

# Omni-R1: Towards the Unified Generative Paradigm for Multimodal Reasoning

Dongjie Cheng<sup>1</sup>, Yongqi Li<sup>1\*</sup>, Zhixin Ma<sup>2</sup>, Hongru Cai<sup>1</sup>,  
Yupeng Hu<sup>3</sup>, Wenjie Wang<sup>4</sup>, Liqiang Nie<sup>5</sup>, Wenjie Li<sup>1</sup>

<sup>1</sup>The Hong Kong Polytechnic University, <sup>2</sup>Singapore Management University,

<sup>3</sup>Shandong University, <sup>4</sup>University of Science and Technology of China

<sup>5</sup>Harbin Institute of Technology (Shenzhen)

dong-jie.cheng@connect.polyu.hk, liyongqi0@gmail.com


## Abstract

Multimodal Large Language Models (MLLMs) are making significant progress in multimodal reasoning. Early approaches focus on pure text-based reasoning. More recent studies have incorporated multimodal information into the reasoning steps; however, they often follow a single task-specific reasoning pattern, which limits their generalizability across various multimodal tasks. In fact, there are numerous multimodal tasks requiring diverse reasoning skills, such as zooming in on a specific region or marking an object within an image. To address this, we propose unified generative multimodal reasoning, which unifies diverse multimodal reasoning skills by generating intermediate images during the reasoning process. We instantiate this paradigm with Omni-R1, a two-stage SFT+RL framework featuring perception alignment loss and perception reward, thereby enabling functional image generation. Additionally, we introduce Omni-R1-Zero, which eliminates the need for multimodal annotations by bootstrapping step-wise visualizations from text-only reasoning data. Empirical results show that Omni-R1 achieves unified generative reasoning across a wide range of multimodal tasks, and Omni-R1-Zero can match or even surpass Omni-R1 on average, suggesting a promising direction for generative multimodal reasoning. Codes and checkpoints are available at [Omni-R1](#).

## 1 Introduction

Multimodal Large Language Models (MLLMs) are evolving from basic perception to more advanced reasoning capabilities (Meng et al., 2025). Early multimodal reasoning methods primarily focus on pure text-based reasoning, where MLLMs perform textual reasoning before generating the final answer. More recently, some studies (Hu et al., 2024) highlight that the intermediate reasoning steps also require the involvement of visual information. For

**Question:** Is the person on the left side of a vehicle?



**Answer:** Based on the Zoom-in image, yes, the person is on the left side of a vehicle.

Figure 1: An example illustrating the necessity of incorporating visual information into the intermediate reasoning steps for multimodal tasks.

example, as illustrated in Figure 1, a spatial relation problem can be more easily solved by zooming in on the critical region during the intermediate reasoning process.

There have been some interleaved-modal reasoning methods (Zheng et al., 2025; Li et al., 2025b; Chen et al., 2025; Chern et al., 2025) to incorporate multimodal information into the reasoning steps. For example, in DeepEyes (Zheng et al., 2025), the method utilizes external tools to zoom in on a specific region of an image, enabling the model to focus on particular areas of the input image for VQA tasks. In MVoT (Li et al., 2025b), the approach involves imagining possible paths by visualizing intermediate states to solve spatial reasoning tasks. Despite the effectiveness, current interleaved-modal reasoning methods still rely on a single specific multimodal reasoning pattern to address a particular multimodal task.

To address the above issue, we propose to develop a general and unified multimodal reasoning paradigm capable of handling a broad range of multimodal tasks. 1) First, there are numerous multimodal tasks that require diverse reasoning skills in practice. As shown in Figure 2, skills such as zoom-in and grounding are essential for general VQA tasks, while visual prediction skills are

\*Corresponding author

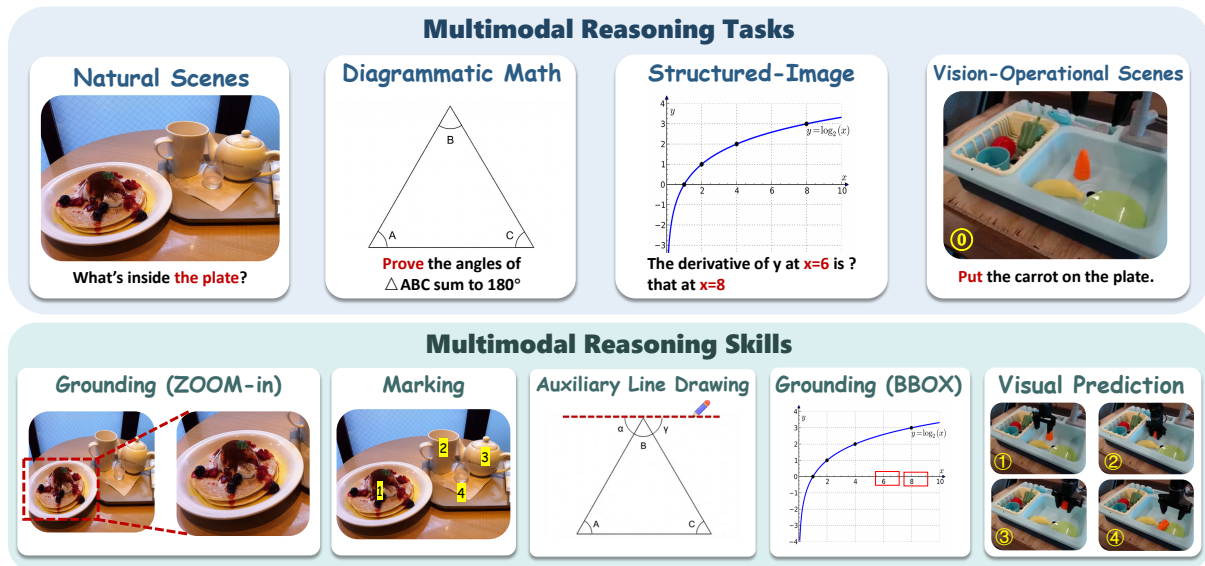


Figure 2: The illustration of various multimodal tasks and the corresponding diverse multimodal reasoning skills, with multimodal tasks shown in the top row and the required reasoning skills summarized in the bottom row.

clearly beneficial for visual-operational scenes. 2) Second, it is necessary to facilitate various multimodal reasoning skills within a single MLLM in a unified paradigm. Inspired by the recent development of Omni models (Team, 2024; Chern et al., 2024; Deng et al., 2025) that exhibit multimodal generation capabilities, we propose that all multimodal reasoning skills can be unified and internalized within a generative paradigm, i.e., **unified generative multimodal reasoning**. For example, zoom-in can be realized by generating a magnified view of a specific image region, while grounding can be achieved by generating annotated bounding boxes over the objects in the original image.

Despite this promising vision of the generative multimodal reasoning paradigm, putting it into practice remains challenging due to the following aspects. 1) Functional image generation. Unlike typical images, the images required for various reasoning skills are functional and often unnatural, such as those containing marked numbers over specific objects. This poses a significant challenge for MLLMs. 2) Costly interleaved-modal reasoning annotations. The training of generative multimodal reasoning requires step-by-step annotations of the interleaved reasoning process. Such interleaved data is costly and rarely available at scale.

In this work, we propose Omni-R1, a framework that unifies various multimodal reasoning skills in a generative manner without relying on extensive supervision. Specifically, we introduce standard Omni-R1, which requires a two-stage optimization

procedure with supervised fine-tuning (SFT) and reinforcement learning (RL) with a small amount of reasoning annotations. To address the issue of functional image generation, we introduce a perception alignment loss for SFT and a perception-calibrated reward for RL, which ensures that the model’s optimization is supervised by visual perception. Furthermore, we present Omni-R1-Zero, which eliminates the need for costly interleaved-modal reasoning annotations. Instead, it introduces a bootstrapping step-wise visualization method to synthesize interleaved-modal reasoning data. Empirical results suggest that Omni-R1 enables unified generative reasoning across various multimodal reasoning tasks, while Omni-R1-Zero even surpasses Omni-R1 to some extent.

