

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 VIPER: EMPOWERING THE SELF-EVOLUTION OF VI- SUAL PERCEPTION ABILITIES IN VISION-LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

The limited capacity for fine-grained visual perception presents a critical bottleneck for Vision-Language Models (VLMs) in real-world applications. Addressing this is challenging due to the scarcity of high-quality data and the limitations of existing methods: supervised fine-tuning (SFT) often compromises general capabilities, while reinforcement fine-tuning (RFT) prioritizes textual reasoning over visual perception. To bridge this gap, we propose a novel two-stage task that structures visual perception learning as a coarse-to-fine progressive process. Based on this task formulation, we develop **ViPER**, a self-bootstrapping framework specifically designed to enable iterative evolution through self-critiquing and self-prediction. By synergistically integrating image-level and instance-level reconstruction with a two-stage reinforcement learning strategy, ViPER establishes a closed-loop training paradigm, where internally synthesized data directly fuel the enhancement of perceptual ability. Applied to the Qwen2.5-VL family, ViPER produces the **Qwen-Viper** series. With an average gain of **1.7%** on seven comprehensive benchmarks spanning various tasks and up to **6.0%** on fine-grained perception, Qwen-Viper consistently demonstrates superior performance across different vision-language scenarios while maintaining generalizability. Beyond enabling self-improvement in perceptual capabilities, ViPER provides concrete evidence for the reciprocal relationship between generation and understanding, a breakthrough to developing more autonomous and capable VLMs.

1 INTRODUCTION

Multimodal Large Language Models (MLLMs) have exhibited strong capabilities and wide-ranging application value across diverse domains (Caffagni et al., 2024). In contrast to Large Language Models (LLMs), MLLMs extend beyond text to multiple modalities, enabling them to handle more complex, information-rich tasks and establishing them as critical for embodied AI and world models (Yin et al., 2024). As an important type of MLLMs, Vision-Language Models (VLMs) are typically architected with a visual encoder to process pixel-level information and a language model backbone for semantic understanding, establishing a connection between the two modalities (Li et al., 2024; Wang et al., 2025b). Consequently, the overall performance of a VLM depends on the coordinated functioning of its visual encoder and language backbone.

Recent progress, such as o1 (El-Kishky, 2024) and Deepseek-R1 (Ren et al., 2025), has advanced the reasoning ability of LLMs, leading to increased focus on “slow thinking” (Wang et al., 2025c). Accordingly, the field has led to the development of multimodal reasoning models (Huang et al., 2025a; Li et al., 2025). Although Chain-of-Thought methods (Wei et al., 2022) have improved the reasoning ability of these models, the expressive limitations of the text present inherent constraints (Lyu et al., 2024; Stechly et al., 2024). For many vision-language tasks, especially those demanding fine-grained perception, performance cannot be attributed solely to the power of the language model backbone. These tasks hinge critically on sophisticated visual understanding coupled with language reasoning, a dual requirement that poses a significant challenge for current VLMs (Shao et al., 2024; Hu et al., 2024). This interdependence means that post-training efforts which enhance only the visual or linguistic component in isolation yield marginal improvements, underscoring the need for approaches that enable their co-evolution.

Several recent studies have proposed task-specific architectures to address the challenges of fine-grained visual perception. However, these approaches often exhibit notable limitations. A widely adopted strategy (Udandarao et al., 2024; Hou et al., 2025; Zhang et al., 2025c; Cao & Ou, 2025) involves scaling up training data by distilling knowledge from superior models. This paradigm is computationally inefficient due to the high cost of data synthesis, and large-scale distillation can lead to marginal performance gains while compromising generalization, resulting in an unsustainable trade-off. In contrast, another line of research employs reinforcement learning (RL) within a “thinking-with-image” paradigm, where performance is improved through iterative tool use and multi-step interaction (Zheng et al., 2025b; Shao et al., 2024; Hu et al., 2024; Gao et al., 2024; Zhou et al., 2024). While this improves system-level outcomes, it introduces substantial latency due to multi-round interactions. Moreover, the reasoning process often becomes superficial, emphasizing tool manipulation over visual comprehension, and fails to improve the model’s underlying perceptual capabilities on downstream tasks (Su et al., 2025a; Zhang et al., 2025a; Sun et al., 2024).

To address these challenges, we formulate the enhancement of visual perception as a two-stage structured process. The first stage cultivates holistic image reasoning and static scene understanding via self-critical caption refinement, teaching the model to “see widely”. The second stage focus on fine-grained perception and dynamic change awareness by inferring visual operations from image differences, training the model to “focus accurately”. This deliberate coarse-to-fine progression structures the model’s learning trajectory from global understanding to local precision.

Correspondingly, we introduce **ViPER**, a self-evolutionary framework that concretizes the two-stage task by integrating data synthesis and the RL method. ViPER consists of two components: an automated data synthesis module and a closely integrated two-stage RL method. By incorporating both image-level and instance-level reconstruction, ViPER’s data synthesis establishes a bidirectional mapping between visual and textual modalities at dual-granularities, intrinsically using the generation process to solidify perceptual understanding. The self-synthesized data is seamlessly channeled into a tightly coupled two-stage RL process. This integration creates a self-reinforcing loop where generation and learning are intertwined, eliminating the need for external bootstrapping and thereby establishing a self-evolving paradigm.

Based on ViPER, we constructed the *Viper10K* dataset and yield **Qwen-Viper** series perceptually enhanced from Qwen2.5-VL. Experiments revealed that during training, the models spontaneously developed a “thinking-with-image” capability and learned to redirect attention to critical details. On seven comprehensive benchmarks encompassing single-image, multi-image, and hallucination tasks, the Qwen-Viper series demonstrated **consistent gains** over the baseline. Notably, on fine-grained perception tasks, the 3B and 7B models achieved significant improvements of up to **4.4%** and **6.0%**, respectively. These results validate that ViPER enables VLMs to undergo self-evolution on perception-centric tasks.

In summary, our main contributions are threefold:

- I. We designed a **progressive two-stage task** that guides VLMs from holistic scene reasoning to fine-grained visual understanding, establishing a structured learning paradigm for perceptual self-improvement.
- II. We propose **ViPER**, a self-evolutionary paradigm that bootstraps VLM perception through a closed loop of data construction and RFT, producing *Viper10K* and the Qwen-Viper models.¹.
- III. We demonstrate that Qwen-Viper achieves **consistent improvements** on perception-intensive tasks, revealing the reciprocal relationship between generation and understanding by practice.

2 RELATED WORK

2.1 VISION-LANGUAGE MODEL

The rapid advancement of LLMs has catalyzed substantial progress in VLMs (Yang et al., 2025; Liu et al., 2024a;b). Representative proprietary models include GPT-4o (Hurst et al., 2024), Gemini 2.5 (Comanici et al., 2025), and Claude3 (Anthropic, 2024). Among open-source models, LLaVA (Liu et al., 2023) has gained widespread adoption through its three-stage architecture,

¹We open-source our implementation codes and data at <https://anonymous.4open.science/r/ViPER-1216>

108 which comprises a Vision Transformer (ViT) (Dosovitskiy et al., 2020), a connector module and
 109 an LLM. Subsequent research has largely extended and refined this architecture. For example,
 110 LLaVA-OneVision (Li et al., 2024) introduced a dynamic resolution mechanism that adaptively ad-
 111 justs the number of visual tokens based on the input image resolution. Similarly, Qwen2-VL (Wang
 112 et al., 2024a) also adopts this encoding strategy and incorporates the M-RoPE positional encoding
 113 mechanism, enabling unified processing of 1D text, 2D images and 3D video data. Its successor,
 114 Qwen2.5-VL (Bai et al., 2025) further integrates sparse attention mechanisms within the visual en-
 115 coder architecture. DeepSeek-VL Lu et al. (2024) optimizes encoding through dual encoders that
 116 separate visual and textual information for various downstream tasks. The InternVL series (Zhu
 117 et al., 2025; Wang et al., 2025b) proposed using a thumbnail to integrate global image information
 118 without substantially increasing the number of tokens.

119 2.2 REINFORCEMENT LEARNING-ENHANCED MULTI-MODAL REASONING

120 The integration of reinforcement learning (RL) has substantially enhanced the reasoning capabilities
 121 of LLMs, as demonstrated by models such as DeepSeek-R1 (Guo et al., 2025) and Kimi-K1.5 (Team
 122 et al., 2025). This progress is driven by effective RL frameworks, including DPO (Rafailov et al.,
 123 2023), PPO (Schulman et al., 2017), GRPO (Guo et al., 2025), GSPO (Zheng et al., 2025a), and
 124 DAPO (Yu et al., 2025a). Recent work has successfully extended these methodologies to VLMs,
 125 leading to promising results in systems such as Vision-R1 (Huang et al., 2025a) and VLM-R1 (Shen
 126 et al., 2025). Alongside end-to-end RL training paradigm, a significant research direction focuses
 127 on tool-augmented methods for multi-modal reasoning (Zhang et al., 2025b; Huang et al., 2025b; Su
 128 et al., 2025b). The success of multi-turn tool-calling in advanced agents such as OpenAI’s O3 (Ope-
 129 nAI, 2025) has inspired numerous efforts to equip models with specialized visual tools. Repre-
 130 sentative works include DeepEyes (Zheng et al., 2025b), Chain-of-Focus (Zhang et al., 2025d), and
 131 Mini-o3 (Lai et al., 2025) aiming to improve the models’ ability to dynamically invoke tools for
 132 solving complex reasoning tasks.

133 In contrast to these methods, our proposed method ViPER integrates data construction with post-
 134 training process to form a closed-loop cycle. The framework employs dual-granularity reconstruc-
 135 tion processes to promote understanding by generation, enabling self-driven evolution without ex-
 136 ternal supervision or large-scale data.

