

# AStar: Boosting Multimodal Reasoning with Automated Structured Thinking

Jinyang Wu<sup>1</sup>, Mingkuan Feng<sup>1</sup>, Guocheng Zhai<sup>1</sup>, Shuai Zhang<sup>1,\*</sup>, Zheng Lian<sup>2</sup>, Fangrui Lv<sup>1</sup>  
Pengpeng Shao<sup>1</sup>, Ruihan Jin<sup>1</sup>, Zhengqi Wen<sup>1</sup>, Jianhua Tao<sup>1,\*</sup>

<sup>1</sup>Department of Automation, BNRist, Tsinghua University

<sup>2</sup>Institute of Automation, Chinese Academy of Sciences  
wu-jy23@mails.tsinghua.edu.cn

## Abstract

Multimodal large language models excel across diverse domains but struggle with complex visual reasoning tasks. To enhance their reasoning capabilities, current approaches typically rely on explicit search or post-training techniques. However, search-based methods suffer from computational inefficiency due to extensive solution space exploration, while post-training methods demand substantial data, computational resources, and often exhibit training instability. To address these challenges, we propose **AStar**, a training-free, **A**utomatic **S**tructured thinking paradigm for multimodal reasoning. Specifically, we introduce novel “thought cards”, a lightweight library of high-level reasoning patterns abstracted from prior samples. For each test problem, AStar adaptively retrieves the optimal thought cards and seamlessly integrates these external explicit guidelines with the model’s internal implicit reasoning capabilities. Compared to previous methods, AStar eliminates computationally expensive explicit search and avoids additional complex post-training processes, enabling a more efficient reasoning approach. Extensive experiments demonstrate that our framework achieves 53.9% accuracy on MathVerse (surpassing GPT-4o’s 50.2%) and 32.7% on MathVision (outperforming GPT-4o’s 30.4%). Further analysis reveals the remarkable transferability of our method: thought cards generated from mathematical reasoning can also be applied to other reasoning tasks, even benefiting general visual perception and understanding. AStar serves as a plug-and-play test-time inference method, compatible with other post-training techniques, providing an important complement to existing multimodal reasoning approaches.

## 1 Introduction

Multimodal Large Language Models (MLLMs) have demonstrated impressive capabilities across diverse tasks and domains (Wang et al. 2024d; OpenAI 2024; Qiao et al. 2024), yet they struggle with complex visual reasoning tasks that require processing multimodal information with sophisticated problem-solving strategies (Zhang et al. 2024; Wang et al. 2024a). Recent advances in System 2 slow-thinking reasoning models like OpenAI-o1 (OpenAI 2024), DeepSeek-R1 (Guo et al. 2025), and Kimi-K1.5 (Team et al. 2025) have inspired growing interest in incorporating structured long Chain-of-Thought (CoT) (Wei et al.

\*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

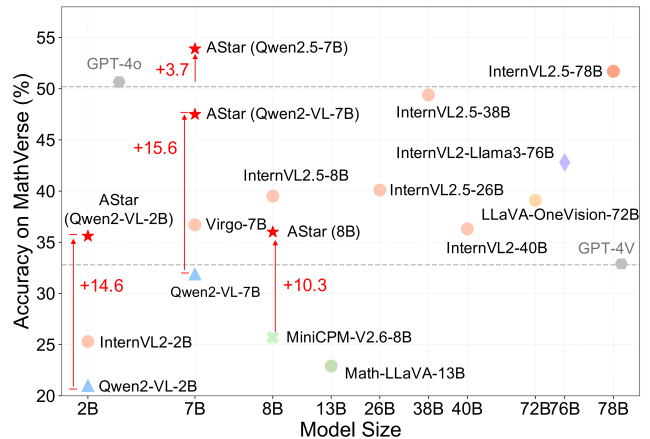


Figure 1: Evaluation results on MathVerse. AStar makes 7B models competent problem-solvers, surpassing GPT-4o. Our approach demonstrates consistent effectiveness across multiple model architectures and scales.

2022) thinking into MLLMs (Xu et al. 2024; Dong et al. 2024b). These approaches address the limitations of conventional MLLMs that often rely on simple “direct prediction” modes due to the scarcity of high-quality long-chain reasoning data (Xu et al. 2024; Luo et al. 2025). Current methods can be divided into two primary categories: (i) *explicit search methods* (Dong et al. 2024a; Yao et al. 2024): leverage algorithms like beam search or Monte Carlo Tree Search (MCTS) with specialized reward models to guide solution path exploration; (ii) *post-training methods* (Wang et al. 2024d, 2025): develop structured long CoT reasoning capabilities through common techniques such as Supervised Fine-Tuning (SFT) (Zhang et al. 2024; Luo et al. 2025) or Reinforcement Learning (RL) like Proximal Policy Optimization (PPO) (Schulman et al. 2017) and Group Relative Policy Optimization (GRPO) (Guo et al. 2025).

However, both methods face some limitations. (i) Search-based methods (Dong et al. 2024a) suffer from computational inefficiency due to extensive exploration across solution spaces; (ii) SFT-based post-training methods (Luo et al. 2025; Yao et al. 2024; Hu et al. 2024) typically require substantial training data ( $\geq 100K$ ) and computational resources

