


AdaThinkDrive: Adaptive Thinking via Reinforcement Learning for Autonomous Driving

Yuechen Luo^{1,3*}, Fang Li^{2,3*}, Shaoqing Xu^{2,3*,‡}, Zhiyi Lai³, Lei Yang⁴, Qimao Chen^{1,3}, Ziang Luo¹,
 Zixun Xie^{3,5}, Shengyin Jiang³, Jiaxin Liu^{1,3}, Long Chen³, Bing Wang³, Zhi-xin Yang^{2,‡}


Abstract—While reasoning technology like Chain-of-Thought (CoT) has been widely adopted in Vision-Language-Action (VLA) models, it demonstrates promising capabilities in end-to-end autonomous driving. However, recent efforts to integrate CoT reasoning often fall short in simple scenarios, introducing unnecessary computational overhead without improving decision quality. To address this, we propose AdaThinkDrive, a novel VLA framework with a dual-mode reasoning mechanism inspired by fast and slow thinking. First, our framework is pretrained on large-scale autonomous driving (AD) scenarios using both question-answering (QA) and trajectory datasets to acquire world knowledge and driving commonsense. During supervised fine-tuning (SFT), we introduce a two-mode dataset—fast answering (w/o CoT) and slow thinking (with CoT), enabling the model to distinguish between scenarios that require reasoning. Furthermore, an *Adaptive Think Reward* strategy is proposed in conjunction with the *Group Relative Policy Optimization (GRPO)*, which rewards the model for selectively applying CoT by comparing trajectory quality across different reasoning modes. Extensive experiments on the Navsim benchmark show that AdaThinkDrive achieves a PDMS of 90.3, surpassing the best vision-only baseline by 1.7 points. Moreover, ablations show that AdaThinkDrive surpasses both the never-Think and always-Think baselines, improving PDMS by 2.0 and 1.4, respectively. It also reduces inference time by 14% compared to the always-Think baseline, demonstrating its ability to balance accuracy and efficiency through adaptive reasoning.

I. INTRODUCTION

In recent years, autonomous driving systems have been shifting from traditional modular pipelines to end-to-end architectures. Although modular approaches offer engineering flexibility, they suffer from information loss across components, resulting in cumulative errors and limited generalization in complex and long-tail scenarios [1], [2]. End-to-end methods mitigate this by jointly optimizing perception, prediction, and planning within a unified model [3], [4], but their reliance on limited supervised data still constrains robustness. To address this, recent research has explored Vision-Language Models (VLMs) [5], leveraging large-scale driving datasets for pretraining to enhance scene understanding capabilities [6], [7], [8], [9], [10], [11]. Current VLM-based methods fall into two categories: meta-action approaches [12], [13], which generate high-level guidance, and planning-based approaches [8], [14], which directly predict trajectories via language modeling. CoT has been



(a) The performance of CoT on InternVL3-8B/2B across scenes of varying complexity, where both denote VLM models after SFT. Scene complexity increases progressively from Level 1 (simple) to Level 3 (challenging).



(b) AdaThinkDrive learns to assess scene complexity and select response modes adaptively. It generates trajectory with Think for challenging scenes and trajectory w/o Think for simple scenes, balancing accuracy and efficiency.

Fig. 1: Impact and Design of Adaptive Reasoning in Trajectory Prediction.

increasingly adopted for the latter to produce structured outputs [15], improving both interpretability and trajectory quality [8]. However, its application to VLA in autonomous driving is still in its infancy.

To explore the potential of CoT, we conducted a comparative study on reasoning performance in VLA models under varying scene complexities. Specifically, the driving scenes were categorized into three complexity levels, as shown in Figure 1a. We observe that for both InternVL3 8B and 2B models[16], the Non-Think model achieves better performance in simple scenarios (Level 1), whereas the Think model consistently outperforms as scene complexity increases (Levels 2 and 3). These findings reveal a critical limitation of existing CoT approaches: a tendency to over-reason in simple scenarios. While CoT reasoning provides substantial benefits in complex and challenging settings, it can introduce unnecessary cognitive steps and heightened uncertainty in simpler scenarios. These findings highlight that the optimal reasoning strategy is not universal but rather dependent on scene complexity. Consequently, enabling models to employ reasoning selectively based on scene complexity

¹Tsinghua University, ²University of Macau, ³Xiaomi EV,

⁴Nanyang Technological University, ⁵Peking University

Email: luo-yc24@mails.tsinghua.edu.cn

*Equal contribution. ‡Project Leader. † Corresponding author.

naturally becomes essential for improving both decision accuracy and inference efficiency in autonomous driving.

Following this, we propose **AdaThinkDrive**, a Vision-Language-Action (VLA) framework with a “Fast answering / Slow thinking” mechanism for end-to-end trajectory prediction (Figure 1b). We begin with a systematic analysis of the NAVSIM benchmark [17], evaluating the performance of existing methods across various scenario complexities. Motivated by this, we design a three-stage adaptive reasoning strategy that enables the model to automatically decide when to reason and when to act directly, guided by a learnable reward mechanism. In implementation, we first pretrain the model on large-scale driving data, then perform SFT using a customized dual-mode Navsim planning dataset to enable the model to generate both Think and Non-Think outputs. Finally, we adopt GRPO as the reinforcement learning algorithm and construct a reward structure that jointly considers trajectory accuracy, action rationality, and reasoning simplicity. This allows AdaThinkDrive to reach an optimal balance between planning performance and computational efficiency. We summarize our main contributions as follows:

- We conduct comparative studies of CoT in VLA across varying levels of scenario complexity. By evaluating think and non-think paradigms on the NAVSIM benchmark, we reveal that *over-reasoning in simple scenarios* appears to be a **key limitation** of existing CoT approaches, highlighting the need for adaptive reasoning strategies.
- We propose **AdaThinkDrive**, an end-to-end VLA framework with a “fast answering / slow thinking” mechanism that adaptively switches between direct prediction and explicit reasoning based on scene complexity. Furthermore, we design an **Adaptive Think Reward** strategy based on GRPO to guide the model in deciding when to reason and when to act directly.
- On the Navsim benchmark, AdaThinkDrive achieves a PDMS of 90.3, outperforming the leading vision-only baseline by 1.7 points. Furthermore, the model demonstrates its adaptive reasoning capability by selectively employing CoT in 96% of challenging scenarios, while defaulting to direct trajectory prediction in 84% of simple scenarios. Additionally, this adaptive approach reduces inference time by 14% compared to an always-think baseline, confirming the framework’s ability to effectively balance high performance with computational efficiency.

