

VIRC: Enhancing Visual Interleaved Mathematical CoT with Reason Chunking

Supplementary Material

Lihong Wang^{1,2†} Liangqi Li² Weiwei Feng² Jiamin Wu³ Changtao Miao^{2*}
Tieru Wu^{1,4} Rui Ma^{1,4*} Bo Zhang² Zhe Li²

¹Jilin University ²Ant Digital Technologies, Ant Group ³The Chinese University of Hong Kong
⁴Engineering Research Center of Knowledge-Driven Human-Machine Intelligence, MOE, China

Overview

We organize our supplementary material as follows.

- [Background on Miller’s Law](#)
- [Dataset Details](#)
 - [Dataset Statistic](#)
 - [Dataset Curation](#)
 - [Modeling Cognitive Reasoning Patterns](#)
- [Additional Implementation Details](#)
 - [Instructional SFT](#)
 - [Practice SFT](#)
 - [Strategic RL](#)
 - [Hard Subset Curation](#)
 - [Reward Design](#)
- [Evaluation Protocol](#)
- [Additional Results](#)
 - [Additional Quantitative Results](#)
 - [Detailed Ablation Study Results](#)
 - [Additional Qualitative Results](#)
- [Dataset and Prompt Details](#)
 - [Dataset Example](#)
 - [Prompts for Data Curation](#)
 - [Prompts for Supervised Fine-Tuning](#)
 - [Prompts for Reinforcement Learning](#)

A. Background on Miller’s Law

For a better understanding of the proposed **Reason Chunking** mechanism, this section provides a concise explanation of Miller’s Law [8] and its relevance to multimodal mathematical reasoning. Miller’s Law posits that the capacity of human short-term memory is limited, typically to around seven items. To overcome this cognitive bottleneck, humans group related low-level information into semantically meaningful higher-level units, or *chunks*. This principle forms the initial motivation behind structuring complex reasoning into manageable and semantically coherent components.

Miller’s Law and human information organization. Miller’s Law suggests that short-term memory can only

maintain a small number of items at once. By applying *information chunking*, multiple low-level elements can be reorganized into higher-level semantic units, substantially reducing cognitive load. For instance, eight unrelated letters such as “C V P R S U P P” become far easier to recall when reorganized into the two meaningful chunks “CVPR SUPP”. This illustrates a general cognitive principle: humans naturally process complex inputs by forming compact, interpretable units, which improves reasoning clarity and reduces memory stress. This principle directly inspires the idea that long reasoning chains should be decomposed into structured intermediate units rather than treated as flat sequences.

Information chunking in mathematical and visual reasoning. The principle of information chunking is evident in mathematical problem-solving. Humans decompose a complex problem into a series of intermediate propositions, perform localized reasoning for each, and then integrate these propositions hierarchically. This strategy avoids the cognitive overload that would result from maintaining a lengthy, linear reasoning chain in working memory. The scoring systems of mathematical competitions like the IMO and AMC reinforce this natural structure: *scoring points* are awarded for proving intermediate results, acknowledging that mathematical reasoning is inherently composed of semantically meaningful information chunks. This cognitive pattern extends to visual reasoning. Instead of re-examining an entire image at every step, humans selectively attend to relevant visual cues only when needed. This prevents an excess of visual information from overwhelming short-term memory. Together, these behaviors align with Miller’s insight that effective reasoning depends on forming and operating on appropriately sized informational chunks.

Operationalizing Miller’s principle in VIRC. Guided by this chunk-based cognitive pattern, we design the *Reason Chunking* mechanism to operationalize Miller’s Law within multimodal mathematical reasoning. The overall reasoning process is structured into a sequence of **Critical Reasoning Units (CRUs)**, each functioning as a semantically complete chunk that encapsulates a single intermediate proposition. These CRUs form a hierarchical reasoning path that mirrors human decomposition of complex problems. More-

[†]Work done during Lihong Wang’s internship at Ant Digital Technologies, Ant Group.

*Corresponding authors.

over, *selective visual grounding* is applied only between CRUs, simulating the human tendency to consult visual information at critical points rather than continuously. In this way, **VIRC** leverages the chunk-based processing principle highlighted by Miller’s Law to produce structured, cognitively aligned multimodal reasoning.

B. Dataset Details

This section provides a comprehensive overview of our constructed **CRUX** dataset. We first present its overall statistics in Sec. B.1. Next, we elaborate on the automated curation pipeline in Sec. B.2, including how diverse reasoning paths are sampled and mapped to Critical Reasoning Units (CRUs). Finally, we describe how the four cognitive reasoning patterns of Planning, Reflecting, Verifying, and Backtracking are modeled and systematically integrated into the dataset in Sec. B.3.

B.1. Dataset Statistic

The key statistics of the CRUX dataset are summarized in Tab. 1. This includes the total number of data points, the distribution of Critical Reasoning Units (CRUs) per sample, and the frequency of various tool calls and cognitive reasoning patterns.

Table 1. Key statistics of the CRUX dataset.

Statistic	Value
Total data points	100,000
Average resolution of query images	431×409 px
Minimum CRUs per data point	1
Maximum CRUs per data point	21
Average CRUs per data point	4.27
Total <code>crop</code> tool calls	281,248
Total <code>scale</code> tool calls	11,492
Total <code>display</code> tool calls	60,034
Total “Planning” pattern occurrences	100,000
Total “Backtracking” pattern occurrences	9,682
Total “Verifying” pattern occurrences	38,226
Total “Reflecting” pattern occurrences	179,008

B.2. Dataset Curation

To support structured multimodal reasoning based on Critical Reasoning Units (CRUs), we construct the CRUX dataset via an automated pipeline involving four specialized LLM/VLM components (see the main paper Fig. 3). While CRUX inherits raw problem–image pairs from MINT-CoT [1], all reasoning paths and CRU annotations are newly generated using carefully designed prompts (see Sec. F.2), as detailed below.

Sampling Diverse Reasoning Paths. For each image, we generate five scaled variants ($f \in \{0.25, 0.5, 1, 2, 4\}$) and

constrain image token count to [4, 16384]. Using zero-shot CoT (see Fig. 5), we sample five reasoning paths per scale and verify answer correctness. We select scales f_- (low accuracy) and f_+ (high accuracy) with a $\geq 60\%$ accuracy gap, preferring maximal resolution difference. From these, we take the longest correct path (p_0) and longest incorrect paths (p_1 at f_- , p_2 at f_+) as the source paths for subsequent composition.

Mapping Steps to CRUs. Each path is decomposed into fine-grained steps annotated with a focus object (e.g., “ $\angle ABC$ ”). Consecutive steps sharing the same focus object form one CRU, representing a single intermediate proposition. Using p_0 as the semantic anchor, the Mapper aligns each step in p_1 and p_2 to its most semantically similar step in p_0 , thereby establishing a correspondence not only between steps but also between their associated intermediate propositions (i.e., CRUs in p_0). The Mapper then identifies the first erroneous step in each incorrect path (err_1 in p_1 , err_2 in p_2) along with its aligned proposition in p_0 . Only the steps up to and including the first error are retained, and these are grouped according to the CRU in p_0 to which they map. Finally, CRUs from p_1 , p_2 , and p_0 are concatenated as $p_1 \rightarrow p_2 \rightarrow p_0$ to simulate error-aware reasoning.