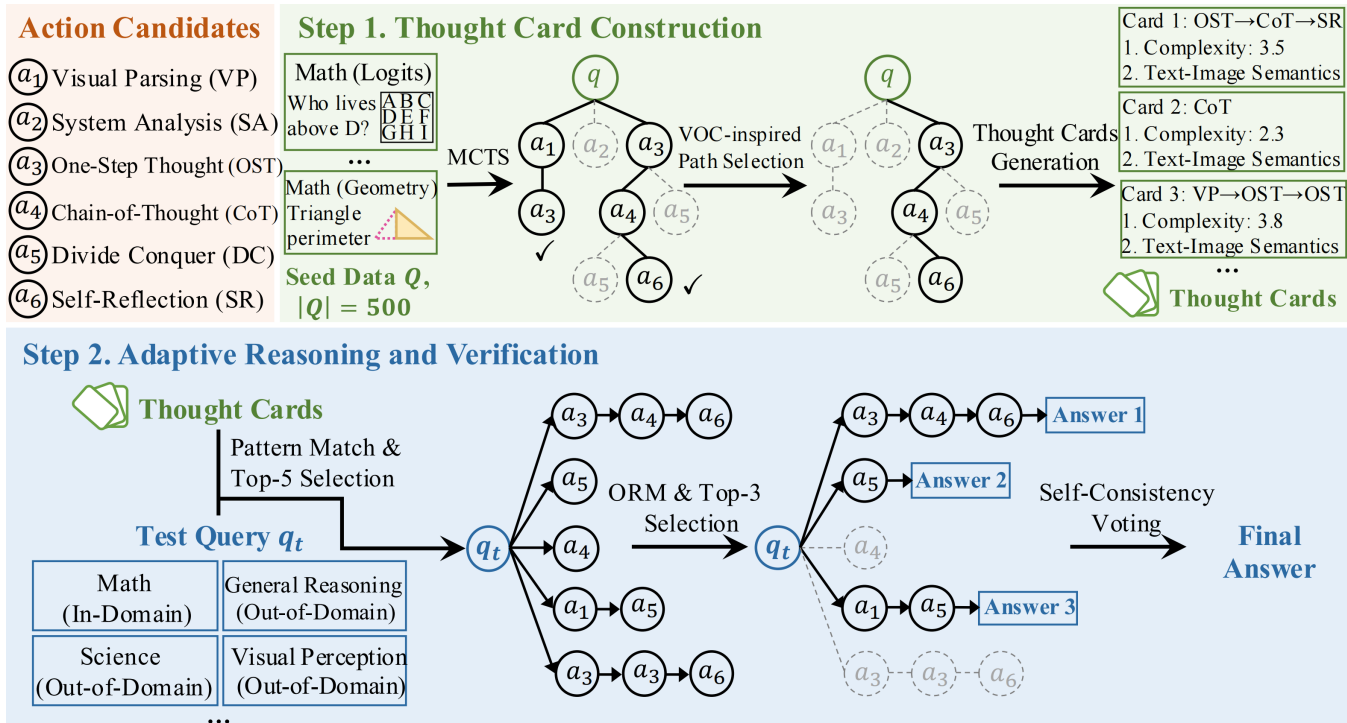


Figure 2: Flowchart of AStar. Building on the action candidates, this framework consists of two key components: (1) Thought Card Construction; and (2) Adaptive Reasoning and Verification.

to implicitly extract reasoning patterns, leading to inefficient data utilization. Their reliance on proprietary models like GPT-4o for data synthesis also creates accessibility barriers for researchers outside major enterprises; (iii) RL-based post-training methods (Peng et al. 2025; Liu et al. 2025; Zhang et al. 2025a) encounter challenges with exploration depth and training stability. As shown in Yue et al. (2025), these methods primarily function by biasing the model’s output distribution toward reward-maximizing paths without introducing external knowledge. This inherently limits their exploration capacity, resulting in a narrower reasoning capability boundary compared to base models.

To address these limitations, we propose **AStar**, an automated structured thinking paradigm for multimodal reasoning. Specifically, we introduce a novel module called “thought cards”, a lightweight, generalizable library that stores abstract, high-level reasoning patterns. These patterns are derived from small-scale prior samples using MCTS and can be effectively generalized across different domains. For each test problem, we design an adaptive retrieval mechanism to select the five most suitable thought cards based on the problem’s characteristics. Leveraging these guidelines, AStar performs visual reasoning and derives the final solution through unified self-consistency checks and outcome reward scoring. Compared to existing methods, AStar eliminates the computational inefficiency of search-based approaches and removes the need for large-scale training data required by post-training methods, enabling a data- and

computation-efficient reasoning framework.

Our method offers four key advantages. (i) **Performance Enhancement**: By leveraging high-level thought cards, we adaptively obtain reasoning guidelines based on problem attributes, achieving remarkable performance improvement (Figure 1); (ii) **Efficiency Improvement**: Our training-free framework achieves comparable performance to powerful baselines with only small-scale prior samples, enhancing both data and computational efficiency; (iii) **Flexibility**: AStar serves as a plug-and-play test-time inference method that seamlessly integrates with other post-training techniques like SFT and PPO; (iv) **Transferability**: The thought cards generated from mathematical domains exhibit strong transferability to other reasoning tasks and also benefit general visual perception and understanding. We summarize our core contributions as follows:

- We propose AStar, a training-free reasoning framework that integrates MLLMs’ internal implicit reasoning capabilities with external explicit guidelines. To achieve this, we introduce “thought cards”, a lightweight mechanism for storing high-level reasoning patterns, which can be adaptively instantiated for specific test problems.
- Extensive experiments demonstrate AStar’s effectiveness and transferability. AStar with a 7B backbone outperforms large-scale models like GPT-4o across various benchmarks. Notably, thought cards from mathematical reasoning exhibit remarkable generalization to scientific reasoning, visual perception, and understanding tasks.

- AStar is a plug-and-play solution that can be integrated with popular post-training techniques. It serves as a valuable complement to existing multimodal reasoning methods, offering enhanced flexibility and scalability.

## 2 Methodology

**Overview of AStar.** This section introduces our proposed method AStar in detail. As shown in Figure 2 and algorithm 1, our approach consists of two key steps:

- *Thought Card Construction:* Leverage MCTS to systematically construct thought cards, which serve as reference insights and effectively guide subsequent problem-solving during inference.
- *Adaptive Reasoning and Verification:* Dynamically select and execute optimal reasoning patterns based on problem complexity, followed by robust solution verification.

**Visual Reasoning Action Definition.** Understanding human complex reasoning is crucial for modeling cognitive processes (Jaffe et al. 2023). Existing studies distinguish between two cognitive systems: System 1 and System 2 (Da Silva 2023). While System 1 represents fast, intuitive, yet error-prone thinking, System 2 involves slow, deliberative thinking with superior performance. With the emergence of advanced models like OpenAI o1, developing efficient System 2 to emulate human cognitive processes has gained significant research attention (Yao et al. 2024; Thawakar et al. 2025; Yang et al. 2024). Inspired by this, we introduce six vision-language reasoning actions to bridge the gap between model reasoning and human cognition: *Visual Parsing (VP,  $a_1$ )*, *System Analysis (SA,  $a_2$ )*, *One-Step Thought (OST,  $a_3$ )*, *Chain-of-Thought (CoT,  $a_4$ )*, *Divide and Conquer (DC,  $a_5$ )*, *Self-Reflection (SR,  $a_6$ )*. Detailed descriptions and rationale are provided in Appendix C.1.