II. RELATED WORK

A. VLA for Autonomous Driving

In recent years, Vision-Language Models (VLMs) have gained increasing attention in autonomous driving, integrating visual and textual inputs for unified perception, planning, and decision-making. Current approaches broadly fall into two paradigms. The first focuses on scene understanding and high-level reasoning [11], [12], [13], [18]. For instance, Senna [12] interprets sensory inputs to produce meta-actions guiding downstream planners, although

improvements to actual driving performance remain limited. The second paradigm directly predicts driving trajectories from raw inputs [6], [8], [7], [9], [10], [19], [20]. To enhance interpretability and accuracy, recent methods increasingly adopt intermediate reasoning (Chain-of-Thought, CoT), revealing internal decision processes. EMMA [8], ReasonPlan [20], and Sce2DriveX [19] demonstrate that domain-specific reasoning significantly improves trajectory predictions. However, our analysis indicates that CoT benefits primarily complex scenarios, offering minimal or even negative impacts in simpler scenarios.

B. Efficient Reasoning Models

With the rising popularity of long Chain-of-Thought (Long CoT) in large language models, such as DeepSeek[21], lengthy inference processes have significantly increased computational costs. AdaptThink[22] addresses this challenge through comparative experiments, showing that direct answers are more accurate for simple tasks, while reasoning enhances performance on challenging tasks, thereby improving both efficiency and accuracy. Current mainstream approaches to adaptive CoT triggering mainly leverage reinforcement learning, emphasizing token-level control and reward design. These methods generally fall into three categories: (1) concise reasoning [23], [24], which encourages brevity through reward shaping or strict length constraints; (2) dynamic early stopping [25], allowing models to terminate reasoning adaptively; and (3) on-demand reasoning [26], [27], [22], [28], allowing models to decide whether to reason based on task complexity autonomously. In the context of autonomous driving, determining “when to think slowly and when to respond quickly” is particularly crucial. In simple scenarios, such as highway cruising, accurate predictions can be made without extensive reasoning. Conversely, in challenging scenarios like intersections or crowded environments, the model must analyze the scene carefully, identify critical agents, and then generate informed trajectories. In this work, we aim to efficiently enable the model to activate slow thinking when necessary and adaptively switch between reasoning modes.

III. METHODS

In this section, we demonstrate the design of our proposed AdaThinkDrive, which includes: (1) Data preparation, including pre-training data and hybrid SFT data. (2) Processing of a two-stage Supervised fine-tuning module that provides effective initialization; (3) Adaptive complex-aware thinking via a reinforcement learning strategy that boosts the efficiency and accuracy of model. The whole framework is demonstrated in Figure 2.

A. Problem Formulation

The input query q includes a front-view image q_{cam} , high-level navigation commands q_{com} (e.g., *Move Forward*, *Turn Left*, *Turn Right*), ego state information q_{ego} (e.g., *velocity and acceleration*), and the historical trajectory of the last three frames $q_{\text{his}} = \{h_{t-3}, h_{t-2}, h_{t-1}\}$. AdaThinkDrive operates with two reasoning modes $\mathcal{M} =$

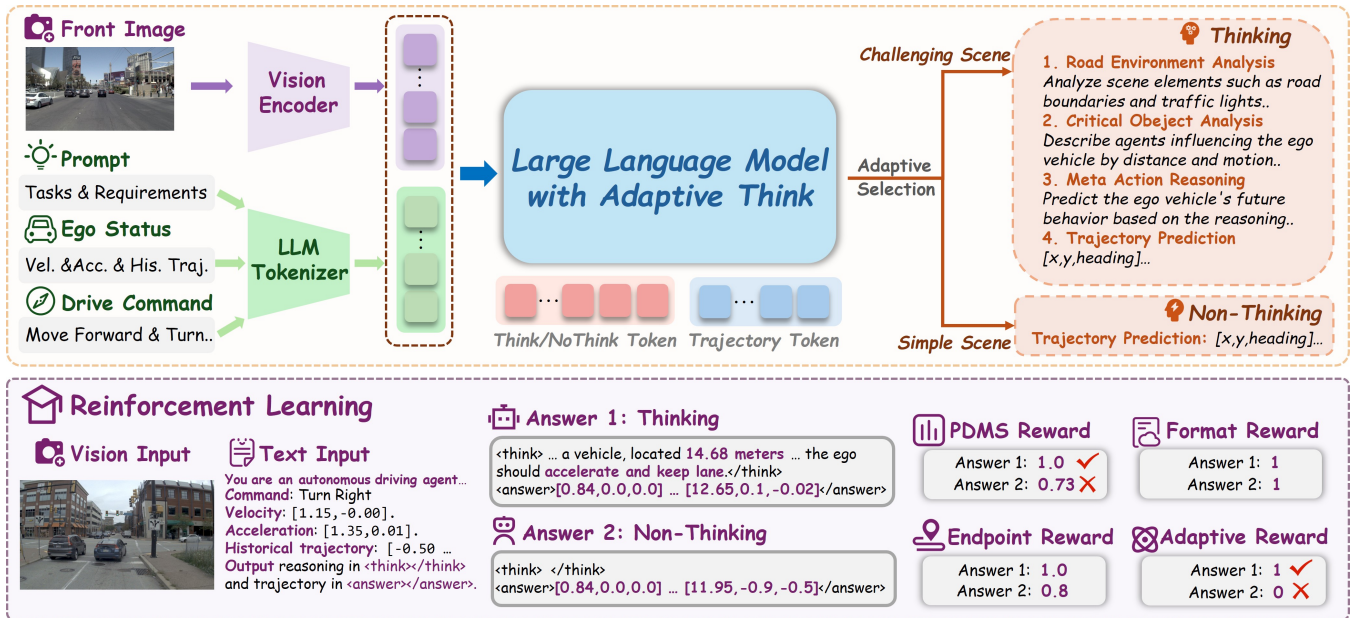


Fig. 2: We present **AdaThinkDrive**, an end-to-end autonomous driving framework that adaptively selects between “Thinking” and “Non-Thinking” modes depending on scene complexity. Given vision and text inputs, the VLM dynamically determines its output mode through an adaptive reasoning mechanism. During the reinforcement learning of the three-stage training process, multiple reward including PDMS, format, and endpoint are combined with the proposed Adaptive Think Reward.

