
Unveiling Chain of Step Reasoning for Vision-Language Models with Fine-grained Rewards

Honghao Chen^{1,2,3*} Xingzhou Lou^{1,2*} Xiaokun Feng^{1,2*} Kaiqi Huang^{1,2†} Xinlong Wang^{3†}

¹Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³Beijing Academy of Artificial Intelligence

Abstract

Chain of thought reasoning has demonstrated remarkable success in large language models, yet its adaptation to vision-language reasoning remains an open challenge with unclear best practices. Existing attempts typically employ reasoning chains at a coarse-grained level, which struggles to perform fine-grained structured reasoning and, more importantly, are difficult to evaluate the reward and quality of intermediate reasoning. In this work, we delve into chain of step reasoning for vision-language models, enabling assessing reasoning step quality accurately and leading to effective reinforcement learning and inference-time scaling with fine-grained rewards. We present a simple, effective, and fully transparent framework, including the step-level reasoning data, process reward model (PRM), and reinforcement learning training. With the proposed approaches, our models set strong baselines with consistent improvements on challenging vision-language benchmarks. More importantly, we conduct a thorough empirical analysis and ablation study, unveiling the impact of each component and several intriguing properties of inference-time scaling. We believe this paper serves as a baseline for vision-language models and offers insights into more complex multimodal reasoning. Our dataset, PRM, and code at <https://github.com/baaivision/CoS>.

1 Introduction

The complex reasoning capabilities of large language models (LLMs) have recently garnered widespread attention [57, 8, 47, 68, 1]. By generating a detailed chain of thought (CoT) process before producing the final answer, LLMs enable effective test-time scaling and significantly enhancing the reasoning and logical capabilities. The integration of CoT reasoning with large-scale reinforcement learning [43] has further led to revolutionary advancements, exemplified by groundbreaking models such as OpenAI-o1 [17] and DeepSeek-R1 [12]. Inspired by the success of NLP, many studies focus on exploring complex reasoning methods in vision-language models (VLMs), achieving various advances [66, 58, 10, 55, 34, 38, 29]. Compared to LLMs, the presence of visual information in VLMs introduces greater complexity and plays a more critical role, emphasizing the necessity of leveraging visual cues for precise step-by-step reasoning and problem-solving.

Despite the advances, enabling VLMs to perform human-level reasoning remains a key challenge [61]. A fundamental difference is that humans typically engage in systematic and structured thinking [48], whereas the CoT outputs of current VLMs are constrained to coarse-grained reasoning, making it difficult to execute systematic and structured inference [58]. Specifically, current VLMs typically output a lengthy thought without a unified structure or clear step-by-step progression. This form

*Equal contribution. Work done during Honghao’s internship at BAAI.

†Corresponding authors.

of reasoning tends to become disorganized or verbose and, more importantly, makes it difficult to evaluate the quality of reasoning. As a result, selecting high-quality reasoning trajectories for efficient training (e.g., reinforcement learning) or inference becomes challenging. It is natural to ask whether step-by-step fine-grained reasoning can facilitate better training or inference. However, answering this is non-trivial, as it involves two major challenges: **i)** defining step, i.e., how to decompose the overall reasoning chain into logically coherent and progressive steps; **ii)** evaluating step, i.e., how to provide fine-grained reward signals for step-wise reasoning to facilitate training and inference.

To address these challenges and answer the above question, we propose Chain-of-Step reasoning and introduce corresponding supervised fine-tuning (SFT) and reinforcement learning strategies to enhance VLMs’ reasoning capabilities. We start by defining step: the entire reasoning chain is decomposed into consecutive steps, with each step consisting of three components: Name, Thought, and Reflection, separated by special tags. The Name component provides a concise summary of the step, the Thought component contains detailed reasoning specific to that step, and the Reflection component establishes connections between the current reasoning step, the visual information, and previous steps. Our design boasts three essential advantages: **i)** structuring reasoning into discrete steps to ensure clarity; **ii)** maintaining a concise and well-organized structure; **iii)** incorporating visual content along with reflective reasoning. We use GPT-4o [16] to construct ShareGPT-Step-300K, a dataset of 300K structured step-wise reasoning samples following the designed template. Our dataset encompasses a diverse range of tasks, utilizing 17 datasets that demand various reasoning skills, such as scientific reasoning [18, 31], mathematical reasoning [51, 64] and world knowledge [41].

Through SFT on ShareGPT-Step-300K, VLMs can output step-by-step reasoning chains. However, reinforcement learning is more crucial for enhancing the model’s reasoning capabilities [66], and providing accurate rewards is of vital importance [34]. Therefore, we delve into the precise evaluation of reasoning steps. We collect process-annotated data using two mainstream methods: Monte Carlo estimation and LLM-as-Judge [67]. Specifically, we employ Math-Shepherd [52] and GPT-4o [16] to annotate the step-wise reasoning trajectories, which we then use to train a process reward model (PRM). Based on PRM, we can assess the quality of intermediate reasoning steps with high accuracy, thereby providing fine-grained rewards.

Building upon the above framework, we can effectively choose reasoning paths to conduct iterative Direct Preference Optimization (DPO) [7] and inference-time scaling. Experiments across multiple benchmarks demonstrate that our approach significantly enhances reasoning capabilities, outperforming strong baselines with consistent improvements. Our empirical analysis and ablations indicate that incorporating step-level evaluation is more reasonable and effective than relying solely on answer-level evaluation. Interestingly, we find that longer reasoning is not necessarily better for VLMs, and reasoning steps’ quality is much more important than reasoning length. Different from prior arts, our work highlights the importance of introducing fine-grained reasoning chains in complex multimodal reasoning and leveraging fine-grained reward-based reinforcement learning to enhance VLMs. We hope our findings provide new insights for the community.

2 Related Work

VLM Reasoning. As the application of VLMs in complex tasks such as mathematics [30, 51] and science [61, 18, 31] grows, the ability to reason through complex problems becomes increasingly important. For instance, in solving mathematical diagram problems, Mavis [65] constructs a corresponding chain-of-thought (COT) [57] reasoning dataset and optimized the model through instruction fine-tuning [40]. Visual-CoT [42] focuses on enhancing the model’s ability to correspond visual questions with key areas using CoT reasoning. Insight-V [10] designs a multi-agent system consisting of a reasoning agent dedicated to performing long-chain reasoning and a summary agent trained to judge and summarize reasoning results. The success of CoT in improving VLMs’ complex visual reasoning reveals the importance of enabling the model to analyze intermediate reasoning steps in detail [66]. Unlike most existing methods [21, 42, 9] that use reasoning chains at a coarse-grained level, our work introduces a fine-grained structured reasoning approach, offering greater potential for enhancing the reasoning capabilities of VLMs.

Inference-Time Scaling. Inference time scaling [45] refers to the process wherein a model utilizes additional reasoning time to solve problems more effectively. This System-2 style thinking ability [39] is recognized as an effective way to enhance the complex reasoning capabilities of LLMs and

has recently garnered significant attention [69, 12]. Traditional inference time scaling methods include majority voting [14], best-N search [56], and sentence-level beam search [49]. Despite their widespread application, majority voting is typically only for problems with standard answers, whereas best-N search and sentence-level beam search present challenges in terms of accuracy assessment and the quality of individual responses [58]. Recently, OpenAI has developed verifiers to supervise and select reasoning paths during the inference process [25], while Math-Shepherd [52] evaluates intermediate reasoning steps based on their likelihood of leading to correct answers. Inspired by the success of these LLMs, LLaVA-CoT [58] and URSA [34] have started exploring the inference time scaling capabilities of VLMs. Procedural reasoning datasets have been constructed and used for the supervised fine-tuning of the models. Different from those prior arts, we utilize PRM to provide granular rewards and facilitate step-level beam search to enhance the inference-time scaling of VLMs.

