

LaViPlan : Language-Guided Visual Path Planning with RLVR

Hayeon Oh

Electronics and Telecommunications Research Institute

oph516@etri.re.kr

Abstract

*Out-of-distribution (OOD) scenarios in autonomous driving pose critical challenges, as planners often fail to generalize beyond their training experience, leading to unsafe or unexpected behavior. Vision-Language Models (VLMs) have shown promise in handling such scenarios by providing high-level scene understanding and user-aligned decisions. However, existing VLMs often exhibit a misalignment between their language-based reasoning and the low-level trajectories required for action-level planning. In this paper, we propose **LaViPlan**, a framework that leverages Reinforcement Learning with Verifiable Rewards (RLVR) to fine-tune VLMs using planning-oriented metrics. Experimental results show that LaViPlan improves planning performance across both in-domain and out-of-domain datasets. While linguistic fidelity slightly decreases after RLVR-based fine-tuning, qualitative evaluation indicates that the outputs remain coherent. We also conduct ablation studies to analyze the effects of sampling ratio and reasoning guidance, highlighting how these design choices influence performance. These findings demonstrate the potential of RLVR as a post-training paradigm for aligning language-guided reasoning with action-level planning in autonomous driving.*

1. Introduction

Out-of-distribution (OOD) scenarios in autonomous driving refer to rare or novel situations that deviate from the training domain, often leading to unexpected and unsafe behavior by learned planning policies. These scenarios pose critical challenges, particularly when the planner fails to generalize beyond its supervised experience.

To address this, recent research has explored the integration of VLMs [1, 5, 18] into autonomous driving. VLMs have demonstrated strong generalization capabilities across diverse tasks and modalities, making them a promising approach for handling OOD scenarios. Early research showed that VLMs could identify unseen driving contexts and generate high-level decisions [13, 14].

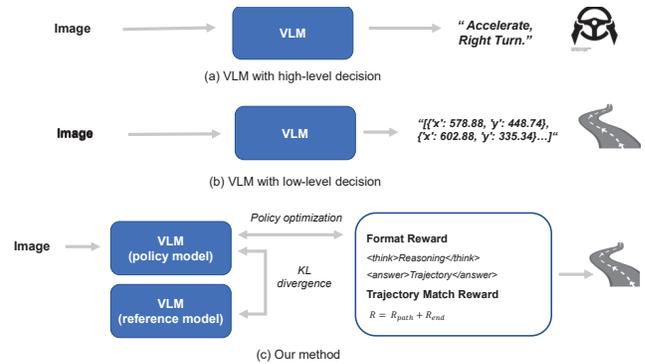


Figure 1. (a) VLMs generating high-level commands (e.g., “Accelerate, Right Turn”) based on scene understanding, but lacking direct trajectory grounding. (b) VLMs producing low-level outputs such as trajectories without explicit reasoning, often leading to semantically inconsistent or context-unaware behavior. (c) Our proposed method introduces a differentiable connection between vision-language reasoning and action space using RLVR. The reward consists of a format-based reasoning verification and trajectory alignment, optimized under KL regularization between the policy and reference model.

However, while VLMs can recognize and describe OOD scenes, their final decisions—especially in the form of predicted trajectories—can often be misaligned with the visual reasoning they produce. This issue reflects a broader challenge in aligning language-based reasoning with action-level planning, which we refer to as the *vision-language-action misalignment*.

To mitigate misalignment issues in other domains, GRPO-based reinforcement fine-tuning has been proposed, showing particular effectiveness in improving performance on numerical prediction tasks [19, 29]. By leveraging multiple candidate outputs and preference-based rewards instead of relying solely on a single ground-truth target, GRPO provides richer supervisory signals that enhance both accuracy and generalization beyond the training distribution. This success suggests that GRPO could be a promising approach to address the misalignment in VLM for autonomous driving.

Building on this insight, we propose to address the misalignment by leveraging Reinforcement Learning with Verifiable Rewards (RLVR), where planning-oriented metrics serve as verifiable reward signals. Our method aims to steer VLMs toward context-aware decision-making that is consistent with their situational reasoning. Our key contributions are summarized as follows:

- We propose a reinforcement learning framework that explicitly optimizes planning-oriented metrics in VLMs, demonstrating a step toward aligning language-guided reasoning with action-oriented planning tasks in autonomous driving.
- Through both quantitative and qualitative analyses, we reveal that RLVR shifts the model’s generation from linguistically faithful outputs to functionally accurate trajectories, indicating a trade-off between semantic similarity and task-specific reasoning.
- Experimental results demonstrate that RLVR requires significantly fewer training samples compared to supervised fine-tuning while still achieving performance gains, showing that including hard cases during RL yields better generalization.

2. Related Works

2.1. VLMs for Autonomous Driving

End-to-end autonomous driving has recently attracted attention due to its simplicity, efficiency, and ability to mitigate suboptimality arising from misaligned objectives between modular components. By eliminating intermediate representations, this paradigm reduces information loss and computational overhead. However, visual abstraction in end-to-end systems can oversimplify complex scene information, potentially discarding critical cues. Achieving robust generalization across diverse driving scenarios also remains challenging, especially in long-tail scenarios where labeled data is limited or sparse.

To address these limitations, recent studies have explored integrating Vision-Language Models (VLMs) into autonomous driving. VLMs leverage multi-modal understanding to reason about previously unseen situations. Typically, these models combine pre-trained large language models (LLMs) with visual encoders: driving instructions or ego vehicle states are fed as textual input to the LLM, while single- or multi-view images are processed by the visual encoder [13, 15, 24, 27, 28, 30, 34, 35]. Recent extensions incorporate 3D perceptual positional embeddings [7, 32, 38] and counterfactual learning [25, 32] to enable context-aware decision-making in complex driving environments.

2.2. Preference Learning for Alignment

While VLMs demonstrate strong multi-modal reasoning capabilities, their outputs may still be misaligned with downstream action-level tasks, such as trajectory prediction. One common approach to address this issue is Reinforcement Learning from Human Feedback (RLHF), which aligns VLMs with human preferences using Proximal Policy Optimization (PPO) [11, 16, 26]. However, RLHF requires multiple components—including a reference model, reward model, critic, and a newly trained generative model—and relies heavily on human-in-the-loop feedback, resulting in substantial computational and labor costs. This limitation is particularly critical for precise tasks like path planning in autonomous driving.