$\{\text{Thinking}, \text{Non-Thinking}\}$. Given a query q , the model jointly determines a reasoning mode $m \in \mathcal{M}$ and an answer $o \in \mathcal{O}$ according to the joint distribution

$$\mathcal{P}(m, o | q) = \mathcal{P}(m | q) \mathcal{P}(o | q, m). \quad (1)$$

For each query q , the selected mode maximizes the expected task-specific utility $\mathcal{U}(q, o)$:

$$m(q) = \arg \max_{m \in \mathcal{M}} \mathbb{E}_{o \sim \mathcal{P}(o | q, m)} [\mathcal{U}(q, o)]. \quad (2)$$

The overall objective is to learn a policy $\pi : \mathcal{P} \rightarrow \mathcal{M}$ that selects a mode for each query so as to maximize the expected utility over query distributions $\Theta = \{(\mathcal{D}_i, \mathcal{U}_i)\}_{i=1}^N$:

$$\max_{\pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q \sim \mathcal{D}_i} [\mathbb{E}_{o \sim \mathcal{P}(o | q, \pi(q))} [\mathcal{U}_i(q, o)]] . \quad (3)$$

B. Data Preparation

To equip the model with foundational driving knowledge and an understanding of when CoT reasoning may be beneficial, we first perform a data preparation stage as follows.

a) *Pre-training Data*: To adapt general VLMs into autonomous driving, we assembled a diverse collection of open-source driving QA datasets, including DriveLM [29], LingoQA [18], ImpromptuVLA [30], NuScenes-QA [31] NuInstruct [32], and OminiDrive [7]. In addition, we constructed a multi-turn Q&A reasoning dataset for NAVSIM following CoT paradigm during SFT stage, including road boundary estimation, critical object identification, ego action prediction and related scene understanding subtasks.

b) *Hybrid SFT Data*: The SFT dataset consists of both reasoning-intensive and direct-answer examples. Reasoning data is generated by an auxiliary model combined with rule-based methods, ensuring high-quality reasoning traces.

For detailed scene descriptions like traffic light states and weather conditions, we automatically generate fine-grained annotations with Qwen2.5-VL-72B. Furthermore, to capture interactive scene dynamics, we identify dynamic agents expected to interact with the ego vehicle and categorize them into three types: Closest In-Path Object (CIPO-1) for those in the ego lane; CIPO-2 for those likely to merge, determined by lane geometry and relative location; and Motion Interaction for those with future trajectories predicted to intersect the ego trajectory, as illustrated in Figure 3. For static elements such as road boundaries, we utilize the NAVSIM map to reconstruct lane topology and find critical boundary features along the ego’s future path.

Furthermore, for each query q in the dataset, we generate both a Think-style response $\{q, o^{\text{Think}}\}$, which retains the full reasoning process $\langle \text{think} \rangle \text{reasoning} \langle / \text{think} \rangle$, and a Non-Think-style response $\{q, o^{\text{Non-Think}}\}$, which omits explicit reasoning but maintains structural consistency. For all cases, we directly fill the trajectory into $\langle \text{answer} \rangle \text{trajectory} \langle / \text{answer} \rangle$. Collectively, we denote the resulting supervised dataset as $\mathcal{D}^{\text{SFT}} = \{\{q, o^{\text{Think}}\}, \{q, o^{\text{Non-Think}}\}\}_{q \in \mathcal{Q}}$. This SFT data serves as a “warm-up”, equipping the model with foundational capability to distinguish between two response styles. The demo pipeline can be found in Figure 2.

c) *Scene Categorization*: To support adaptive reasoning, we categorize the NAVSIM training and validation sets into three levels of increasing complexity: Level 1, Level 2, and Level 3. This division follows the same criteria used in constructing the Think-style and non-Think-style datasets, which are based on two factors: whether the ego vehicle is near road boundaries and whether critical objects are present

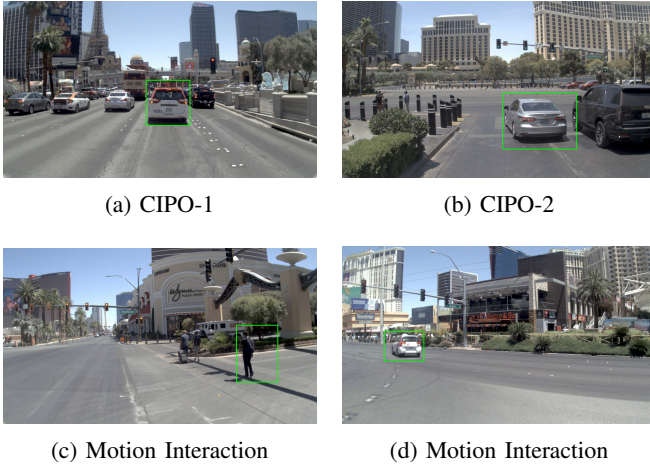


Fig. 3: Visualization of dynamic agents for Think-style CoT supervision. (a)-(b) depict CIPO-1 agents (occupying the ego lane) and CIPO-2 agents (likely to merge), while (c)-(d) show Motion Interaction cases where agents’ future trajectories intersect with the ego vehicle’s trajectory.

that may influence driving decisions. Level 1 includes scenes with neither condition, Level 2 includes scenes with only one, and Level 3 includes scenes with both. We define the dataset as $\mathcal{D} = \{\mathcal{D}^+, \mathcal{D}^-\}$, where \mathcal{D}^- corresponds to Level 1 and \mathcal{D}^+ corresponds to Level 2 and 3. These two categories serve as auxiliary labels to provide a principled initialization for downstream reinforcement learning.

C. Two-Stage Supervised Fine-tuning Processing

To build a model with driving knowledge and trajectory planning capabilities, we perform a two-stage fine-tuning procedure. The first stage injects general driving knowledge, while the second focuses on equipping models with the ability to adhere to trajectory generation and output formats.

In the first stage, the model is pretrained on a large corpus of driving-related Q&A pairs, enhancing its understanding of driving domain cognition. This phase addresses tasks like understanding the drivable area, object localization, and traffic semantics.

Subsequently, the second stage introduces the trajectory prediction task. For each query $q = (q_{cam}, q_{com}, q_{ego}, q_{his})$, two outputs are generated: $o^{Thinking}$, which includes reasoning, and $o^{Non-Thinking}$, which contains only the final trajectory.