The contributions are summarized as follows:

- We identify a range of distinctive multimodal reasoning skills for multimodal tasks, and propose a generative multimodal reasoning paradigm to unify them within one MLLM.
- We introduce Omni-R1, which facilitates the unified multimodal reasoning paradigm and stabilizes functional image generation via the designed perception alignment loss and perception-calibrated reward.
- We present Omni-R1-Zero, which eliminates the need for interleaved-modal reasoning annotations while achieving performance comparable to models trained with supervised annotations.

## 2 Related Work

### 2.1 Text-based Multimodal Reasoning

Prior work focuses on text-based multimodal reasoning, where images serve as context, and the model generates textual rationales and the final answer. KAM-CoT (Mondal et al., 2024) augmented multimodal chain-of-thought with structured external knowledge. More recently, reinforcement learning has been proven to be effective in multimodal reasoning. MM-EUREKA (Meng et al., 2025) trained MLLMs with rule-based, verifiable rewards, enhancing their reasoning abilities in math tasks. Visual-Thinker-R1-Zero (Zhou et al., 2025) further applied such rule-based reinforcement learning directly to a non-SFT MLLM on math data, achieving improved visual reasoning performance.

### 2.2 Interleaved-modal Reasoning

ICOT (Gao et al., 2025) introduced the concept of *interleaved-modal reasoning*, where the model augments its trajectory by interleaving additional multimodal evidence. DeepEyes (Zheng et al., 2025) used tools to acquire and append new multimodal observations during reasoning, enabling fine-grained improvements through reinforcement learning. More recently, interleaving has been realized within a single omni-model as interleaved generation. MVoT (Li et al., 2025b) trained an omni-MLLM to generate image-text interleaved rationale in a spatial reasoning task. To support such behaviors with supervision, Zebra-CoT (Li et al., 2025a) released a 182K interleaved text-image reasoning dataset for the exploration of interleaved-modal reasoning, especially in omni-models. There is another work with the same name, Omni-R1 (Zhong et al., 2025), which uses RL to optimize the selection of informative keyframes and pixel-level grounding for video segmentation.

## 3 Omni-Bench

To better evaluate MLLM’s multimodal reasoning capabilities, we construct Omni-Bench, which consists of various multimodal tasks that require diverse multimodal reasoning skills.

**Uni-Tasks.** Uni-Tasks characterize the multimodal reasoning scenarios included in Omni-Bench. *Natural-Scene Perception* focuses on natural-scene images and requires evidence localization for answering. *Diagrammatic Math* involves diagram-grounded visual arithmetic and ge-

Task	Source	#Samples
Natural-Scene Perception	V* (Wu and Xie, 2024)	100
Structured-Image	ArxivQA (Li et al., 2024) ChartQA (Masry et al., 2022)	100 100
Diagrammatic Math	Geometry3k (Lu et al., 2021) MathVista (Lu et al., 2024)	100 100
Vision-Operational Scenes	ViC-Bench (Wu et al., 2025)	300

Table 1: The composition of Omni-Bench. A total of 800 samples span four Uni-Tasks.

ometric reasoning. *Structured-Image* targets structured image inputs, such as figures and charts, that combine text and graphics. *Vision-Operational Scenes*, such as visual games and embodied planning, require complex visual operations or even predicting the transition of scene states. For each Uni-Task, we randomly sampled data from established datasets to form a slice. Table 1 summarizes the source datasets and the number of prompts in each slice.

**Uni-Skills.** Across these scenarios, generative multimodal reasoning should require a set of recurring visual skills for functional image generation. We summarize them as four Uni-Skills: *Grounding* localizes task-relevant evidence and may zoom in for closer inspection when needed. *Auxiliary line drawing* draws helper lines to make geometric relations or alignment constraints explicit. *Marking* highlights or enumerates relevant instances so they can be referred to unambiguously. *Visual prediction* anticipates the next visual state by performing a one-step transition.

**Evaluation.** For evaluation, we extract the final answer span *Ans* and use an LLM judge (Zheng et al., 2023; Raina et al., 2024) with a fixed evaluation template to compare it against the gold reference, producing a binary correctness decision. Here, the LLM judge acts only as a deterministic answer-matching checker rather than an open-ended evaluator.

## 4 Method

As shown in Figure 3, we present two training frameworks, **Omni-R1** and **Omni-R1-Zero**. Both aim to learn generative multimodal reasoning while unifying diverse reasoning skills and remaining effective under limited or even zero interleaved supervision.

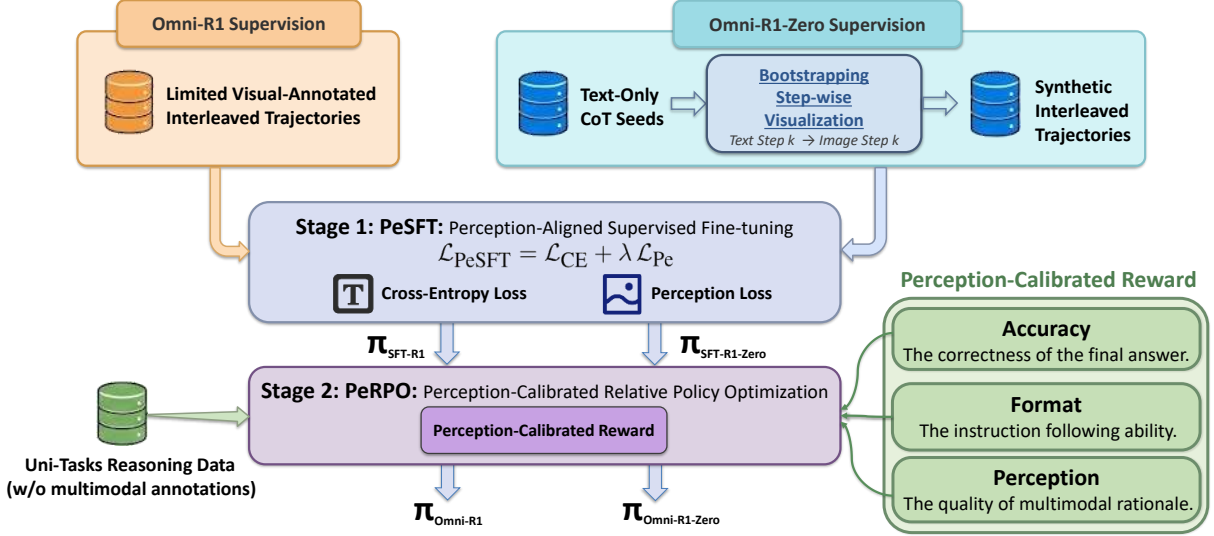


Figure 3: **Training pipeline for Omni-R1 and Omni-R1-Zero.** (Top) Data initialization uses either limited human annotations (orange) or synthetic bootstrapped trajectories (blue). (Middle) Stage 1 (PeSFT) performs supervised fine-tuning with joint cross-entropy and perception losses. (Bottom) Stage 2 (PeRPO) refines the policy using unified tasks without multimodal annotation (left) and a composite reward (Accuracy, Format, Perception) for final alignment.

## 4.1 Omni-R1

Omni-R1 optimizes the policy in two stages: 1) **Perception-Aligned Supervised Fine-Tuning (PeSFT)**. This stage learns the interleaved reasoning format and Uni-Skills prediction using cross-entropy loss, and optimizes the image tokens representation with the perception loss. 2) **Perception-Calibrated Relative Policy Optimization (PeRPO)**. It extends training to Uni-Tasks without multimodal annotations by refining the policy model with a perception-calibrated reward.

### 4.1.1 PeSFT

We denote the multimodal reasoning token trajectory for input  $x$  as  $y_{1:T}$ , where  $T$  indicates the token count. The output  $y$  adheres to a unified template across various scenarios, as detailed in Appendix A. PeSFT optimizes a cross-entropy loss to enforce this unified generation format, while applying a perception loss to align the generation of visual tokens.

