

CogAD: Cognitive-Hierarchy Guided End-to-End Autonomous Driving

Zhennan Wang* Jianing Teng* Canqun Xiang Kangliang Chen
 Xing Pan Lu Deng Weihao Gu†
 HAOMO.AI Technology Co., Ltd

wangzhennan2017@email.szu.edu.cn*, tengjn@alumni.sysu.edu.cn*
 guwh22@mails.tsinghua.edu.cn†

Abstract

While end-to-end autonomous driving has advanced significantly, prevailing methods remain fundamentally misaligned with human cognitive principles in both perception and planning. In this paper, we propose CogAD, a novel end-to-end autonomous driving model that emulates the hierarchical cognition mechanisms of human drivers. CogAD implements dual hierarchical mechanisms: global-to-local context processing for human-like perception and intent-conditioned multi-mode trajectory generation for cognitively-inspired planning. The proposed method demonstrates three principal advantages: comprehensive environmental understanding through hierarchical perception, robust planning exploration enabled by multi-level planning, and diverse yet reasonable multi-modal trajectory generation facilitated by dual-level uncertainty modeling. Extensive experiments on nuScenes and Bench2Drive demonstrate that CogAD achieves state-of-the-art performance in end-to-end planning, exhibiting particular superiority in long-tail scenarios and robust generalization to complex real-world driving conditions.

1. Introduction

In recent years, end-to-end autonomous driving (E2E-AD) has been a topic of interest [19, 23, 24, 45, 54]. E2E-AD methods elevate the performance upper bound of autonomous driving, while broadening the architectural design space for models [54]. Although significant progress has been made in the field, current approaches exhibit substantial discrepancies from human driving processes in both perception and planning patterns, from the perspective of cognitive psychology.

In the perception pattern, we argue that the perception of human drivers is hierarchical. Cognitive psychology supports the pre-attentive two-stage theory of human visual

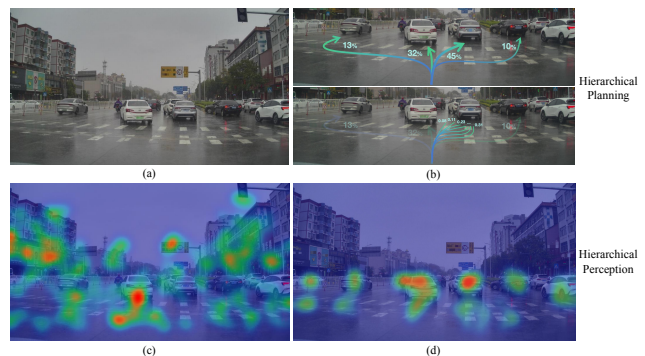


Figure 1. Human drivers scan surrounds (c) before focusing on key objects (d), and plan hierarchically from intent to trajectory (b).

perception [27], which posits that individuals first perform global scanning of the scene to identify contextual or spatial cues, followed by selective attention allocation to critical local objects for detailed analysis. This mechanism aligns with neurobiological evidence [13] demonstrating that early visual areas prioritize scene-wide features, while higher-order regions refine task-specific processing through context-dependent modulation. In terms of driving environmental perception, human drivers typically initiate their environmental perception by establishing holistic scene comprehension of the traffic scenario, encompassing all contextual elements as illustrated in Fig. 1 (c), followed by selective attention prioritization towards critical dynamic objects such as proximal vehicles, pedestrians, and lane markings, as illustrated in Fig. 1 (d).

In the planning pattern, we argue that the planning of human drivers is hierarchical. The BDI (Belief-Desire-Intention) cognitive model [2, 39] decomposes human action planning into hierarchically organized processes, distinguishing between high-level intentional states (e.g., goal setting) and low-level behavioral executions (e.g., motor sequence implementation). In terms of driving planning, human drivers typically follow a hierarchical planning process: intentional planning at the higher level (e.g., lane

* Equal contribution; † Corresponding author.

change) followed by trajectory planning at the lower level (e.g., generating a specific trajectory). For instance, drivers first establish a global intent, which then constrains the generation of local trajectories, as illustrated in Fig. 1 (b).

From the perspective of hierarchical cognitive psychology, recent methods like [19, 24, 54] fail to incorporate hierarchical perception, while others like [23, 33, 45] lack hierarchical planning, as detailed in Sec. 2. These analyses highlight a critical gap in aligning E2E-AD models with human drivers’ cognitive principles. To bridge this gap, we propose CogAD, an end-to-end autonomous driving model that emulates human hierarchical cognition. For hierarchical perception, we design a sequential interaction mechanism where the ego-vehicle first processes global BEV features to capture environmental context, then focuses on critical instance-level elements. For hierarchical planning, CogAD first plans high-level driving intents, then generates corresponding low-level trajectories, mirroring human multi-level planning.

The advantages of CogAD are summarized as follows: (a) CogAD captures hierarchical environmental representations, encompassing both holistic scene features and critical element attributes; (b) CogAD achieves comprehensive planning space exploration through dual-level uncertainty modeling at both intent and trajectory levels; (c) CogAD inherently supports multi-modal trajectory planning, generating diverse and plausible trajectories for downstream tasks. In summary, the main contributions of this paper are three-fold:

- We propose a hierarchical scene-instance perception paradigm that significantly enhances the ego vehicle’s scene understanding capabilities.
- We develop a hierarchical intent-trajectory planning mechanism that simultaneously enhances both behavioral diversity and motion rationality in end-to-end autonomous driving.
- CogAD achieves state-of-the-art performance in both open-loop and closed-loop evaluation, with particularly significant improvements in long-tail scenarios compared to prior methods.

2. Related Work

Significant progress has been made in end-to-end autonomous driving (E2E-AD) approaches, pioneered by prior works such as [1, 10, 36].

