

VAD: Vectorized Scene Representation for Efficient Autonomous Driving

Bo Jiang^{1,*}, Shaoyu Chen^{1,*}, Qing Xu², Bencheng Liao^{1,◇}, Jiajie Chen²,
Helong Zhou², Qian Zhang², Wenyu Liu¹, Chang Huang², Xinggang Wang^{1,⊠}

¹ Huazhong University of Science & Technology ² Horizon Robotics

{bjiang, shaoyuchen, bcliao, liuwu, xgwang}@hust.edu.cn

{qing.xu, jiajie.chen, helong.zhou, qian01.zhang, chang.huang}@horizon.ai

Abstract

Autonomous driving requires a comprehensive understanding of the surrounding environment for reliable trajectory planning. Previous works rely on dense rasterized scene representation (e.g., agent occupancy and semantic map) to perform planning, which is computationally intensive and misses the instance-level structure information. In this paper, we propose VAD, an end-to-end vectorized paradigm for autonomous driving, which models the driving scene as a fully vectorized representation. The proposed vectorized paradigm has two significant advantages. On one hand, VAD exploits the vectorized agent motion and map elements as explicit instance-level planning constraints which effectively improves planning safety. On the other hand, VAD runs much faster than previous end-to-end planning methods by getting rid of computation-intensive rasterized representation and hand-designed post-processing steps. VAD achieves state-of-the-art end-to-end planning performance on the nuScenes dataset, outperforming the previous best method by a large margin. Our base model, VAD-Base, greatly reduces the average collision rate by 29.0% and runs 2.5× faster. Besides, a lightweight variant, VAD-Tiny, greatly improves the inference speed (up to 9.3×) while achieving comparable planning performance. We believe the excellent performance and the high efficiency of VAD are critical for the real-world deployment of an autonomous driving system. Code and models are available at <https://github.com/hustvl/VAD> for facilitating future research.

1. Introduction

Autonomous driving requires both comprehensive scene understanding for ensuring safety and high efficiency for

* Equal contribution; ◇ This work was done during the internship of Bo Jiang, Shaoyu Chen, and Bencheng Liao at Horizon Robotics; ⊠ Corresponding author: xgwang@hust.edu.cn

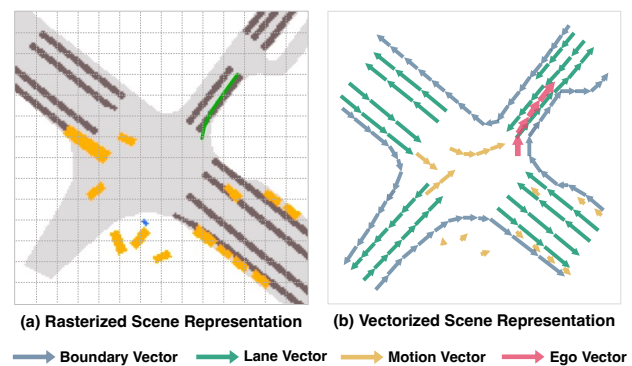


Figure 1. Previous paradigms mainly rely on rasterized representation (a) for planning (e.g., semantic map, occupancy map, flow map, and cost map), which is computationally intensive. The proposed VAD is fully based on the vectorized scene representation (b) for end-to-end planning. VAD leverages instance-level structure information as planning constraints and guidance, achieving promising performance and efficiency.

real-world deployment. An autonomous vehicle needs to efficiently perceive the driving scene and perform reasonable planning based on the scene information.

Traditional autonomous driving methods [23, 14, 7, 48] adopt a modular paradigm, where perception and planning are decoupled into standalone modules. The disadvantage is, the planning module cannot access the original sensor data, which contains rich semantic information. And since planning is fully based on preceding perception results, the error in perception may severely influence planning and can not be recognized and cured in the planning stage, which leads to the safety problem.

Recently, end-to-end autonomous driving methods [19, 21, 2, 11] take sensor data as input for perception and output planning results with one holistic model. Some works [40, 9, 41] directly output planning results based on the sensor data without learning scene representation, which lacks interpretability and is difficult to optimize. Most works [19, 21, 2] transform the sensor data into rasterized

scene representation (e.g., semantic map, occupancy map, flow map, and cost map) for planning. Though straightforward, rasterized representation is computationally intensive and misses critical instance-level structure information.

In this work, we propose VAD (Vectorized Autonomous Driving), an end-to-end vectorized paradigm for autonomous driving. VAD models the scene in a fully vectorized way (i.e., vectorized agent motion and map), getting rid of computationally intensive rasterized representation.

We argue that vectorized scene representation is superior to rasterized one. Vectorized map (represented as boundary vectors and lane vectors) provides road structure information (e.g., traffic flow, drivable boundary, and lane direction), and helps the autonomous vehicle narrow down the trajectory search space and plan a reasonable future trajectory. The motion of traffic participants (represented as agent motion vectors) provides instance-level restriction for collision avoidance. What’s more, vectorized scene representation is efficient in terms of computation, which is important for real-world applications.

VAD takes full advantage of the vectorized information to guide planning both implicitly and explicitly. On one hand, VAD adopts map queries and agent queries to implicitly learn instance-level map features and agent motion features from sensor data, and extracts guidance information for planning via query interaction. On the other hand, VAD proposes three instance-level planning constraints based on the explicit vectorized scene representation: the ego-agent collision constraint for maintaining a safe distance between the ego vehicle and other dynamic agents both laterally and longitudinally; the ego-boundary overstepping constraint for pushing the planning trajectory away from the road boundary; and the ego-lane direction constraint for regularizing the future motion direction of the autonomous vehicle with vectorized lane direction. Our proposed framework and the vectorized planning constraints effectively improve the planning performance, without incurring large computational overhead.

Without fancy tricks or hand-designed post-processing steps, VAD achieves state-of-the-art (SOTA) end-to-end planning performance and the best efficiency on the challenging nuScenes [1] dataset. Compared with the previous SOTA method UniAD [21], our base model, VAD-Base, greatly reduces the average planning displacement error by 30.1% (1.03m v.s. 0.72m) and the average collision rate by 29.0% (0.31% v.s. 0.22%), while running 2.5× faster (1.8 FPS v.s. 4.5 FPS). The lightweight variant, VAD-Tiny, runs 9.3× faster (1.8 FPS v.s. 16.8 FPS) while achieving comparable planning performance, the average planning displacement error is 0.78m and the average collision rate is 0.38%. We also demonstrate the effectiveness of our design choices through thorough ablations.

