

Drive-R1: Bridging Reasoning and Planning in VLMs for Autonomous Driving with Reinforcement Learning

Anonymous ICCV submission

Paper ID 2

Abstract

Large vision-language models (VLMs) for autonomous driving (AD) are evolving beyond perception and cognition tasks toward motion planning. However, we identify two critical challenges in this direction: (1) VLMs tend to learn shortcuts by relying heavily on history input information, achieving seemingly strong planning results without genuinely understanding the visual inputs; and (2) the chain-of-thought (COT) reasoning processes are always misaligned with the motion planning outcomes, and how to effectively leverage the complex reasoning capability to enhance planning remains largely underexplored. In this paper, we start from a small-scale domain-specific VLM and propose **Drive-R1**, designed to bridge the scenario reasoning and motion planning for AD. **Drive-R1** first undergoes the supervised finetuning on a elaborate dataset containing both long and short COT data. **Drive-R1** is encouraged to reason step-by-step from visual input to final planning decisions. Subsequently, **Drive-R1** is trained within a reinforcement learning framework that incentivizes the discovery of reasoning paths that are more informative for planning, guided by rewards based on predicted trajectories and meta actions. Experimental evaluations on the nuScenes and DriveLM-nuScenes benchmarks demonstrate that **Drive-R1** achieves superior performance compared to existing state-of-the-art VLMs. We believe that **Drive-R1** presents a promising direction for bridging reasoning and planning in AD, offering methodological insights for future research and applications.

1. Introduction

Autonomous driving (AD) systems aim to enable vehicles to perceive, understand, and interact with their environments in a safe and intelligent manner. Among the core modules in AD pipelines, motion planning plays a central role in determining the future actions, balancing the safety, efficiency, and comfort in real-world driving sce-

narios. Given observations of the environment and other agents, trajectory prediction directly influence the subsequent low-level control.

Traditional motion planning methods often rely on manually crafted rules [3, 8, 35] that operate under simplified assumptions of the environment and agent behaviors. While these approaches offer interpretability and robustness in structured scenarios, they typically struggle to handle uncertainty, multi-agent interaction, and diverse traffic patterns. Recently, deep learning-based methods [11, 12, 17] have shown remarkable success in trajectory prediction by leveraging large-scale driving datasets. These methods, comprised of encoder-decoder architectures or spatio-temporal transformers, model the complex agent dynamics and social interactions. The trajectory prediction lacks the interpretability and still faces limitations in reasoning under ambiguous contexts, adapting to open-world conditions and long-tailed events. The emergence of large vision-language models (VLMs) have introduced new opportunities for enhancing AD systems. Recent methods [28–30, 38, 40, 44] have demonstrated promising results in scene perception, description, and decision-making with analysis in open form visual question answers task. Further, the methods [14, 18, 28, 42] extend the perception and cognition tasks to motion planning task, with some output interpretable decision processes.

However, several fundamental limitations remain insufficiently addressed in current VLM-based planning systems.

1) **The utilization of visual-grounded response in motion planning is limited or even entirely absent.** Recent VLM-based approaches [15, 40] achieve strong open-loop metrics by predicting trajectories from image-text inputs, often with short or no chain-of-thought (COT) reasoning. Early GPT-driver [28] revealed that transforming all perceptual and historical information into textual inputs and using a pure large language model alone can already produce competitive planning performance. To further probe this, we train a general VLM to predict trajectories without COT supervision. At the test time, we ablate the visual input entirely and find that the model performs comparably or even better than

Question: Frame 1: <image> Frame 2: <image> Frame 3: <image> Frame 4: <image> Frame 5: <image> Frame 6: <image>. These six images are the front view, front left view, front right view, back view, back left view and back right view of the ego vehicle. The inputs are : 1. Historical Trajectory (last 2 seconds): $[(-1.41,-8.98), (-0.96,-6.75), (-0.51,-4.52), (-0.05,-2.29)]$. 2.Ego-States: - Velocity (v_x,v_y) : $(0.03,2.28)$ - Heading Angular Velocity (v_{yaw}) : (-0.01) - Acceleration (a_x,a_y) : $(-0.40,0.06)$ - Can Bus: $(0.32,1.38)$ - Heading Speed: (2.35) - Steering: (2.22) . 3. Mission Goal: LEFT. Develop a safe and feasible 3-second route using 6 new waypoints.



Answer: $[(-0.42,2.41), (-1.00,4.61), (-1.61,6.79), (-2.24,8.89), (-2.88,10.94), (-3.54,12.94)]$

No - Images

Answer: $[(-0.42,2.38), (-1.00,4.77), (-1.61,7.15), (-2.23,9.52), (-2.85,11.88), (-3.47,14.24)]$

GT: $[(-0.32,2.37), (-0.90,4.85), (-1.69,7.46), (-2.71,10.15), (-3.86,12.92), (-5.15,15.69)]$

Figure 1. Inference results with and without visual inputs from the model which is trained to predict trajectory without chain of thoughts.

with full multi-model inputs. The observation indicates that VLMs for motion planning under-utilize the visual modality and heavily rely on textual priors, especially historical motion cues, raising concerns about their visual grounding and generalization. 2) **The CoT reasoning traces are always misaligned with the motion planning outcomes.** Leveraging the reasoning capability to enhance planning performance remains largely underexplored. Recent methods [14, 40] engages in sequential question-answering to arrive at the final trajectory prediction. While such methods introduce interpretable intermediate steps, the reasoning remains loosely coupled with planning. We further observe that even when a domain-specific (DS) driving model is trained on motion planning datasets with CoT reasoning, it often falls into a reasoning trap. First, the reasoning patterns learned from CoT data especially designed for complex scenarios may introduce unnecessary analysis in simple cases, leading to overthinking and ultimately injecting noise into the planning output. Second, even manually annotated CoT cannot guarantee precise alignment with the ground-truth trajectories, as natural language reasoning tends to be coarse-grained and ambiguous compared to the fine-grained numerical representation required for planning [18].

