

Ground What You See: Hallucination-Resistant MLLMs via Caption Feedback, Diversity-Aware Sampling, and Conflict Regularization

Miao Pan^{1*}, Wangjie Gan^{1*}, Jintao Chen^{1†}, Wenqi Zhang¹,
Sun Bing², Jianwei Yin¹, Xuhong Zhang^{1,3†}

¹School of Software Technology, Zhejiang University

²National Certification Technology (Hangzhou) Co., Ltd

³Ningbo Global Innovation Center, Zhejiang University
{chenjintao,zhangxuhong}@zju.edu.cn

Abstract

Multimodal large language models (MLLMs) have achieved significant results in various tasks, but their practical application is still severely constrained by hallucination issues, which are particularly prominent in reinforcement learning (RL) optimization processes. This paper systematically analyzes the causes of hallucinations in MLLM under RL training, identifying three key factors: (1) The model relies heavily on chained visual reasoning to guide decision-making during RL training. Thus, error and irrelevant information in visual reasoning can easily cause hallucinations, including inaccurate initial visual descriptions that anchor subsequent inferences to incorrect information, as well as redundant and broad inferential information; (2) Insufficient exploration diversity during the policy optimization phase, causing the model to output overly confident results; (3) The destructive conflict between different samples during optimization is a key factor that leads to false associations and unstable parameter updates. To address these issues, we propose a solution framework comprising three core modules. First, to improve the accuracy of visual localization, we add planning and caption stages before thinking and answer stages. To enhance initial visual descriptions ability, we allow LLMs to respond based solely on the caption and provide corresponding caption reward based on the quality of the response. Second, to enhance exploration capabilities, we classify samples based on the mean and variance of the reward distribution and select samples with high reward variance for training, thereby increasing the model’s focus on diverse samples. Finally, to mitigate conflicts between training samples, we identify neural tangent kernel (NTK) similarity as the key factor. Rather than minimizing it uniformly, we regulate NTK similarity by grouping sample pairs based on a similarity threshold. An InfoNCE loss is then applied to pull dissimilar pairs closer and push overly similar ones apart, guiding interactions toward a balanced range. The experimental results demonstrate that the proposed method significantly reduces the hallucination rate and effectively improves the inference accuracy of MLLMs.

Code — <https://github.com/omniAI-Lab/Ground-What-You-See.git>

*These authors contributed equally.

†They are co-corresponding authors

Introduction

Multimodal large language models (MLLMs) (Hurst et al. 2024; Team et al. 2024) have demonstrated strong performance across a range of multimodal tasks. However, despite their impressive performance, these models are prone to hallucination (Dona et al. 2025), producing responses that are linguistically fluent but factually inconsistent with the visual input (Leng et al. 2024). This limitation poses a significant challenge for real-world deployment, particularly in safety-critical applications.

While hallucination is a common issue across MLLMs, such problems can become even more pronounced during reinforcement learning (RL) optimization (Kaelbling, Littman, and Moore 2018). In particular, RL-based training tends to amplify *Flaw Repetition*, where models fall into loops of semantically redundant reasoning, and introduce *Think-Answer Mismatch*, where final answers deviate from preceding thought processes (Yao et al. 2025). To better understand the underlying factors behind these behaviors, we identify three major challenges that arise in RL training.

First, in multimodal domains, reasoning steps are essential as they guide the interpretation of visual information (Yao et al. 2025). RL training may introduce hallucinations due to **Visual Misinterpretation** (Christiano et al. 2017). Hallucination occurs when the initial visual description inaccurately reflects the visual content, leading the model to propagate this misinformation through subsequent reasoning and ultimately generate incorrect answers (Figure 1(a)-Q1). Another form arises when the model broadly or imprecisely attends to the input, leading it to produce redundant or irrelevant reasoning that is mistakenly treated as key evidence. This misattribution obscures truly informative cues and ultimately results in incorrect conclusions (Figure 1(a)-Q2).

Second, during RL training, **Limited Exploration** diversity (MLLM output distribution) may lead to hallucinations. As shown in Figure 1(b), samples with higher reward variance consistently achieve better generalization performance compared to those with lower reward variance. This suggests that focusing primarily on samples with lower reward variance tends to reinforce dominant output patterns and reduce output diversity. As a result, the model’s exploration