3 Method

In this section, we present the pipeline of CoS as illustrated in Fig. 1. In Section 3.1, we introduce the design of our structured reasoning template. In Section 3.2, we describe the composition of the SFT dataset ShareGPT-Step-300K. Then we explain the acquisition of process annotation data in Section 3.3. Last, in Section 3.4, we elaborate on the details of the RL training pipeline with fine-grained rewards.

3.1 Structured Reasoning

We structure the entire thought into multiple consecutive cognitive steps, ensuring that these steps logically flow in a natural and coherent manner towards the final answer. Each step consists of the following three components:

- **Name.** A summary and overview of the current step, providing a high-level explanation. This outlines the key aspects of the step to be executed.
- **Thought.** A detailed description of the thought process in this step, explaining what is being observed, inferred, or analyzed.
- **Reflection.** Establishing connections between the visual content and previous steps. Since VLMs are prone to generating content that contradicts the image, the reflection is introduced to create links with the visual content and prior steps, thus alleviating hallucination.

We use special tokens to establish this reasoning format, rather than simply relying on additional prompts or external prompt engineering, because depending on prompt-based output format is unstable and would require more data filtering and cleaning. In contrast, ensuring stable output of this format is more critical for splitting and evaluating steps.

CoS begins with the special token `<|reasoning_start|>`. We take reasoning step as the basic unit for structured and fine-grained reasoning. Steps are connected by `<|reasoning_proceed|>`. The reasoning stops after outputting the token `<|reasoning_end|>`, then the model derives the final answer from previous reasoning. The model generates reasoning steps and answers in an autoregressive manner, where the number of reasoning steps and their specific content are self-generated without any restrictions. This structured output and its length are autonomously adapted and adjusted

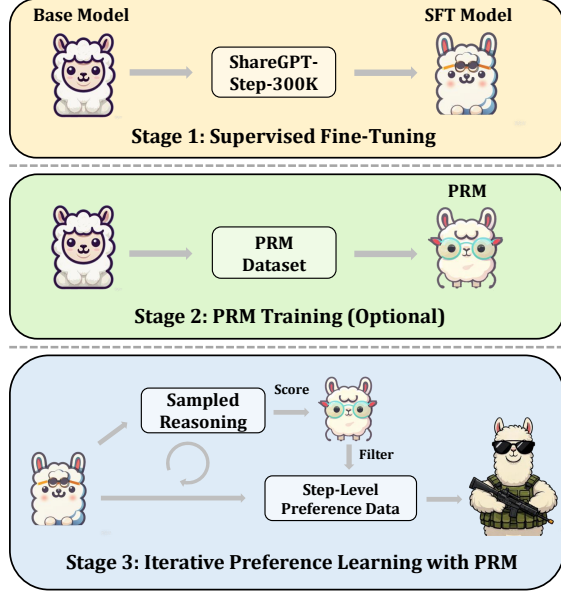


Figure 1: **Overview of our framework.** We conduct SFT to enable the model to generate step-by-step reasoning paths. Then we train a PRM to provide fine-grained RL reward, and enhance the complex reasoning ability through step-level preference learning. Note that PRM only needs to be trained once for different models.

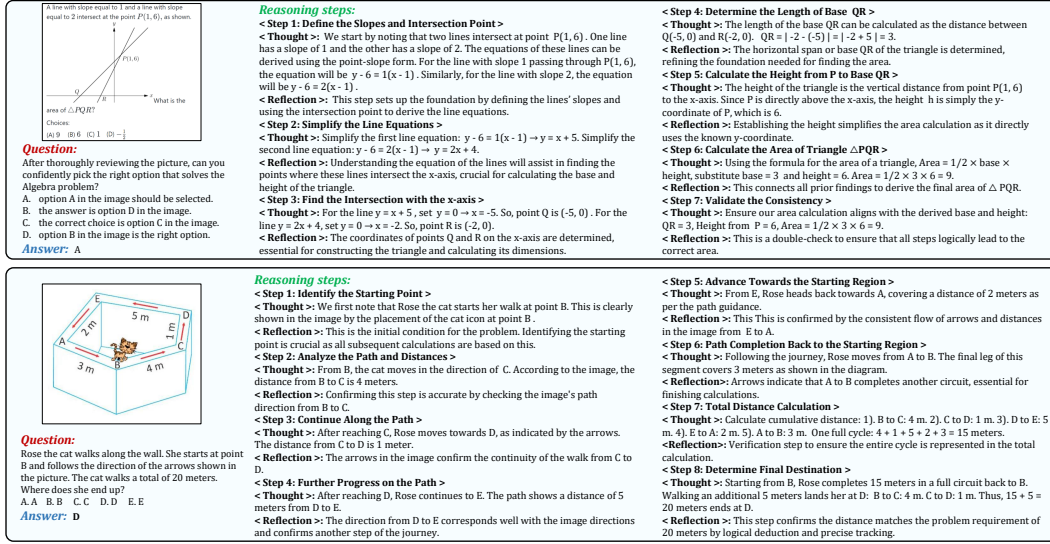


Figure 2: **Qualitative examples.** We depict the reasoning process on cases that need to understand both visual content and mathematics. Our CoS has accurate understandings of both aspects, and more importantly, combine them to derive multimodal reasoning. Note that special tokens are omitted.

by the model, thereby enhancing its ability to perform complex reasoning. We depict two qualitative examples in Fig. 2 and the details of special tokens can be found in Appendix A.

3.2 ShareGPT-Step-300K

Given that most VQA, multi-modal math, and reasoning datasets only provide short answer annotations, we design a method to leverage these short annotations to generate long, step-by-step reasoning data. Specifically, we provide the question and its ground-truth answers as references, and prompts GPT-4o to generate step-by-step reasoning. This can be seen as reasoning from the results. The reference answer can greatly reduce the difficulty of the question, making it easier to generate high-quality thinking steps.

To ensure diversity, our data is collected from various multimodal datasets, encompassing 17 tasks that require diverse multimodal reasoning capabilities. These tasks can be categorized into four major types including mathematical reasoning [51, 64, 3, 26, 24], scientific reasoning [18, 31, 23, 5], chart&document analysis [35, 32, 37, 36], and world knowledge [41, 54, 22, 44]. A summary of the dataset sources can be found in Appendix D.

After generating the data, we conducted meticulous data cleaning and filtering to remove a large amount of data that did not meet the required formatting standards. Ultimately, we obtained 300K high-quality step-level data instances, which we named as ShareGPT-Step-300K. We plan to release this dataset to advance research in the community on fine-grained step-level reasoning.

3.3 Step-Level Annotation

Currently, mainstream methods for evaluating reasoning steps can be categorized into three types [67]: **i)** Human, which relies on human annotators to produce high-quality data but incurs substantial costs; **ii)** Monte-Carlo (MC) Estimation, which infers the correctness of steps from the outcomes of a vast number of sampled solutions; **iii)** LLM-as-Judge, which leverages LLM as critic by directly prompting the LLM to evaluate the correctness of steps. In this work, for the sake of efficiency and scalability, we employ MC Estimation and LLM-as-Judge to annotate step-level data.

