

Hydra-NeXt: Robust Closed-Loop Driving with Open-Loop Training

Zhenxin Li^{1,2*} Shihao Wang^{3*} Shiyi Lan⁴ Zhiding Yu⁴ Zuxuan Wu^{1,2†} Jose M. Alvarez⁴

¹Institute of Trustworthy Embodied AI, Fudan University

²Shanghai Collaborative Innovation Center of Intelligent Visual Computing

³The Hong Kong Polytechnic University ⁴NVIDIA

Abstract

End-to-end autonomous driving research currently faces a critical challenge in bridging the gap between open-loop training and closed-loop deployment. Current approaches are trained to predict trajectories in an open-loop environment, which struggle with quick reactions to other agents in closed-loop environments and risk generating kinematically infeasible plans due to the gap between open-loop training and closed-loop driving. In this paper, we introduce Hydra-NeXt, a novel multi-branch planning framework that unifies trajectory prediction, control prediction, and a trajectory refinement network in one model. Unlike current open-loop trajectory prediction models that only handle general-case planning, Hydra-NeXt further utilizes a control decoder to focus on short-term actions, which enables faster responses to dynamic situations and reactive agents. Moreover, we propose the Trajectory Refinement module to augment and refine the planning decisions by effectively adhering to kinematic constraints in closed-loop environments. This unified approach bridges the gap between open-loop training and closed-loop driving, demonstrating superior performance of 65.89 Driving Score (DS) and 48.20% Success Rate (SR) on the Bench2Drive dataset without relying on external experts for data collection. Hydra-NeXt surpasses the previous state-of-the-art by 22.98 DS and 17.49 SR, marking a significant advancement in autonomous driving. Code will be available at <https://github.com/woxihuanjiangguo/Hydra-NeXt>.

1. Introduction

End-to-end autonomous driving (E2E AD) [7, 9, 20, 22, 23, 27, 31, 41, 44, 45, 47] has emerged as a trending alternative to traditional perception-planning pipelines. Current research in E2E AD is divided into open-loop and closed-loop environment-oriented approaches.

An **open-loop environment** [3, 4, 12, 22] refers to an autonomous system’s performance being assessed without

*Work done during internship at NVIDIA.

†Corresponding author.

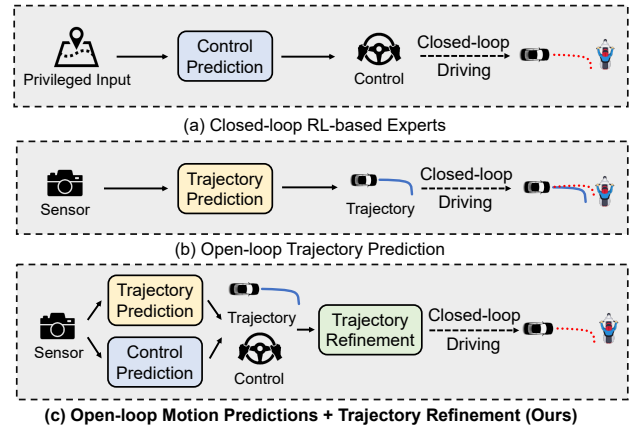


Figure 1. **Different Paradigms of Autonomous Driving.** Closed-loop RL-based experts learn from environment interactions, ensuring responsiveness to agents and kinematic feasibility by predicting control signals. End-to-end methods use open-loop trajectory prediction, often neglecting ego-agent interaction and kinematic constraints. Our approach integrates trajectory prediction, control prediction, and a trajectory refinement module to bridge the gap between open-loop training and closed-loop driving.

environmental feedback. The main advantage of this approach is its simplicity and efficiency, as it avoids the complexity of real-time feedback, making it cost-effective and suitable for large-scale training and testing. Benchmarks like NAVSIM [12] provide controlled environments where agents follow predefined trajectories, allowing for precise evaluation of metrics like collision rates, progress, and traffic-rule following. However, the lack of reactive behaviors and dynamic interactions with other agents or changing conditions makes it difficult to deploy models trained in open-loop environments in real-world driving.

A **closed-loop environment** in autonomous driving involves real-time feedback between the agent’s actions and the environment, typically via a simulator [1, 4, 14]. The primary advantage of closed-loop methods is their ability to simulate dynamic, interactive environments, allowing for more realistic testing of driving policies. However, they face significant challenges: These systems often rely on synthetic

simulator data, leading to a domain gap between training and real-world deployment. For instance, the previous state-of-the-art E2E AD method [24] depends heavily on reinforcement learning (RL) teacher feature embeddings, unavailable in real-world data. In addition, the driving behavior learned may deviate from human patterns [15, 30], and evaluation metrics focus mainly on collision rates, overlooking other important factors such as smoothness and driving comfort [26].

The gap between open-loop and closed-loop environment has led to different approaches in each domain, and this separation continues to grow. Fig. 1 shows that closed-loop methods learn to control through reinforcement learning in simulations [24, 25, 30, 50], while open-loop methods focus on imitating expert trajectories [7, 22, 27]. The gap between the two domains stem from their training methods and output representations, which can be potentially addressed through two approaches. We could either adapt models trained in an open-loop environment to a closed-loop environment or bring closed-loop models to open environments. We argue that the first option, *From Open-loop to Closed-loop*, is more promising due to the availability of massive real-world data. In contrast, the second option, *From Closed-loop to Open-loop*, faces sim-to-real domain gaps [19, 21, 29], which pose the biggest challenge to deploy a model in a real car. Hence, we focus on addressing the limitations of open-loop models when driving in closed-loop environments, which involve **handling reactive agents** and **adhering to kinematic constraints**.