During fine-tuning, the model is supervised on both outputs, aiming to maximize the conditional likelihood:

$$\mathcal{L}_{SFT} = \mathbb{E}_{(q,o) \sim \mathcal{D}^{SFT}} [-\log \pi_{\theta}(o | q)]. \quad (4)$$

This training strategy enables the model to learn both *Thinking* and *Non-Thinking* reasoning modes under a unified interface, while remaining unbiased toward either style. As a result, the model can generate both response types for any query q , enabling adaptive reasoning in the GRPO phase.

D. Adaptive Thinking via Reinforcement Learning

After SFT processing, AdaThinkDrive acquires the initial ability to support two different reasoning modes on the

same query q without collapsing. However, our goal is to enable the policy model to adaptively select the most appropriate reasoning mode $m(q)$ to improve its efficiency. To this end, we introduce a reinforcement learning phase in addition to supervised learning to explicitly teach the model how to adaptively select between reasoning modes, while simultaneously improving the model’s planning ability.

To enable the model to learn not only how but also when to reason, and to effectively leverage the reasoning process for accurate trajectory prediction, we design four complementary reward components: PDMS Reward, Format Reward, Endpoint Reward, and Adaptive Think Reward.

1) *PDMS Reward*: Evaluation metric *Predictive Driver Model Score (PDMS)* [17] for predicted trajectory is used for trajectory reward \mathcal{R}_{traj} , which is a Continuous value from 0 to 1.

2) *Format Reward*: This reward \mathcal{R}_{fmt} enforces compliance with the prescribed output format, covering both the correct use of `<think>...</think>` and `<answer>...</answer>` tags and the standardized representation of predicted trajectories. It provides discrete feedback for violations in either component, thereby ensuring consistent structural and content formatting.

3) *Endpoint Reward*: To encourage accurate alignment between predicted and ground-truth trajectory endpoints, we adopt a piecewise reward $\mathcal{R}_{endpoint}$ based on the L1 distance of the final point. A full score of 1.0 is given when the deviation is below 2 meters, decreasing stepwise to 0.8 (<4 m), 0.6 (<6 m), 0.4 (<10 m), 0.2 (< 15m), and 0.0 otherwise. This design penalizes large errors while remaining sensitive to small deviations near the trajectory endpoint.

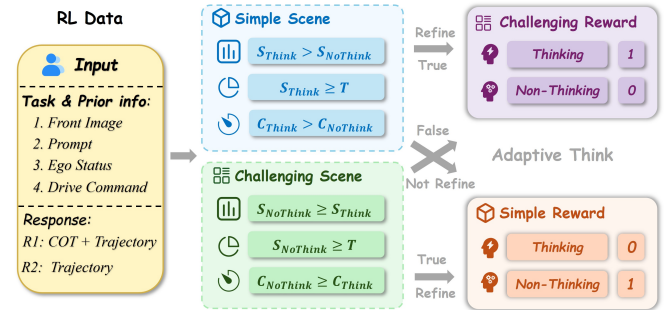


Fig. 4: Adaptive Think Reward: A Dynamic Reasoning Control Strategy. This reward adjusts the model’s reasoning behavior by identifying misclassified scenes. When scene-specific conditions are satisfied, it assigns rewards to either *Thinking* or *Non-thinking* responses accordingly.

4) *Adaptive Think Reward*: During reinforcement learning training, the policy model’s capability evolves dynamically. The same scenarios may be perceived as simple or challenging at different stages of training. To address for this, *Adaptive Think Reward* $\mathcal{R}_{adaptive}$ is designed to teach the model when to think, preventing over-reliance on static, manually defined scene tags $\{\mathcal{D}^+, \mathcal{D}^-\}$. These manual tags serve as an initial basis for reasoning, helping the model avoid “collapse”, where it only outputs *Thinking* or *Non-Thinking* without adapting to the reasoning needs of the

situation. The dynamic adjustment mechanism allows the model to correct scene tags based on actual reasoning needs, gradually learning adaptive thinking and improving both decision-making efficiency and prediction accuracy.

Adaptive Think Reward guides the model to adjust reasoning behavior dynamically using multiple rollouts. The detailed process is illustrated in Figure 4 and Algorithm 1.

Algorithm 1 Adaptive Think Reward

Input:

S_{Think} : Average PDMS of Thinking rollouts
 S_{Nothink} : Average PDMS of Non-thinking rollouts
 C_{Think} : Count of Thinking rollouts
 C_{Nothink} : Count of Non-thinking rollouts
 T : Confidence threshold (default: 0.9)
 D : Scene complexity label (0: Simple, 1: Challenging)

Output:

$\text{Reward}_{\text{Thinking}}, \text{Reward}_{\text{Non-Thinking}}$
Case $D = 0$ (Simple scene):
 1: **if** $S_{\text{Think}} > S_{\text{Nothink}} \ \& \ S_{\text{Think}} > T \ \& \ C_{\text{Think}} > C_{\text{Nothink}}$ **then**
 // Corrected to Challenging
 2: $\text{Reward}_{\text{Thinking}} \leftarrow 1, \text{Reward}_{\text{Non-Thinking}} \leftarrow 0$
 3: **else** // Maintained as Simple
 4: $\text{Reward}_{\text{Thinking}} \leftarrow 0, \text{Reward}_{\text{Non-Thinking}} \leftarrow 1$
 5: **end if**
Case $D = 1$ (Challenging scene):
 6: **if** $S_{\text{Nothink}} > S_{\text{Think}} \ \& \ S_{\text{Nothink}} > T \ \& \ C_{\text{Nothink}} > C_{\text{Think}}$ **then**
 // Corrected to Simple
 7: $\text{Reward}_{\text{Thinking}} \leftarrow 0, \text{Reward}_{\text{Non-Thinking}} \leftarrow 1$
 8: **else** // Maintained as Challenging
 9: $\text{Reward}_{\text{Thinking}} \leftarrow 1, \text{Reward}_{\text{Non-Thinking}} \leftarrow 0$
 10: **end if**
 11: **return** $\text{Reward}_{\text{Thinking}}, \text{Reward}_{\text{Non-Thinking}}$

The overall reward in the reinforcement learning process is computed by integrating four specifically designed reward components, which present as follows:

$$\mathcal{R}(q, a) = \mathcal{R}_{\text{traj}} + \mathcal{R}_{\text{fmt}} + \mathcal{R}_{\text{endpoint}} + \mathcal{R}_{\text{adaptive}}. \quad (5)$$