Existing approaches exhibit distinct limitations in hierarchical feature utilization. Scene-centric methods [9, 31, 32, 54, 55] emphasize scene-level features while systematically overlooking instance-level characteristics. Specifically, ParaDrive [54] and BEV-Planner [32] solely employ BEV features extracted from visual inputs, whereas TransFuser-based architectures [9, 31, 55] integrate BEV features obtained through multimodal lidar-camera fusion. Con-

versely, instance-oriented approaches [11, 24, 29, 44, 45, 60] prioritize instance-specific attributes but fail to integrate scene-level contextual information. For instance, VAD [24] establishes interactions between ego queries and map/agent queries without BEV feature integration, while SparseDrive [45] restricts feature interactions to current and historical instance-level representations. Although DiFSD [44] implements hierarchical interactions within its planning module, these remain confined to inter-instance communication. Despite leveraging both scene-level context and instance-level details, methods like [8, 19, 23, 57] do not exhibit a clear perceptual hierarchy in feature interaction. Notably, UniAD [19], FusionAD [57], and PPAD [8] adopt an inverted interaction sequence—prioritizing instance-level processing before BEV feature integration—contrary to human cognitive patterns. DriveTransformer [23] employs iterative refinement between instance and scene representations, yet lacks explicit hierarchical separation. In contrast, CogAD establishes a cognitively-inspired perception hierarchy through a scene-prioritized interaction paradigm followed by instance-level refinement, achieving principled decoupling of hierarchical representations.

Recent research in E2E-AD reveals notable limitations in hierarchical planning with uncertainty modeling across different abstraction levels. Current approaches predominantly focus on deterministic trajectory generation [8, 9, 11, 14, 18, 19, 24, 29, 32, 41, 42, 54, 57, 60], fundamentally neglecting uncertainty quantification. While multimodal trajectory prediction methods [23, 44–46, 59] incorporate trajectory-level uncertainty, they fail to address higher-order intent uncertainty. Notably, SparseDrive [45] employs learnable embeddings to represent the planning modes and adopts the winner-takes-all strategy for loss computation, analogous to our trajectory uncertainty modeling. HiPAD [46] implements temporal-spatial hierarchical planning with driving-style modeling but lacks explicit intent representation. Alternative approaches [33, 53] extract diverse motion patterns that capture intent-level uncertainty, neglecting trajectory-level uncertainty. In particular, DriveWM [53] employs human-specified high-level commands to represent intent in a heuristic manner. Similar to our intent uncertainty module, DiffusionDrive [33] constructs planning anchors by offline K-Means clustering on the training set, yet omit trajectory-level uncertainty. Other methods generating multimodal trajectories via large-scale sampling [5, 31, 55] conflate intent and trajectory uncertainties. Some non-end-to-end autonomous driving approaches also incorporate hierarchical planning design. Methods like [26, 37] adopt a temporal hierarchical paradigm, diverging fundamentally from our uncertainty-aware intent-trajectory dual hierarchy. Other approaches proposed in [12, 34, 49, 52] similarly employ hierarchical planning, yet critically depend on manually engineered command definitions rather than

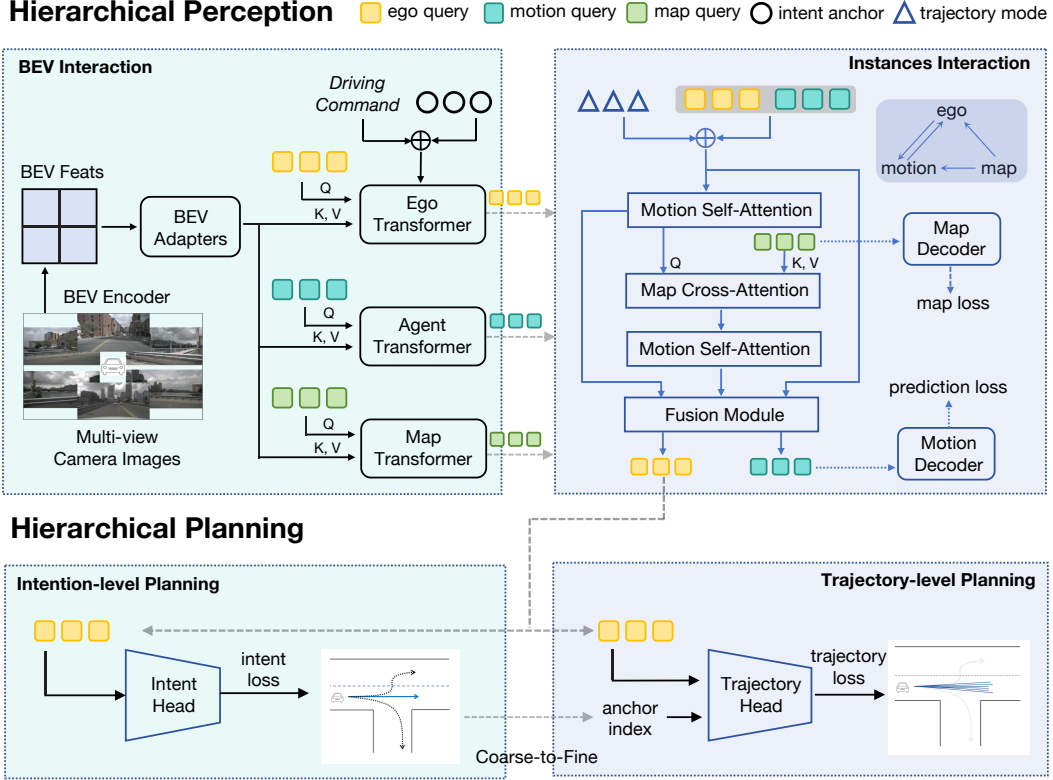


Figure 2. The overall framework of CogAD. CogAD extracts BEV features into task-specific queries, then performs cross-task instance feature interaction, forming a hierarchical perception paradigm. Meanwhile, CogAD implements intent-level planning and subsequently conducts trajectory-level planning, establishing a hierarchical planning mechanism.

data-driven representations. Contemporary LLM-based planners [7, 15, 20, 25, 35, 38, 47, 48, 56] demonstrate conceptual parallels through meta-action prediction but suffer from manually defined commands and computational inefficiency. Distinct from these approaches, our CogAD introduces a novel bi-level uncertainty-aware planning mechanism that decouples intent uncertainty from trajectory uncertainty, achieving superior performance with optimized computational efficiency.