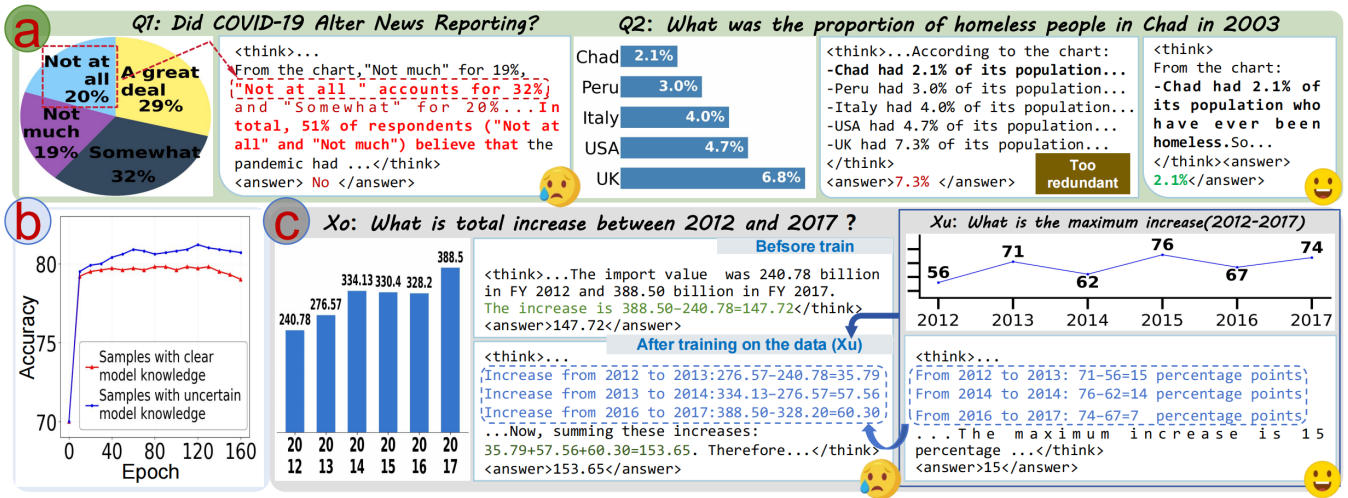


Figure 1: Three Types of Hallucination in RL-Tuned Multimodal LLMs: Visual Misinterpretation, Limited Exploration, and Sample Conflict

space becomes narrower, increasing the risk of hallucinations caused by overfitting and the over-inference leading to false associations, especially in ambiguous contexts.

Third, training **Samples Conflict** may contribute to hallucinations in MLLMs. As shown in Figure 1(c), updating the model on one sample x_u may unintentionally affect the model’s predictions on other, unrelated sample x_o . Such conflicts may introduce spurious associations between unrelated samples in the model’s reasoning process, leading to overconfident yet incorrect predictions for inputs. This significantly heightens the risk of hallucinations across diverse scenarios. However, the key factors underlying such conflicts remain underexplored in existing research.

To address these hallucination challenges, we introduce three key components. First, to ensure that model reasoning remains grounded in visual evidence, we add **planning and caption stages** before thinking and answer stages. In addition, we introduce a **Caption Reward**, which quantitatively measures the consistency between the generated caption and the visual input. This is achieved by verifying whether a language model can correctly answer the question using only the generated caption and the question itself. Secondly, to mitigate false correlations caused by insufficient exploration in RL, we analyzed how reward signals influence the model’s output distribution. We found that *positive rewards make the distribution more peaked, while negative rewards make the distribution flatter*. Based on this, we determined that medium samples (high reward variance) are the most valuable for training, as they enable effective exploration in RL, guiding the model from uncertainty toward confident predictions. To enhance learning diversity, we retained only these **high-variance samples** for RL updates. Finally, to mitigate sample conflicts, we identify **NTK similarity** (Jacot, Gabriel, and Hongler 2018) as the primary factor. When NTK similarity is too high, conflicting samples can exert excessive influence on each other; when too low, the influence among mutually beneficial samples be-

comes too weak. Therefore, instead of uniformly minimizing NTK similarity, we propose to regulate it by grouping sample pairs based on a threshold: pairs with excessively high similarity are treated as negatives, while those with low similarity are treated as positives. An **InfoNCE loss** is then applied to pull positive pairs closer and push negative pairs apart, guiding NTK similarity into a balanced range.

- We introduce a Visual-Grounded Reasoning mechanism by adding planning and caption stages before the thinking and answer stages. In addition, we incorporate a Caption Reward that quantifies visual-text alignment by verifying whether the caption alone enables a language model to answer the question correctly.
- We address hallucinations caused by limited exploration diversity in RL training by categorizing samples according to the mean and variance of their reward scores and selecting those with high reward variance for training.
- We introduce a Conflict-Aware Regularization mechanism by measuring the NTK similarity between samples. In addition, we incorporate an InfoNCE loss that guides the similarity values—determined by a threshold τ —toward an appropriate range.

Related work

Reinforcement Learning for MLLMs. Reinforcement learning is rapidly advancing multimodal large models, enhancing their reasoning and task performance. Researchers, inspired by the format and accuracy reward mechanisms in DeepSeek R1-Zero’s Equation 13, are now trying to apply this successful approach to multimodal tasks (Pan et al. 2025; Zhou et al. 2025; Meng et al. 2025; Liu et al. 2025c). Recent applications demonstrate diverse strategies for integrating RL, particularly for complex, multi-step problems and multimodal information. For instance, Vision-R1 (Huang et al. 2025) used Progressive Thinking Suppression Training (PTST) to extend Chain-of-Thought

(CoT) and decouple rewards, while R1-VL (Zhang et al. 2025) introduced StepGRPO with StepRAR and StepRVR for consistency in complex tasks. For multimodal integration, LMM-R1 (Peng et al. 2025) used a two-stage training (text then image-text), and R1-Onevision (Yang et al. 2025) formalized image descriptions to integrate visual information. In terms of cross-modal reward design, Video-R1 (Feng et al. 2025) added a temporal consistency reward for video understanding. Specialized applications include Reason-RFT, which customizes rewards for mathematical problems, Q-Insight (Li et al. 2025) integrated verifiable score accuracy for image quality assessment and additional (Lu et al. 2025; Liu et al. 2025b; Pan and Liu 2025; Kang et al. 2025; Deng et al. 2025). However, previous studies have largely ignored the illusions caused by visual extraction errors and how to design reward signals to guide the model to generate accurate and vision-based captions. This limits the effectiveness of reinforcement learning-based alignment in multimodal settings.