To overcome these challenges, Group Relative Policy Optimization (GRPO) [10] has been proposed. GRPO enhances VLMs’ reasoning and arithmetic capabilities by learning directly from comparisons between the policy model and a reference model, without requiring an explicit reward model or critic. The authors advocate for rule-based reward signals, arguing that neural reward models may be vulnerable to reward hacking during large-scale reinforcement learning. Unlike standard or preference-based RL [22, 26], GRPO directly optimizes verifiable, task-specific metrics. This aligns the policy with downstream objectives, such as path planning, rather than linguistic mimicry, and reduces both human supervision and computational cost.

Building on GRPO, Reinforcement Learning with Verifiable Reward (RLVR) has been successfully applied to tasks where correctness can be objectively evaluated, including object detection, classification [19], mathematics [6], and code generation. RLVR has also been extended to robotics, such as manipulator control [29, 33] and autonomous driving [14, 21], where reward signals can be verified via simulations or pre-defined behavioral criteria. Theoretical analyses show that RLVR with GRPO improves the success rate of reward-maximizing outputs through a dynamic process converging to a fixed point and guarantees preference amplification over the reference model [20, 31].

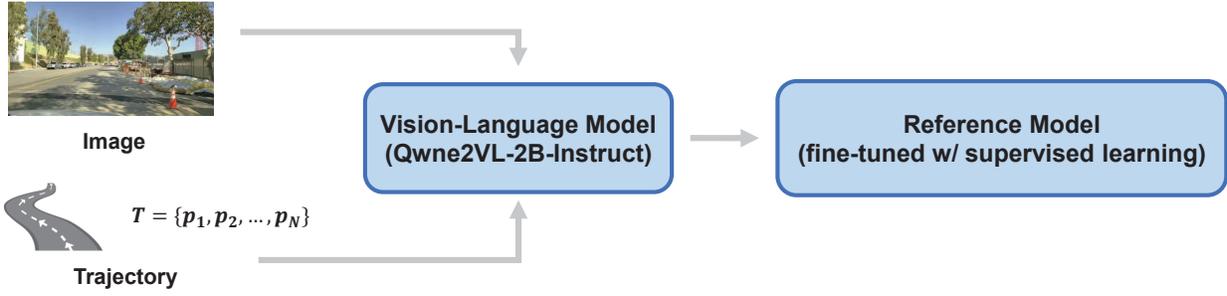
In this work, we extend RLVR with GRPO to autonomous driving by leveraging planning-oriented metrics, such as Average Displacement Error (ADE) and Final Displacement Error (FDE), as verifiable reward signals. Our approach fine-tunes VLMs to produce trajectories that are consistent with situational reasoning, improving both out-of-distribution generalization and alignment between language-based planning and action-level execution.

3. Methodology

3.1. Preliminary

Reinforcement Learning with Verifiable Rewards RLVR aims to optimize a policy π_θ by maximizing the expected

(a) Phase 1 : Supervised Fine-Tuning



(b) Phase 2 : Reinforcement Learning with Verifiable Reward

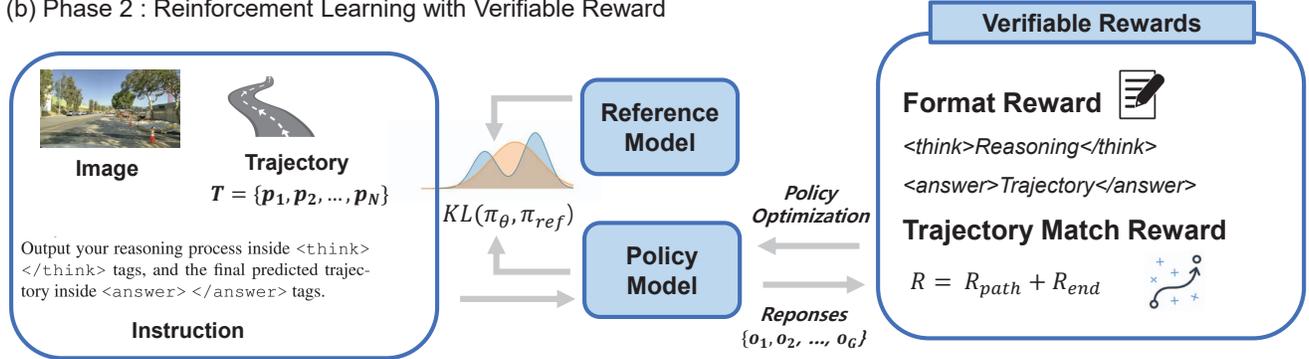


Figure 2. Overview of proposed method. (a) In Phase 1, the Vision-Language Model is fine-tuned with supervised learning using paired image-instruction-trajectory data. (b) In Phase 2, reinforcement fine-tuning with verifiable rewards based on format accuracy of responses and trajectory alignment. The policy model is optimized via KL divergence from a supervised reference model with group size of G .

reward while maintaining proximity to a reference policy π_{ref} to prevent over-optimization. The objective function is formulated as:

$$\max_{\pi_\theta} \mathbb{E}_{o \sim \pi_\theta(q)} [R_{RLVR}(q, o)] \quad (1)$$

$$= [R(q, o) - \beta \text{KL}(\pi_\theta(o | q) \| \pi_{ref}(o | q))] \quad (2)$$

where q represents the input query, o denotes the generated trajectory, $R(q, o)$ is the reward function, and β is a regularization coefficient that controls the trade-off between reward maximization and KL divergence penalty. The KL divergence term ensures that the learned policy does not deviate significantly from the reference policy, maintaining stability during training.

Group Relative Policy Optimization GRPO extends the Proximal Policy Optimization (PPO)[26] by incorporating group-relative advantage estimation. The advantage function is computed using group statistics to reduce variance:

$$\hat{A}_{i,t} = \frac{R_i - \text{mean}(\{R_i\}_{i=1}^G)}{\text{std}(\{R_i\}_{i=1}^G)} \quad (3)$$

where G represents the number of generated trajectories, i.e. the group size, R_i is the reward for the i -th trajectory,

and the advantage $\hat{A}_{i,t}$ is normalized using the mean and standard deviation of rewards within the group. This normalization helps stabilize training by reducing the impact of reward scale variations.