3. Method

3.1. Overview

As depicted in Fig. 2, the proposed CogAD implements a unified multi-task architecture that integrates perception, prediction, and planning in a coordinated manner. To comprehensively capture environmental information, we design a hierarchical feature interaction framework. First, we establish interactions between task-specific instance queries (ego, agent, and map) and dense Bird’s Eye View (BEV) features while incorporating intent uncertainty in the ego-vehicle’s planning queries. Subsequently, cross-task instance interaction is conducted among task-specific sparse representations with inter-dependencies, where trajectory uncertainty is explicitly modeled in the ego-vehicle’s planning instance.

Next, we introduce the hierarchical BEV-instance perception paradigm and the hierarchical intent-trajectory planning mechanism.

3.2. Hierarchical Perception

Interaction with BEV Features. BEV features provide a unified scene representation that captures global context and geometric relationships. We therefore initiate the feature interaction process by bridging the task-specific instance queries with BEV features. For BEV features extraction, we implement the BEVFormer [30] framework due to its effective spatial-temporal fusion mechanism, which captures both geometric layouts and temporal motion dynamics. We develop a dedicated BEV adapter for each task to enhance the task-specific adaptability of BEV features.

$$\begin{aligned}
 Q_{ego} &= E_{ego} + E_{cmd} + E_{intent} \\
 I_{ego} &= Transformer(Q = Q_{ego}, \\
 &K, V = Adapter_{ego}(F_{bev}))
 \end{aligned}
 \tag{1}$$

Eq. (1) models the interaction between planning queries and BEV features F_{bev} . $Transformer$ represents the standard transformer module [51]. CogAD integrates ego embedding E_{ego} , high-level driving commands (e.g., lane-keeping/turning) E_{cmd} , and intent anchors E_{intent} to

initialize the ego-vehicle queries Q_{ego} , which are then fed into the Ego Transformer module shown in Fig. 2. The intent anchors encode the ego vehicle’s coarse planning space, with implementation details provided in Sec. 3.3.

CogAD also performs online mapping and dynamic obstacle detection via the interaction between BEV features and map/motion queries, constituting mission-critical modules that are algorithmically coupled with the ego planning task. Online mapping provides topological constraints for planning. CogAD employs multiple learnable spatial interaction queries to decode structured road topology from BEV features through multi-head cross attention, implemented through the MapTransformer module in Fig. 2. This interaction process can be formally defined as:

$$I_{map} = Transformer(Q = Q_{map}, K, V = Adapter_{map}(F_{bev})) \quad (2)$$

Obstacle detection provides precise spatial localization of surrounding agents, thereby ensuring safety-critical planning. CogAD initiates with interaction queries that extract spatiotemporal information from BEV features via the DetTransformer module in Fig. 2, whereby geometric attributes (position, heading angle) and semantic properties (classification confidence) are jointly predicted via a lightweight multi-layer perceptron. The interaction process can be formulated as follows:

$$I_{agent} = Transformer(Q = Q_{agent}, K, V = Adapter_{agent}(F_{bev})) \quad (3)$$

Cross-Task Instance Interaction. Although BEV features facilitate a unified scene understanding, they exhibit limitations in representing the multimodal future trajectories of dynamic agents and ego-vehicle. Additionally, gradient conflicts [6, 21] are likely to arise among the loss functions of multiple tasks, within the BEV space. To address these limitations, CogAD involves inter-instances interactions across different tasks, as shown in the upper right part of Fig. 2.

In the planning module, we account for trajectory uncertainty by introducing motion mode embeddings, as elaborated in Sec. 3.3. Regarding obstacles, we enhance their representations by incorporating dynamic attributes through motion prediction, aiming to avoid potential collisions with the ego-vehicle’s planned trajectory. Similarly, we introduce motion mode embeddings to facilitate multi-modal motion prediction. We align with the philosophy advocated by SparseDrive [45], which emphasizes that motion prediction and planning should account for bidirectional interactions. To this end, the instance queries for ego-vehicle planning and motion prediction are concatenated and then fed into a self-attention (SA) module, thereby facilitating bidirectional information interaction between the two tasks. This design

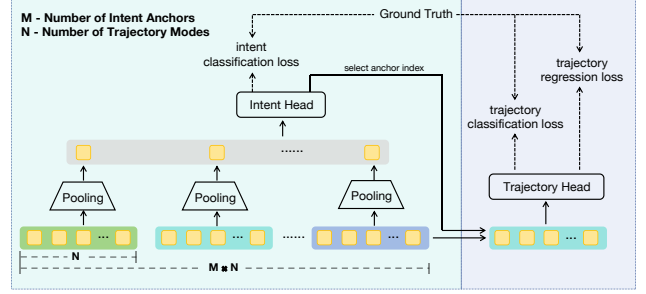


Figure 3. Intent-level and Trajectory-level uncertainty modeling.

enables the planning task to consider potential future behaviors of surrounding agents, while the motion prediction task also benefits from the planned trajectory of the ego-vehicle. This interaction process can be formulated as follows:

$$I_{ego} = I_{ego} + E_{mode}, \quad I_{mot} = I_{agent} + E_{mode} \quad (4)$$

$$I'_{ego}, I'_{mot} = Transformer(Q, K, V = [I_{ego}, I_{mot}])$$

Moreover, given the critical role of map information in both tasks (e.g., lane-keeping assistance, identifying road curvature, and providing critical boundary information), CogAD employs cross-attention (CA) to separately distill map features into the instance embeddings of each task. After acquiring map information, both the planning and motion prediction instance embeddings update their respective feature representations. To ensure mutual awareness of potential future trajectories between the ego-vehicle and surrounding agents, CogAD employs self-attention to model bidirectional interactions based on the updated embeddings.

$$I''_{ego}, I''_{mot} = Transformer(Q, K, V = Transformer(Q = [I'_{ego}, I'_{mot}], K, V = I_{map})) \quad (5)$$

Meanwhile, to guarantee effective utilization of multi-stage features and enhance the absorption of supervised signals, we adopt skip connections that directly transport the outputs from the BEV interaction (I_{ego}, I_{mot}), the initial ego-agent self-attention interaction (I'_{ego}, I'_{mot}), and the final instance embeddings (I''_{ego}, I''_{mot}) to the corresponding task heads.