We utilize GRPO [33] as the training algorithm. For each query q , a set of candidate outputs $\{o_1, o_2, \dots, o_G\}$ is sampled from the old policy π_{old} , and the current policy π_θ is optimized based on their reward signals. To ensure stable training and avoid drastic policy shifts, GRPO incorporates truncated importance weights and a KL divergence regularization term: the former suppresses excessive policy updates, while the latter constrains the current policy from deviating too far from the reference policy π_{ref} . The final optimization objective for GRPO is defined as follows:

$$\mathcal{J}(\theta) = \mathbb{E}_{q, \{o_i\} \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{G} \sum_{i=1}^G \mathcal{J}_i - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right], \quad (6)$$

$$\mathcal{J}_i = \min(c_i A_i, \text{clip}(c_i, 1 - \epsilon, 1 + \epsilon) A_i). \quad (7)$$

where $c_i = \frac{\pi_\theta(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}$, ϵ and β are hyper-parameters, and the relative advantage A_i is calculated based on the normalized reward difference of a set of candidate outputs.

Through reinforcement learning, the policy model can develop adaptive reasoning strategies that adjust dynamically to the scenarios of different complexity.

A. Dataset and Metric

1) *Dataset*: We conduct comprehensive experiments and evaluations on NAVSIM, a planning-oriented autonomous driving dataset built on the OpenScene platform. In addition to the reasoning data collected from NAVSIM, we further leverage several open-source datasets, such as DriveLM, ImpromptuVLA, and LingoQA, by reformatting their visual VQA pairs to better support CoT reasoning.

2) *Metric*: The NAVSIM benchmark provides a non-reactive simulation environment and employs the Predictive Driver Model Score (PDMS) as closed-loop planning metric:

$$\text{PDMS} = \text{NC} \times \text{DAC} \times \left(\frac{5 \times \text{EP} + 5 \times \text{TTC} + 2 \times \text{C}}{12} \right). \quad (8)$$

where PDMS integrates five sub-metrics: No At-Fault Collision (NC), Drivable Area Compliance (DAC), Time-to-Collision (TTC), Comfort (Comf.), and Ego Progress (EP) to produce a comprehensive closed-loop planning score.

B. Implementation Details

We use InternVL3-8B [16] as the base model. The training consists of three stages. In the first stage, we fine-tune on a large-scale driving knowledge dataset for 2 epochs with a learning rate of 1×10^{-5} and batch size 1. The second stage fine-tunes on a curated Navsim planning dataset with Think and Non-Think annotations for 2 epochs with a learning rate of 4×10^{-5} and batch size 2. The third stage applies reinforcement learning for 2 epochs with a learning rate of 2×10^{-6} , batch size 4, and 64 NVIDIA H20 GPUs. The threshold T for Adaptive Think reward is set to 0.9.

C. Performance Comparison

1) *AdaThinkDrive Performance*: Table I presents the performance comparison between AdaThinkDrive and current leading methods on the NAVSIM benchmark. Under the vision-only setting, AdaThinkDrive achieves a PDMS of 90.3, establishing a new state-of-the-art (SOTA). Compared with the previous best vision-only method, Hydra-NeXt, AdaThinkDrive improves PDMS by 1.7, demonstrating significant advances in modeling capability and trajectory prediction accuracy. Moreover, despite relying solely on vision input, AdaThinkDrive performs comparably to the multi-modal approach GoalFlow, further validating the effectiveness of its adaptive reasoning mechanism and its strong generalization ability in complex driving scenarios. Finally, in best-of-N planning, we use the Navsim reference trajectory evaluator to select the optimal trajectory from four generated candidates, ultimately achieving the highest PDMS of 93.0.

2) *Quantitative Evaluation of Adaptive Think*: We first compare AdaThinkDrive (Table II) against several models:

- **Think/Non-Think SFT**: SFT model trained to always or never generate CoT.
- **Think/Non-Think RL**: RL model fine-tuned from the corresponding Think or Non-Think SFT model.

TABLE I: The Performance Comparison on NAVSIM using Closed-Loop Metrics.

Method	Image	Lidar	NC \uparrow	DAC \uparrow	TTC \uparrow	CF \uparrow	EP \uparrow	PDMS \uparrow
Constant Velocity			68.0	57.8	50.0	100	19.4	20.6
Ego Status MLP			93.0	77.3	83.6	100	62.8	65.6
UniAD [3]	✓		97.8	91.9	92.9	100	78.8	83.4
LFT [34]	✓		97.4	92.8	92.4	100	79.0	83.8
TransFuser [34]	✓	✓	97.7	92.8	92.8	100	84.0	84.0
PARA-Drive [35]	✓		97.9	92.4	93.0	99.8	79.3	84.0
DRAMA [36]	✓	✓	98.0	93.1	94.8	100	80.1	85.5
Hydra-MDP- \mathcal{V}_{8192} -W-EP [37]	✓	✓	98.3	96.0	94.6	100	78.7	86.5
DiffusionDrive [38]	✓	✓	98.2	96.2	94.7	100	82.2	88.1
WoTE [39]	✓	✓	98.5	96.8	94.9	99.9	81.9	88.3
Hydra-NeXt [40]	✓		98.1	97.7	94.6	100	81.8	88.6
GoalFlow [41]	✓	✓	98.4	98.3	94.6	100	85.0	90.3
AdaThinkDrive (Ours)	✓		98.4	97.8	95.2	100	84.4	90.3
AdaThinkDrive (Best-of-N)	✓		99.1	98.8	97.2	100	87.9	93.0

TABLE II: Comparison of Think and Non-Think Models in SFT and RL (InternVL3-8B). We compare the performance of SFT and RL models in Thinking (w) and Non-Thinking (w/o) modes. "Ours" refers to adaptive think mode.

Model	Mode	NC \uparrow	DAC \uparrow	TTC \uparrow	CF \uparrow	EP \uparrow	PDMS \uparrow
Non-Think SFT	w/o	98.1	92.1	94.0	100	77.1	83.3
Think SFT	w	98.5	94.4	94.9	100	79.9	86.2
Non-Think RL	w/o	98.2	96.1	94.3	100	83.5	88.3
Think RL	w	98.2	96.4	94.6	100	84.2	88.9
AdaThinkDrive	Ours	98.4	97.8	95.2	100	84.4	90.3

TABLE III: Comparison of inference time and PDMS among the Non-Think Model, Think Model, and AdaThinkDrive. Inference time denotes the average time to predict 4-second trajectories on the NAVSIM Test dataset.