The GRPO objective function combines the clipped surrogate objective from PPO with the group-relative advantages:

$$\mathcal{J}_{GRPO} = \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left[\min \left\{ \frac{\pi_\theta(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{old}}(o_{i,t}|q, o_{i,<t})} \hat{A}_{i,t}, \right. \right. \\ \left. \left. \text{clip} \left(\frac{\pi_\theta(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{old}}(o_{i,t}|q, o_{i,<t})}, 1 - \varepsilon, 1 + \varepsilon \right) \hat{A}_{i,t} \right\} \right. \\ \left. - \beta \text{KL}[\pi_\theta \| \pi_{ref}] \right] \quad (4)$$

where $\pi_{\theta_{old}}$ represents the policy from the previous iteration, $|o_i|$ is the length of the i -th output response, ε is the clipping parameter, and the min operation implements the conservative policy update mechanism characteristic of PPO. The KL divergence term provides additional regularization to maintain training stability.

3.2. RLVR with Planning-Oriented Metrics

In this work, we apply GRPO to autonomous driving by designing a reward function that directly reflects planning performance, as shown in Fig. 2. We build upon the insight that GRPO enables learning pairwise preferences among multiple candidate trajectories. Unlike single-reward updates, this approach exposes the policy to richer and more diverse supervisory signals, providing multiple perspectives on the range of candidate trajectories. This diversity encourages the policy to learn robust behaviors that are not overfitted to specific situations, which ultimately improves generalization across a wide range of scenarios. Therefore, we define the reward using commonly adopted trajectory evaluation metrics: ADE and FDE.

$$R_{\text{planning}} = -\log\left(1 + \frac{1}{N} \sum_{i=1}^N \|\hat{p}_i - p_i\|_2\right) - \log(1 + \|\hat{p}_N - p_N\|_2) \quad (5)$$

The first term encourages the predicted trajectory $\hat{T} = \{\hat{p}_1, \hat{p}_2, \dots, \hat{p}_N\}$, where each $\hat{p}_i \in \mathbb{R}^2$ represents an (x, y) position in image plane to stay close to the ground-truth trajectory $T = \{p_1, p_2, \dots, p_N\}$, where each $p_i \in \mathbb{R}^2$ over the full horizon (ADE), while the second term penalizes deviations at the final timestep (FDE). We apply logarithmic smoothing for numerical stability and better learning dynamics. In addition to planning accuracy, we incorporate a formatting reward R_{format} , which encourages adherence to the expected output format with reasoning and response.

$$R = R_{\text{format}} + R_{\text{planning}} \quad (6)$$

4. Experiment

In-Domain Dataset. The ROADWork dataset [9] serves as an OOD benchmark focused on road construction scenarios. Among the entire dataset, 5,430 samples contain image, scene description, and a list of trajectories in image plane. To demonstrate that reinforcement fine-tuning can yield performance improvements even with data less than supervised fine-tuning, we divided the dataset into two subsets: 4,344 samples for supervised fine-tuning and 1,086 samples (20% of full samples) for reinforcement fine-tuning. Training details for supervised fine-tuning and reinforcement fine-tuning can be seen in the Appendix A.

Furthermore, inspired by KITTI [8], we split the dataset into two subsets based on the variance of x-coordinates in the trajectory. One subset comprises trajectories with low lateral variance representing straight trajectory (**Easy**). In contrast, the other subset comprises trajectories with higher lateral variance and more curved trajectories involving left

and right turns (**Hard**). Details of the dataset splitting procedure are provided in the Appendix B.

Out-of-Domain Dataset. The CODA-LM [2] extends the CODA dataset [17] which includes various corner cases such as road construction and adverse weather conditions by augmenting it with natural language captions. The generalization ability of our method under OOD conditions was assessed using the corner-case dataset not limited to construction zones.

What is the best baseline model? We choose Qwen2VL-2B-Instruct as our baseline. This decision is supported by existing studies [3, 14, 19] that integrate RLVR with diverse tasks and consistently adopt this model, thereby demonstrating its validity and robustness in multimodal learning contexts.

Table 1. Instruction for reinforcement fine-tuning and example of response.

Trajectory Prediction Prompt: Given only the image, predict the autonomous vehicle’s future trajectory as a list of 20 (x, y) image coordinates. You must output your reasoning process inside <think> </think> tags, and the final predicted trajectory inside <answer> </answer> tags. Strictly follow the output format below: <think> Reasoning based on visual cues in the image. </think>
<answer>[{'x': x1, 'y': y1}, ..., {'x': x20, 'y': y20}]</answer>

Example of Response: <think>tubular markers on right side of road. work vehicle on right side of road. worker on right sidewalk.</think><answer>[{'x': 578.88, 'y': 448.74}, {'x': 578.56, 'y': 442.6}, ..., {'x': 602.88, 'y': 335.34}]</answer>

Dataset and Prompt Design for Reinforcement Fine-Tuning. The dataset for reinforcement fine-tuning consists of a reasoning description enclosed within <think></think> tags, which includes visual reasoning process from the image, and a predicted trajectory represented as a list of 20 (x, y) image coordinates enclosed within <answer></answer> tags as shown in Tab. 1. Specifically, the <think></think> section provides the rationale behind the prediction, while the <answer></answer> section contains the autonomous

vehicle’s future trajectory as a sequence of coordinate pairs.

4.1. In-domain Evaluation

Table 2. Comparison of model performance on ROADWork dataset. ADE and FDE are presented as percentages by normalizing with image resolution for better readability. The models with supervised fine-tuning were trained on the full dataset of 5K samples, while LaViPlan was fine-tuned with 4K samples for supervised learning and an additional 1K samples for reinforcement fine-tuning. **Bolded** values indicate the best performance and N/A means no available result.