3.3. Hierarchical Planning

Intent Uncertainty. Intent uncertainty is introduced during the interaction process between the ego-vehicle and the BEV features. Intent uncertainty characterizes the inherent ambiguity associated with the high-level intent or targeted objective of an autonomous vehicle in dynamic environments. In driving scenarios, this can involve uncertainty regarding the lane-change intentions or trajectory convergence points toward which the ego-vehicle is actively steering. Modeling intent uncertainty is crucial for precisely

Table 1. Planning performance on nuScenes dataset. We use * for LLM-based method.

Method	L2 (m) ↓				Collision (%) ↓				Latency (ms)
	1s	2s	3s	Avg.	1s	2s	3s	Avg.	
UniAD [19]	0.48	0.74	1.07	0.76	0.12	0.13	0.28	0.17	555.6
VAD [24]	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22	224.3
ParaDrive [54]	0.25	0.46	0.74	0.48	0.14	0.23	0.39	0.25	-
SparseDrive [45]	0.29	0.55	0.91	0.58	0.01	0.02	0.13	0.06	136.9
DiffusionDrive [33]	0.27	0.54	0.90	0.57	0.03	0.05	0.16	0.08	-
Senna* [25]	0.37	0.54	0.86	0.59	0.09	0.12	0.33	0.18	-
TOKEN* [47]	0.26	0.71	1.47	0.81	-	-	-	-	-
VLM-AD* [56]	0.30	0.54	0.80	0.55	0.11	0.15	0.38	0.21	-
CogAD	0.24	0.45	0.74	0.48	0.00	0.02	0.15	0.06	92.9

forecasting the ego-vehicle’s future behavior and improving planning robustness under uncertain conditions [4].

Unlike many LLM-based autonomous driving methods [25, 48] that rely on hand-crafted meta-actions to represent intent, CogAD leverages a data-driven approach to obtain intent points. Following practices similar to those in VADv2 [5] and Hydra-MDP [31], we employ K-means clustering to generate trajectory anchors, which serve as intent prototypes. To simplify the implementation, we utilize online K-means for constructing trajectory anchors. Notably, as our focus is on high-level intent discovery rather than fine-grained trajectory clustering, the cluster count is significantly lower than these previous methods [5, 31]. After the hierarchical perception, the ego queries are pooled along the anchor dimension and subsequently processed through the intent head to compute the intent classification loss or derive the selected anchor index, as illustrated in the left part of Fig. 3.

Trajectory Uncertainty. Trajectory uncertainty is incorporated into the interaction process between planning instances and motion prediction instances. Trajectory uncertainty, which refers to the unpredictability of a vehicle’s future path due to factors such as environmental dynamics, diverse behaviors of road users, and subjective driving behavior, exists in both planning and motion prediction tasks. From an alternative perspective, trajectory uncertainty can be conceptualized as the uncertainty pertaining to the offset from the trajectory anchors within the latent space.

Drawing inspiration from MultiPath++ [50], our hypothesis posits that trajectory uncertainty follows some intrinsic motion patterns. To capture these motion patterns, we directly employ learnable mode embeddings to represent them. We argue that these motion patterns demonstrate consistency across both planning and motion prediction tasks, given that both tasks fundamentally involve forecasting the future trajectories of agents. Therefore, we utilize shared mode

embeddings for both planning and motion prediction tasks. Our experimental results indicate that this strategy not only enhances model performance, but also effectively avoids mode collapse [4, 17, 33, 40]. Finally, based on the anchor selected by the intent head, the trajectory head generates multi-mode trajectories. During training, a winner-takes-all strategy [28, 50] is employed to compute the regression loss and classification loss, as shown on the right part of Fig. 3.

3.4. Training and Inference

Training. The system enables joint optimization of all task objectives via differentiable end-to-end learning. In particular, the planning module incorporates two core loss components: an intent grounding loss corresponding to the intent uncertainty and a multi-mode trajectory imitation loss corresponding to the trajectory uncertainty. For intent grounding, the intent anchor that exhibits minimum spatial deviation from the ground truth ego trajectory is designated as the label. To model the multi-mode imitation behavior, we employ the winner-takes-all strategy [28] to select the prediction closest to the ground-truth, following motion prediction paradigms [24, 50]. In addition, we incorporate the planning constraints utilized in VAD [24]. The overall training loss is the weighted sum of the aforementioned components:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{map}} + \lambda_2 \mathcal{L}_{\text{det}} + \lambda_3 \mathcal{L}_{\text{mot}} + \lambda_4 \mathcal{L}_{\text{plan_intent}} + \lambda_5 \mathcal{L}_{\text{plan_wta}} + \lambda_6 \mathcal{L}_{\text{plan_constr}} + \lambda_7 \mathcal{L}_{\text{k-means}} \quad (6)$$

where \mathcal{L}_{map} denotes the online mapping loss, \mathcal{L}_{det} denotes the obstacle detection loss, \mathcal{L}_{mot} denotes the motion prediction loss, $\mathcal{L}_{\text{plan_intent}}$ denotes the intent grounding loss, $\mathcal{L}_{\text{plan_wta}}$ denotes the multi-mode trajectory imitation loss, $\mathcal{L}_{\text{plan_constr}}$ denotes the overall planning constraints, and the $\mathcal{L}_{\text{k-means}}$ drives the online K-means clustering to generate intent anchors. Eq. (6) is optimized by joint training all CogAD components, with the loss weighting parameters set to $\lambda_1 = \lambda_2 = 2.0$, $\lambda_3 = 0.2$, and $\lambda_4 = \lambda_5 = \lambda_6 = \lambda_7 = 1.0$.

Table 2. Long-tail scenario performance comparison.