Hallucination in MLLMs. Hallucination, generating fluent but factually incorrect content, remains a critical challenge in large-scale models, particularly in multimodal and multi-step reasoning. In vision-language models, hallucination is amplified as attention to visual tokens degrades (Liu et al. 2025a), revealing a reliance on language priors. This aligns with findings that high-reward samples disproportionately influence updates, while semantically overlapping, moderately incorrect answers propagate misleading gradients (Zhu et al. 2025). Existing mitigation strategies include visual grounding via retrieval (Lee and Song 2025), hallucination-aware decoding (Leng et al. 2024), and representation-space filtering (Ghosh et al. 2024). Crucially, GRPO exacerbates hallucination in MLLMs by reinforcing semantically plausible but visually ungrounded responses, especially with sparse or ambiguous supervision.

Preliminary

Group Relative Policy Optimization (GRPO) (Shao et al. 2024). We apply GRPO to optimize the MLLM using fine-grained rewards. In each training iteration, for a sample q , the model π_θ generates G candidate responses, where each $y_i = (y_{i,1}, \dots, y_{i,T})$ is a sequence of tokens. For each candidate, a scalar reward r_i is assigned based on answer quality. The sequence-level reward is normalized to obtain a per-token advantage $A_{i,t} = \frac{r_i - \text{mean}(r)}{\text{std}(r)}$. The likelihood ratio $\rho_{i,t} = \frac{\pi_\theta(y_{i,t}|q, y_{i,<t})}{\pi_{\text{ref}}(y_{i,t}|q, y_{i,<t})}$ is used to compare the current and reference policies at each token. The loss is as follows:

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{q, \{y_i\}} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} [\beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) + \min(\rho_{i,t} \cdot A_{i,t}, \text{clip}(\rho_{i,t}, 1 - \varepsilon, 1 + \varepsilon) \cdot A_{i,t})], \quad (1)$$

where β controls the KL penalty and ε is the clipping threshold. This objective encourages the model to generate high-reward responses while constraining deviation from the reference policy.

Method

To address hallucination in RL training, we introduce three key components, illustrated in Figure 2. First, to ensure reasoning remains grounded in visual evidence, we extend the standard reasoning process by introducing two additional stages—planning and caption—before the thinking and answer stages. We further propose a Caption Reward that evaluates the consistency between generated captions and visual input by testing whether an LLM, given the caption and question, can produce the correct answer. Second, we observe that in RL training, positive rewards sharpen the output distribution, while negative rewards flatten it. This reveals that medium samples—those with high reward variance—are key to effective exploration. They initially induce flattening when the model is uncertain, and later lead to sharpening as predictions improve. This transition from exploration to confident exploitation reflects an ideal learning trajectory. To select high-quality samples that support this process, we compute the reward variance across multiple responses per input and retain only the high-variance samples for RL. Finally, we observe that harmful gradient conflicts are closely tied to NTK similarity (Jacot, Gabriel, and Hongler 2018) between training samples. To mitigate such interference without suppressing beneficial interactions, we regulate rather than eliminate NTK similarity. For each sample, we compute its NTK similarity with others and partition pairs based on a threshold. An InfoNCE loss is then applied to encourage alignment among dissimilar pairs and discourage convergence among overly similar ones, guiding NTK similarity into a balanced range.

Visual-Grounded Reasoning Enhancement

Redefine the Reasoning Paradigm. As illustrated in Figure 1(a), a primary source of hallucinations arises when the model fails to accurately localize relevant visual cues. To strengthen visually grounded reasoning, we add planning and caption stages before the thinking and answer stages. In the planning stage, we perform early localization of problem-relevant visual regions. The caption stage generates concise textual descriptions for these regions, providing focused intermediate guidance. The thinking stage incorporates these descriptions into multi-step reasoning, while the answer stage synthesizes them into the final response.

Caption Reward. To address the issue of hallucinations caused by inconsistencies between the early-generated text signals and the visual signals, we propose the *Caption Reward*. This mechanism ensures that the generated text signals are consistent with the visual signals, thereby preventing erroneous reasoning and mitigating hallucinations in later stages. First, we extract the textual caption generated by the model, which serves as the linguistic description of the visual input. This caption is then combined with the question and fed into a separate large language model (LLM) to generate a final answer. If the LLM is able to correctly answer the question only based on the caption, it implies that the model’s extracted caption is both accurate and effective in representing the visual content. If the LLM answers correctly using only the caption, a positive reward is given; otherwise, the model receives zero reward.