### 2.1 Thought Card Construction

Based on the above action definitions, we introduce “*thought cards*” as reasoning templates to guide inference in Section 2.2. Using only 500 prior samples, we derive their reasoning paths (Phase 1) and then distill them into high-level thought cards (Phase 2). These cards provide structured guidance for efficient problem-adaptive reasoning during inference.

**Phase 1: Acquire reasoning paths for seed data.** As shown in Figure 2, we employ MCTS to iteratively optimize the solution search process, generating high-quality reasoning paths for the seed dataset. This design leverages MCTS’s systematic exploration and MLLMs’ inherent reasoning capabilities (Yin et al. 2023; Wang et al. 2024d). We formulate each multimodal reasoning problem  $q \in \mathbb{Q}$  (consisting of input questions and images) as a tree search problem, where  $q$  represents the root node and subsequent nodes denote reasoning steps (actions and corresponding outcomes) generated by a policy MLLM  $\pi_\theta$ . We define the state  $S_{t-1}$  as the trajectory  $q, s_1, \dots, s_{t-1}$ , where  $S_0 = q$ . The next step is sampled as  $s_t \sim \pi_\theta(S_{t-1})$ . To guide tree expansion, we define  $Q(s)$  as the reward value for node  $s$ . Initially, all unexplored nodes are assigned  $Q(s_i) = 0$ . They are updated

---

### Algorithm 1: AStar Algorithm

---

**Input:** a policy model  $\pi_\theta$ ; a multimodal test query  $q_t$ ; a set of seed data  $\mathbb{Q}$

**Output:** the optimal reasoning trajectory  $y_t$

```

1: Initialize action space  $A = \{a_1, a_2, a_3, a_4, a_5, a_6\}$ 
2: Initialize repository  $D \leftarrow []$ ; Cards  $\mathbb{C} \leftarrow \{\}$ 
3: // 2.1. Thought Card Construction
4: for  $q \in \mathbb{Q}$  do
5:   Acquire reasoning paths  $\{q, P\} \leftarrow \text{MCTS}(\pi_\theta; q)$ 
6:   if found at least one valid reasoning path then
7:     Find  $p_{\text{best}}$  from  $P$  (Equation (3))
8:     Add  $\{q, p_{\text{best}}\}$  into  $D$ 
9:     Update  $\mathbb{C}$  from  $D$ 
10:  end if
11: end for
12: // 2.2. Adaptive Reasoning and Verification
13:  $\mathbb{C}_{q_t} \leftarrow \text{Card\_Match}(\mathbb{C}; q_t)$ 
14:  $y_t \leftarrow \text{Reason\_And\_Verify}(\pi_\theta; q_t; \mathbb{C}_{q_t})$ 

```

---

using a weighted average between the parent’s current value and its child node’s value:

$$Q(p) \leftarrow (1 - \alpha)Q(p) + \alpha Q(s) \quad (1)$$

where  $\alpha$  is a discount factor for future rewards. For terminal nodes, we use the likelihood of self-consistency majority vote as reward value, enabling supervision-free generalization. Specifically, this phase comprises 4 MCTS operations:

(1) *Selection.* This operation identifies promising nodes for expansion. Starting from the root node, we iteratively select child nodes using the Upper Confidence Bounds applied to Trees (UCT) (Kocsis and Szepesvári 2006) until reaching a leaf node:

$$UCT(s) = Q(s) + w \sqrt{\frac{\ln N(p)}{N(s)}} \quad (2)$$

where  $Q(s)$  is the reward value for node  $s$ ,  $N(s)$  is the visit count,  $p$  is the parent node, and  $w$  is the exploration weight. The node with the highest UCT value is selected for subsequent phases, balancing exploration and exploitation.

(2) *Expansion.* The selected node  $s$  is expanded by sampling  $n$  actions from  $\pi_\theta$  and generating corresponding reasoning outcomes. These  $n$  child nodes are added to the tree and stored in an external memory structure.

(3) *Simulation.* Starting from the selected node, we iteratively sample and expand nodes until reaching a terminal state (maximum depth or final answer node).

(4) *Backpropagation.* Upon simulation completion, node information is updated along the simulation path  $s_0, \dots, s_d$ . Visit counts are incremented ( $N(s) \leftarrow N(s) + 1$ ), and node value  $Q(s)$  is propagated backward to its parent node  $p$  using Equation (1). These updated values are used to guide subsequent UCT-based node selection.

**Phase 2: Distill paths into thought cards.** After executing MCTS, we obtain a tree structure for each seed dataset question, yielding multiple valid reasoning paths that constitute the path set  $P$ . Inspired by the concept of Value of Computation (VOC) (Russell and Wefald 1991), which optimizes

the trade-off between computational benefits and costs, we propose a VOC-inspired selection metric to identify the optimal reasoning trajectory from candidate solutions:

$$Score(q, p_q) = k \cdot R(p_q|q) - (1 - k) \cdot C(p_q) \quad (3)$$

where  $q$  is the task input,  $p_q \in P$  denotes a candidate reasoning path, and  $k$  balances benefits against computational costs. Here,  $R(p_q|q)$  denotes the path’s final reward (defined as the leaf node’s Q-value), while  $C(p_q)$  is the reasoning cost (defined as the number of actions in the sequence).

Then, for each question in the seed dataset, we select the path  $p_{best}$  with the highest  $Score(q, p_q)$  to build a *Question-path repository*  $D$  with one-to-one mappings. Inspired by metareasoning principles (Russell and Wefald 1991), which advocate for adaptive reasoning strategies, we distill these question-path pairs into abstract thought cards  $\mathbb{C}$  as high-level guidelines. Each card is distilled based on two attributes: problem complexity (PC) (Lee and Heyworth 2000) and text-image semantics (TIS) (Radford et al. 2021). PC represents prior known conditions derived from input image-text pairs  $(i, t)$ , obtained using a lightweight 2B MLLM. TIS refers to joint semantic representations based on the CLIP model encodings of  $(i, t)$ :

$$E_q(i, t) = \frac{E_I(i) + E_T(t)}{2} \quad (4)$$

Finally, each thought card contains a high-level thought template (e.g.,  $a_1 \rightarrow a_2 \rightarrow a_4$ ), along with the average problem complexity and text-image semantics of multiple questions sharing this template. We provide a detailed example of thought cards in Appendix C.