Method	L2 (m) ↓				CR (%) ↓
	1s	2s	3s	Avg.	Avg.
3-point turn (zero-shot)					
VAD [24]	0.71	1.66	3.24	1.57	0.00
PARA-Drive [54]	0.50	1.38	2.76	1.29	5.33
TOKEN [47]	0.39	1.29	2.60	1.18	4.00
DiMA [16]	0.36	1.18	2.37	1.05	0.00
CogAD	0.41	0.81	1.47	0.90	0.00
Resume from stop					
VAD [24]	0.60	1.72	2.83	1.42	0.00
PARA-Drive [54]	0.14	0.79	2.30	0.85	0.00
TOKEN [47]	0.13	0.70	1.58	0.65	0.00
DiMA [16]	0.15	0.65	1.34	0.66	0.00
CogAD	0.10	0.27	0.64	0.34	0.00
Overtake					
VAD [24]	0.46	1.16	2.17	1.06	2.49
PARA-Drive [54]	0.27	0.89	1.94	0.85	2.30
TOKEN [47]	0.29	0.77	1.63	0.74	0.00
DiMA [16]	0.24	0.72	1.50	0.66	1.29
CogAD	0.28	0.53	0.86	0.56	0.28

Inference. Benefiting from our hierarchical uncertainty modeling, the inference stage supports multiple operational modes. (a) Deterministic Trajectory: Select the highest-confidence intent and trajectory at both levels. (b) Intent Sampling: Apply probabilistic sampling at the intent level while selecting the highest-confidence trajectory at the trajectory level. (c) Trajectory Sampling: Choose the highest-confidence intent and apply probabilistic sampling at the trajectory level. (d) Dual-Level Sampling: Employ probabilistic sampling at both intent and trajectory levels. The probabilistic sampling mechanism enhances trajectory diversity, which is critical in closed-loop experiments to mitigate the risk of the ego vehicle becoming trapped in suboptimal positions.

4. Experiments

Our experiments are conducted on the nuScenes dataset [3] and the Bench2Drive benchmark [22] for open-loop and closed-loop evaluations, respectively. We follow TOKEN [47] and manually construct a long-tail scenario validation set curated from nuScenes for comprehensive evaluation, including three scenarios: 1) executing 3-point turns; 2) resuming motion after a full stop; 3) overtaking parked cars through the oncoming lane. CogAD intentionally excludes any form of ego state or historical trajectory as input to avoid overfitting [32, 58]. Moreover, we set the number of intent anchors to 30 and the trajectory modes to 6. For fairness, we compare only against methods that do not utilize this information and adhere to the same training protocol of

Table 3. Closed-loop evaluation results on Bench2Drive. DS: Driving Score. SR: Success Rate.

Method	DS ↑	SR (%) ↑	Latency(ms)
UniAD [19]	45.81	16.36	663.4
VAD [24]	42.35	15.00	278.3
GenAD [60]	44.81	15.91	-
MomAD [43]	44.54	16.71	-
CogAD	48.30	24.00	121.3

60 epochs on nuScenes and 6 epochs on Bench2Drive.

4.1. Main Results

Open-loop Planning Results. We compare CogAD against recent state-of-the-art (sota) end-to-end methods, including LLM-based planners. As shown in Tab. 1, CogAD demonstrates significant advantages in both performance and efficiency. Specifically, CogAD achieves a sota average collision rate of 0.06, substantially outperforming previous methods, including those enhanced with LLM. Meanwhile, CogAD maintains competitively low L2 error and high computational efficiency. To our knowledge, CogAD is the first to achieve 0% collision rate at the 1s planning horizon without relying on ego-state or historical ego information. CogAD demonstrates superior overall performance in terms of L2 error, collision rate, and inference latency. Notably, this is achieved without employing resource-intensive techniques, such as the massive parameter counts or the post-selection refinement. These results highlight the effectiveness and efficiency of CogAD in real-world autonomous driving scenarios.

Long-tail Scenario Results. Following TOKEN [47], we construct a long-tail scenario validation set curated from nuScenes for comprehensive evaluation. As shown in Tab. 2, CogAD achieves significant performance improvements over previous methods. Notably, CogAD establishes new sota results in L2 error for all three long-tail scenarios. In the challenging "3-point turn", a zero-shot case absent from training data, CogAD achieves a 10% reduction in L2 error compared to previous sota result. For the "overtake" scenario, CogAD achieves a competitively low collision rate that nearly matches TOKEN [47]'s performance, while simultaneously outperforming it in average L2 error metric by 0.18m.

Closed-loop Planning Results. As shown in Tab. 3, the closed-loop evaluation of CogAD on the challenging Bench2Drive benchmark demonstrates its superior performance. Compared to other methods, CogAD achieves state-of-the-art performance in Driving Score and Success Rate, alongside competitive latency.

Table 4. Ablation for designs in Hierarchical Perception and Hierarchical Planning modules.

Ablation Setting		L2 (m) ↓				Collision (%) ↓				
	BEV	Instance	1s	2s	3s	Avg.	1s	2s	3s	Avg.
	Interaction	Interaction								
Hierarchical Perception	-	✓	0.27	0.49	0.81	0.52	0.05	0.11	0.23	0.13
	✓	-	0.26	0.45	0.75	0.49	0.03	0.09	0.23	0.12
	✓	✓	0.24	0.45	0.74	0.48	0.00	0.02	0.15	0.06
	Intent	Trajectory	1s	2s	3s	Avg.	1s	2s	3s	Avg.
	Uncertain	Uncertain								
Hierarchical Planning	-	✓	0.26	0.49	0.82	0.52	0.05	0.11	0.32	0.16
	✓	-	0.22	0.41	0.69	0.44	0.09	0.13	0.27	0.16
	✓	✓	0.24	0.45	0.74	0.48	0.00	0.02	0.15	0.06

Table 5. Ablation for the design of model details.

BEV Adapter	Mode Sharing	Ego-Motion Interaction	L2 (m) ↓				Collision (%) ↓			
			1s	2s	3s	Avg.	1s	2s	3s	Avg.
-	✓	✓	0.25	0.47	0.80	0.51	0.47	0.55	0.78	0.60
✓	-	✓	0.22	0.41	0.68	0.44	0.05	0.15	0.18	0.13
✓	✓	-	0.25	0.46	0.77	0.49	0.05	0.10	0.25	0.13
✓	✓	✓	0.24	0.45	0.74	0.48	0.00	0.02	0.15	0.06

4.2. Ablation Study

Effect of Designs in Hierarchical Perception. To validate the effectiveness of the hierarchical perception paradigm, we conduct two ablation studies: one retaining only the ego queries’ interaction with BEV features, and the other preserving solely the ego-to-instance queries interaction. As shown in Tab. 4, interaction with BEV features benefits both the L2 error and collision rate, while instance-level interaction predominantly reduces the collision rate. These results collectively validate the effectiveness of our hierarchical perception paradigm.