	ADE ↓		FDE ↓	
	Easy	Hard	Easy	Hard
<i>Baseline</i>				
Vision-Language Models				
Qwen2VL-2B	52.44	52.77	102.39	105.05
Qwen2VL-7B	60.73	60.71	66.61	67.57
Qwen2.5-VL-3B	16.37	16.40	20.60	20.77
LLaMA3.2-11B	59.27	58.88	74.16	71.44
Domain-Specific Models				
Senna	N/A	N/A	N/A	N/A
DriveLM (w/ LLaMA-Adapter)	37.10	38.40	56.99	56.90
<i>Supervised Fine-tuning</i>				
Vision-Language Models				
Qwen2VL-2B	4.52	5.66	4.46	6.46
Qwen2VL-7B	4.80	6.04	5.08	7.35
Qwen2.5-VL-3B	4.97	6.22	5.07	7.34
LLaMA3.2-Vision-11B	4.52	5.46	5.20	7.10
Domain-Specific Models				
Senna	5.71	5.73	6.58	7.46
DriveLM (w/ LLaMA-Adapter)	6.73	7.79	6.87	8.43
<i>Reinforcement Fine-tuning</i>				
LaViPlan (ours)	3.62	4.83	3.85	6.09

4.1.1. Overall Results

Tab. 2 shows overall result across multiple foundation models: baseline, supervised fine-tuning, and our proposed method. They are evaluated with ADE and FDE on both easy and hard subsets. In the baseline, VLMs without additional task-specific supervision lack the capability to generalize directly to trajectory planning. In contrast, supervised fine-tuning improves performance and our proposed method achieves the best performance across all metrics. Compared to supervised fine-tuning models, LaViPlan further reduces the errors. This demonstrates the effectiveness of reward optimization in refining planning-oriented outputs beyond what can be achieved with standard supervised learning. In summary, these results indicate that (1) VLMs are insufficient for planning tasks without supervised fine-tuning, (2) supervised fine-tuning is essential for grounding planning behaviors, and (3) RLVR with planning-oriented

metrics leads to further performance gains by explicitly optimizing for planning.

4.1.2. From Linguistic Consistency to Functional Reasoning.



Figure 3. Qualitative comparison of scene reasoning between the SFT (top) and LaViPlan (bottom) models across six scenarios in [9]. While SFT tends to produce more verbose, human-like descriptions that closely resemble natural annotations, LaViPlan emphasizes core, task-relevant elements such as cones, barriers, and work vehicles. This shift in reasoning reflects a transition from linguistic fidelity to planning-oriented abstraction, enabling the model to better prioritize hazards and navigable space—even at the cost of similarity to ground-truth descriptions.

We evaluate the linguistic effects of our method on the combined easy and hard subsets using BERTScore [37] and Natural Language Inference (NLI) [4], as shown in Tab. 3. Compared to supervised fine-tuning (SFT), reinforcement fine-tuning (LaViPlan) leads to a decrease in BERTScore, as well as a reduction in entailment and an increase in neutral and contradiction labels. While these results indicate a divergence from ground-truth human-written descriptions, they do not necessarily imply degraded reasoning ability. Rather, we observe a meaningful shift in the model’s reasoning style. As shown in Fig. 3, scene descriptions from LaViPlan tend to be more concise and hazard-focused emphasizing critical elements such as barriers, cones, and work vehicles in terms of planning. In contrast, SFT tends to replicate verbose and linguistically rich but often functionally redundant information. In several scenarios, LaViPlan omits less relevant descriptors (e.g., sidewalks or drum lines) in favor of spatial cues that better support trajectory-level decisions. This trade-off suggests that conventional language metrics may underrepresent reasoning utility in autonomous driving contexts. Despite lower alignment with human-like phrasing, LaViPlan yields improved planning performance, indicating that more effective and actionable

reasoning can emerge even as linguistic similarity declines. We advocate for a shift toward task-aware evaluation of vision-language models in safety-critical domains, where functional relevance should take precedence over semantic mimicry.

Table 3. Comparison of Scene Understanding. SFT refers to supervised fine-tuning and RFT refers to reinforcement fine-tuning

	SFT	LaViPlan	Δ
BERTScore \uparrow	0.748	0.712	-4.81%
Entailment (%) \uparrow	44.8	37.7	-15.85%
Neutral (%) \downarrow	25.3	29.1	+15.02%
Contradiction (%) \downarrow	29.9	33.1	+10.70%

4.2. Out-of-domain Evaluation

We evaluate the proposed method on CODA-LM [2], a zero-shot dataset composed of diverse corner cases involving out-of-distribution (OOD) road hazards. Since CODA-LM does not provide ground-truth trajectories, we evaluate the method based on 2D bounding box in the dataset.

4.2.1. Evaluation Metrics

To quantitatively assess the safety of the predicted trajectories, we evaluate three core indicators that capture a distinct aspect of safety. The detailed procedure for obtaining F, C, and P can be found in the Appendix C:

- **Fail Rate (F):** A trajectory is considered a failure if it invades any 2D bounding box in the scene. This reflects how often unsafe behaviors occur.
- **Collision Count (C):** The average number of bounding boxes that are violated per trajectory across all scenes. This captures the extent of interaction with multiple objects
- **Penetration Length (P):** When a collision occurs, this metric measures the extent of trajectory penetration into the bounding box area, providing a proxy for the severity of the failure. This quantifies the depth of intrusion, emphasizing the potential danger in collision.

Inspired by [12], we compute an overall safety score by applying min-max normalization to each metric, followed by a weighted average:

$$\text{SafetyScore}_i = \sum_{j \in \{F, C, P\}} w_j \cdot \left(1 - \frac{x_{ij} - \min_k x_{kj}}{\max_k x_{kj} - \min_k x_{kj}} \right) \quad (7)$$

The safety score for each model i is computed as a weighted sum over metrics $j \in \{F, C, P\}$ — failure rate, collision count, and penetration length — where each metric is min-max normalized across all models k . Then, we employ multiple weighting schemes to capture different evaluation perspectives. This approach enables a balanced and

comprehensive analysis of model robustness and trade-offs under various out-of-distribution scenarios.

- **Balanced** ($w_f = 0.4, w_c = 0.3, w_p = 0.3$): This evaluation reflects a balanced trade-off between safety and driving success under general OOD scenarios.
- **Safety-Focused** ($w_f = 0.3, w_c = 0.2, w_p = 0.5$): Each indicator is first min-max normalized to the $[0, 1]$ range, and then aggregated using predefined weights. This approach highlights the relative performance differences across models and is useful for identifying overall safety trends.
- **Performance-Focused** ($w_f = 0.5, w_c = 0.3, w_p = 0.2$): This scheme prioritizes driving success by assigning a high weight to the trajectory failure rate, while giving smaller weights to collision count and penetration length, respectively. It emphasizes task completion while still accounting for safety.
- **Equal Weight** ($w_f = 0.33, w_c = 0.33, w_p = 0.34$): All three safety indicators are weighted equally (approximately 33% each). This evaluation provides a neutral assessment without emphasizing any specific aspect, making it suitable for overall robustness analysis.

Table 4. Safety scores under different weighting schemes for failure (F), collision (C), and penetration (P).

