

DriveX: Omni Scene Modeling for Learning Generalizable World Knowledge in Autonomous Driving

Chen Shi^{1◊} Shaoshuai Shi² Kehua Sheng² Bo Zhang² Li Jiang^{1†}

¹The Chinese University of Hong Kong, Shenzhen

²Didi Chuxing, China

chenshi@link.cuhk.edu.cn jiangli@cuhk.edu.cn

Abstract

Data-driven learning has advanced autonomous driving, yet task-specific models struggle with out-of-distribution scenarios due to their narrow optimization objectives and reliance on costly annotated data. We present DriveX, a self-supervised world model that learns generalizable scene dynamics and holistic representations (geometric, semantic, and motion) from large-scale driving videos. DriveX introduces Omni Scene Modeling (OSM), a module that unifies multimodal supervision—3D point cloud forecasting, 2D semantic representation, and image generation—to capture comprehensive scene evolution. To simplify learning complex dynamics, we propose a decoupled latent world modeling strategy that separates world representation learning from future state decoding, augmented by dynamic-aware ray sampling to enhance motion modeling. For downstream adaptation, we design Future Spatial Attention (FSA), a unified paradigm that dynamically aggregates spatiotemporal features from DriveX’s predictions to enhance task-specific inference. Extensive experiments demonstrate DriveX’s effectiveness: it achieves significant improvements in 3D future point cloud prediction over prior work, while attaining state-of-the-art results on diverse tasks including occupancy prediction, flow estimation, and end-to-end driving. These results validate DriveX’s capability as a general-purpose world model, paving the way for robust and unified autonomous driving frameworks.

1. Introduction

Data-driven learning techniques have significantly advanced autonomous driving, enhancing core tasks like perception [27, 48, 51, 66, 71] and planning [22, 23, 37, 52]. Despite this impressive progress, task-specific models still struggle with complex and out-of-distribution scenarios,

◊: Intern of Voyager Research; †: Corresponding author.

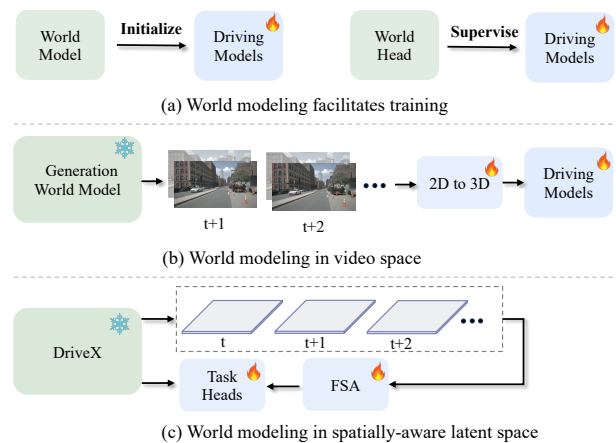


Figure 1. Comparison of different methods for integrating world models in autonomous driving systems. (a) World modeling serves as a pretraining task or an auxiliary supervision for driving models, without fully capturing general world knowledge. (b) Future prediction is performed in video space to assist downstream tasks, requiring additional 2D-to-3D transformation. (c) Our proposed DriveX framework decodes future states in latent 3D space, learning task-agnostic features that seamlessly benefit driving tasks.

primarily due to their task-specific optimization and reliance on limited human-annotated training data. World models, which predict future environment states based on historical observations and ego actions, have gained increasing attention for their ability to model complex driving environments [1, 11, 12, 32, 44, 55, 56, 62, 63]. By learning generalizable world knowledge from large-scale driving videos, world models hold the potential to improve robustness and safety across various autonomous driving tasks.

Existing world modeling approaches in autonomous driving can be broadly categorized into three strategies, as shown in Figure 1. The first strategy involves using world modeling as a general pretraining step [1, 44, 63], followed by fine-tuning for specific driving tasks. The second strategy treats world modeling as an auxiliary head [12, 31, 34, 62, 72] during task-specific training. While effective, both

strategies struggle to retain general world knowledge and may not fully exploit the world model’s potential for detailed spatial and temporal prediction, as the model parameters are ultimately adapted for specific downstream tasks. The third strategy uses world model as a generative model to predict future video sequences [11, 21, 32, 55, 56] by utilizing video generation techniques [17, 24, 54], which are subsequently used for task-specific purposes. Although this approach explicitly provides valuable predictive context by functioning as a world simulator, these generative models frequently struggle with maintaining consistent temporal-spatial coherence and lack adequate 3D spatial information, limiting their practical applications. These challenges motivate us to explore a fundamental question: how can we develop a world model framework that extracts generalizable world representations to benefit various downstream autonomous driving tasks within a unified paradigm?

We argue that an effective world model should excel in two key dimensions: 1) it should perceive the environment by capturing both geometric structure and visual semantics—critical for ensuring that learned representations can be effectively transferred across diverse tasks; and 2) it should accurately model world dynamics, which is essential for decision-making tasks such as planning. In this work, we propose DriveX, a novel world model framework that not only preserves essential world knowledge within a spatially-aware bird’s-eye view (BEV) latent space—facilitating general knowledge transfer across multiple downstream tasks—but also captures precise world evolution. As illustrated in Figure 1(c), DriveX demonstrates its extensibility by seamlessly integrating with various autonomous driving tasks.

DriveX distinguishes itself from existing world modeling approaches through two key aspects. First, existing methods face a critical trade-off: prior work either focuses on geometric cues (e.g., point cloud forecasting [26, 63, 70]) or semantic occupancy prediction using manual labels [43, 62, 72], sacrificing either visual details or scalable supervision. To overcome these limitations, we introduce the Omni Scene Modeling (OSM) module, which encodes rich multimodal world knowledge into a bird’s-eye view (BEV) latent space. Specifically, OSM integrates 3D point cloud forecasting for geometric fidelity, 2D semantic representations for class-aware reasoning, and image generation for retaining visual texture information. As shown in Table 1, when used as a pretrained feature extractor for end-to-end driving, DriveX maintains robust performance even with frozen parameters. In contrast, ViDAR[63] and DriveWorld[43] experience significant performance degradation when parameters are frozen, highlighting the generalizable knowledge captured by our approach. Moreover, OSM operates in a fully self-supervised manner, leveraging either inherently available labels or those generated by

Method	Finetuned	Fixed	Gap
ViDAR [63]	83.6	77.0	-6.6
DriveWorld [43]†	83.7	82.2	-1.5
DriveX	83.7	83.6	-0.1

Table 1. End-to-end driving performance on NAVSIM [8] test set. “Finetuned” refers to training the world model simultaneously, while “Fixed” represents freezing the world model. We observe negligible performance gap in DriveX. †: re-implemented by us.

foundation models [46, 69] for scalable learning.