## 2.2 Adaptive Reasoning and Verification

During inference, given a multimodal test query  $q_t$ , we compute its PC and TIS, and perform nearest neighbor matching against pre-constructed thought cards  $\mathbb{C}$  to identify the five most relevant cards that best align with its complexity level and semantics. The selection process involves ranking cards based on similarity for each attribute independently:

$$R_{TIS}(c) = \text{Rank} \left( E_{q_t}(i_t, t_t)^\top E_q(i_c, t_c) \right), \forall c \in \mathbb{C} \quad (5)$$

$$R_{PC}(c) = \text{Rank} \left( \frac{1}{|PC_{q_t} - PC_c|} \right), \forall c \in \mathbb{C} \quad (6)$$

where  $\text{Rank}(\cdot)$  assigns rankings in descending order, with higher values receiving better rankings (e.g., the highest value is ranked 1, the lowest value is ranked  $|\mathbb{C}|$ ). We then combine these rankings to compute a total ranking score:

$$R(q_t, c) = R_{TIS}(c) + R_{PC}(c) \quad (7)$$

Finally, we select the five thought cards with better combined rankings:

$$NN_5(q_t, \mathbb{C}) = \arg \min_{\mathbb{C}_{q_t} \subseteq \mathbb{C}, |\mathbb{C}_{q_t}|=5} \sum_{c \in \mathbb{C}_{q_t}} R(q_t, c) \quad (8)$$

where  $\mathbb{C}_{q_t} \subseteq \mathbb{C}$  contains the five thought cards with the best combined rankings. We instantiate these templates for the test query to obtain five candidate solutions. To identify the

best reasoning trajectory, we employ both self-consistency checks and text-domain outcome reward models due to the scarcity of visual-domain verification models. Details are described in Appendix B.4.

In summary, our approach adaptively selects high-level reasoning guidelines based on problem attributes, seamlessly integrating the model’s internal implicit reasoning with external explicit guidance. This dynamic selection mechanism enhances flexibility and efficiency while maintaining robust performance across diverse problem types.

## 3 Experiments

In this section, we first describe our experimental setup, and then evaluate the effectiveness of AStar across four key aspects: (1) performance enhancement, (2) efficiency improvement, (3) flexibility, and (4) transferability. Finally, we conduct ablation studies and analyze impact of seed dataset size.

### 3.1 Experimental Setup

**Datasets.** We conduct extensive experiments across various 4 domains and 8 datasets: (1) mathematical reasoning: MathVista (Lu et al. 2023), MathVerse (Zhang et al. 2025b), and MathVision (Wang et al. 2024a); (2) general reasoning: MMMU (Yue et al. 2024); (3) domain-specific scientific reasoning: GAOKAO-MM (Zong and Qiu 2024); (4) visual perception and understanding: ChartQA (Masry et al. 2022), MMStar (Chen et al. 2024a), and BLINK (Fu et al. 2025). For all benchmarks, we use the official evaluation metrics.

**Models.** To demonstrate AStar’s versatility, we evaluate its effectiveness on both LLM and MLLM, including Qwen2.5-7B (Qwen Team 2024), and Qwen2-VL-2/7B (Bai et al. 2025). This design aims to validate that AStar can seamlessly leverage pre-trained LLMs and MLLMs as its inference backbone without modifications.

**Baselines.** We evaluate AStar against four representative baseline categories: (1) open-source general MLLMs like Qwen2-VL (Wang et al. 2024b) and InternVL2 series (Chen et al. 2024b); (2) open-source reasoning MLLMs, including SFT-based methods URSA (Luo et al. 2025) and MathLLaVA (Shi et al. 2024), and RL-based methods R1-VL-7B (Zhang et al. 2025a), LMM-R1 (Peng et al. 2025), and MM-Eureka (Meng et al. 2025); (3) search-based methods like Mulberry (Yao et al. 2024); and (4) closed-source MLLMs like GPT-4o (OpenAI 2024).

### 3.2 Performance Enhancement

**Main Results.** Table 1 presents the performance of AStar across three representative multimodal reasoning benchmarks. We have four key findings:

- AStar consistently outperforms both general and math-specialized MLLMs. With the Qwen2.5-7B reasoning backbone, AStar achieves 53.9% accuracy on MathVerse, surpassing large-scale CoT-trained URSA-8B by 8.2%, and expensive GRPO-trained R1-VL-7B by 13.9%.

- AStar demonstrates strong performance across diverse reasoning types. It reaches 59.5% on logical reasoning (LOG), outperforming the powerful URSA-8B by 37.9%.

Model	MathVerse						MathVista					MathVision	
	ALL	VI	VD	VO	TD	TO	ALL	ARI	LOG	STA	VQA	ALL	LOG
Random	12.4	12.4	12.4	12.4	12.4	12.4	17.9	13.8	13.4	14.3	26.3	7.2	7.6
Human	64.9	61.4	68.3	66.7	71.2	41.7	60.3	59.2	40.7	63.9	55.9	68.8	61.3
GPT-4o	50.2	39.6	42.5	39.3	54.7	50.0	60.1	58.4	27.0	69.0	49.2	30.4	29.4
<i>Open-Source General MLLMs</i>													
MiniGPT4-7B (Zhu et al. 2023)	12.2	12.5	14.8	8.7	12.3	13.4	23.1	32.0	10.8	17.9	30.2	10.8	12.7
LLaVA-1.5-13B (Liu et al. 2024)	12.7	12.6	12.7	9.0	17.1	22.6	27.7	28.6	10.8	22.9	30.2	11.1	13.5
SPHINX-8x7B (Lin et al. 2023)	22.8	21.1	19.6	18.3	33.3	23.1	42.6	43.0	14.4	50.8	43.3	15.8	17.9
InternVL2-8B (Chen et al. 2024b)	35.9	32.2	30.9	27.7	39.0	36.0	58.3	56.4	10.8	68.8	49.7	18.4	15.3
Qwen2-VL-7B (Wang et al. 2024b)	33.6	31.3	30.3	28.1	37.4	35.0	58.9	57.5	24.3	43.1	58.1	17.2	12.7
<i>Open-Source Reasoning MLLMs (Large-Scale Post-Training)</i>													
Math-LLaVA-13B (Shi et al. 2024)	22.9	24.5	21.7	16.1	27.3	27.0	46.6	40.7	23.3	42.3	33.5	15.7	16.0
Math-PUMA-7B (Zhuang et al. 2024)	33.6	33.4	31.6	26.0	42.1	39.8	47.9	46.2	21.6	55.8	30.2	14.0	13.7
MultiMath-7B (Peng et al. 2024)	27.7	28.1	25.9	15.0	34.8	35.3	50.0	42.2	23.3	64.9	49.2	16.3	17.9
URSA-8B (Luo et al. 2025)	45.7	46.4	43.9	28.6	55.3	51.8	59.8	53.5	21.6	57.1	40.2	26.2	24.8
R1-VL-7B (Zhang et al. 2025a)	40.0	37.3	33.6	39.8	45.0	40.7	63.5	54.7	28.5	61.2	<b>60.9</b>	27.1	23.6
AStar (Qwen2.5-7B, Training-free)	<b>53.9</b>	<b>49.7</b>	<b>64.4</b>	<b>48.6</b>	<b>56.4</b>	<b>56.1</b>	<b>64.2</b>	<b>63.8</b>	<b>59.5</b>	<b>69.1</b>	<b>60.9</b>	<b>32.7</b>	<b>39.4</b>
AStar (Qwen2-VL-7B, Training-free)	47.5	41.8	51.6	<b>48.6</b>	49.3	42.2	61.7	60.6	28.6	68.4	59.6	27.9	29.4