Effect of Designs in Hierarchical Planning. Tab. 4 demonstrates the effectiveness of our hierarchical planning design. The absence of intent uncertainty degrades CogAD’s planning module to a conventional multi-mode ego-trajectory planning method, resulting in a significantly higher collision rate. Besides, preserving only intent uncertainty while omitting trajectory uncertainty also yields suboptimal results, as this configuration eliminates the model’s capacity to generate refined planning trajectories. Our visualization analysis confirms that hierarchical planning delivers critical advantages in collision avoidance, enabling maneuvers like deceleration for yielding and lane-changing to bypass obstacles. This explains the results in Tab. 4: Collision Rate decreases by 62.5% versus non-hierarchical planning counterparts, while L2 Error remains statistically comparable (showing minimal variation).

Ablation for Model Design Details. We validate the effectiveness of model design details, including the usage of BEV adapter, Mode Sharing, and Ego-Motion interaction. Mode Sharing indicates whether the ego query and motion query utilize distinct or shared trajectory mode embeddings. In the Ego-Motion interaction ablation, we remove the bidirectional interaction between ego query and motion query, retaining only the one-way interaction where ego query cross-attends motion query, which destroys the integrity of interactions between instances. As shown in Tab. 5, BEV adapter significantly improves the planning performance, which decreases the collision rate by nearly 90%, demonstrating its clear advantage over using shared BEV features alone. The removal of either Mode Sharing or bidirectional Ego-Motion interaction maintains comparable L2 error in CogAD, but elevates the collision rate from 0.06% to 0.13% (116.6% rise).

4.3. Qualitative Results

In Fig. 4, we visualize the planning results of CogAD on nuScenes. Both the surrounding camera images and the prediction results in BEV space are provided accordingly. CogAD exhibits the capability to generate hierarchically dual-level multimodal trajectories. The rightmost column displays the highest-confidence trajectories corresponding to the top-10 intents, demonstrating that intent-level multimodal planning is diverse yet valid. The second column illustrates trajectory-level multi-mode planning conditioned on the dominant intent, with visualizations of the top-3 candidate trajectories. Multi-mode planning produces

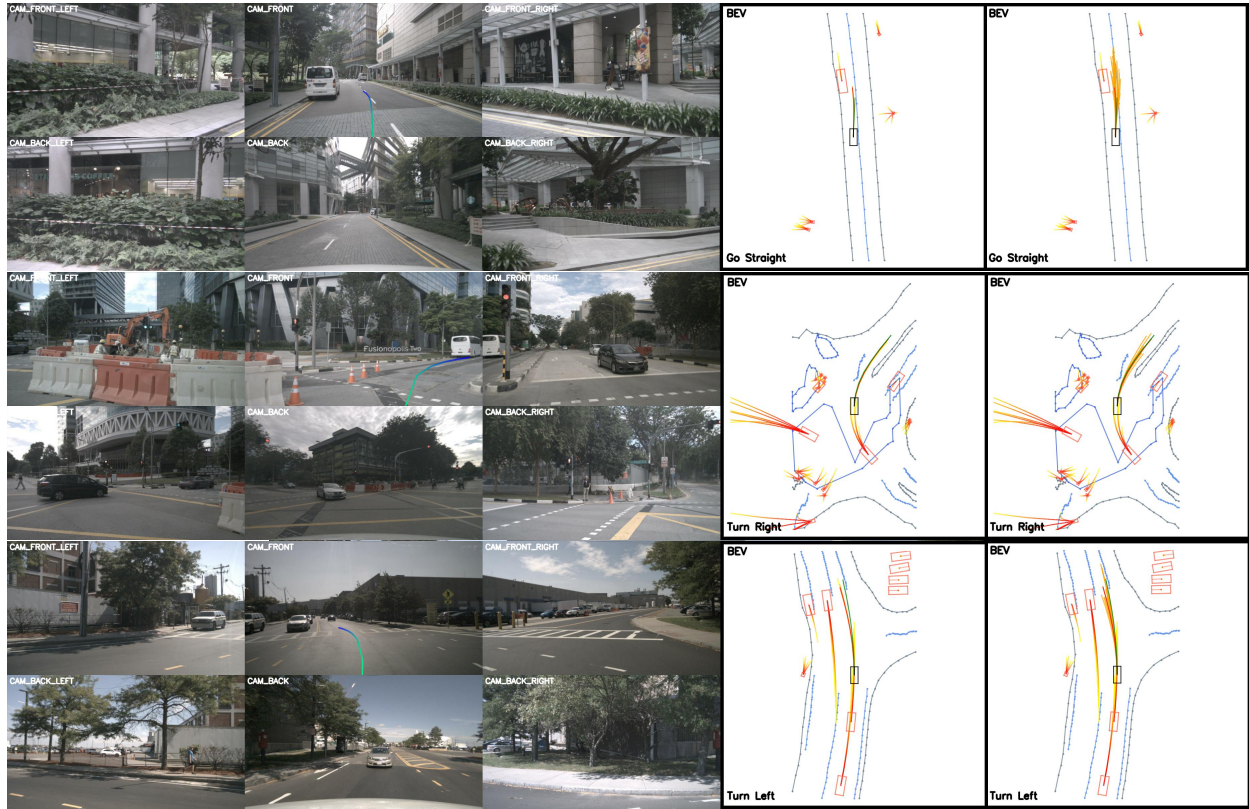


Figure 4. Qualitative results of CogAD on nuScenes. The 2nd column shows the Top-3 multi-mode trajectories of the highest-probability intent, with colors of red, orange, and yellow, respectively. The 3rd column displays the Top-10 multi-intent trajectories with the highest-probability mode. GT trajectory is drawn in green.



(a) Slow down and yield to merging traffic



(b) Change to the left lane to pass the bicycle ahead

Figure 5. Qualitative results of CogAD on Bench2Drive.

multiple refined trajectory options that maintain motion consistency. By combining multi-intent diversity and multi-mode refinement, CogAD generates multiple behaviorally diverse yet physically plausible planning candidates.

Our qualitative analysis demonstrates that CogAD exhibits behaviors analogous to human cognition, as evidenced by the highest-probability planned trajectory visualized in Fig. 5. These behaviors include strategic speed reduction for yielding, proactive lane changes for obstacle avoidance, and context-aware trajectory planning.