As the state-of-the-art open-loop E2E method on the NAVSIM benchmark [12], Hydra-MDP shows promising results for non-reactive agents, so the challenge is to extend this ability to reactive agents. Non-reactive agents maintain their behavior even if the ego vehicle deviates from the ground truth trajectory. In contrast, reactive agents adjust in response to the ego’s predicted trajectory, potentially causing the E2E AD model’s collision predictions to fail. The general issue caused by open-loop training is that most open-loop metrics [4, 12, 22] focus on waypoints rather than control signals. While waypoints are useful for general navigation, they do not account for immediate changes in the behavior of other agents, leading to delays in reaction time when unexpected events occur. In contrast, control signals directly influence the vehicle’s actions, such as braking or steering, allowing for faster adjustments in response to sudden changes. Although existing literature [47] also predicts control with waypoints directly, it performs poorly in interactive situations [26], which may stem from a lack of multi-modal decisions and insufficient exploration of their nuanced relationships. Further, the predicted waypoints may not satisfy kinematic constraints in the closed-loop environment, which can cause compounding errors [11] and result in unsafe behaviors in dynamic scenarios.

Therefore, we introduce Hydra-NeXt, featuring a multi-

branch planning framework. Hydra-NeXt consists of Multi-head Motion Decoders for trajectory and control prediction, and a Trajectory Refinement network for kinematics-based proposal selection. In addition to the trajectory decoder to handle general-case planning [31], the control decoder enhances the capability for short-term actions. Additionally, the Trajectory Refinement module applies kinematic constraints to improve predictions from both decoders, effectively combining and refining the planning decisions.

Our contributions can be summarized as follow:

1. We propose Hydra-NeXt, which is the first framework to unify control, trajectory prediction, and trajectory refinement. Hydra-NeXt demonstrates robust closed-loop driving performance with only open-loop learning.
2. We benchmark Hydra-NeXt against other E2E AD solutions on the Bench2Drive dataset [26] based on the simulator CARLA v2 [1]. Hydra-NeXt outperforms previous state-of-the-art methods by a significant margin (+22.9 Driving Score, +17.5 Success Rate) by the CARLA v2 evaluation protocol without relying on external experts for data collection. Hydra-NeXt also generalizes well to real-world data, achieving a new state-of-the-art on the open-loop planning benchmark NAVSIM [12].

2. Related Work

2.1. Open-loop End-to-end Autonomous Driving

To explore end-to-end planning with real-world vehicle data [3], researchers have integrated modularized neural networks into end-to-end fully-differentiable stacks [20, 22, 27, 45]. This approach quickly gained significant interest in the AV community as it shows a promising direction to scalable autonomy. However, recent works [11, 12, 32] discover heavy biases of current E2E planning datasets [12, 32], the imitation learning paradigm [11, 32], and evaluation protocols [12, 32], rendering open-loop E2E methods as unreliable. To address these issues, a more advanced benchmark NAVSIM [12] is proposed to filter homogeneous planning data and benchmark with a series of rule-based metrics for collision avoidance, map adherence, and more. The state-of-the-art open-loop planner Hydra-MDP [31], which learns these rule-based metrics via knowledge distillation, achieves greater reliability than imitation-based E2E planners. Building on Hydra-MDP, we develop a model suited for dynamic scenarios, transitioning from open-loop trajectory planning to a closed-loop framework using control prediction and kinematics-based trajectory refinement.

2.2. Closed-loop Driving Benchmarks and Policies

Compared with open-loop E2E autonomous driving, closed-loop driving has a longer history. Benchmarks such as nuPlan [4], CARLA v1 [14], CARLA v2 [1], and Bench2Drive [26] focus on different aspects of driving poli-




Figure 2. **Roadmap from Hydra-MDP to Hydra-NeXt**. DriveAdapter [24] was the previous state-of-the-art method on the Bench2Drive benchmark.

cies, from trajectory planning [4] to E2E driving [1, 14, 26]. Recently, CARLA v2 [1] poses significant challenges as it requires the handling of interactive and dynamic scenarios. Bench2Drive [26] optimizes the variances produced by the CARLA v2 evaluation protocol and collects data from an RL-based expert [30] on diverse, shorter routes, making it suitable to assess policies in varied conditions.

On these benchmarks, expert-level policies include rule-based path planners [2, 11] and RL-based policies [15, 30, 50]. These policies take privileged inputs such as ground-truth perception and map data, while E2E AD methods [7, 9, 23, 41, 44, 47] only use sensor observations. These E2E policies, as well as previous open-loop E2E methods, fall behind expert-level policies when it comes to dynamic and interactive scenarios [1]. Moreover, Bench2Drive shows that these methods rely on expert feature embeddings [24] for better performance. In our paper, we aim to close the gap between the two worlds: open-loop planning and closed-loop driving by enhancing a top open-loop planner [31] with control and kinematic components.

2.3. Diffusion-based Driving Policies

Diffusion policies [6, 8, 48] have been used widely for generating robot behaviors. They prove effective in capturing multi-modal action distributions and generating smooth trajectories. In autonomous driving, recent works leverage diffusion models to predict trajectory waypoints given privileged input of surrounding traffic scenes [16, 28, 46, 49]. Our approach differs by employing the diffusion policy to gener-

ate multiple smooth and high-frequency control sequences, which serve as additional proposals for trajectory refinement.

3. Trajectory Decoder: Hydra-MDP

Before introducing Hydra-NeXt, we start from Hydra-MDP [31], a multi-modal planner with a perception network and a trajectory decoder, which learns from both human demonstrations and rule-based open-loop metrics.

Perception Network: Given front-view and back-view images, we use an image backbone to extract multi-view features. These features are flattened into a sequence of environment tokens F_{env} to represent the surroundings. To focus on motion planning, we apply a minimal design in the perception network without auxiliary perception tasks [7, 9, 22, 27, 45], although they can facilitate planning.