For MC Estimation, we use math-shepherd [52] to annotate steps automatically. Specifically, given a problem with a golden answer and a step-by-step solution, to annotate the correctness of a specific step, we utilize a fine-tuned VLM to decode N subsequent reasoning paths from this step. We further validate whether the decoded final answer matches with the golden answer. If a reasoning step can

deduce more correct answers than another, it would be assigned a higher correctness score. Here we use N=16 by default.

For LLM-as-Judge, we utilize GPT-4o [16] to perform the annotations. Each step is evaluated with one of three possible ratings: **i) Good**: the step is logical, directly relevant to the question, and most importantly, correct; **ii) Neutral**: the step is relevant but contains minor flaws or gaps; **iii) Bad**: the step is flawed, wrong, or shows a misunderstanding of the question or context. For solutions where the final answer is correct, both Good and Neutral ratings lead to a correct label. Conversely, for solutions with incorrect final answer, only steps annotated as Good are ultimately marked as correct. The prompt is provided in Appendix B.

We use both math-shepherd and GPT-4o to generate 100K step-annotated data, respectively. These step-annotated data are then utilized to train a PRM, which can provide fine-grained reward signals for reinforcement learning and inference-time scaling.

3.4 RL with Fine-grained Rewards

As illustrated in Fig. 1, our training pipeline is designed to be straightforward and efficient. Starting from a well-pretrained VLM, we first perform supervised fine-tuning on the ShareGPT-Step-300K dataset to enable the model to generate step-by-step reasoning paths. Subsequently, we train a PRM to provide fine-grained rewards. Leveraging the PRM, we efficiently select reasoning paths for iterative Direct Preference Optimization, thereby further enhancing the model’s reasoning capabilities.

Process reward model training. In RLHF, the reward models are typically trained with binary preferences [2]. In reasoning tasks where step-annotation of solutions is accessible, reward models can be trained under the supervision of such ground-truth correctness. Given a question q , process reward model (PRM) estimates the correctness of individual reasoning steps. The PRM r_{process} is trained with the following objective:

$$\mathcal{L}_{\text{PRM}} = \mathbb{E}_{q, p^{(k)}, y_k \sim \mathcal{D}} [\text{Loss}(y_k, r_{\text{process}}(q, p^{(k)}))] \quad (1)$$

where y^k is the label for the partial solution $p^{(k)}$ and Loss is the loss function. Here we adopt binary cross entropy loss by default. We train a InternVL-2.5-MPO-38B [55] on the step-annotated data as introduced in Section 3.3 to get our PRM. We ablate on the choice of PRM in Section 4.3.4.

Iterative direct preference optimization. Preference learning has gained increasing focus in VLMs [10]. To further enhance the quality of reasoning chains, we apply RL using the iterative DPO [7] algorithm to better align the model’s reasoning process towards high-quality reasoning paths. We use SFT model as the policy model π_θ , generating 16 reasoning paths for each question. Each reasoning path is evaluated with the trained PRM to assess its quality. We calculate the average score of reasoning steps and the score of final answer respectively, and take the weighted sum of them as the final score. The score is normalized to $[0, 1]$. Then we select positive-negative pairs depending on the scores. Note that the score of the positive path must exceed that of the negative path by a threshold t to ensure sufficient distinction between them, thereby avoiding cases where the differentiation is minimal.

Formally, we denote the dataset as $\mathcal{D}_{\text{DPO}} = (x, y_+, y_-)$, where x is the question, y_+ and y_- are the positive and negative responses. The DPO objective is defined as below:

$$\mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_+, y_-) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(y_+ | x)}{\pi_{\text{ref}}(y_+ | x)} - \beta \log \frac{\pi_\theta(y_- | x)}{\pi_{\text{ref}}(y_- | x)} \right) \right] \quad (2)$$

where π_θ is the policy model to be optimized and π_{ref} is the reference model and they are both initialized with SFT model. σ is the logistic function and β is set to 0.1. After training, we use the trained policy model π_θ to generate next-round data $\mathcal{D}_{\text{DPO}}^{(2)}$, while keeping π_{ref} fixed (i.e., the SFT model). We repeat this procedure for multiple rounds to progressively enhance the the model’s reasoning.

4 Experiment

In this section, we first introduce the setting and implementation details in Section 4.1. Then we present a comprehensive comparison with state-of-the-art VLMs across multiple reasoning benchmarks in Section 4.2. In Section 4.3, we conduct extensive ablation studies and analyses to validate the effectiveness of our approach and explore several interesting properties.

Table 1: **Comparison with open-source VLMs.** We evaluate our method on six benchmarks covering both general and task-specific reasoning capacities. Our CoS has consistently strong performances across these benchmarks, surpassing different baselines and other state-of-the-art VLMs by large margins. The items in **bold** and underlined respectively represent the first or second highest scores.

Method	Size	MathVista	MMStar	MMMU	M3CoT	AI2D	ChartQA	Average
LLaVA-1.5 [27]	13B	27.6	32.8	-	39.5	-	-	-
IXC-2.5 [63]	7B	63.7	59.9	42.9	-	81.5	82.2	-
Ovis1.5-LLaMA3 [33]	8B	63.0	57.3	48.3	-	-	76.4	-
LLaVA-CoT [58]	11B	54.8	57.6	-	-	78.7	-	-
LlamaV-o1 [50]	11B	54.4	59.5	-	-	81.2	-	-
Vision-R1-LlamaV [15]	11B	62.7	61.4	-	-	-	83.9	-
R1-VL [62]	7B	63.5	60.0	-	-	-	83.9	-
R1-Onevision [59]	7B	64.1	-	-	-	-	-	-
MiniCPM-V-2-6 [60]	8B	60.6	57.5	49.8	56.0	82.1	79.4	64.2
LLaVA-OneVision [21]	7B	63.2	<u>61.7</u>	48.8	52.3	81.4	80.0	64.6
Qwen2-VL [53]	7B	58.2	60.7	53.7	57.8	83.0	77.4	65.1
Insight-V [10]	7B	59.9	61.5	50.2	61.5	79.8	81.5	65.7
InternVL-2.5 [6]	8B	62.2	59.6	<u>54.1</u>	62.4	<u>84.5</u>	84.9	68.0
LLaVA-NeXT-LLaMA3 [28]	8B	45.9	43.1	36.9	45.6	71.5	69.4	52.1
+ SFT	8B	51.4	54.7	39.6	67.4	76.1	75.7	60.8
+ Iterative DPO (CoS-LLaVA)	8B	54.7	58.9	41.8	71.7	79.2	79.1	64.2
InternVL2.5-MPO [55]	8B	65.0	60.7	53.8	67.5	84.2	85.0	69.4
+ SFT	8B	<u>65.9</u>	61.0	53.7	<u>75.7</u>	81.6	88.3	<u>71.0</u>
+ Iterative DPO (CoS)	8B	67.8	63.5	55.5	81.0	84.9	<u>87.4</u>	73.4

4.1 Experimental Setup

Implementation details. To verify the effectiveness and generalizability of our approach, we employ LLaVA-NeXT [28] and InternVL-2.5-MPO [55] as our base VLMs, respectively. For each base VLM, we first conduct one epoch of supervised fine-tuning on ShareGPT-Step-300K. Then, based on the SFT model, we apply three rounds of iterative DPO to obtain the final model. For each round, we compile approximately 20K preference pairs as training data for DPO. The training recipes can be found in Appendix C. For PRM, we select InternVL-2.5-MPO-38B [55] as the base model. We trained it for two epochs on 200K process-annotated data, using a learning rate of 1e-6 and a batch size of 128 to get the PRM. For the training time, SFT takes about 9 hours on a single A800 node, and three rounds of iterative DPO takes about 6 hours in total on a single A800 node.