5. Conclusion

We present CogAD, a hierarchical end-to-end planning method inspired by human driving processes through the lens of cognitive science. For hierarchical perception, CogAD incorporates BEV-Instance interaction and cross-task instance interaction, which significantly enhance the ego vehicle’s scene understanding capabilities. For hierarchical planning, CogAD utilizes intent anchors and trajectory modes to ensure behavioral diversity and geometric precision in generated trajectories. Extensive evaluations demonstrate that our approach achieves state-of-the-art performance across both open-loop and closed-loop benchmarks.

References

- [1] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [2] Michael Bratman. Intention, plans, and practical reason. 1987.
- [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 multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [4] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *arXiv preprint arXiv:1910.05449*, 2019.
- [5] Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Vad2: End-to-end vectorized autonomous driving via probabilistic planning. *arXiv preprint arXiv:2402.13243*, 2024.
- [6] Xuesong Chen, Shaoshuai Shi, Tao Ma, Jingqiu Zhou, Simon See, Ka Chun Cheung, and Hongsheng Li. M3net: Multimodal multi-task learning for 3d detection, segmentation, and occupancy prediction in autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2275–2283, 2025.
- [7] Yuan Chen, Zi-han Ding, Ziqin Wang, Yan Wang, Lijun Zhang, and Si Liu. Asynchronous large language model enhanced planner for autonomous driving. In *European Conference on Computer Vision*, pages 22–38. Springer, 2024.
- [8] Zhili Chen, Maosheng Ye, Shuangjie Xu, Tongyi Cao, and Qifeng Chen. Ppad: Iterative interactions of prediction and planning for end-to-end autonomous driving. In *European Conference on Computer Vision*, pages 239–256. Springer, 2024.
- [9] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE transactions on pattern analysis and machine intelligence*, 45(11):12878–12895, 2022.
- [10] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4693–4700. IEEE, 2018.
- [11] Simon Doll, Niklas Hanselmann, Lukas Schneider, Richard Schulz, Marius Cordts, Markus Enzweiler, and Hendrik Lensch. Dualad: Disentangling the dynamic and static world for end-to-end driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14728–14737, 2024.
- [12] Guodong Du, Yuan Zou, Xudong Zhang, Zirui Li, and Qi Liu. Hierarchical motion planning and tracking for autonomous vehicles using global heuristic based potential field and reinforcement learning based predictive control. *IEEE Transactions on Intelligent Transportation Systems*, 24(8):8304–8323, 2023.
- [13] Daniel J Felleman and David C Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1):1–47, 1991.
- [14] 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.
- [15] Ziang Guo, Konstantin Gubernatorov, Selamawit Asfaw, Zakhar Yagudin, and Dzmitry Tsetserukou. Vdt-auto: End-to-end autonomous driving with vlm-guided diffusion transformers. *arXiv preprint arXiv:2502.20108*, 2025.
- [16] Deepti Hegde, Rajeev Yasarla, Hong Cai, Shizhong Han, Apratim Bhattacharyya, Shweta Mahajan, Litian Liu, Rishiek Garrepalli, Vishal M Patel, and Fatih Porikli. Distilling multimodal large language models for autonomous driving. *arXiv preprint arXiv:2501.09757*, 2025.
- [17] Joey Hong, Benjamin Sapp, and James Philbin. Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8454–8462, 2019.
- [18] 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 *European Conference on Computer Vision*, pages 533–549. Springer, 2022.
- [19] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023.
- [20] Yidong Huang, Jacob Sansom, Ziqiao Ma, Felix Gervits, and Joyce Chai. Drivlme: Enhancing llm-based autonomous driving agents with embodied and social experiences. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3153–3160. IEEE, 2024.
- [21] Zhijian Huang, Sihao Lin, Guiyu Liu, Mukun Luo, Chaoqiang Ye, Hang Xu, Xiaojun Chang, and Xiaodan Liang. Fuller: Unified multi-modality multi-task 3d perception via multi-level gradient calibration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3502–3511, 2023.
- [22] Xiaosong Jia, Zhenjie Yang, Qifeng Li, Zhiyuan Zhang, and Junchi Yan. Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. *arXiv preprint arXiv:2406.03877*, 2024.
- [23] Xiaosong Jia, Junqi You, Zhiyuan Zhang, and Junchi Yan. Drivetransformer: Unified transformer for scalable end-to-end autonomous driving. *arXiv preprint arXiv:2503.07656*, 2025.
- [24] 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 *Proceedings of the*

