses range images and a custom post-processing step to mitigate discretization artifacts. SCPNet [86] introduces multi-scale feature aggregation and dense-to-sparse distillation to enhance representation learning. Camera-based methods have gained traction due to their lower cost and broader applicability. MonoScene [7] pioneers monocular 3D semantic completion by projecting 2D features into 3D space. StereoScene [41] further bridges stereo geometry and BEV features for improved semantic and geometric consistency. Subsequent works such as TPVFormer [32] introduce tri-perspective views for richer 3D representations, while SurroundOcc [83] and OccFormer [107] leverage multi-view images and transformer-based aggregation for dense occupancy estimation. Despite these advances, most existing methods rely on limited modalities, hindering their robustness in complex driving scenarios. In contrast, our approach first generates pixel-aligned multi-modal outputs of RGB, semantic, and metric depth from a reference image, then leverages these complementary modalities to predict 3D semantic occupancy. This joint generation strategy fully exploits both geometric and contextual information, leading to more accurate and consistent 3D scene understanding.

A.2. Vision-Language Models

Recent years have witnessed growing interest in integrating Vision-Language Models (VLMs) and Large Language Models (LLMs) into autonomous driving systems, capitalizing on their strong reasoning capabilities and world knowledge [2, 15, 33, 50, 51, 65, 72, 91, 103, 110, 113, 114]. These models often exhibit emergent reasoning skills, such

as chain-of-thought (CoT) reasoning, which can be harnessed for interpretable decision-making. For example, DriveVLM [72] combines LLM-based trajectory proposals with an end-to-end refinement module for final planning. SSR [51] addresses the limited spatial understanding of VLMs by converting depth maps into textual rationales and distilling them into lightweight embeddings for plug-and-play spatial reasoning. DriveGPT4 [91] uses LLMs in an iterative QA framework to explain driving behaviors and predict control signals. Doe-1 [113] reformulates driving as a next-token prediction task using a multi-modal generative model, while EMMA [33] encodes non-sensor inputs and outputs as natural language to fully leverage pre-trained LLMs. AutoVLA [118] utilizes raw visual inputs and language instructions for semantic reasoning and trajectory planning. FSDrive [103] introduces a pre-training paradigm for spatio-temporal reasoning, and Impromptu-VLA [11] provides a large-scale annotated dataset and baseline models for vision-language-action tasks in unstructured driving scenarios. In this work, we introduce OmniNWM-VLA, a semantic-geometric reasoning agent built upon Qwen-2.5-VL [2]. It integrates multi-modal inputs of RGB, depth, and semantic maps via a Mamba-based [20] fusion block, enabling comprehensive scene understanding and trajectory planning.

B. Architecture Overview of OmniNWM-VLA

As shown in Figure 11, OmniNWM-VLA is a semantic-geometric reasoning agent designed to bridge high-level scene understanding with precise trajectory planning in autonomous driving. Built upon Qwen-2.5-VL [2], it integrates multi-modal inputs of RGB, depth, and semantic maps through a Mamba-based [20] fusion block, enabling comprehensive spatial reasoning and trajectory prediction. The agent outputs not only waypoint coordinates but also heading angles, which are essential for generating normalized Plücker ray-maps used in panoramic video generation. In this way, our model supports precise control through Plücker ray-maps, bridging the gap between high-level reasoning and low-level action execution in a unified framework.

Tri-Modal Mamba-based Interpreter. At the core of OmniNWM-VLA is the Tri-Modal Mamba-based Interpreter (Tri-MIDI), a plug-and-play module that fuses visual, geometric, and semantic information into a unified latent representation. This design allows the model to reason jointly over appearance (RGB), 3D structure (depth), and scene semantics (segmentation), which is critical for forecasting geo-




Figure 11. **Architecture overview of OmniNWM-VLA.** The Tri-Modal Mamba-based Interpreter (Tri-MIDI) fuses visual, geometric, and semantic information into a unified latent representation to forecast waypoint coordinates and heading angles.

metrically consistent and semantically plausible trajectories.

Given aligned multi-view inputs—RGB images X_V , metric depth maps X_D , and semantic segmentation maps X_S —we first embed them into latent tokens P_V, P_D, P_S to preserve spatial consistency across views. Each modality is encoded separately using pre-trained encoders similar to the previous approach [51]. Specifically, we leverage CLIP [62] for RGB images, SigLIP [104] for depth maps, and SegFormer [88] for semantic maps:

$$H_\alpha = E_\alpha(P_\alpha), \quad \alpha \in \{V, D, S\}. \quad (15)$$

The features are then projected into a common embedding space via modality-specific MLPs:

$$Z_\alpha = \phi_\alpha(H_\alpha), \quad \alpha \in \{V, D, S\}. \quad (16)$$

A Mamba-based state-space model f_{LM} then performs cross-modal fusion, guided by a textual query X_T , to produce a fused latent representation:

$$H_R = f_{LM}([Z_V; Z_D; Z_S], X_T). \quad (17)$$

Using the Tokenized Rationale (TOR) mechanism [3], we insert special tokens into the query sequence to anchor intermediate reasoning steps. The output tokens at these positions are projected into the latent space of the downstream VLM:

$$Z_{TOR} = \phi_R(H_{R_{tor}}). \quad (18)$$

These tokens enrich the VLM’s input with spatially and semantically grounded cues, enabling it to predict trajectories that include both 2D waypoints and heading angles. This dual-output format is crucial for constructing normalized Plücker ray-maps, which provide pixel-level control over panoramic video generation in the OmniNWM framework.

By integrating multi-modal perception with sequential reasoning, OmniNWM-VLA supports closed-loop evaluation and long-horizon planning, effectively connecting high-level decision-making with low-level action execution in a unified driving world model pipeline.

C. Semantic and Depth Data Curation

To construct semantic labels for the NuScenes dataset [6], we finetune a pretrained DDRNet model [27] on diverse driving scenes to enhance robustness. Training data includes Cityscapes [12], Mapillary Vistas [61], Waymo Open [68], Woodscape [101], and BDD100k [102] datasets. Image-to-image translation techniques [35] are employed to synthesize night-scene imagery, further improving model generalization. The finetuned DDRNet generates semantic segmentation labels for the NuScenes dataset at 12 Hz.

For metric depth supervision, we follow established practices [21]: sparse LiDAR-projected depth and MVS-reconstructed depth maps, along with corresponding RGB images, are fed into a depth completion model [52] to generate dense, accurate depth maps.

D. Training Objectives

Our framework comprises three core components: the Vision-Language-Action (VLA) planner, the Panoramic Diffusion Transformer (PDiT) for multi-modal video generation, and the 3D semantic occupancy prediction module. Each component is trained with specialized objectives tailored to its respective task.

VLA Planning Objective. The OmniNWM-VLA planner is built upon Qwen-2.5-VL [2] and inherits its training methodology. The model is optimized using the standard causal language modeling objective, where the loss is computed as the negative log-likelihood of predicting the next token in the sequence:

$$\mathcal{L}_{VLA} = -\mathbb{E}_{(x,y) \sim \mathcal{D}} \sum_{t=1}^T \log P(y_t | y_{<t}, x), \quad (19)$$

where x represents the multi-modal input (concatenated RGB, depth, and semantic tokens processed by Tri-MIDI), y is the target sequence containing both trajectory waypoints and heading angles, and T is the sequence length. This objective enables the model to learn coherent multi-step reasoning for trajectory planning while maintaining compatibility with pre-trained VLMs.

Panoramic Video Generation Objective. For the PDiT-based panoramic video generation, we adopt the flow matching approach similar to OpenSora [115]. Let X_0 denote the video latent representation and $X_1 \sim \mathcal{N}(0, 1)$ represent Gaussian noise. The model f_θ takes as input an interpolated latent $X_t = (1-t)X_0 + tX_1$ and is trained to predict the velocity component $X_0 - X_1$. The corresponding loss function is formulated as:

$$\mathcal{L}_{PDiT} = \mathbb{E}_{t, X_0, X_1} [\|f_\theta(X_t, t, y) - (X_0 - X_1)\|], \quad (20)$$

where y represents the conditioning input (historical trajectory and reference image), and the timestep t is first sampled

from a logit-normal distribution and then scaled according to the shape of X_0 . This objective enables stable training and high-quality generation of panoramic RGB, semantic, and depth videos.

3D Semantic Occupancy Objective. For the 3D semantic occupancy prediction module, we follow the learning objective established in previous semantic scene completion works [7, 41]. The overall loss function combines multiple supervision signals:

$$\mathcal{L}_{\text{Occ}} = \mathcal{L}_{\text{depth}} + \lambda_{\text{sem}}\mathcal{L}_{\text{sem}} + \lambda_{\text{geo}}\mathcal{L}_{\text{geo}} + \lambda_{\text{ce}}\mathcal{L}_{\text{ce}}, \quad (21)$$

where $\mathcal{L}_{\text{depth}}$ is the Binary cross-entropy loss that encourages accurate sparse depth distribution estimation. $\mathcal{L}_{\text{sem}}$ is the standard semantic loss for voxel-wise semantic classification. $\mathcal{L}_{\text{geo}}$ is the geometry loss that enforces structural consistency in the 3D volume. $\mathcal{L}_{\text{ce}}$ is the class-weighted cross-entropy loss to handle class imbalance. The λ coefficients are balancing hyperparameters tuned to ensure stable training across different supervision signals. This multi-objective approach ensures that the generated occupancy volumes are both semantically accurate and geometrically plausible.