Model	Balanced \uparrow	Safety-Focused \uparrow	Performance-Focused \uparrow	Equal \uparrow
Baseline	0.40	0.30	0.50	0.33
SFT (5k)	0.60	0.59	0.56	0.63
LaViPlan	0.64	0.73	0.56	0.70

Table 4 presents the safety scores of each model under four different weighting schemes. Across all evaluation settings, the proposed method consistently achieves the highest scores, indicating superior safety performance in both driving success and physical impact mitigation. In particular, it shows a significant advantage under the *Safety-Focused*, minimizing penetration depth and collisions. Although SFT (fine-tuned with whole dataset) performs relatively well, especially under the *Performance-Focused*, it lags behind LaViPlan in overall safety. The baseline model consistently scores the lowest, confirming its limited capability in handling OOD scenarios safely. The qualitative results of SFT and ours can be found in Appendix D.

4.3. Ablation Studies

4.3.1. Supervised vs. Reinforcement Fine-Tuning

We analyze the impact of reinforcement fine-tuning by comparing models with and without reinforcement-based optimization. The observed gains over supervised fine-tuning, as reported in Tab. 5, indicate that the reinforcement fine-tuning captures complementary learning signals beyond supervised losses. Interestingly, the effect size varies

Table 5. Impact of Reinforcement Fine-Tuning (RFT) on performance after Supervised Fine-Tuning (SFT), showing consistent improvements across easy and hard scenarios.

	ADE ↓		FDE ↓	
	Easy	Hard	Easy	Hard
Baseline	52.44	52.77	102.39	105.05
SFT (4k)	4.12	5.31	4.44	6.51
LaViPlan	3.62	4.83	3.85	6.09
Δ	-12.1%	-9.1%	-13.3%	-6.5%

with task complexity, implying that the reinforcement fine-tuning provides stronger gradient signals in low-uncertainty regimes while remaining effective under harder conditions.

4.3.2. Effective Sample Set for GRPO

To evaluate the impact of sample difficulty on reinforcement fine-tuning, we fix the SFT:RFT ratio at 4:1 (corresponding to 4,344 and 1,086 samples, respectively), and vary the proportion of easy and hard samples within this fixed budget.

In-Domain Dataset. As shown in Tab. 6, allocating a larger portion of hard scenarios (up to 40%) leads to consistent improvements across both ADE and FDE, particularly for hard cases. These results suggest that reinforcement fine-tuning benefits from challenging samples that provide richer learning signals, thereby enhancing the model’s generalization ability in complex planning situations. Notably, the 6:4 ratio achieves the best overall performance, as it strikes an effective balance between stable learning from easy cases and the additional supervision obtained from hard cases. While increasing the hard portion further (e.g., 7:3) introduces more diverse challenges, it can undermine the stability of learning in easy cases, resulting in less consistent improvements across metrics.

Table 6. Performance impact of different easy-to-hard data ratio when the total number of training samples is held constant with in-domain dataset.

Ratio	ADE ↓		FDE ↓	
	Easy	Hard	Easy	Hard
9:1	3.84 (-6.8%)	5.05 (-4.9%)	4.09 (-7.9%)	6.31 (-3.1%)
7:3	5.55 (+34.7%)	6.70 (+26.2%)	4.05 (-8.8%)	6.16 (-5.4%)
6:4	3.62 (-12.1%)	4.83 (-9.1%)	3.85 (-13.3%)	6.09 (-6.5%)

Out-of-Domain Dataset. The ground-truth trajectory does not always align with the safest path, and such misalignment becomes more prominent in difficult scenarios. In this setting, the 7:3 ratio yields the most favorable outcomes as shown in Tab. 7, especially under safety-focused and balanced metrics. This is because the larger share of hard samples exposes the model to safety-critical conflicts, encouraging it to learn beyond simply imitating labeled trajec-

tories and to acquire strategies that prioritize safe driving. However, further increasing the proportion of hard samples to 6:4 biases the training distribution excessively toward complex cases, which diminishes generalization in performance-focused metrics.

Table 7. Performance impact of different easy-to-hard data ratio with out-of-domain dataset and absolute difference from baseline.

Ratio / Model	Balanced ↑	Safety-Focused ↑	Performance-Focused ↑	Equal ↑
SFT (5K)	0.60 (+0.20)	0.59 (+0.29)	0.56 (+0.06)	0.63 (+0.30)
LaViPlan (9:1)	0.58 (+0.18)	0.62 (+0.32)	0.51 (+0.01)	0.63 (+0.30)
LaViPlan (7:3)	0.64 (+0.24)	0.73 (+0.43)	0.56 (+0.06)	0.70 (+0.37)
LaViPlan (6:4)	0.45 (+0.05)	0.49 (+0.19)	0.39 (-0.11)	0.51 (+0.18)

In short, our findings indicate that the effectiveness of easy-to-hard sampling is highly domain-sensitive, suggesting that in-domain performance benefits from a balanced exposure to both easy and hard cases, while out-of-domain robustness requires greater emphasis on hard cases to capture safety-critical cues.

4.3.3. Impact of Reasoning in Two-Phased Fine-tuning

We ablate the use of explicit reasoning in both supervised fine-tuning (phase 1) and subsequent reinforcement fine-tuning (phase 2) to investigate its effect, as shown in Tab. 8 and Tab. 9.

Reasoning Impact on Supervised Fine-Tuning. In supervised fine-tuning (phase 1), models trained with reasoning outperformed those without, although the margin was relatively small. Both versions achieved significant gains over the instruction-tuned baseline, with over 90% reductions in ADE and FDE across all subsets. This indicates that supervised fine-tuning plays a major role in grounding planning behavior, and the addition of reasoning further enhances the model’s ability to contextualize scenes.

Table 8. Ablation Study: Effect of Reasoning on Supervised Fine-Tuning (phase 1).

	ADE ↓		FDE ↓	
	Easy	Hard	Easy	Hard
Baseline	52.44	52.77	102.39	105.05
w/o reasoning	4.14 (-92.10%)	5.46 (-89.66%)	4.40 (-95.70%)	6.65 (-93.67%)
w/ reasoning	4.12 (-92.14%)	5.31 (-89.94%)	4.44 (-95.66%)	6.51 (-93.80%)

Reasoning Impact on Reinforcement Fine-Tuning. In reinforcement fine-tuning (phase 2), where reinforcement fine-tuning is applied on top of the already fine-tuned models from phase 1, the inclusion of reasoning continues to offer additional benefits. Although the overall improvements are more modest than those observed in phase 1, models with reasoning consistently outperform those without across both ADE and FDE. The performance gap is slightly larger in the hard subset, suggesting that reasoning becomes more useful in complex scenarios. These findings imply that

explicit reasoning, when preserved throughout both fine-tuning stages, provides stable gains by promoting structured understanding aligned with planning objectives.