137 3 METHODOLOGY

140 In this section, we first formalize the two-stage tasks in Section § 3.1. Then we elaborated on the pro-
 141 posed self-evolutionary framework ViPER in Section § 3.2 and introduced its two key components:
 142 the data synthesis module and the coupled two-stage reinforcement learning strategy.

144 3.1 TASK FORMULATION

146 **Caption Self-Refining:** The first stage involves the self-refinement of captions. This task is designed
 147 to train a unified VLM to revise errors or biases in its own generated textual descriptions through
 148 deconstructing visual information and self-reflection.

150 Specifically, for each training sample $(I, C_g, R) \in \mathcal{D}$, the model performs the task of self-
 151 correction. It takes the image I and the initial caption C_g which is generated by itself with
 152 parameters θ_0 as input, analyses inaccuracy in the original caption and outputs a set of refinement
 153 actions:

$$R_{\text{pred}} = f(I, C_g; \theta). \quad (1)$$

154 The goal of the training is to minimize the discrepancy between the predicted set of refinement points
 155 R_{pred} and the ground-truth set R by adjusting the parameters θ . This is formulated as minimizing
 156 the expectation of a discrepancy metric $\delta(R, R_{\text{pred}})$ over the training distribution:

$$\min_{\theta} \mathbb{E}_{(I, C_g, R) \sim \mathcal{D}} [\delta(R, R_{\text{pred}})]. \quad (2)$$

160 Through this process, the VLM learns to iteratively refine its own generated captions through an
 161 image-based reasoning process. This task is designed to enhance the model’s ability for self-
 162 reflection and correction in visual perception, without relying on external correction mechanisms.

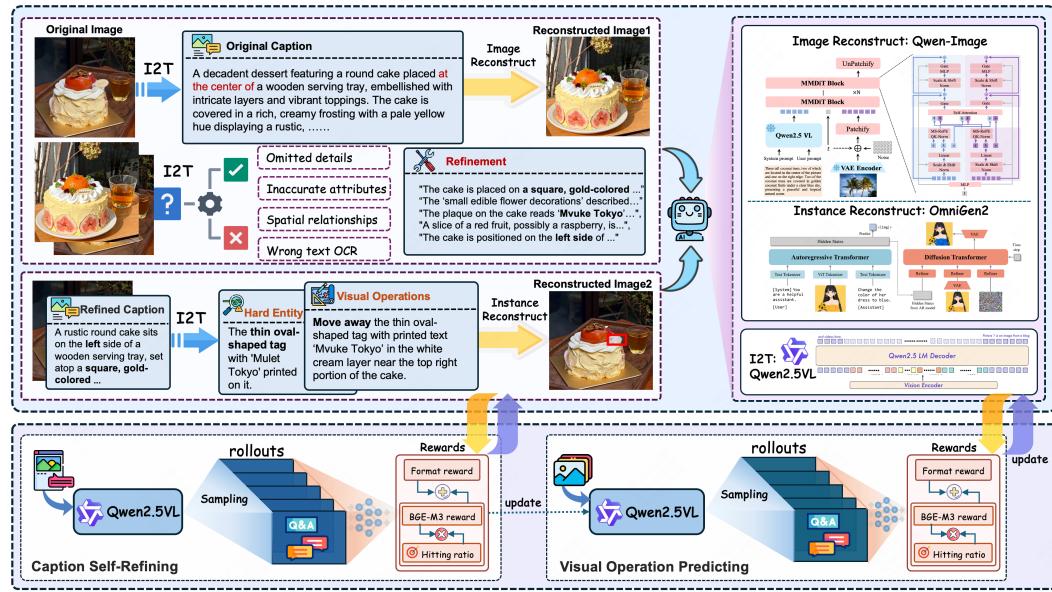


Figure 1: An overview of ViPER. The upper part illustrates the two-stage data synthesis framework, whose core component is a bidirectional vision-language mapping module composed of a VLM and a diffusion model. The lower part depicts the corresponding two-stage RL process: the first stage focuses on caption self-refinement, and the second stage on visual-operation predicting. The Qwen2.5-VL models are trained on *Viper10K* constructed by the ViPER framework and result in Qwen-Viper, thereby achieving substantial self-improvement.

Visual-Operation Predicting: The second phase aims to train the VLM to predict visual operations based on a pair of highly similar images with slight differences in visual information. The reconstructed image is generated by applying an editing operation to the original image, and the instruction itself is generated by the same VLM with parameters θ_1 . For each training sample $(I_{\text{orig}}, I_{\text{recon}}, \text{Ops}) \in \mathcal{D}$, the VLM predicts the visual operations based on careful perception and reasoning across the image pair:

$$\text{Ops}_{\text{pred}} = f(I_{\text{orig}}, I_{\text{recon}}; \theta), \quad (3)$$

Where $\text{Ops}_{\text{pred}} \in \mathcal{T}$ is the model's predicted visual operations. The objective of the training process is to minimize the discrepancy between the model's predicted visual operations and the actual visual operation instructions:

$$\min_{\theta} \mathbb{E}_{(I_{\text{orig}}, I_{\text{recon}}, \text{Ops}) \sim \mathcal{D}} [\delta(\text{Ops}, \text{Ops}_{\text{pred}})], \quad (4)$$

Where $\delta : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}^+$ is a text sequence difference metric. Through this process, the VLM learns to predict accurate visual operation instructions based on the differences between image pairs, thereby enhancing the model's ability to focus on key information and understand visual operations. The two-stage tasks, encompassing both static scene understanding and dynamic visual reasoning, not only enhance the model's image-based planning and reasoning abilities but also strengthen its fine-grained visual perception.

3.2 ViPER: THE SELF-EVOLUTIONARY FRAMEWORK

In this section, we introduced the unified framework **ViPER**, which integrates the data synthesis process with a corresponding two-stage RL strategy, and elaborated on how it enables self-evolutionary enhancement of perception capabilities.

3.2.1 THE DATA SYNTHESIS MODULE

A central part of ViPER is an automated data synthesis module, which functions as an integrated system for two-stage training data generation. This system seamlessly connects the upstream *Cap-*

216 *Caption Self-Refining* task to the downstream *Visual-Operation Predicting* task. The structure of this
 217 module is illustrated in the upper section of Figure 1.

218 The first-stage data synthesis is centered on an image-level reconstruction process that creates a
 219 closed-loop feedback mechanism for the VLM. The VLM first generates a static description from
 220 the original image. A diffusion model then reconstructs an image based on this description, which,
 221 due to the inherent information loss in visual-to-text conversion, exhibits local discrepancies from
 222 the original. These visual differences serve as a form of visual feedback, guiding the VLM to
 223 critique and refine its initial description by rectifying errors in object attributes, text, and spatial
 224 relationships. Thus, the diffusion model acts not merely as a generator, but as a critic that enables
 225 the VLM to iteratively optimize its visual grounding.

226 In contrast, the second-stage synthesis shifts the focus to instance-level reconstruction. Leveraging
 227 the refined caption from the upstream task, the VLM selects hard entities and generates correspond-
 228 ing visual operation instructions through a hand-crafted heuristic rules, as detailed in Appendix A.2.
 229 A diffusion model then executes these instructions to edit the original image, producing a recon-
 230 structed version. The instructions used naturally serve as ground truth for the *Visual-Operation*
 231 *Predicting* task, where the VLM learns to infer the operation from the image pair. Crucially, by
 232 concretizing instructional intent into observable visual variations, the generative model enables the
 233 VLM to learn fine-grained visual reasoning from self-induced scene changes.

234 The introduction of the generative model imbues the process with profound significance: It exter-
 235 nalizes the VLM’s internal reasoning into a concrete visual snapshot. This act of mapping critical
 236 cognitive steps back into the visual modality provides a tangible “cognitive anchor” for the VLM.
 237 Effectively, it endows the model with a form of image imagination, allowing it to perceive, critique,
 238 and refine its own understanding by confronting the visual consequences of its own textual descrip-
 239 tions. Specifically, we selected the Qwen2.5-VL-7B (Bai et al., 2025) model for the VLM in the
 240 framework, maintaining consistency with the trained baseline, while the diffusion models in two
 241 stages are implemented using Qwen-Image (Wu et al., 2025a) and OmniGen2 (Wu et al., 2025b),
 242 respectively. In line with the closed-loop design, the VLM within the data synthesis module is dy-
 243 namically updated from training checkpoints, enabling a co-evolution of the model and its training
 244 data. Based on our framework, we constructed a 10K perception-intensive vision-language dataset
 245 **Viper10K**, including 7K *Caption Self-Refining* data and 3K *Visual-Operation Predicting* data. The
 246 detailed implementation and instructions for the dataset construction are thoroughly presented in the
 247 Appendix A to ensure reproducibility and to facilitate future research within the community.

248 3.2.2 TWO-STAGE REINFORCEMENT LEARNING

249 To align with the progressive cognitive demands of the two-stage task, we designed a phased rein-
 250 forcement learning approach, as shown in the lower part of Figure 1. Since all training data is self-
 251 synthesized by the model undergoing training, distribution shift from heterogeneous data sources
 252 is eliminated. This self-sourcing strategy renders the entire RL process free from any cold-start re-
 253 quirement. The training proceeds sequentially: the first phase utilizes the *Caption Self-Refining* data,
 254 followed by the second phase focused on the *Visual-Operation Predicting* task. A unified reward
 255 computation mechanism is applied consistently across both stages.