Our key contributions are summarized as follows:

- We propose VAD, an end-to-end vectorized paradigm for autonomous driving. VAD models the driving scene as a fully vectorized representation, getting rid of computationally intensive dense rasterized representation and hand-designed post-processing steps.
- VAD implicitly and explicitly utilizes the vectorized scene information to improve planning safety, via query interaction and vectorized planning constraints.
- VAD achieves SOTA end-to-end planning performance, outperforming previous methods by a large margin. Not only that, because of the vectorized scene representation and our concise model design, VAD greatly improves the inference speed, which is critical for the real-world deployment of an autonomous driving system.

It’s our belief that autonomous driving can be performed in a fully vectorized manner with high efficiency. We hope the impressive performance of VAD can reveal the potential of vectorized paradigm to the community.

2. Related Work

Perception. Accurate perception of the driving scene is the basis for autonomous driving. We mainly introduce some camera-based 3D object detection and online mapping methods which are most relevant to this paper. DETR3D [47] uses 3D queries to sample corresponding image features and accomplish detection without non-maximum suppression. PETR [31] introduces 3D positional encoding to image features and uses detection queries to learn object features via attention [46] mechanism. Recently, bird’s-eye view (BEV) representation has become popular and has greatly contributed to the field of perception [26, 51, 17, 29, 8, 34, 28]. LSS [39] is a pioneering work that introduces depth prediction to project features from perspective view to BEV. BEVFormer [26] proposes spatial and temporal attention for better encoding the BEV feature map and achieves remarkable detection performance with camera input only. FIERY [17] and BEVerse [51] use the BEV feature map to predict dense map segmentation. HDMapNet [25] converts lane segmentation to vectorized map with post-processing steps. VectorMapNet [32] predicts map elements in an autoregressive way. MapTR [29] recognizes the permutation invariance of the map instance points and can predict all map elements simultaneously. LaneGAP [28] models the lane graph in a novel path-wise manner, which well preserves the continuity of the lane and encodes traffic information for planning. We leverage a group of BEV queries, agent queries, and map queries to accomplish scene perception following BEVFormer [26] and MapTR [29], and further use these

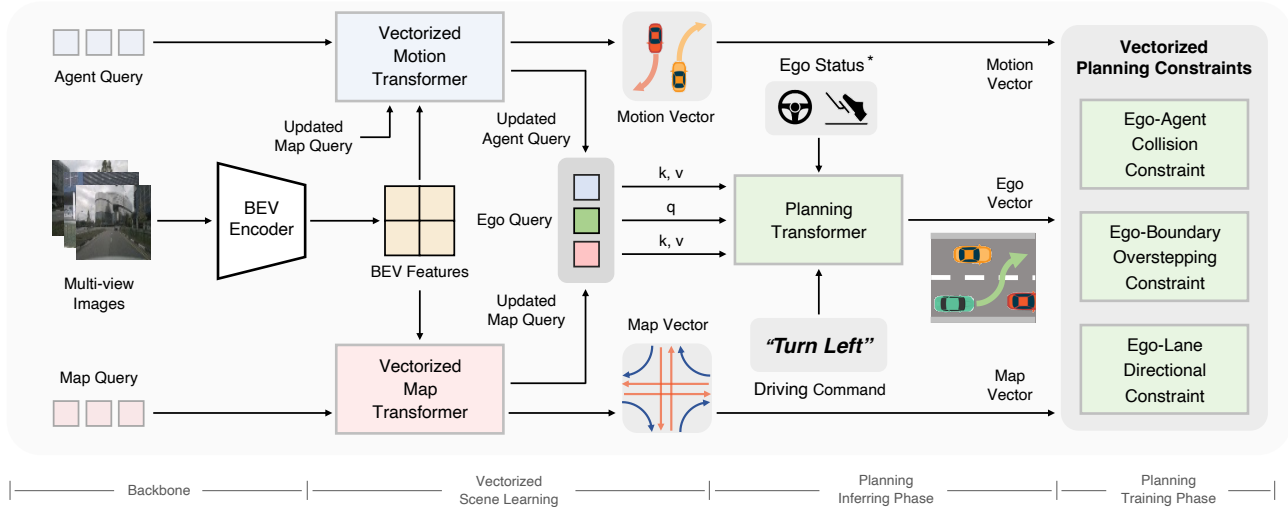


Figure 2. **Overall architecture of VAD.** The full pipeline of VAD is divided into four phases. Backbone includes an image feature extractor and a BEV encoder to project the image features to the BEV features. Vectorized Scene Learning aims to encode the scene information into agent queries and map queries, as well as represent the scene with motion vectors and map vectors. In the inferring phase of planning, VAD utilizes an ego query to extract map and agent information via query interaction and outputs the planning trajectory (represented as ego vector). The proposed vectorized planning constraints regularize the planning trajectory in the training phase. *: optional.

query features and perception results in the motion prediction and planning stage. Details are shown in Sec. 3.

Motion Prediction. Traditional motion prediction takes perception ground truth (*e.g.*, agent history trajectories and HD map) as input. Some works [3, 38] render the driving scene as BEV images and adopt CNN-based networks to predict future motion. Some other works [13, 33, 37] use vectorized representation, and adopt GNN [27] or Transformer [46, 33, 37] to accomplish learning and prediction. Recent end-to-end works [17, 51, 15, 22] jointly perform perception and motion prediction. Some works [17, 51, 20] see future motion as dense occupancy and flow instead of agent-level future waypoints. ViP3D [15] predicts future motion based on the tracking results and HD map. PIP [22] proposes an interaction scheme between dynamic agents and static vectorized map, and achieves SOTA performance without relying on HD map. VAD learns vectorized agent motion by interacting between dynamic agents and static map elements, inspired by [22].