Table 1: Performance comparison on MathVista, MathVerse, and MathVision. The best model results are highlighted in **bold**. For MathVerse, we show 6 categories: ALL (overall accuracy), VI (vision intensive), VD (vision dominant), VO (vision only), TD (text dominant), and TO (text only). For MathVista, we present 5 categories: ALL (overall accuracy), ARI (arithmetic reasoning), LOG (logical reasoning), STA (statistical reasoning), and VQA (visual question answering). For MathVision, we present 2 categories: ALL (overall accuracy) and LOG (logical reasoning).

Similar gains are observed in statistical reasoning (STA: 69.1%) and visual question answering (VQA: 60.9%).

- AStar’s adaptive reasoning benefits are universal across varying multimodal information distributions. It maintains consistent gains from vision-dominant (VD: 64.4%) to text-dominant (TD: 56.4%) scenarios, showcasing robust performance regardless of modality balance. Its strong performance in text-only (TO) scenarios highlights the paradigm’s versatility and ability to generalize to pure text domains.

- On the more challenging MathVision benchmark (Wang et al. 2024a), AStar achieves 32.7% overall accuracy, surpassing GPT-4o (30.4%). Notably, in logical reasoning, AStar attains 39.4% accuracy, significantly outperforming GPT-4o (29.4%) by 10.0%. This stems from our adaptive decomposition framework that transforms complex reasoning into manageable sub-problems.

**High Performance with Efficient Models.** As shown in Figure 1, we intuitively demonstrate AStar’s effectiveness across different model architectures and scales. Notably, AStar (Qwen2.5-7B) achieves 53.9% accuracy on MathVerse, surpassing GPT-4o’s 50.2%. Even smaller models like AStar+Qwen2-VL-2B outperform larger baselines like Qwen2-VL-7B and InternVL2-40B. This demonstrates AStar’s capacity to boost smaller models to competitive performance against significantly larger architectures. Detailed results are provided in Appendix E.2.

### 3.3 Efficiency Improvement

To demonstrate AStar’s efficiency, we compare it against search-based (Mulberry) and post-training methods (URSA, LMM-R1, MM-Eureka, R1-VL). Table 2 evaluates resource requirements, accessibility, and performance on visual understanding (MMStar) and mathematical reasoning (MathVerse, MathVision) benchmarks.

AStar achieves superior performance with exceptional efficiency. Our approach requires only 0.5K prior samples and 50 minutes of preprocessing time. Our approach requires only 0.5K prior samples and 50 minutes for thought card generation (see Figure 2), eliminating the need for any training process or model parameter updates. This results in a 520-fold and 2200-fold reduction in data requirements compared to Mulberry (260K) and URSA-8B (1100K), while eliminating URSA-8B’s 3-day training requirement. Despite these significant efficiency gains, AStar consistently outperforms baselines: MMStar (62.3% vs. 61.3% for Mulberry), MathVerse (53.9% vs. 45.7% for URSA-8B), and MathVision (32.7% vs. 26.9% for MM-Eureka-7B). This stems from our explicit reasoning pattern extraction that captures problem-solving strategies without extensive computational overhead. We think our framework may provide insights for resource-constrained researchers.

### 3.4 Flexibility

As a test-time inference framework, AStar can be effectively integrated with post-training methods. We demonstrate this

Methods	Type	OS Only	Training-Free	Prior Data ↓	Pre. Time ↓	MMStar ↑	MathVerse ↑	MathVision ↑
Mulberry-7B	Search	✗	✗	260K	-	61.3	44.9	26.4
URSA-8B	SFT	✗	✗	1100K	3 days	55.4	45.7	26.2
LMM-R1-3B	PPO	✓	✗	55.3K	-	58.0	41.8	26.9
R1-VL-7B	GRPO	✓	✗	260K	-	60.0	40.0	27.1
MM-Eureka-7B	GRPO	✓	✗	15K	-	59.4	50.3	26.9
<b>Ours (Qwen2.5-7B)</b>	-	✓	✓	<b>0.5K</b>	<b>50 mins</b>	<b>62.3</b>	<b>53.9</b>	<b>32.7</b>
<b>Ours (Qwen2-VL-7B)</b>	-	✓	✓	<b>0.5K</b>	<b>50 mins</b>	<b>62.0</b>	<b>47.5</b>	<b>27.9</b>

Table 2: Efficiency comparison. ‘OS Only’ indicates exclusive use of open-source models. ‘Pre. Time’ denotes preprocessing or training time, specifically thought card construction time for our method and dataset construction/training time for others.