Benchmarks. We evaluate our method on several representative benchmarks, spreading a wide range of tasks that demand the model’s complex reasoning abilities. MathVista [30] is a well-known dataset for multi-modal reasoning, focusing on mathematics skills such as plane geometry, function and puzzle. MMMU [61] evaluates the model’s multimodal reasoning capability across various university-level disciplines. ChartQA [35] aims at assessing logical reasoning with charts. MMstar [4] focuses on various multi-modal problems that must rely on visual contents. M3CoT [5] evaluates the model’s chain-of-thought reasoning abilities. AI2D [18] is designed to test scientific knowledge and reasoning.

4.2 Main Results

We conducted a comprehensive comparison between our method and leading open-source VLMs on the aforementioned benchmarks. As shown in Table 1, our CoS achieves a substantial improvement over different baselines. Specifically, our method yields an average increase of 4.9% and 3.2% when applied to LLaVA-NeXT and InternVL-2.5-MPO, demonstrating the effectiveness and generalizability of our framework.

For LLaVA-NeXT, both SFT and DPO lead to significant performance gains, indicating that for relatively weaker VLMs, both SFT and reinforcement learning (RL) play crucial roles. In contrast, for InternVL2.5-MPO, the performance improvement from SFT alone is relatively limited. However, after applying iterative DPO, the model can achieve a notable performance boost, reaching state-of-the-art. This indicates that for stronger VLMs, the benefits of further improvement through SFT are marginal, while RL is more essential and effective for enhancing strong model’s reasoning capabilities.

4.3 Analysis and Ablation Study

4.3.1 Reasoning Step Matters

Firstly, we dive into the impact of intermediate reasoning steps and final answers on evaluating reasoning quality. To be specific, we use LLaVA-NeXt [28] fine-tuned on the ShareGPT-Step-300K as base model and sample 16 reasoning paths and corresponding answers for each question in M3CoT. We then utilize a trained PRM to score these reasoning steps and answers. We use the weighted sum of the step score and answer score as the selection criterion to report best-of-16 accuracies.

As shown in Fig. 3, we progressively increase the weight of step scores from 0% (i.e., relying solely on the final answer for evaluation) to 100% (i.e., relying exclusively on reasoning steps). The model’s best-of-16 accuracy exhibits a trend of initially increasing and then decreasing. This indicates that relying solely on either the final answer or the intermediate reasoning steps is not optimal; instead, a balanced evaluation that considers both is preferable. This conclusion also holds for reinforcement learning training as shown in Section 4.3.3. Unless mentioned, we set the weight of step scores to 20% by default in our experiments.

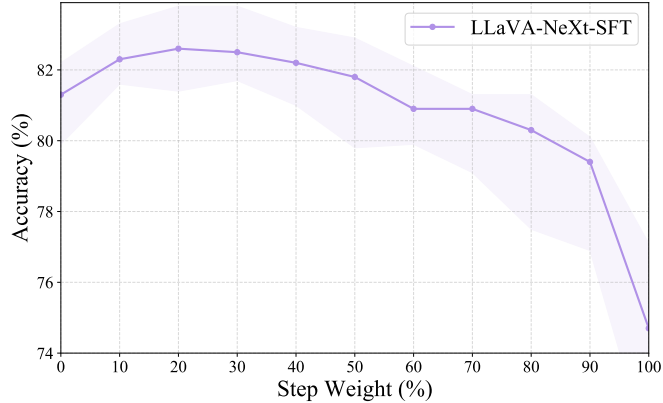


Figure 3: **Ablation on step weight.** We progressively increase the weight of step score from 0% to 100% and report the best-of-16 accuracy on M3CoT. Considering both step’s and answer’s quality leads to optimal performances.

4.3.2 Inference-Time Scaling

Step-level beam search. In order to further enhance the reasoning ability during inference, we design a step-level beam search by making full use of the structured output of CoS. The specific implementation steps are as follows:

- Sample N responses for the first step.
- Use PRM to score the N first steps.
- Continue to sample the next step based on the best step.
- Repeat the above steps until the final answer is generated.

It is worth noting that step-level beam search’s cost and sample efficiency is exactly the same as that of Best-of-N (i.e., independently sample N responses, score with PRM, and select the best). An illustration is shown in Fig. 4.

Next, we explore the inference-time scaling properties of CoS reasoning. We compare four methods: Pass@N, Self-Consistency, PRM, and PRM with step-level beam search. Pass@N refers to the probability that at least one out of N sampled solutions is correct, which can be regarded as the upper bound of inference-time performance. Self-Consistency refers to selecting the most frequent answer among N sampled solutions. PRM refers to selecting the Best-of-N solution based on the scores provided by the PRM. PRM-BS refers to selecting the Best-of-N solution obtained through step-level beam search. We incrementally increase the number of samples N from 1 to 64 to evaluate the performance of these methods. As shown in Fig. 5, the answer selected by PRM has a significant advantage over Self Consistency. Moreover, PRM-BS has achieved consistent improvement over PRM. Considering that the sampling efficiency of PRM-BS and PRM is completely the same, this further illustrates the advantages of our method in the inference-time.

4.3.3 Ablation on RL training

We proceed to investigate the impact of rewards on the RL training. Specifically, we utilize LLaVA-NeXt-SFT as our base model and employ three methodologies to offline select training data for

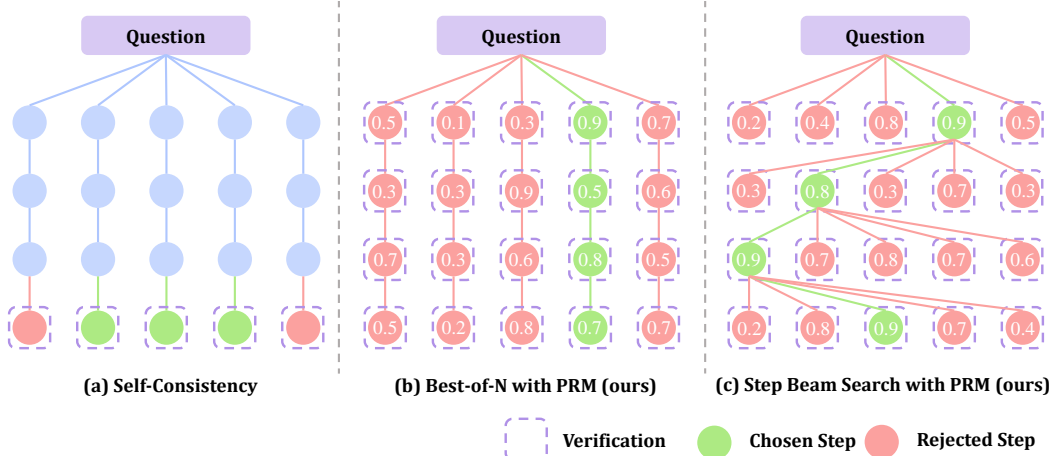


Figure 4: **Illustration of inference-time scaling.** Our structured and fine-grained designs enable us to execute (b) Best-of-N with PRM and (c) Step Beam Search with PRM. Both of them outperforms Self-Consistency by significant margins.