Planning. Recently, learning-based planning methods prevail. Some works [40, 9, 41] omit intermediate stages such as perception and motion prediction, and directly predict planning trajectories or control signals. Although this idea is straightforward and simple, they lack interpretability and are difficult to optimize. Reinforcement learning is quite up to the planning task and has become a promising research direction [45, 5, 4]. Explicit dense cost map has great interpretability and is widely used [2, 19, 11, 44]. The

cost maps are constructed from the perception or motion prediction results, or come from a learning-based module. And hand-crafted rules are often adopted to select the best planning trajectory with minimum cost. The construction of a dense cost map is computationally intensive and the using of hand-crafted rules brings robustness and generalization problems. UniAD [21] effectively incorporates the information provided by various preceding tasks to assist planning in a goal-oriented spirit, and achieves remarkable performance in perception, prediction, and planning. PlanT [43] takes perception ground truth as input and encodes the scene in object-level representation for planning. In this paper, we explore the potential of vectorized scene representation for planning and get rid of dense maps or hand-designed post-processing steps.

3. Method

Overview. The overall framework of VAD is depicted in Fig. 2. Given multi-frame and multi-view image input, VAD first encodes the image features with a backbone network and utilizes a group of BEV queries to project the image features to the BEV features [26, 39, 52]. Second, VAD utilizes a group of agent queries and map queries to learn the vectorized scene representation, including vectorized map and vectorized agent motion (Sec. 3.1). Third, planning is performed based on the scene information (Sec. 3.2). Specifically, VAD uses an ego query to learn the implicit scene information through interaction with agent queries and map queries. Based on the ego query, ego status fea-

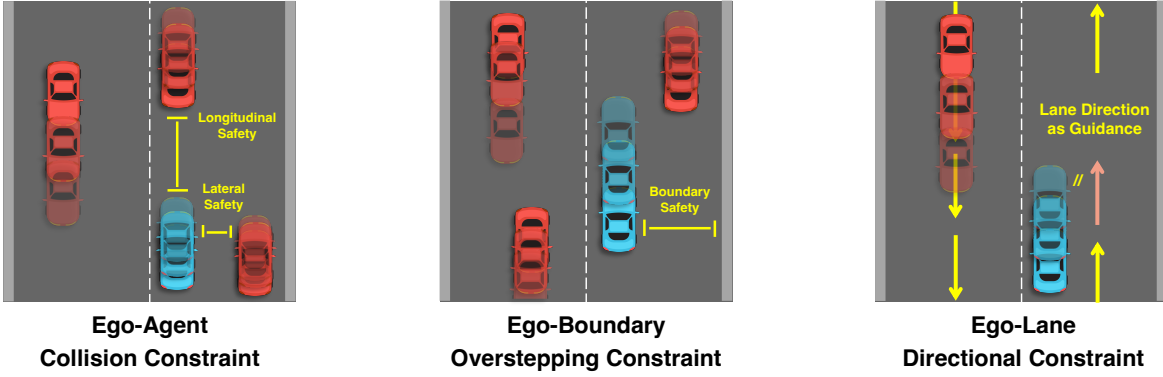


Figure 3. **Illustration of vectorized planning constraints.** Ego-agent collision constraint aims to keep longitudinal safety and lateral safety between the ego vehicle and other agents. Ego-boundary overstepping constraint punishes the predictions when the planning trajectory gets too close to the lane boundary. Ego-lane directional constraint leverages the direction of the closet lane vector from the ego car (the pink lane in the right sub-figure) as prior to regularize the motion direction of planning.

tures, and high-level driving command, the Planning Head outputs the planning trajectory. Besides, VAD introduces three vectorized planning constraints to restrict the planning trajectory at the instance level (Sec. 3.3). VAD is fully differentiable and trained in an end-to-end manner (Sec. 3.4).

3.1. Vectorized Scene Learning

Perceiving traffic agents and map elements are important in driving scene understanding. VAD encodes the scene information into query features and represents the scene by map vectors and agent motion vectors.

Vectorized Map. Previous works [19, 21] use rasterized semantic maps to guide the planning, which misses critical instance-level structure information of the map. VAD utilizes a group of map queries [29] Q_m to extract map information from BEV feature map and predicts map vectors $\hat{V}_m \in \mathbb{R}^{N_m \times N_p \times 2}$ and the class score for each map vector, where N_m and N_p denote the number of predicted map vectors and the points contained in each map vector. Three kinds of map elements are considered: lane divider, road boundary, and pedestrian crossing. The Lane divider provides direction information, and the road boundary indicates the drivable area. Map queries and map vectors are both leveraged to improve the planning performance (Sec. 3.2 and Sec. 3.3).

Vectorized Agent Motion. VAD first adopts a group of agent queries Q_a to learn agent-level features from the shared BEV feature map via deformable attention [53]. The agent’s attributes (location, class score, orientation, *etc.*) are decoded from the agent queries by an MLP-based decoder head. To enrich the agent features for motion prediction, VAD performs agent-agent and agent-map inter-

action [22, 37] via attention mechanism. Then VAD predicts future trajectories of each agent, represented as multi-modality motion vectors $\hat{V}_a \in \mathbb{R}^{N_a \times N_k \times T_f \times 2}$. N_a, N_k , and T_f denote the number of predicted agents, the number of modalities, and the number of future timestamps. Each modality of the motion vector indicates a kind of driving intention. VAD outputs a probability score for each modality. The agent motion vectors are used to restrict the ego planning trajectory and avoid collision (Sec. 3.3). Meanwhile, the agent queries are sent into the planning module as scene information (Sec. 3.2).