To address the aforementioned challenges and bridge the gap between scenario reasoning and trajectory-level motion planning in AD, we introduce **Drive-R1** tailored for vision-language reasoning and trajectory prediction. We begin with a general VLM, InternVL2 [5], and adapt it to the AD domain by post-training on a large-scale, self-collected dataset comprising 3 million samples. This DS model is endowed with strong perception and scene understanding capabilities, forming a foundation for downstream planning tasks.

To enable reasoning-aware planning, we construct a

structured annotation pipeline that generates CoT data according to key domains in real-world AD [22], including traffic knowledge understanding, general element recognition, traffic graph generation, target attribute comprehension, and ego decision-making and planning. The resulting CoT dataset contains approximately 4,000 samples, categorized into short and long CoT based on the complexity of the driving scenarios: short CoT correspond to relatively simple situations that require minimal deliberation, whereas long CoT are designed for complex, multi-agent, or rule-intensive scenes demanding richer step-by-step reasoning. During the supervised learning stage, **Drive-R1** is trained on the elaborate dataset to learn to reason from visual observations toward final planning outputs in an interpretable and structured manner. This stage is crucial for encouraging grounded reasoning and mitigating the tendency to overfit to historical trajectory patterns or exploit dataset shortcuts.

To further align the textual reasoning and numerical trajectory planning, we introduce the reinforcement learning (RL) inspired by the success of recent RL approaches [9, 13]. Specifically, **Drive-R1** employs the Group Relative Policy Optimization (GRPO), which performs optimization over a set of candidate solutions. The relative optimization mechanism is particularly suitable for motion planning, where multiple plausible trajectories may exist under the same driving scenario. By leveraging comparisons across diverse candidates, GRPO encourages the model to discover reasoning paths that generalize well across variations, rather than overfitting to a single deterministic trajectory, thereby enhancing both planning robustness and generalization. The reward design in GRPO integrates four components: trajectory accuracy, meta-action correctness, repetition penalty, and output format compliance. Among them, the trajectory reward captures outcome-level planning

quality, while the meta-action reward reflects the reasoning process quality. These two reward signals are complementary, further promoting effective alignment between reasoning and planning within the **Drive-R1** framework.

We conduct extensive experiments on both the nuScenes [1] dataset and the DriveLM-nuScenes [38] dataset. Our proposed **Drive-R1** achieves state-of-the-art performance on the trajectory prediction task, demonstrating its effectiveness in visual-grounded motion planning. Furthermore, we perform comprehensive ablation studies on DriveLM-nuScenes, investigating the impact of various components, including the GRPO on model in different phases, the number of rollouts, and the influence of different reward functions. Our contribution can be summarized as follows:

- We identify two key challenges in applying VLMs to motion planning: (i) the over-reliance on historical textual inputs leads to shortcut learning, weakening the visual grounding; and (ii) the misalignment between reasoning chains and planning outputs hinders effective integration of interpretability and decision quality.
- We propose **Drive-R1**, a DS VLM tailored for AD, which connects visual-grounded reasoning to trajectory planning. Our approach incorporates supervised learning on a carefully constructed dataset containing both long and short CoT annotations, followed by RL with GRPO to further align reasoning quality with planning performance.
- We conduct extensive experiments on nuScenes and DriveLM-nuScenes, where **Drive-R1** achieves state-of-the-art results on trajectory prediction.

While our work represents a straightforward exploration of integrating VLM into the motion planning pipeline, the insights gained from **Drive-R1** may offer valuable guidance for future efforts toward the practical deployment of VLMs in AD systems.

2. Related Work

2.1. Vision-language Models for Autonomous Driving

The integration of VLMs into AD has recently gained significant attention, aiming to unify perception, reasoning, and planning within a single framework. Existing works in this field can be broadly divided into two categories: scene reasoning-oriented models, and planning and control-oriented models. The first focuses on scene understanding and reasoning [7, 26, 29, 30, 38], where VLMs are used to analyze visual environments through natural language, often leveraging question-answering or chain-of-thought reasoning to enhance transparency and trustworthiness. Planning and control-oriented models [4, 31, 37, 40, 42, 44], on the other hand, aim to directly generate actionable out-

puts such as trajectories or control signals from visual and linguistic inputs. These systems often leverage large-scale data and unified modeling to perform planning implicitly within the language model, with or without intermediate reasoning steps. In this paper, we focus on trajectory prediction and find that models can achieve competitive planning performance even with limited or no visual input, suggesting a potential over-reliance on linguistic or historical features and insufficient grounding in visual observations.