Second, we introduce a decoupled latent world modeling strategy that separates learning into two key components: world representation learning and latent future decoding. This hierarchical approach first builds a robust spatial representation before modeling temporal evolution, effectively capturing sparse motion signals. In the world representation learning component, we train a world encoder to generate structured BEV features from multi-view historical images. These BEV features are optimized by our OSM module and thus are generalizable across driving tasks. In the latent future decoding component, we train a future decoder to model world evolution in the latent BEV space. Given the pretrained world encoder, the future decoder predicts motion flow and future BEV features based on given ego actions, with self-supervised alignment against the encoded future video frames by the world encoder. To further enhance motion modeling, we also apply OSM loss on the predicted latent features and introduce a dynamic-aware ray sampling strategy that prioritizes motion-salient regions.

To fully exploit DriveX’s predictive capabilities in downstream tasks, we introduce the Future Spatial Attention (FSA) paradigm, a lightweight yet effective module that seamlessly integrates DriveX into various driving applications. By acting as a “world simulator”, DriveX predicts future latent BEV features over the next few seconds, providing a rich predictive context. FSA then employs task-specific queries to dynamically aggregate relevant contextual information from these features. This streamlined approach improves state-of-the-art performance across diverse tasks such as occupancy prediction, occupancy flow estimation, and end-to-end driving, highlighting DriveX’s strong generalization and predictive capabilities.

In summary, our contributions are four-fold: 1) We introduce DriveX, a self-supervised world model that learns generalizable world knowledge in latent space through Omni Scene Modeling with multimodal supervision; 2) We propose a decoupled latent world modeling strategy that separates world representation learning from latent future decoding, enhanced by dynamic-aware ray sampling to prioritize motion-salient regions; 3) We present the Future Spatial Attention (FSA) paradigm, which leverages the pretrained world model to enhance diverse driving tasks within a unified framework; 4) DriveX achieves significant improvements across multiple tasks, achieving a rel-

ative 6% reduction in Chamfer Distance for 3-second future point cloud prediction, along with notable improvements in downstream tasks such as occupancy prediction, occupancy flow estimation, and end-to-end driving.

2. Related Work

Point Cloud Forecasting Point cloud forecasting, which aims to predict future point clouds from historical observations, has emerged as a crucial self-supervised learning paradigm in autonomous driving. Early approaches [42, 57, 58] typically project past point clouds into dense range images and leverage LSTM or 3D convolutional networks for future range image prediction. Recent work 4D-Occ [26] decomposes the task into 4D occupancy forecasting and employs differential depth rendering [25, 45] to generate point cloud predictions from occupancy grids. COPILOT4D [70] extends this line of work by incorporating diffusion models [19, 49]. In a different direction, ViDAR [63] explores a vision-centric approach by predicting future point clouds from past visual inputs. Unlike existing researches that primarily utilize point cloud forecasting as a pretraining strategy for encoder in downstream tasks, our work adopts a fundamentally distinct perspective: we integrate point cloud forecasting as a component of the world model to learn more generalizable representations.

World Models World models [13, 28], inspired by human cognition, enable intelligent agents to predict future outcomes based on past experiences and current actions. Learning world models has been widely studied to enhance policy function in reinforcement learning [13–16, 47, 60], but their adoption in real-world applications remains limited. Another line of research emphasizes world simulation, aiming to generate high-fidelity future videos in gaming [2, 4, 50] and real-world scenarios [5, 20, 29, 64]. However, these methods often fail to provide consistent and precise spatiotemporal representation. In this work, we target real-world autonomous driving and construct world models directly in the spatial BEV space, enabling transferable representations that support diverse downstream driving tasks.

World Modeling for Autonomous Driving As a complex and safety-critical real-world application, autonomous driving has increasingly attracted interest in world modeling. Leveraging advances in video generation, several approaches [9–11, 21, 32, 55, 56] have demonstrated the capability to synthesize realistic driving sequences across diverse scenarios. However, the generated videos lack comprehensive 3D scene understanding, restricting their applications in downstream tasks. Alternative methods [33, 44, 62, 72] explore future prediction in occupancy space, achieving remarkable progress but relying on costly semantic occupancy annotations. Meanwhile, research on applying world models to real-world driving applications mainly focuses on using world models as auxiliary supervi-

sion [31, 34, 34, 63] for autonomous model training, failing to capture task-agnostic world features. Recent works such as LAW [34] and SSR [31] have also explored world modeling in latent BEV space to eliminate the need for manual labels, but are limited to the end-to-end driving task. To address these challenges, we present DriveX, a spatially-aware world model that learns general world knowledge in a self-supervised manner and offers seamless integration with various driving tasks via spatial attention.

3. Methodology

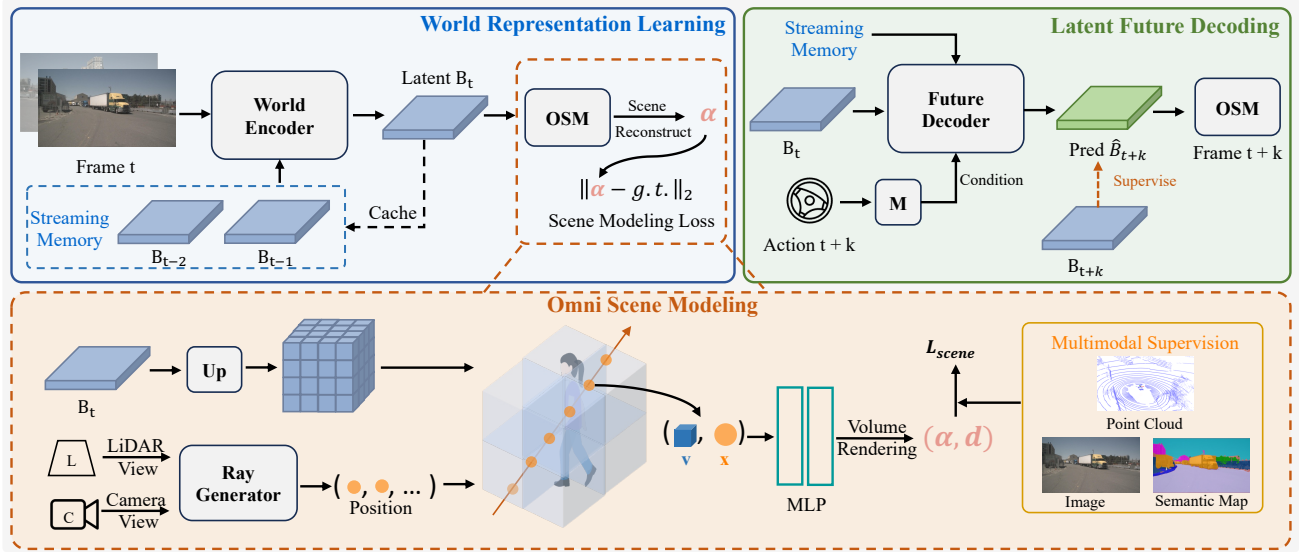
In autonomous driving scenarios, world models are typically designed to predict future environmental states from past observations and ego actions, enabling generalized scene understanding. Some approaches leverage world modeling as a pre-training step [26, 63, 70] or as an auxiliary task [17, 24, 54] but often compromise transferability due to reliance on specialized supervision. Another line of world models [11, 32, 55, 56] explores to generate future video sequences but faces challenges in ensuring the spatial and temporal consistency.