3.2. Planning via Interaction

Ego-Agent Interaction. VAD utilizes a randomly initialized ego query Q_{ego} to learn the implicit scene features which are valuable for planning. In order to learn the location and motion information of other dynamic traffic participants, the ego query first interacts with the agent queries through a Transformer decoder [46], in which ego query serves as query of attention q , and agent queries serve as key k and value v . The ego position p_{ego} and agent positions p_a predicted by the perception module are encoded by a single layer MLP PE_1 , and serve as query position embedding q_{pos} and key position embedding k_{pos} . The positional embeddings provide information on the relative position relationship between agents and the ego vehicle. The above process can be formulated as:

$$\begin{aligned}
 Q'_{ego} &= \text{TransformerDecoder}(q, k, v, q_{pos}, k_{pos}), \\
 q &= Q_{ego}, k = v = Q_a, \\
 q_{pos} &= PE_1(p_{ego}), k_{pos} = PE_1(p_a).
 \end{aligned} \tag{1}$$

Ego-Map Interaction. After interacting with agent queries, the updated ego query Q'_{ego} further interacts with the map queries Q_m in a similar way. The only difference is we use a different MLP PE_2 to encode the positions of the ego vehicle and the map elements. The output ego query Q''_{ego} contains both dynamic and static information of the driving scene. The process is formulated as:

$$\begin{aligned} Q''_{\text{ego}} &= \text{TransformerDecoder}(q, k, v, q_{\text{pos}}, k_{\text{pos}}), \\ q &= Q'_{\text{ego}}, k = v = Q_m, \\ q_{\text{pos}} &= \text{PE}_2(p_{\text{ego}}), k_{\text{pos}} = \text{PE}_2(p_m). \end{aligned} \quad (2)$$

Planning Head. Because VAD performs HD-map-free planning, a high-level driving command c is required for navigation. Following the common practice [19, 21], VAD uses three kinds of driving commands: *turn left*, *turn right* and *go straight*. So the planning head takes the updated ego queries ($Q'_{\text{ego}}, Q''_{\text{ego}}$) and the current status of the ego vehicle s_{ego} (optional) as ego features f_{ego} , as well as the driving command c as inputs, and outputs the planning trajectory $\hat{V}_{\text{ego}} \in \mathbb{R}^{T_f \times 2}$. VAD adopts a simple MLP-based planning head. The decoding process is formulated as follows:

$$\begin{aligned} \hat{V}_{\text{ego}} &= \text{PlanHead}(\text{ft} = f_{\text{ego}}, \text{cmd} = c), \\ f_{\text{ego}} &= [Q'_{\text{ego}}, Q''_{\text{ego}}, s_{\text{ego}}]. \end{aligned} \quad (3)$$

where [...] denotes concatenation operation, ft denotes features used for decoding, and cmd denotes the navigation driving command.

3.3. Vectorized Planning Constraint

Based on the learned map vector and motion vector, VAD regularizes the planning trajectory \hat{V}_{ego} with instance-level vectorized constraints during the training phase, as shown in Fig. 3.

Ego-Agent Collision Constraint. Ego-agent collision constraint explicitly considers the compatibility of the ego planning trajectory and the future trajectory of other vehicles, in order to improve planning safety and avoid collision. Unlike previous works [19, 21] that adopt dense occupancy maps, we utilize vectorized motion trajectories which both keep great interpretability and require less computation. Specifically, we first filter out low-confidence agent predictions by a threshold ϵ_a . For multi-modality motion prediction, we use the trajectory with the highest confidence score as the final prediction. We consider collision constraint as a safe boundary for the ego vehicle both laterally and longitudinally. Multiple cars may be close to each other (e.g., driving side by side) in the lateral direction, but

a longer safety distance is required in the longitudinal direction. So we adopt different agent distance thresholds δ_X and δ_Y for different directions. For each future timestamp, we find the closest agent within a certain range δ_a in both directions. Then for each direction $i \in \{X, Y\}$, if the distance d_a^i with the closet agent is less than the threshold δ_i , then the loss item of this constraint $\mathcal{L}_{\text{col}}^i = \delta_i - d_a^i$, otherwise it is 0. The loss for ego-agent collision constraint can be formulated as:

$$\begin{aligned} \mathcal{L}_{\text{col}} &= \frac{1}{T_f} \sum_{t=1}^{T_f} \sum_i \mathcal{L}_{\text{col}}^{it}, i \in \{X, Y\}, \\ \mathcal{L}_{\text{col}}^{it} &= \begin{cases} \delta_i - d_a^{it}, & \text{if } d_a^{it} < \delta_i \\ 0, & \text{if } d_a^{it} \geq \delta_i. \end{cases} \end{aligned} \quad (4)$$

Ego-Boundary Overstepping Constraint. This constraint aims to push the planning trajectory away from the road boundary so that the trajectory can be kept in the drivable area. We first filter out low-confidence map predictions according to a threshold ϵ_m . Then for each future timestamp, we calculate the distance d_{bd}^t between the planning waypoint and its closest map boundary line. Then the loss for this constraint is formulated as:

$$\begin{aligned} \mathcal{L}_{\text{bd}} &= \frac{1}{T_f} \sum_{t=1}^{T_f} \mathcal{L}_{\text{bd}}^t, \\ \mathcal{L}_{\text{bd}}^t &= \begin{cases} \delta_{bd} - d_{bd}^t, & \text{if } d_{bd}^t < \delta_{bd} \\ 0, & \text{if } d_{bd}^t \geq \delta_{bd}, \end{cases} \end{aligned} \quad (5)$$

where δ_{bd} is the map boundary threshold.

Ego-Lane Directional Constraint. Ego-lane directional constraint comes from a prior that the vehicle's motion direction should be consistent with the lane direction where the vehicle locates. The directional constraint leverages the vectorized lane direction to regularize the motion direction of our planning trajectory. Specifically, first, we filter out low-confidence map predictions according to ϵ_m . Then we find the closest lane divider vector $\hat{v}_m \in \mathbb{R}^{T_f \times 2 \times 2}$ (within a certain range δ_{dir}) from our planning waypoint at each future timestamp. Finally, the loss for this constraint is the angular difference averaged over time between the lane vector and the ego vector:

$$\mathcal{L}_{\text{dir}} = \frac{1}{T_f} \sum_{t=1}^{T_f} \text{F}_{\text{ang}}(\hat{v}_m^t, \hat{v}_{\text{ego}}^t), \quad (6)$$

in which $\hat{v}_{\text{ego}} \in \mathbb{R}^{T_f \times 2 \times 2}$ is the planning ego vectors. \hat{v}_{ego}^t denotes the ego vector starting from the planning waypoint at the previous timestamp $t-1$ and pointing to the planning