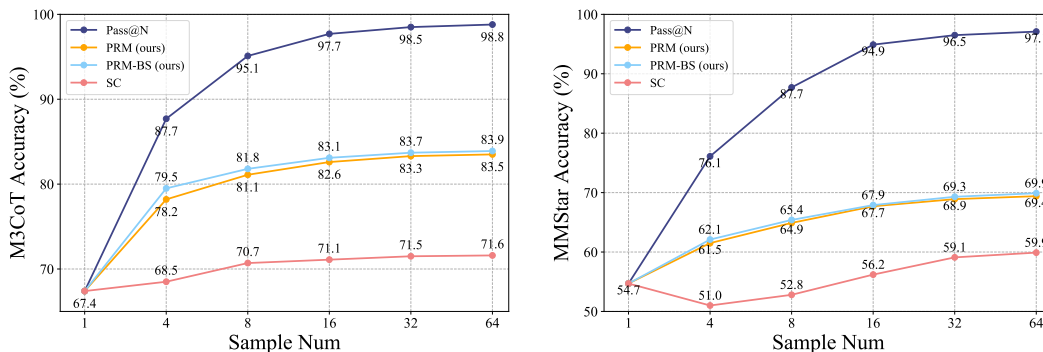


Figure 5: **Pass@N and Best-of-N comparison on M3CoT and MMStar.** Here SC refers to Self-Consistency, PRM-BS refers to step-level beam search. We increase N from 1 to 64, and our method demonstrates appealing inference-time property.

iterative DPO: **i) Outcome.** This approach uses the ground-truth labels of the questions as the gold criterion for judging the quality of responses. Responses with correct answers are selected as positive samples, while those with incorrect answers are negative ones; **ii) PRM scores for answers.** This can be seen as a lower bound on performance compared to the outcome method, as the accuracy of PRM’s evaluation of answers is inherently lower than that of direct evaluation using ground-truth labels; **iii) PRM scores for both reasoning steps and answers.** This method comprehensively considers the quality of step and answer to selects positive and negative pairs.

As shown in Table 2, solely relying on answer scoring for selection underperforms the outcome-based method, which is an anticipated result. The strategy that takes into account both the intermediate steps and the answer quality achieves the best performances, surpassing outcome-based method with clear gains. This underscores the significance of fine-grained rewards in assessing reasoning quality, thereby enhancing the efficacy of RL training.

Table 2: **Ablation on the rewards for RL.** Taking both the steps and answer quality into account leads to the best performances.

Method	MathVista	MMStar	M3CoT
LLaVA-NeXt-SFT	51.4	54.7	67.4
Answer (PRM)	53.1	57.3	69.7
Outcome	53.5	58.1	70.0
Step&Answer (PRM)	54.7	58.9	71.7

Table 3: **Ablation on the PRM.** Stronger PRM leads to more accurate evaluation for reasoning steps.

Base Model	Question	Step Acc	Answer Acc
LLaVA-NeXt-8B	Seen	85.1	90.5
LLaVA-NeXt-8B	Unseen	83.7	84.7
InternVL-2.5-8B	Unseen	85.7	91.1
InternVL-2.5-38B	Unseen	87.3	88.7

4.3.4 Ablation on PRM

To evaluate the accuracy of the PRM, we reserved 10K process-annotated reasoning data for assessment. Among these, 5K questions were from the PRM training dataset, while the other 5K were unseen during PRM’s training, allowing us to evaluate PRM’s ability to assess in-domain data and generalization capability to unseen questions. We employed three base models: LLaVA-NeXt-8B, InternVL2.5-8B, and InternVL2.5-38B. As demonstrated in Table 3, the stronger the base model, the more accurate its judgment of intermediate steps for the PRM. Given the critical importance of step evaluation in our framework, we select InternVL2.5-38B as the base model for our PRM.

4.3.5 Reasoning Length

It has been observed that the length of reasoning tends to increase progressively when LLMs are trained with RL on mathematical problems, and their performance improves accordingly [12]. However, things become different for visual complex reasoning. Existing visual complex reasoning problems do not require lengthy intermediate processes as pure mathematics does, but rather rely on the utilization of visual information and the connection of relevant knowledge to trigger reasoning (e.g., geometric shapes). Therefore, we investigate how the length of VLM reasoning steps changes with training. We take LLaVA-NeXt-SFT as the base model and calculate the average reasoning step after each round of outcome DPO and PRM DPO, respectively. We depict the change curve of reason step length in Fig. 6.

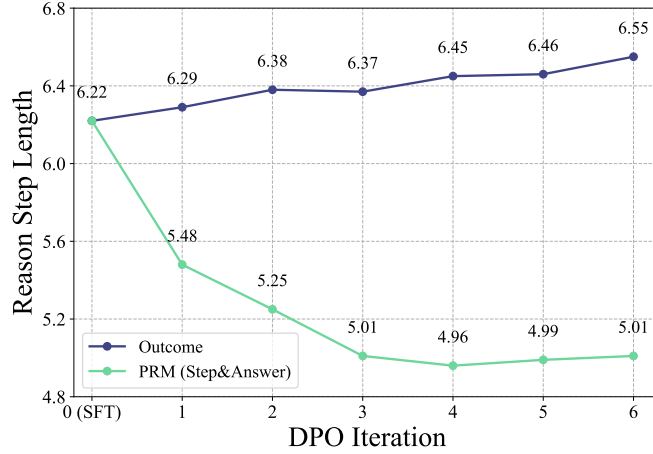


Figure 6: **Reasoning step length curve.** Increasing reasoning length is not necessarily better. For visual complex reasoning, ensuring high-quality steps is of vital importance.

For outcome DPO, we select reasoning paths solely based on the correctness of the final results. It can be observed that as the training iteration progresses, the reasoning step length progressively increases. This indicates that the model tends to keep extending the reasoning length if we don’t consider the quality of intermediate steps. Interestingly, the situation is entirely the opposite for PRM DPO. The model tends to initially reduce the length of reasoning to improve the quality of reasoning, and only after stabilizing does it slowly increase the length of reasoning. Considering that PRM DPO yields better results than outcome DPO as shown in Table 2, this phenomenon indicates that for visual complex reasoning, increasing reasoning length is not necessarily better. It is more crucial to first ensure that the model can output high-quality reasoning steps.

For outcome DPO, we select reasoning paths solely based on the correctness of the final results. It can be observed that as the training iteration progresses, the reasoning step length progressively increases. This indicates that the model tends to keep extending the reasoning length if we don’t consider the quality of intermediate steps. Interestingly, the situation is entirely the opposite for PRM DPO. The model tends to initially reduce the length of reasoning to improve the quality of reasoning, and only after stabilizing does it slowly increase the length of reasoning. Considering that PRM DPO yields better results than outcome DPO as shown in Table 2, this phenomenon indicates that for visual complex reasoning, increasing reasoning length is not necessarily better. It is more crucial to first ensure that the model can output high-quality reasoning steps.