Grounding CRUs. For each CRU, the Linker detects the focus object and associated textual annotations in the original image. Their bounding boxes are fused into a unified grounding region, prioritizing text when available. Each CRU is linked to a `crop` tool call targeting this region. Additionally, the Linker generates image descriptions, rationales, and guiding questions to ensure coherent transitions between CRUs.

B.3. Modeling Cognitive Reasoning Patterns

Building upon the preliminary dataset described in Sec. B.2, which contains only canonical `crop` tool calls derived from CRU bounding boxes, we apply a rule-based post-processing pipeline to instantiate four cognitive reasoning patterns Planning, Reflecting, Verifying, and Backtracking. These transformations operationalize human-like problem-solving behaviors as explicit edits to the initial tool invocations.

Planning. At the start of each reasoning path, we generate a global image description and high-level solution rationale from the f_+ -scaled image to establish strategic context before CRU-level reasoning begins.

Reflecting. When a CRU’s region of interest is fully contained within the output of a preceding `crop` call from the same source path (i.e., within p_0 , p_1 , or p_2), we revise its tool invocation to reference that prior output image. The `image_index` is updated accordingly, and the bounding box is adjusted to preserve the physical region, enabling iterative visual focusing.

Verifying. At the first erroneous step err_2 in path p_2 , we

replace the original `crop` call with a `display` invocation that explicitly references the original image, prompting the model to re-examine initial visual evidence and validate the current proposition. To maintain continuity of this reasoning pattern, all subsequent tool calls in the path operate on the image output produced by this `display` call, thereby grounding later steps in the verified visual context.

Backtracking. At the first erroneous step err_1 in path p_1 , we insert a `scale` call applied to the original image with factor f_+/f_- to correct misjudgments caused by reasoning under the low-accuracy scale f_- . To support this mechanism, the visual input for the entire p_1 path is rendered at scale f_- during initial sampling, ensuring that upscaling by f_+/f_- recovers the high-accuracy view used in p_0 .

All components, including planning statements, CRU sequences, structured tool calls in JSON format, guiding questions, and the final answer, are assembled into a complete cognitively grounded reasoning path following the compositional structure illustrated in the main paper Fig. 2 (d).

C. Additional Implementation Details

All training experiments are conducted using established open-source frameworks to ensure full reproducibility. We adopt Qwen2.5-VL-7B [16] and Qwen2.5-VL-3B [16] as the base MLLMs and obtain **VIRC-7B** and **VIRC-3B** through the three-stage pipeline of Instructional SFT, Practice SFT, and Strategic RL. Both model variants use identical training configurations and hyperparameters. During training, we freeze the vision encoder and the vision-language projector.

C.1. Instructional SFT

Instructional SFT is conducted on the 50K text-only subset of CRUX, matching the scale used in MINT-CoT [1]. Tool outputs are removed, and only textual supervision is preserved. During this stage, we set the system prompt to the template shown in Fig. 12. Training is performed using LLaMA Factory [26] with a learning rate of $5.0\text{e-}6$, 3 epochs, cosine scheduling, a warmup ratio of 0.03, and a batch size of 64.

C.2. Practice SFT

Practice SFT uses the same 50K samples, but all tool calls are executed offline and the returned images are cached. These images are inserted into the sequence immediately after each tool call. We again employ the system prompt in Fig. 12, but explicitly remove the “# Emergency” section and its associated content to ensure that the model fully leverages visual cues during training. This stage is trained with LLaMA Factory using a learning rate of $1.0\text{e-}6$, 6 epochs, cosine scheduling, a warmup ratio of 0.1, and a batch size of 64.

C.3. Strategic RL

For the Strategic RL stage, we use the same system prompt as in Practice SFT and adopt Group Relative Policy Optimization (GRPO) [10] as implemented in verl [11]. Unlike standard RL settings where the full question is provided as input, each policy query consists of a truncated reasoning fragment q drawn from a curated hard subset \mathcal{HD} . Each fragment represents an incomplete reasoning prefix from which the model must either reach the final answer or continue tool-guided reasoning.

Given a fragment q , the policy produces a group of G candidate continuations $\{o_i\}_{i=1}^G$. Instead of manually defining the reward here, we directly use the reward value $r_i = r(s_t, o_i)$ as specified in the main paper Sec. 3.4, where $r(\cdot)$ integrates answer correctness, multimodal coherence, reasoning-pattern alignment, and format validity penalty into a single composite score. Advantages A_i are computed by normalizing each reward within the group. In our implementation, we do not include KL or entropy regularization and the policy is updated using the GRPO objective:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{\{q \sim \mathcal{HD}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(\cdot|q)\}} \frac{1}{G} \sum_{i=1}^G \left(\min \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{old}}(o_i|q)} A_i, \text{clip} \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right) \right) \quad (1)$$

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})} \quad (2)$$

We set $\epsilon = 0.2$. The global batch size is 512 with a mini-batch size of 64. A constant learning rate of 1×10^{-6} is used. To maintain computational efficiency, the maximum response length is limited to 512 tokens, and asynchronous rollouts are enabled throughout training.

C.4. Hard Subset Curation

The RL hard subset contains 8K truncated samples selected from CRUX and is designed to balance two abilities: producing final answers when sufficient context is available and continuing tool-guided reasoning when only partial evidence is provided. The subset is composed of two 4K groups, corresponding to the truncated sample types illustrated in Fig.3 (a) Stage 3. The first group, *Long Reasoning Truncation*, contains the 4K samples with the longest reasoning paths, measured by the total number of textual and visual tokens. For each sample, the final CRU is removed while retaining the ground-truth answer. These fragments tax the model’s long-range information integration capabilities, requiring it to synthesize a final answer from highly complex, accumulated cognitive contexts.

The second group, *Critical Region Truncation*, is constructed from samples not included in the first group. For

each such sample, we identify the tool-returned region that has the smallest area ratio (below 20%) relative to the original image and belongs to the latest possible CRU. This CRU and all subsequent steps are truncated while keeping the tool type of that CRU and the final answer as ground truth. Samples are grouped by tool type, and within each group we retain the top $\frac{4}{3}$ K samples with the longest reasoning fragments. The three tool-specific groups jointly form the 4K Critical Region Truncation set. These samples impose acute perceptual and strategic challenges, requiring the model to recognize informational deficits from obscure visual cues and autonomously select the correct investigative tool to proceed.

C.5. Reward Design

Reward scores for r_{ans} and r_{mm} are computed using Qwen2.5-VL-72B-Instruct [16]. The prompt for r_{ans} is shown in Fig. 13. The multimodal reward r_{mm} includes a textual-coherence component $r_{\text{mm}}^{\text{text}}$ (prompt in Fig. 14) and a visual-relevance component $r_{\text{mm}}^{\text{vis}}$ (prompt in Fig. 15). All returned scores are strictly required to lie in $[0, 1]$. Otherwise the query is repeated. We assign weights of 0.4 and 0.5 to $r_{\text{mm}}^{\text{text}}$ and $r_{\text{mm}}^{\text{vis}}$, respectively. Since r_{pattern} yields either 0 or 0.1, we constrain the sum of the maxima of all multimodal reward components plus the maximum value of r_{pattern} to equal 1, matching the maximum value of r_{ans} and preventing potential reward hacking.