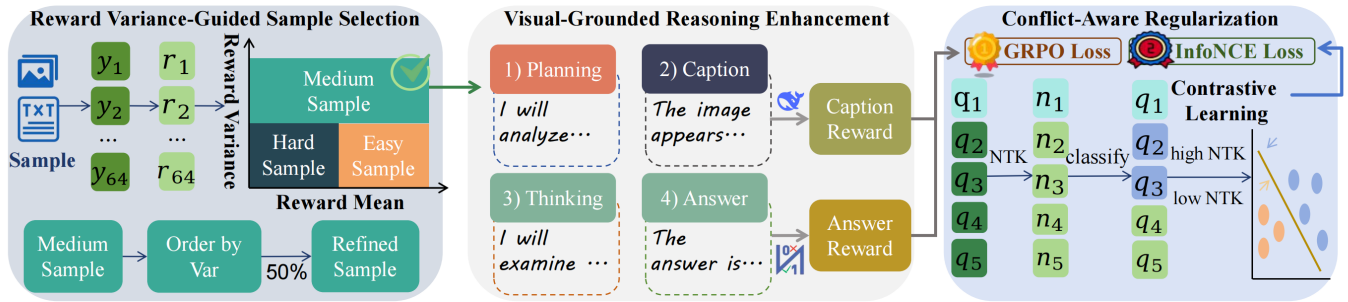


Figure 2: The proposed framework for robust visual reasoning, composed of three components: Reward Variance-Guided Sample Selection, Visual-Grounded Reasoning Enhancement, and Conflict-Aware Regularization.

Reward Variance-Guided Sample Selection

To investigate how reinforcement learning drives output diversity in MLLMs, we examine how different training samples influence the model’s distributional dynamics. Within a batch, samples naturally vary in how well the model initially understands them—reflected in their reward values. We categorize these samples into two types: **samples with clear model knowledge**, which receive high rewards, indicating that the model assigns high confidence to the correct answer; and **samples with uncertain model knowledge**, which receive low rewards, suggesting that the model incorrectly assigns high confidence to a wrong answer.

To quantify how the model’s output distribution changes for each sample, we generate an answer for every training instance, assign it a reward score based on output quality, and perform mean normalization to compute the per-sample advantage $A_{i,t}$. By taking the gradient of GRPO loss with respect to the logits z_v , we obtain:

$$\frac{\partial \mathcal{L}_i}{\partial z_v} \propto \begin{cases} -\pi_{y_t} \cdot (1 - \pi_{y_t}) \cdot A_{i,t} & \text{if } v = y_t \\ \pi_v \cdot \pi_{y_t} \cdot A_{i,t} & \text{if } v \neq y_t \end{cases} \quad (2)$$

where z_v denotes the logit of token v , π_v is its corresponding softmax probability, y_t is the sampled token at position t , and $A_{i,t}$ is the normalized advantage of sample i at time t . This expression reveals that the *sign* of $A_{i,t}$ governs the direction of change in the output distribution: when $A_{i,t} > 0$, corresponding to samples with clear model knowledge, the gradient update **sharpens the distribution**—it increases the probability of the sampled token while decreasing those of others. Conversely, when $A_{i,t} < 0$, corresponding to samples with uncertain model knowledge, the gradient update **flattens the distribution**—it lowers the probability of the sampled token while increasing those of unsampled tokens.

If a batch contains mostly high-reward samples, repeated sharpening can lead to *premature overfitting* and degraded generalization. Conversely, if dominated by low-reward samples, the model may enter a regime of *persistent uncertainty*, yielding overly flat distributions and ineffective exploration. The most valuable training signals arise from **samples with high reward variance**—those the model sometimes gets right and sometimes wrong. Initially, such samples receive negative rewards that flatten the distribution due to incorrect high-confidence outputs. As training progresses and predictions improve, positive rewards gradually

sharpen the distribution. This transition from exploration to confident exploitation reflects an ideal learning trajectory.

Sample Classification. For each input query, we generate 64 responses and compute the mean and variance of their computed reward scores. Based on these statistics, samples are categorized into three types: (i) easy samples with high mean and low variance, (ii) hard samples with low mean and low variance, and (iii) medium samples with high variance. To prioritize informative training signals, we use reward variance as a selection score and retain only the top 50% high-variance samples for RL training.

Conflict-Aware Regularization

To understand how updates on one sample influence others, we investigate how training samples (x_u, y_u) conflict with prior knowledge associated with x_o through the NTK-based learning dynamics. Specifically, we leverage formulation:

$$\begin{aligned} & \Delta \log \pi^t(\mathbf{y}_u | \mathbf{x}_o) \\ &= \eta \cdot \mathcal{A}^t(x_o) \cdot \mathcal{K}^t(x_o, x_u) \cdot \nabla_z \log \pi^t(y_u | x_u) \cdot A_{u,t} \quad (3) \\ &+ \mathcal{O}\left(\eta^2 \|\nabla_{\theta} \mathbf{z}(\mathbf{x}_u)\|_{\text{op}}^2\right), \end{aligned}$$

where η is the learning rate, and $\Delta \log \pi^t(\mathbf{y}_u | \mathbf{x}_o)$ denotes the change in log-probability at x_o after a policy update using sample (x_u, y_u) . Here, $\mathcal{A}^t(x_o) = I - 1 \cdot \pi^t(x_o)^\top$ is the Jacobian of log-probability w.r.t. logits, and $\mathcal{K}^t(x_o, x_u) = \nabla_{\theta} \mathbf{z}(x_o)^\top \nabla_{\theta} \mathbf{z}(x_u)$ is the empirical neural tangent kernel. The vector $\nabla_z \log \pi^t(y_u | x_u)$ denotes the derivative of the log-probability with respect to the logits z , and $A_{u,t}$ is the advantage. The second-order term is typically small under gradient clipping.