2.2. Reinforcement Learning

RL has played a pivotal role in the recent evolution of large language models, particularly in aligning model outputs with human preferences or task-specific objectives. Early developments, such as proximal policy optimization [36] and direct policy optimization [34] have been widely adopted in general-purpose LLMs to improve response helpfulness and safety from human feedback, demonstrating that large models could benefit from post-training optimization beyond supervised learning, enabling them to reason and act in more aligned and consistent ways. Recent Group Relative Policy Optimization (GRPO) [9] proposes a group-wise relative optimization strategy, which compares the relative merits of multiple output candidates instead of optimizing based solely on absolute reward values. GRPO has shown strong potential in complex reasoning tasks by encouraging models to explore interpretable thought processes rather than shortcutting to answers. Building on this, RL has been extended to VLMs to enhance their ability to perform multi-step reasoning grounded in visual inputs [13, 18, 20]. AlhpaDrive [18] introduced RL to high-level planning and ReCogDrive [23] claimed the gap between the discrete language space and the continuous action space. Aligning textual reasoning with numerical outputs like trajectories in AD presents unique challenges, requiring designs that balance process-level and result-level precision. In this paper, our work still pursues the alignment between reasoning and planning in the discrete space.

3. Drive-R1

Drive-R1 aims to bridge scenario-level reasoning and trajectory planning for AD through a combination of SFT and RL. We begin by introducing the construction of the Reasoning-Planning chain of thought (RP-CoT) dataset, which encodes intermediate reasoning steps aligned with planning outcomes. Then we detail the supervised training phase, highlighting the initial capabilities the model must acquire to address the challenges discussed above. Finally, we describe the RL procedure, which leverages carefully designed reward functions to further align textual reasoning with numerical trajectory prediction, enhancing both interpretability and planning performance.

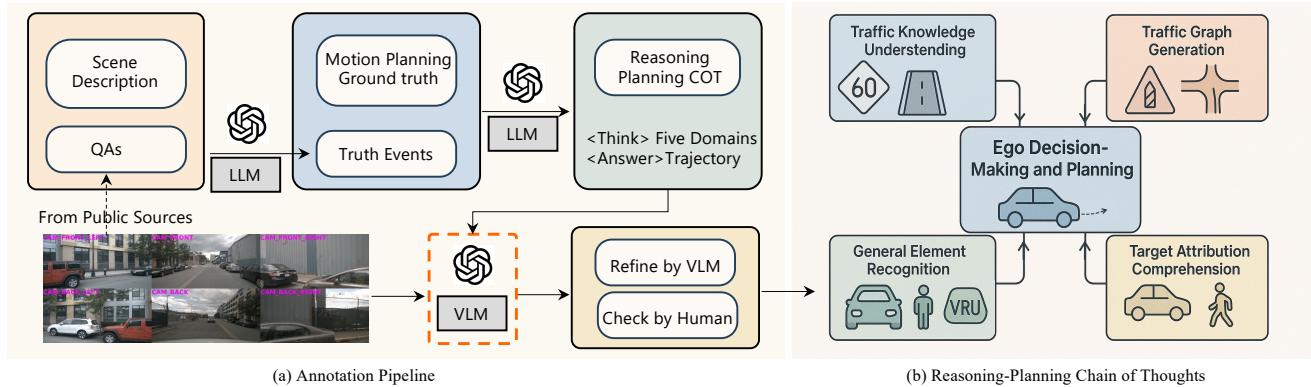


Figure 2. The RP-COT data annotation pipeline.

3.1. RP-COT Data Annotation

Following the five key domains identified in [22] as fundamental to motion planning, i.e., traffic knowledge understanding, general element recognition, traffic graph generation, target attribute comprehension, and ego decision-making and planning, we construct the RP-CoT dataset. RP-CoT is designed to bridge high-level scenario reasoning with low-level trajectory outputs in AD. The scenes are selected from nuScenes [1]. Each annotation sample in RP-CoT includes step-by-step textual reasoning that reflects a structured understanding of the driving scene, ultimately grounded in a precise trajectory decision.

As shown in Fig.2, the annotation pipeline is semi-automatic. We begin by collecting driving scenes from publicly available sources [16, 33, 38], which are annotated with scene descriptions and question-answer (QA) pairs. According to the hundreds of QAs of the scene, ChatGPT first generates the truth events, which are structured representations of the underlying reasoning rationale. Next, based on the ground-truth events and the motion planning information (history trajectory, ego status, meta action), ChatGPT generates RP-COT data through the five domains. Each sample includes the <think></think> section that explains reasoning steps and the <trajectory></trajectory> section that specifies the future trajectory (6 points with 3 seconds). To ensure the generated RP-CoTs are grounded in visual reality, the VLM, GPT-4o, is employed to refine these outputs by aligning them with scene content. Finally, all annotations are checked by human annotators to guarantee consistency, correctness, and planning validity.

Our annotation pipeline systematically decomposes the visual-linguistic information into reasoning stages aligned with the aforementioned domains. This structured format enables the model to learn interpretable reasoning paths that progressively lead to planning actions, laying a strong foundation for subsequent learning stages.

3.2. Supervised Fine-tuning Phase

As discussed above, the utilization of visual-grounded response in motion planning is limited or even absent. As shown in Table 1, directly training a general-purpose VLM (Internvl2-4B) to output trajectories without CoT supervision can yield surprisingly competitive performance. However, we observe a counterintuitive outcome: the model performs better when visual inputs are ignored, indicating a strong reliance on historical textual context over visual perception. We attribute this phenomenon to two key factors: (1) the model lacks sufficient familiarity with DS tasks in AD and (2) the model are more sensitive to historical motion cues than to scene-level visual information. To address this, we first perform full-parameter finetuning of an InternVL2-4B model on a large-scale, DS dataset comprising 3 million AD QAs, which are collected from public sources [2, 10, 19, 21, 25–29, 32, 38, 41, 44]. The DS model from the first SFT stage significantly mitigates the overreliance on historical information and enhances its general understanding of AD scenarios. Nevertheless, the gap between visual-informed and vision-agnostic reasoning remains narrow. Further, we incorporate the previously constructed RP-CoT dataset into the second SFT stage. Through supervised CoT supervision, the model is encouraged to form visual-grounded reasoning paths across key domains, solving the dependency on textual history information and thereby promoting more robust, perception-aware planning behavior.