256 **Reward mechanism:** For each model output O , we split it into a set of text sequences and use
 257 the BGE-M3 model (Chen et al., 2024a) to calculate the semantic similarity matrix between these
 258 sequences and the ground truth. The reward is applied as follows:

$$259 R_{\text{format}} = \begin{cases} 1 & \text{if } O \text{ matches the expected format} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

$$260 R_{\text{correct}} = \left(\frac{\sum_{i=1}^N \mathbb{1} [\max_{j=1}^M \text{sim}(s_i, g_j) \geq \tau]}{N} \right) \times \left(\frac{\sum_{i=1}^N (\mathbb{1} [\max_{j=1}^M \text{sim}(s_i, g_j) \geq \tau] \cdot L(s_i))}{\sum_{i=1}^N L(s_i)} \right), \quad (6)$$

$$261 R = w_f \cdot R_{\text{format}} + w_c \cdot R_{\text{correct}}, \quad (7)$$

262 Here, $G = \{g_1, g_2, \dots, g_M\}$ is the ground truth set, $S = \{s_1, s_2, \dots, s_N\}$ is the set of text se-
 263 quences obtained by splitting the output O , $L(s_i)$ is the character length of the split sentence s_i ,

270 $\text{sim}(s_i, g_j)$ is the semantic similarity between the sentence s_i and the ground truth g_j , and τ is the
 271 similarity threshold, set to 0.85. The weight coefficients w_f and w_c are set as $w_f = 0.05$ and
 272 $w_c = 0.95$.

273 **Optimization process:** We adopt a variation of Group Relative Policy Optimization (GRPO) al-
 274 gorithm, which has been proven effective in tasks involving mathematically structured reasoning
 275 chains (Ren et al., 2025). Following DAPO, we decouple the lower and higher clipping range to
 276 promote the diversity of the system as well as avoiding entropy collapse, and remove the KL penalty
 277 term to encourage bolder policy updates. Specifically, for each input question q , we first sample
 278 a set of outputs $\{o_1, o_2, \dots, o_n\}$, and then apply the reward function in Equation 7 to compute the
 279 reward score for each output. The policy model is optimized by maximizing the following objective:
 280

$$\begin{aligned} \mathcal{J}(\theta) = & \mathbb{E}[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)] \\ & \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left\{ \min \left[r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}}) \hat{A}_{i,t} \right] \right\}, \end{aligned} \quad (8)$$

285 where,

$$r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t} | q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} | q, o_{i,<t})}, \quad \hat{A}_{i,t} = \frac{R_i - \text{mean}(\{R_i\}_{i=1}^G)}{\text{std}(\{R_i\}_{i=1}^G)}. \quad (9)$$

289 4 EXPERIMENTS

290 4.1 EXPERIMENTAL SETUP

293 We selected Qwen2.5-VL-3B and Qwen2.5-VL-7B as base models and performed two-stage RL
 294 training on *Viper10K*, resulting in Qwen-Viper series. Identical hyperparameters were used for both
 295 stages: a batch size of 128, 5 rollouts per prompt at temperature 1.0, and rewards computed by the
 296 BGE-M3 model following the Equation 7. More details are presented in Appendix A.1.

298 4.2 MAIN RESULTS

300 4.2.1 BENCHMARKS

301 In our work, we use seven multimodal benchmarks for evaluation: (1) **MMStar** (Chen et al., 2024b)
 302 emphasizes visual dependency, ensuring that answer inference requires visual grounding while min-
 303 imizing potential data leakage. (2) **RealWorldQA** (X.AI, 2024) focuses on image understanding
 304 and verifiable reasoning in real-world scenarios. (3) **MME-RW(en)** (Zhang et al., 2025e) is the
 305 English subset of MME-RW, improving over previous benchmarks in scale, resolution, and task
 306 complexity for real-world applications. (4) **BLINK** (Fu et al., 2024) examines core visual per-
 307 ception abilities such as depth estimation, visual correspondence, and multi-view reasoning. (5)
 308 **MANTIS Eval** (Jiang et al., 2024) is designed to assess multi-image reasoning, including refer-
 309 ence resolution, comparison, and temporal understanding. (6) **HallusionBench** (Guan et al., 2023)
 310 evaluates model robustness against visual and language hallucinations. (7) **CRPE (relation)** (Wang
 311 et al., 2024b) evaluates models on subject–predicate–object structures, requiring recognition of both
 312 entities and their relations. The statics of each benchmark are detailed in the Appendix B.

313 4.2.2 EXPERIMENTAL RESULTS

315 We evaluated our method on seven diverse vision-language benchmarks encompassing single-image,
 316 multi-image, and hallucination tasks. The selected benchmarks include both general and reality
 317 perception-oriented visual question answering (VQA) tasks, aligning with our method’s focus. As
 318 shown in Table 1, ViPER yielded average performance gains of 1.7% and 1.6% across all bench-
 319 marks for the 3B and 7B models, respectively.

320 To analyze the specific capability improvements, we conducted fine-grained evaluations across six
 321 subdomains, and results are shown in Table 2.² The improvements were most pronounced in Fine-
 322 grained Perception, where Qwen-Viper-3B and Qwen-Viper-7B achieved gains of 4.4% and 6.0%

323 ²The data of all subdomains are from MMStar.

324 Table 1: Experimental results on multiple benchmarks encompassing single-image, multi-image
 325 and hallucination tasks, $\Delta \uparrow$ denotes the absolute gain of ViPER over the base model.

Model	MMStar	RealWorldQA	MME-RW (en)	BLINK (val)	Mantis Eval	Hallusion Bench(Avg)	CRPE (relation)	Overall
GPT-4o-0513(Hurst et al., 2024)	64.7	75.4	45.2	68.0	–	55.0	76.6	–
LLaVA-OneVision-0.5B(Li et al., 2024)	37.7	55.6	–	52.1	39.6	27.9	–	–
InternVL2.5-1B(Chen et al., 2024c)	50.1	57.5	44.2	42.0	51.2	39.0	60.9	49.3
InternVL3-1B(Zhu et al., 2025)	51.5	58.2	46.0	42.9	50.2	41.4	64.0	50.6
InternVL2.5-2B(Chen et al., 2024c)	53.7	60.1	48.8	44.0	54.8	42.6	70.2	53.5
InternVL3-2B(Zhu et al., 2025)	60.7	64.3	53.8	50.3	65.9	42.5	71.5	58.4
Qwen2.5-VL-3B(Bai et al., 2025)	55.9	65.4	53.1	47.6	68.7	46.3	73.6	58.7
Qwen-Viper-3B	57.8	67.7	54.6	49.2	70.0	49.1	74.3	60.4
$\Delta \uparrow$	1.9	2.3	1.5	1.6	1.3	2.8	0.7	1.7
MiniCPM-V2.6(Yao et al., 2024)	57.5	65.0	–	53.0	69.0	48.1	75.2	–
LLaVA-OneVison-7B(Li et al., 2024)	61.7	66.3	–	48.2	64.2	46.8	–	–
InternVL2.5-8B(Chen et al., 2024c)	62.8	70.1	59.1	54.8	67.7	50.1	78.4	63.3
InternVL3-8B(Zhu et al., 2025)	68.2	70.8	62.0	55.5	70.1	49.9	76.3	64.7
VL-Rethinker-7B(Wang et al., 2025a)	63.4	68.1	57.1	56.7	74.2	53.7	77.2	64.3
Qwen2.5-VL-7B(Bai et al., 2025)	63.9	68.5	57.4	56.4	75.1	52.9	76.4	64.4
+Viper10K (SFT)	64.5	68.2	57.9	56.8	75.6	53.2	76.1	64.6
Qwen-Viper-7B	66.2	71.4	59.0	57.6	75.6	54.4	77.6	66.0
$\Delta \uparrow$	2.3	2.9	1.6	1.2	0.5	1.5	1.2	1.6

344 Table 2: Performances across different domains. $\Delta \uparrow$ denotes the absolute gain of ViPER over the
 345 base model.

Model	Coarse Perception	Fine-grained Perception	Instance Reasoning	Logic Reasoning	Math	Science & Technology
Qwen2.5-VL-3B	68.8	48.4	62.4	56.4	62.0	37.2
Qwen-Viper-3B	70.0	52.8	64.4	57.6	62.8	39.2
$\Delta \uparrow$	1.2	4.4	2.0	1.2	0.8	2.0
Qwen2.5-VL-7B	73.6	55.6	73.2	69.6	66.8	44.4
Qwen-Viper-7B	75.2	61.6	74.8	70.4	67.2	48.0
$\Delta \uparrow$	1.6	6.0	1.6	0.8	0.4	3.6

356 respectively, underscoring ViPER’s primary effect on detailed visual understanding(e.g., object location/counting, attribute recognition). Stable gains were also observed in Coarse Perception and Instance Reasoning, confirming broader perceptual strengthening. Notably, despite no direct training
 357 on knowledge tasks, a significant 3.6% average improvement emerged in the Science & Technology
 358 domain for the 7B model. This compellingly suggests that enhanced perception enables more accurate
 359 integration of visual cues with parametric knowledge, demonstrating that model capabilities are
 360 deeply intertwined rather than modular.

363 Results on hallucination benchmarks further support this observation. Compared to the base models,
 364 Qwen-Viper exhibited lower hallucination rates, indicating that enhanced visual perception enables
 365 more faithful processing of image information and partially mitigates the negative impact of linguistic
 366 priors. Furthermore, performance gains on multi-image benchmarks confirm that the perception
 367 enhancements are fundamental and generalizable to out-of-domain tasks, enabling more effective
 368 understanding and reasoning across complex visual contexts.