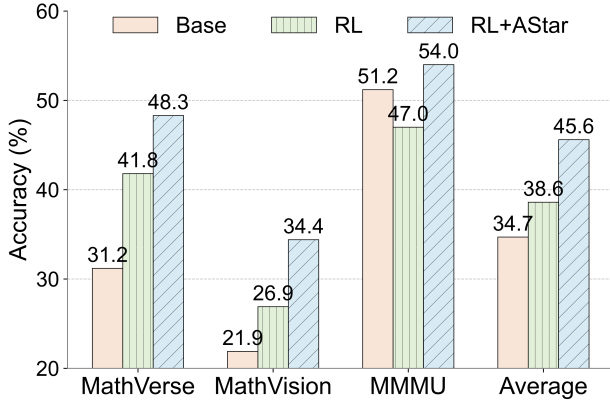


Figure 3: Flexibility verification. AStar is a plug-and-play framework that can be integrated with RL-trained models (from Qwen2.5-VL-3B) for further improvement.

by applying our reasoning paradigm to PPO-trained LMM-R1 (Peng et al. 2025) from Qwen2.5-VL-3B. Figure 3 shows that AStar consistently enhances performance across all benchmarks when combined with post-trained models. The combination (RL+AStar) achieves 48.3% accuracy on MathVerse, improving upon the base RL model (41.8%) by 6.5%. This synergy indicates AStar captures complementary reasoning patterns beyond post-training, highlighting its plug-and-play adaptability. Additional integration with SFT and GRPO are provided in the Appendix E.5.

### 3.5 Transferability

Recent work has highlighted that distributional shifts severely affect MLLM reliability (Yin et al. 2023). While these models excel on in-domain tasks, their performance often degrades in out-of-domain (OOD) scenarios. This challenge is compounded by the difficulty of acquiring sufficient domain-specific training data and necessary computational resources.

To evaluate AStar’s cross-domain transferability, we investigate whether thought cards constructed from mathematical domains can generalize to non-mathematical tasks. Since our thought cards derive from math domains (details in Appendix D.3), we conduct OOD evaluations on general reasoning (MMMU), domain-specific reasoning (GAOKAO), visual perception (MMStar, BLINK), and chart

Model Setting	MathVision	MathVerse	Average
AStar (Full)	<b>32.7</b>	<b>53.9</b>	<b>43.3</b>
w/o Thought Cards	24.8	42.9	33.8 <sup>-9.5↓</sup>
w/o Card Matching	27.9	47.2	37.6 <sup>-5.7↓</sup>
w/ Random Selection	30.9	51.7	41.3 <sup>-2.0↓</sup>
w/ Self-Consistency	31.5	52.0	41.8 <sup>-1.5↓</sup>

Table 3: Ablation study. Using Qwen2.5-7B as the backbone, we systematically evaluate the impact of: (1) removing thought cards (replaced with random action combinations), (2) disabling card matching (replaced with random cards), and (3) replacing verification with simpler alternatives (replaced with random selection or self-consistency).

understanding (ChartQA). Figure 4 shows AStar’s consistent improvements across non-mathematical domains. AStar with Qwen2-VL-2B achieves notable gains on GAOKAO (44.3% vs. 35.3% SC@5), and MMStar (52.0% vs. 48.8%), while Qwen2-VL-7B shows similar improvements. Remarkably, even for GPT-4o, AStar provides consistent enhancements across benchmarks like MMMU (73.2% vs. 70.3%) and GAOKAO-MM (52.2% vs. 47.8%). This demonstrates that mathematical thought cards can successfully transfer to diverse domains, validating that high-level abstract reasoning patterns provide robust cross-domain generalization. We further explore weak-to-strong generalization in Appendix E.4, where reasoning guidance from smaller models (Qwen2-VL-7B) even enhances GPT-4o’s performance, supporting our framework’s universal applicability.

### 3.6 Ablation Study and Analysis

**Ablation Study.** As shown in Table 3, we remove or replace each key component in AStar to understand its individual performance impacts:

- *Structured reasoning patterns matter.* Replacing thought cards with random action sequences results in a 9.5% performance drop, highlighting the importance of systematically extracted reasoning patterns.

- *Problem-pattern matching is crucial.* Different matching metrics yield varying performance, demonstrating that retrieving appropriate thought cards based on problem characteristics is vital for optimal performance. This validates




Figure 4: Transferability verification. Beyond math reasoning, we conduct extensive cross-domain experiments on diverse tasks, including general reasoning (MMMU), domain-specific reasoning (GAOKAO), visual perception and comprehension (MMStar and BLINK), and chart understanding (ChartQA). Results on ChartQA are consistently scaled for better visualization, which does not affect the conclusions. Our framework exhibits remarkable cross-domain generalization.

Seed Data	MathVision $\uparrow$	MathVerse $\uparrow$	Average $\uparrow$
50	23.8	43.2	33.5
100	28.9	49.8	39.4
200	30.8	51.8	41.3
500	32.7	53.9	43.3
1000	33.4	54.8	44.1

Table 4: Impact of seed dataset size

our adaptive selection mechanism.

◦ *Verification provides consistent gains.* While replacing our verification module with simpler alternatives shows modest degradation (e.g., 1.5% on self-consistency), the results indicate that thought cards enable robust solution generation even with basic verification strategies.

**Impact of Seed Dataset Size.** As shown in Table 4, the performance varies with the size of the seed dataset. On average, the accuracy improves from 33.5% (with 50 samples) to 44.1% (with 1,000 samples). Notably, AStar achieves competitive accuracy even with just 100 samples, highlighting its strong practicality in resource-constrained settings. In this paper, we select 500 seed samples by default, balancing performance and computational efficiency.

**Other Analysis.** We report pass@16 results on MathVision with Qwen2-VL-7B, achieving 26.2% (Base) and 35.9% (AStar), indicating enhanced model exploration through high-quality reasoning path planning. Additional results are provided in the Appendix, including intergration with SFT and GRPO, weak-to-strong generalization, sampling space visualization, and case study.

## 4 Related Work

**Multimodal Reasoning.** MLLMs have demonstrated robust capabilities across diverse domains, yet complex multimodal reasoning remains challenging. While recent methods (Xu et al. 2024; Thawakar et al. 2025) implement structured reasoning to enhance CoT capabilities, they face crit-

ical limitations: (i) explicit search methods (Dong et al. 2024a; Yao et al. 2024) suffer from computational inefficiency, and (ii) post-training methods (Wang et al. 2025; Zhang et al. 2025a) demand substantial resources, relying on SFT with limited generalization or RL (e.g., GRPO) lacking external knowledge integration (Yue et al. 2025; Wu et al. 2025). Moreover, existing methods employ rigid structures that limit flexibility, overlooking adaptive reasoning’s importance (Wang et al. 2024c). AStar leverages high-level thought cards to integrate internal model capabilities with external guidelines, enabling task-specific reasoning path generation that balances performance with efficiency.