Method	L2 (m) ↓				Collision (%) ↓				Latency (ms)	FPS
	1s	2s	3s	Avg.	1s	2s	3s	Avg.		
NMP [†] [49]	-	-	2.31	-	-	-	1.92	-	-	-
SA-NMP [†] [49]	-	-	2.05	-	-	-	1.59	-	-	-
FF [†] [18]	0.55	1.20	2.54	1.43	0.06	0.17	1.07	0.43	-	-
EO [†] [24]	0.67	1.36	2.78	1.60	0.04	0.09	0.88	0.33	-	-
ST-P3 [19]	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71	628.3	1.6
UniAD [21]	0.48	0.96	1.65	1.03	0.05	0.17	0.71	0.31	555.6	1.8
VAD-Tiny	0.46	0.76	1.12	0.78	0.21	0.35	0.58	0.38	59.5	16.8
VAD-Base	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22	224.3	4.5
VAD-Tiny [‡]	0.20	0.38	0.65	0.41	0.10	0.12	0.27	0.16	59.5	16.8
VAD-Base [‡]	0.17	0.34	0.60	0.37	0.07	0.10	0.24	0.14	224.3	4.5

Table 1. **Open-loop planning performance.** VAD achieves the best end-to-end planning performance and the fastest inference speed on the nuScenes [1] val dataset. LiDAR-based methods are denoted with †. ‡ denotes taking ego status information as input. In open-loop evaluation, we deactivate ego status information for a fair comparison. FPS of ST-P3 and VAD are measured on one NVIDIA Geforce RTX 3090 GPU. FPS of UniAD is measured on one NVIDIA Tesla A100 GPU.

waypoint at the current timestamp t . $F_{\text{ang}}(v_1, v_2)$ denotes the angular difference between vector v_1 and vector v_2 .

3.4. End-to-End Learning

Vectorized Scene Learning Loss. Vectorized scene learning includes vectorized map learning and vectorized motion prediction. For vectorized map learning, Manhattan distance is adopted to calculate the regression loss between the predicted map points and the ground truth map points. Besides, focal loss [30] is used as the map classification loss. The overall map loss is denoted as \mathcal{L}_{map} .

For vectorized motion prediction, we use l_1 loss as the regression loss to predict agent attributes (location, orientation, size, *etc.*), and focal loss [30] to predict agent classes. For each agent who has matched with a ground truth agent, we predict N_k future trajectories and use the trajectory which has the minimum final displacement error (minFDE) as a representative prediction. Then we calculate l_1 loss between this representative trajectory and the ground truth trajectory as the motion regression loss. And focal loss is adopted as the multi-modal motion classification loss. The overall motion prediction loss is denoted as \mathcal{L}_{mot} .

Vectorized Constraint Loss. The vectorized constraint loss is composed of three constraints proposed in Sec. 3.3, *i.e.*, ego-agent collision constraint \mathcal{L}_{col} , ego-boundary overstepping constraint \mathcal{L}_{bd} , and ego-lane directional constraint \mathcal{L}_{dir} , which regularize the planning trajectory \hat{V}_{ego} with vectorized scene representation.

Imitation Learning Loss. The imitation learning loss \mathcal{L}_{imi} is an l_1 loss between the planning trajectory \hat{V}_{ego} and the ground truth ego trajectory V_{ego} , aiming at guiding the

planning trajectory with expert driving behavior. \mathcal{L}_{imi} is formulated as follows:

$$\mathcal{L}_{\text{imi}} = \frac{1}{T_f} \sum_{t=1}^{T_f} \|V_{\text{ego}}^t - \hat{V}_{\text{ego}}^t\|_1. \quad (7)$$

VAD is end-to-end trainable based on the proposed vectorized planning constraint. The overall loss for end-to-end learning is the weighted sum of vectorized scene learning loss, vectorized planning constraint loss, and imitation learning loss:

$$\mathcal{L} = \omega_1 \mathcal{L}_{\text{map}} + \omega_2 \mathcal{L}_{\text{mot}} + \omega_3 \mathcal{L}_{\text{col}} + \omega_4 \mathcal{L}_{\text{bd}} + \omega_5 \mathcal{L}_{\text{dir}} + \omega_6 \mathcal{L}_{\text{imi}}. \quad (8)$$

4. Experiments

We conduct experiments on the challenging public nuScenes [1] dataset, which contains 1000 driving scenes, and each scene roughly lasts for 20 seconds. nuScenes provides 1.4M 3D bounding boxes of 23 categories in total. The scene images are captured by 6 cameras covering 360° FOV horizontally, and the keyframes are annotated at 2Hz. Following previous works [19, 21], Displacement Error (DE) and Collision Rate (CR) are adopted to comprehensively evaluate the planning performance. For the closed-loop setting, we adopt CARLA simulator [12] and the Town05 [42] benchmark for simulation. Following previous works [42, 19], Route Completion (RC) and Driving Score (DS) are used to evaluate the planning performance.

4.1. Implementation Details

VAD uses 2-second history information and plans a 3-second future trajectory. ResNet50 [16] is adopted as

ID	Agent Inter.	Map Inter.	Overstep. Const.	Dir. Const.	Col. Const.	L2 (m) ↓				Collision (%) ↓			
						1s	2s	3s	Avg.	1s	2s	3s	Avg.
1	✓	-	✓	✓	✓	0.52	0.84	1.22	0.86	0.12	0.24	0.51	0.29
2	-	✓	✓	✓	✓	0.49	0.80	1.16	0.82	0.07	0.23	0.49	0.26
3	✓	✓	-	-	-	0.43	0.73	1.11	0.76	0.10	0.22	0.52	0.28
4	✓	✓	✓	-	-	0.46	0.78	1.16	0.80	0.07	0.18	0.47	0.24
5	✓	✓	-	✓	-	0.42	0.72	1.10	0.75	0.10	0.20	0.45	0.25
6	✓	✓	-	-	✓	0.44	0.75	1.11	0.77	0.12	0.23	0.44	0.26
7	✓	✓	✓	✓	✓	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22

Table 2. **Ablation for design choices.** "Agent Inter." and "Map Inter." indicate ego-agent and ego-map query interaction in the planning module. "OverStep. Const.", "Dir. Const.", and "Col. Const." indicate Ego-Boundary Overstepping Constraint, Ego-Lane Directional Constraint, and Ego-Agent Collision Constraint, respectively.