On the other hand, the COT reasoning traces are always misaligned with the motion planning outcomes. The experimental results in Table 1 show that applying long CoT supervision during the SFT stage lead to a decline in performance compared to directly supervising the final trajectory output. Interestingly, similar observations occur in other domains. Recent researches [39, 43] report that for tasks involving spatial reasoning or numerical sensitivity, models trained with CoT supervision often underperform compared

Table 1. The preliminary experimental results validated on 799 samples from DriveLM-nuScenes [27]. BA and DS are the base and domain-specific models. WI denotes inference without image inputs.

Models	RP-COT		Training Phase		L2(m)↓				Collision↓			
	Long	Short	SFT	RFT	1s	2s	3s	Avg	1s	2s	3s	Avg
BA	×	✓	✓	×	0.25	0.55	0.97	0.59	0.00	0.09	0.56	0.22
BA-WI	×	✓	✓	×	0.25	0.53	0.90	0.56	0.00	0.03	0.48	0.18
DS	×	✓	✓	×	0.18	0.42	0.76	0.45	0.00	0.03	0.46	0.16
DS	✓	×	✓	×	0.24	0.53	0.91	0.56	0.00	0.19	0.61	0.27
DS	✓	✓	✓	×	0.19	0.39	0.67	0.41	0.00	0.03	0.29	0.11
BA	-	-	×	✓	0.37	0.75	1.22	0.78	0.00	0.16	0.84	0.33
DS	-	-	×	✓	0.26	0.55	0.93	0.58	0.00	0.19	0.61	0.27
DS	✓	✓	✓	✓	0.17	0.35	0.60	0.37	0.00	0.00	0.30	0.10

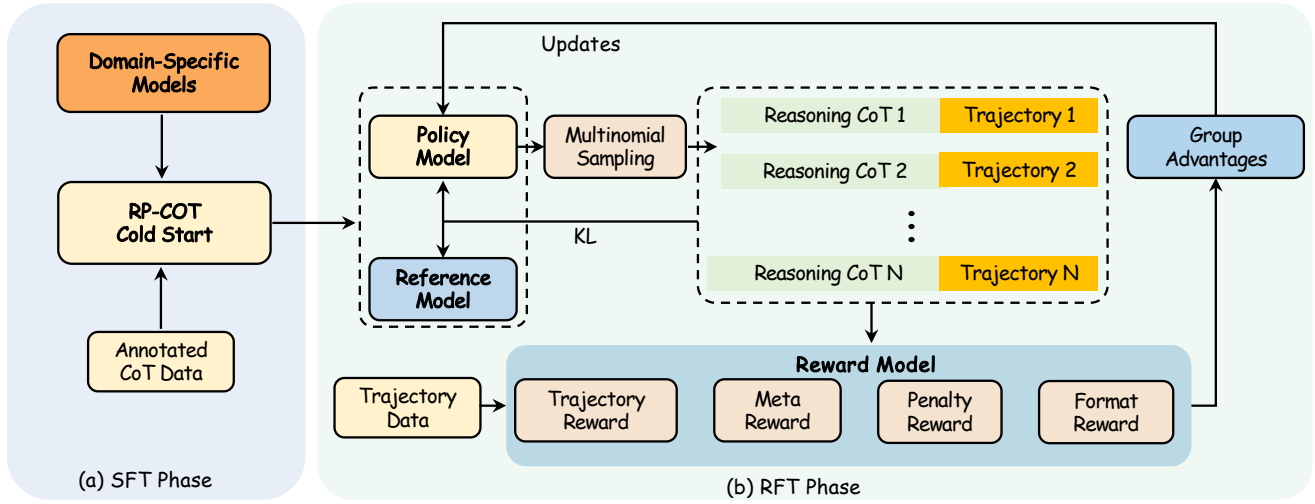


Figure 3. The overview of the proposed **Drive-R1**, which comprises the supervised fine-tuning (SFT) and reinforcement fine-tuning (RFT) phases.

to those trained with direct answer supervision. We hypothesize that the observed performance degradation may stem from two primary factors: (1) the limited representation capacity of small-scale models, which restricts their ability to accurately encode and utilize complex reasoning paths [24] and (2) the differing tolerance to errors of the models between textual and numerical outputs. Specifically, reasoning texts generated during CoT supervision may contain semantic inconsistencies or hallucinations, either due to imperfect annotation quality or intrinsic limitations of the model. While such errors may have negligible impact on the interpretability or plausibility of the textual reasoning itself, they can propagate to the numerical prediction stage, e.g., trajectory prediction, where small deviations are amplified into significant planning errors.

To mitigate the negative impact of indiscriminate CoT supervision, we introduce a fast-and-slow thinking strategy in the second SFT stage. The core idea is to adapt the complexity of reasoning supervision to the difficulty of each driving scenario. Specifically, we categorize CoT supervision into short CoT and long CoT, depending on the rea-

soning demand: short CoT corresponds to relatively simple scenarios requiring minimal deliberation, while long CoT is designed for complex, multi-agent, or rule-intensive scenes that demand richer step-by-step reasoning. We begin by training a model to directly generate trajectory outputs without CoT supervision. This specific model is then used to assess the reasoning complexity of each scene, serving as a proxy for scenario difficulty. Scenes with low planning metrics are assigned short CoT supervision, while those with high planning metrics are paired with long CoT annotations. As shown in Table 1, models fine-tuned with this adaptive fast-and-slow thinking strategy achieve the best overall performance, validating its effectiveness in balancing long and short CoT.