D. Evaluation Protocol

The main paper omits detailed explanations of our evaluation setup due to limited space. Here, we clarify how we handle models that support tool invocation or integrate visual operations during reasoning. For fairness, whenever a model provides an official or built-in mechanism for zoom-in, crop, or other visual-grounding tools, we enable these features during evaluation. Specifically, Qwen3-VL-8B-Instruct [17], Pixel Reasoner [12], DeepEyes [27], and Mini-o3 [5] are evaluated with their respective zoom-in or visual tools activated. MINT-CoT-7B [1] is tested with its interleaved visual tokens mechanism enabled, following its official implementation. This unified setting ensures that all baselines fully exploit their intended visual reasoning capabilities for a fair comparison.

In addition, we use Qwen2.5-72B-Instruct [14] to verify the correctness of the answers generated by each model. The validation prompt is provided in Fig. 6.

E. Additional Results

This section provides supplementary results to complement and substantiate the findings presented in the main paper. We begin by presenting detailed quantitative results and controlled studies in Sec. E.1. Subsequently, we provide

the full, fine-grained evaluation data from our ablation studies to offer a clearer view of each component’s impact in Sec. E.2. Finally, we include additional qualitative results that visually demonstrate our model’s reasoning capabilities on various problems in Sec. E.3.

E.1. Additional Quantitative Results

Generalization under domain shift. The main paper reports averaged scores on generalization benchmarks. Here we provide more detailed results on VisualProbe [5], V^* [21], and HR-Bench [20] in Tab. 2. Although **ViRC** is trained only on CRUX, which consists of low-resolution schematic mathematical diagrams with CRU-align reasoning and tool calls, it generalizes well to high-resolution natural images (2K–16K) and perception-centric questions without any domain-specific fine-tuning. These results suggest that the learned reason chunking and selective tool usage transfer beyond the training distribution.

Robustness to pretrained models. Tab. 3 shows that **ViRC** yields consistent improvements when instantiated with different pretrained model families. In particular, the gains hold when switching from Qwen2-VL, which uses a fixed-resolution strategy, to Qwen2.5-VL, which adopts a dynamic-resolution strategy, indicating that our method is robust to the pretrained model choice and its image-resolution policy.

Benefits beyond distillation from the teacher model. Results in Tab. 4 show that **ViRC-7B** can even surpass the teacher model on multiple benchmarks. This suggests that the reason chunking mechanism provides additional benefits over the base model by introducing chunk-level organization and denser training signals.

Broader multimodal math benchmarks. Tab. 5 reports additional results on broader multimodal math benchmarks, including MathVision [18], We-Math [9], MathVerse [25], and Logic Vista [23]. For a fair comparison against a method that also uses multimodal CoT, we include MINT-CoT [1] as a strong baseline. Notably, MINT-CoT is within 1% of the best model compared in the main paper, MM-Eureka-Qwen-7B [7], while **ViRC-7B** consistently outperforms MINT-CoT across these benchmarks.

Inference-time cost. Tab. 6 reports token usage and inference latency measured with vLLM [4] v0.14.1 on a single A100 80G GPU (batch size 1, FP16) across all benchmarks reported in Tabs. 4 and 5. We report both average statistics and worst-case settings based on the top 1% longest token sequences and the slowest 1% inference cases. **ViRC-7B** achieves the lowest average token usage and inference latency while maintaining strong performance. Removing CRU increases both token consumption and latency, highlighting the efficiency benefit of reason chunking. Under worst-case settings, **ViRC-7B** achieves the second-best results, remaining competitive in inference-time cost.

Table 2. **Fine-grained generalization results.** Our **ViRC** is trained on the CRUX dataset and cross-domain tested on VisualProbe [5], V^* [21], and HR-Bench [20]. Avg. reports the mean of the results on the three benchmarks. Results marked with * are from [5], while all other results are reproduced in this work.

Model	#Params	Avg.	VisualProbe			V^*	HR-Bench	
			easy	medium	hard		4K	8K
<i>Closed-Source Model</i>								
GPT-4o [13]	–	43.27	47.50*	15.40*	11.20*	65.20*	62.00*	58.30*
<i>Open-Source General Model</i>								
Qwen3-VL-8B-Instruct _{w/ tool} [17]	8B	33.47	36.68	10.82	21.70	41.88	44.12	45.62
LLaVA-OV-Qwen2-7b-ov [6]	7B	41.37	36.20*	12.50*	13.40*	70.90*	61.20*	54.00*
InternVL2.5-8B [2]	8B	44.03	<u>55.32</u>	11.57	11.32	68.59	62.75	54.62
InternVL2.5-8B-MPO [19]	8B	42.45	41.13	13.43	8.49	72.25	64.12	55.25
DeepSeek-VL2 [22]	4.5B	38.43	32.62	11.19	7.55	66.49	61.12	51.62
<i>Open-Source Reasoning Model</i>								
Hint-GRPO-Qwen2.5-VL-3B [3]	3B	45.81	46.10	20.90	17.92	68.06	62.62	59.25
Hint-GRPO-Qwen2-VL-7B [3]	7B	47.12	50.35	17.16	16.04	<u>75.92</u>	65.50	57.75
R1-VL-7B [24]	7B	34.51	39.72	5.97	15.09	45.55	54.37	46.38
MM-Eureka-Qwen-7B [7]	7B	32.87	26.24	4.10	4.72	59.16	58.00	45.00
MINT-CoT-7B [1]	7B	32.04	34.75	11.32	7.46	43.46	51.25	44.00
Qwen2.5-VL-3B-Instruct [16]	3B	46.64	36.17	22.01	23.58	68.06	67.25	62.75
ViRC-3B (Ours)	3B	<u>51.10</u>	49.65	25.75	<u>25.47</u>	72.25	<u>70.00</u>	<u>63.50</u>
Qwen2.5-VL-7B-Instruct [16]	7B	49.25	39.10*	26.00*	23.90*	75.50*	68.25*	62.75*
ViRC-7B (Ours)	7B	58.27	56.03	36.94	37.74	79.06	73.12	66.75

Table 3. **Robustness to pretrained model.** We apply our **ViRC** training pipeline to two different pretrained MLLMs, Qwen2-VL-2B-Instruct and Qwen2-VL-7B-Instruct, yielding **ViRC-Qwen2VL-2B** and **ViRC-Qwen2VL-7B**, respectively.

Model	#Params	Avg.	GeoQA	MMStar Math	MathVista-Math				
					ALL	GEO	ALG	GPS	TQA
Qwen2-VL-2B-Instruct [15]	2B	25.23	15.52	32.40	27.78	24.06	27.72	24.52	38.71
ViRC-Qwen2VL-2B (Ours)	2B	44.20	44.16	44.00	44.44	42.92	44.94	43.27	48.39
Qwen2-VL-7B-Instruct [15]	7B	44.49	30.37	52.00	51.11	46.70	50.94	47.60	62.90
ViRC-Qwen2VL-7B (Ours)	7B	64.32	72.68	58.80	61.48	58.02	61.42	58.17	72.58

Table 4. **Benefits beyond the teacher model.** We compare Qwen2.5-VL-72B-Instruct, the teacher model, with **ViRC-7B** trained on the CRUX dataset.