To advance the utility of the world model in driving tasks, we propose DriveX, as depicted in Figure 2. DriveX leverages multi-view historical camera video and ego action data to predict future world states within a spatially-aware latent space. To learn general world representation, we propose the Omni Scene Modeling module (described in Sec 3.1), which employs a multimodal supervision strategy to simultaneously encode geometric, semantic, and visual details in a self-supervised manner. To further improve world dynamics modeling, we design a decoupled learning strategy that enables precise motion signal learning based on well-trained BEV latent representations (See Sec 3.2). Finally, we present the Future Spatial Attention (FSA) mechanism (See Sec 3.3), which seamlessly integrates DriveX with various downstream driving tasks.


3.1. World Representation Learning

To encode environmental information into latent representations, prior works have explored point cloud reconstruction [26, 63] and occupancy prediction as proxy tasks [43, 62, 72]. However, these specialized objectives fail to fully capture comprehensive scene information (as demonstrated in Table 1). To address this limitation, we design the Omni Scene Modeling (OSM) module that promotes the learning of generalizable representations. Leveraging foundation models [46, 69] and paired multi-view images and point cloud data, this learning process is fully self-supervised.

Temporal Streaming World Encoder. Given a sequence of multi-view images $I_{1:t}$ up to the current timestamp t , we apply a world encoder to compress the multi-view historical observations into a latent BEV representation $B_t \in \mathbb{R}^{H \times W \times C}$, where H , W and C are the width, height and



(a) World Model Learning



(b) Future Spatial Attention

Figure 2. **Top:** An overview of our DriveX framework. The learning process consists of two stages: world representation learning, where the model learns temporal and geometric semantics through Omni Scene Modeling, and latent future decoding, where the model predicts future states in the learned latent space. Both stages are trained in a self-supervised manner. **Bottom:** Illustration of the FSA paradigm. Various driving tasks can dynamically aggregate information from predicted latent features through FSA.

number of channels of the feature map, respectively. We follow the BEVFormer [35] paradigm for 2D-to-3D view transformation. Taking previous latent BEV features B_{t-1} in streaming memory and images I_t from N_{view} cameras as inputs, the world encoder extracts visual information and transforms it into grid-structured BEV features, as follows:

$$B_t = \text{TCA}(E(I_t), B_{t-1}), \quad (1)$$

where E denotes the combination of a 2D encoder [18, 39, 53] and a 2D-to-3D transformation network [35]. TCA represents a Temporal Cross-Attention module [74] that achieves interaction between current states and historical latent information. The encoder operates in a streaming manner, continuously caching BEV features in the streaming memory to maintain long-term temporal dependencies.

In this way, we obtain a compact and structured representation B_t , where each grid cell corresponds to a specific location in the physical world. The spatial priors of grids in B_t facilitate scene dynamics modeling by enabling independent processing of each grid cell. Furthermore, such structured representation ensures strong compatibility with diverse practical applications (detailed in Section 3.3).

Omni Scene Modeling. A well-designed training signal is essential for the world encoder to capture generalizable latent BEV features. To achieve this, we design a omni scene modeling module to compel the world encoder to develop a comprehensive understanding of the environment. Specifically, we categorize the environmental understanding into three aspects: geometric structure, semantic information, and visual details. Then we introduce a multimodal supervision strategy, combining 3D point cloud reconstruction, 2D semantic prediction, and image generation to improve the richness of the latent representation.

To be specific, let $I_t \in \mathbb{R}^{N_{\text{view}} \times H_{\text{img}} \times W_{\text{img}} \times 3}$ and $P_t \in \mathbb{R}^{N_p \times 3}$ represent the multi-view images with shape $(W_{\text{img}}, H_{\text{img}})$ and point clouds with N_p points, respectively. Firstly, we cast N_r rays from image and LiDAR views by sampling corresponding pixels and points. Each ray can be parameterized as $r_i = o_i + td_i$, where o_i denotes the camera or LiDAR center and d_i indicates the ray direction pointing to the sampled pixel or point. Then we sample n waypoints $x_i \in \mathbb{R}^{n \times 3}$ along the ray r_i . The coordinates of these waypoints are computed as follows:

$$x_{i,j} = o_i + \lambda_j d_i, \quad (2)$$

where i and j denotes the index of ray and waypoint, respectively, and $\lambda = \{\lambda_1, \lambda_2, \dots, \lambda_n\}$ represents predefined distances.

The waypoint features $V_t \in \mathbb{R}^{N_r \times n \times C^*}$ are obtained by first transforming the BEV features B_t into 3D space through a Channel-to-Height module, yielding volumetric features $F_t \in \mathbb{R}^{H \times W \times Z \times C^*}$, followed by a trilinear interpolation operation:

$$\begin{aligned} F_t &= \text{Reshape}(\text{Conv}(B_t), [H, W, Z, C^*]), \\ V_t &= \text{Interp}(F_t, X_t), \end{aligned} \quad (3)$$

where Conv denotes 2D convolution layer with input channels C and output channels ZC^* , and $X_t = \{x_i\}_{i=1}^{N_r}$ represents the collection of ray waypoints. Subsequently, for each waypoint $x_{i,j}$, we utilize a MLP to predict its attribution filed $\alpha_{i,j} \in \mathbb{R}^{N_s + N_c}$ and density field $\sigma_{i,j} \in \mathbb{R}$:

$$(\alpha_{i,j}, \sigma_{i,j}) = \text{MLP}(\text{Concat}(v_{i,j}, \text{PE}(x_{i,j}))), \quad (4)$$

where $v_{i,j}$ represents the waypoint features from V_t , N_s and N_c indicate the number of semantic classes and RGB color channels, respectively. PE is the cosine position encoding. Finally, the attribute prediction α_i for r_i is accumulated through volume rendering process [45, 61]:

$$\begin{aligned} \tau_j &= \exp\left(-\sum_{l=1}^{j-1} \sigma_{i,l} \delta_l\right), \\ \alpha_i &= \sum_{j=1}^n \tau_j (1 - \exp(-\sigma_{i,j} \delta_j)) \alpha_{i,j}, \end{aligned} \quad (5)$$

where $\delta_j = \lambda_j - \lambda_{j-1}$ ($\lambda_0 = o_i$) represents the distance between two adjacent waypoints, and α_i can be decomposed into semantic prediction $s_i \in \mathbb{R}^{N_s}$ and color prediction $c_i \in \mathbb{R}^{N_c}$. Following the same formulation, the expected depth d_i of r_i is estimated by replacing $\alpha_{i,j}$ with the distance λ_j in Eq.(5).