The conflict effect is primarily governed by the NTK term $\mathcal{K}^t(x_o, x_u)$. When the NTK term is high, even for unrelated queries, reward-driven updates at (x_u, y_u) can substantially alter the model’s predictions at x_o . Consequently, the model may begin to exhibit overgeneralization, where responses meant for specific contexts are inappropriately applied elsewhere. Notably, the NTK similarity is quite different from true semantic similarity. Therefore, model may produce confident yet incorrect outputs that are similar via NTK but unrelated in meaning.

Conflict Mitigation via InfoNCE Loss. To mitigate such unintended conflict, we propose to regulate NTK similarity

rather than eliminate it entirely. Excessively reducing NTK similarity between training samples may suppress beneficial interactions and hinder generalization. To strike a balance, we aim to regulate NTK-based similarity rather than minimize it uniformly. Since the NTK is difficult to compute directly, we approximate NTK similarity using the cosine similarity between the log-probability gradients of two samples at the final layer. For each sample, we compare it against other samples in the same batch and classify them into positive or negative pairs based on a threshold τ : samples with lower similarity are treated as positives (encouraged to **align**), while those with excessive similarity are treated as negatives (encouraged to **diverge**). The InfoNCE loss is then applied to adjust these relations, guiding the NTK similarity values into a balanced range that reduces harmful interference while preserving cooperative signals.

$$\mathcal{L} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\sum_{j \in \mathcal{P}(i)} \exp(\text{sim}(f_i, f_j))}{\sum_{k=1}^B \mathbb{I}_{[k \neq i]} \exp(\text{sim}(f_i, f_k))}, \quad (4)$$

where f_i is log-prob gradient with respect to the final layer of sample x_i and $\text{sim}(\cdot, \cdot)$ denotes cosine similarity. For a sample i , $\mathcal{P}(i)$ is the set of unsimilar samples for i .

Experiment

Experiment Setup

Benchmark. To comprehensively evaluate the multimodal reasoning, perception, and hallucination resistance of large models, we adopt a diverse set of benchmarks covering spatial, temporal, and mathematical capabilities. MMVU (Zhao et al. 2025) evaluates expert-level multidisciplinary video understanding. VideoMMMU (Hu et al. 2025) assesses knowledge-based question-answering over time. VideoHalluciner (Wang et al. 2024) and POPE (Li et al. 2023) evaluate resistance to hallucinations in video and image understanding, respectively. MathVista (Lu et al. 2023) presents problems requiring logical, algebraic, and scientific inference, including quantitative reasoning. MMBench (Liu et al. 2024) comprehensively assesses perception and cognitive abilities. This suite jointly evaluates models’ visual, temporal, analytical, and hallucination-resistant capabilities across diverse scenarios.

Baseline. To evaluate our method, we compare it against a range of representative baselines. GPT-4o (Hurst et al. 2024) and Gemini-1.5-Pro (Team et al. 2024) serve as strong closed-source performance references. Open-source video models, including LLaVA-OneVision-7B (Li et al. 2024) and VILA-1.5-8B, are relevant for temporal visual understanding. General-purpose vision-language foundations such as MiniCPM-V2.6-8B (Yao et al. 2024), InternVL2.5-8B (Chen et al. 2024), and LLaMA3.2-11B (Grattafiori et al. 2024) are used as base models. Reasoning-focused baselines like R1-OneVision-7B (Yang et al. 2025), R1-VL (Zhang et al. 2025), Vision-R1-7B (Huang et al. 2025), and Video-R1-7B (Feng et al. 2025) incorporate structured prompting or alignment for improved multi-step reasoning. Finally, we include our models: the base Qwen-VL-2.5-7B (Bai et al.

2025), its SFT and GRPO variants, and our enhanced version (Ours), to analyze the impact of visual grounding and contrastive regularization.

Main Result

As shown in Table 1, our model achieves state-of-the-art performance among all open-source models across image, video, and hallucination benchmarks, particularly excelling in video reasoning tasks. On MMVU—a knowledge-intensive video benchmark—our model attains 65.6% accuracy, surpassing all open-source models and ranking second only to GPT-4o, highlighting its strong video inference capabilities. While performance on MathVista (focused on abstract mathematical reasoning) is moderate, our model delivers leading results on image-heavy benchmarks such as MMBench (88.6%), demonstrating broad effectiveness in visual understanding.

In addition, standard GRPO training reduces performance on the hallucination benchmark, demonstrating that RL can introduce hallucinations. Our model achieves top results on VideoHalluciner (50.8%), indicating improved faithfulness and reduced hallucination risks compared to prior methods. Notably, supervised fine-tuning brings limited improvements, whereas RL yields consistent gains across all reasoning benchmarks, confirming its importance in enhancing both reasoning ability and response reliability.

To complement the main results, we evaluate whether the Caption Reward improves visual grounding, whether selecting moderate samples enhances exploration diversity in RL, and whether InfoNCE loss mitigates interference between conflicting samples.

Effect of Caption Reward on Hallucination Reduction

To evaluate the effectiveness of Caption Reward, we conduct an ablation study comparing models trained with and without this component across standard multimodal reasoning benchmarks. In particular, we aim to quantify its impact on hallucination reduction and visual grounding.

The results, as shown in Table 2, indicate that the model achieves the best performance when adopting our proposed output stages and incorporating the Caption Reward. Conversely, performance decreases when the Caption Reward is not applied, highlighting its effectiveness. Additionally, when the model does not follow our proposed output stages, performance further declines, underscoring the importance of the output stages we introduced.