Rasterized Representation	Vectorized Representation	Dir. Const.	Overstep. Const.	L2 (m) ↓				Collision (%) ↓			
				1s	2s	3s	Avg.	1s	2s	3s	Avg.
✓	-	-	-	0.43	0.72	1.08	0.74	0.22	0.36	0.58	0.39
-	✓	-	-	0.44	0.75	1.11	0.77	0.12	0.23	0.44	0.26
-	✓	✓	✓	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22

Table 3. **Ablation about map representation (rasterized / vectorized).** "OverStep. Const." and "Dir. Const." indicate Ego-Boundary Overstepping Constraint and Ego-Lane Directional Constraint, respectively.

Method	Town05 Short		Town05 Long	
	DS ↑	RC ↑	DS ↑	RC ↑
CILRS [10]	7.47	13.40	3.68	7.19
LBC [6]	30.97	55.01	7.05	32.09
Transfuser [†] [42]	54.52	78.41	33.15	56.36
ST-P3 [19]	55.14	86.74	11.45	83.15
VAD-Base	64.29	87.26	30.31	75.20

Table 4. **Closed-loop simulation results.** VAD achieves the best closed-loop vision-only end-to-end planning performance on CARLA [12]. [†]: LiDAR-based method.

the default backbone network for encoding image features. VAD performs vectorized mapping and motion prediction for a $60\text{m} \times 30\text{m}$ perception range longitudinally and laterally. We have two variants of VAD, which are VAD-Tiny and VAD-Base. VAD-Base is the default model for the experiments. The default number for BEV query, map query, and agent query is 200×200 , 100×20 , and 300, respectively. There is a total of 100 map vector queries, each containing 20 map points. The feature dimension and the default hidden size are 256. Compared with VAD-Base, VAD-Tiny has fewer BEV queries, which is 100×100 . The number of BEV encoder layer and decoder layer of motion and map modules is reduced from 6 to 3, and the input image size is reduced from 1280×720 to 640×360 .

As for training, the confidence thresholds ϵ_a and ϵ_m are set to 0.5, the distance thresholds δ_a , δ_{bd} and δ_{dir} are 3.0m,

Module	Latency (ms)	Proportion
Backbone	23.2	39.0%
BEV Encoder	12.3	20.7%
Motion Module	11.5	19.3%
Map Module	9.1	15.3%
Planning Module	3.4	5.7%
Total	59.5	100.0%

Table 5. **Module runtime.** The inference speed is measured for VAD-Tiny on one NVIDIA GeForce RTX 3090 GPU.

1.0m, and 2.0m, respectively. The agent safety threshold δ_X and δ_Y are set to 1.5m and 3.0m. We use AdamW [36] optimizer and Cosine Annealing [35] scheduler to train VAD with weight decay 0.01 and initial learning rate 2×10^{-4} . VAD is trained for 60 epochs on 8 NVIDIA GeForce RTX 3090 GPUs with batch size 1 per GPU.

VAD-Base is adopted for the closed-loop evaluation. The input image size is 640×320 . Following previous works [42, 19], the navigation information includes a sparse goal location and a corresponding discrete navigational command. This navigation information is encoded by an MLP and sent to the planning head as one of the input features. Besides, we add a traffic light classification branch to recognize the traffic signal. Specifically, it consists of a Resnet50 network and an MLP-based classification head. The input of this branch is the cropped front-view image,