370 4.3 IN-DEPTH ANALYSIS

372 4.3.1 INFLUENCE OF COLD-START

374 Conventional RL methods often rely on high-quality cold-start data (Ren et al., 2025; Lee et al.,
 375 2024; Yu et al., 2025b). Integrating a cold-start phase with RL training typically leads to more
 376 stable and significant improvements compared to direct RL fine-tuning. However, owing to the self-
 377 bootstrapping nature of the ViPER method, the data used for RL training is entirely self-generated,
 thereby eliminating the distribution discrepancy typically introduced by external models.

378 We conducted two experimental setups: a
 379 three-stage supervised SFT + RL process that
 380 includes a cold-start phase and a direct two-
 381 stage RL process without cold-start. In the
 382 cold-start stage, we used the Gemini-2.5-Pro
 383 model to generate 1K chain-of-thought data
 384 samples containing caption refinement and vi-
 385 sual operation prediction annotations, which
 386 were then used to SFT of Qwen2.5-VL.

As shown in Figure 2, although the reward curve of the training process without cold-start began at a significantly lower level compared to that with cold-start, its reward growth trend caught up with and surpassed the latter after 300 steps, eventually converging to a marginally higher final reward. The experimental phenomenon indicates that incorporating a cold-start phase in ViPER does not enhance model performance and may even constrain its potential for exploration and self-evolution. Thus, the proposed ViPER method effectively eliminates the dependency on high-quality cold-start data.

4.3.2 SFT vs. RL

We conducted supervised fine-tuning of Qwen2.5-VL-7B on the Viper10K dataset to systematically compare the efficacy of SFT and RL when trained on the same synthesized data. As summarized in Table 1, the SFT-based approach resulted in only marginal performance gains. We attribute this outcome to several factors: First, the limited scale of the Viper10K dataset constrains the degree of gradient-based adaptation achievable through pure SFT, thereby restricting its potential for substantial model improvement. Second, as the synthesized data lacks explicit chain-of-thought reasoning annotations, the SFT paradigm fails to facilitate the emergence of sophisticated visual reasoning capabilities in the model. Furthermore, the presence of subtle generator-specific artifacts introduces a latent bias, which SFT inadvertently amplifies during training.

Collectively, these experimental findings and analyses indicate that the success of the VIPER framework hinges not merely on data design, but more critically, on the tight coupling between data construction and downstream reinforcement learning. This underscores the distinct advantage of the RL approach in our proposed methodology.

4.3.3 TWO-STAGE RL VS. MIXED RL

We further analyze the advantages of the two-stage RL strategy. A common practice in reinforcement learning is to mix data from different types of tasks during training, typically to prevent the model from overfitting to a specific task and suffering from performance degradation on others.

Our two-stage task follows a coarse-to-fine progressive process: the *Caption Self-Refining* stage focuses on global image information processing and emphasizes the understanding of static scenes, while the *Visual-Operation Predicting* stage targets more fine-grained visual cues and highlights the comprehension of dynamic changes. Therefore, a multi-stage RL approach that incrementally enhances the model presents an intuitively sound training paradigm.

430 To evaluate the training strategy, we compared our
431 two-stage RL approach against a mixed RL process, where data from both stages were randomly
blended during training on the *Viper10K* dataset. As illustrated in Figure 3, the two-stage RL strategy

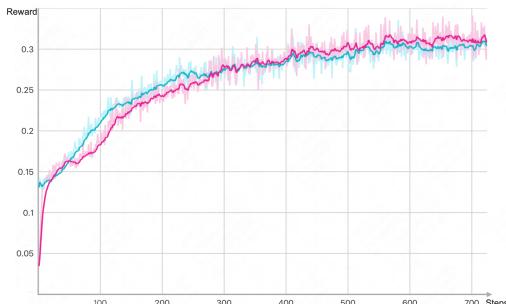


Figure 2: Reward curves **with/without** cold-start.

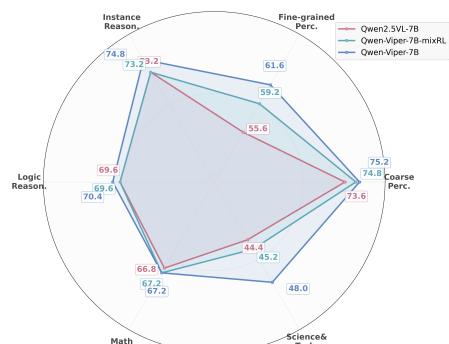
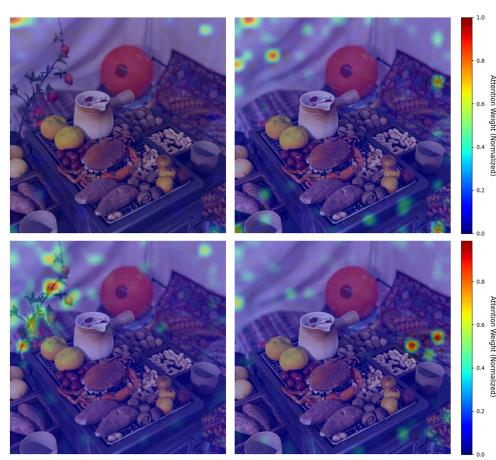


Figure 3: Improvements of the two-stage RL vs. the mixed RL across six domains. The two-stage strategy outperforms the mixed RL.



(a) Word cloud of chain-of-thought tokens from the model after *Caption Self-Refining* RL training. Verbs indicating visual operations are highlighted with colored boxes.



(b) Attention heatmaps for the same set of input images. The top two panels show the attention heatmaps from Qwen2.5-VL-7B, while the bottom corresponds to those from Qwen-Viper-7B.

Figure 4: Visualization of word cloud and attention.

significantly outperforms the mixed training approach, demonstrating the necessity of the proposed phased training strategy.

To further investigate the effects before and after the two-stage training, we conducted a more in-depth analysis. For the *Caption Self-Refining* stage, we intriguingly observed the emergence of visual thinking during training. We visualized the word cloud of the chain-of-thought tokens after training, as shown in Figure 4a. The trained model spontaneously produced high-frequency operation verbs such as “scan” “zoom in” “look closely at” and “focus on”, demonstrating a “thinking-with-images” reasoning pattern. For the *Visual-Operation Predicting* stage, we tracked changes in the model’s attention distribution for identical visual input before and after training to analyze the impact of this training stage through the lens of attention mechanisms. The results are presented in Figure 4b. The model after the second stage showed significantly increased concentration on critical information, effectively translating the image operations described in the first-stage reasoning chains into shifts in attention patterns. This shift enables the model to spontaneously focus on and perform detailed analysis of local regions.

4.3.4 ABLATION STUDY

To validate the necessity of each stage, we conducted ablation studies. Specifically, we performed RL training solely with the *Caption Self-Refining* stage and separately with only the *Visual-Operation Predicting* stage, then compared their results against the fully trained Qwen-Viper model that underwent both stages. The detailed experimental outcomes are presented in Figure 5.

The ablation results confirm that while each stage contributes uniquely, their integration is key to significant gains. Training solely on the *Caption Self-Refining* stage fosters a foundational, global visual reasoning capability. By learning to generate and critique holistic descriptions, the model builds a robust scaffold for scene understanding, leading to balanced but moderate gains. In contrast, exclusive training on the *Visual-Operation Predicting* drives localized perceptual sensibility, which forces the model to master fine-

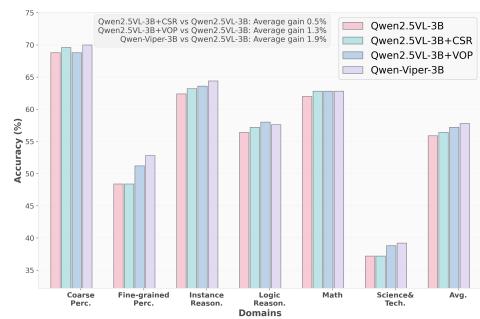


Figure 5: Ablation results on *Caption Self-Refining* and *Visual-Operation Predicting* Tasks.

486 grained attribute and relationship analysis, resulting in sharper improvements on corresponding
 487 tasks. Crucially, the full framework’s superiority stems from a coordinated mechanism: the first
 488 stage establishes reliable global context, which in turn primes and grounds the intensive local analy-
 489 sis of the second stage. This coherent progression from global to local enables a more profound and
 490 integrated self-evolution than either stage could achieve in isolation.

491

492

493 4.3.5 MECHANISMS FOR MITIGATING NOISES FROM GENERATIVE MODELS

494

495 ViPER’s framework centrally leverages diffusion models, which are employed not only for caption-
 496 to-image generation but also for highly challenging fine-grained visual editing tasks. This reliance
 497 imposes stringent requirements on the precision of diffusion models and consequently establishes a
 498 critical performance bottleneck. Current multimodal generative models often exhibit limitations in
 499 instruction-following capabilities, particularly in fine-grained compositional generation tasks, which
 500 may lead to data collapse and compromise training stability. To address the potential data noise
 501 introduced by generative models, we implement a dual-mechanism strategy:

502

503 First, for fine-grained editing tasks, we introduce VLM-as-Judge to filter generated outputs that fail
 504 to meet specified requirements. Implementation details and corresponding prompts are documented
 505 in Appendix A.4. Nevertheless, while VLM-as-Judge effectively ensures semantic fidelity, it
 506 provides inadequate control over non-semantic attributes such as lighting conditions and textural
 507 patterns. As a result, reconstructed images often inherit style transfer effects and textural artifacts
 508 characteristic of the generative model. To circumvent the impact of such generator-specific artifacts,
 509 we strategically design the reward function in the reinforcement learning phase to neutralize
 510 their influence. Although the VLM may occasionally produce self-critical content containing
 511 hallucinatory refinements due to biases introduced during data construction, it will never generate
 512 the same hallucinatory content during training because of the absence of reconstructed images.
 513 The reward of RL focuses exclusively on recall based on semantic similarity. Consequently, these
 514 minimally hallucinated refinements remain unmatched by the model’s rollouts, thereby yielding
 515 zero advantage estimates and no gradient updates. This design systematically addresses the impact
 516 of data noise on the VLM’s training process.