- IEEE/CVF International Conference on Computer Vision*, pages 8340–8350, 2023.
- [25] Bo Jiang, Shaoyu Chen, Bencheng Liao, Xingyu Zhang, Wei Yin, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Senna: Bridging large vision-language models and end-to-end autonomous driving. *arXiv preprint arXiv:2410.22313*, 2024.
- [26] Dongchan Kim, Gihoon Kim, Hayoung Kim, and Kunsoo Huh. A hierarchical motion planning framework for autonomous driving in structured highway environments. *IEEE Access*, 10:20102–20117, 2022.
- [27] Christof Koch and Shimon Ullman. Shifts in selective visual attention: towards the underlying neural circuitry. In *Matters of intelligence: Conceptual structures in cognitive neuroscience*, pages 115–141. Springer, 1987.
- [28] Stefan Lee, Senthil Purushwalkam Shiva Prakash, Michael Cogswell, Viresh Ranjan, David Crandall, and Dhruv Batra. Stochastic multiple choice learning for training diverse deep ensembles. *Advances in Neural Information Processing Systems*, 29, 2016.
- [29] 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. *arXiv preprint arXiv:2406.08481*, 2024.
- [30] 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. *arXiv preprint arXiv:2203.17270*, 2022.
- [31] Zhenxin Li, Kailin Li, Shihao Wang, Shiyi Lan, Zhiding Yu, Yishen Ji, Zhiqi Li, Ziyue Zhu, Jan Kautz, Zuxuan Wu, et al. Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation. *arXiv preprint arXiv:2406.06978*, 2024.
- [32] Zhiqi Li, Zhiding Yu, Shiyi Lan, Jiahao Li, Jan Kautz, Tong Lu, and Jose M Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14864–14873, 2024.
- [33] 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. *arXiv preprint arXiv:2411.15139*, 2024.
- [34] Yang Lu, Xin Xu, Xinglong Zhang, Lilin Qian, and Xing Zhou. Hierarchical reinforcement learning for autonomous decision making and motion planning of intelligent vehicles. *IEEE Access*, 8:209776–209789, 2020.
- [35] Jianbiao Mei, Yukai Ma, Xuemeng Yang, Licheng Wen, Xinyu Cai, Xin Li, Daocheng Fu, Bo Zhang, Pinlong Cai, Min Dou, et al. Continuously learning, adapting, and improving: A dual-process approach to autonomous driving. *arXiv preprint arXiv:2405.15324*, 2024.
- [36] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [37] Yao Qi, Binbing He, Rendong Wang, Le Wang, and Youchun Xu. Hierarchical motion planning for autonomous vehicles in unstructured dynamic environments. *IEEE Robotics and Automation Letters*, 8(2):496–503, 2022.
- [38] Kangan Qian, Zhikun Ma, Yangfan He, Ziang Luo, Tianyu Shi, Tianze Zhu, Jiayin Li, Jianhui Wang, Ziyu Chen, Xiao He, et al. Fasionad: Fast and slow fusion thinking systems for human-like autonomous driving with adaptive feedback. *arXiv preprint arXiv:2411.18013*, 2024.
- [39] Anand S Rao, Michael P Georgeff, et al. Bdi agents: from theory to practice. In *Iemas*, pages 312–319, 1995.
- [40] Nicholas Rhinehart, Kris M Kitani, and Paul Vernaza. R2p2: A reparameterized pushforward policy for diverse, precise generative path forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 772–788, 2018.
- [41] 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 *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 414–430. Springer, 2020.
- [42] Yinzhe Shen, Ömer Şahin Taş, Kaiwen Wang, Royden Wagner, and Christoph Stiller. Divide and merge: Motion and semantic learning in end-to-end autonomous driving. *arXiv preprint arXiv:2502.07631*, 2025.
- [43] Ziyang Song, Caiyan Jia, Lin Liu, Hongyu Pan, Yongchang Zhang, Junming Wang, Xingyu Zhang, Shaoqing Xu, Lei Yang, and Yadan Luo. Don’t shake the wheel: Momentum-aware planning in end-to-end autonomous driving. *arXiv preprint arXiv:2503.03125*, 2025.
- [44] Haisheng Su, Wei Wu, and Junchi Yan. Difs: Ego-centric fully sparse paradigm with uncertainty denoising and iterative refinement for efficient end-to-end autonomous driving. *arXiv preprint arXiv:2409.09777*, 2024.
- [45] Wenchao Sun, Xuewu Lin, Yining Shi, Chuang Zhang, Haoran Wu, and Sifa Zheng. Sparsedrive: End-to-end autonomous driving via sparse scene representation. *arXiv preprint arXiv:2405.19620*, 2024.
- [46] Yingqi Tang, Zhuoran Xu, Zhaotie Meng, and Erkang Cheng. Hip-ad: Hierarchical and multi-granularity planning with deformable attention for autonomous driving in a single decoder. *arXiv preprint arXiv:2503.08612*, 2025.
- [47] Ran Tian, Boyi Li, Xinshuo Weng, Yuxiao Chen, Edward Schmerling, Yue Wang, Boris Ivanovic, and Marco Pavone. Tokenize the world into object-level knowledge to address long-tail events in autonomous driving. *arXiv preprint arXiv:2407.00959*, 2024.
- [48] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Yang Wang, Zhiyong Zhao, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*, 2024.
- [49] Nam Dinh Van, Muhammad Sualeh, Dohyeong Kim, and Gon-Woo Kim. A hierarchical control system for autonomous driving towards urban challenges. *Applied Sciences*, 10(10): 3543, 2020.
- [50] Balakrishnan Varadarajan, Ahmed Hefny, Avikalp Srivastava, Khaled S Refaat, Nigamaa Nayakanti, Andre Cornman,

- Kan Chen, Bertrand Douillard, Chi Pang Lam, Dragomir Anguelov, et al. Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 7814–7821. IEEE, 2022.
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [52] Jingke Wang, Yue Wang, Dongkun Zhang, Yezhou Yang, and Rong Xiong. Learning hierarchical behavior and motion planning for autonomous driving. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2235–2242. IEEE, 2020.
- [53] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024.
- [54] Xinshuo Weng, Boris Ivanovic, Yan Wang, Yue Wang, and Marco Pavone. Para-drive: Parallelized architecture for real-time autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15449–15458, 2024.
- [55] Zebin Xing, Xingyu Zhang, Yang Hu, Bo Jiang, Tong He, Qian Zhang, Xiaoxiao Long, and Wei Yin. Goalflow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving. *arXiv preprint arXiv:2503.05689*, 2025.
- [56] Yi Xu, Yuxin Hu, Zaiwei Zhang, Gregory P Meyer, Siva Karthik Mustikovela, Siddhartha Srinivasa, Eric M Wolff, and Xin Huang. Vlm-ad: End-to-end autonomous driving through vision-language model supervision. *arXiv preprint arXiv:2412.14446*, 2024.
- [57] Tengju Ye, Wei Jing, Chunyong Hu, Shikun Huang, Lingping Gao, Fangzhen Li, Jingke Wang, Ke Guo, Wencong Xiao, Weibo Mao, et al. Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving. *arXiv preprint arXiv:2308.01006*, 2023.
- [58] 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 nuscenec. *arXiv preprint arXiv:2305.10430*, 2023.
- [59] Diankun Zhang, Guoan Wang, Runwen Zhu, Jianbo Zhao, Xiwu Chen, Siyu Zhang, Jiahao Gong, Qibin Zhou, Wenyuan Zhang, Ningzi Wang, et al. Sparsead: Sparse query-centric paradigm for efficient end-to-end autonomous driving. *arXiv preprint arXiv:2404.06892*, 2024.
- [60] Wenzhao Zheng, Ruiqi Song, Xianda Guo, Chenming Zhang, and Long Chen. Genad: Generative end-to-end autonomous driving. In *European Conference on Computer Vision*, pages 87–104. Springer, 2024.