**Cross-Entropy loss.** We apply standard cross-entropy loss over all tokens in the generative reasoning trajectory,

$$\mathcal{L}_{\text{CE}} = \mathbb{E} \left[ - \sum_{t=1}^T \log \pi_{\theta}(y_t | x, y_{<t}) \right], \quad (1)$$

which trains the model to follow or imitate the ideal reasoning process. Overall,  $\mathcal{L}_{\text{CE}}$  encourages the model to reproduce the interleaved generative

reasoning format of the trajectory. For image-token segments, it also serves as local supervision for Uni-Skills, guiding the model to generate the desired functional images given the current context.

**Perception loss.** We utilize a visual codebook  $\mathbf{E} \in \mathbb{R}^{K \times D}$  as the reference for image tokens within the unified template, where  $K$  denotes the codebook size and  $D$  the embedding dimension. Let  $t \in \Omega$  be the index of an image token in  $y$ , and  $c_t \in \{1, \dots, K\}$  be its corresponding ground-truth index in  $\mathbf{E}$ . We define the perception loss as:

$$\mathcal{L}_{\text{Pe}} = \frac{1}{|\Omega|} \sum_{t \in \Omega} \|W h_t - \mathbf{E}[c_t]\|_2^2. \quad (2)$$

where  $h_t \in \mathbb{R}^H$  is the final-layer hidden state and  $W \in \mathbb{R}^{D \times H}$  is a linear projection. With  $\mathbf{E}$  frozen, this loss aligns hidden states with the codebook’s geometry, acting as a perceptual prior to stabilize autoregressive image-token generation.

**Overall training objective.** Putting the two terms together, the overall objective is

$$\mathcal{L}_{\text{PeSFT}} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{Pe}}, \quad \lambda = 1. \quad (3)$$

where we weight the perception term by a coefficient  $\lambda$  and set  $\lambda = 1$  in all experiments.

### 4.1.2 PeRPO

PeRPO is built upon a group-relative RL objective and aims to optimize long, interleaved multimodal sequences. It introduces a perception-calibrated

reward to more accurately align RL feedback with the quality of intermediate functional image generations.

We formulate PeRPO as an optimization over the policy model’s ( $\pi_\theta$ ) autoregressive trajectory generation, integrating a perception-calibrated reward  $\mathcal{R}$  with a group-relative PPO objective.

**Perception-calibrated reward.** We define the reward  $\mathcal{R}$  as a weighted sum of three components,  $\mathcal{R}_{Acc}$ ,  $\mathcal{R}_{Fmt}$ , and  $\mathcal{R}_{Pe}$ .

**Accuracy ( $\mathcal{R}_{Acc}$ ).** We compute an accuracy reward by parsing the final answer and comparing it to the ground truth with a rule-based verifier that checks numeric correctness, symbolic/textual equivalence, and domain-specific patterns. The accuracy reward is set to a value in  $[0, 1]$  based on the degree of match to the ground truth. Implementation details of the verifier are provided in Appendix B.

**Format ( $\mathcal{R}_{Fmt}$ ).** We set  $\mathcal{R}_{Fmt} = 1$  if the generated trajectory contains a well-formed reasoning part and a final answer part that can be reliably parsed by the verifier, and  $\mathcal{R}_{Fmt} = 0$  otherwise.

**Perception ( $\mathcal{R}_{Pe}$ ).** We measure perceptual coherence of intermediate image generations via 2D Total Variation (TV) on codebook embeddings. Let  $\mathbf{E} \in \mathbb{R}^{K \times D}$  be a frozen visual codebook. Given a visual token index  $c_t \in \{1, \dots, K\}$ , where  $K$  is the codebook size and  $D$  is the codebook embedding dimension, we retrieve its embedding by

$$\mathbf{e}_t = \mathbf{E}[c_t] \in \mathbb{R}^D. \quad (4)$$

The 1-D image-token segment with  $N = H_q W_q$  tokens, where  $H_q$  and  $W_q$  is the height/width of the quantized image grid in VQVAE of the model, is reshaped into a  $H_q \times W_q$  grid. The embedding matrix of this image can be written as

$$\mathbf{Z} \in \mathbb{R}^{H_q \times W_q \times D}. \quad (5)$$

We define the index sets  $\mathcal{N}_h$  and  $\mathcal{N}_v$  as

$$\begin{aligned} \mathcal{N}_h &= \{(i, j) \mid 1 \leq i \leq H_q, 1 \leq j < W_q\}, \\ \mathcal{N}_v &= \{(i, j) \mid 1 \leq i < H_q, 1 \leq j \leq W_q\}. \end{aligned} \quad (6)$$

Then

$$E_h = \frac{1}{|\mathcal{N}_h|} \sum_{(i,j) \in \mathcal{N}_h} \|\mathbf{z}_{i,j+1} - \mathbf{z}_{i,j}\|_2^2, \quad (7)$$

and

$$E_v = \frac{1}{|\mathcal{N}_v|} \sum_{(i,j) \in \mathcal{N}_v} \|\mathbf{z}_{i+1,j} - \mathbf{z}_{i,j}\|_2^2. \quad (8)$$

Finally, we compute the 2-D TV energy as

$$E_{2D} = \frac{1}{D} (E_h + E_v). \quad (9)$$

We map energy to a normalized score

$$s_r = \frac{1}{1 + E_{2D}/\tau}, \quad (10)$$

where  $\tau > 0$  controls sensitivity. Let  $\mathcal{S}_{seg}$  be the set of image-token segments in a trajectory. We aggregate segment scores by

$$\mathcal{R}_{Pe} = \frac{1}{|\mathcal{S}_{seg}|} \sum_{r \in \mathcal{S}_{seg}} s_r. \quad (11)$$

We combine the three reward components as

$$\mathcal{R} = \alpha \mathcal{R}_{Acc} + \beta \mathcal{R}_{Fmt} + \gamma \mathcal{R}_{Pe}, \quad (12)$$

where  $\alpha, \beta, \gamma$  balance answer accuracy, format compliance, and perceptual consistency. We use  $\mathcal{R}$  as the scalar reward  $r_i$  for each candidate  $i$  when computing group-relative advantages.

**Policy optimization.** In each optimization step, we sample a set of reasoning trajectory candidates  $\mathcal{G} = \{y^{(i)}\}$  for each prompt  $x$  using the behavioral policy  $\pi_{old}$ . Each  $y^{(i)}$  is an interleaved multimodal sequence of length  $T_i$  and is assigned a scalar reward  $r^{(i)}$  according to Eq. 12. To facilitate meaningful gradient updates, we filter the samples to retain only mixed-outcome (non-degenerate) groups (i.e.,  $\sigma_{\mathcal{G}} > 0$ ) before updating the policy network to  $\pi_\theta$ . Subsequently, we compute the group-relative advantage based on the reward:

$$A^{(i)} = \frac{r^{(i)} - \mu_{\mathcal{G}}}{\sigma_{\mathcal{G}} + \delta}, \quad (13)$$

where  $\mu_{\mathcal{G}}$  and  $\sigma_{\mathcal{G}}$  denote the mean and standard deviation of the rewards  $\{r^{(j)}\}_{j \in \mathcal{G}}$ , and  $\delta > 0$  is a small constant for numerical stability. Let  $m_{i,t} \in \{0, 1\}$  be the response mask for candidate  $y^{(i)}$ , where  $t \in \{1, \dots, T_i\}$  indexes token positions and  $m_{i,t} = 1$  indicates a response position to be optimized. We denote the number of optimized (response) positions by  $L_i = \sum_{t=1}^{T_i} m_{i,t}$ .