4.4 Ablation on Reasoning Pattern

In order to demonstrate our reasoning design’s advantages, we add a direct quantitative comparison between the proposed method and simpler prompting strategies in Table below. Specifically, No Reason denotes a setting where the model performs direct question answering without any explicit reasoning; Direct Prompt means directly prompting VLMs to reason step by step (as in [19]); CoS corresponds to our proposed method using step-level structured reasoning with fine-grained rewards. All methods are trained using the same SFT and RL datasets to ensure a fair comparison. For No

Table 4: **Ablation on the reasoning pattern.** No Reason denotes a setting where the model performs direct question answering without any explicit reasoning; Direct Prompt means directly prompting VLMs to reason step by step. The corresponding SFT dataset’s format has been adjusted accordingly to make a fair comparison.

Method	Reason	Reward	MathVista	MMStar	M3CoT	Average
LLaVA-NeXt	-	-	45.9	43.1	45.6	44.9
No Reason SFT	×	-	49.1	53.9	62.1	55.0
No Reason RL	×	outcome	51.5	56.4	63.4	57.1 (+2.1)
Direct Prompt SFT	✓	-	51.2	54.9	66.3	57.5
Direct Prompt RL	✓	outcome	53.1	58.2	69.3	60.2 (+2.7)
CoS SFT	✓	-	51.4	54.7	67.4	57.8
CoS RL	✓	PRM	54.7	58.9	71.7	61.8 (+4.0)

Table 5: **Ablation for GRPO experiments.** Our Cos has consistently better performance than the outcome-reward counterpart.

Method	Reward	MathVista	MMStar	M3CoT	Average
LLaVA-NeXt-SFT	-	51.4	54.7	67.4	57.8
outcome GRPO	outcome	54.3	57.9	71.4	61.2
CoS GRPO	PRM	56.3	59.1	73.7	63.0

Reason, we use only the question-answer pairs (without intermediate reasoning thought) from the ShareGPT-Step-300K dataset for SFT. For Direct Prompt, we concatenate all steps of reasoning (i.e., "thoughts") from each instance into a single reasoning sequence. Since both No Reason and Direct Prompt can not output strict step-level thought, we cannot assess their intermediate reasoning quality, so we apply the outcome reward (i.e., correctness of the final answer) during RL training.

There are two key observations: i) Models trained with reasoning-augmented SFT data (i.e., Direct Prompt and Cos) clearly outperform models trained with QA-only data (No Reason) on reasoning benchmarks; ii) Our Cos yields much greater improvements during RL compared to outcome-only reward settings. This ablation study directly supports our central motivation for fine-grained reasoning and strengthens the empirical foundation of our conclusions.

4.5 GRPO-based ablation

Recently, GRPO [43] has become the mainstream algorithm in current reinforcement learning and has demonstrated advantages across multiple tasks. Therefore, we further conducted an ablation experiments based on the GRPO algorithm to validate the effectiveness of our proposed method. For GRPO experiments, we deploy the 38B PRM as the reward server. We set batchsize to 128, rollout.n to 5, lr to 1e-5 and run 100 steps starting from the SFT model. As shown in Table 5, our fine-grained reward RL outperforms outcome reward RL consistently using GRPO, further strengthening our method and analysis.

5 Conclusion

In this work, we introduce Chain of Step (CoS) reasoning for vision-language models, enabling assessing reasoning step quality accurately and leading to effective reinforcement learning and inference-time scaling with fine-grained rewards. Experimental results across multiple benchmarks demonstrate the effectiveness of CoS. More importantly, we conduct extensive empirical analysis and ablations to unveil CoS’s appealing properties. We hope this paper offers insights into more complex multi-modal reasoning.

Acknowledgments

This work is supported in part by the National Science and Technology Major Project, Grant No.2022ZD0116403.

References

- [1] Shengnan An, Bo Zhou, Zeqi Lin, Qiang Fu, Bei Chen, Nanning Zheng, Weizhu Chen, and Jian-Guang Lou. Skill-based few-shot selection for in-context learning. *arXiv preprint arXiv:2305.14210*, 2023.
- [2] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [3] Jie Cao and Jing Xiao. An augmented benchmark dataset for geometric question answering through dual parallel text encoding. In *Proceedings of the 29th international conference on computational linguistics*, pages 1511–1520, 2022.
- [4] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language models? *arXiv preprint arXiv:2403.20330*, 2024.
- [5] Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu, and Wanxiang Che. M³ cot: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought. *arXiv preprint arXiv:2405.16473*, 2024.
- [6] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024.
- [7] Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*, 2024.
- [8] Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. A survey of chain of thought reasoning: Advances, frontiers and future. corr. abs/2309.15402, 2023. doi: 10.48550. *arXiv preprint ARXIV.2309.15402*, 2023.
- [9] Linger Deng, Yuliang Liu, Bohan Li, Dongliang Luo, Liang Wu, Chengquan Zhang, Pengyuan Lyu, Ziyang Zhang, Gang Zhang, Errui Ding, et al. R-cot: Reverse chain-of-thought problem generation for geometric reasoning in large multimodal models. *arXiv preprint arXiv:2410.17885*, 2024.
- [10] Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkang Yang, Winston Hu, Yongming Rao, and Ziwei Liu. Insight-v: Exploring long-chain visual reasoning with multimodal large language models. *arXiv preprint arXiv:2411.14432*, 2024.
- [11] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- [12] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [13] Valentin Hartmann, Anshuman Suri, Vincent Bindschaedler, David Evans, Shruti Tople, and Robert West. Sok: Memorization in general-purpose large language models. *arXiv preprint arXiv:2310.18362*, 2023.
- [14] Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*, 2022.
- [15] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025.
- [16] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- [17] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Hel-yar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- [18] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer, 2016.

- [19] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- [20] Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms. *arXiv preprint arXiv:2406.18629*, 2024.
- [21] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024.
- [22] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023.
- [23] Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. Multimodal arxiv: A dataset for improving scientific comprehension of large vision-language models. *arXiv preprint arXiv:2403.00231*, 2024.
- [24] Zhenwen Liang, Kehan Guo, Gang Liu, Taicheng Guo, Yujun Zhou, Tianyu Yang, Jiajun Jiao, Renjie Pi, Jipeng Zhang, and Xiangliang Zhang. Scemqa: A scientific college entrance level multimodal question answering benchmark. *arXiv preprint arXiv:2402.05138*, 2024.
- [25] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
- [26] Adam Dahlgren Lindström and Savitha Sam Abraham. Clevr-math: A dataset for compositional language, visual and mathematical reasoning. *arXiv preprint arXiv:2208.05358*, 2022.
- [27] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024.
- [28] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Lllavanext: Improved reasoning, ocr, and world knowledge, 2024.
- [29] Zuyan Liu, Yuhao Dong, Yongming Rao, Jie Zhou, and Jiwen Lu. Chain-of-spot: Interactive reasoning improves large vision-language models. *arXiv preprint arXiv:2403.12966*, 2024.
- [30] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023.
- [31] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022.
- [32] Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv preprint arXiv:2209.14610*, 2022.
- [33] Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. Ovis: Structural embedding alignment for multimodal large language model. *arXiv preprint arXiv:2405.20797*, 2024.
- [34] Ruilin Luo, Zhuofan Zheng, Yifan Wang, Yiyao Yu, Xinzhe Ni, Zicheng Lin, Jin Zeng, and Yujiu Yang. Ursa: Understanding and verifying chain-of-thought reasoning in multimodal mathematics. *arXiv preprint arXiv:2501.04686*, 2025.
- [35] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- [36] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706, 2022.
- [37] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021.