517

518 5 CONCLUSION

519

520 To break through the bottleneck of VLMs in perception-intensive tasks, we introduce a self-
 521 evolutionary framework ViPER, which establishes a closed-loop cycle of data construction and
 522 reinforcement fine-tuning. ViPER is built around a novel two-stage perceptual task that leverages
 523 image-level reconstruction and instance-level reconstruction to bootstrap visual perception and com-
 524 prehension. This framework addresses the scarcity of high-quality perceptual data by enabling the
 525 model to autonomously generate its own training samples, while the subsequent post-training phase
 526 continuously refines the model’s capabilities using the self-synthesized data. In this way, ViPER
 527 facilitates VLMs’ self-evolutionary enhancement of visual perception, moving beyond the limita-
 528 tions of text-centric reasoning approaches. Extensive experiments demonstrate the effectiveness of
 529 ViPER and provide insights into its self-consistent design, strongly validating the role of generation
 530 in advancing perceptual understanding.

531

532

533

534

535

536

537

538

539

540 REFERENCES
541

542 Anthropic. The claude 3 model family: Opus, sonnet, haiku. Technical re-
543 port, Anthropic PBC, 2024. URL https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.
544 Model Card.

545

546 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
547 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
548 2025.

549

550 Davide Caffagni et al. The revolution of multimodal large language models. In *Findings of the Asso-
551 ciation for Computational Linguistics: ACL 2024L*, 2024. URL <https://aclanthology.org/2024.findings-acl.807/>.

552

553 Yuqi Cao and Minghao Ou. Ame: Aligned manifold entropy for robust vision-language distillation.
554 *arXiv preprint arXiv:2508.08644*, 2025. URL <https://arxiv.org/abs/2508.08644>.

555

556 Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. M3-embedding:
557 Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge
558 distillation. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 2318–
559 2335, 2024a.

560

561 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong
562 Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. Are we on the right way
563 for evaluating large vision-language models? In A. Globerson, L. Mackey, D. Bel-
564 grave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural In-
565 formation Processing Systems*, volume 37, pp. 27056–27087. Curran Associates, Inc.,
566 2024b. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/2f8ee6a3d766b426d2618e555b5aeb39-Paper-Conference.pdf.

567

568 Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shen-
569 glong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source
570 multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*,
571 2024c.

572

573 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
574 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the
575 frontier with advanced reasoning, multimodality, long context, and next generation agentic ca-
576 pabilities. *arXiv preprint arXiv:2507.06261*, 2025.

577

578 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
579 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
580 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint
arXiv:2010.11929*, 2020.

581

582 Ahmed El-Kishky. Openai o1 system card. *ArXiv*, abs/2412.16720, 2024. URL <https://arxiv.org/abs/2412.16720>.

583

584 Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A.
585 Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but
586 not perceive. *ArXiv*, abs/2404.12390, 2024.

587

588 Zhi Gao et al. Multi-modal agent tuning: Building a vlm-driven agent for efficient tool usage.
589 *arXiv:2412.15606*, 2024.

590

591 Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang
592 Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An
593 advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-
language models. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition
(CVPR)*, pp. 14375–14385, 2023.

594 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 595 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 596 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

597

598 Yujie Hou et al. Vision-language model knowledge distillation method for iqa. *arXiv preprint*
 599 *arXiv:2507.15680*, 2025. URL <https://arxiv.org/abs/2507.15680>.

600 Yuxin Hu et al. Visual sketchpad: Sketching as a visual chain of thought for multimodal lms.
 601 *NeurIPS*, 2024.

602

603 Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and
 604 Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models.
 605 *arXiv preprint arXiv:2503.06749*, 2025a.

606 Zeyi Huang, Yuyang Ji, Anirudh Sundara Rajan, Zefan Cai, Wen Xiao, Haohan Wang, Junjie Hu,
 607 and Yong Jae Lee. Visualtoolagent (vista): A reinforcement learning framework for visual tool
 608 selection. *arXiv preprint arXiv:2505.20289*, 2025b.

609

610 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 611 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 612 *arXiv:2410.21276*, 2024.

613 Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max W.F. Ku, Qian Liu, and Wenhui Chen. Mantis:
 614 Interleaved multi-image instruction tuning. *Trans. Mach. Learn. Res.*, 2024, 2024. URL <https://openreview.net/forum?id=skLtdUVaJa>.

615

616 Xin Lai, Junyi Li, Wei Li, Tao Liu, Tianjian Li, and Hengshuang Zhao. Mini-o3: Scaling up
 617 reasoning patterns and interaction turns for visual search. *arXiv preprint arXiv:2509.07969*, 2025.

618

619 Hyung Won Lee et al. Rlaif: Scaling reinforcement learning from human feedback with ai feedback.
 620 In *ICLR 2024*, 2024. URL <https://openreview.net/forum?id=AAxIs3D2ZZ>.

621

622 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan
 623 Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint*
 624 *arXiv:2408.03326*, 2024.

625 Yunxin Li, Zhenyu Liu, Zitao Li, Xuanyu Zhang, Zhenran Xu, Xinyu Chen, Haoyuan Shi, Shenyuan
 626 Jiang, Xintong Wang, Jifang Wang, et al. Perception, reason, think, and plan: A survey on large
 627 multimodal reasoning models. *arXiv preprint arXiv:2505.04921*, 2025.

628

629 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
 630 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*
 631 *arXiv:2412.19437*, 2024a.

632

633 Haotian Liu, Chunyuan Li, Qingsheng Wu, and Yong Jae Lee. Visual instruction tuning. 36:34892–
 634 34916, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/6dcf277ea32ce3288914faf369fe6de0-Paper-Conference.pdf.

635

636 Shihong Liu, Samuel Yu, Zhiqiu Lin, Deepak Pathak, and Deva Ramanan. Language models as
 637 black-box optimizers for vision-language models. In *Proceedings of the IEEE/CVF Conference*
 638 *on Computer Vision and Pattern Recognition*, pp. 12687–12697, 2024b.

639

640 Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng
 641 Ren, Zhusu Li, Hao Yang, Yaofeng Sun, Chengqi Deng, Hanwei Xu, Zhenda Xie, and Chong
 642 Ruan. Deepseek-vl: Towards real-world vision-language understanding, 2024. URL <https://arxiv.org/abs/2403.05525>.

643

644 Qiao Lyu et al. Towards faithful model explanation in nlp: A survey. *Computational Linguistics*,
 645 2024. URL <https://aclanthology.org/2024.c1-2.6.pdf>.

646

647 OpenAI. Openai o3 and o4-mini system card. Technical report, OpenAI, April 2025. URL
<https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf>. System Card.

648 Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and
 649 Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model.
 650 *arXiv preprint arXiv:2305.18290*, 2023. URL <https://arxiv.org/abs/2305.18290>.
 651

652 Zehui Ren et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.
 653 *Nature*, 2025. doi: 10.1038/s41586-025-09422-z.
 654

655 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 656 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
 657

658 Hao Shao et al. Visual cot: Advancing multi-modal language models with step-by-step visual rea-
 659 soning. *arXiv:2403.16999*, 2024.
 660

661 Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun
 662 Zhang, Kangjia Zhao, Qianqian Zhang, et al. Vlm-r1: A stable and generalizable r1-style large
 663 vision-language model. *arXiv preprint arXiv:2504.07615*, 2025.
 664

665 Kacper Stechly et al. Chain of thoughtlessness? an analysis of cot in planning. In *NeurIPS*, 2024.
 666

667 Alex Su, Haozhe Wang, Weiming Ren, Fangzhen Lin, and Wenhui Chen. Pixel reasoner: In-
 668 centivizing pixel-space reasoning with curiosity-driven reinforcement learning. *arXiv preprint
 669 arXiv:2505.15966*, 2025a.
 670

671 Zhaochen Su, Linjie Li, Mingyang Song, Yunzhuo Hao, Zhengyuan Yang, Jun Zhang, Guanjie
 672 Chen, Jiawei Gu, Juntao Li, Xiaoye Qu, et al. Openthinking: Learning to think with images via
 673 visual tool reinforcement learning. *arXiv preprint arXiv:2505.08617*, 2025b.
 674

675 Kai Sun, Yushi Bai, Ji Qi, Lei Hou, and Juanzi Li. Mm-math: Advancing multimodal math eval-
 676 uation with process evaluation and fine-grained classification. *arXiv preprint arXiv:2404.05091*,
 677 2024.
 678

679 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
 680 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with
 681 llms. *arXiv preprint arXiv:2501.12599*, 2025.
 682

683 Vishaal Udandarao et al. Active data curation effectively distills large-scale multimodal models. In
 684 *CVPR*, 2024. URL <https://arxiv.org/abs/2411.18674>.
 685

686 Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhui Chen. VI-
 687 rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning.
 688 *arXiv preprint arXiv:2504.08837*, 2025a.
 689

690 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
 691 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the
 692 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.
 693

694 Weiyun Wang, Yiming Ren, Haowen Luo, Tiantong Li, Chenxiang Yan, Zhe Chen, Wenhui Wang,
 695 Qingyun Li, Lewei Lu, Xizhou Zhu, et al. The all-seeing project v2: Towards general relation
 696 comprehension of the open world. *arXiv preprint arXiv:2402.19474*, 2024b.
 697

698 Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu,
 699 Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3. 5: Advancing open-source multimodal
 700 models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025b.
 701

702 Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai
 703 He, Kuan Wang, Jianfeng Gao, et al. Reinforcement learning for reasoning in large language
 704 models with one training example. *arXiv preprint arXiv:2504.20571*, 2025c.
 705

706 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 707 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 708 neural information processing systems*, 35:24824–24837, 2022.
 709

710 Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai
 711 Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*,
 712 2025a.