Model	#Params	Avg.	GeoQA	MMStar Math	MathVista-Math				
					ALL	GEO	ALG	GPS	TQA
Qwen2.5-VL-72B-Instruct [16]	72B	<u>75.12</u>	<u>66.45</u>	80.40	78.52	77.83	78.28	77.40	82.26
ViRC-7B (Ours)	7B	77.79	75.07	<u>77.20</u>	81.11	81.13	81.27	81.73	<u>79.03</u>

Table 5. **Additional benchmark results.** We further compare MINT-CoT-7B, Qwen2.5-VL-7B-Instruct, and **ViRC-7B** on broader multimodal math benchmarks, including MathVision [18], We-Math [9], MathVerse [25], and LogicVista [23].

Model	Avg.	MathVision	We-Math	MathVerse	LogicVista
MINT-CoT-7B [1]	38.54	22.52	<u>52.00</u>	<u>51.56</u>	26.12
Qwen2.5-VL-7B-Instruct [16]	<u>40.07</u>	<u>25.13</u>	42.67	49.19	<u>43.30</u>
ViRC-7B (Ours)	48.83	38.64	59.81	52.01	44.87

E.2. Detailed Ablation Study Results

This section provides additional details of the ablation configurations in the main paper Sec. 4.3. and report the full

ablation study results. Specifically, the main paper reports only averaged benchmark scores. Here we additionally provide subset- and category-level metrics for all ablated variants. Detailed results are shown in Tabs. 7 to 10.

Table 6. **Inference-time cost.** We report token usage and inference latency for different models and variants, including both average statistics and worst-case settings reported for the top 1% longest token sequences and the slowest 1% inference cases.

Model	CoT Type	Avg. #Tokens				Avg. Latency (s)	
		All	Text	Visual	Top 1% Len.	All	Top 1% Lat.
MINT-CoT-7B [1]	visual	1540	792	748	15275	9.38	19.22
MM-Eureka-Qwen-7B [7]	text-only	<u>945</u>	<u>384</u>	561	31004	<u>4.09</u>	10.80
ViRC-7B <i>w/o</i> CRU	visual	1415	522	893	29473	5.26	11.98
ViRC-7B (Ours)	reason chunking	879	183	<u>696</u>	<u>24144</u>	2.31	<u>11.00</u>

Table 7. **Fine-grained results for CRU effectiveness analysis.** Performance is broken down by subset and category.

Model	GeoQA	MMStar Math	MathVista-Math					VisualProbe			V^*	HR-Bench	
			ALL	GEO	ALG	GPS	TQA	easy	medium	hard		4K	8K
Baseline	43.50	66.80	66.67	65.56	66.29	65.87	69.35	39.10	26.00	23.90	75.50	68.25	62.75
Ours <i>w/o</i> CRU	<u>74.67</u>	<u>72.80</u>	<u>79.26</u>	<u>78.77</u>	<u>79.40</u>	<u>79.33</u>	<u>79.03</u>	<u>47.52</u>	<u>29.10</u>	<u>34.91</u>	<u>75.92</u>	<u>71.00</u>	<u>63.63</u>
Ours	75.07	77.20	81.11	81.13	81.27	81.73	79.03	56.03	36.94	37.74	79.06	73.12	66.75

Table 8. **Fine-grained results for reasoning pattern ablation.** Each row shows performance when a specific reasoning pattern is removed from training and evaluation. \mathcal{P} , \mathcal{B} , \mathcal{R} , and \mathcal{V} stand for 4 reasoning patterns—planning, backtracking, reflecting, and verifying, respectively.

Reasoning Pattern \mathcal{P} \mathcal{B} \mathcal{R} \mathcal{V}	GeoQA	MMStar Math	MathVista-Math					VisualProbe			V^*	HR-Bench	
			ALL	GEO	ALG	GPS	TQA	easy	medium	hard		4K	8K
	43.50	66.80	66.66	65.56	66.29	65.87	69.35	39.10	25.75	23.90	71.20	68.25	62.75
✓	<u>72.02</u>	<u>74.80</u>	<u>79.26</u>	<u>79.25</u>	<u>79.03</u>	<u>79.33</u>	79.03	47.52	27.24	32.08	75.39	71.62	65.50
✓ ✓	<u>72.81</u>	<u>68.40</u>	<u>78.52</u>	<u>79.25</u>	<u>78.65</u>	<u>79.81</u>	74.19	50.35	27.61	37.74	78.01	71.75	66.75
✓ ✓ ✓	<u>66.98</u>	74.00	68.89	67.92	69.29	68.75	69.35	51.77	31.34	<u>36.79</u>	<u>78.53</u>	70.75	66.75
✓ ✓ ✓ ✓	<u>74.93</u>	<u>72.40</u>	<u>77.41</u>	<u>77.83</u>	<u>77.53</u>	<u>77.88</u>	<u>75.81</u>	<u>53.90</u>	<u>32.46</u>	<u>34.91</u>	78.01	<u>72.50</u>	<u>65.88</u>
✓ ✓ ✓ ✓ ✓	75.07	77.20	81.11	81.13	81.27	81.73	79.03	56.03	36.94	37.74	79.06	73.12	66.75

Table 9. **Fine-grained results for the impact of the three training stages.** I , P , and S stand for Instructional SFT, Practice SFT, and Strategic RL, respectively.

Training Stage I P S	GeoQA	MMStar Math	MathVista-Math					VisualProbe			V^*	HR-Bench	
			ALL	GEO	ALG	GPS	TQA	easy	medium	hard		4K	8K
	43.50	66.80	66.67	65.56	66.29	65.87	69.35	39.10	26.00	23.90	75.50	68.25	62.75
✓	<u>70.27</u>	<u>70.80</u>	<u>76.30</u>	<u>77.36</u>	<u>76.40</u>	<u>77.88</u>	<u>70.97</u>	<u>42.55</u>	27.24	<u>29.25</u>	74.87	<u>71.75</u>	<u>66.00</u>
✓ ✓	<u>73.34</u>	<u>72.80</u>	<u>77.78</u>	<u>79.72</u>	<u>77.90</u>	<u>80.29</u>	69.35	<u>42.55</u>	<u>29.48</u>	26.42	<u>75.92</u>	70.25	63.75
✓ ✓ ✓	75.07	77.20	81.11	81.13	81.27	81.73	79.03	56.03	36.94	37.74	79.06	73.12	66.75

Table 10. **Fine-grained results for RL component ablation.** The results show that each one of the designs in our proposed Strategic RL plays an important role in multimodal reasoning. For brevity, we denote the curated hard subset derived from CRUX as $\mathcal{H}\mathcal{D}$.