Using the computed advantages, we maximize a PPO-clip objective with a KL regularizer to a fixed reference policy  $\pi_{ref}$ :

$$\begin{aligned} \mathcal{J}(\theta) &= \mathbb{E}_{x, \{y^{(i)}\} \sim \pi_{old}} \left[ \frac{1}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} \frac{1}{L_i} \sum_{t=1}^{T_i} m_{i,t} \left( \mathcal{L}_{clip}^{(i,t)}(\theta) \right. \right. \\ &\quad \left. \left. - \beta_{KL} D_{KL}^{(i,t)}(\pi_\theta, \pi_{ref}) \right) \right]. \end{aligned} \quad (14)$$

Here  $\mathcal{L}_{clip}^{(i,t)}(\theta)$  is the PPO clipped surrogate loss,

$$\mathcal{L}_{clip}^{(i,t)}(\theta) = \min \left\{ \rho^{(i,t)}(\theta) A^{(i)}, \text{clip}(\rho^{(i,t)}(\theta), l, u) A^{(i)} \right\}. \quad (15)$$

with  $l = 1 - \varepsilon_{\text{low}}$ ,  $u = 1 + \varepsilon_{\text{high}}$ ,  $\beta_{\text{KL}} > 0$ , and  $\text{clip}(z, l, u) = \min(\max(z, l), u)$ . For the autoregressive prefix  $y_{<t}^{(i)} = (y_1^{(i)}, \dots, y_{t-1}^{(i)})$ , we define the token-level likelihood ratio and KL term as

$$\begin{aligned} \rho^{(i,t)}(\theta) &= \frac{\pi_{\theta}(y_t^{(i)} | x, y_{<t}^{(i)})}{\pi_{\text{old}}(y_t^{(i)} | x, y_{<t}^{(i)})}, \\ D_{\text{KL}}^{(i,t)}(\pi_{\theta}, \pi_{\text{ref}}) &= \log \frac{\pi_{\theta}(y_t^{(i)} | x, y_{<t}^{(i)})}{\pi_{\text{ref}}(y_t^{(i)} | x, y_{<t}^{(i)})}. \end{aligned} \quad (16)$$

## 4.2 Omni-R1-Zero

We consider the setting where human-annotated interleaved multimodal rationales are unavailable, and bootstrap interleaved supervision from text-only CoT to activate generative multimodal reasoning.

### 4.2.1 Bootstrapping Step-wise Visualization

As illustrated in Figure 3, Omni-R1-Zero starts from text-only Chain-of-Thought (CoT) seeds and constructs synthetic interleaved trajectories via step-wise visualization, generating one image for each reasoning step (Text Step  $k \rightarrow$  Image Step  $k$ ) and assembling them with the same control-token template as Omni-R1.

These synthetic traces are not intended to be perfect supervision, but rather to teach the interleaved reasoning format and expose intermediate multimodal states.

### 4.2.2 Optimization

Omni-R1-Zero reuses the same two-stage pipeline as Omni-R1: PeSFT on the synthetic interleaved data, followed by PeRPO with the same objective described in Section 4.1. This optimization improves correctness and stabilizes intermediate image generations, enabling interleaved multimodal traces to emerge without human-annotated rationales.

## 5 Experiment

### 5.1 Experimental Setup

**Benchmarks.** We evaluated Omni-R1 and Omni-R1-Zero on Omni-Bench, which covers four Uni-Tasks, namely Natural-Scene Perception, Diagrammatic Math, Structured-Image, and Vision-Operational. We additionally reported results on standard general multimodal benchmarks, including MME (Fu et al., 2023), MM-Vet (Yu et al., 2024), V\*-Bench (Wu and Xie, 2024), POPE (Li et al., 2023), MMVP (Tong et al., 2024), and

BLINK (Fu et al., 2024). Detailed dataset descriptions are provided in the Appendix C.

**Baselines.** We compared against a diverse set of multimodal reasoning baselines, including Zebra-CoT (Li et al., 2025a), Bagel (Text CoT) (Deng et al., 2025), MM-Eureka (Meng et al., 2025), MINT-CoT (Chen et al., 2025), and Deep-Eyes (Zheng et al., 2025). These baselines cover native multimodal interleaved generation, supervised interleaved reasoning, text-based multimodal reasoning, and tool-based multimodal reasoning.

**Evaluation.** For Omni-Bench, we followed the evaluation protocol described in Section 3. For general benchmarks, we conducted evaluation using the widely adopted VLMEvalKit (Duan et al., 2024). Unless otherwise specified, we reported accuracy, and for POPE (Li et al., 2023), we reported the F1-score.

**Implementation.** We instantiate Omni-R1 on two Unified Multimodal Model (UMM) backbones, Anole (Chern et al., 2024) and Bagel (Deng et al., 2025), and Omni-R1-Zero on Anole (Chern et al., 2024). For Omni-R1, we report the results of Omni-R1-M and Omni-R1-L, which were trained with 30 and 60 PeRPO steps, respectively. As for Omni-R1-Zero, we report Omni-R1-Zero-T, Omni-R1-Zero-S, and Omni-R1-Zero-M trained with 10, 20, and 30 PeRPO steps, respectively. Data usage is described in Appendix C, and implementation details are provided in Appendix D.

## 5.2 Main Results

We presented and discussed results on Omni-Bench and general multimodal benchmarks separately.

**Analysis on Omni-Bench.** Table 2 shows that Omni-R1 and Omni-R1-Zero perform competitively across the four Uni-Task slices, with consistent improvements over the corresponding base models. Notably, Omni-R1-M (Bagel) achieves the strongest overall performance, showing that our framework is competitive with text-only and tool-based reasoning baselines. Meanwhile, Omni-R1-Zero-T/S/M (Anole) consistently outperform the base model across different RL budgets, indicating that PeRPO can successfully induce generative multimodal reasoning.

**Analysis on general benchmarks.** Table 3 shows that, on the Anole (Chern et al., 2024) backbone, both Omni-R1 and Omni-R1-Zero outper-

Backbone	Model	Natural	Structured	Diagrammatic	Vision-Op.	Average	$\Delta\text{Avg} (\%)$
Qwen-2.5-VL-7B (Bai et al., 2025)	Qwen-2.5-VL (Base) (Bai et al., 2025)	0.710	<u>0.657</u>	0.388	0.063	0.374	-
	MM-Eureka (Meng et al., 2025)	0.670	0.608	<u>0.489</u>	<b>0.103</b>	<u>0.397</u>	+6.1
	DeepEyes (Zheng et al., 2025)	<b>0.780</b>	0.578	0.433	0.080	0.380	+1.7
	MINT-CoT (Chen et al., 2025)	0.580	0.647	0.488	0.017	0.363	-3.0
Anole (Base) (Chern et al., 2024)	Anole (Base) (Chern et al., 2024)	0.180	0.140	0.055	0.027	0.081	-
	Zebra-CoT (Li et al., 2025a)	0.398	0.130	0.073	0.077	0.129	+59.2
<i>Omni-R1</i>							
Anole-7B (Chern et al., 2024)	<b>Omni-R1-M</b>	0.440	0.135	0.115	0.093	0.153	+87.7
	<b>Omni-R1-L</b>	0.420	0.175	0.090	0.090	0.153	+87.7
<i>Omni-R1-Zero</i>							
Anole-7B (Chern et al., 2024)	<b>Omni-R1-Zero-T</b>	0.450	0.165	0.095	0.087	0.154	+89.2
	<b>Omni-R1-Zero-S</b>	0.340	0.145	0.090	<u>0.097</u>	0.138	+69.2
	<b>Omni-R1-Zero-M</b>	0.410	0.170	0.120	0.093	0.159	+95.4
Bagel-7B (Deng et al., 2025)	Bagel (Text CoT, Base) (Deng et al., 2025)	0.440	0.562	0.317	0.030	0.286	-
	Zebra-CoT (Li et al., 2025a)	0.500	0.612	0.383	0.060	0.333	+16.5
<i>Omni-R1</i>							
	<b>Omni-R1-M</b>	0.560	<b>0.706</b>	<b>0.498</b>	0.083	<b>0.402</b>	+40.7

Table 2: Results on Omni-Bench. Column headers correspond to the four Uni-Tasks: Natural (Natural-Scene Perception), Structured (Structured-Image), Diagrammatic (Diagrammatic Math), and Vision-Op. (Vision-Operational). Bold numbers indicate the best result in each column, and underlined numbers indicate the second-best result. *Average* is computed as a sample-weighted average over slices.  $\Delta\text{Avg} (\%)$  is computed w.r.t. the base model within the same backbone group.