Table 9. Ablation Study: Effect of Reasoning on Reinforcement Fine-Tuning (phase 2).

	ADE ↓		FDE ↓	
	Easy	Hard	Easy	Hard
Baseline	4.12	5.31	4.44	6.51
w/o reasoning	3.71 (-9.95%)	4.95 (-6.78%)	3.96 (-10.81%)	6.26 (-3.84%)
w/ reasoning	3.62 (-12.14%)	4.83 (-9.04%)	3.85 (-13.29%)	6.09 (-6.45%)

5. Discussion

Sparse Reward RLVR can guarantee performance improvement, but its efficiency may suffer when rewards are sparse. In our setting, the planning reward is sparse, as ADE is only computed after the entire rollout, emphasizing the importance of step-wise feedback during policy optimization. This sparsity may have limited GRPO’s potential, reducing the effectiveness of intermediate decision-making signals [36]. Moreover, sparse rewards can slow convergence and increase variance in policy updates, suggesting that integrating denser or auxiliary feedback could further improve learning efficiency and the quality of generated trajectories [23].

GRPO as Policy Regularizer, RLVR as Affordance. GRPO-based fine-tuning is not a silver bullet and relies on the presence of a reasonably strong SFT model as a prerequisite. This is because the policy model is optimized under a divergence constraint that keeps it close to the baseline model. Moreover, the reward function designed for reinforcement fine-tuning defines the affordance of the model, shaping which behaviors are reinforced. Therefore, beyond incorporating existing metrics into the reward design, future work should consider reward formulations that explicitly account for safety, success rate, and other critical factors, rather than solely imitating the ground-truth trajectory.

6. Conclusion

In this work, we present LaViPlan, a reinforcement learning framework guided by planning-oriented metrics for fine-tuning VLMs on autonomous driving tasks. Our findings indicate that RLVR enhances the zero-shot scene understanding capabilities of VLMs and may help mitigate misalignment between vision, language, and action. However, the extent to which these benefits generalize to world models—particularly in counterfactual reasoning under actions outside the training distribution in sequential decision-making tasks—remains an open question.

7. Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (RS-2023-00236245, Development of Perception/Planning AI SW for Seamless Autonomous Driving in Adverse Weather/Unstructured Environment)

References

- [1] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. 1
- [2] Kai Chen, Yanze Li, Wenhua Zhang, Yanxin Liu, Pengxiang Li, Ruiyuan Gao, Lanqing Hong, Meng Tian, Xinhai Zhao, Zhenguo Li, Dit-Yan Yeung, Huchuan Lu, and Xu Jia. Automated evaluation of large vision-language models on self-driving corner cases. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 7817–7826, 2025. 4, 6
- [3] Liang Chen, Lei Li, Haozhe Zhao, and Yifan Song. Vinci-r1-v: Reinforcing super generalization ability in vision-language models with less than \$3, 2025. 4
- [4] Yanran Chen and Steffen Eger. Menli: Robust evaluation metrics from natural language inference. *Transactions of the Association for Computational Linguistics*, 11:804–825, 2023. 5
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24185–24198, 2024. 1
- [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. 2
- [7] Haoyu Fu, Diankun Zhang, Zongchuang Zhao, Jianfeng Cui, Dingkan Liang, Chong Zhang, Dingyuan Zhang, Hongwei Xie, Bing Wang, and Xiang Bai. Orion: A holistic end-to-end autonomous driving framework by vision-language instructed action generation. *arXiv preprint arXiv:2503.19755*, 2025. 2
- [8] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 4
- [9] Anurag Ghosh, Shen Zheng, Robert Tamburo, Khiem Vuong, Juan Alvarez-Padilla, Hailiang Zhu, Michael Cardei, Nicholas Dunn, Christoph Mertz, and Srinivasa G Narasimhan. Roadwork: A dataset and benchmark for learning to recognize, observe, analyze and drive through work zones. In *ICCV*, 2025. 4, 5