**Tree-based Search.** Tree structures have shown significant potential in language models (Qi et al. 2024; Wu et al. 2024). Recent research applies them to identify effective reasoning paths for MLLMs. Recent research applies them to identify effective reasoning paths for MLLMs. AR-MCTS (Dong et al. 2024a) integrates MCTS with active retrieval but requires extensive iterations. Mulberry (Yao et al. 2024) distills 260K reasoning chains from GPT-4o using tree structures, yet demands substantial computational resources and high-capacity teacher models. These methods fail to balance performance with efficiency. We incorporate high-level problem-solving guidelines into MCTS, achieving competitive performance with higher efficiency. While using MCTS as our reasoning organization structure, our framework also accommodates other advanced structures in future work.

## 5 Conclusion

We introduce AStar, a training-free structured thinking paradigm for multimodal reasoning. By leveraging *thought cards* and adaptively retrieving problem-specific guidelines, our method effectively integrates MLLMs’ internal capabilities with external explicit guidance. AStar-7B achieves 53.9% accuracy on MathVerse and 32.7% on MathVision, surpassing GPT-4o. Further analysis reveals that thought cards from math domains also benefit other tasks like visual perception. As a plug-and-play test-time framework, AStar offers a promising pathway for multimodal reasoning.

## Acknowledgments

This work is supported by the Excellent Youth Program of State Key Laboratory of Multimodal Artificial Intelligence Systems (No. MAIS2024311) and Youth Science Fund Project of National Natural Science Foundation of China (No. 62201572).

## References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Chen, L.; Li, J.; Dong, X.; Zhang, P.; Zang, Y.; Chen, Z.; Duan, H.; Wang, J.; Qiao, Y.; Lin, D.; and Zhao, F. 2024a. Are We on the Right Way for Evaluating Large Vision-Language Models? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Da Silva, S. 2023. System 1 vs. System 2 Thinking. *Psych*, 5(4): 1057–1076.
- Dong, G.; Zhang, C.; Deng, M.; Zhu, Y.; Dou, Z.; and Wen, J.-R. 2024a. Progressive multimodal reasoning via active retrieval. *arXiv preprint arXiv:2412.14835*.
- Dong, Y.; Liu, Z.; Sun, H.-L.; Yang, J.; Hu, W.; Rao, Y.; and Liu, Z. 2024b. Insight-V: Exploring Long-Chain Visual Reasoning with Multimodal Large Language Models. *arXiv preprint arXiv:2411.14432*.
- Fu, X.; Hu, Y.; Li, B.; Feng, Y.; Wang, H.; Lin, X.; Roth, D.; Smith, N. A.; Ma, W.-C.; and Krishna, R. 2025. BLINK: Multimodal Large Language Models Can See but Not Perceive. In Leonardis, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; and Varol, G., eds., *Computer Vision – ECCV 2024*, 148–166. Cham: Springer Nature Switzerland. ISBN 978-3-031-73337-6.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hu, Y.; Shi, W.; Fu, X.; Roth, D.; Ostendorf, M.; Zettlemoyer, L.; Smith, N. A.; and Krishna, R. 2024. Visual Sketchpad: Sketching as a Visual Chain of Thought for Multimodal Language Models. *arXiv preprint arXiv:2406.09403*.
- Jaffe, P. I.; Poldrack, R. A.; Schafer, R. J.; and et al. 2023. Modelling human behaviour in cognitive tasks with latent dynamical systems. *Nature Human Behaviour*, 7: 986–1000.
- Kocsis, L.; and Szepesvári, C. 2006. Bandit Based Monte-Carlo Planning. In Fürnkranz, J.; Scheffer, T.; and Spiliopoulou, M., eds., *Machine Learning: ECML 2006*, 282–293. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-46056-5.
- Lee, F.-L.; and Heyworth, R. 2000. Problem complexity: A measure of problem difficulty in algebra by using computer. *Education Journal*, 28(1): 85–108.
- Lin, Z.; Liu, C.; Zhang, R.; Gao, P.; Qiu, L.; Xiao, H.; Qiu, H.; Lin, C.; Shao, W.; Chen, K.; et al. 2023. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2311.07575*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, Z.; Zhang, Y.; Liu, F.; Zhang, C.; Sun, Y.; and Wang, J. 2025. OThink-MR1: Stimulating multimodal generalized reasoning capabilities through dynamic reinforcement learning. *arXiv preprint arXiv:2503.16081*.
- Lu, P.; Bansal, H.; Xia, T.; Liu, J.; Li, C.; Hajishirzi, H.; Cheng, H.; Chang, K.-W.; Galley, M.; and Gao, J. 2023. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*.
- Luo, R.; Zheng, Z.; Wang, Y.; Yu, Y.; Ni, X.; Lin, Z.; Zeng, J.; and Yang, Y. 2025. URSA: Understanding and Verifying Chain-of-thought Reasoning in Multimodal Mathematics. *arXiv preprint arXiv:2501.04686*.
- Masry, A.; Do, X. L.; Tan, J. Q.; Joty, S.; and Hoque, E. 2022. ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Findings of the Association for Computational Linguistics: ACL 2022*, 2263–2279. Dublin, Ireland: Association for Computational Linguistics.
- Meng, F.; Du, L.; Liu, Z.; Zhou, Z.; Lu, Q.; Fu, D.; Han, T.; Shi, B.; Wang, W.; He, J.; et al. 2025. MM-Eureka: Exploring the Frontiers of Multimodal Reasoning with Rule-based Reinforcement Learning. *arXiv preprint arXiv:2503.07365*.
- OpenAI. 2024. GPT-4o System Card. <https://openai.com/research/gpt-4o-system-card>. Accessed: 2024-05.
- OpenAI. 2024. Learning to Reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>. Accessed: 2024-09.
- Peng, S.; Fu, D.; Gao, L.; Zhong, X.; Fu, H.; and Tang, Z. 2024. Multimath: Bridging visual and mathematical reasoning for large language models. *arXiv preprint arXiv:2409.00147*.
- Peng, Y.; Zhang, G.; Zhang, M.; You, Z.; Liu, J.; Zhu, Q.; Yang, K.; Xu, X.; Geng, X.; and Yang, X. 2025. Lmm-r1: Empowering 3b llms with strong reasoning abilities through two-stage rule-based rl. *arXiv preprint arXiv:2503.07536*.
- Qi, Z.; Ma, M.; Xu, J.; Zhang, L. L.; Yang, F.; and Yang, M. 2024. Mutual reasoning makes smaller llms stronger problem-solvers. *arXiv preprint arXiv:2408.06195*.
- Qiao, R.; Tan, Q.; Dong, G.; et al. 2024. We-Math: Does Your Large Multimodal Model Achieve Human-like Mathematical Reasoning? *arXiv preprint arXiv:2407.01284*.