To better understand how reward signals evolve during RL training, we visualize both Caption and Answer Rewards on medium and hard samples in Figure 3. Caption Reward starts low, indicating weak visual grounding, but rises steadily on medium samples and gradually aligns with Answer Reward—suggesting that learning better captions directly supports accurate answering. In contrast, Caption Reward on hard samples improves slowly and remains low, limiting its effectiveness as a learning signal.

These trends highlight the central role of Caption Reward: it not only reflects visual grounding quality but also acts as

Category	Model	MMVU	VideoMMMU	VideoHalluciner	MathVista	POPE	MMBench(en)
Closed-source	GPT-4o	75.4	61.2	53.3	61.3	86.9	83.4
	Gemini-1.5-Pro	-	60.6	37.8	63.9	-	73.9
Open-source Video Models	LLaVA-OneVision-7B	49.2	31.2	44.6	62.6	86.4	81.7
	VILA-1.5-8B	31.5	20.8	14.8	36.7	70.6	57.6
Open-source Base Models	MiniCPM-V2.6-8B	52.4	49.8	48.4	60.6	84.4	81.5
	InternVL2.5-8B	54.9	44.2	50.5	64.4	85.9	84.6
	LLaMA3.2-11B	-	-	-	51.5	86.2	65.8
Open-source Reasoning Models	R1-OneVision-7B	55.2	44.1	42.4	64.1	83.1	82.3
	R1-VL-7B	59.7	42.9	43.6	63.5	85.7	86.2
	Vision-R1-7B	57.6	39.7	44.2	73.5	86.4	83.8
	Video-R1-7B	62.5	48.1	45.1	71.0	85.5	86.5
Our Models	Qwen-VL-2.5-7B	57.6	43.9	46.5	63.7	84.4	86.3
	Qwen-VL-2.5-7B-SFT	62.7	46.0	43.5	54.7	82.2	83.9
	Qwen-VL-2.5-7B-GRPO	62.1	47.3	42.3	69.3	83.6	86.8
	Qwen-VL-2.5-7B-Ours	65.6	50.0	50.8	69.4	88.7	88.6

Table 1: Benchmark Performance Comparison of Closed-source, Open-source, and Proposed Models across Video, Reasoning, Math, and Hallucination-related Tasks. Bolded scores indicate the best performance among all open-source and our models.

Method	MMVU	VideoHalluciner	POPE
w/o Caption	62.6	47.6	85.2
w/o Caption Reward	63.3	48.6	86.6
w/ (Caption + Caption Reward)	65.6	50.8	88.7

Table 2: Impact of Caption and Caption Reward on Qwen-VL-2.5-7B Performance over MMVU, VideoHalluciner and POPE.

a guiding signal that shapes the answer distribution. Its improvement precedes and facilitates gains in Answer Reward, especially on learnable (medium) samples, underscoring its value in promoting stable and effective multimodal learning.

RL Exploration Diversity of Different Sample Types

Sample Type	MMVU	VideoHalluciner	POPE
Full Sample	64.3	49.4	87.5
Easy Sample	62.2	47.7	86.6
Hard Sample	63.5	48.4	88.2
Medium Sample	65.6	50.8	88.7

Table 3: Impact of Sample Type on Qwen-VL-2.5-7B Performance over MMVU, VideoHalluciner and POPE.

Table 3 indicates that the model trained with medium samples performs the best on all benchmarks, achieving results that are comparable to, or even superior to, those achieved by models trained on the full dataset. In contrast, training with easy samples leads to moderate performance due to the model’s tendency to overfit, as it has a strong grasp of these samples. On the other hand, training with hard sam-




Figure 3: Caption and Answer Reward curves during RL training on medium and hard samples.

ples results in poor performance because the model struggles to comprehend these samples, preventing meaningful knowledge acquisition. For medium samples, however, the model is at a critical threshold where it can neither fully grasp nor completely fail to understand them. Through further reinforcement learning, the model successfully learns new knowledge from this subset, leading to superior performance.

To better understand how training samples influence model behavior, we analyze policy entropy as a proxy for action diversity. We randomly select 100 ChartQA (Masry et al. 2022) test examples and compute the average token-level entropy for each. As shown in Figure 4, we compare models trained on easy, hard, and medium samples. For the medium category, we further sort samples by reward variance and construct three subsets containing the top 30%,

Method	MMVU	VideoHalluciner	POPE
w/o InfoNCELoss	63.8	48.3	86.8
w/ InfoNCELoss	65.6	50.8	88.7

Table 4: Impact of InfoNCELoss on Qwen-VL-2.5-7B Performance over MMVU, VideoHalluciner and POPE.

60%, and 100% high-variance samples to examine the impact of targeted sample selection.




Figure 4: Policy entropy distribution across models trained on easy, medium, and hard samples.

As illustrated in Figure 4, models trained solely on easy samples produce low-entropy outputs, indicating overconfident and potentially overfitted policies. Conversely, models trained on hard samples exhibit high entropy with large variability, suggesting unstable or unfocused decision-making due to the complexity of training inputs. Notably, the medium-sample models yield moderate entropy with a tighter distribution, indicating a balance between diversity and reliability. Moreover, models trained on high-variance subsets (30% and 60%) of medium samples achieve comparable or more stable entropy patterns than using the full medium set, suggesting that carefully selecting uncertain yet learnable samples within the medium group can further enhance training efficiency and policy quality.