To enable comprehensive scene understanding, we train the world encoder using a multimodal supervision framework, deriving the scene modeling loss $\mathcal{L}_{\text{scene}}$ from both camera and LiDAR perspectives. Specifically, for camera-view rays, we optimize the scene modeling loss $\mathcal{L}_{\text{scene}}^{\text{camera}}$ that integrates semantic segmentation (\mathcal{L}_{sem}) and color reconstruction (\mathcal{L}_{rgb}) objectives. For LiDAR-view rays, we optimize the scene modeling loss $\mathcal{L}_{\text{scene}}^{\text{LiDAR}}$ with only depth estimation ($\mathcal{L}_{\text{depth}}$):

$$\begin{aligned} \mathcal{L}_{\text{scene}}^{\text{camera}} &= \mathcal{L}_{\text{sem}} + \mathcal{L}_{\text{rgb}}, \\ \mathcal{L}_{\text{scene}}^{\text{LiDAR}} &= \mathcal{L}_{\text{depth}}, \end{aligned} \quad (6)$$

where we adopt cross-entropy loss for \mathcal{L}_{sem} and L2 Loss for both \mathcal{L}_{rgb} and $\mathcal{L}_{\text{depth}}$. The ground truth depth and color values are directly acquired from raw sensor data, while semantic labels are automatically generated using foundation models Grounded SAM [46] and OpenSeeD [69], enabling scalable training without manual annotation.

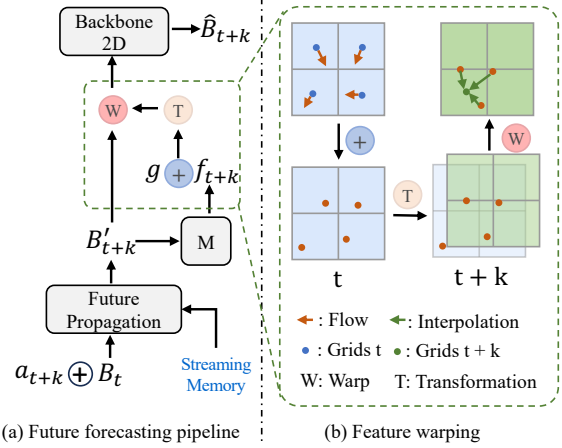


Figure 3. Illustration of the flow-based future forecasting pipeline. (a) The detailed structures of future decoder, consisting of a future propagation module, a motion head “M”, and a refinement 2D convolutional backbone. (b) Grid points with predicted flows are transformed to time $t + k$ coordinates, followed by distance-based interpolation to obtain the future grid features.

3.2. Decoupled Latent World Modeling

Instead of jointly learning representation and temporal dynamics, we propose a decoupled strategy, where the world decoder is first trained to generate representative BEV features that reflect a comprehensive understanding of past observations. Building on these spatially-aware features, we then introduce a dynamic future decoder module to model temporal evolution in the latent BEV space, without being burdened by feature learning. Specifically, we employ an explicit flow-based future forecasting strategy to generate future world states in the latent BEV space, along with a dynamic-aware ray sampling strategy to improve learning efficiency in potentially dynamic regions.

Flow-based Future Forecasting. As illustrated in Figure 3, the future decoder predicts future states based on the latent spatial features B_t and action condition, *i.e.*, ego-motion. Specifically, the ego-motion to target time $t + k$ is first encoded into an embedding $a_{t+k} \in \mathbb{R}^C$ through a multi-layer perceptron (MLP), which is added to B_t to incorporate action information. Then, a future propagation module consisting of stacked deformable attention blocks processes these condition-augmented features and interacts with historical latent features in streaming memory, producing future features B'_{t+k} . The corresponding future grid positions g_f are computed by adding current grid positions $g \in \mathbb{R}^{H \times W \times 2}$ with flow fields $f_{t+k} \in \mathbb{R}^{H \times W \times 2}$, where f_{t+k} is predicted from B'_{t+k} via a motion head. Since g_f are expressed in the frame- t coordinate system, we transform them to the target frame $t + k$ as follows:

$$g'_f = T_t^{t+k} g_f, \quad (7)$$

where T_t^{t+k} is the ego-motion transformation matrix from

the current frame to frame $t + k$. Finally, distance-based interpolation is applied to warp the features into grid-structure future state \hat{B}_{t+k} . In particular, for each grid in \hat{B}_{t+k} , we select N nearest neighbor grids in B'_{t+k} using g'_f , and aggregate the grid features through inverse distance weighted summation. A 2D convolutional backbone is employed to further refine the interpolated features. During training, we supervise the predicted features \hat{B}_{t+k} using ground truth latent representations B_{t+k} extracted by the world encoder:

$$\mathcal{L}_{\text{latent}} = \|B_{t+k} - \hat{B}_{t+k}\|. \quad (8)$$

In this way, we can predict future states at arbitrary timestamps. Compared to autoregressive methods [63, 72], this direct prediction strategy not only avoids error accumulation but also offers greater flexibility in practice.

Dynamic-aware Ray Sampling. In addition to latent space supervision, we incorporate scene modeling loss $\mathcal{L}_{\text{scene}}$ at future timestamps. However, since the majority of rays belong to static background elements, salient moving objects may be overlooked during training. To bolster the model’s understanding of world dynamics, we present a dynamic-aware ray sampling strategy, which leverages an off-the-shelf tracker to identify regions of interest (RoIs) with high dynamic activity, preserving rays within these regions to prioritize motion-salient areas.

Learning Objective To summarize, the overall objective for the future decoder contains latent space alignment and future scene modeling:

$$\mathcal{L}_{\text{future}} = \sum_{k=1}^F (\omega_l \mathcal{L}_{\text{latent}}^k + \omega_s \mathcal{L}_{\text{scene}}^k), \quad (9)$$

where ω_l and ω_s are the weight coefficients, and F represents the number of future timestamps used for supervision. Notably, the supervision signal B_{t+k} in $\mathcal{L}_{\text{latent}}$ is generated from raw images using Eq.(1), while $\mathcal{L}_{\text{scene}}$ operates without manual annotations, ensuring a fully self-supervised training paradigm in future decoding.