Trajectory Decoder: The trajectory decoder π_{traj} generates a trajectory T for routing based on environment tokens F_{env} and a discrete trajectory vocabulary \mathcal{V} [7, 37, 38]. Following Hydra-MDP [31], we create a trajectory vocabulary \mathcal{V} with 4096 discrete trajectories, embed them into latent queries Q_{traj} , and attend to environment tokens F_{env} in a transformer decoder [43]. With only imitation learning, this approach achieves 49.0 DS as shown in Fig. 2, outperforming the previous state-of-the-art DriveAdapter by 6.1 DS. Further, we implement an open-loop metric system on Bench2Drive using the following metrics:

- **Collision:** The collision metric checks if the ego vehicle intersects with other agents in the Bird’s-Eye View (BEV) space [12, 22, 27, 32]. We compute collisions at a higher frequency (*i.e.* 10Hz) by interpolating trajectory waypoints following NAVSIM [12].
- **Soft Lane Keeping:** To ensure lane adherence while allowing lane changing, we design a soft lane keeping metric that bounds the ego’s driving area following [27]. We penalize trajectories with excessive angular differences between trajectory segments and lane segments.
- **Ego Progress:** The Ego Progress metric is included to discourage passive driving behaviors and promote driving progress. Similar to [4, 12], we project the trajectory waypoints onto the lane segments that the expert travels through, and normalize the ego distance by the expert’s.

These metrics are used as rule-based teachers in Hydra-MDP. Fig. 2 shows the performance enhancement brought by rule-based heads, which appear to be limited compared with observations on NAVSIM [12]. This is likely caused by randomly disappearing agents in Bench2Drive¹, which confuse the rule-based heads. *e.g.* A trajectory that extends far into the future may be labeled as safe if its future waypoints do not intersect with a disappeared agent, even though the agent is currently visible and close to the ego vehicle. In Sec. 4, we

¹<https://github.com/Thinklab-SJTU/Bench2Drive/issues/31>

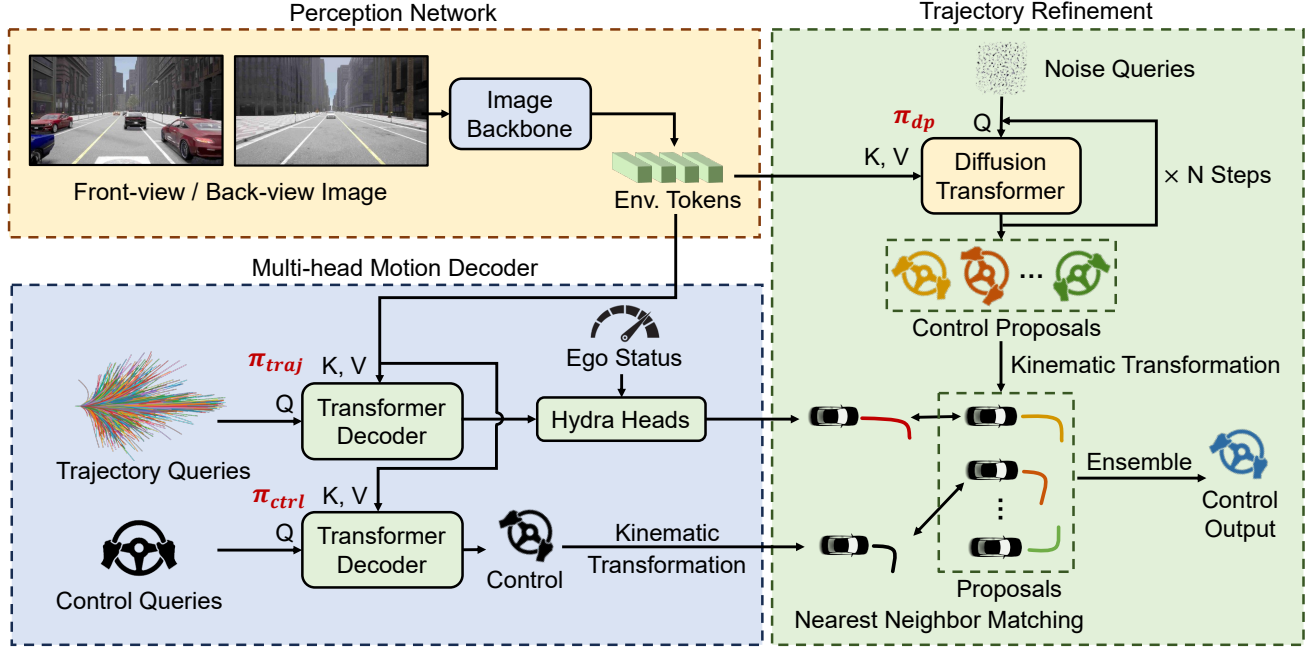


Figure 3. **The Overall Architecture of Hydra-NeXt**. The Perception Network first processes raw sensor observations into environment tokens. The Multi-head Motion Decoder (Trajectory Decoder π_{traj} and Control Decoder π_{ctrl}) is responsible for generating a trajectory and a control tuple based on these tokens, which are refined by the Trajectory Refinement Module π_{dp} into final control signals.

propose Hydra-NeXt, an extended version of Hydra-MDP addressing the limitations in closed-loop driving.

4. Hydra-NeXt

In this section, we elaborate on Hydra-NeXt, an E2E framework with robust closed-loop abilities. Apart from the Trajectory Decoder π_{traj} , Hydra-NeXt has two other policies: **Control Decoder** π_{ctrl} , and **Trajectory Refinement** π_{dp} .

4.1. Control Decoder π_{ctrl}

Our goal is to enable quick responses to reactive agents through direct control output, instead of relying on trajectory waypoints for rapid decisions. This is because waypoints are typically transformed into control signals through specific controllers such as a PID Controller or Model Predictive Control [5], which can lead to errors during path following. Further, TCP [47] points out that ensembling the trajectory and the control signal can boost driving performance, but it only considers a single modality for both the trajectory and the control prediction branches. This limitation makes it challenging to address uncertainties in environments [7] and susceptible to interpolate between different action modes [23].

Based on these findings, we add a second classification-based control decoder π_{ctrl} to generate control signals ($C^1, \dots, C^{t_{ctrl}}$) for t_{ctrl} timesteps. In the CALRA simulator [1, 14], each control signal is a tuple of (*brake*, *throttle*, *steer*) used to direct the vehicle. Meanwhile, we incorporate the idea of discretization into the con-

trol decoder π_{ctrl} to handle uncertainty, following the practice in the RL-based driving policy Think2Drive [30]. The use of learning-based control prediction sidesteps the need to convert trajectories into control signals through traditional controllers, enabling faster responses to reactive agents in the closed-loop environment.