Evaluating the Effectiveness of InfoNCE Loss in Reducing Sample Conflicts

As shown in Table 4, incorporating InfoNCE Loss consistently improves performance across all benchmarks. These improvements highlight the effectiveness of InfoNCE in enhancing model robustness against hallucination.

To better understand the mechanism behind the improvement, we first analyze the impact of the NTK similarity threshold τ used to guide the InfoNCE loss. As shown in Figure 5(a), model accuracy peaks when $\tau = 0.54$, indicating that this value provides a balanced criterion for distinguishing beneficial and conflicting training interactions. τ guides the InfoNCE Loss in learning to distinguish between samples that facilitate or hinder learning.

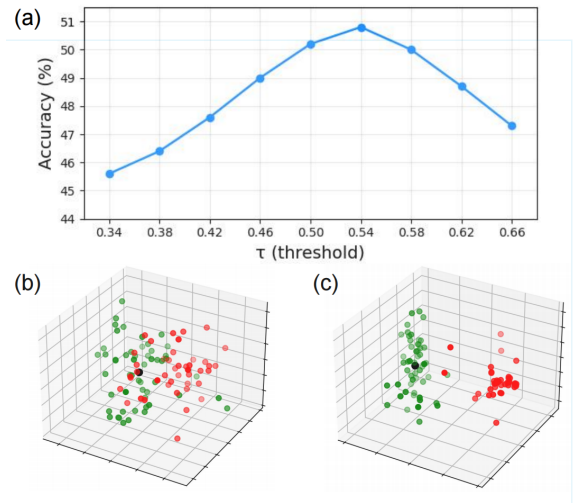


Figure 5: (a) Accuracy under different NTK thresholds τ . (b) and (c) Reshaping Effect on Final Layer Sample Representations Before and After InfoNCE Application.

We then investigate how this NTK-guided contrastive loss reshapes the model’s internal representations. Specifically, we randomly select 129 samples and designate one as the anchor. Based on their NTK similarity to the anchor and the chosen threshold τ , we construct positive and negative pairs for the InfoNCE Loss. Figure 5(b) and (c) show the final-layer representations before and after training. Prior to training, the samples are dispersed without clear structure. After training, cooperative samples (green), treated as positives—form a compact neighborhood around the anchor, while conflicting samples (red)—treated as negatives—are pushed apart. This structured separation emerges from contrastive optimization rather than from the initial threshold itself, and reflects improved functional organization and generalization.

Conclusion

This work presents a systematic study of hallucination in MLLMs trained with RL. We identify three primary causes—visual representation errors, limited exploration diversity, and sample-level conflicts—that undermine the reliability of model reasoning. To tackle these challenges, we introduce three components: a Caption Reward to enhance visual grounding, a reward-variance-based sampling strategy to encourage diverse exploration, and contrastive regularization to reduce interference between training samples. Our approach enhances visual faithfulness, improves robustness, and stabilizes learning dynamics of MLLMs without compromising generalization. Extensive evaluations on image and video benchmarks demonstrate that the proposed framework consistently reduces hallucination and strengthens multimodal reasoning under RL training.

Acknowledgments

This work is supported by National Key R&D Program of China under Grant No. 2024YFB3908400. This work is also supported by the Key R&D Program of Ningbo under Grant No.2024Z115.