- [10] 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
- [11] Zilin Huang, Zihao Sheng, Chengyuan Ma, and Sikai Chen. Human as ai mentor: Enhanced human-in-the-loop reinforcement learning for safe and efficient autonomous driving. *Communications in Transportation Research*, 4:100127, 2024. 2
- [12] Xiaosong Jia, Zhenjie Yang, Qifeng Li, Zhiyuan Zhang, and Junchi Yan. Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. *Advances in Neural Information Processing Systems*, 37:819–844, 2024. 6
- [13] Bo Jiang, Shaoyu Chen, Bencheng Liao, Xingyu Zhang, Wei Yin, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. Senna: Bridging large vision-language models and end-to-end autonomous driving. *arXiv preprint arXiv:2410.22313*, 2024. 1, 2
- [14] 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, 4
- [15] Bu Jin, Xinyu Liu, Yupeng Zheng, Pengfei Li, Hao Zhao, Tong Zhang, Yuhang Zheng, Guyue Zhou, and Jingjing Liu. Adapt: Action-aware driving caption transformer. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7554–7561. IEEE, 2023. 2
- [16] Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Shane Gu. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*, 2023. 2
- [17] 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*, pages 406–423. Springer, 2022. 4
- [18] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, pages 34892–34916. Curran Associates, Inc., 2023. 1
- [19] Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rl: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*, 2025. 1, 2, 4
- [20] Youssef Mroueh. Reinforcement learning with verifiable rewards: Grpo’s effective loss, dynamics, and success amplification. *arXiv preprint arXiv:2503.06639*, 2025. 2
- [21] Zhenghao Peng, Wenjie Luo, Yiren Lu, Tianyi Shen, Cole Gulino, Ari Seff, and Justin Fu. Improving agent behaviors with rl fine-tuning for autonomous driving. In *European Conference on Computer Vision*, pages 165–181. Springer, 2024. 2
- [22] 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. 2
- [23] Desik Rengarajan, Gargi Vaidya, Akshay Sarvesh, Dileep Kalathil, and Srinivas Shakkottai. Reinforcement learning with sparse rewards using guidance from offline demonstration. *arXiv preprint arXiv:2202.04628*, 2022. 8
- [24] Katrin Renz, Long Chen, Ana-Maria Marcu, Jan Hünemann, 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. 2
- [25] Katrin Renz, Long Chen, Elahe Arani, and Oleg Sinavski. Simlingo: Vision-only closed-loop autonomous driving with language-action alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 11993–12003, 2025. 2
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 2, 3
- [27] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15120–15130, 2024. 2
- [28] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *Computer Vision – ECCV 2024*, pages 256–274, Cham, 2025. Springer Nature Switzerland. 2
- [29] Zirui Song, Guangxian Ouyang, Mingzhe Li, Yuheng Ji, Chenxi Wang, Zixiang Xu, Zeyu Zhang, Xiaoqing Zhang, Qian Jiang, Zhenhao Chen, et al. Manipulm-r1: Reinforcement learning for reasoning in embodied manipulation with large vision-language models. *arXiv preprint arXiv:2505.16517*, 2025. 1, 2
- [30] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Yang Wang, Zhiyong Zhao, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivelm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*, 2024. 2
- [31] Milan Vojnovic and Se-Young Yun. What is the alignment objective of grpo? *arXiv preprint arXiv:2502.18548*, 2025. 2
- [32] Shihao Wang, Zhiding Yu, Xiaohui Jiang, Shiyi Lan, Min Shi, Nadine Chang, Jan Kautz, Ying Li, and Jose M Alvarez. Omnidrive: A holistic vision-language dataset for autonomous driving with counterfactual reasoning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22442–22452, 2025. 2
- [33] Jialong Wu, Shaofeng Yin, Ningya Feng, and Mingsheng Long. Rlv-r-world: Training world models with reinforcement learning. *arXiv preprint arXiv:2505.13934*, 2025. 2
- [34] 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. [2](#)
- [35] Zhenhua Xu, Yan Bai, Yujia Zhang, Zhuoling Li, Fei Xia, Kwan-Yee K Wong, Jianqiang Wang, and Hengshuang Zhao. Drivegpt4-v2: Harnessing large language model capabilities for enhanced closed-loop autonomous driving. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17261–17270, 2025. [2](#)
- [36] Jingyi Zhang, Jiaying Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv preprint arXiv:2503.12937*, 2025. [8](#)
- [37] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019. [5](#)
- [38] Xingcheng Zhou, Xuyuan Han, Feng Yang, Yunpu Ma, and Alois C Knoll. Opendrivevl: Towards end-to-end autonomous driving with large vision language action model. *arXiv preprint arXiv:2503.23463*, 2025. [2](#)

Appendix

A. Training Details

Table 1. Hyperparameters for language world model training.

Hyperparameter		
SFT	Global Batch size	128
	Batch size per GPU	4
	LoRA rank [2]	64
	LoRA α	16
	Epoch	1
	Learning rate	1×10^{-5}
	Weight decay	0.1
RLVR	Max response length	1024
	Batch size	128
	PPO (GRPO) mini batch size	4
	KL loss coefficient	0.04
	Group size	4
	Learning rate	5×10^{-6}
Sampling	Top- p	0.95
	Temperature	1.2
	Repetition Penalty	1.2

Through an empirical study, we observed that for effective policy model optimization using the GRPO algorithm [1], the responses of the reference model must exhibit sufficient diversity within a limited group size. Accordingly, we increased the repetition penalty and temperature slightly compared to the default settings.

B. Split Dataset

B.1. Dataset Construction and Splitting Strategy

To enable both supervised and reinforcement fine-tuning for visual path planning, we split the dataset into Supervised Fine-Tuning (SFT) and Reinforcement Fine-Tuning (RFT) subsets using a structured and variance-aware strategy. We first load the complete set of trajectory metadata D and determine the target ratios for SFT versus RFT (e.g., 4:1), as well as the distribution of turning and straight trajectories within each subset (e.g., 6:4 or 4:6 depending on the setting). For each trajectory $d \in D$, we calculate the variance of its x -coordinates, which serves as a proxy for trajectory

Algorithm 1 Splitting Dataset into SFT and RFT data for Visual Path Planning

- 1: Load entire trajectory metadata list D
 - 2: Compute desired split sizes for:
 - SFT vs. RFT (e.g., 4:1 ratio)
 - Straight vs. Turn trajectories (e.g., 6:4 or 4:6 depending on SFT/RFT)
 - 3: For each sample $d \in D$, compute variance of x -coordinates over its trajectory
 - 4: Sort all samples in descending order of x -variance
 - 5: Split into two groups:
 - Top- N samples with high variance \rightarrow turning samples
 - Remaining samples \rightarrow straight samples
 - 6: From turning samples:
 - Assign the first N_1 to RFT-turn
 - Assign the remaining N_2 to SFT-turn
 - 7: From straight samples:
 - Assign first M_1 to SFT-straight
 - Assign next M_2 to RFT-straight
 - 8: For each sample in the SFT and RFT sets:
 - Convert sample format into instruction-following format with:
 - An image reference
 - A natural language prompt
 - A reasoning + answer pair
 - Append to respective output list (SFT or RFT)
 - 9: Save both lists as JSON files
-

curvature. We then sort the trajectories in descending order of x -variance. Based on this ordering, we classify the top- N trajectories with high variance as turning trajectories and the remaining as straight trajectories. From these, we allocate subsets to SFT and RFT splits: a portion of turning trajectories to RFT-turn and the remainder to SFT-turn, and similarly, straight trajectories to SFT-straight and RFT-straight. Each selected sample is then converted into an instruction-following format consisting of an image reference, a natural language prompt, and a corresponding reasoning-answer pair. These formatted samples are saved as separate JSON files for the SFT and RFT datasets.

B.2. Validation Set Construction

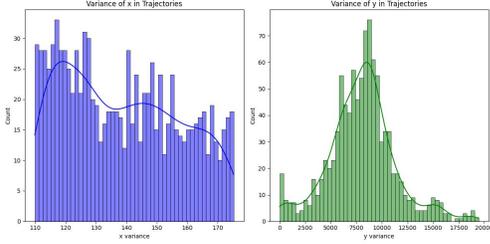


Figure 1. Distribution of trajectories in D_{easy} . trajectories exhibit lower x-variance than D_{hard} as can be seen in Fig. 1.