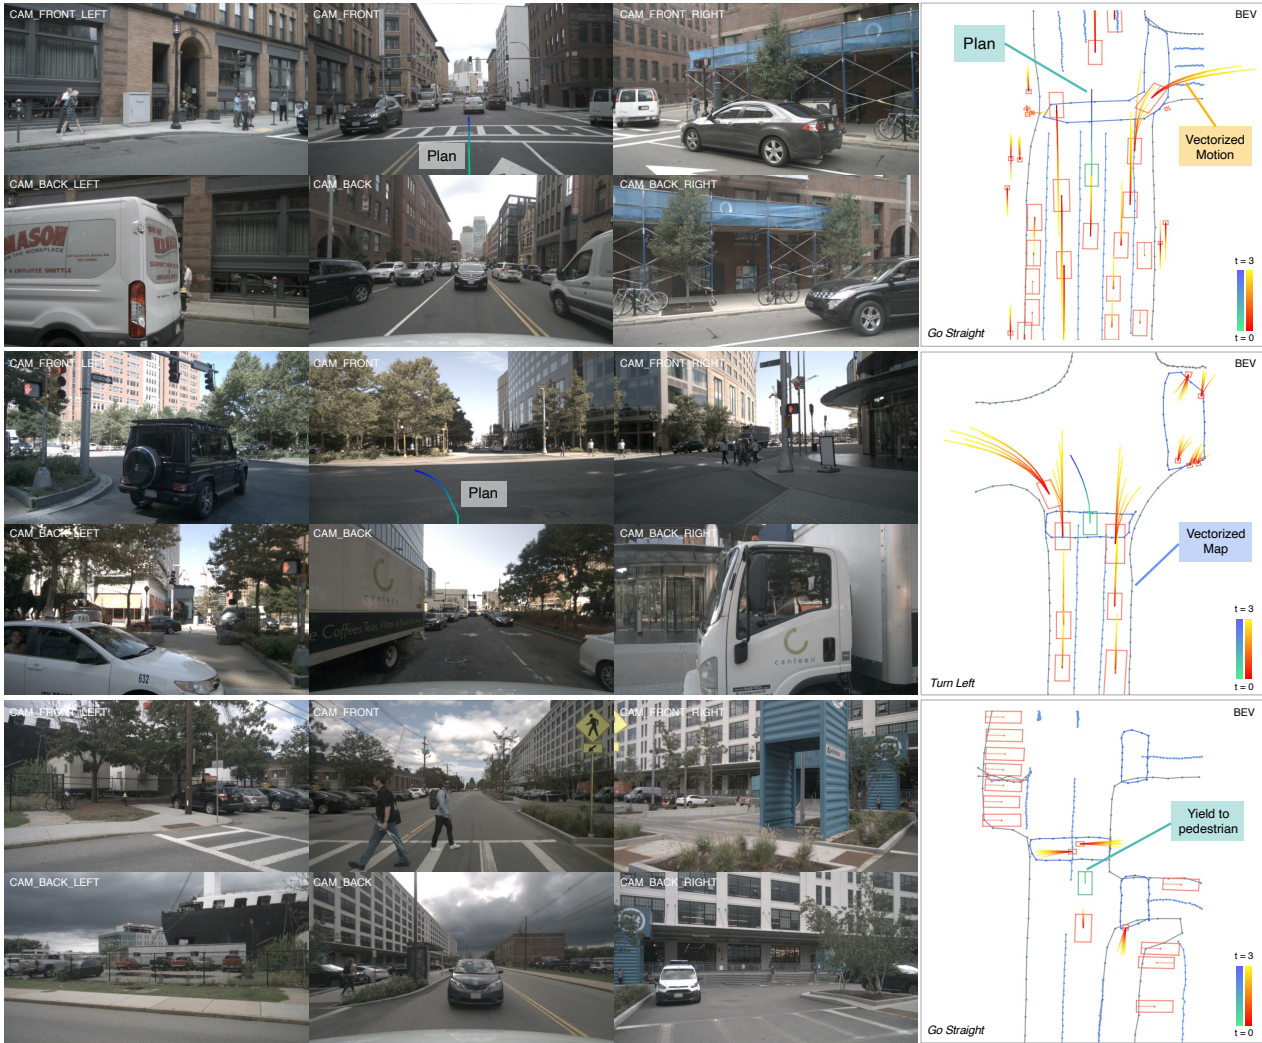


Figure 4. **Qualitative results of VAD.** VAD outputs planning results based on the vectorized scene representation in an end-to-end manner, not requiring any dense maps or hand-designed post-processing steps.

corresponding to the upper middle part of the image. The image feature map is flattened and also sent to the planning head to help the model realize the traffic light information.

4.2. Main Results

Open-loop planning results. As shown in Tab. 1, VAD shows great advantages in both performance and speed compared with the previous SOTA method [21]. On one hand, VAD-Tiny and VAD-Base greatly reduce the average planning displacement error by 0.25m and 0.31m. Meanwhile, VAD-Base greatly reduces the average collision rates by 29.0%. On the other hand, because VAD does not need many auxiliary tasks (*e.g.*, tracking and occupancy prediction) and tedious post-processing steps, it achieves the fastest inference speed based on the vectorized scene rep-

resentation. VAD-Tiny runs $9.3\times$ faster while keeping a comparable planning performance. VAD-Base achieves the best planning performance and still runs $2.5\times$ faster. It is worth noticing that in the main results, VAD omits ego status features to avoid shortcut learning in the open-loop planning [50], but the results of VAD using ego status features are still preserved in Tab. 1 for reference.

Closed-loop planning results. VAD outperforms previous SOTA vision-only end-to-end planning methods [42, 19] on the Town05 Short benchmark. Compared to ST-P3 [19], VAD greatly improves DS by 9.15 and has a better RC. On the Town05 Long benchmark, VAD achieves 30.31 DS, which is close to the LiDAR-based method [42], while significantly improving RC from 56.36 to 75.20. ST-P3 [19] obtains better RC but has a much worse DS.

4.3. Ablation Study

Effectiveness of designs. Tab. 2 shows the effectiveness of our design choices. First, because map can provide critical guidance for planning, the planning distance error is much larger without ego-map interaction (ID 1). Second, the ego-agent interaction and ego-map interaction provide implicit scene features for the ego query so that the ego car can realize others' driving intentions and plan safely. The collision rate becomes much higher without interaction (ID 1-2). Finally, the collision rate can be reduced with any of the vectorized planning constraints (ID 4-6). When utilizing the three constraints together, VAD achieves the lowest collision rate and the best planning accuracy (ID 7).

Rasterized map representation. We show the results of a VAD variant with rasterized map representation in Tab. 3. Specifically, this VAD variant utilizes map queries to perform BEV map segmentation instead of vectorized map detection, and the updated map queries are used in the planning transformer the same as VAD. As shown in Tab. 3, VAD with rasterized map representation suffers from a much higher collision rate.

Runtime of each module. We evaluate the runtime of each module of VAD-Tiny, and the results are shown in Tab. 5. Backbone and BEV Encoder take most of the runtime for feature extraction and transformation. Then motion module and map module take 34.6% of the total runtime to accomplish multi-agent vectorized motion prediction and vectorized map prediction. The runtime of the planning module is only 3.4ms, thanks to the sparse vectorized representation and concise model design.

4.4. Qualitative Results

We show three vectorized scene learning and planning results of VAD in Fig. 4. For a better understanding of the scene, we also provide raw surrounding camera images and project the planning trajectories to the front camera image. VAD can predict multi-modality agent motions and map elements accurately, as well as plan the ego future movements reasonably according to the vectorized scene representation. More visualizations are available in Appendix.

5. Conclusion

In this paper, we explore the fully vectorized representation of the driving scene, and how to effectively incorporate the vectorized scene information for better planning performance. The resulting end-to-end autonomous driving paradigm is termed VAD. VAD achieves both high performance and high efficiency, which are vital for the safety and deployment of an autonomous driving system. We hope the

impressive performance of VAD can reveal the potential of vectorized paradigm to the community.

VAD predicts multi-modality motion trajectories for other dynamic agents, We use the most confident prediction in our collision constraint to improve planning safety. How to utilize the multi-modality motion predictions for planning, is worthy of future discussion. Besides, how to incorporate other traffic information (*e.g.*, lane graph, road sign, traffic light, and speed limit) into this autonomous driving system, also deserves further exploration.

Acknowledgement

This work was in part supported by the National Natural Science Foundation of China (No. 62276108).