- [38] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *Advances in Neural Information Processing Systems*, 36:25081–25094, 2023.
- [39] Maxwell Nye, Michael Tessler, Josh Tenenbaum, and Brenden M Lake. Improving coherence and consistency in neural sequence models with dual-system, neuro-symbolic reasoning. *Advances in Neural Information Processing Systems*, 34:25192–25204, 2021.
- [40] Mengjie Ren, Boxi Cao, Hongyu Lin, Cao Liu, Xianpei Han, Ke Zeng, Guanglu Wan, Xunliang Cai, and Le Sun. Learning or self-aligning? rethinking instruction fine-tuning. *arXiv preprint arXiv:2402.18243*, 2024.
- [41] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *European conference on computer vision*, pages 146–162. Springer, 2022.
- [42] Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 37:8612–8642, 2025.
- [43] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- [44] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019.
- [45] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
- [46] Till Speicher, Bishwamittra Ghosh, Mohammad Aflah Khan, Qinyuan Wu, Vedant Nanda, Soumi Das, Krishna P Gummadi, and Evimaria Terzi. Rethinking memorization in llms: On learning by rote vs. with understanding.
- [47] Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183*, 2024.
- [48] Winnie Street, John Oliver Siy, Geoff Keeling, Adrien Baranes, Benjamin Barnett, Michael McKibben, Tatenda Kanyere, Alison Lentz, Robin IM Dunbar, et al. Llms achieve adult human performance on higher-order theory of mind tasks. *arXiv preprint arXiv:2405.18870*, 2024.
- [49] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- [50] Omkar Thawakar, Dinura Dissanayake, Ketan More, Ritesh Thawkar, Ahmed Heakl, Noor Ahsan, Yuhao Li, Mohammed Zumri, Jean Lahoud, Rao Muhammad Anwer, et al. Llamav-o1: Rethinking step-by-step visual reasoning in llms. *arXiv preprint arXiv:2501.06186*, 2025.
- [51] Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. *arXiv preprint arXiv:2402.14804*, 2024.
- [52] Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv preprint arXiv:2312.08935*, 2023.
- [53] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- [54] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2413–2427, 2017.
- [55] Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, et al. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*, 2024.

- [56] Xiaofei Wang, Jinhua Li, and Yifan Zhang. Improved value alignment in large language models using variational best-of-n techniques. 2024.
- [57] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [58] Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. Llava-o1: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*, 2024.
- [59] Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal reasoning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025.
- [60] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024.
- [61] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567, 2024.
- [62] Jingyi Zhang, Jiaying Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv preprint arXiv:2503.12937*, 2025.
- [63] Pan Zhang, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Rui Qian, Lin Chen, Qipeng Guo, Haodong Duan, Bin Wang, Linke Ouyang, et al. Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output. *arXiv preprint arXiv:2407.03320*, 2024.
- [64] Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Yu Qiao, et al. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? In *European Conference on Computer Vision*, pages 169–186. Springer, 2024.
- [65] Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Ziyu Guo, Shicheng Li, Yichi Zhang, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Bin Wei, et al. Mavis: Mathematical visual instruction tuning with an automatic data engine. *arXiv preprint arXiv:2407.08739*, 2024.
- [66] Ruohong Zhang, Bowen Zhang, Yanghao Li, Haotian Zhang, Zhiqing Sun, Zhe Gan, Yinfei Yang, Ruoming Pang, and Yiming Yang. Improve vision language model chain-of-thought reasoning. *arXiv preprint arXiv:2410.16198*, 2024.
- [67] Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical reasoning. *arXiv preprint arXiv:2501.07301*, 2025.
- [68] Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*, 2022.
- [69] Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, Yanjun Lyu, Peng Shu, Xiaowei Yu, et al. Evaluation of openai o1: Opportunities and challenges of agi. *arXiv preprint arXiv:2409.18486*, 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Please refer to the introduction and abstract.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the limitations in the Appendix F.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[NA\]](#)

Justification: This paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: Please refer to Section 3 and Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have made the code and dataset publicly available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please refer to Section 4 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Please refer to Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Please refer to Section 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research is conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please refer to the Appendix G.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper currently does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

```

<|reasoning_start|>
<|reasoning_step_start|>
<|reasoning_step_name_start|>Name1<|reasoning_step_name_end|>
<|reasoning_step_thought_start|>Thought1<|reasoning_step_thought_end|>
<|reasoning_step_reflection_start|>Reflection1<|reasoning_step_reflection_end|>
<|reasoning_step_end|>
<|reasoning_proceed|>
<|reasoning_step_start|>
<|reasoning_step_name_start|>Name2<|reasoning_step_name_end|>
<|reasoning_step_thought_start|>Thought2<|reasoning_step_thought_end|>
<|reasoning_step_reflection_start|>Reflection2 <|reasoning_step_reflection_end|>
<|reasoning_step_end|>
<|reasoning_proceed|>
...
<|reasoning_proceed|>
...
<|reasoning_end|>
<|answer_start|>The answer is A.<|answer_end|>

```

Figure 7: **Illustration of structured reasoning.** We take reasoning step (special tokens marked in blue) as the basic unit to conduct structured and fine-grained reasoning.

A Special Tokens

Our CoS uses special tokens to realize structured step-by-step reasoning. We list all 11 special tokens and their corresponding definitions in Table 6. For the initialization of these special tokens, we use the average value of the token embeddings of its corresponding definition. The reasoning structure is illustrated in Fig. 7.

B Prompts

In this section, we provide a demonstration of the prompts used for SFT and PRM data construction. The prompt used to create ShareGPT-Step-300K is illustrated in Fig. 8. The prompt used to utilize GPT-4o to label process-level annotation is depicted in Fig. 9.

C Training Recipes

In this section, we provide the specific hyper-parameter settings for the three stages in Table 7. All SFT and iterative DPO experiments are conducted on $8 \times$ NVIDIA-A800 GPUs by default. For PRM, it is trained on $16 \times$ NVIDIA-A800 GPUs. Additionally, we provide important parameters used in data construction. During the generation of positive and negative example pairs and the Math-Shepherd [52] Sampling period, we set the temperature to 1.0, n return sequences to 16, and top p to 0.95.

Table 6: **Special tokens we utilized and their definitions.**

Special token	Definition
< reasoning_start >	Start token of the reasoning process
< reasoning_end >	End token of the reasoning process
< reasoning_step_start >	Start token of a single reasoning step
< reasoning_step_end >	End token of a single reasoning step
< reasoning_step_name_start >	Start token of the reasoning step name
< reasoning_step_name_end >	End token of the reasoning step name
< reasoning_step_thought_start >	Start token of the reasoning thought content
< reasoning_step_thought_end >	End token of the reasoning thought content
< reasoning_step_reflection_start >	Start token of the reasoning reflection content
< reasoning_step_reflection_end >	End token of the reasoning reflection content
< reasoning_proceed >	Delimiter token between reasoning steps

Table 7: **Hyperparameter setting and training recipes.**

Hyper-parameter	SFT	PRM Training	DPO
Learning Rate	5e-6	1e-6	1e-6
Epoch	1	2	3 (iteration)
Warm-up Ratio	0.03	0.03	0.03
Weight Decay	0.05	0.05	0.05
Batch Size	128	128	256
Drop-path Rate	0.1	0.4	0.1
Update Parts	Projector, LLM	Projector, LLM	ViT, Projector, LLM
Data Size	300K	200K	20K

D Dataset Sources

To ensure diversity, our data is collected from various multimodal datasets, encompassing 17 tasks that require diverse multimodal reasoning capabilities. As shown in Table 8, these tasks can be categorized into four major types including mathematical reasoning [51, 64, 3, 26, 24], scientific reasoning [18, 31, 23, 5], chart&document analysis [35, 32, 37, 36], and world knowledge [41, 54, 22, 44].