In particular, we first randomly initialize t_{ctrl} control queries Q_{ctrl} , which correspond to the t_{ctrl} current and future steps for which we want to predict control signals. Similarly, Q_{ctrl} attends to F_{env} in a transformer decoder. After this, Q_{ctrl} is processed by three separate MLP layers for each control signal (*i.e.* *brake*, *throttle*, *steer*). Since the expert data collected by Think2Drive only consists of discrete control signals, the control decoder can be easily trained using cross-entropy loss functions without further discretizing the control signal ground truths. We select the control tuples with the highest likelihood at each step ($C^1, \dots, C^{t_{ctrl}}$) as the final output of π_{ctrl} . The design choices of the control decoder are discussed in Sec. 5.4, such as architectures and loss functions.

4.2. Trajectory Refinement π_{dp}

Another important aspect of closed-loop driving is to ensure the smoothness of driving. Benchmarks like nuPlan [4] and NAVSIM [12] post-process trajectories using an LQR tracker [42] and a kinematic bicycle model [39, 40] to adhere to kinematic constraints. Nevertheless, this idea has been neglected by existing open-loop E2E AD methods, which mainly focuses on cloning trajectory waypoints without con-

sidering kinematic constraints in closed-loop environments. This oversight becomes particularly critical when the vehicle experiences control loss in closed-loop environments [1, 26] or encounters large compounding errors [11] when quickly recovering to predicted waypoints is difficult.

To address this issue, we introduce the Trajectory Refinement network to augment and then combine the planning output T and C while following kinematic constraints. Inspired by PDM-Lite [2], we use a kinematic bicycle model [39, 40] that operates on high-frequency control signals for fine-grained waypoint rollouts [2]. Unlike predicted waypoints from π_{traj} , these waypoints are kinematically feasible. However, we empirically find that existing control policies struggle to produce a long and smooth control sequence at a high frequency (see Sec. 5.4). Therefore, we resort to the Diffusion Policy [8], which proves capable of generating smooth and diverse robot behaviors. During training, a diffusion transformer π_{dp} is trained to predict noise added to the ground-truth control sequence following the common practice of Denoising Diffusion Probabilistic Models (DDPMs) [18]. During inference, π_{dp} gradually denoises N random noise queries Q_{dp} into smooth action sequences $\{(\tilde{C}_i^1, \dots, \tilde{C}_i^{t_{dp}})\}_{i=1}^N$ through numerous denoising iterations. Each proposal spans across t_{dp} timesteps.

With these smooth control sequence proposals, we devise a process named Nearest Neighbor Matching for choosing the best proposal in a kinematically feasible way. Algorithm 1 depicts the nearest neighbor matching process for proposal selection in detail. Specifically, we first use a kinematic bicycle model to transform control sequences from π_{dp} and π_{ctrl} into trajectory waypoints. After the transformation, we select two nearest control candidates to match current candidates based on L2 distances. The new control candidates can be viewed as a kinematically feasible version of the predictions given by the Multi-head Motion Decoder. Finally, we ensemble the candidates \mathcal{C} into the final control signal C^* by averaging the throttle and steer values, which is a simplified ensembling method from TCP [47]. The brake is set to 1 if the condition $\sum_{C \in \mathcal{C}} C.brake \geq \tau$ holds; otherwise, it is set to 0. Note that the brake values produced by the candidates are binary and τ is a predefined brake threshold.

5. Experiments

5.1. Dataset and Metrics

We use the E2E driving benchmark Bench2Drive [26] for training and evaluating Hydra-NeXt. The training data in this paper uses the official training dataset of Bench2Drive, which includes 2 million frames encompassing 44 interactive scenarios. These data are collected by Think2Drive [30], an RL-based expert model on CARLA v2 [1]. For evaluation, the E2E AD model is deployed in the CARLA simulator [14] to perform closed-loop driving on 220 short routes designed

Algorithm 1: Nearest Neighbor Matching

Input: trajectory T , control sequence $(C^1, \dots, C^{t_{ctrl}})$, N control proposals $\{(\tilde{C}_i^1, \dots, \tilde{C}_i^{t_{dp}})\}_{i=1}^N$
Output: control candidates \mathcal{C}
Definitions : KB: Kinematic Bicycle Model
PID: PID controller
L₂: L₂ Distance

```

1  $\mathcal{C} \leftarrow \{\text{PID}(T), C^1\}$ 
2  $T_{ctrl} \leftarrow \text{KB}(C^1, \dots, C^{t_{ctrl}})$ 
3 for  $i \leftarrow 1$  to  $N$  do
4    $T_{dp}^i \leftarrow \text{KB}(\tilde{C}_i^1, \dots, \tilde{C}_i^{t_{dp}})$ 
5 end
6  $i \leftarrow \arg \min_i \text{L}_2(T_{dp}^i, T)$ 
7  $j \leftarrow \arg \min_j \text{L}_2(T_{dp}^j, T_{ctrl})$ 
8  $\mathcal{C} \leftarrow \mathcal{C} \cup \{\tilde{C}_i^1, \tilde{C}_j^1\}$ 

```

by Bench2Drive. These short routes assess the AD model’s abilities on different scenarios, while ensuring low variance in the final score.

Bench2Drive includes several metrics, including Driving Score (DS), Success Rate (SR), Efficiency, Comfort, and Multi-ability Results. The calculation of DS can be based on two evaluation protocols: CARLA v2 and Bench2Drive. The former accumulates the minimum speed infractions into the DS, whereas the Bench2Drive protocol separates this into the metric Efficiency. Meanwhile, the latter protocol relaxes the time constraint of the closed-loop evaluation, which reduces the difficulty for the model to complete the routes. Without further notations, we default to the CARLA v2 evaluation protocol in ablation studies to reflect the comprehensive performance of the method using the DS metric.