Method	Infer Time \downarrow	PDMS \uparrow
Non-Think RL	0.68	88.3
Think RL	0.86	88.9
AdaThinkDrive	0.74	90.3

Notably, AdaThinkDrive achieves the best overall performance, outperforming the Non-Think RL and Think RL baselines by 2.0 and 1.4 PDMS, respectively. As shown in Table IV, it achieves 2.2 higher PDMS than Non-Think RL in Level 1 (simple scenarios) and 2.0 higher than Think RL in Level 3 (challenging scenarios). These improvements highlight AdaThinkDrive’s ability to combine the advantages of both strategies: skipping reasoning in simple scenarios to enhance efficiency, while leveraging structured reasoning in challenging ones for greater accuracy. Behavior analysis across different scene complexity (Figure 5) further confirms that AdaThinkDrive prefers the Non-Think mode in simple scenes and increasingly adopts the Think mode in challenging ones, demonstrating its capacity for dynamic reasoning control. In addition, Table III shows our model’s inference time is 9% higher than Non-Think RL baseline, while achieving a notable improvement of 2.0 PDMS in accuracy. Moreover, it shows a 14% reduction in inference time compared to Think RL baseline. Overall, these results validate the effectiveness of adaptive reasoning in balancing

TABLE IV: Comparison of Think and Non-Think RL Models in Simple (Level 1) and Challenging (Level 3) Scenarios.

Model	Level	NC \uparrow	DAC \uparrow	TTC \uparrow	CF \uparrow	EP \uparrow	PDMS \uparrow
Non-Think RL	Level 1	98.5	96.3	95.0	100	83.0	88.5
AdaThinkDrive	Level 1	98.8	98.1	96.1	100	84.5	90.7
Think RL	Level 3	98.7	92.9	95.4	100	84.7	87.8
AdaThinkDrive	Level 3	99.4	94.6	96.3	100	86.6	89.8




Fig. 5: The ratio of Think vs. Non-Think choices by AdaThinkDrive across different NAVSIM Test dataset levels. Scene complexity increases progressively from Level 1 (simple) to Level 3 (challenging).

accuracy and efficiency across diverse driving scenarios.

3) *Qualitative Analysis of Adaptive Think*: Figure 6 shows qualitative comparisons between AdaThinkDrive and baselines in both simple and challenging scenarios. In the simple scenario, the Think model misclassifies a distant object as critical, leading to unnecessary reasoning and a trajectory that strays from the drivable area. AdaThinkDrive skips redundant reasoning and directly outputs a smooth, accurate trajectory. In the challenging case, the Non-Think RL model fails to assess the distance to the lead vehicle, resulting in a risky plan. In contrast, AdaThinkDrive identifies the critical object and generates a safe trajectory. These examples demonstrate its ability to adaptive reasoning to scene complexity, improving both safety and decision quality.

D. Ablation Studies

1) *Ablation study on AdaThinkDrive*: Table V presents the ablation results of the three-stage training pipeline for AdaThinkDrive. Using only NAVSIM trajectory data for

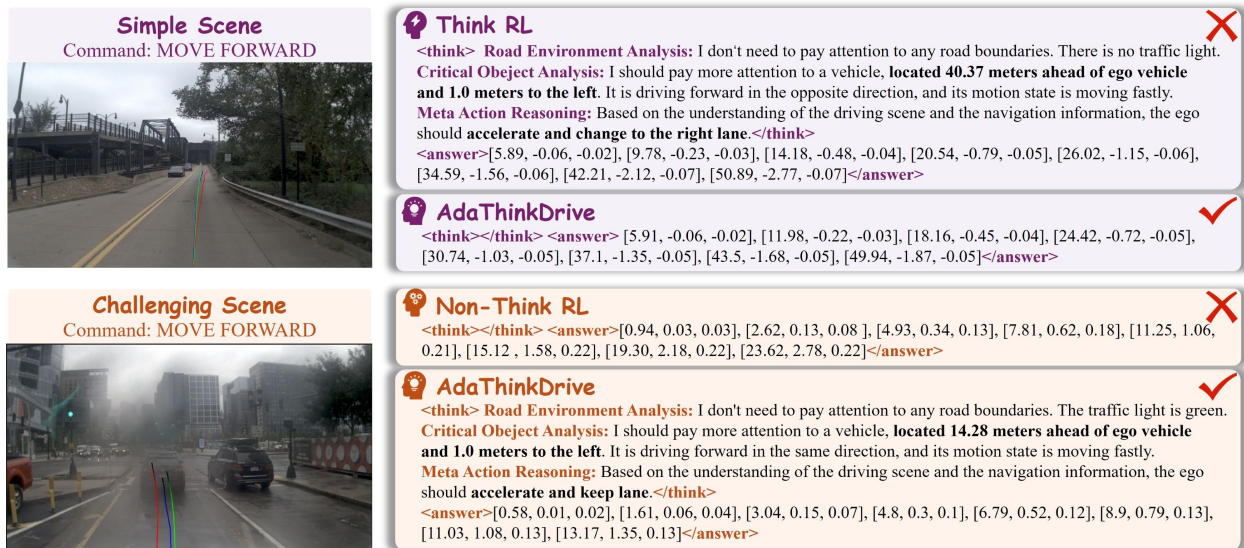


Fig. 6: Qualitative comparisons of trajectory predictions on the NAVSIM dataset. We compare AdaThinkDrive with the Think RL model in the simple scenario (top), and with the Non-Think RL model in a challenging scenario (bottom). Each example shows the input image, driving command, reasoning content (if any), and the predicted trajectory. Green, blue, and red represent the ground-truth, AdaThinkDrive planning, and Think/Non-Think Model planning, respectively.

TABLE V: Ablation Study on AdaThinkDrive Components. We evaluate the effect of pre-training, supervised fine-tuning, and reinforcement learning on driving performance using NAVSIM evaluation metrics.