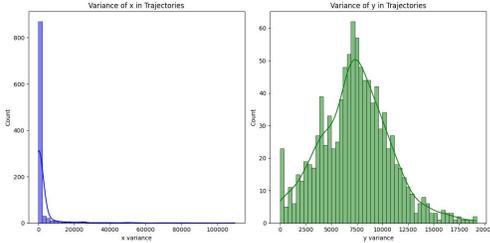


Figure 2. Distribution of trajectories in D_{hard} . trajectories exhibit lower x-variance than D_{easy} as can be seen in Fig. 2.

For evaluation, we construct validation sets focused on moderate and hard scenarios. We begin with a dense set of validation trajectories D_{dense} about 12K samples and remove any overlap with standard annotations $D_{standard}$ to obtain candidate validation samples D_{val} . We compute the x -variance for each sample and sort them accordingly. To form the easy set D_{easy} , we select a middle slice (e.g., 1K samples centered around the median x -variance). For the hard set D_{hard} , we randomly sample 0.7K trajectories from the top 70% (high-variance) and 0.3K from the bottom 10% (low-variance) of the sorted list. Both validation subsets are saved as JSON files for downstream evaluation.

Table 2. Summary of Validation Subsets Construction

Subset	Sampling Strategy	#Trajectories
Easy (D_{easy})	Median-centered 1K slice of x -variance	1K
Hard (D_{hard})	0.7K from top 70% + 0.3K from bottom 10%	1K

Algorithm 2 Construct Validation Sets (Easy, Hard)

- 1: Load dense validation trajectories D_{dense}
- 2: Load standard validation annotations $D_{standard}$
- 3: $D_{val} \leftarrow D_{dense} \setminus D_{standard}$
- 4: For each $d \in D_{val}$, compute variance of x over trajectory
- 5: $D_{sorted} \leftarrow$ sort D_{val} by x -variance (descending)
- 6: $N \leftarrow |D_{sorted}|$
- 7: **Easy Set:** $D_{easy} \leftarrow D_{sorted}[N/2 - 500 : N/2 + 500]$
- 8: **Hard Set:**
- 9: Randomly sample 700 from top 70% of D_{sorted}
- 10: Randomly sample 300 from bottom 10% of D_{sorted}
- 11: $D_{hard} \leftarrow$ combined samples
- 12: Save D_{easy} and D_{hard} as JSON

C. Pseudo code of OOD Evaluation

Algorithm 3 Procedure for Out-of-Distribution evaluation with as pseudo code.

Require: Dataset $\mathcal{D} = \{(I_i, B_i)\}_{i=1}^N$ where I_i : image, B_i : 2D bounding boxes

Require: Model \mathcal{M} predicting trajectory $\mathcal{T} = \{(x_j, y_j)\}_{j=1}^{20}$

Ensure: Fail Rate, Avg 2D BBox Collision, Avg Penetration Length

- 1: **Initialize:** $C_{fail} \leftarrow 0, C_{bbox} \leftarrow 0, L_{pen} \leftarrow 0, N \leftarrow 0, T_{bbox} \leftarrow 0$
- 2: **for** each (I, B) in \mathcal{D} **do**
- 3: $\mathcal{T} \leftarrow \mathcal{M}(I)$ ▷ 20-point trajectory
- 4: $traj \leftarrow$ LineString(\mathcal{T}), $c \leftarrow 0$
- 5: **for** each $b \in B$ **do**
- 6: $p \leftarrow$ Polygon(b)
- 7: **if** $traj.intersects(p)$ **then**
- 8: $c \leftarrow c + 1$
- 9: $L_{pen} \leftarrow L_{pen} + \text{length}(traj \cap p)$
- 10: **end if**
- 11: $T_{bbox} \leftarrow T_{bbox} + 1$
- 12: **end for**
- 13: **if** $c > 0$ **then** $C_{fail} \leftarrow C_{fail} + 1$
- 14: **end if**
- 15: $C_{bbox} \leftarrow C_{bbox} + c, N \leftarrow N + 1$
- 16: **end for**
- 17: **return** $\frac{C_{fail}}{N}, \frac{C_{bbox}}{N}, \frac{L_{pen}}{N}$

Trajectory Fail Rate: The proportion of predicted trajectories that intersect with at least one 2D bounding box in the scene:

$$\text{Fail Rate} = \frac{\text{Number of collision trajectories}}{\text{Total number of samples}} \quad (1)$$

Average 2D BBox Collision: The average number of

bounding boxes that each predicted trajectory collides with:

$$\text{Collision Count} = \frac{\text{Total bbox collisions}}{\text{Total number of samples}} \quad (2)$$

Penetration Length: The average geometric length of trajectory segments that penetrate into bounding boxes:

$$\text{Avg Penetration Length} = \frac{\sum_{i=1}^N \sum_{j=1}^{|B_i|} \text{length}(\mathcal{T}_i \cap \text{bbox}_j)}{\text{Total number of samples}} \quad (3)$$

where $\mathcal{T}_i \cap \text{bbox}_j$ represents the intersection between trajectory \mathcal{T}_i and bounding box bbox_j .

D. Qualitative Results on OOD Benchmark

We qualitatively evaluate the proposed method on CODA-LM, a zero-shot dataset composed of diverse corner cases involving out-of-distribution (OOD) road hazards. Since CODA-LM does not provide ground-truth trajectories, we rely on visual inspection of the predicted paths to assess plausibility, safety, and contextual awareness. As shown in Fig. 3, the model fine-tuned with RLVR produces more realistic and contextually appropriate trajectories compared to the one trained via supervised fine-tuning. In scenarios involving roadwork, blocked lanes, and ambiguous path constraints, the supervised model often generates overly linear or risk-prone trajectories that lack appropriate deviation or caution. In contrast, the proposed method consistently generates smoother, safer, and more context-aware trajectories, often adjusting path curvature to avoid obstacles such as cones, barriers, and vehicles. These results demonstrate the model’s capacity to generalize to previously unseen situations by aligning its reasoning and path generation with latent planning cues in the scene, even without explicit supervision. While the absence of quantitative metrics in CODA-LM limits numerical evaluation, these qualitative outcomes suggest that reinforcement fine-tuning enhances the model’s ability to adapt to driving contexts.

References

- [1] 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. 1
- [2] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 1



Figure 3. Qualitative results comparing supervised fine-tuning (top) and RLVR-based reinforcement fine-tuning (bottom) across diverse out-of-distribution (OOD) hazard scenarios. Examples include construction zones, adverse weather (e.g., rain and poor lighting), unexpected pedestrian behavior, and road obstacles. The proposed method generates more accurate and context-aware trajectories under complex conditions, indicating better robustness in real-world hazard cases.