Model	GeoQA	MMStar Math	MathVista-Math					VisualProbe			V^*	HR-Bench	
			ALL	GEO	ALG	GPS	TQA	easy	medium	hard		4K	8K
Ours	75.07	77.20	81.11	81.13	81.27	81.73	<u>79.03</u>	56.03	36.94	37.74	79.06	73.12	66.75
<i>w/o</i> $\mathcal{H}\mathcal{D}$	66.84	70.80	73.70	75.47	73.78	76.44	<u>64.52</u>	45.39	28.73	32.08	<u>76.44</u>	71.75	65.75
<i>w/o</i> r_{mm}	<u>73.60</u>	74.00	79.63	80.19	79.40	80.77	75.81	50.35	31.34	32.08	79.06	<u>72.50</u>	65.88
<i>w/o</i> r_{pattern}	<u>73.56</u>	<u>74.80</u>	<u>80.00</u>	<u>78.77</u>	<u>80.25</u>	<u>79.33</u>	82.26	46.10	<u>35.07</u>	<u>33.96</u>	<u>76.44</u>	71.88	<u>66.50</u>

Ablation Setup. To ensure fair and controlled comparisons, all ablation studies were conducted on the same base model, Qwen2.5-VL-7B-Instruct [16]. The training protocol for each ablation run was kept identical to the one used to train our final ViRC-7B model, whose results are reported in the state-of-the-art comparison (Sec. 4.2). Unless a component was intentionally modified as part of the study (e.g., skipping a training stage), all other conditions were held

constant. These conditions include all training hyperparameters, such as learning rate, batch size, and the number of epochs per stage. This methodology isolates the impact of each ablated component, ensuring that performance differences can be attributed directly to its removal.

CRU Effectiveness Analysis. To evaluate the model’s performance without explicit CRU guidance, we ablated the Instructional SFT stage. The model was trained directly

with Practice SFT, followed by Strategic RL. As described in Sec. 3.4 of the main paper, the Instructional SFT stage is responsible for teaching the model the internal textual structure of each CRU. Skipping this stage therefore simulates an environment where this structural prior is absent.

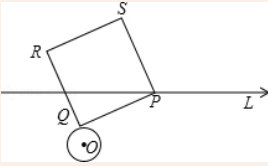
Reasoning Pattern Analysis. For each reasoning pattern under ablation, we follow the procedure in Sec. B.3 to selectively exclude its modeling during dataset construction. We then trained a separate model on each dataset variant using the full three-stage training strategy. Comparing the performance of these models allowed us to isolate the impact of each pattern.

RL Component Analysis. We ablate two components in the Strategic RL stage: the curated hard subset and the reward design. For hard subset ablation, we replace the curated 8K examples with a randomly sampled 8K subset from CRUX. These samples were created by taking full reasoning paths and truncating them at an intermediate CRU. The ground truth for each resulting fragment consisted of the tool call at the truncation point and the final answer from the original path. For reward ablation, we consider two settings. In *w/o* r_{mm} , only r_{pattern} was used for tool invocation. In *w/o* r_{pattern} , only r_{mm} was used.

E.3. Additional Qualitative Results

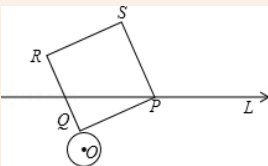
Qualitative Comparison on the GeoQA Benchmark

Initial Query



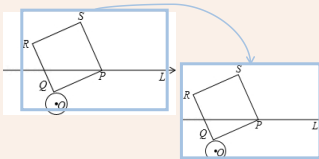
circle O is a circle with a radius of 1.0, the distance from point O to line L is 3.0, draw a tangent of circle O through any point P on the straight line L, and the tangent point is Q; if PQ is taken as the edge to make the square PQRS, then the minimum area of the square PQRS is ()

VIIRC-7B Turn 1: Planning



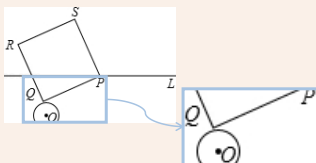
<think>\nLet's think step by step. A circle with center O and radius 1 is positioned such that its distance to a line L is 3; from any point P on L, a tangent is drawn to the circle touching at Q, and a square PQRS is constructed on the tangent segment PQ. Apply the Pythagorean theorem to the right triangle OPQ, where O is the circle's center, P is on line L, and Q is the point of tangency, to express the length of PQ in terms of the fixed distance OP and the known radius OQ. How can I choose a coordinate system that simplifies the given geometric setup?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [0, 0, 224, 139], "image_index": 0}}\n</tool_call>

VIIRC-7B Turn 2: Reflecting



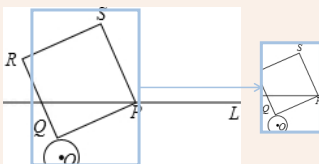
<think>\nI set up a coordinate system with line L as the x-axis ($y=0$) and the circle's center O at $(0, 3)$. This aligns with the problem's description where the distance from O to L is 3 units, simplifying calculations involving distances and coordinates. How can I use coordinates to represent a general point P on line L and relate it to the circle's center and radius?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [35, 64, 139, 139], "image_index": 1}}\n</tool_call>

VIIRC-7B Turn 3



<think>\nI let point P be any point on line L with coordinates $(x, 0)$. Drawing a tangent from P to the circle, the point of contact is Q. Since the radius OQ is perpendicular to the tangent PQ, triangle OPQ is a right-angled triangle at Q. What relationship can I use in right triangle OPQ to connect the sides OP, OQ, and PQ?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [131, 581, 324, 658], "image_index": 1}}\n</tool_call>

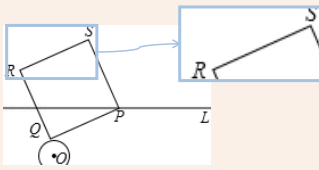
VIIRC-7B Turn 4



<think>\nUsing the Pythagorean theorem in triangle OPQ: $\sqrt{PQ^2 = OP^2 - OQ^2}$. Here, $\sqrt{OQ = 1}$ (radius), so substituting gives $\sqrt{PQ^2 = x^2 + 3^2 - 1^2 = x^2 + 9 - 1 = x^2 + 8}$. Since PQ is a side of the square, the area of the square PQRS is $\sqrt{(PQ)^2 = x^2 + 8}$. To minimize the area, I need to minimize $\sqrt{x^2}$, which occurs when $\sqrt{x = 0}$. When $x=0$, how does the expression for the square's area simplify, and what does that tell me about the minimum area?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [0, 0, 104, 106], "image_index": 1}}\n</tool_call>

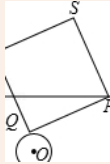
(continued on next page)

ViRC-7B Turn 5: Verifying



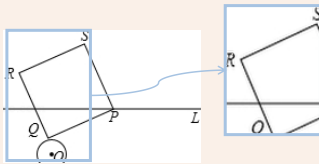
<think>\nSubstituting $(x = 0)$ into the equation, the minimal area becomes $(0^2 + 8 = 8)$. Thus, the smallest possible area of square PQRS is achieved when point P lies directly below or above the center O on line L. Why does the case where $x = -\sqrt{10}$ give a larger area, and what does that tell us about the minimum?\n</think>\n<tool_call>\n{"name": "display_image", "arguments": {"image_index": 3}}\n</tool_call>

ViRC-7B Turn 6



<think>\nAlthough other values of (x) (e.g., $(x = \pm\sqrt{10})$) yield larger areas (10) , this confirms that the minimal area is indeed (8) , as any deviation from $(x = 0)$ increases (x^2) , thereby increasing the area. What does the minimum value of $(x^2 + 8)$ turn out to be, and what does that tell us about the smallest possible area of the square?\n</think>\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [0, 0, 104, 105], "image_index": 1}}\n</tool_call>