3.3. Reinforcement Learning Phase

DeepSeek-R1 [9] demonstrates that RL frameworks like GRPO can effectively elicit long CoT reasoning abilities of large language models. However, subsequent studies [6, 46] have shown that the reasoning paths produced by RL-finetuned models already exist with high probab-

ity in the output distribution of base model, i.e., problems solvable by the RL model can also be addressed by the base model through sufficient sampling. Building upon these insights, we adopt GRPO not as a means to unlock fundamentally new capabilities, but rather as a post-training alignment mechanism to improve the efficiency and consistency for further aligning the reasoning and planning.

3.3.1. Algorithm

Specifically, for each question q , GRPO [9] samples a group of candidate outputs $\{o_1, o_2, \dots, o_G\}$ from the old policy $\pi_{\theta_{\text{old}}}$, and subsequently updates the current policy π_{θ} by maximizing the following objective:

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

$$w_i = \min \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}, \text{clip} \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}, 1 - \epsilon, 1 + \epsilon \right) \right), \quad (2)$$

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}, \quad (3)$$

where the KL loss is calculated by $\mathbb{D}_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}) = \frac{\pi_{\text{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - \log \frac{\pi_{\text{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - 1$, $\{r_1, r_2, \dots, r_G\}$ are the rewards of the candidate outputs.

3.3.2. Rewards

To better align the intermediate reasoning steps with final motion planning outcomes, we design a composite reward function that balances both process-level and outcome-level results. The total rewards comprise the following components:

Trajectory Reward measures the accuracy of the predicted trajectory by computing the L2 distance between the predicted trajectory $\hat{\tau}$ and the ground truth τ . The raw distance is then mapped using a sigmoid-based transformation:

$$R_{\text{traj}} = \frac{2e^{-d}}{1+e^{-d}}, d = \|\hat{\tau} - \tau\|_2.$$

Meta-Action Reward assesses the high-level planning decisions in the reasoning section, including the short-term lateral and longitudinal decisions. Each contributing 0.5 to the total reward score.

Repetition Penalty penalizes the generation of redundant or repetitive reasoning steps within the CoT to encourage concise and efficient planning rationale [45].

Format Reward ensures structural correctness of the output format.

Importantly, our analysis shows that the result-oriented trajectory reward and the process-oriented meta-action reward are positively correlated.

3.3.3. Training

Through extensive experiments, we observe that effective RL in the context of motion planning is highly dependent on the model’s prior alignment with the AD domain. When applied to models without sufficient domain adaptation, reinforcement signals often result in unstable updates or limited policy improvement, suggesting that the capacity to interpret structured driving scenarios is a prerequisite for successful policy refinement. Consequently, we perform RL on a model that has been supervised with two-stage fine-tuning, as introduced in the SFT phase. Building on such warm-up, RL further amplifies the synergy between visual-grounded reasoning and motion planning, leading to the most significant performance gains observed in our experiments.

4. Experiments

4.1. Datasets and Baselines

In the first SFT phase, the domain specific data are collected from a diverse set of AD datasets, including [2, 10, 19, 21, 25–29, 32, 38, 41, 44] with 3 million samples. The QAs are built following the five key domains, and include the single-view, multi-view, and sequential image inputs. In the second SFT phase, RP-COT dataset are construct from the annotations in [16, 33, 38] with the number of samples 4,072. When compared on the 6019 validation samples on nuScenes [1], the numbers of long and short RP-COT are 24058 and 4072. When compared on the 799 validation samples on DriveLM-nuScenes [38], the numbers of long and short RP-COT are 2036 and 2036. In the RL phase, the samples are selected from those in 4072 RP-COT datasets.


We benchmark **Drive-R1** against both end-to-end and vision-language planning baselines. The former includes ST-P3 [11], UniAD [12] and their modified versions augmented with ego-status inputs [17]. The latter set of baselines includes DriveVLM [40], RDA-Driver [14], OmniDrive [42], and EMMA [15]. Notably, prior approaches typically output direct trajectory predictions either without reasoning or with short CoT supervision. In contrast, **Drive-R1** produces both reasoning chains and trajectory outcomes in a unified manner, enabling interpretable and context-sensitive planning.

4.2. Implementation and Metrics

The SFT training is conducted based on the official code-base of InternVL2 [5]. The first-stage SFT is trained on 32 V100 nodes with a batch size of 256, while the second-stage SFT is trained on 16 V100 nodes with a batch size of 128. The RL phase is implemented using the ms-swift framework [47] and trained on 2 V100 nodes with a batchsize of 16 and a rollout of 6. The context length is set to 4096. For evaluation, we adopt the L2 distance and collision rate metrics, following ST-P3 [11].

Table 2. Overall comparison with baselines on the nuScenes [1] validation.