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 2, 6
- [2] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *CVPR*, 2021. 1, 3
- [3] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *arXiv preprint arXiv:1910.05449*, 2019. 3
- [4] Raphael Chekroun, Marin Toromanoff, Sascha Hornauer, and Fabien Moutarde. Gri: General reinforced imitation and its application to vision-based autonomous driving. *arXiv preprint arXiv:2111.08575*, 2021. 3
- [5] Dian Chen, Vladlen Koltun, and Philipp Krähenbühl. Learning to drive from a world on rails. In *ICCV*, 2021. 3
- [6] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. 2020. 7
- [7] Long Chen, Lukas Platinsky, Stefanie Speichert, Błażej Osiniński, Oliver Scheel, Yawei Ye, Hugo Grimmer, Luca Del Pero, and Peter Ondruska. What data do we need for training an av motion planner? In *ICRA*, 2021. 1
- [8] Shaoyu Chen, Tianheng Cheng, Xinggang Wang, Wenming Meng, Qian Zhang, and Wenyu Liu. Efficient and robust 2d-to-bev representation learning via geometry-guided kernel transformer. *arXiv preprint arXiv:2206.04584*, 2022. 2
- [9] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *ICCV*, 2019. 1, 3
- [10] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. 2019. 7
- [11] Alexander Cui, Sergio Casas, Abbas Sadat, Renjie Liao, and Raquel Urtasun. Lookout: Diverse multi-future prediction and planning for self-driving. In *ICCV*, 2021. 1, 3
- [12] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *CoRL*, 2017. 6, 7

- [13] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *CVPR*, 2020. 3
- [14] David González, Joshué Pérez, Vicente Milanés, and Fawzi Nashashibi. A review of motion planning techniques for automated vehicles. *T-ITS*, 2015. 1
- [15] Junru Gu, Chenxu Hu, Tianyuan Zhang, Xuanyao Chen, Yilun Wang, Yue Wang, and Hang Zhao. Vip3d: End-to-end visual trajectory prediction via 3d agent queries. *arXiv preprint arXiv:2208.01582*, 2022. 3
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 6
- [17] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras. In *ICCV*, 2021. 2, 3
- [18] Peiyun Hu, Aaron Huang, John Dolan, David Held, and Deva Ramanan. Safe local motion planning with self-supervised freespace forecasting. In *CVPR*, 2021. 6
- [19] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *ECCV*, 2022. 1, 3, 4, 5, 6, 7, 8
- [20] Yihan Hu, Wenxin Shao, Bo Jiang, Jiajie Chen, Siqi Chai, Zhening Yang, Jingyu Qian, Helong Zhou, and Qiang Liu. Hope: Hierarchical spatial-temporal network for occupancy flow prediction. *arXiv preprint arXiv:2206.10118*, 2022. 3
- [21] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Goal-oriented autonomous driving. *arXiv preprint arXiv:2212.10156*, 2022. 1, 2, 3, 4, 5, 6, 8
- [22] Bo Jiang, Shaoyu Chen, Xinggang Wang, Bencheng Liao, Tianheng Cheng, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, and Chang Huang. Perceive, interact, predict: Learning dynamic and static clues for end-to-end motion prediction. *arXiv preprint arXiv:2212.02181*, 2022. 3, 4
- [23] Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen, Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to drive in a day. In *ICRA*, 2019. 1
- [24] Tarasha Khurana, Peiyun Hu, Achal Dave, Jason Ziglar, David Held, and Deva Ramanan. Differentiable raycasting for self-supervised occupancy forecasting. 2022. 6
- [25] Qi Li, Yue Wang, Yilun Wang, and Hang Zhao. Hdmapnet: An online hd map construction and evaluation framework. In *ICRA*, 2022. 2
- [26] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022. 2, 3
- [27] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, 2020. 3
- [28] Bencheng Liao, Shaoyu Chen, Bo Jiang, Tianheng Cheng, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Lane graph as path: Continuity-preserving path-wise modeling for online lane graph construction. *arXiv preprint arXiv:2303.08815*, 2023. 2
- [29] Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang Huang. Maptr: Structured modeling and learning for online vectorized hd map construction. *arXiv preprint arXiv:2208.14437*, 2022. 2, 4
- [30] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 6
- [31] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *arXiv preprint arXiv:2203.05625*, 2022. 2
- [32] Yicheng Liu, Yue Wang, Yilun Wang, and Hang Zhao. Vectormapnet: End-to-end vectorized hd map learning. *arXiv preprint arXiv:2206.08920*, 2022. 2
- [33] Yicheng Liu, Jinghui Zhang, Liangji Fang, Qinlong Jiang, and Bolei Zhou. Multimodal motion prediction with stacked transformers. In *CVPR*, 2021. 3
- [34] Zhi Liu, Shaoyu Chen, Xiaojie Guo, Xinggang Wang, Tianheng Cheng, Hongmei Zhu, Qian Zhang, Wenyu Liu, and Yi Zhang. Vision-based uneven bev representation learning with polar rasterization and surface estimation. *CoRL*, 2022. 2
- [35] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 7
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7
- [37] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting multiple agent trajectories. *arXiv preprint arXiv:2106.08417*, 2021. 3, 4
- [38] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Covernet: Multimodal behavior prediction using trajectory sets. In *CVPR*, 2020. 3
- [39] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, 2020. 2, 3
- [40] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *NeurIPS*, 1988. 1, 3
- [41] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *CVPR*, 2021. 1, 3
- [42] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. 2021. 6, 7, 8
- [43] Katrin Renz, Kashyap Chitta, Otniel-Bogdan Mercea, A Koepke, Zeynep Akata, and Andreas Geiger. Plant: Explainable planning transformers via object-level representations. *arXiv preprint arXiv:2210.14222*, 2022. 3

- [44] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *ECCV*, 2020. 3
- [45] Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *CVPR*, 2020. 3
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 2, 3, 4
- [47] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *CoRL*, 2022. 2
- [48] Wenda Xu, Qian Wang, and John M Dolan. Autonomous vehicle motion planning via recurrent spline optimization. In *ICRA*, 2021. 1
- [49] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019. 6
- [50] 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. 8
- [51] Yunpeng Zhang, Zheng Zhu, Wenzhao Zheng, Junjie Huang, Guan Huang, Jie Zhou, and Jiwen Lu. Beverage: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving. *arXiv preprint arXiv:2205.09743*, 2022. 2, 3
- [52] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *CVPR*, 2022. 3
- [53] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 4