Furthermore, we evaluate Hydra-NeXt on the real-world NAVSIM Benchmark [12], which evaluates a 4-second trajectory using open-loop metrics: No at-fault Collisions (NC), Drivable Area Compliance (DAC), Time-to-collision (TTC), Comfort (C), and Ego Progress (EP). The PDM score (PDMS) is an aggregate of these sub-metrics.

5.2. Implementation Details

The implementation of Hydra-NeXt is largely consistent with open-loop E2E baselines [22, 27] on Bench2Drive. First, we train Hydra-NeXt on the training data for 20 epochs with a total batch size of 256, using 8 NVIDIA V100 GPUs. The AdamW [35] optimizer is used with the Cosine Annealing Scheduler [36] at a learning rate of 2×10^{-4} and a weight decay of 0.01. During training, data augmentations such as random cropping and photometric distortion are applied to the input images, which are first resized to 800×450 .

Method	Sensors	Open-loop Metric L2 (meter) ↓	Closed-loop Metrics (CARLA v2)	
			Driving Score ↑	Success Rate (%) ↑
UniAD-Base [22]	6 Cameras	0.73	37.72	9.54
VAD [27]	6 Cameras	0.91	39.42	10.00
Hydra-MDP [31]	Front/Back Cameras	0.82	52.80	30.73
TCP* [47]	Front Camera	1.70	23.63	7.72
TCP-Ctrl* [47]	Front Camera	-	18.63	5.45
TCP-Traj* [47]	Front Camera	1.70	36.78	26.82
ThinkTwice* [25]	6 Cameras	0.95	39.88	28.14
DriveAdapter* [24]	6 Cameras	1.01	42.91	30.71
Hydra-NeXt	Front/Back Cameras	0.92	65.89	48.20

Table 1. **Open-loop and Closed-loop Performance of E2E-AD Methods on the Bench2Drive Benchmark with the CARLA v2 Evaluation Protocol.** The CARLA v2 protocol calculates the Driving Score (DS) by aggregating all infractions multiplicatively, including minimum speed infractions. * The model benefits from expert feature distillation.

Method	Sensors	Closed-loop Metrics (Bench2Drive)			
		Driving Score ↑	Success Rate (%) ↑	Efficiency ↑	Comfort ↑
UniAD-Base [22]	6 Cameras	45.81	16.36	129.21	43.58
VAD [27]	6 Cameras	42.35	15.00	157.94	46.01
Hydra-MDP [31]	Front/Back Cameras	59.95	29.82	186.83	18.62
TCP* [47]	Front Camera	40.70	15.00	54.26	47.80
TCP-Ctrl* [47]	Front Camera	30.47	7.27	55.97	51.51
TCP-Traj* [47]	Front Camera	59.90	30.00	76.54	18.08
ThinkTwice* [25]	6 Cameras	62.44	31.23	69.33	16.22
DriveAdapter* [24]	6 Cameras	64.22	33.08	70.22	16.01
Hydra-NeXt	Front/Back Cameras	73.86	50.00	197.76	20.68

Table 2. **Closed-loop Performance of E2E-AD Methods on the Bench2Drive Benchmark with the Bench2Drive Evaluation Protocol.** * The model benefits from expert feature distillation.

Hydra-NeXt employs an ImageNet-pretrained [13] ResNet-50 [17] as the image backbone to extract front-view and back-view image features. Following VAD [27], the trajectory decoder π_{traj} predicts a 3-second trajectory at 2Hz. The frequencies of π_{ctrl} and π_{dp} are set to 2Hz and 10Hz by default, which will be further analyzed in Sec. 5.4. The ego status feature includes the current longitudinal velocity and a one-hot navigation command. In Trajectory Refinement, the brake threshold τ is set to half of the candidates set size 2 for moderate behaviors, while the kinematic bicycle model is based on the implementation of PDM-Lite [2] and the denoising timestep of π_{dp} is set to 100 [8]. For NAVSIM, we extend Hydra-MDP [31] by incorporating acceleration and steering rate prediction as a substitute for control signals. Details can be found in the appendix.

5.3. Main Results

As shown in Tab. 1 and Tab. 2, Hydra-NeXt surpasses all E2E methods on Bench2Drive in key metrics such as the DS and SR distinctively. Specifically, under the CARLA v2 protocol, Hydra-NeXt outperforms the previous state-of-the-art, DriveAdapter [24], by 22.98 DS and 17.49% SR. Under the Bench2Drive protocol, it achieves improvements of 9.64

DS and 16.92 SR, despite DriveAdapter utilizing expert features from Think2Drive [30]. Additionally, Hydra-NeXt demonstrates higher efficiency than the baselines, especially those using expert features. Nevertheless, Hydra-NeXt falls behind TCP [47] in terms of Comfort, which is likely due to frequent braking during interactive situations.

For the Multi-ability Results shown in Tab. 3, Hydra-NeXt shows an improvement of 11.14% over DriveAdapter in the average performance. It is worth noting that Hydra-NeXt exhibits superior performance in interactive scenarios such as merging (+11.18%), overtaking (+38.06%), and emergency braking (+12.91%), which indicates its proficiency in handling reactive agents. However, Hydra-NeXt falls behind DriveAdapter in scenarios where the vehicle should adhere to traffic signs. This phenomenon may be caused by causal confusion [24] when encountering mixed expert behaviors before traffic signs, and is worth future investigations. Notably, all methods achieve less satisfactory results in yielding to specialized vehicles (*e.g.* ambulances), which may be caused by the scarcity of such training data.