Models	L2(m)↓				Collision↓			
	1s	2s	3s	Avg	1s	2s	3s	Avg
nuScenes								
ST-P3 [11]	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71
UniAD [12]	0.48	0.96	1.65	1.03	0.05	0.17	0.71	0.31
UniAD-E [12]	0.20	0.42	0.75	0.46	0.02	0.25	0.84	0.37
VAD-E [17]	0.17	0.34	0.60	0.37	0.07	0.10	0.24	0.14
DriveVLM [40]	0.18	0.34	0.68	0.40	0.10	0.22	0.45	0.27
RDA-Driver [14]	0.17	0.37	0.69	0.40	0.01	0.05	0.26	0.10
OmniDrive [42]	0.14	0.29	0.55	0.33	0.00	0.13	0.78	0.30
EMMA [15]	0.14	0.29	0.54	0.32	-	-	-	-
Drive-R1 (Ours)	0.14	0.28	0.50	0.31	0.02	0.06	0.19	0.09
nuScenes-DriveLM								
ST-P3 [11]	1.28	2.03	2.81	2.04	0.14	0.72	1.28	0.71
GPT-Driver [28]	0.22	0.43	0.73	0.46	0.00	0.13	0.46	0.19
RDA-Driver [14]	0.18	0.38	0.68	0.41	0.00	0.06	0.36	0.14
Drive-R1 (Ours)	0.17	0.35	0.60	0.37	0.00	0.00	0.30	0.10

**System Prompt:**

1. Role Definition 2. Inputs Definition 3. Task Definition: Thought Process, Trajectory Prediction

Prompt:

The Inputs are: 1. Frames. 2. Historical Trajectory (last 2 seconds) 3. Ego-States 4. Mission Goal.

DSModel + GRPO :

<think>...approaching an intersection with a traffic light.no traffic signs or road barriers. The key objects include a white sedan, a black sedan, and a white truck. The white sedan is moving straight ahead, while the black sedan is stationary and is in an adjacent lane. The trajectory planning involves a main straight path with a right turn at the intersection. The ego vehicle should decelerate slightly to ensure a safe turn. The historical trajectory shows a consistent path, supporting the decision to maintain speed and execute a right turn. The ego-states indicate a slight steering adjustment to align with the turn. </think>

<trajectory>[(0.10,2.60), (0.36,5.22), (0.70,7.77), (1.06,10.12), (1.45,12.29), (1.82,14.34)]</trajectory>

Drive-R1:

<think> The ego vehicle ... making a right turn at an intersection, and continuing forward. There are no traffic signals or signs directly affecting the ego vehicle's immediate path. The key objects include a black sedan, a white sedan, and a white truck, and pedestrians. The relationship between the vehicles and the lanes is straightforward. The white sedan is in the ego lane. The behavior of these objects is consistent, with no indications of sudden stops or lane changes. The trajectory is planned to ensure a smooth transition into the right lane while maintaining a safe distance from the moving vehicles. The ego vehicle's lateral meta action is to change lanes to the right, and the longitudinal meta action is to maintain speed.</think>

<trajectory>[(0.15,2.54), (0.47,4.99), (0.98,7.44), (1.04,9.74), (1.82,12.00), (2.30,13.20)]</trajectory>

GT: [(0.06,2.08), (0.26,4.36), (0.58,6.36), (0.98,8.16), (1.64,10.49), (2.39,12.79)]

Figure 4. An inference results from different models.

4.3. Results

Table 2 presents a comprehensive comparison between **Drive-R1** and existing representative baselines. **Drive-R1** achieves the lowest average L2 error of 0.31,

marginally outperforming EMMA (0.32) [15] and OmniDrive (0.33) [42], while also exhibiting the lowest average collision rate of 0.09. In contrast, several end-to-end methods demonstrate competitive L2 metrics, yet suf-

Table 3. Ablation Studies of the reward designs and the number of rollouts in GRPO. T., F., R., M. represent the trajectory reward, format reward, repetition penalty, and meta-action reward.

Rewards			Rollouts	L2(m)↓				Collision↓			
T.&F.	R.	M.		1s	2s	3s	Avg	1s	2s	3s	Avg
✓	×	×	6	0.17	0.35	0.60	0.37	0.06	0.06	0.42	0.18
✓	✓	×	6	0.17	0.34	0.59	0.37	0.06	0.06	0.30	0.14
✓	✓	✓	6	0.17	0.35	0.60	0.37	0.00	0.00	0.30	0.10
✓	✓	✓	12	0.17	0.35	0.60	0.37	0.00	0.03	0.25	0.11
✓	✓	✓	24	0.17	0.35	0.60	0.37	0.00	0.03	0.21	0.08

fer from relatively higher collision rates. This suggests that while these models may fit the trajectory well numerically, they may lack robustness in safety-critical aspects of planning. Among VLM-based baselines, **Drive-R1** consistently achieves better planning quality and safety. Notably, compared with RDA-Driver [14] and OmniDrive [42], our model demonstrates both improved trajectory precision and reduced collision risks, validating the effectiveness of reasoning-aligned trajectory generation.

4.4. Ablation Studies

We conduct extensive ablation experiments on the DriveLM-nuScenes [38] to investigate the effects of RP-COT input types and RL configurations.

Effect of COT Length in SFT Stage. As shown in Table 1, we evaluate the influence of long and short RP-COTs in the second SFT stage. Models trained with only short RP-COTs or long RP-COTs underperform those trained with both long and short RP-COTs, suggesting that applying a uniform CoT strategy across diverse scenarios is suboptimal. Instead, combining both short and long CoTs better equips the model to handle a wider variety of AD contexts, leveraging both concise and elaborate reasoning chains.

Effectiveness of RL on Different Model Bases. In Table 1, we further assess how RL impacts different model variants. The DS model, trained on the first SFT stage, benefits from RL than the base model. Further, when incorporating long and short RP-COT pattern, the model significantly benefit from the RL phase, reducing both trajectory deviation and collision rates. This underscores the necessity of prior domain alignment before performing RL fine-tuning.