3.3. World Model Integration with Task Heads

The fully-trained DriveX model can predict future world evolutions in the latent BEV space, offering valuable insights for a range of autonomous driving tasks. Notably, the generated generic and spatially-aware BEV features are readily compatible with various driving models. As shown in the bottom part of Figure 2, various downstream tasks can directly employ the BEV features from the frozen world model through a lightweight adapter, which enhances B_t with latent image features from world encoder using a stack of deformable attention layers. Then, we introduce a Future Spatial Attention (FSA) paradigm that enables future predictions to be incorporated into various autonomous driving models within a unified framework. To illustrate the integration of the DriveX model, we focus on two representative tasks: occupancy prediction and end-to-end driving.

Method	Modality	Chamfer Distance (m ²) ↓					
		0.5s	1.0s	1.5s	2.0s	2.5s	3.0s
4D-Occ* [26]	L	0.91	1.13	1.30	1.53	1.72	2.11
ViDAR [63]	C	1.01	1.12	1.25	1.38	1.54	1.73
HERMES [73]	C	-	0.78	-	0.95	-	1.17
DriveX-B	C	0.55	0.66	0.75	0.86	0.97	1.10

Table 2. Point cloud forecasting performance comparison on nuScenes [3] dataset. L and C denote LiDAR and Camera inputs, respectively. *: re-evaluated by [62].

Action Prediction. Because the world model’s future predictions are conditional on the ego vehicle’s future actions, we employ a straightforward strategy to predict these actions for models that don’t inherently include this functionality, such as perception models. Specifically, we follow AD-MLP [67] to predict actions based on ego states and historical trajectories, focusing on short-term predictions to simplify the task compared to full trajectory planning. For the end-to-end driving task, we directly adopt the planning head’s output as the predicted actions.

Future Spatial Attention Paradigm. To integrate the predicted future BEV features into downstream tasks, we propose a unified Future Spatial Attention (FSA) paradigm. As shown in Figure 2, this paradigm uses task-specific queries to directly aggregate world state information from the predicted future BEV features via standard spatial attention. Concretely, given a query q from the task model and its 3D reference point p , the spatial attention can be formulated as:

$$q := q + \sum_{k=1}^K \sum_{j=1}^J A_{kj} W \hat{B}_{t+k} [T_t^{t+k}(p + \Delta p_{kj})], \quad (10)$$

where K and J denote the number of future timesteps and sampling positions, respectively. W represents the learnable projection matrix, A_{kj} indicates the normalized attention weight, and Δp_{kj} is sample offset. Unlike conventional deformable attention, we transform the sample point $p + \Delta p_{kj}$ to the frame- $(t+k)$ coordinate system for feature sampling from \hat{B}_{t+k} . For simplicity, we omit the explicit expression of attention heads. The key to adapting spatial attention for various downstream tasks lies in the selection of Δp_{kj} , which enables the DriveX model to seamlessly integrate with diverse autonomous driving tasks while preserving the integrity of existing architectures.

4. Experiments

4.1. Experimental Settings

Dataset. We conduct experiments on two multimodal autonomous driving datasets: nuScenes [3] and NAVSIM [8]. The nuScenes dataset comprises 1,000 driving sequences and supports various autonomous driving tasks. For occupancy prediction and occupancy flow estimation experiments, we utilize the annotations from the FlowOcc3D

Method	Backbone	Flow Head	Flow Head			others	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	drive. suf.	other flat	sidewalk	terrain	manmade	vegetation
			mIoU	IoU _{geo}	mAVE ↓																	
BEVFormer[35]	InternT	✗	36.93	68.49	-	8.56	42.94	19.34	47.02	49.59	20.36	22.62	24.69	19.77	30.11	35.24	82.11	40.86	50.62	54.81	42.56	36.55
CVT-Occ [65]	Res101	✗	40.34	-	-	9.45	49.46	23.57	49.18	55.36	23.10	27.85	28.88	29.07	34.97	40.98	81.44	40.92	51.37	54.25	45.94	39.71
ViewFormer [30]	InternT	✗	43.61	72.46	-	13.82	50.32	29.49	49.24	54.52	24.34	32.72	31.09	31.49	34.44	41.62	85.47	51.27	59.03	62.15	48.33	42.06
FB-OCC [36]	InternT	✗	38.69	69.95	-	11.40	41.42	24.27	46.01	49.38	24.56	27.06	28.09	25.61	32.23	38.46	80.97	42.99	50.95	56.15	40.55	37.61
ViewFormer (ViDAR [63])†	InternT	✗	43.72	72.88	-	13.52	50.69	29.56	49.77	54.72	25.24	32.41	31.59	31.20	34.54	42.03	85.40	51.01	58.96	62.28	48.10	42.17
ViewFormer [30]	InternT	✗	43.61	72.46	-	13.82	50.32	29.49	49.24	54.52	24.34	32.72	31.09	31.49	34.44	41.62	85.47	51.27	59.03	62.15	48.33	42.06
ViewFormer (DriveX-S)	InternT	✗	44.38	73.61	-	14.00	51.18	30.10	50.13	55.44	26.82	33.12	32.37	32.07	35.60	43.18	85.78	50.93	59.16	62.23	49.80	42.59
<i>Improvement</i>			+0.77	+1.15																		
BEVFormer* [35]	InternT	✓	33.61	67.41	0.695	8.44	40.80	12.57	37.70	45.76	18.44	12.14	22.52	21.25	25.97	29.40	80.92	38.19	48.56	52.58	40.98	35.07
FB-OCC* [36]	InternT	✓	37.36	69.73	0.433	10.95	40.08	23.14	42.87	47.17	21.43	23.84	27.24	24.50	31.54	37.36	80.31	42.26	50.14	55.44	39.98	36.84
ViewFormer (ViDAR [63])†	InternT	✓	42.64	72.76	0.404	13.73	49.23	29.04	46.98	52.79	21.80	29.85	30.73	30.55	34.42	41.22	85.17	50.45	58.28	61.09	47.82	41.72
ViewFormer [30]	InternT	✓	42.54	72.36	0.412	13.63	49.35	28.74	46.63	52.71	21.04	29.63	30.34	30.53	34.01	41.04	85.23	50.63	58.68	61.63	47.72	41.56
ViewFormer (DriveX-S)	InternT	✓	43.47	73.44	0.385	13.72	49.46	30.64	48.85	54.99	27.13	31.90	32.62	31.29	34.29	41.16	84.89	49.91	58.00	60.75	47.50	41.88
<i>Improvement</i>			+0.93	+1.08	-0.027																	

Table 3. 3D occupancy and occupancy flow prediction performance comparison on FlowOcc3D [30]. We report results for models with and without a flow head. Performance gains achieved by our DriveX model are highlighted. *: reported by [30]. †: re-implemented by us.