Backbone	Model	MME-P	MME-R	MM-Vet	V*	POPE	MMVP	BLINK	Average	$\Delta\text{Avg} (\%)$
Anole (Base) (Chern et al., 2024)	Anole (Base) (Chern et al., 2024)	224.55	54.64	8.95	21.99	27.74	39.00	21.99	22.25	- (-41.6)
	Zebra-CoT (Li et al., 2025a)	623.70	231.07	7.20	24.61	48.88	37.00	29.04	38.12	+71.3 (-)
<i>Omni-R1</i>										
Anole-7B (Chern et al., 2024)	<b>Omni-R1-M</b>	676.67	259.29	11.70	35.60	62.91	45.33	<b>35.61</b>	46.23	+107.8 (+21.3)
	<b>Omni-R1-L</b>	<u>775.17</u>	<u>279.64</u>	11.97	30.90	61.84	<b>50.67</b>	<u>35.25</u>	<u>48.29</u>	+117.0 (+26.7)
<i>Omni-R1-Zero</i>										
Anole-7B (Chern et al., 2024)	<b>Omni-R1-Zero-T</b>	642.38	277.86	<u>17.16</u>	35.08	<b>66.19</b>	36.67	31.35	45.73	+105.5 (+20.0)
	<b>Omni-R1-Zero-S</b>	<u>699.84</u>	<b>295.71</b>	<b>22.02</b>	<b>37.61</b>	<u>66.06</u>	<u>48.67</u>	33.04	<b>50.19</b>	+125.5 (+31.6)
	<b>Omni-R1-Zero-M</b>	631.27	232.86	14.56	<u>36.96</u>	65.82	47.00	30.41	45.16	+102.9 (+18.4)

Table 3: Results on general benchmarks. Bold numbers indicate the best result in each column, underlined numbers indicate the second-best result. *Average* is computed by first scaling MME-P and MME-R by 0.1 and 0.25, respectively, and then averaging over all seven metrics.  $\Delta\text{Avg} (\%)$  is the relative change w.r.t. Anole (Base), and the value in parentheses is computed w.r.t. Zebra-CoT.

form the baselines on standard multimodal benchmarks, indicating stronger general reasoning capability beyond Omni-Bench. On perception-oriented evaluations, including MME-P (Fu et al., 2023), MMVP (Tong et al., 2024), and BLINK (Fu et al., 2024), Omni-R1 variants are generally stronger, reflecting improved visual perception with multiple reasoning skills. By contrast, Omni-R1-Zero is more competitive on the reasoning-oriented MME-R (Fu et al., 2023) and also performs strongly on the comprehensive MM-Vet (Yu et al., 2024). It also remains stronger on POPE (Li et al., 2023), suggesting better evidence consistency and reduced hallucination.

### 5.3 In-depth Analysis

While the main results establish consistent gains, they do not explain which components drive them or how the models behave. We therefore per-

Model	Natural	Structured	Diagrammatic	Vision-Op.	Avg.
<b>Omni-R1-Zero</b>	<b>0.410</b>	<b>0.170</b>	<b>0.120</b>	<b>0.093</b>	<b>0.159</b>
w/o PerPO	0.390	0.130	0.075	0.033	0.113
w/o PerReward	0.330	0.150	0.085	0.080	0.130

Table 4: Ablation study on Omni-Bench for Omni-R1-Zero. Results are reported on four Uni-Task slices. *Avg.* denotes the average over all samples.

formed ablations and qualitative analyses on the Anole (Chern et al., 2024) backbone to provide more insights.

**Ablation study.** We conducted an ablation study on Omni-R1-Zero to disentangle the effects of PerPO and the perception-calibrated reward, by either removing PerPO entirely (w/o PerPO) or disabling the perception reward while keeping PerPO (w/o PerReward). As shown in Table 4, removing PerPO results in the largest drop, with the regression concentrated on the Vision-Op. and Diagram-

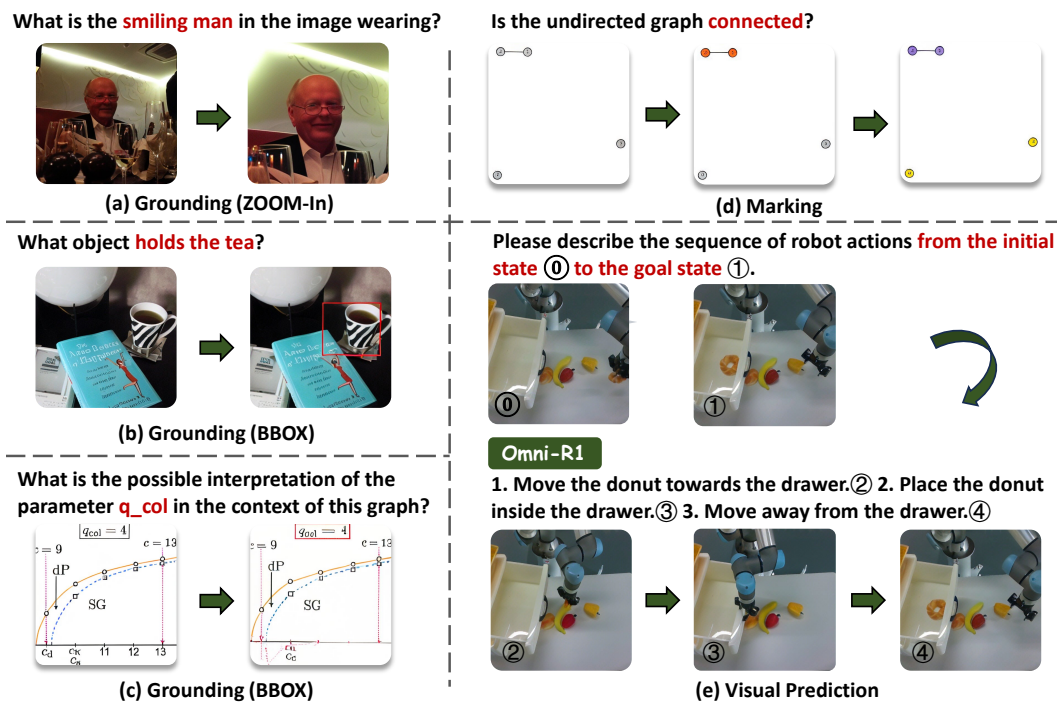


Figure 4: Case studies of Omni-R1 (Anole)’s multimodal reasoning skills, including (a) Grounding (Zoom-In) for attribute recognition, (b) Grounding (BBOX) for target localization, and (c) Grounding (BBOX) for parameter localization. (d) Marking for graph connectivity verification and (e) Visual prediction for robotic manipulation sequences.




Figure 5: The t-SNE visualization of generated images from Omni-R1 and Omni-R1-Zero, respectively. Filled markers indicate correct predictions and empty markers indicate incorrect predictions.

matic tasks. This suggests that PerPO is crucial for developing multi-step and complex multimodal reasoning. Disabling the perception-calibrated reward yields a smaller but consistent degradation, noticeably on Natural tasks. This suggests that perception-calibrated reward improves visual evidence utilization and stabilizes policy optimization.

**Analysis on generated multimodal patterns.** We studied how Omni-R1 and Omni-R1-Zero shape the distribution of generated multimodal patterns differently and how these patterns relate to reasoning. In Figure 5, Omni-R1 concentrates into a few compact modes, whereas Omni-R1-Zero exhibits a more dispersed and multi-modal structure.