- Qwen Team. 2024. Qwen2.5: A Party of Foundation Models.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmlR.
- Russell, S.; and Wefald, E. 1991. Principles of metareasoning. *Artificial Intelligence*, 49(1): 361–395.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shi, W.; Hu, Z.; Bin, Y.; Liu, J.; Yang, Y.; Ng, S.-K.; Bing, L.; and Lee, R. K.-W. 2024. Math-LLaVA: Bootstrapping Mathematical Reasoning for Multimodal Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 4663–4680. Miami, Florida, USA: Association for Computational Linguistics.
- Team, K.; Du, A.; Gao, B.; Xing, B.; Jiang, C.; Chen, C.; Li, C.; Xiao, C.; Du, C.; Liao, C.; et al. 2025. Kimi k1.5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*.
- Thawakar, O.; Dissanayake, D.; More, K.; Thawkar, R.; Heakl, A.; Ahsan, N.; Li, Y.; Zumri, M.; Lahoud, J.; Anwer, R. M.; et al. 2025. LlamaV-o1: Rethinking Step-by-step Visual Reasoning in LLMs. *arXiv preprint arXiv:2501.06186*.
- Wang, K.; Pan, J.; Shi, W.; Lu, Z.; Ren, H.; Zhou, A.; Zhan, M.; and Li, H. 2024a. Measuring Multimodal Mathematical Reasoning with MATH-Vision Dataset. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wang, W.; Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Zhu, J.; Zhu, X.; Lu, L.; Qiao, Y.; et al. 2024c. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*.
- Wang, Y.; Chen, W.; Han, X.; Lin, X.; Zhao, H.; Liu, Y.; Zhai, B.; Yuan, J.; You, Q.; and Yang, H. 2024d. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*.
- Wang, Y.; Wu, S.; Zhang, Y.; Wang, W.; Liu, Z.; Luo, J.; and Fei, H. 2025. Multimodal chain-of-thought reasoning: A comprehensive survey. *arXiv preprint arXiv:2503.12605*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, J.; Feng, M.; Zhang, S.; Che, F.; Wen, Z.; and Tao, J. 2024. Beyond examples: High-level automated reasoning paradigm in in-context learning via mcts. *arXiv preprint arXiv:2411.18478*.
- Wu, J.; Liao, C.; Feng, M.; Zhang, S.; Wen, Z.; Shao, P.; Xu, H.; and Tao, J. 2025. Thought-augmented policy optimization: Bridging external guidance and internal capabilities. *arXiv preprint arXiv:2505.15692*.
- Xu, G.; Jin, P.; Hao, L.; Song, Y.; Sun, L.; and Yuan, L. 2024. LLaVA-o1: Let Vision Language Models Reason Step-by-Step. *arXiv preprint arXiv:2411.10440*.
- Yang, L.; Yu, Z.; Zhang, T.; Cao, S.; Xu, M.; Zhang, W.; Gonzalez, J. E.; and Cui, B. 2024. Buffer of Thoughts: Thought-Augmented Reasoning with Large Language Models. *Advances in Neural Information Processing Systems*.
- Yao, H.; Huang, J.; Wu, W.; Zhang, J.; Wang, Y.; Liu, S.; Wang, Y.; Song, Y.; Feng, H.; Shen, L.; et al. 2024. Mulberry: Empowering MLLM with o1-like Reasoning and Reflection via Collective Monte Carlo Tree Search. *arXiv preprint arXiv:2412.18319*.
- Yin, S.; Fu, C.; Zhao, S.; Li, K.; Sun, X.; Xu, T.; and Chen, E. 2023. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*.
- Yue, X.; Ni, Y.; Zhang, K.; Zheng, T.; Liu, R.; Zhang, G.; Stevens, S.; Jiang, D.; Ren, W.; Sun, Y.; Wei, C.; Yu, B.; Yuan, R.; Sun, R.; Yin, M.; Zheng, B.; Yang, Z.; Liu, Y.; Huang, W.; Sun, H.; Su, Y.; and Chen, W. 2024. MMMU: A Massive Multi-discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI. In *Proceedings of CVPR*.
- Yue, Y.; Chen, Z.; Lu, R.; Zhao, A.; Wang, Z.; Song, S.; and Huang, G. 2025. Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model? *arXiv preprint arXiv:2504.13837*.
- Zhang, J.; Huang, J.; Yao, H.; Liu, S.; Zhang, X.; Lu, S.; and Tao, D. 2025a. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv preprint arXiv:2503.12937*.
- Zhang, R.; Jiang, D.; Zhang, Y.; Lin, H.; Guo, Z.; Qiu, P.; Zhou, A.; Lu, P.; Chang, K.-W.; Qiao, Y.; et al. 2025b. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? In *European Conference on Computer Vision*, 169–186. Springer.
- Zhang, Z.; Zhang, A.; Li, M.; hai zhao; Karypis, G.; and Smola, A. 2024. Multimodal Chain-of-Thought Reasoning in Language Models. *Transactions on Machine Learning Research*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zhuang, W.; Huang, X.; Zhang, X.; and Zeng, J. 2024. Mathpuma: Progressive upward multimodal alignment to enhance mathematical reasoning. *arXiv preprint arXiv:2408.08640*.
- Zong, Y.; and Qiu, X. 2024. GAOKAO-MM: A Chinese Human-Level Benchmark for Multimodal Models Evaluation. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 8817–8825. Bangkok, Thailand: Association for Computational Linguistics.