Method	Modality	NC ↑	DAC ↑	TTC ↑	Comf. ↑	EP ↑	PDMS ↑
Constant Velocity [8]	Camera	68.0	57.8	50.0	100	19.4	20.6
Ego Status MLP [8]	Camera	93.0	77.3	83.6	100	62.8	65.6
Transfuser* [6]	Camera	97.4	92.8	92.4	100	79.0	83.8
PARA-Drive * [59]	Camera	97.9	92.4	93.0	99.8	79.3	84.0
UniAD* [22]	Camera	97.8	91.9	92.9	100	78.8	83.4
DiffusionDrive-C [38]†	Camera	97.8	92.2	92.6	99.9	78.9	83.6
DriveX-S	Camera	97.5	94.0	93.0	100	79.7	84.5
<i>Improvement</i>	-	-0.3	+1.8	+0.4	+0.1	+0.8	+0.9

Table 4. End-to-end planning performance on NAVSIM test split. “PDMS” indicates the predictive driver model score. *: reported by [8]. †: our re-implementation of the camera-based DiffusionDrive model. The best results are marked in **bold**.

benchmark [30]. The NAVSIM dataset, built upon OpenScene [7], provides 115k filtered samples for challenging end-to-end driving. We generate image semantic labels for both datasets using 2D foundation models [46, 69], resulting in 15 classes as defined in [68].

Network Architecture. The world encoder consists of an InternImage-Tiny [53] backbone with an FPN neck [39], followed by six encoder layers [35] to extract BEV-space world representations from multi-view images and historical latent states. The future decoder consists of a future propagation module with three deformable attention blocks [35], a motion head, and a 2D CNN backbone for refinement. We develop two variants of the world model: DriveX-S, which uses a smaller input image size of 256×704 , and DriveX-B, which employs a larger size of 900×1600 . For downstream task verifications, we adopt the ViewFormer [30] head as our baseline for occupancy and flow prediction task, and the DiffusionDrive [38] head for end-to-end driving task. The adapter and FSA are configured with 4 and 3 attention layers, respectively.

Training. The world model adopts a two-step training strategy, with the world encoder first trained for 40 epochs, followed by the future decoder trained for 24 epochs. Both stages utilize the AdamW optimizer [40] and are conducted on 32 NVIDIA H20 GPUs with a total batch size of 32. In

the future decoding training phase, we predict states at 0.5-second intervals over a 3-second horizon. The loss weights ω_l and ω_s in Eq.(9) are set to 1.0 and 0.5, respectively. For ray sampling, we first utilize DetZero [41] to identify moving objects and aggregate point clouds from the time range [-2s, +2.5s] to reconstruct the scene. We then select up to 8000 rays from LiDAR view through voxel-based down-sampling and 10000 rays from image view based on pixel semantics. Additionally, for dynamic-aware ray sampling, we sample up to 2,000 rays within the bounding boxes of moving objects. For downstream tasks, we follow the official training protocols while configuring the world model to predict 2 future frames for 3D occupancy and flow prediction, and 3 frames for end-to-end driving.

4.2. Main Results

In this section, we first evaluate our DriveX model in point cloud forecasting task to demonstrate its superior capability in future forecasting. We then benchmark its effectiveness in improving the performance of occupancy prediction, occupancy flow estimation, as well as end-to-end driving.

Point Cloud Forecasting. Following [62, 73], we evaluate on points within the range of [-51.2m, 51.2m] on both X and Y axes. As presented in Table 2, our DriveX model outperforms the previous state-of-the-art method HERMES [73] across all timestamps, achieving a reduction of 0.12m^2 and 0.09m^2 in Chamfer Distance at 1.0s and 2.0s horizons.

Occupancy and Flow Prediction. We verify our DriveX model on the 3D occupancy and flow prediction tasks, employing the state-of-the-art ViewFormer [30] as baseline. As shown in Table 3, for occupancy prediction, the DriveX model boosts ViewFormer’s performance by +0.77 mIoU and +1.15 IoU_{geo}, surpassing all previous methods. Notably, while conventional pretrained-based world models like ViDAR [62] show marginal improvements on such a strong baseline, DriveX still delivers substantial gains, which we attribute to its ability to learn a more compre-

hensive understanding of the world. To further investigate its effectiveness in modeling world dynamics, we add a flow head for occupancy motion estimation. The model achieves an mAVE of 0.385, reducing the error by 6.5%. The performance gains also remain remarkable for both mIoU (+0.93) and IoU_{geo} (+1.08). These improvements highlight the capability of the DriveX model in enhancing downstream models via future scene prediction.

End-to-end Driving. Safely planning for autonomous driving requires accurate anticipation of future consequences. By leveraging a world model, we can comprehensively evaluate possible futures to enhance performance. As shown in Table 4, our DriveX model, based on the re-implemented DiffusionDrive-C [38], boosts the overall PDMS to 84.5, achieving state-of-the-art performance among camera-only end-to-end driving models.

4.3. Ablation Study

In this section, we conduct a series of ablation studies on nuScenes [3] dataset to investigate the individual components of our DriveX framework. By default, we use an input image resolution of 256×704 for efficient training. All models are trained with a batch size of 32 on 32 GPUs.

Omni scene modeling. We first ablate the effects of our omni scene modeling module through occupancy and flow prediction task. As presented in Table 5, DriveX captures only geometric information when supervised solely by depth estimation. Adding semantic supervision enriches the model’s semantic understanding, achieving a mIoU of 42.53. However, since the semantic labels from auto-labeling inevitably suffer from information loss, incorporating color reconstruction preserves additional generalizable features, ultimately improving the mIoU to 43.47.

Effect of prediction horizon. Table 6 verifies that the downstream performance gain stems from DriveX’s predictive capability. As indicated in the 2nd row, DriveX achieves significant improvements of +0.9 in mIoU and -0.03 in mAVE compared to the baseline without future prediction. While extending the prediction horizon to 4 frames brings further gains, we adopt 2 as the default setting to balance performance and computational cost.