**Question:** Which of these would appear shinier when polished?  
A. Metal spoon  
B. Wooden spoon

**Omni-R1-Zero**

**Step 1:**  
The image suggest that a metal spoon is silver in color and a wooden spoon is brown in color.<image>

**Step 2:**  
Therefore, it can be inferred that the metal spoon will appear shinier when polished.<image>

**Answer:** A. Metal spoon

Figure 6: Omni-R1-Zero’s generative reasoning process on a commonsense multimodal question.

In both projections, correct instances cluster more tightly than incorrect ones, suggesting that trace supervision promotes canonical and stable generations, while reward guidance encourages broader exploration that can still support correct decisions.

**Case study.** To better understand how our model behaves, we presented several examples of the reasoning trajectories produced by our models. Figure 4 shows that Omni-R1 produces task-conditioned intermediate visual evidence throughout generation. This supports our claim that Uni-Skills are flexibly composed within a unified generative process. Figure 6 shows that Omni-R1-Zero,

despite having no supervised multimodal traces, can still generate intermediate visual evidence before getting the final answer. More complete generative multimodal reasoning cases can be found in Appendix E.

## 6 Conclusion and Future Work

We revisit multimodal reasoning as generative multimodal reasoning, identifying two practical bottlenecks: functional image generation and costly interleaved-modal reasoning annotations. Omni-R1 targets the former, whereas Omni-R1-Zero targets the latter. Empirically, results on both unified and standard evaluations support these choices: Omni-R1 is particularly effective when reliable intermediate visual evidence is crucial, and Omni-R1-Zero still yields consistent gains despite limited multimodal trace supervision. A key direction is to develop scalable supervision signals for generative multimodal reasoning without annotation, thereby unlocking the full potential of zero-shot settings.

## Limitations

This work demonstrates the potential of unified generative multimodal reasoning, but there are also certain limitations: 1) Inference cost. Similar to the reasoning behind LLMs, our approach enhances performance by allocating additional test-time compute for visual generation, which may introduce latency in certain settings. 2) Interpretability of Omni-R1-Zero. Omni-R1 generally exhibits a higher level of interpretability than text-only reasoning. In Omni-R1-Zero, this advantage may be weakened if the visual generations in the interleaved reasoning traces become abstract, reducing its interpretability to a level closer to that of text-only reasoning.

## Ethics Statement

The datasets used in our experiments are publicly released open-source benchmarks and data, and our work does not involve collecting or annotating new human-labeled data. As a result, privacy risks are avoided, and no personally identifiable information is intentionally included. Moreover, our experiments do not involve safety-critical applications, so trustworthiness-related risks are avoided in our setting. The scientific artifacts we use are available for research under permissive licenses, and our use in this paper is consistent with their intended purposes and license terms.

## Acknowledgments

The work described in this paper was supported by the Research Grants Council of Hong Kong (PolyU/15213323, PolyU/15205325, PolyU/15207122, PolyU/15209724) and PolyU internal grants (BDWP).

## References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. *Qwen2.5-vl technical report*. *Preprint*, arXiv:2502.13923.
- Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu, and Wanxiang Che. 2024. M3cot: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8199–8221.
- Xinyan Chen, Renrui Zhang, Dongzhi Jiang, Aojun Zhou, Shilin Yan, Weifeng Lin, and Hongsheng Li. 2025. Mint-cot: Enabling interleaved visual tokens in mathematical chain-of-thought reasoning. *arXiv preprint arXiv:2506.05331*.
- Ethan Chern, Zhulin Hu, Steffi Chern, Siqi Kou, Jiadi Su, Yan Ma, Zhijie Deng, and Pengfei Liu. 2025. Thinking with generated images. *arXiv preprint arXiv:2505.22525*.
- Ethan Chern, Jiadi Su, Yan Ma, and Pengfei Liu. 2024. Anole: An open, autoregressive, native large multimodal models for interleaved image-text generation. *arXiv preprint arXiv:2407.06135*.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. 2025. *Emerging properties in unified multimodal pretraining*. *Preprint*, arXiv:2505.14683.
- Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, and 1 others. 2024. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 11198–11201.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.

- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. 2024. BLINK: Multimodal large language models can see but not perceive. *arXiv preprint arXiv:2404.12390*.
- Jun Gao, Yongqi Li, Ziqiang Cao, and Wenjie Li. 2025. Interleaved-modal chain-of-thought. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19520–19529.
- Yushi Hu, Weijia Shi, Xingyu Fu, Dan Roth, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, and Ranjay Krishna. 2024. Visual sketchpad: Sketching as a visual chain of thought for multimodal language models. *Advances in Neural Information Processing Systems*, 37:139348–139379.
- Ang Li, Charles Wang, Deqing Fu, Kaiyu Yue, Zikui Cai, Wang Bill Zhu, Ollie Liu, Peng Guo, Willie Neiswanger, Furong Huang, and 1 others. 2025a. Zebra-cot: A dataset for interleaved vision language reasoning. *arXiv preprint arXiv:2507.16746*.
- Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, and Furu Wei. 2025b. Imagine while reasoning in space: Multimodal visualization-of-thought. *arXiv preprint arXiv:2501.07542*.
- Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. 2024. [Multimodal ArXiv: A dataset for improving scientific comprehension of large vision-language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14369–14387, Bangkok, Thailand. Association for Computational Linguistics.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Evaluating object hallucination in large vision–language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Singapore. Association for Computational Linguistics.
- Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. 2025. [Flow-grpo: Training flow matching models via online rl](#). *Preprint*, arXiv:2505.05470.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. [Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts](#). *arXiv preprint arXiv:2310.02255*. ICLR 2024.
- Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, and Song-Chun Zhu. 2021. Inter-gps: Interpretable geometry problem solving with formal language and symbolic reasoning. In *The 59th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. [Chartqa: A benchmark for question answering about charts with visual and logical reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279.
- Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Tiancheng Han, Botian Shi, Wenhai Wang, Junjun He, Kaipeng Zhang, Ping Luo, Yu Qiao, Qiaosheng Zhang, and Wenqi Shao. 2025. Mm-eureka: Exploring the frontiers of multimodal reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2503.07365*.
- Debjyoti Mondal, Suraj Modi, Subhadarshi Panda, Rituraj Singh, and Godawari Sudhakar Rao. 2024. Kamcot: Knowledge augmented multimodal chain-of-thoughts reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 18798–18806.
- Vyas Raina, Adian Liusie, and Mark Gales. 2024. [Is llm-as-a-judge robust? investigating universal adversarial attacks on zero-shot llm assessment](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7499–7517.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Chameleon Team. 2024. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. Eyes wide shut? exploring the visual shortcomings of multimodal LLMs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galouédec. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Penghao Wu and Saining Xie. 2024. [V\\*: Guided visual search as a core mechanism in multimodal llms](#). In

Xuecheng Wu, Jiaying Liu, Danlei Huang, Xiaoyu Li, Yifan Wang, Chen Chen, Liya Ma, Xuezhi Cao, and Junxiao Xue. 2025. **Vic-bench: Benchmarking visual-interleaved chain-of-thought capability in mllms with free-style intermediate state representations.** *arXiv preprint arXiv:2505.14404*.

Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2024. **MM-vet: Evaluating large multimodal models for integrated capabilities.** In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 57730–57754. PMLR.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. **Judging llm-as-a-judge with mt-bench and chatbot arena.** *arXiv preprint arXiv:2306.05685*.

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

Hao Zhong, Muzhi Zhu, Zongze Du, Zheng Huang, Canyu Zhao, Mingyu Liu, Wen Wang, Hao Chen, and Chunhua Shen. 2025. **Omni-r1: Reinforcement learning for omnimodal reasoning via two-system collaboration.** *ArXiv*, abs/2505.20256.

Hengguang Zhou, Xirui Li, Ruochen Wang, Minhao Cheng, Tianyi Zhou, and Cho-Jui Hsieh. 2025. **R1-zero’s “aha moment” in visual reasoning on a 2b non-sft model.** *arXiv preprint arXiv:2503.05132*.