Impact of Reward Design and Rollout Numbers. In Table 3, we evaluate how various reward components and rollout counts affect model performance. The inclusion of meta-action rewards and repetition penalties leads to consistent improvements in collision rates (e.g., 0.14 to 0.10), highlighting their effectiveness in guiding safer planning behavior. However, for models with relatively small capacity, simply increasing the number of rollouts does not always yield stable or consistent performance gains. For instance, although the collision rate decreases from 0.10 to 0.11 and to 0.08 when the number of rollouts increases from

6 to 12 and to 24, we observe that the training becomes unstable beyond a certain number of steps under the 24-rollout setting. It is worth noting that the reported result at 24-rollout is extracted before the onset of training collapse.

5. Conclusion

In this work, we present **Drive-R1**, which bridges the structured chain-of-thought reasoning and the trajectory-level motion planning. To address the insufficient visual grounding and the misalignment between reasoning traces and planning outputs observed in existing VLM-based approaches, we construct a domain-specific VLM and augment it with a systematically annotated CoT dataset spanning five essential reasoning domains. Furthermore, we incorporate a RL phase based on GRPO to optimize planning quality for aligning the reasoning process with trajectory outcomes. Comprehensive experiments conducted on nuScenes and DriveLM-nuScenes benchmarks validate the effectiveness of our proposed method. Drive-R1 achieves state-of-the-art performance on trajectory prediction tasks while offering interpretable and structured reasoning capabilities. **Drive-R1** represents an early exploration toward bridging high-level cognitive reasoning and low-level trajectory planning in AD. In addition, we conduct extensive experiments on large-scale in-house datasets using Ascend 910 hardware platforms, which further verify the generalizability and robustness of the Drive-R1 framework. We believe that the insights gained from **Drive-R1** may offer valuable guidance for future efforts toward the practical deployment of VLMs in AD systems.

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multi-modal dataset for autonomous driving. In *Conference on Computer Vision and Pattern Recognition*, 2020. 3, 4, 6, 7
- [2] Xu Cao, Tong Zhou, Yunsheng Ma, Wenqian Ye, Can Cui, Kun Tang, Zhipeng Cao, Kaizhao Liang, Ziran Wang, James M Rehg, et al. Maplm: A real-world large-scale vision-language benchmark for map and traffic scene understanding. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 21819–21830, 2024. 4, 6
- [3] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE international conference on computer vision*, pages 2722–2730, 2015. 1
- [4] 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*. Springer, 2024. 3
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Conference on Computer Vision and Pattern Recognition*, 2024. 2, 6
- [6] Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *arXiv preprint arXiv:2501.17161*, 2025. 5
- [7] Xinpeng Ding, Jianhua Han, Hang Xu, Xiaodan Liang, Wei Zhang, and Xiaomeng Li. Holistic autonomous driving understanding by bird’s-eye-view injected multi-modal large models. In *Conference on Computer Vision and Pattern Recognition*, pages 13668–13677, 2024. 3
- [8] Haoyang Fan, Fan Zhu, Changchun Liu, Liangliang Zhang, Li Zhuang, Dong Li, Weicheng Zhu, Jiangtao Hu, Hongye Li, and Qi Kong. Baidu apollo em motion planner. *arXiv preprint arXiv:1807.08048*, 2018. 1
- [9] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 2, 3, 5, 6
- [10] Yunfei Guo, Fei Yin, Xiao-hui Li, Xudong Yan, Tao Xue, Shuqi Mei, and Cheng-Lin Liu. Visual traffic knowledge graph generation from scene images. In *International Conference on Computer Vision*, pages 21604–21613, 2023. 4, 6
- [11] 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*, 2022. 1, 6, 7
- [12] 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 *Conference on Computer Vision and Pattern Recognition*, 2023. 1, 6, 7
- [13] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025. 2, 3
- [14] Zhijian Huang, Tao Tang, Shaoxiang Chen, Sihao Lin, Zequn Jie, Lin Ma, Guangrun Wang, and Xiaodan Liang. Making large language models better planners with reasoning-decision alignment. In *European Conference on Computer Vision*, pages 73–90. Springer, 2024. 1, 2, 6, 7, 8
- [15] Jyh-Jing Hwang, Runsheng Xu, Hubert Lin, Wei-Chih Hung, Jingwei Ji, Kristy Choi, Di Huang, Tong He, Paul Covington, Benjamin Sapp, et al. Emma: End-to-end multimodal model for autonomous driving. *arXiv preprint arXiv:2410.23262*, 2024. 1, 6, 7
- [16] Yuichi Inoue, Yuki Yada, Kotaro Tanahashi, and Yu Yamaguchi. Nuscenesc-mqa: Integrated evaluation of captions and qa for autonomous driving datasets using markup annotations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 930–938, 2024. 4, 6
- [17] 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 *International Conference on Computer Vision*, 2023. 1, 6, 7
- [18] Bo Jiang, Shaoyu Chen, Qian Zhang, Wenyu Liu, and Xinggang Wang. Alphadrive: Unleashing the power of vlms in autonomous driving via reinforcement learning and reasoning. *arXiv preprint arXiv:2503.07608*, 2025. 1, 2, 3
- [19] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles. In *Conference on Computer Vision and Pattern Recognition*, 2019. 4, 6
- [20] Yuxiang Lai, Jake Zhong, Ming Li, Shitian Zhao, and Xiaofeng Yang. Med-r1: Reinforcement learning for generalizable medical reasoning in vision-language models. *arXiv preprint arXiv:2503.13939*, 2025. 3
- [21] Kaican Li, Kai Chen, Haoyu Wang, Lanqing Hong, Chaoqiang Ye, Jianhua Han, Yukuai Chen, Wei Zhang, Chunjing Xu, Dit-Yan Yeung, et al. Coda: A real-world road corner case dataset for object detection in autonomous driving. In *European Conference on Computer Vision*. Springer, 2022. 4, 6
- [22] Yue Li, Meng Tian, Zhenyu Lin, Jiangtong Zhu, Dechang Zhu, Haiqiang Liu, Zining Wang, Yueyi Zhang, Zhiwei Xiong, and Xinhai Zhao. Fine-grained evaluation of large vision-language models in autonomous driving. *arXiv preprint arXiv:2503.21505*, 2025. 2, 4
- [23] Yongkang Li, Kaixin Xiong, Xiangyu Guo, Fang Li, Sixu Yan, Gangwei Xu, Lijun Zhou, Long Chen, Haiyang Sun, Bing Wang, et al. Recogdrive: A reinforced cognitive framework for end-to-end autonomous driving. *arXiv preprint arXiv:2506.08052*, 2025. 3
- [24] Yuetai Li, Xiang Yue, Zhangchen Xu, Fengqing Jiang, Luyao Niu, Bill Yuchen Lin, Bhaskar Ramasubramanian, and Radha Poovendran. Small models struggle to learn from strong reasoners. *arXiv preprint arXiv:2502.12143*, 2025. 5
- [25] Yuhang Lu, Yichen Yao, Jiadong Tu, Jiangnan Shao, Yuexin Ma, and Xinge Zhu. Can lvlms obtain a driver’s license? a benchmark towards reliable agi for autonomous driving. *arXiv preprint arXiv:2409.02914*, 2024. 4, 6
- [26] Yingzi Ma, Yulong Cao, Jiachen Sun, Marco Pavone, and Chaowei Xiao. Dolphins: Multimodal language model for