Decoupled training strategy. We ablate the effects of our decouple training strategy in Table 7. For a fair comparison, we train a joint training variant of our DriveX model for 64 epochs, matching the total training duration of the decoupled version. As shown in 1st and 2nd rows, our decoupled approach significantly improves prediction accuracy, reducing the Chamfer Distance by 0.44m^2 for 3.0s predictions. This validates that separating the world model into world encoding and latent future predicting leads to a better modeling of world evolution. Moreover, the integration of dynamic-aware ray sampling strategy enables the model to focus on motion-critical regions, further re-

Sup-D	Sup-S	Sup-C	mIoU	IoU_{geo}	mAVE ↓
✓			3.96	59.7	1.388
✓	✓		42.53	73.1	0.396
✓	✓	✓	43.47	73.44	0.385

Table 5. Ablation study on components of omni scene modeling module. “Sup-D”, “Sup-S” and “Sup-C” denote supervision from depth estimation, semantic segmentation, and color reconstruction, respectively.

Future frames	mIoU	IoU_{geo}	mAVE ↓
0	42.57	72.60	0.415
2	43.47	73.44	0.385
4	43.59	73.42	0.387

Table 6. Effects of prediction horizon in FSA .

Joint	Decoupled	D-sample	Chamfer Distance (m^2) ↓			
			1.0s	2.0s	3.0s	Avg.
✓			1.09	1.41	1.75	1.41
	✓		0.83	1.03	1.31	1.06
	✓	✓	0.80	0.99	1.28	1.02

Table 7. Ablation study on decoupled training strategy. “D-Sample” refers dynamic-aware ray sampling.

Strategy	Chamfer Distance (m^2) ↓			
	1.0s	2.0s	3.0s	Avg.
Autoregressive-pred	0.99	1.21	1.76	1.32
Direct-pred	0.80	0.99	1.28	1.02

Table 8. Comparison of autoregressive and direct prediction strategies for future state prediction.

ducing the average Chamfer Distance to 1.02m^2 . Additionally, we compare direct prediction with autoregressive strategies for future decoding in Table 8, where direct prediction demonstrates superior performance, diminishing the average Chamfer Distance from 1.32m^2 to 1.02m^2 .

5. Conclusion

In this paper, we present DriveX, a self-supervised world model that learns generalizable scene dynamics and holistic representations from driving videos. DriveX introduces Omni Scene Modeling, unifying multimodal supervision for comprehensive scene understanding, and employs a decoupled latent modeling strategy to simplify the learning of complex dynamics. Additionally, we develop the Future Spatial Attention paradigm to enhance downstream task adaptability. Extensive experiments on the nuScenes and NAVSIM datasets demonstrate the effectiveness of our approach, paving the way for developing more practical and scalable world models in autonomous driving systems.

Acknowledgment

This work is supported by the Guangdong Basic and Applied Basic Research Foundation (Grant No. 2025A1515011434) and CCF-DiDi GAIA Collaborative Research Funds.

References

- [1] Ben Agro, Quinlan Sykora, Sergio Casas, Thomas Gilles, and Raquel Urtasun. Uno: Unsupervised occupancy fields for perception and forecasting. In *CVPR*, 2024. 1
- [2] Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos J Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. *NeurIPS*, 2024. 3
- [3] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 6, 8
- [4] Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. *arXiv preprint arXiv:2411.00769*, 2024. 3
- [5] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *CVPR*, 2024. 3
- [6] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *TPAMI*, 2022. 7
- [7] OpenScene Contributors. Openscene: The largest up-to-date 3d occupancy prediction benchmark in autonomous driving. <https://github.com/OpenDriveLab/OpenScene>, 2023. 7
- [8] Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, Andreas Geiger, and Kashyap Chitta. Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking. In *NeurIPS*, 2024. 2, 6, 7
- [9] Ruiyuan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, and Qiang Xu. Magicdrive: Street view generation with diverse 3d geometry control. *arXiv preprint arXiv:2310.02601*, 2023. 3
- [10] Ruiyuan Gao, Kai Chen, Zhihao Li, Lanqing Hong, Zhenguo Li, and Qiang Xu. Magicdrive3d: Controllable 3d generation for any-view rendering in street scenes. *arXiv preprint arXiv:2405.14475*, 2024.
- [11] 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. *NeurIPS*, 2024. 1, 2, 3
- [12] Mingzhe Guo, Zhipeng Zhang, Yuan He, Ke Wang, and Liping Jing. End-to-end autonomous driving without costly modularization and 3d manual annotation. *arXiv preprint arXiv:2406.17680*, 2024. 1
- [13] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *NeurIPS*, 2018. 3
- [14] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *ICLR*, 2020.
- [15] Danijar Hafner, Timothy P Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *ICLR*, 2021.
- [16] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023. 3
- [17] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. *NeurIPS*, 2022. 2, 3
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33, 2020. 3
- [20] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *ICLR*, 2023. 3
- [21] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023. 2, 3
- [22] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhui Wang, et al. Planning-oriented autonomous driving. In *CVPR*, 2023. 1, 7
- [23] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *ICCV*, 2023. 1
- [24] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023. 2, 3
- [25] Tarasha Khurana, Peiyun Hu, Achal Dave, Jason Ziglar, David Held, and Deva Ramanan. Differentiable raycasting for self-supervised occupancy forecasting. In *ECCV*, 2022. 3
- [26] Tarasha Khurana, Peiyun Hu, David Held, and Deva Ramanan. Point cloud forecasting as a proxy for 4d occupancy forecasting. In *CVPR*, 2023. 2, 3, 6
- [27] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019. 1
- [28] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 2022. 3
- [29] Jiachen Li, Weixi Feng, Tsu-Jui Fu, Xinyi Wang, Sugato Basu, Wenhui Chen, and William Yang Wang. T2v-turbo: Breaking the quality bottleneck of video consistency model with mixed reward feedback. *NeurIPS*, 2024. 3