## A Problem Formulation for Generative Multimodal Reasoning

### A.1 Trajectory Formulation

Given a multimodal input  $x^M$  (e.g., an image) and a textual input  $x^T$  (e.g., the question), we model generative multimodal reasoning as an interleaved trajectory of textual rationales and executable visual actions. The trajectory is represented as

$$\mathcal{Y} = \left( x^M, x^T, \{(rat_\ell^T, a_\ell, rat_\ell^M)\}_{\ell=1}^L, Ans \right), \quad (17)$$

where  $L$  denotes the number of reasoning steps,  $rat_\ell^T$  is a textual rationale,  $a_\ell$  is an atomic visual action,  $rat_\ell^M$  is the post-action visual rationale (an image), and  $Ans$  is the final textual answer.

**Policy and execution.** At each step  $\ell$ , the policy  $\pi_\theta$  produces a textual rationale and selects an action:

$$(rat_\ell^T, a_\ell) \sim \pi_\theta(\cdot | x^T, rat_{\ell-1}^M, \{(rat_j^T, a_j)\}_{j<\ell}). \quad (18)$$

The selected action is then executed by a renderer/executor  $R$  that deterministically updates the visual state and produces exactly one post-action image:

$$rat_\ell^M = R(rat_{\ell-1}^M, a_\ell), \quad \ell = 1, \dots, L, \quad (19)$$

with initialization

$$rat_0^M := x^M. \quad (20)$$

Finally, the answer is generated conditioned on the full interaction history, e.g.,

$$Ans \sim \pi_\theta(\cdot | x^T, rat_L^M, \{(rat_\ell^T, a_\ell)\}_{\ell=1}^L). \quad (21)$$

**Atomic action space.** We use a fixed set of atomic visual actions:

$$\mathcal{A} = \{\text{ZOOM-in, BBOX, MARK, LINE, PRED}\}, \quad (22)$$

whose argument formats and semantics are specified next.

### A.2 Atomic Action Protocol Details

This subsection specifies the atomic visual action vocabulary  $\mathcal{A}$  used in the interleaved trajectories (Eq. 17–22). Each action  $a_\ell \in \mathcal{A}$  is applied to the current visual state  $rat_{\ell-1}^M$  and yields exactly one post-action image  $rat_\ell^M$  via the executor  $R$  (Eq. 19).

**Coordinate convention.** All spatial arguments are in normalized image coordinates w.r.t.  $rat_{\ell-1}^M$ :  $(0, 0)$  is the top-left and  $(1, 1)$  is the bottom-right corner. For boxes, we use  $(x, y, w, h)$  for the top-left corner and width/height in  $[0, 1]$ . Out-of-range or degenerate arguments are treated as format errors.

**ZOOM-in** $(x, y, w, h)$ . Produces  $rat_\ell^M$  by cropping the specified region from  $rat_{\ell-1}^M$  and resizing it to a fixed resolution (e.g.,  $512 \times 512$ ).

**BBOX** $(x, y, w, h)$ . Produces  $rat_\ell^M$  by overlaying a bounding box on  $rat_{\ell-1}^M$  at the specified region.

**MARK** $(x, y, id)$ . Produces  $rat_\ell^M$  by rendering a visual highlight on  $rat_{\ell-1}^M$  to draw attention to a target instance or region. The highlight may take various forms (e.g., a colored semi-transparent overlay, halo, marker, or contour) anchored at  $(x, y)$ . Optionally, a lightweight identifier  $id$  may be included to facilitate references in  $rat_\ell^T$ .

**LINE**( $x_1, y_1, x_2, y_2$ ). Produces  $rat_\ell^M$  by drawing a line between two normalized endpoints on  $rat_{\ell-1}^M$ .

**PRED**( $\Delta$ ). Produces  $rat_\ell^M$  by applying a one-step, task-dependent state transition to  $rat_{\ell-1}^M$ , where  $\Delta$  is a serialized delta (text) specifying the update. *Note*: PRED updates the visual state (the post-action image) rather than directly outputting the final answer.

## B Rule-based verifier for $\mathcal{R}_{Acc}$

This appendix summarizes the rule-based verifier used to compute the accuracy reward  $\mathcal{R}_{Acc}$ . The verifier is deterministic and lightweight: it extracts a final answer from the generated trajectory and compares it to the ground truth using a small set of ordered checks. The design emphasizes reliability (avoiding false positives) and robustness to superficial formatting variation.

**Answer extraction and normalization.** The verifier first locates a dedicated final-answer segment using an explicit marker (e.g., `Final Answer:`). The extracted string is then lightly normalized to remove non-semantic variation such as whitespace and punctuation differences, and to strip common answer-introduction phrases (e.g., “Answer is ...”). This step improves the stability of the evaluation without relaxing the core matching criteria.

**Prioritized matching.** The verifier applies checks in a reliability-first order: (i) *numeric matching* when both answers can be interpreted as numbers (including common formats such as percentages); (ii) *symbolic/mathematical equivalence* when the answer contains explicit mathematical expressions; (iii) *conservative textual matching* for general free-form strings; and (iv) *minimal domain-specific handling* for a small set of high-risk patterns (e.g., multiple-choice labels or chemistry-style strings) to reduce known false-positive cases. When multiple checks are applicable, the verifier prefers the most reliable one.

**Scoring.** The accuracy reward  $\mathcal{R}_{Acc} \in [0, 1]$  reflects the degree of match. Checks based on numeric or symbolic equivalence typically yield binary outcomes, while textual matching may assign partial credit only when the match is sufficiently close under conservative criteria. The format reward  $\mathcal{R}_{Fmt}$  is computed independently and only

indicates whether the trajectory is well-formed and parsable.

## C Dataset Details

**Training data.** For Omni-R1, we sample Zebra-CoT (Li et al., 2025a), an 182K image-text interleaved multimodal reasoning dataset, as the supervised data for the PeSFT stage. For Omni-R1-Zero, the PeSFT-stage data is bootstrapped from a small text-only CoT subset of M3CoT (Chen et al., 2024) using the synthetic interleaved construction in Section 4.2, yielding 791 interleaved samples. For the PeRPO stage, we subsample approximately 10% of Zebra-CoT and additionally include 1K ArxivQA (Li et al., 2024) and 250 Geometry3K (Lu et al., 2021) training samples, which do not provide interleaved multimodal reasoning annotations.

**General benchmarks.** We report detailed information on the general benchmarks, including task formulations, dataset sizes, and evaluation metrics. Detailed description is shown in Table 5.

## D Training Details

**Implementation Details.** Our implementation is built upon the open-source libraries verl (Sheng et al., 2024), Transformers (Wolf et al., 2020), Flow-GRPO (Liu et al., 2025) and trl (von Werra et al., 2020), following their respective licenses.

**Hyperparameter Settings.** We summarize the hyperparameter settings used in our experiments in Table 6. The configurations are organized into four groups: (i) Objective, (ii) Batching, (iii) Optimization, and (iv) Rollout.

## E Additional Results

To qualitatively illustrate how Omni-R1 performs generative multimodal reasoning, we present additional full interleaved trajectories in Figures 7–9. These examples span diverse tasks and capability requirements, showing how the model generates intermediate visual evidence and composes UniSkills across different scenarios.