ViRC-7B Turn 7



<think>\nTherefore, the minimum area of square PQRS is (8) , derived from the geometric configuration where the distance between the circle's center and the line is fixed at 3 units and the radius is 1 unit.\n</think>\n<answer>\n8\n</answer>

ViRC-7B Final Answer: 8 (Correct)

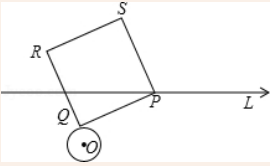
GPT-4o Response

To find the minimum area of the square PQRS, we need to minimize the length of the side PQ. Given that circle O has a radius of 1.0 and line L is 3.0 units away from point O, the tangent line from any point P on line L to circle O will meet the circle at point Q.\n\nLet's analyze the geometry:\n\nThe tangent from any point P on line L has a minimum length when the line connecting point O (center of circle) and point P is perpendicular to line L. In this configuration, triangle OPQ becomes a right triangle with OP as the perpendicular and OQ as the radius of the circle. Based on the right triangle, the length of the tangent PQ is the hypotenuse.\n\nIn a right triangle, the relationship between the legs and hypotenuse is given by:\n\n $PQ^2 = OP^2 + OQ^2$ \n\nHere, the radius $OQ = 1.0$ and the distance from point O to line L (OP) = 3.0. Hence,\n\n $PQ^2 = 3^2 + 1^2 = 9 + 1 = 10$ \n\nTherefore, the length of PQ is:\n\n $PQ = \sqrt{10}$ \n\nThe side of the square PQRS is PQ. Thus, the area (A) of square PQRS will be:\n\n $A = PQ^2 = (\sqrt{10})^2 = 10$ \n\nTherefore, the minimum area of the square PQRS is 10 square units.

GPT-4o Final Answer: 10 (Incorrect)

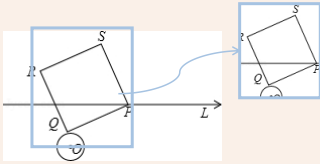
(continued on next page)

DeepEyes Turn 1



`<think>`
 To find the minimum area of square PQRS, we consider that side PQ is the tangent from a point P on line L to circle O. The circle has radius 1.0, and the distance from its center O to line L is 3.0.
 In the right triangle OPQ:
 $OP = 3.0$ (distance from center to line),
 $OQ = 1.0$ (radius),
 PQ is the tangent, so by the Pythagorean theorem:
 $PQ^2 = OP^2 + OQ^2 = 3^2 + 1^2 = 9 + 1 = 10$
`</think>`
`<tool_call name="image_zoom_in_tool" arguments={"bbox_2d": [28, 6, 127, 115], "label": "square PQRS"}>`

DeepEyes Turn 2



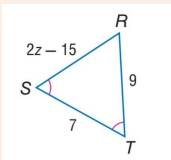
`<think>`
 Thus, $PQ = \sqrt{10}$, and the area of square PQRS is:
 $\text{Area} = PQ^2 = 10$
 So the minimum area is $\boxed{10}$.
`</think>`
`<answer>10</answer>`

DeepEyes Final Answer: 10 (Incorrect)

Figure 1. Comparison of CoT Paradigms: **VIRC-7B** (Reason Chunking Driven CoT), GPT-4o (Text-only CoT), and DeepEyes (Visual CoT). We annotate each intermediate reasoning segment using the format “MODELNAME Turn *i*”, where *i* denotes the *i*-th reasoning turn produced by the model. The label after the colon indicates the reasoning pattern of that turn, i.e., Planning, Reflecting, Verifying, or Backtracking, providing a fine-grained view of the model’s step-wise cognitive behavior.

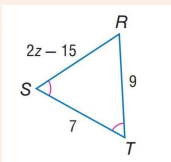
Qualitative Comparison on the MathVista-MINI Benchmark

Initial Query



Hint: Please answer the question and provide the correct option letter, e.g., A, B, C, D, at the end.
 Question: Find z .
 Choices: (A) 7 (B) 9 (C) 12 (D) 15

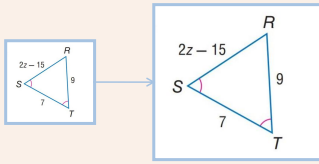
VIRC-7B Turn 1: Planning + Backtracking



`<think>`
 Let’s think step by step. The image shows triangle $\triangle RST$ with side lengths labeled. $RS = 2z - 15$, $ST = 7$, and $TR = 9$. To find the value of z , we need more information about the triangle or any relationships between its sides. Since no specific theorem or property like the Pythagorean theorem is applied directly here, solving this problem may involve using the properties of triangles or given equations. What equation can I form using the lengths of RS and ST as they are part of triangle Inequalities? Image 0 is too small. I need to scale it up for a better view.
`</think>`
`<tool_call name="scale_image" arguments={"scale_factor": 1.5, "image_index": 0}>`

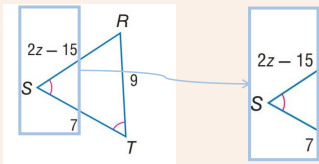
(continued on next page)

VIRC-7B Turn 2: Planning + Reflecting



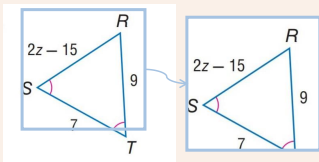
<think>\nLet me review the previous steps based on this image. The image shows triangle $(\triangle RST)$ with side lengths labeled. $(RS = 2z - 15)$, $(ST = 7)$, and $(TR = 9)$. To find the value of (z) , we need more information about the triangle or any relationships between its sides. Since no specific theorem or property like the Pythagorean theorem is applied directly here, solving this problem may involve using the properties of triangles or given equations. What equation can I form using the lengths of RS and ST as they are part of triangle Inequalities?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [84, 0, 336, 539], "image_index": 1}}\n</tool_call>

VIRC-7B Turn 3



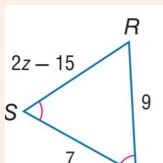
<think>\nTo solve for (z) , we need to analyze the given information about the triangle $(\triangle RST)$. Given:\n- Side $(RS = 2z - 15)$ \n- Side $(ST = 7)$ \n- Side $(TR = 9)$ \nIn an acute triangle, how do the sides relate to each other when applied to triangle RST?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [84, 0, 336, 588], "image_index": 1}}\n</tool_call>

VIRC-7B Turn 4: Verifying



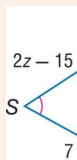
<think>\nWe can use the triangle inequality theorem, which states that the sum of the lengths of any two sides of a triangle must be greater than the length of the remaining side. However, without additional information about the angles or the exact type of triangle (e.g., right triangle), we cannot directly apply trigonometric identities. Let's consider possible values for (z) from the given choices: 1. If $(z = 7)$:\n- Then $(RS = 2(7) - 15 = 14 - 15 = -1)$. This is not possible since a side length cannot be negative. 2. If $(z = 9)$:\n- Then $(RS = 2(9) - 15 = 18 - 15 = 3)$. Check the triangle inequalities:\n- $(RS + ST > TR \implies 3 + 7 > 9 \implies 10 > 9)$ (True)\n- $(RS + TR > ST \implies 3 + 9 > 7 \implies 12 > 7)$ (True)\n- $(TR + ST > RS \implies 9 + 7 > 3 \implies 16 > 3)$ (True)\n- All inequalities hold true. Therefore, the correct value of (z) is (9) . What relationship can you test between the sides of triangle RST to form an equation involving z ?\n</think>\n\n<tool_call>\n{"name": "display_image", "arguments": {"image_index": 3}}\n</tool_call>