- [30] Jinke Li, Xiao He, Chonghua Zhou, Xiaoqiang Cheng, Yang Wen, and Dan Zhang. Viewformer: Exploring spatiotemporal modeling for multi-view 3d occupancy perception via view-guided transformers. In *ECCV*, 2024. 7
- [31] Peidong Li and Dixiao Cui. Navigation-guided sparse scene representation for end-to-end autonomous driving. In *ICLR*, 2025. 1, 3
- [32] Xiaofan Li, Yifu Zhang, and Xiaoqing Ye. Drivingdiffusion: Layout-guided multi-view driving scenarios video generation with latent diffusion model. In *ECCV*, 2024. 1, 2, 3
- [33] Xiang Li, Pengfei Li, Yupeng Zheng, Wei Sun, Yan Wang, and yilun chen. Semi-supervised vision-centric 3d occupancy world model for autonomous driving. In *ICLR*, 2025. 3
- [34] Yingyan Li, Lue Fan, Jiawei He, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang, and Tieniu Tan. Enhancing end-to-end autonomous driving with latent world model. In *ICLR*, 2025. 1, 3
- [35] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, 2022. 4, 7
- [36] Zhiqi Li, Zhiding Yu, Wenhai Wang, Anima Anandkumar, Tong Lu, and Jose M Alvarez. Fb-bev: Bev representation from forward-backward view transformations. In *ICCV*, 2023. 7
- [37] Zhiqi Li, Zhiding Yu, Shiyi Lan, Jiahua Li, Jan Kautz, Tong Lu, and Jose M Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? In *CVPR*, 2024. 1
- [38] Bencheng Liao, Shaoyu Chen, Haoran Yin, Bo Jiang, Cheng Wang, Sixu Yan, Xinbang Zhang, Xiangyu Li, Ying Zhang, Qian Zhang, et al. Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving. In *CVPR*, 2025. 7, 8
- [39] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 4, 7
- [40] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 7
- [41] Tao Ma, Xuemeng Yang, Hongbin Zhou, Xin Li, Botian Shi, Junjie Liu, Yuchen Yang, Zhizheng Liu, Liang He, Yu Qiao, et al. Detzero: Rethinking offboard 3d object detection with long-term sequential point clouds. In *ICCV*, 2023. 7
- [42] Benedikt Mersch, Xieyuanli Chen, Jens Behley, and Cyrill Stachniss. Self-supervised point cloud prediction using 3d spatio-temporal convolutional networks. In *CoRL*, 2022. 3
- [43] Chen Min, Dawei Zhao, Liang Xiao, Jian Zhao, Xinli Xu, Zheng Zhu, Lei Jin, Jianshu Li, Yulan Guo, Junliang Xing, Liping Jing, Yiming Nie, and Bin Dai. Driveworld: 4d pre-trained scene understanding via world models for autonomous driving. In *CVPR*, 2024. 2, 3
- [44] Chen Min, Dawei Zhao, Liang Xiao, Jian Zhao, Xinli Xu, Zheng Zhu, Lei Jin, Jianshu Li, Yulan Guo, Junliang Xing, et al. Driveworld: 4d pre-trained scene understanding via world models for autonomous driving. In *CVPR*, 2024. 1, 3
- [45] Mingjie Pan, Jiaming Liu, Renrui Zhang, Peixiang Huang, Xiaoqi Li, Hongwei Xie, Bing Wang, Li Liu, and Shanghang Zhang. Renderocc: Vision-centric 3d occupancy prediction with 2d rendering supervision. In *ICRA*, 2024. 3, 5
- [46] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 2, 3, 5, 7
- [47] Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak Pathak. Planning to explore via self-supervised world models. In *ICML*, 2020. 3
- [48] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pvrnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. *IJCV*, 2023. 1
- [49] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 3
- [50] Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. *arXiv preprint arXiv:2408.14837*, 2024. 3
- [51] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xianguyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *ICCV*, 2023. 1
- [52] Shihao Wang, Zhiding Yu, Xiaohui Jiang, Shiyi Lan, Min Shi, Nadine Chang, Jan Kautz, Ying Li, and Jose M Alvarez. Omnidrive: A holistic llm-agent framework for autonomous driving with 3d perception, reasoning and planning. *arXiv preprint arXiv:2405.01533*, 2024. 1
- [53] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *CVPR*, 2023. 4, 7
- [54] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *NeurIPS*, 2024. 2, 3
- [55] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiayang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. In *ECCV*, 2024. 1, 2, 3
- [56] 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
- [57] Xinshuo Weng, Jianren Wang, Sergey Levine, Kris Kitani, and Nicholas Rhinehart. Inverting the pose forecasting pipeline with spf2: Sequential pointcloud forecasting for sequential pose forecasting. In *CoRL*, 2021. 3
- [58] Xinshuo Weng, Junyu Nan, Kuan-Hui Lee, Rowan McAllister, Adrien Gaidon, Nicholas Rhinehart, and Kris M Kitani. S2net: Stochastic sequential pointcloud forecasting. In *ECCV*, 2022. 3
- [59] Xinshuo Weng, Boris Ivanovic, Yan Wang, Yue Wang, and Marco Pavone. Para-drive: Parallelized architecture for real-time autonomous driving. In *CVPR*, 2024. 7

- [60] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Daydreamer: World models for physical robot learning. In *CoRL*, 2023. 3
- [61] Hao Yang, Haiyang Wang, Di Dai, and Liwei Wang. Pred: pre-training via semantic rendering on lidar point clouds. *NeurIPS*, 2024. 5
- [62] Yu Yang, Jianbiao Mei, Yukai Ma, Siliang Du, Wenqing Chen, Yijie Qian, Yuxiang Feng, and Yong Liu. Driving in the occupancy world: Vision-centric 4d occupancy forecasting and planning via world models for autonomous driving. *arXiv preprint arXiv:2408.14197*, 2024. 1, 2, 3, 6, 7
- [63] Zetong Yang, Li Chen, Yanan Sun, and Hongyang Li. Visual point cloud forecasting enables scalable autonomous driving. In *CVPR*, 2024. 1, 2, 3, 6, 7
- [64] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. In *ICLR*, 2025. 3
- [65] Zhangchen Ye, Tao Jiang, Chenfeng Xu, Yiming Li, and Hang Zhao. Cvt-occ: Cost volume temporal fusion for 3d occupancy prediction. In *ECCV*, 2024. 7
- [66] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *CVPR*, 2021. 1
- [67] Jiang-Tian Zhai, Ze Feng, Jinhao Du, Yongqiang Mao, Jiang-Jiang Liu, Zichang Tan, Yifu Zhang, Xiaoqing Ye, and Jingdong Wang. Rethinking the open-loop evaluation of end-to-end autonomous driving in nusenes. *arXiv preprint arXiv:2305.10430*, 2023. 6
- [68] Chubin Zhang, Juncheng Yan, Yi Wei, Jiaxin Li, Li Liu, Yansong Tang, Yueqi Duan, and Jiwen Lu. Occnerf: Advancing 3d occupancy prediction in lidar-free environments. *arXiv preprint arXiv:2312.09243*, 2023. 7
- [69] Hao Zhang, Feng Li, Xueyan Zou, Shilong Liu, Chunyuan Li, Jianwei Yang, and Lei Zhang. A simple framework for open-vocabulary segmentation and detection. In *ICCV*, 2023. 2, 3, 5, 7
- [70] Lunjun Zhang, Yuwen Xiong, Ze Yang, Sergio Casas, Rui Hu, and Raquel Urtasun. Copilot4d: Learning unsupervised world models for autonomous driving via discrete diffusion. In *ICLR*, 2024. 2, 3
- [71] Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. In *ICCV*, 2023. 1
- [72] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. In *ECCV*, 2024. 1, 2, 3, 6
- [73] 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. 6, 7
- [74] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *NeurIPS*, 2020. 4