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Deng, H.; Zou, D.; Ma, R.; Luo, H.; Cao, Y.; and Kang, Y. 2025. Boosting the generalization and reasoning of vision language models with curriculum reinforcement learning. *arXiv preprint arXiv:2503.07065*.
- Dona, M. A. M.; Cabrero-Daniel, B.; Yu, Y.; and Berger, C. 2025. BetterCheck: Towards Safeguarding VLMs for Automotive Perception Systems. *arXiv preprint arXiv:2507.17722*.
- Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wu, J.; Zhang, X.; Wang, B.; and Yue, X. 2025. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*.
- Ghosh, S.; Evuru, C. K. R.; Kumar, S.; Tyagi, U.; Nieto, O.; Jin, Z.; and Manocha, D. 2024. Visual description grounding reduces hallucinations and boosts reasoning in vlms. *arXiv preprint arXiv:2405.15683*.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hu, K.; Wu, P.; Pu, F.; Xiao, W.; Zhang, Y.; Yue, X.; Li, B.; and Liu, Z. 2025. Video-MMMU: Evaluating Knowledge Acquisition from Multi-Discipline Professional Videos. *arXiv preprint arXiv:2501.13826*.
- Huang, W.; Jia, B.; Zhai, Z.; Cao, S.; Ye, Z.; Zhao, F.; Xu, Z.; Hu, Y.; and Lin, S. 2025. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jacot, A.; Gabriel, F.; and Hongler, C. 2018. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31.
- Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. 2018. Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 237–285.
- Kang, H.; Sachdeva, E.; Gupta, P.; Bae, S.; and Lee, K. 2025. GFlowVLM: Enhancing Multi-step Reasoning in Vision-Language Models with Generative Flow Networks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3815–3825.
- Lee, J.; and Song, M. 2025. Retrieval Visual Contrastive Decoding to Mitigate Object Hallucinations in Large Vision-Language Models. *arXiv preprint arXiv:2505.20569*.
- Leng, S.; Zhang, H.; Chen, G.; Li, X.; Lu, S.; Miao, C.; and Bing, L. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13872–13882.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Li, W.; Zhang, X.; Zhao, S.; Zhang, Y.; Li, J.; Zhang, L.; and Zhang, J. 2025. Q-insight: Understanding image quality via visual reinforcement learning. *arXiv preprint arXiv:2503.22679*.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J.-R. 2023. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Liu, C.; Xu, Z.; Wei, Q.; Wu, J.; Zou, J.; Wang, X. E.; Zhou, Y.; and Liu, S. 2025a. More Thinking, Less Seeing? Assessing Amplified Hallucination in Multimodal Reasoning Models. *arXiv preprint arXiv:2505.21523*.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; et al. 2024. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, 216–233. Springer.
- Liu, Y.; Peng, B.; Zhong, Z.; Yue, Z.; Lu, F.; Yu, B.; and Jia, J. 2025b. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *arXiv preprint arXiv:2503.06520*.
- Liu, Z.; Zhang, Y.; Liu, F.; Zhang, C.; Sun, Y.; and Wang, J. 2025c. OThink-MR1: Stimulating multimodal generalized reasoning capabilities via dynamic reinforcement learning. *arXiv preprint arXiv:2503.16081*.
- Lu, P.; Bansal, H.; Xia, T.; Liu, J.; Li, C.; Hajishirzi, H.; Cheng, H.; Chang, K.-W.; Galley, M.; and Gao, J. 2023. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*.
- Lu, Z.; Chai, Y.; Guo, Y.; Yin, X.; Liu, L.; Wang, H.; Xiao, H.; Ren, S.; Xiong, G.; and Li, H. 2025. UI-R1: Enhancing Efficient Action Prediction of GUI Agents by Reinforcement Learning. *arXiv preprint arXiv:2503.21620*.
- Masry, A.; Long, D. X.; Tan, J. Q.; Joty, S.; and Hoque, E. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.

Meng, F.; Du, L.; Liu, Z.; Zhou, Z.; Lu, Q.; Fu, D.; Shi, B.; Wang, W.; He, J.; Zhang, K.; et al. 2025. Mm-eureka: Exploring visual aha moment with rule-based large-scale reinforcement learning. *CoRR*.

Pan, J.; Liu, C.; Wu, J.; Liu, F.; Zhu, J.; Li, H. B.; Chen, C.; Ouyang, C.; and Rueckert, D. 2025. Medvlm-r1: Incentivizing medical reasoning capability of vision-language models (vlms) via reinforcement learning. *arXiv preprint arXiv:2502.19634*.

Pan, Z.; and Liu, H. 2025. Metaspatial: Reinforcing 3d spatial reasoning in vlms for the metaverse. *arXiv preprint arXiv:2503.18470*.

Peng, Y.; Zhang, G.; Zhang, M.; You, Z.; Liu, J.; Zhu, Q.; Yang, K.; Xu, X.; Geng, X.; and Yang, X. 2025. Lmm-r1: Empowering 3b lmms with strong reasoning abilities through two-stage rule-based rl. *arXiv preprint arXiv:2503.07536*.

Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.

Team, G.; Georgiev, P.; Lei, V. I.; Burnell, R.; Bai, L.; Gulati, A.; Tanzer, G.; Vincent, D.; Pan, Z.; Wang, S.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

Wang, Y.; Wang, Y.; Zhao, D.; Xie, C.; and Zheng, Z. 2024. Videohalluciner: Evaluating intrinsic and extrinsic hallucinations in large video-language models. *arXiv preprint arXiv:2406.16338*.

Yang, Y.; He, X.; Pan, H.; Jiang, X.; Deng, Y.; Yang, X.; Lu, H.; Yin, D.; Rao, F.; Zhu, M.; et al. 2025. R1-onevision: Advancing generalized multimodal reasoning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*.

Yao, Y.; Yu, T.; Zhang, A.; Wang, C.; Cui, J.; Zhu, H.; Cai, T.; Li, H.; Zhao, W.; He, Z.; et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.

Yao, Z.; Liu, Y.; Chen, Y.; Chen, J.; Fang, J.; Hou, L.; Li, J.; and Chua, T.-S. 2025. Are Reasoning Models More Prone to Hallucination? *arXiv preprint arXiv:2505.23646*.

Zhang, J.; Huang, J.; Yao, H.; Liu, S.; Zhang, X.; Lu, S.; and Tao, D. 2025. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv preprint arXiv:2503.12937*.

Zhao, Y.; Zhang, H.; Xie, L.; Hu, T.; Gan, G.; Long, Y.; Hu, Z.; Chen, W.; Li, C.; Xu, Z.; et al. 2025. Mmvu: Measuring expert-level multi-discipline video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8475–8489.

Zhou, H.; Li, X.; Wang, R.; Cheng, M.; Zhou, T.; and Hsieh, C.-J. 2025. R1-Zero's "Aha Moment" in Visual Reasoning on a 2B Non-SFT Model. *arXiv preprint arXiv:2503.05132*.

Zhu, X.; Xia, M.; Wei, Z.; Chen, W.-L.; Chen, D.; and Meng, Y. 2025. The Surprising Effectiveness of Negative Reinforcement in LLM Reasoning. *arXiv preprint arXiv:2506.01347*.