Moreover, Tab. 4 shows that Hydra-NeXt achieves a new state-of-the-art on the NAVSIM Benchmark. Though NAVSIM only evaluates the predicted trajectory from π_{traj} ,

Method	Ability (%) \uparrow					
	Merging	Overtaking	Emergency Brake	Give Way	Traffic Sign	Mean
UniAD-Base [22]	14.10	17.78	21.67	10.00	14.21	15.55
VAD [27]	8.11	24.44	18.64	20.00	19.15	18.07
Hydra-MDP [31]	19.23	20.00	45.00	50.00	41.58	35.16
TCP* [47]	16.18	20.00	20.00	10.00	6.99	14.63
TCP-Ctrl* [47]	10.29	4.44	10.00	10.00	6.45	8.23
TCP-Traj* [47]	8.89	24.29	51.67	40.00	46.28	34.22
ThinkTwice* [25]	27.38	18.42	35.82	50.00	54.23	37.17
DriveAdapter* [24]	28.82	26.38	48.76	50.00	57.21	42.08
Hydra-NeXt	40.00	64.44	61.67	50.00	50.00	53.22

Table 3. Multi-Ability Results of E2E-AD Methods on the Bench2Drive Benchmark. * denotes expert feature distillation.

Method	NC \uparrow	DAC \uparrow	TTC \uparrow	EP \uparrow	C \uparrow	PDMS \uparrow
Transfuser [9]	97.7	92.8	92.8	79.2	100	84.0
Hydra-MDP [31]	98.3	96.0	94.6	78.7	100	86.5
DiffusionDrive [33]	98.2	96.2	94.7	82.2	100	88.1
Hydra-NeXt	98.1	97.7	94.6	81.8	100	88.6

Table 4. Performance of E2E-AD Methods on NAVSIM.

the extension of learning targets in π_{ctrl} and π_{dp} can benefit the trajectory prediction, leading to an improvement of 2.1 PDMS over Hydra-MDP [31].

5.4. Ablation Study

Ablation on Different Components and their Designs.

Fig. 2 shows the detailed ablations of each component in Hydra-NeXt. Operating on perspective image tokens, Hydra-MDP achieves a 52.8 DS, which already surpasses DriveAdapter by around 10 DS. Nevertheless, the trajectory prediction remains insufficient when facing more complex agent interactions, which promotes the utilization of a control prediction network π_{ctrl} for rapid reactions. The output of π_{ctrl} is ensembled with the trajectory in a similar fashion to Sec. 4.2, where the brake threshold is correspondingly set to 1 given two proposals. This gives us a large increase of 12.7% in the SR, indicating fewer collisions with reactive agents. Furthermore, using multi-step control predictions into future frames as in [47] provides auxiliary supervision (+1.8 DS). Applying a binary focal loss [34] for brake prediction balances the training (+1.8% SR) and simply post-processing the final control output, such as slowing down during turns, leads to further enhancements (+2.5% SR). Finally, we build the Trajectory Refinement network on top of the previous network. Involving the Diffusion Policy and applies a straightforward ensembling leads to an increase of around 3 DS, while applying the Nearest Neighbor Matching (Algorithm 1) to follow kinematic constraints boosts the final performance to 65.9 DS and 48.2% SR.

Design of Trajectory Refinement. To illustrate the necessity of using the Diffusion Policy π_{dp} as the proposal generator, we first replace π_{dp} with a discrete control decoder π'_{ctrl}

Policy in Traj. Refine.	Driving Score \uparrow	Success Rate (%) \uparrow
π'_{ctrl}	60.15	42.25
π_{dp}	65.89	48.20

Table 5. Ablation of the Policy in Trajectory Refinement.

π_{ctrl}	π_{dp}	Driving Score \uparrow	Success Rate (%) \uparrow
2 Hz	2 Hz	61.95	39.05
2 Hz	10 Hz	65.89	48.20
10 Hz	10 Hz	64.66	44.45

Table 6. Ablation of the Prediction Frequency of the Control Decoder and the Diffusion Policy.

Proposal Number N	Driving Score \uparrow	Success Rate (%) \uparrow
5	65.52	46.36
10	65.89	48.20
20	66.09	46.08

Table 7. Ablation of the Proposal Number N in Trajectory Refinement.

with a similar architecture to π_{ctrl} , which can also generate N proposals based on top- N confidence scores. This leads to a degradation of 5.78 DS as shown in Tab. 5, , which potentially stems from the poor ability of π'_{ctrl} to generate long and smooth control sequences. Further, we investigate the ability of π_{dp} in modeling high-frequency control sequences. Tab. 6 indicates a performance degradation when using a low-frequency π_{dp} . On the contrary, a low-frequency π_{ctrl} is better than a high-frequency one. This echos the previous finding and proves the effectiveness of employing π_{dp} as the generator. Finally, Tab. 7 examines the impact of varying the number of proposals. Increasing the proposal number N from 5 to 20 leads to saturating performance on DS, while SR mildly fluctuates with a larger N . The fluctuation may result from over-confident matching between π_{dp} and the other decoders, while we can potentially benefit from small differences between their predictions for exploration.

5.5. Runtime Efficiency and Visualization

Runtime Efficiency. Tab. 8 compares the runtime efficiency of Hydra-NeXt with E2E baselines. Since π_{traj} performs