E. Training Details

Panoramic Diffusion Transformer (PDiT). The training of our PDiT follows a progressive three-stage curriculum to ensure stable convergence and robust generalization:

Stage 1: Single-view Control (10K iterations). The model is initially trained on single-view sequences with Plücker ray-map control signals. We use a fixed frame length of 17 and an input resolution of 224×400 to establish basic generation capabilities.

Stage 2: Multi-view and Multi-modal Extension (3K iterations). The model is extended to handle 6 panoramic camera views while incorporating joint generation of RGB, semantic, and depth modalities.

Stage 3: Variable-length and Resolution Training (3K iterations). Following [115], we introduce variability in sequence length (17 or 33 frames) and resolution (224×400 or 448×800) to improve adaptability to diverse driving scenarios and computational constraints.

Throughout all stages, we employ the AdamW optimizer [55] with a learning rate of 1×10^{-4} and weight decay of 0.01. The model is trained on 48 NVIDIA A800 GPUs with a total batch size of 48.

3D Semantic Occupancy Objective. The occupancy generation module is implemented in PyTorch and trained for 24 epochs with a batch size of 8. We use the AdamW optimizer with a learning rate of 1×10^{-4} and weight decay of 0.01. Data augmentation includes random horizontal flipping and color jittering for RGB inputs.

OmniNWM-VLA. Our OmniNWM-VLA planner follows the same training procedure as Qwen-2.5-VL [2], with modi-




Figure 12. **Average rewards for specific scenarios** across VLA planning baselines. Trajectories are evaluated using our occupancy-grounded reward function.

fications to accommodate autonomous driving requirements. The model uses a 3-B parameter architecture and is trained on a curated dataset of driving scenarios with multi-modal annotations. Unlike previous methods that operate at 2Hz, our model is trained at a control frequency of 12Hz to enable finer-grained trajectory planning. The training objective combines next-token prediction for trajectory waypoints and heading angles with multi-modal understanding tasks.

Overall Training Strategy. While each component can be trained independently, we employ a staged training strategy: first pre-training the PDiT and occupancy modules, then fine-tuning the VLA planner with the generated multi-modal outputs. This approach ensures that each module reaches optimal performance before being integrated into the full OmniNWM pipeline for closed-loop evaluation.

F. More Visualization Results

Trajectory Planning Evaluation. Figure 12 presents the average rewards for specific scenarios on the NuScenes validation set across VLA planning baselines, evaluated using our occupancy-grounded driving reward. The results indicate that our reward function effectively distinguishes between the planning trajectories generated by different VLAs when the ego vehicle navigates near drivable area boundaries of median strips, or oncoming trucks.

Precise Panoramic Camera Control. As shown in the supplementary figures (Figures 13, 14, 15 and 16), including a case of reversing trajectories (Figure 17), OmniNWM enables precise panoramic camera control using a normalized Plücker ray-map representation. By converting diverse input trajectories into a unified pixel-level format, our method facilitates accurate viewpoint manipulation while preserving geometric consistency across camera views. This control mechanism is vital for generating realistic, physically plausible driving scenarios.

Long-term Navigation Stability. As shown in Figure 18, OmniNWM achieves robust long-term generation beyond ground-truth sequence lengths, producing coherent sequences of up to 321 frames. This capability is enabled by our flexible forcing strategy, which injects independent multi-level noise during training to mitigate error accumulation in auto-regressive inference. The results validate the model’s stability for extended closed-loop navigation simulations.

Comprehensive Multi-modal Generation. Figure 19 showcases OmniNWM’s joint generation of panoramic RGB, semantic segmentation, metric depth, and 3D occupancy outputs. The pixel-level alignment across modalities ensures geometric and semantic consistency, providing a comprehensive environment representation for downstream tasks such as trajectory planning and reward computation.




Figure 13. **Precise panoramic camera control via normalized Plücker ray-maps.** Given the same conditional frame, OmniNWM generates consistent multi-view videos by interpreting different input trajectories into pixel-level control signals.




Figure 14. Additional examples of panoramic camera control.



Figure 15. Additional examples of panoramic camera control.




Figure 16. Additional examples of panoramic camera control.




Figure 17. Precise panoramic camera control of reversing trajectories via normalized Plücker ray-maps.




Figure 18. **Long-term navigation sequence (321 frames) generated through flexible forcing strategy.** The model maintains temporal coherence and structural integrity beyond training sequence lengths, enabling extended closed-loop evaluation.




Figure 19. **Joint multi-modal generation results**, including panoramic RGB frames, semantic segmentation maps, metric depth estimations, and corresponding 3D semantic occupancy, demonstrating comprehensive scene generation capabilities.