Model	NC \uparrow	DAC \uparrow	TTC \uparrow	CF \uparrow	EP \uparrow	PDMS \uparrow
SFT	98.5	94.4	94.9	100	79.9	86.2
Pre+SFT	98.9	95.3	96.0	100	80.6	87.5
Pre+SFT+RL	98.4	97.8	95.2	100	84.4	90.3

SFT, the model achieves a PDMS of 86.2. Adding pretraining on a large-scale driving QA dataset boosts the score to 87.5, an increase of 1.3. Incorporating Adaptive Think reinforcement learning and the Adaptive Think Reward further improves performance to 90.3 PDMS, a gain of 2.8. These results demonstrate that both pretraining and the adaptive reinforcement learning strategy play a crucial role in enhancing the model’s understanding and reasoning capabilities.

2) *Comparison of Reward Effectiveness*: Table VI illustrates the impact of different reward combinations on PDMS. With only basic PDMS and format rewards, the model reaches 88.1. Adding the Endpoint Reward slightly raises it to 89.1. Incorporating Adaptive Think Reward further improves PDMS to 90.3, showing that adaptive reasoning is essential for improving planning efficiency and accuracy, and enhancing decision-making across diverse scenarios.

V. CONCLUSION

In this paper, we argue that reasoning in simple scenarios often introduces computational overhead without improving decision quality. To overcome this, We introduced AdaThinkDrive, a vision-language-action framework that enables the agent to adaptively learn when to think. Our key contribution is a reinforcement learning framework guided by an Adaptive Think Reward, which aligns reasoning behavior

TABLE VI: Ablation Studies of Reward Designs in GRPO. We evaluate the effect of PDMS Reward (P.), Format Reward (F.), Endpoint Reward (E.), and Adaptive Think Reward (A.) on driving performance using NAVSIM evaluation metrics.

ID	P. & F.	E.	A.	NC \uparrow	DAC \uparrow	TTC \uparrow	CF \uparrow	EP \uparrow	PDMS \uparrow
1	✓			96.7	95.2	90.7	100	87.8	88.1
2	✓	✓		98.3	96.6	94.5	100	84.4	89.1
3	✓	✓	✓	98.4	97.8	95.2	100	84.4	90.3

with scene complexity. Experimental results on the NAVSIM benchmark demonstrate that AdaThinkDrive achieves SOTA performance. These findings highlight the importance of adaptive thinking for achieving both accurate and efficient decision-making in autonomous systems.

VI. ACKNOWLEDGEMENT

This work was funded in part by the National Natural Science Foundation of China under Grant 62461160260, National Key Research and Development Program of China under grant 2023YFE0205800, the Science and Technology Development Fund, Macau SAR (Grant no. 0075/2023/AMJ, 0092/2024/AFJ, 0003/2023/RIB1, 001/2024/SKL), the Guangdong Department of Science and Technology (Grant no. 2023A0505030003), and the University of Macau (Grant No.: MYRG-GRG2024-00299-FST-UMDF, MYRG-GRG2025-00312-FST).

REFERENCES

- [1] L. Chen, P. Wu, K. Chitta, B. Jaeger, A. Geiger, and H. Li, “End-to-end autonomous driving: Challenges and frontiers,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [2] S. Jiang, Z. Huang, K. Qian, Z. Luo, T. Zhu, Y. Zhong, Y. Tang, M. Kong, Y. Wang, S. Jiao *et al.*, “A survey on vision-language-action models for autonomous driving,” *arXiv preprint arXiv:2506.24044*, 2025.