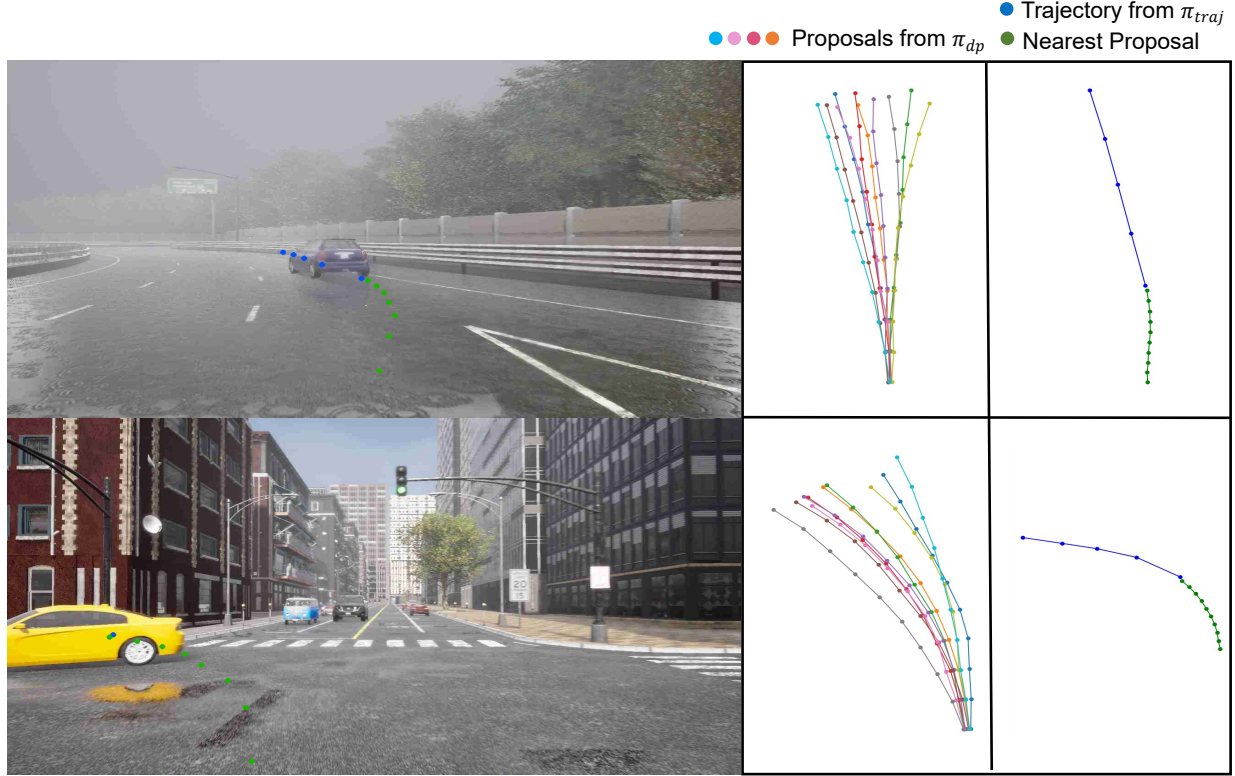


Figure 4. **Visualizations of Trajectory Refinement.** The figure shows the front-view image, trajectories transformed from control proposals, and the selected nearest proposal to the predicted trajectory.

Method		Latency (ms)
UniAD-Base [22]		555.6
VAD [27]		224.3
DriveAdapter [24]		894.0

Method	π_{traj}	π_{ctrl}	π_{dp}	\mathcal{F}	Latency (ms)
Hydra-NeXt	✓				245.4
	✓	✓			250.6
	✓	✓	✓		998.6
	✓	✓	✓	✓	528.3

Table 8. **Analysis of Runtime Efficiency.** The latency of Hydra-NeXt and VAD are benchmarked on an NVIDIA RTX 3090, while UniAD and DriveAdapter are on NVIDIA Tesla A100 and A6000, respectively. \mathcal{F} refers to Flash-attention [10] for acceleration.

planning based on perspective image tokens without constructing an explicit BEV feature [22], it achieves higher efficiency than the previous UniAD. The control decoder π_{ctrl} is also lightweight. Nevertheless, involving the Diffusion Policy π_{dp} for refining trajectories results in an efficiency degradation, which is an inherent limitation of the iterative denoising process [8] and deserves further optimization.

Visualization. Fig. 4 depicts the process of Trajectory Refinement in a highway merging scenario and an unprotected left turn. Multiple control sequences are proposed by π_{dp} and transformed into trajectories. The best-matching pro-

posal can smooth the trajectory produced by π_{traj} with a large curvature, ensuring kinematic feasibility.

6. Conclusion

In this paper, we first analyzed the development of current end-to-end autonomous driving systems, focusing on the discrepancy between open-loop and closed-loop environments and the challenges of bringing an open-loop trained model to closed-loop driving. These challenges include handling reactive agents and adhering to kinematic constraints. To address these issues, we propose Hydra-NeXt, a unified approach for trajectory and control prediction. By integrating Hydra-MDP with the Multi-head Motion Decoder and the Trajectory Refinement network, Hydra-NeXt obtains stronger closed-loop driving performance. Specifically, Hydra-NeXt demonstrates superior performance of 65.89 Driving Score (DS) and 48.20% Success Rate (SR) on the closed-loop driving benchmark Bench2Drive and outperforms previous state-of-the-art methods by a significant margin. Hydra-NeXt also achieves a new state-of-the-art on the real-world E2E planning benchmark NAVSIM.

Acknowledgement This work was supported in part by the National Natural Science Foundation of China (Grant 62472098). This work was supported by the Science and Technology Commission of Shanghai Municipality (No. 24511103100).