- driving. In *European Conference on Computer Vision*, 2024. 3
- [27] Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. Drama: Joint risk localization and captioning in driving. In *Winter Conference on Applications of Computer Vision*, 2023. 5
- [28] Jiageng Mao, Yuxi Qian, Junjie Ye, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. 1, 7
- [29] Ana-Maria Marcu, Long Chen, Jan Hünemann, Alice Karnsund, Benoit Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, et al. Lingoqa: Visual question answering for autonomous driving. In *European Conference on Computer Vision*, 2024. 3, 4, 6
- [30] Ming Nie, Renyuan Peng, Chunwei Wang, Xinyue Cai, Jianhua Han, Hang Xu, and Li Zhang. Reason2drive: Towards interpretable and chain-based reasoning for autonomous driving. In *European Conference on Computer Vision*. Springer, 2024. 1, 3
- [31] Chenbin Pan, Burhaneddin Yaman, Tommaso Nesti, Abhirup Mallik, Alessandro G Allievi, Senem Velipasalar, and Liu Ren. Vlp: Vision language planning for autonomous driving. In *Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [32] Chirag Parikh, Rohit Saluja, CV Jawahar, and Ravi Kiran Sarvadevabhatla. Idd-x: A multi-view dataset for ego-relative important object localization and explanation in dense and unstructured traffic. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14815–14821. IEEE, 2024. 4, 6
- [33] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *AAAI Conference on Artificial Intelligence*, 2024. 4, 6
- [34] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741, 2023. 3
- [35] Axel Sauer, Nikolay Savinov, and Andreas Geiger. Conditional affordance learning for driving in urban environments. In *Conference on robot learning*, pages 237–252. PMLR, 2018. 1
- [36] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 3
- [37] Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L Waslander, Yu Liu, and Hongsheng Li. Lmdrive: Closed-loop end-to-end driving with large language models. In *Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [38] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *European Conference on Computer Vision*, 2024. 1, 3, 4, 6, 8
- [39] Huajie Tan, Yuheng Ji, Xiaoshuai Hao, Minglan Lin, Pengwei Wang, Zhongyuan Wang, and Shanghang Zhang. Reason-rft: Reinforcement fine-tuning for visual reasoning. *arXiv preprint arXiv:2503.20752*, 2025. 4
- [40] 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. 1, 2, 3, 6, 7
- [41] Huijie Wang, Tianyu Li, Yang Li, Li Chen, Chonghao Sima, Zhenbo Liu, Bangjun Wang, Peijin Jia, Yuting Wang, Shengyin Jiang, et al. Openlane-v2: A topology reasoning benchmark for unified 3d hd mapping. *Advances in Neural Information Processing Systems*, 36:18873–18884, 2023. 4, 6
- [42] Shihao Wang, Zhiding Yu, Xiaohui Jiang, Shiyi Lan, Min Shi, Nadine Chang, Jan Kautz, Ying Li, and Jose M Alvarez. Omnidrive: A holistic llm-agent framework for autonomous driving with 3d perception, reasoning and planning. *arXiv preprint arXiv:2405.01533*, 2024. 1, 3, 6, 7, 8
- [43] Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, et al. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*, 2024. 4
- [44] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 1, 3, 4, 6
- [45] Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-of-thought reasoning in llms. *arXiv preprint arXiv:2502.03373*, 2025. 6
- [46] Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv preprint arXiv:2504.13837*, 2025. 5
- [47] Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, et al. Swift: a scalable lightweight infrastructure for fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 29733–29735, 2025. 6