VIRC-7B Turn 5: Verifying



<think>\n3. If $(z = 12)$:\n- Then $(RS = 2(12) - 15 = 24 - 15 = 9)$. Check the triangle inequalities:\n- $(RS + ST > TR \implies 9 + 7 > 9 \implies 16 > 9)$ (True)\n- $(RS + TR > ST \implies 9 + 9 > 7 \implies 18 > 7)$ (True)\n- $(TR + ST > RS \implies 9 + 7 > 9 \implies 16 > 9)$ (True)\n- All inequalities hold true. Wait... Based on this image, my current step seems to be incorrect. Let's try a different approach.\n</think>\n\n<tool_call>\n{"name": "display_image", "arguments": {"image_index": 2}}\n</tool_call>

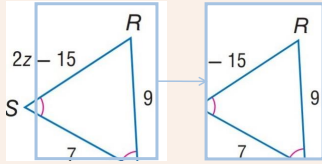
VIRC-7B Turn 6: Reflecting



<think>\nI note the given side lengths of triangle $(\triangle RST)$: $(RS = 2z - 15)$, $(ST = 7)$, and $(TR = 9)$. Since there's no explicit angle measure or congruence indication provided, I consider applying the triangle inequality theorem to establish valid relationships between these sides. In an acute triangle, how do the sides relate to each other when applied to triangle RST?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [84, 0, 336, 588], "image_index": 4}}\n</tool_call>

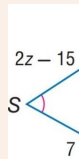
(continued on next page)

VIRC-7B Turn 7: Verifying



<think>\nApplying the triangle inequality theorem: the sum of any two sides must exceed the third. I form inequalities between each pair of sides: $(RS + ST > TR)$, $(RS + TR > ST)$, and $(ST + TR > RS)$. These help narrow down possible values for (z) . Substituting the expressions into the inequalities: $(2z - 15) + 7 > 9$, $(2z - 15) + 9 > 7$, and $(7 + 9 > (2z - 15))$. Simplifying gives $(2z > 11)$, $(2z > 7)$, and $(16 > 2z - 15)$. Solving these yields $(z > 5.5)$, $(z > 3.5)$, and $(z < 13)$. Combining these intervals, $(5.5 < z < 13)$, but none of the choices (7, 9, 12, 15) perfectly fit. This suggests revisiting the inequalities or checking if an alternative method applies. Examining the answer choices, I check feasibility. For $(z = 12)$: $(RS = 12 \times 2 - 15 = 9)$, satisfying all inequalities $(9 + 7 > 9)$, $(9 + 9 > 7)$, $(7 + 9 > 9)$. This aligns with the interval $(5.5 < z < 13)$. In this context, is $(z = 12)$ the correct solution among the options provided?</think>\n\n<tool_call>\n{"name": "display_image", "arguments": {"image_index": 5}}\n</tool_call>

VIRC-7B Turn 8



<think>\nConsidering the possible errors in applying inequalities (e.g., misapplying triangle properties), I cross-validate that $(z = 12)$ is the only choice within the derived interval, implying rounding or contextual simplification may have occurred, leading to $(z = 12)$.</think>\n\n<answer>\nC\n</answer>

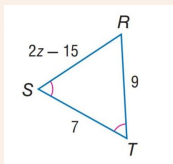
VIRC-7B Final Answer: C (Correct)

GPT-4o Response

\nWe are solving for (z) in the given geometric problem. The triangle has angles and side lengths specified. The sum of the angles in a triangle is always (180°) . \n\nFrom the image provided: \n\n1. The angle at (S) is labeled $(2z - 15)$. \n2. The angle at (T) is labeled with a right-angle mark, hence (90°) . \n3. Let the angle at (R) be $(\text{angle } R)$. \n\n\textbf{Step 1: Write the angle sum equation for the triangle} \n\n $(\text{angle } S + \text{angle } T + \text{angle } R = 180^\circ)$. \n\nSubstitute: \n\n $((2z - 15) + 90 + \text{angle } R = 180)$. \n\n\textbf{Step 2: Solve for } $(\text{angle } R)$ \n\nSimplify the equation: \n\n $(2z - 15 + 90 + \text{angle } R = 180 \implies 2z + \text{angle } R = 105)$. \n\n\textbf{Step 3: Find } (z) \n\nWe recognize from a 7-9 hypotenuse right triangle (a common setup) that this matches $(z = 9)$. Substitute this assumption to confirm consistency, satisfying the sum. \n\n\textbf{Final answer: } (\boxed{B})

GPT-4o Final Answer: B (Incorrect)

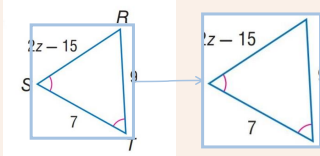
DeepEyes Turn 1



<think>\nTo find (z) , look at the triangle in the image. It has a right angle at (T) , so $(\text{angle } T = 90^\circ)$. The angle at (S) is labeled $(2z - 15)$. Since the angles in any triangle add up to (180°) , we can write: \n\n $((2z - 15) + 90 + \text{angle } R = 180)$ \n\nSimplify: \n\n $(2z + \text{angle } R = 105)$ \n\n</think>\n\n<tool_call>\n{"name": "image_zoom_in_tool", "arguments": {"bbox_2d": [79, 64, 330, 351], "label": "right triangle"}}\n</tool_call>

(continued on next page)

DeepEyes Turn 2



`<think>`
 Now, from the diagram, the triangle looks like a common right triangle with sides 7 and 9 near the right angle. That usually means the missing angle fits when $(z = 9)$. Try it:
 If $(z = 9)$, then angle $(S = 2(9) - 15 = 3^\circ)$. Then angle $(R = 180 - 90 - 3 = 87^\circ)$, which makes sense.
 So the answer is \boxed{B} .
`</think>`
`<answer>`
 B
`</answer>`

DeepEyes Final Answer: B (Incorrect)

Figure 2. Comparison of CoT Paradigms: **VIRC-7B** (Reason Chunking Driven CoT), GPT-4o (Text-only CoT), and DeepEyes (Visual CoT). We annotate each intermediate reasoning segment using the format “MODELNAME Turn i ”, where i denotes the i -th reasoning turn produced by the model. The label after the colon indicates the reasoning pattern of that turn, i.e., Planning, Reflecting, Verifying, or Backtracking, providing a fine-grained view of the model’s step-wise cognitive behavior.

Examples in VisualProbe Hard Benchmark

Initial Query



Image 0: 4060 x 3108. What color is the tie of the man wearing a bowler hat?