References

- [1] Carla: Carla autonomous driving leaderboard (2022). <https://leaderboard.carla.org/>, 2022. 1, 2, 3, 4, 5
- [2] Jens Beißwenger. PDM-Lite: A rule-based planner for carla leaderboard 2.0. <https://github.com/OpenDriveLab/DriveLM/blob/DriveLM-CARLA/docs/report.pdf>, 2024. 3, 5, 6
- [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 *CVPR*, pages 11621–11631, 2020. 1, 2
- [4] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. *arXiv preprint arXiv:2106.11810*, 2021. 1, 2, 3, 4
- [5] Eduardo F. Camacho and Carlos Bordons Alba. *Model Predictive Control*. Springer Science & Business Media, 2013. 4
- [6] Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Applications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*, 2024. 3
- [7] Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Vadv2: End-to-end vectorized autonomous driving via probabilistic planning. *arXiv preprint arXiv:2402.13243*, 2024. 1, 2, 3, 4
- [8] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023. 3, 5, 6, 8
- [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 TPAMI*, 2022. 1, 3, 7
- [10] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 8
- [11] Daniel Dauner, Marcel Hallgarten, Andreas Geiger, and Kashyap Chitta. Parting with misconceptions about learning-based vehicle motion planning. In *CoRL*, pages 1268–1281. PMLR, 2023. 2, 3, 5
- [12] Daniel Dauner, Marcel Hallgarten, Tianyu Li, Xinshuo Weng, Zhiyu Huang, Zetong Yang, Hongyang Li, Igor Gilitschenski, Boris Ivanovic, Marco Pavone, Andreas Geiger, and Kashyap Chitta. Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking. In *NeurIPS*, 2024. 1, 2, 3, 4, 5
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009. 6
- [14] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *CoRL*, pages 1–16. PMLR, 2017. 1, 2, 3, 4, 5
- [15] Dechen Gao, Shuangyu Cai, Hanchu Zhou, Hang Wang, Iman Soltani, and Junshan Zhang. Cardreamer: Open-source learning platform for world model based autonomous driving. *arXiv preprint arXiv:2405.09111*, 2024. 2, 3
- [16] Zhiming Guo, Xing Gao, Jianlan Zhou, Xinyu Cai, and Botion Shi. Scenedm: Scene-level multi-agent trajectory generation with consistent diffusion models. *arXiv preprint arXiv:2311.15736*, 2023. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 5
- [19] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Melissa Mozifian, Florian Golemo, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, et al. Sim2real in robotics and automation: Applications and challenges. *IEEE transactions on automation science and engineering*, 18(2): 398–400, 2021. 2
- [20] 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*, pages 533–549. Springer, 2022. 1, 2
- [21] Xuemin Hu, Shen Li, Tingyu Huang, Bo Tang, Rouxing Huai, and Long Chen. How simulation helps autonomous driving: A survey of sim2real, digital twins, and parallel intelligence. *IEEE T-IV*, 2023. 2
- [22] 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 *CVPR*, pages 17853–17862, 2023. 1, 2, 3, 5, 6, 7, 8
- [23] Bernhard Jaeger, Kashyap Chitta, and Andreas Geiger. Hidden biases of end-to-end driving models. In *ICCV*, pages 8240–8249, 2023. 1, 3, 4
- [24] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *ICCV*, pages 7953–7963, 2023. 2, 3, 6, 7, 8
- [25] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *CVPR*, pages 21983–21994, 2023. 2, 6, 7
- [26] Xiaosong Jia, Zhenjie Yang, Qifeng Li, Zhiyuan Zhang, and Junchi Yan. Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. In *NeurIPS*, 2024. 2, 3, 5
- [27] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *ICCV*, pages 8340–8350, 2023. 1, 2, 3, 5, 6, 7, 8
- [28] Chiyu Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, Dragomir Anguelov, et al. Motiondiffuser: Controllable multi-agent motion prediction using diffusion. In *CVPR*, pages 9644–9653, 2023. 3

- [29] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Sim2real predictivity: Does evaluation in simulation predict real-world performance? *IEEE Robotics and Automation Letters*, 5(4):6670–6677, 2020. [2](#)
- [30] Qifeng Li, Xiaosong Jia, Shaobo Wang, and Junchi Yan. Think2drive: Efficient reinforcement learning by thinking in latent world model for quasi-realistic autonomous driving (in carla-v2). In *ECCV*, 2024. [2](#), [3](#), [4](#), [5](#), [6](#)
- [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. [1](#), [2](#), [3](#), [6](#), [7](#)
- [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 *CVPR*, pages 14864–14873, 2024. [2](#), [3](#)
- [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. [7](#)
- [34] T Lin. Focal loss for dense object detection. *arXiv preprint arXiv:1708.02002*, 2017. [7](#)
- [35] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [5](#)
- [36] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [5](#)
- [37] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Covernet: Multimodal behavior prediction using trajectory sets. In *CVPR*, pages 14074–14083, 2020. [3](#)
- [38] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, pages 194–210. Springer, 2020. [3](#)
- [39] Philip Polack, Florent Alth  , Brigitte d’Andr  a Novel, and Arnaud de La Fortelle. The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles? In *IEEE T-IV*, pages 812–818. IEEE, 2017. [4](#), [5](#)
- [40] Rajesh Rajamani. *Vehicle dynamics and control*. Springer Science & Business Media, 2011. [4](#), [5](#)
- [41] Katrin Renz, Long Chen, Ana-Maria Marcu, Jan H  nermann, Benoit Hanotte, Alice Karnsund, Jamie Shotton, Elahe Arani, and Oleg Sinavski. Carllava: Vision language models for camera-only closed-loop driving. *arXiv preprint arXiv:2406.10165*, 2024. [1](#), [3](#)
- [42] Yuval Tassa, Nicolas Mansard, and Emo Todorov. Control-limited differential dynamic programming. In *ICRA*, pages 1168–1175. IEEE, 2014. [4](#)
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. [3](#)
- [44] Wenhao Wang, Jiangwei Xie, ChuanYang Hu, Haoming Zou, Jianan Fan, Wenwen Tong, Yang Wen, Silei Wu, Hanming Deng, Zhiqi Li, et al. Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving. *arXiv preprint arXiv:2312.09245*, 2023. [1](#), [3](#)
- [45] Xinshuo Weng, Boris Ivanovic, Yan Wang, Yue Wang, and Marco Pavone. Para-drive: Parallelized architecture for real-time autonomous driving. In *CVPR*, pages 15449–15458, 2024. [1](#), [2](#), [3](#)
- [46] Theodor Westny, Bj  rn Olofsson, and Erik Frisk. Diffusion-based environment-aware trajectory prediction. *arXiv preprint arXiv:2403.11643*, 2024. [3](#)
- [47] Penghao Wu, Xiaosong Jia, Li Chen, Junchi Yan, Hongyang Li, and Yu Qiao. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. *NeurIPS*, 35:6119–6132, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [48] Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, Tsung-Wei Ke, and Katerina Fragkiadaki. Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation. In *CoRL*, 2023. [3](#)
- [49] Brian Yang, Huangyuan Su, Nikolaos Gkanatsios, Tsung-Wei Ke, Ayush Jain, Jeff Schneider, and Katerina Fragkiadaki. Diffusion-es: Gradient-free planning with diffusion for autonomous driving and zero-shot instruction following. *arXiv preprint arXiv:2402.06559*, 2024. [3](#)
- [50] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *ICCV*, pages 15222–15232, 2021. [2](#), [3](#)