702 Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan
 703 Jiang, Yexin Liu, Junjie Zhou, et al. Omnipgen2: Exploration to advanced multimodal generation.
 704 *arXiv preprint arXiv:2506.18871*, 2025b.

705 X.AI. Realworldqa. Blog post, November 2024. URL <https://x.ai/news/grok-1.5v>.

707 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 708 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 709 *arXiv:2505.09388*, 2025. URL <https://arxiv.org/abs/2505.09388>.

710 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li,
 711 Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint*
 712 *arXiv:2408.01800*, 2024.

713 Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on
 714 multimodal large language models. *National Science Review*, 2024. URL <https://arxiv.org/abs/2306.13549>.

715 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian
 716 Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system
 717 at scale. *arXiv preprint arXiv:2503.14476*, 2025a.

718 Tianyu Yu, Haoye Zhang, Qiming Li, Qixin Xu, Yuan Yao, Da Chen, Xiaoman Lu, Ganqu Cui,
 719 Yunkai Dang, Taiwen He, et al. Rlaif-v: Open-source ai feedback leads to super gpt-4v trust-
 720 worthiness. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp.
 721 19985–19995, 2025b.

722 Di Zhang, Jingdi Lei, Junxian Li, Xunzhi Wang, Yujie Liu, Zonglin Yang, Jiatong Li, Weida Wang,
 723 Suorong Yang, Jianbo Wu, Peng Ye, Wanli Ouyang, and Dongzhan Zhou. Critic-v: Vlm critics
 724 help catch vlm errors in multimodal reasoning. In *Proceedings of the IEEE/CVF Conference on*
 725 *Computer Vision and Pattern Recognition (CVPR)*, pp. 9050–9061, June 2025a.

726 Haoji Zhang, Xin Gu, Jiawen Li, Chixiang Ma, Sule Bai, Chubin Zhang, Bowen Zhang, Zhichao
 727 Zhou, Dongliang He, and Yansong Tang. Thinking with videos: Multimodal tool-augmented
 728 reinforcement learning for long video reasoning. *arXiv preprint arXiv:2508.04416*, 2025b.

729 Xin Zhang et al. Hkd4vlm: A progressive hybrid knowledge distillation framework for robust mul-
 730 timodal hallucination and factuality detection in vlms. *arXiv preprint arXiv:2506.13038*, 2025c.
 731 URL <https://arxiv.org/abs/2506.13038>.

732 Xintong Zhang, Zhi Gao, Bofei Zhang, Pengxiang Li, Xiaowen Zhang, Yang Liu, Tao Yuan, Yuwei
 733 Wu, Yunde Jia, Song-Chun Zhu, et al. Chain-of-focus: Adaptive visual search and zooming for
 734 multimodal reasoning via rl. *arXiv preprint arXiv:2505.15436*, 2025d.

735 YiFan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng
 736 Li, Kun Wang, Qingsong Wen, Zhang Zhang, Liang Wang, and Rong Jin. MME-realworld:
 737 Could your multimodal LLM challenge high-resolution real-world scenarios that are difficult for
 738 humans? In *The Thirteenth International Conference on Learning Representations*, 2025e. URL
 739 <https://openreview.net/forum?id=k5VHHgsRbi>.

740 Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang,
 741 Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy op-
 742 timization. 2025a. URL <https://arxiv.org/abs/2507.18071>.

743 Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and
 744 Xing Yu. Deepeyes: Incentivizing “thinking with images” via reinforcement learning. *arXiv*
 745 *preprint arXiv:2505.14362*, 2025b.

746 Qiji Zhou, Ruochen Zhou, Zike Hu, Panzhong Lu, Siyang Gao, and Yue Zhang. Image-of-thought
 747 prompting for visual reasoning refinement in multimodal large language models. *arXiv preprint*
 748 *arXiv:2405.13872*, 2024.

749 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
 750 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
 751 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

756 **A IMPLEMENTATION DETAILS**
757758 **A.1 EXPERIMENTAL SETUP**
759760 We implemented the two-stage reinforcement learning based on the verl 0.3.1³. The entire training
761 process was conducted on a single node equipped with eight A100 80GB GPUs, leveraging Fully
762 Sharded Data Parallel (FSDP) for efficient large-scale training. Key hyperparameters used in the
763 reinforcement learning process are summarized in Table 3.
764765 **Table 3: Key experimental settings for the RL process**
766

767 Module	768 Parameter	769 Value
770 data	batch_size	128
	max_prompt_length	10240
	max_response_length	4096
772 actor	optim.lr	1×10^{-6}
	mini_batch_size	64
	micro_batch_size_per_gpu	4
	use_kl_loss	False
776 sampling	temperature	1.0
	rollout.n	5
778 hardware	num_gpus	8
	memory_per_gpu	80 GB

781 **A.2 HEURISTIC RULES OF VISUAL OPERATIONS**
782783 For the data construction process of the visual-operation predicting task, we designed heuristic rules
784 for the VLM. First, we predefined four visual operation tasks as toolkits, specifying their applica-
785 ble scenarios and rules. These correspond to four common error types in fine-grained perception
786 scenarios for VLMs: 1. Omission of small-scale details; 2. Confusion of spatial relationships;
787 3. Overemphasis on primary subjects while neglecting contextual backgrounds; 4. Incorrect fine-
788 grained attribute discrimination;
789790 The model selects the most suitable editing task upon thorough and meticulous analysis of the input
791 image:
792793 **Prompt for Heuristic Rules of Visual Operations**794 **Role and Core Principle**795 You are a creative and insightful visual analyst. Your task is to analyze the provided image
796 and determine the most suitable category of subtle edit for creating challenging VQA (Visual
797 Question Answering) data.798 The edit category you choose must be the most suitable to the content of the image, and
799 support fine-grained changes that require careful observation to notice.800 **Task Description**

801 First, conduct a thorough analysis of the provided image. Describe all the details.

802 Next, refer to the "Editing Toolkit" below.

803 Based on your analysis, select the most suitable task category from the toolkit.

804 Finally, articulate your reasoning for this choice, explaining why this category is superior to
805 others for creating a subtle and effective edit on this specific image.806 **Editing Toolkit**807
808
809
³<https://verl.readthedocs.io/en/v0.3.x/start/install.html>

810
 811
 812
 813
 814
 815
 816
 817
 818
 819
 820
 821
 822
 823
 824
 825
 826
 827
 828
 829
 830
 831
 832

- **Removing/Adding Details:** If the image contains many elements or objects with complex ingredients, consider removing or adding one of the small-scale objects.
- **Changing Spatial Relationships:** If an entity has a well-defined spatial relation to its surroundings and can be feasibly repositioned, then alter its location to disrupt and reconstruct the spatial context.
- **Editing Background:** If the image is clearly separated into a foreground and background, and the background is rich with elements, consider making edits to an object there.
- **Attribute Tuning:** If the image contains an object with a property that is easy to change, and the change would be subtle enough to require close inspection to notice, consider editing one of the object’s properties.

Output Format

Your output MUST follow the JSON structure below. Do not include any other commentary or introductory text.

{

Image Analysis: [Your detailed analysis of the image in natural language.]

Edit Task: [The name of the task category you chose from the toolkit.]

Reasoning: [Your detailed justification for why this task category is the most suitable for this image, explaining how it enables a subtle and effective edit and why it is superior to other options.]

}

For the given visual operation tasks, we provide advanced and detailed task descriptions, and require the model to generate precise visual-operation instructions that are optimally tailored to the image content.

A.3 PROMPT IN DUAL-LEVEL RECONSTRUCTION

We provide the key prompts used for dual-level reconstruction during the data synthesis process below to ensure full reproducibility of the pipeline:

Prompt for Caption Self-Refining

You are an expert AI assistant specializing in high-fidelity image caption evaluation.

Your task is to analyze an original image, its caption, and a reconstructed image that was generated based only on that caption. Your goal is to identify how the caption fails to accurately and completely describe the original image.

The core of your analysis is to compare the original image with the reconstructed image. The differences between these two images are a direct result of inaccuracies or omissions in the caption. For every significant difference you find, you must describe the corresponding flaw in the caption.

Analysis Criteria:

Analyze the caption’s flaws based on the visual discrepancies between the original and reconstructed images:

- **Omitted Details:** Does the reconstructed image lack important objects, elements, or background details that are present in the original image? This indicates the caption omitted these details.
- **Misidentified Entities:** Are objects, people, or scenes in the reconstructed image fundamentally different from the original? (e.g., a dog instead of a cat, a city instead of a forest). This indicates the caption misidentified something.
- **Inaccurate Spatial Relationships:** Is the position or arrangement of elements in the reconstructed image incorrect compared to the original? This indicates the caption described spatial relationships inaccurately.

864
865
866
867

- Unmentioned or Incorrect Text: Is there text visible in the original image that is missing or garbled in the reconstructed image? This indicates the caption either omitted or incorrectly transcribed the text.

868
869
870
871
872

Input:

- Original Image: The ground-truth image
- Reconstructed Image: The image generated from the caption
- Current Caption: The text caption being evaluated

873
874
875
876
877

Output Requirements:

- Your response must be a single JSON object. This object must contain a single key, “refinement”, which is a list of strings. Each string in the list must be a concise and precise description of a single factual error or significant omission you identified in the caption.
- Focus exclusively on factual errors and omissions that cause visual discrepancies.
- Crucially, do not describe the differences between the images in your output. Instead, describe the flaw in the caption that caused the difference.
- Do NOT comment on writing style, phrasing, tone, or subjective quality.
- If the caption is perfectly accurate and the reconstructed image is a faithful representation of the original, return a JSON object with an empty list: “refinement”: [].