## F Prompt Templates

To reduce evaluation variance and ensure reproducibility on Omni-Bench, we adopt a fixed LLM-judge prompt template, shown in Table 7. The template enforces a binary decision using only the ground-truth answer and the model output, and

Benchmark	Details
MME-P (Fu et al., 2023)	<b>Task:</b> Binary visual QA focusing on perception and recognition. <b>Size:</b> 1,536 questions. <b>Metric:</b> MME score for the perception split. <b>Notes:</b> Reported as <i>MME-P</i> in Table 3.
MME-R (Fu et al., 2023)	<b>Task:</b> Binary visual QA emphasizing reasoning and cognition. <b>Size:</b> 280 questions. <b>Metric:</b> MME score for the cognition split. <b>Notes:</b> Reported as <i>MME-R</i> in Table 3.
MM-Vet (Yu et al., 2024)	<b>Task:</b> Open-ended multimodal QA covering a wide range of capabilities. <b>Size:</b> 218 examples. <b>Metric:</b> LLM-judge score on a 0–100 scale. <b>Notes:</b> We use the standard evaluation protocol in VLMEvalKit.
V*-Bench (Wu and Xie, 2024)	<b>Task:</b> Detail-sensitive visual understanding and visual search. <b>Size:</b> 191 images. <b>Metric:</b> Accuracy. <b>Notes:</b> Reported as <i>V*</i> in Table 3.
POPE (Li et al., 2023)	<b>Task:</b> Yes/No probing for object hallucination. <b>Size:</b> 3,000 questions. <b>Metric:</b> Overall F1 score. <b>Notes:</b> In Table 3, POPE corresponds to Overall F1 rather than accuracy.
MMVP (Tong et al., 2024)	<b>Task:</b> Diagnostic visual patterns with challenging paired examples. <b>Size:</b> 300 questions. <b>Metric:</b> Accuracy. <b>Notes:</b> Evaluated with the default setup in VLMEvalKit.
BLINK (Fu et al., 2024)	<b>Task:</b> Multi-choice evaluation of core visual perception skills across diverse sub-tasks. <b>Size:</b> 3,807 questions. <b>Metric:</b> Accuracy. <b>Notes:</b> Evaluated with the default setup in VLMEvalKit.

Table 5: Details of general multimodal benchmarks.

standardizes the handling of common edge cases, including numeric formatting, unit conversion, and multiple-choice variants.

Hyperparameter	Setting	Description
<b>Objective</b>		
KL Loss Coefficient	0.003	KL regularization weight
PPO Clip Range (Low)	0.20	lower clipping bound
PPO Clip Range (High)	0.28	upper clipping bound
<b>Batching</b>		
Train Batch Size	64	prompts per update
Gen Batch Size	192	3 × train batch
PPO Mini-batch Size	16	per PPO mini-batch
<b>Optimization</b>		
Learning Rate	$1 \times 10^{-7}$	peak LR
Warmup Steps	5	LR warmup
Weight Decay	0.10	L2 regularization
Entropy Coefficient	0.001	entropy bonus weight
Gradient Clipping	1.0	global-norm clip
Loss Aggregation	seq-mean-token-mean	loss reduction
<b>Rollout</b>		
Responses per Prompt	16	samples per prompt
Temperature	1.0	sampling temperature
Top- $p$	0.95	nucleus threshold

Table 6: Training Hyperparameters. Settings used in our experiments.

● **Omni-R1**


**Question:** *What object holds the coffee drink shown in the image?*



**Omni-R1 :**

THOUGHT 0: I need to examine the provided image to identify the container holding the coffee.

THOUGHT 1: To clearly identify the container, I will focus on the object directly beneath the coffee cup.<image\_1>



THOUGHT 2: The zoomed-in view clearly shows a white, circular dish with a handle, which is characteristic of a plate.

THOUGHT 3: Based on this visual evidence, the object containing the coffee is identifiable.

Final Answer: plate

● **Omni-R1**


**Question:** *What is the smiling man in the image wearing?*



**Omni-R1 :**

THOUGHT 0: To answer what the smiling man is wearing, I need to closely examine his attire and accessories in the provided image.

THOUGHT 1: The original image is a bit blurry for clear detail. To get a clearer view of the man's face and accessories, I will focus on that specific area.<image\_1>



THOUGHT 2: The zoomed-in image clearly shows the man wearing glasses. This detail is crucial for answering the question.

THOUGHT 3: Based on the visual evidence from the detailed view, the item worn by the smiling man is identifiable.

Final Answer: glasses

Figure 7: Additional case results of Omni-R1.

## Omni-R1


**Question:** *Is the undirected graph shown below connected?*




**Omni-R1 :**

THOUGHT 0: ...Let's use BFS starting from node 0. We will keep track of visited nodes and use colors to visualize the process: pale gray for unvisited, gold for the current node being processed, medium purple for previously visited nodes, and tomato red for newly discovered neighbors added to the queue.

THOUGHT 1: We begin the BFS algorithm. Initialize a queue for nodes to visit and a set to store visited nodes. Start with node 0. Add node 0 to the queue and mark it as visited. The first visualization shows the state after initiating the process from node 0. Node 0 is colored gold, indicating it's the current node being processed. Its only neighbor, node 2, is colored tomato red, indicating it has been discovered and added to the queue. Nodes 1 and 3 remain pale gray, indicating they are unvisited.<image\_1>



THOUGHT 2: The first visualization shows node 0 (gold) being processed. According to the BFS algorithm, we now explore its neighbors. Looking at the graph structure, node 0 has no edges connecting it to any other node. Therefore, node 0 has no neighbors. Since node 0 has no neighbors, no new nodes are discovered or added to the queue. The queue becomes empty. The second visualization highlights the final state of the component that was explored starting from node 0. Node 0 is marked as visited (medium purple), while nodes 1 and 3 remain unvisited (pale gray).<image\_2>



THOUGHT 3: The second visualization confirms the result of the BFS starting from node 0. Only node 0 was reached. Nodes 1 and 3 were not visited. This indicates that there is no path from node 0 to nodes 1 or 3. Therefore, the graph is not connected. It consists of multiple disconnected components.  
Final Answer: No

Figure 8: Additional case results of Omni-R1.

## Omni-R1

**Question:** Identify joined collapse areas (4-neighbor policy, no corner-to-corner links)



**Omni-R1 :**

THOUGHT 1: I identify connected holes in the grid using 4-connectivity.<image\_1>




THOUGHT 2: Based on the image, I can see 4 connected components marked with different colors.

Final Answer: There are 4 connected hole components.

## Omni-R1


**Question:** The image shows a jigsaw puzzle with a missing section. Identify which of the four provided options completes the puzzle.



**Omni-R1 :**

THOUGHT 0: The problem requires identifying the correct piece to fill the gray area in the puzzle image.

THOUGHT 1: To determine the correct piece, I need to see what the complete image looks like in the area where the piece is missing. I will use a visual aid showing the full scene.<image\_1>



THOUGHT 2: The visual aid shows the complete image. By comparing the missing area in the original puzzle image with the corresponding section in this complete image, I can identify the correct pattern and details needed for the missing piece.

THOUGHT 3: Comparing the required details from the complete image to the provided options, option B matches the content and shape needed to fill the gap.

Final Answer: B

Figure 9: Additional case results of Omni-R1.

---

**System Prompt:**

You are a strict QA judge.

Decide correctness by comparing **ONLY** the ground-truth answer and the model answer. The question may be multiple-choice or open-ended.

**OUTPUT:**

Return exactly one token with **NO** quotes/punctuation/spaces/code fences: True or False.

**GENERAL RULES:**

- 1) Judge factual/semantic equivalence; ignore phrasing, filler, or reasoning text. If the model's final claim contradicts the ground truth or hedges without committing, return False.
- 2) Numbers: allow formatting differences (1,000 vs 1000), scientific notation, or rounding that preserves the stated value. If units are present, require the same value after conversion; missing/extra incompatible units => False.
- 3) Lists/sets: require the same items; order doesn't matter. Missing or extra items => False.
- 4) Spans/names: accept common synonyms and aliases that uniquely indicate the same entity.
- 5) If ambiguous, empty, multiple conflicting answers, or cannot be judged, return False.

**SPECIAL RULES FOR MULTIPLE-CHOICE** (only when options are provided below):

A) Treat option **LETTERS** and their **NUMERIC ORDINALS** as equivalent (A=1, B=2, C=3, ...), but **ONLY** within this question's options.

B) Treat the **CORRECT OPTION'S FULL TEXT** as equivalent to its letter and numeric index.

Your final output must be only the single token: True or False.

---

**User Prompt:****### Ground Truth**

{Ground Truth}

**### Model Answer**

{Model Answer}

**### Decision**

Return only one of: True, False.

---

Table 7: Template for Omni-Bench's GPT Judge.