Table 8: **SFT dataset sources.** To ensure diversity, we collect our dataset from a wide range of tasks that demands diverse reasoning capabilities.

Task Type	Dataset
Mathematics	GeoQA+ [3], MathVision [51], SceMQA [24], MathVerse [64], CLEVER-Math [26]
Chart	ChartQA [35], TabMWP [32]
Science	ArxivQA [23], M3CoT [5], ScienceQA [31], AI2D [18]
Document	DocVQA [37], InfoVQA [36]
Knowledge	A-OKVQA [41], FVQA [54], TextVQA [44], SEED [22]

E Attempt for Step-wise DPO

Due to our structured and fine-grained design, we can improve the reasoning ability of the model by conducting fine-grained step level beam search during reasoning. Naturally, we wonder whether we could construct similar fine-grained preference data during RL training. Specifically, we use PRM to find the optimal step and a bad one at each step, and then build a set of pair data. Then we sample the next step based on the optimal step, and cycle until we get the final answer. In this way, we build a pair of preference data at each step, which can provide step-wise fine-grained preference rewards. We term this method to be step-wise DPO. This is different from step-dpo [20], which only detect the first error step while we build pairs of every single step.

However, this method does not work as expected, but has a significant negative impact on model performances as shown in Table 9. We believe that the failure of this method lies in that the next sampling based on a good step is probably also a good step. Therefore, the distinction between the constructed chosen and reject pairs is relatively small. As a consequence, in order to avoid outputting negative steps during training, the model will also tend not to output chosen steps. This conjecture is validated in the rewards: the rewards of chosen and reject are both very small negative numbers. Therefore, we mainly conduct fine-grained step-by-step sampling in the inference time.

F Discussion of Table 3

We observed a counterintuitive phenomenon that the larger InternVL-2.5-38B does not achieve a higher Answer Accuracy compared to its smaller counterparts. Our speculation on this phenomenon is as follows:

Table 9: **Ablation on per-step-wise DPO.** The chosen and rejected steps are too similar to form effective preference pair, so the model will refuse to output both, backfiring performances.

Method	M3CoT	MathVista	Chosen Rewards	Reject Rewards
SFT	67.4	51.4	-	-
Outcome	70.0	53.5	-0.07	-0.16
Per-step-wise	64.8	50.7	-1.39	-3.86

Notably, evaluating intermediate reasoning steps is substantially more challenging than evaluating the final answer for three reasons: i) Higher semantic complexity: Step-level evaluation requires a deeper and more holistic understanding of the question. The model must assess whether each step contributes constructively toward the final answer. Similarly, annotating process data is also more cognitively demanding for human annotators. ii) More supervision signals. As shown in Equation 1, PRM is trained to minimize the sum of cross-entropy losses over all intermediate steps and the final answer. A complete reasoning trace contains multiple steps but only a single answer, thus accurately assessing all the intermediate steps is much harder than only assessing the answer. iii) Answer memorization. As highlighted in recent studies [13, 46], LLMs often exhibit "answer hacking"—producing correct answers via memorized question or patterns rather than genuine understanding. In contrast, intermediate steps exhibit higher variance across instances and are less amenable to memorization.

As we know, neural networks are inherently short-cut learners[11]. During training, smaller models (i.e., 8B PRM) tend to prioritize optimizing for final answer accuracy, as it offers a more accessible gradient path to loss reduction. Improving step-level predictions poses a greater optimization challenge. The 38B PRM, with its higher capacity, is better able to model and optimize intermediate reasoning quality. Consequently, it achieves greater improvements in step-level accuracy, which more effectively reduces the total loss—resulting in the pattern observed in Table 3.

G Limitations

We generated process-labeled data through Monte Carlo estimation and LLM-as-Judge and verified the effectiveness of our method through extensive experiments. But it cannot be denied that the process data obtained by these two methods still cannot be guaranteed to be 100% correct. Even the process data labeled by humans cannot guarantee accuracy and quality [25], and we do not use human annotation due to the high cost and inability to scale. In the future, if there is a way to produce higher-quality process labeled data, our method can hopefully shine further.

H Broader Impacts

This paper presents work whose goal is to advance the field of Vision-language Models and Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

You are a thought assistant, i will provide you with an image, a question, and a reference answer. Your responsibility is to develop a sequence of intermediate thinking steps that would logically connect the question to its answer.

Each step must consist of and only consist of the following three parts:

1. **Step Name:** A brief title summarizing the main idea of the step.
2. **Detailed Thinking:** An in-depth description of the thought process in this step, explaining what is being observed, inferred, or analyzed. Note that the process of thinking and reasoning should be presented in as much detail as possible.
3. **Reflection:** A reflection on this step's connection to prior steps, pointing out any potential errors, and if possible, citing evidence in the image that supports or informs this step.

There are some guidelines you should follow:

1. Make sure the steps are clear, sequential, and detailed enough to make each step flows logically into the next.
2. The number of steps must be no less than 5 and no more than 10.
3. Do not state an answer at the beginning. Explain the visual clues that help to derive the answer.
4. Don't state that the reference answer is correct or consistent to your finding. You are writing your own solution.
5. The explanation can help you derive the thinking process.
6. For math problems, you must give specific operation steps and concrete intermediate solution leading to the final answer.
7. Make sure to incorporate concrete observations from the provided materials.

The output format must be strictly organized as follows:

```
### <Step 1: xxx>
<Detailed Thinking>: xxx
<Reflection>: xxx
```

```
### <Step 2: xxx>
<Detailed Thinking>: xxx
<Reflection>: xxx
```

...

```
### <Step x: xxx>
<Detailed Thinking>: xxx
<Reflection>: xxx
```

...

```
### <Answer>
xxx
```

Figure 8: Prompt used for generating step-by-step reasoning data.

You are a judging assistant specializing in evaluating process reward, the following is an image, a question, a reference answer and a step-by-step intermediate answer process for the question. The reference answer is correct while the intermediate answer steps are not necessarily correct.

Your task is to review and critique each step of the intermediate answer process based on its correctness. Assign one of the following ratings to each step:

- **Good:** The step is logical, directly relevant to the question, and most importantly, correct. This step provides substantive progress to solve the question.
- **Neutral:** The step is somewhat relevant but contains minor flaws or gaps in reasoning. Or this step only includes some abstract practices rather than concrete solutions, which is not helpful for solving the question in effect.
- **Bad:** The step is flawed, wrong, or shows a misunderstanding of the question or context. This step has a negative effect on solving the question.

For each step, provide a brief explanation of your evaluation, identifying areas for improvement if Neutral or Bad. Finally, return your output in the form of a step label list. 1 stands for Good; 0 stands for Neutral; -1 stands for Bad.

The final step label list format must be strictly organized as follows:

```
### Step Label: [1,1,0,-1,-1,-1]
```

Figure 9: Prompt used for generating process-level annotation.