886
887

Prompt for Generating Refined Captions

888
889
890
891
892

Role

You are an extremely meticulous and precise text editor.

893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Task

Your sole task is to revise an “Original Caption” based on a list of “Refinement Suggestions”. Your goal is to produce a fluent, coherent “Refined Caption” that fully incorporates all suggestions.

Rules

- **Must Integrate All Suggestions:** You must accurately reflect every point from the “Refinement Suggestions” in the new caption.
- **No New Information:** You are strictly forbidden from adding any new details that are not present in either the original caption or the suggestions.
- **Preserve Correct Information:** All accurate information from the original caption that is not targeted by a suggestion must be retained.
- **Single Text Output:** The final result must be a single, flowing descriptive paragraph, not a list of changes.
- **Strict Output Format:** Your response must contain only the final refined caption text. Do not include any preambles (e.g., “Here is the refined caption:”), explanations, or Markdown formatting.

Input Data

- **Original Caption:** {original_caption}
- **Refinement Suggestions:** {suggestions_str}

Please directly output the refined caption:

918
919**Prompt for Generating Visual Editing Instructions**

920

You are a professional image editing instruction generator. Follow these steps:

921

1. First, carefully observe the image and identify all entities and elements in the image.
2. Analyze the cognitive difficulty of each entity (considering complexity, abstraction level, recognition difficulty, etc.).
3. Determine the entity with the highest cognitive difficulty and suitable to edit.
4. Based on the given task description and specific image content, generate a targeted editing instruction.

922

Return JSON format with these fields:

923

- “entity_to_edit”: the entity chosen to be remove
- “editing_instruction”: Comprehensive editing instruction (see requirements below)

924

Requirements for editing instructions:

925

- Must be concise, clear, and directly usable as a prompt for image editing/generation models
- Must comply with the given task description
- Avoid lengthy explanations, focus on direct editing commands

926

Task Description:

927

{task_description}

928

Please analyze this image and generate corresponding editing instructions.

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

A.4 DATA QUALITY ASSURANCE

946

Precise and diverse visual editing instructions pose significant challenges for current open-source image editing models, which are prone to uncontrollable errors in fine-grained editing tasks. To address this, we have introduced a dual-validation mechanism combining CLIP scores and LLM-as-a-judge. The CLIP score serves as an initial filter to eliminate editing results that deviate significantly from the original image, while the LLM-as-a-judge employs stricter validation rules to ensure that the final edits strictly adhere to the instructions while maintaining high consistency with the original image.

947

948

949

950

951

952

953

954

955

Prompt for Checking Instance-Level Reconstruction

956

Role

957

You are an image validation specialist, and you specialize in rigorously assessing edited images to ensure they precisely adhere to the given instructions.

958

959

960

961

962

963

964

The original image

965

{image1}

966

967

The edited image

968

{image2}

969

970

971

Please evaluate the edited image compared to the original image, focusing only on semantic changes.

Editing instruction

{edit instruction}

What to evaluate:

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

- **Strict Adherence:** Does the edit accurately reflect the user’s instruction in terms of content and objects?

- **No Unintended Semantic Changes:** Are there any new objects, removed objects (that were not supposed to be removed), or changes in the attributes of existing objects that were not part of the instruction?

What to ignore:

Minor, global variations in brightness, contrast, or color saturation. Consider these acceptable side effects unless the instruction was specifically about changing them.

Decision:

- If the edit is semantically accurate and clean, set “is_valid” to true.
- If the edit fails to follow the instruction or introduces unintended semantic changes, set “is_valid” to false and provide a brief explanation in the “reason” field.

Please return your judgment only in JSON format as follows:

{ “is_valid”: boolean, “reason”: “A brief explanation if is_valid is false, otherwise this field can be omitted.” }

1026 B BENCHMARK DETAILS

1028 To provide a clearer understanding of the datasets used in our experiments, we present more detailed
 1029 descriptions of the seven multimodal benchmarks:

1030 (1) **MMStar** (Chen et al., 2024b) is designed to address the limitations of weak visual dependency
 1031 and unintentional data leakage in previous benchmarks. It contains 1,500 human-curated samples
 1032 that require genuine visual reasoning and are resistant to memorization effects. The benchmark
 1033 spans six core capabilities and 18 fine-grained axes, ensuring a balanced and rigorous evaluation of
 1034 large VLMs.

1035 (2) **RealWorldQA** (X.AI, 2024) provides more than 700 images from real-world scenarios, includ-
 1036 ing those captured from vehicles and other everyday contexts. Each image is paired with a question
 1037 and a verifiable answer, making it a concise yet practical dataset to assess real-world grounding and
 1038 reasoning in multimodal models.

1039 (3) **MME-RW** (Zhang et al., 2025e) is currently one of the largest manually annotated multimodal
 1040 benchmarks, comprising tens of thousands of high-quality images collected from public datasets
 1041 and the Internet. Expert annotators generated diverse question–answer pairs covering challenging
 1042 real-world subtasks. Compared with previous datasets, it features higher-resolution images and
 1043 significantly greater task diversity. **MME-RW (en)** denotes its English subset, where all questions
 1044 and answers are provided in English.

1045 (4) **BLINK** (Fu et al., 2024) reformulates 14 classic computer vision tasks into 3,807 multiple-
 1046 choice questions, paired with single or multiple images. The benchmark assesses core perceptual
 1047 abilities, including relative depth estimation, visual correspondence, and multi-view reasoning.

1048 (5) **Mantis Eval** (Jiang et al., 2024) is a benchmark specifically designed to evaluate multi-image
 1049 reasoning in multi-modal models. It consists of carefully curated test cases that require reasoning
 1050 across multiple images, including tasks such as reference resolution, cross-image comparison, and
 1051 temporal understanding. By focusing on multi-image contexts, the benchmark provides a systematic
 1052 way to assess whether models can integrate information beyond single-image perception.

1053 (6) **HallusionBench** (Guan et al., 2023) is built to evaluate multimodal reasoning under
 1054 hallucination-inducing conditions. It contains 346 images paired with 1,129 expert-designed ques-
 1055 tions, many organized into control structures that allow fine-grained analysis of logical consistency
 1056 and error types.

1057 (7) **CRPE** (Wang et al., 2024b) provides a structured evaluation of object recognition and relation
 1058 comprehension. The benchmark consists of single-choice questions divided into four subsets: Ex-
 1059 istence, Subject, Predicate, and Object. This division enables systematic testing of whether models
 1060 can detect the presence of objects, identify entities, and capture subject-predicate-object relations,
 1061 offering insight into fine-grained relational reasoning abilities. In our experiments, we specifically
 1062 adopt the **relation** subset to evaluate the subject-predicate-object reasoning.

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

C DATA SAMPLE

1081

1082

1083

1084

1085

1086

1087

Caption Self-Refining

1088

Original Image & Reconstructed Image

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

Original Caption

1105

A decadent dessert featuring a round cake placed at the center of a wooden serving tray, embellished with intricate layers and vibrant toppings. The cake is covered in a rich, creamy frosting with a pale yellow hue displaying a rustic, textured finish reminiscent of hand-spread icing. Embedded along the sides are slices of ripe pink fruit, possibly figs, their interior pulp and outer skin visible through the frosting. Atop the cake sits an eye-catching glossy orange segment, likely a fruit compote or poached persimmon, crowned by small edible flower decorations with earthy green and burgundy tones. Surrounding the centerpiece is a white whipped cream spiral arranged in a flower-like shape, with a petite, oval-shaped decorative plaque reading “Muletokyo” nestled among the toppings. To the side of the cake, sitting on the same wooden tray, is a clear glass filled with amber-colored liquid, potentially tea or juice, offering a refreshing pairing to the indulgent treat. The background reveals an outdoor setting with hints of vegetation, characterized by thin branches and clusters of leaves, alongside a dark glass panel that reflects faint shapes and shadows.

1118

Refinement

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133



1134
1135

Visual-Operation Predicting

1136
1137

Original Image & Reconstructed Image

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154



Visual-Operation Instructions

1155 Add small chopped green scallions scattered on top of the sesame-covered meat pastries and
1156 on the paper lining in the basket.

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188
1189
1190
1191
1192
1193

D CASE STUDY

1194
1195
1196
1197
1198
1199
1200
1201
1202 In this section, we present concrete cases to visually demonstrate the significant improvements
1203 achieved by our method. For each example, we provide responses from three models: Gemini-2.5-
1204 Pro (by API), Qwen2.5-VL-7B, and Qwen-Viper-7B. For the latter two models, we visualize their
1205 attention distributions during the response generation process, using attention heatmaps as a lens to
1206 observe the operational mechanism of the ViPER method.
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225**Case Study of Visual Counting**1226
1227 **Question:** How many garages does the first building on our right have?
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241**Gemini-2.5-pro****Determining Garage Count**

I'm zeroing in on the user's query: specifically, how many garages are visible on the first structure to the right. My focus is on thoroughly examining the image to isolate and enumerate them.

Focusing On Specifics

Now, I'm focusing on the specifics. I've pinpointed the correct structure and need to accurately count the visible garages. Counting them carefully, I've identified five distinct garage doors on the ground floor. My next step is to formulate the answer concisely and confidently, incorporating a brief descriptive confirmation for clarity.