- [3] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang *et al.*, “Planning-oriented autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17 853–17 862.
- [4] B. Jiang, S. Chen, Q. Xu, B. Liao, J. Chen, H. Zhou, Q. Zhang, W. Liu, C. Huang, and X. Wang, “Vad: Vectorized scene representation for efficient autonomous driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8340–8350.
- [5] J. Zhang, J. Huang, S. Jin, and S. Lu, “Vision-language models for vision tasks: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [6] H. Fu, D. Zhang, Z. Zhao, J. Cui, D. Liang, C. Zhang, D. Zhang, H. Xie, B. Wang, and X. Bai, “Orion: A holistic end-to-end autonomous driving framework by vision-language instructed action generation,” *arXiv preprint arXiv:2503.19755*, 2025.
- [7] S. Wang, Z. Yu, X. Jiang, S. Lan, M. Shi, N. Chang, J. Kautz, Y. Li, and J. M. Alvarez, “OmniDrive: A holistic vision-language dataset for autonomous driving with counterfactual reasoning,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 22 442–22 452.
- [8] J.-J. Hwang, R. Xu, H. Lin, W.-C. Hung, J. Ji, K. Choi, D. Huang, T. He, P. Covington, B. Sapp *et al.*, “Emma: End-to-end multimodal model for autonomous driving,” *arXiv preprint arXiv:2410.23262*, 2024.
- [9] S. Xing, C. Qian, Y. Wang, H. Hua, K. Tian, Y. Zhou, and Z. Tu, “Openemma: Open-source multimodal model for end-to-end autonomous driving,” in *Proceedings of the Winter Conference on Applications of Computer Vision*, 2025, pp. 1001–1009.
- [10] Z. Qiao, H. Li, Z. Cao, and H. X. Liu, “Lightemma: Lightweight end-to-end multimodal model for autonomous driving,” *arXiv preprint arXiv:2505.00284*, 2025.
- [11] X. Tian, J. Gu, B. Li, Y. Liu, Y. Wang, Z. Zhao, K. Zhan, P. Jia, X. Lang, and H. Zhao, “DriveVLM: The convergence of autonomous driving and large vision-language models,” *arXiv preprint arXiv:2402.12289*, 2024.
- [12] B. Jiang, S. Chen, B. Liao, X. Zhang, W. Yin, Q. Zhang, C. Huang, W. Liu, and X. Wang, “Senna: Bridging large vision-language models and end-to-end autonomous driving,” *arXiv preprint arXiv:2410.22313*, 2024.
- [13] B. Jiang, S. Chen, Q. Zhang, W. Liu, and X. Wang, “Alphadrive: Unleashing the power of vlms in autonomous driving via reinforcement learning and reasoning,” *arXiv preprint arXiv:2503.07608*, 2025.
- [14] Y. Li, K. Xiong, X. Guo, F. Li, S. Yan, G. Xu, L. Zhou, L. Chen, H. Sun, B. Wang *et al.*, “Recogdrive: A reinforced cognitive framework for end-to-end autonomous driving,” *arXiv preprint arXiv:2506.08052*, 2025.
- [15] Y. Cui, H. Lin, S. Yang, Y. Wang, Y. Huang, and H. Chen, “Chain-of-thought for autonomous driving: A comprehensive survey and future prospects,” *arXiv preprint arXiv:2505.20223*, 2025.
- [16] J. Zhu, W. Wang, Z. Chen, Z. Liu, S. Ye, L. Gu, H. Tian, Y. Duan, W. Su, J. Shao *et al.*, “InternV3: Exploring advanced training and test-time recipes for open-source multimodal models,” *arXiv preprint arXiv:2504.10479*, 2025.
- [17] D. Dauner, M. Hallgarten, T. Li, X. Weng, Z. Huang, Z. Yang, H. Li, I. Gilitschenski, B. Ivanovic, M. Pavone *et al.*, “Navsim: Data-driven non-reactive autonomous vehicle simulation and benchmarking,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 28 706–28 719, 2024.
- [18] A.-M. Marcu, L. Chen, J. Hünemann, A. Karnsund, B. Hanotte, P. Chidananda, S. Nair, V. Badrinathan, A. Kendall, J. Shotton *et al.*, “Lingoqa: Visual question answering for autonomous driving,” in *European Conference on Computer Vision*. Springer, 2024, pp. 252–269.
- [19] R. Zhao, Q. Yuan, J. Li, H. Hu, Y. Li, C. Zheng, and F. Gao, “Scedrivex: A generalized mllm framework for scene-to-drive learning,” *arXiv preprint arXiv:2502.14917*, 2025.
- [20] X. Liu, Z. Zhong, Y. Guo, Y.-F. Liu, Z. Su, Q. Zhang, J. Wang, Y. Gao, Y. Zheng, Q. Lin *et al.*, “Reasonplan: Unified scene prediction and decision reasoning for closed-loop autonomous driving,” *arXiv preprint arXiv:2505.20024*, 2025.
- [21] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi *et al.*, “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” *arXiv preprint arXiv:2501.12948*, 2025.
- [22] J. Zhang, N. Lin, L. Hou, L. Feng, and J. Li, “Adaptthink: Reasoning models can learn when to think,” *arXiv preprint arXiv:2505.13417*, 2025.
- [23] M. Fatemi, B. Rafiee, M. Tang, and K. Talamadupula, “Concise reasoning via reinforcement learning,” *arXiv preprint arXiv:2504.05185*, 2025.
- [24] J. Yi, J. Wang, and S. Li, “Shorterbetter: Guiding reasoning models to find optimal inference length for efficient reasoning,” *arXiv preprint arXiv:2504.21370*, 2025.
- [25] B. Hou, Y. Zhang, J. Ji, Y. Liu, K. Qian, J. Andreas, and S. Chang, “Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning,” *arXiv preprint arXiv:2504.01296*, 2025.
- [26] G. Fang, X. Ma, and X. Wang, “Thinkless: Llm learns when to think,” *arXiv preprint arXiv:2505.13379*, 2025.
- [27] S. Tu, J. Lin, Q. Zhang, X. Tian, L. Li, X. Lan, and D. Zhao, “Learning when to think: Shaping adaptive reasoning in r1-style models via multi-stage rl,” *arXiv preprint arXiv:2505.10832*, 2025.
- [28] C. Lou, Z. Sun, X. Liang, M. Qu, W. Shen, W. Wang, Y. Li, Q. Yang, and S. Wu, “Adacot: Pareto-optimal adaptive chain-of-thought triggering via reinforcement learning,” *arXiv preprint arXiv:2505.11896*, 2025.
- [29] C. Sima, K. Renz, K. Chitta, L. Chen, H. Zhang, C. Xie, J. Beißwenger, P. Luo, A. Geiger, and H. Li, “DriveVLM: Driving with graph visual question answering,” in *European conference on computer vision*. Springer, 2024, pp. 256–274.
- [30] H. Chi, H.-a. Gao, Z. Liu, J. Liu, C. Liu, J. Li, K. Yang, Y. Yu, Z. Wang, W. Li *et al.*, “Impromptu vla: Open weights and open data for driving vision-language-action models,” *arXiv preprint arXiv:2505.23757*, 2025.
- [31] T. Qian, J. Chen, L. Zhuo, Y. Jiao, and Y.-G. Jiang, “Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 5, 2024, pp. 4542–4550.
- [32] X. Ding, J. Han, H. Xu, X. Liang, W. Zhang, and X. Li, “Holistic autonomous driving understanding by bird’s-eye-view injected multimodal large models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 13 668–13 677.
- [33] Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu *et al.*, “Deepseekmath: Pushing the limits of mathematical reasoning in open language models,” *arXiv preprint arXiv:2402.03300*, 2024.
- [34] K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, “Transfuser: Imitation with transformer-based sensor fusion for autonomous driving,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 11, pp. 12 878–12 895, 2022.
- [35] X. Weng, B. Ivanovic, Y. Wang, Y. Wang, and M. Pavone, “Paradrive: Parallelized architecture for real-time autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 449–15 458.
- [36] C. Yuan, Z. Zhang, J. Sun, S. Sun, Z. Huang, C. D. W. Lee, D. Li, Y. Han, A. Wong, K. P. Tee *et al.*, “Drama: An efficient end-to-end motion planner for autonomous driving with mamba,” *arXiv preprint arXiv:2408.03601*, 2024.
- [37] Z. Li, K. Li, S. Wang, S. Lan, Z. Yu, Y. Ji, Z. Li, Z. Zhu, J. Kautz, Z. Wu *et al.*, “Hydra-mdp: End-to-end multimodal planning with multi-target hydra-distillation,” *arXiv preprint arXiv:2406.06978*, 2024.
- [38] B. Liao, S. Chen, H. Yin, B. Jiang, C. Wang, S. Yan, X. Zhang, X. Li, Y. Zhang, Q. Zhang *et al.*, “Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 12 037–12 047.
- [39] Y. Li, Y. Wang, Y. Liu, J. He, L. Fan, and Z. Zhang, “End-to-end driving with online trajectory evaluation via bev world model,” *arXiv preprint arXiv:2504.01941*, 2025.
- [40] Z. Li, S. Wang, S. Lan, Z. Yu, Z. Wu, and J. M. Alvarez, “Hydra-next: Robust closed-loop driving with open-loop training,” *arXiv preprint arXiv:2503.12030*, 2025.
- [41] Z. Xing, X. Zhang, Y. Hu, B. Jiang, T. He, Q. Zhang, X. Long, and W. Yin, “Goalflow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 1602–1611.