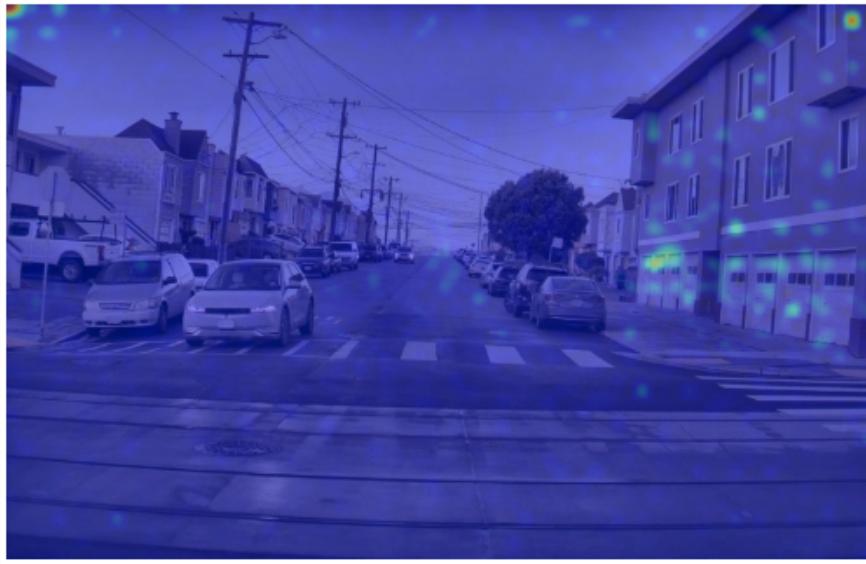
Refining The Response

Now, I'm taking a more analytical approach. I've broken down the user's question into distinct steps: identifying the building, counting the garages, and crafting a clear answer. I've confirmed that the target building is the most prominent on the right with five garage doors. I'm focusing on ensuring my response is both direct and detailed to instill trust.

Based on the image, the first building on the right has **five** garages.

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261

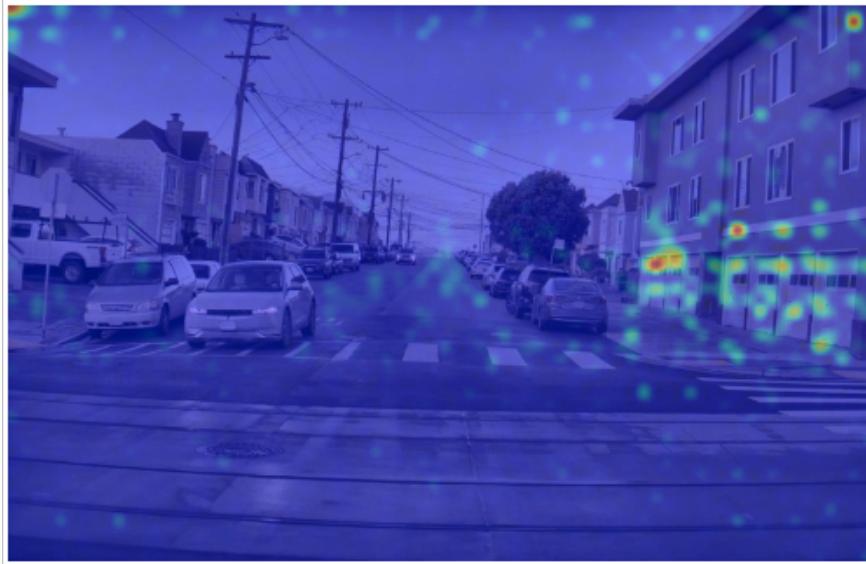
Qwen2.5-VL-7B



1262 The first building on the right side of the image has **four** garage doors. These are clearly
1263 visible at the bottom of the building, just above the street level.
1264

1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283

Qwen-Viper-7B(Ours)



1284 I'll start by examining the image, **focusing on** the structure on the right side of the street.
1285 The image is a view down a residential street with parked cars and buildings on both sides.
1286 My goal is to determine how many garages the first building on the right has.
1287 I'm looking at this building, which appears to be a multi-story structure, likely an apartment
1288 complex or a similar type of building. It's a light grayish-brown color. As my eyes **scan its**
1289 **facade**, I see what looks like a row of white doors. These are clearly visible and are located
1290 directly beneath a row of windows on the ground floor. They are evenly spaced and appear
1291 to be identical in size and shape.

1292 Next, I'll **do a quick visual count** of these doors. Starting from the left, I can see one, two,
1293 three, four, five, six, seven, and eight of them. They are all aligned in a straight line, making
1294 it a simple count. This building on the right has a total of eight white garage doors.

1295

1296

1297 The question specifically asks for the number of garages on the first building on the right,
 1298 and I can confidently state that there are **eight** of them.

1299

1300

1301 Case Study of Fine-grained Perception

1302

1303 **Question:** What is the color scheme of the costumes worn by the men in the parade?

1304

1305 **Options:** A: Yellow, red, and gold, B: Red, blue, and black, C: Green, orange, and white,
 1306 D: Black, white, and gray

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328



 alamy stock photo

KFBVXP
www.alamy.com

1329

1330 Gemini-2.5-pro

1331

1332 Based on the image, the primary colors visible in the costumes are:

1333

1334 Red (outer robes, tassels on hats)

1335

1336 Blue (sashes/belts)

1337

1338 Yellow/Gold (inner garments/trim)

1339

1340 Black (main body of the tunic/vest, hats)

1341

1342 White (gloves, neck trim)

1343

1344 Comparing this to the options:

1345

1346 A: Yellow, red, and gold (Missing prominent black and blue)

1347

1348 B: Red, blue, and black (These are all very prominent colors)

1349

1343 C: Green, orange, and white (Green and orange are not major colors)

1344 D: Black, white, and gray (Missing prominent red, blue, and yellow)

1345

1346 The most accurately descriptive option focusing on the major, dominant colors is B: Red, blue, and black, as these three are clearly primary components of the main garment layers.

1347

1348 However, since Yellow is also very prominent, let's look closer. The yellow is a significant color block. The options are a bit restrictive.

1349

If we have to choose the best fit:

Red, blue, and black are all very strong, dominant colors.

1350

1351 Red, blue, and yellow/gold are also very strong, dominant colors.

1352

1353 Given the choices, B: Red, blue, and black captures three of the most contrasting and major
 1354 colors in the main body of the uniform. However, if the yellow parts are considered "gold,"
 1355 A: Yellow, red, and gold would also be a strong contender, focusing on the brightest colors,
 1356 but it misses the prominent black and blue.

1357

1358 Let's stick with the most prominent and contrasting large color blocks: Red, blue, and black
 1359 appear to be the dominant features.

1360

1361 The correct option is **B: Red, blue, and black**.

1362

Qwen2.5-VL-7B

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385



1385 The color scheme of the costumes worn by the men in the parade is:

1386

1387 **A: Yellow, red, and gold.**

1388

Qwen-Viper-7B(Ours)

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457



I'll start by examining the image of the group of men in what appears to be a traditional parade. My goal is to determine the color scheme of their costumes from the options provided.

First, I'll look at the central figure on the left. He's wearing a hat with a red tassel, a red and black robe, and a blue belt. The next man to his right is wearing a similar outfit but with a red hat and a red and black top. As my eyes scan the group, I can see this pattern continues. The colors that stand out are red, black, and blue, along with some golden elements like the trims.

Next, I'll compare these colors against the given options. Option A mentions yellow, red, and gold, which misses the black and blue as main colors. Option B includes red, blue, and black, which is a perfect match for what I'm seeing. Option C lists green, orange, and white, none of which are present. And option D describes black, white, and gray, which are not part of the costumes at all.

After carefully reviewing the image, the most fitting description is clearly **B: Red, blue, and black**. That's my final answer.

E RESULTS ON ViRL39K

Table 4: Comparison between Qwen2.5-VL-7B and the version RFT on ViRL39K data, Δ denotes the variations in accuracy.

Model	MMStar	RealWorldQA	MME-RW (en)	BLINK (val)	Mantis Eval	Hallusion Bench(Avg)	CRPE (relation)	Overall
Qwen2.5-VL-7B	63.9	68.5	57.4	56.4	75.1	52.9	76.4	64.4
Qwen2.5-VL-7BViRL39K	63.4	68.1	57.1	57.0	74.2	53.7	76.8	64.3
Δ	0.5 \downarrow	0.4 \downarrow	0.3 \downarrow	0.6 \uparrow	0.9 \downarrow	0.8 \uparrow	0.4 \uparrow	0.1 \downarrow

To compare the training data synthesized by the ViPER framework with existing multimodal reasoning datasets, we fine-tuned Qwen2.5-VL-7B using the comprehensive ViRL39K dataset (Wang et al., 2025a) via the GRPO algorithm. ViRL39K provides a curated collection of 38,870 verifiable QAs for Vision-Language RL training, with over 80% are math or chart/diagram reasoning

1458 questions, The results summarized in Table 4 show that the fine-tuned model **did not** achieve su-
1459 perior performance on the evaluated benchmarks. Although modest improvements were observed
1460 on hallucination-oriented tasks, the model slightly underperformed the baseline on single-image
1461 and multi-image benchmarks emphasizing real-world perception. This outcome highlights two key
1462 insights: 1.The reasoning paradigms derived from mathematical and chart-oriented tasks exhibit
1463 **limited generalizability** to perception-intensive visual tasks; 2.It underscores the targeted efficacy
1464 of our proposed method in specifically enhancing visual perceptual capabilities.

1465 F DATA AVAILABILITY STATEMENT

1466 All image data employed in this study are under formal authorization for academic research pur-
1467 poses, ensuring full compliance with legal and ethical standards regarding copyright protection and
1468 privacy preservation.

1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511