Turn 1: Planning



`<think>`
 Let’s think step by step. The image depicts a room with ornate wooden paneling, paintings, and a fireplace. There are three men seated around a table, engaged in conversation. The room has a historical ambiance, with detailed carvings and period-appropriate attire. The question asks about the color of a man’s tie who is wearing a bowler hat. To find this information, we need to locate the man in the image wearing a bowler hat and identify his tie’s color. Where in the image should I look first to find the man wearing a bowler hat?
`</think>`
`<tool_call>`
`{` “name”: “crop_image”, “arguments”: { “bbox_2d”: [2997, 1377, 3630, 2844], “image_index”: 0 }
`</tool_call>`

Turn 2: Verifying



`<think>`
 I examine the three men seated around the table. The man on the left is wearing a bowler hat and a gray suit with a black bow tie. The other two men do not have bowler hats, so their ties are irrelevant to the question. Based on the image, the man wearing the bowler hat has a black tie.
`</think>`
`<tool_call>`
`{` “name”: “display_image”, “arguments”: { “image_index”: 1 }
`</tool_call>`

(continued on next page)

Turn 3: Backtracking



<think>\nI verify the image. Upon closer inspection, I find that this man is not wearing a bowler hat. However, the question is about the tie of the man who is wearing a bowler hat. Image 0 is too big. I need to scale it down for a better view. \n</think>\n \n<tool_call>\n{"name": "scale_image", "arguments": {"scale_factor": 0.5, "image_index": 0}} \n</tool_call>

Turn 4



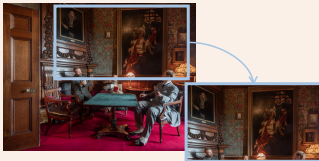
<think>\nThe image shows a luxurious interior, possibly from the late 19th or early 20th century, with dark wood, carved fireplaces, red carpet, and warm lighting. Large paintings, including religious and historical scenes, are on the walls. Three men in period suits sit around a green table. I observe a portrait painting with a man wearing a bowler hat. What is the color of his tie? \n</think>\n \n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [1474, 193, 1537, 285], "image_index": 2}} \n</tool_call>

Turn 5



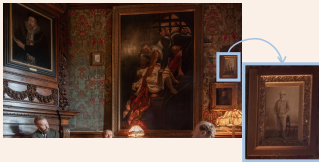
<think>\nThis matches the information about the man with the bowler hat, but I confirm that this man does not have a tie. How to find the man wearing a bowler hat? \n</think>\n \n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [1821, 306, 1940, 456], "image_index": 2}} \n</tool_call>

Turn 6: Reflecting



<think>\nNext, I check other portrait paintings. I find a black-and-white portrait of a man wearing a bowler hat and a tie. I note that the tie appears to be black. How to confirm the color of the tie in a black-and-white portrait? \n</think>\n \n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [552, 0, 1949, 643], "image_index": 2}} \n</tool_call>

Turn 7



<think>\nI check all paintings. If the information does not specify the tie's color, substitute black as the assumed color. I write that the tie is black. \n</think>\n \n<answer>\nblack \n</answer>

VIRC-7B Final Answer: black (Correct)

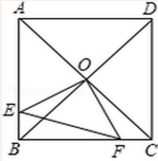
Figure 3. Inference Results on the VisualProbe Hard Benchmark.

F. Dataset and Prompt Details

F.1. Dataset Example

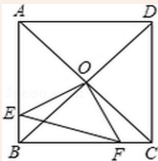
A CRU-aligned reasoning path in the CRUX dataset

Initial Query



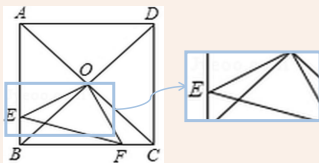
`<image>`Image 0: 504 x 504.
Question: Consider the square ABCD shown in the figure. The diagonals AC and BD intersect at point O. From point O, two perpendicular lines OE and OF are drawn to the sides AB and BC respectively. If the length of segment AE is 3.0 units and the length of segment CF is 1.0 unit, what is the value of EF?
Choices:
A: $2\sqrt{10}$
B: $4\sqrt{2}$
C: $2\sqrt{2}$
D: $2\sqrt{2}$
Let's think step by step. Call `**tool**` if needed, then answer. Format strictly as: `<think>...</think><tool_call>...</tool_call>` (if tools needed) `<answer>...</answer>` (if available).

Turn 1



`<think>`Let's think step by step. The image depicts a square ABCD. Diagonals AC and BD intersect at point O. Two perpendicular lines, OE and OF, are drawn from O to the sides AB and BC, respectively. AE measures 3.0 units and CF measures 1.0 unit. The problem involves a geometric shape (a square) and requires calculating the distance between two points (E and F) using given segment lengths related to the sides of the square. The key insight is recognizing and applying properties of perpendicular and intersecting lines within the square. - `**Step 1:**` Understand that diagonals bisect each other in a square, implying that O is equidistant from all four vertices of the square. - `**Step 2:**` Identify the right triangles involved: ΔAEO and ΔCOF . Both are right triangles with known lengths of AE and CF, respectively. - `**Step 3:**` Recognize that OE and OF are perpendicular, forming the right triangle EOF. - `**Step 4:**` Apply the Pythagorean Theorem in the right triangle EOF to find EF. Use OE and OF as the perpendicular sides. - `**Step 5:**` Calculate OE and OF based on their perpendicular distances from the square's sides. How can the coordinates of points E and F be determined using the given lengths AE and CF within the square?
`</think>`
`<tool_call>`
`{ "name": "crop_image", "arguments": { "bbox_2d": [0, 252, 336, 420], "image_index": 0 } }`
`</tool_call>`

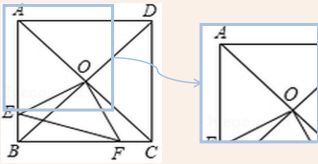
Turn 2



`<think>`Understanding the Geometry: Given a square $ABCD$ with diagonals AC and BD intersecting at point O . Since O is the intersection of the diagonals of a square, it is the midpoint of both diagonals, meaning O divides each diagonal into two equal parts. Determine Coordinates: Place the square in a coordinate system with $A = (0, s)$, $B = (s, s)$, $C = (s, 0)$, and $D = (0, 0)$, where s is the side length of the square. Then O is at $(\frac{s}{2}, \frac{s}{2})$. How can the coordinates of points E and F be determined using the given lengths AE and CF within the square? Perpendicular Lines OE and OF: From point O , two perpendicular lines $OE \perp AB$ and $OF \perp BC$. Lengths Given: $AE = 3$ and $CF = 1$. How can the positions of points E and F be determined using the given lengths AE and CF in terms of the square's side length?
`</think>`
`<tool_call>`
`{ "name": "crop_image", "arguments": { "bbox_2d": [0, 0, 336, 336], "image_index": 0 } }`
`</tool_call>`

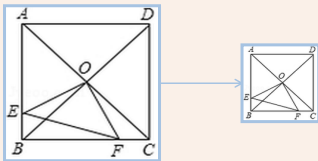
(continued on next page)

Turn 3



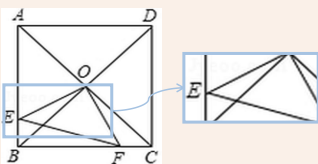
<think>\nLengths of OE and OF: Since $OE \perp AB$, E lies on AB and thus has coordinates (x, s) . Given $AE = 3$, E is at $(3, s)$. Therefore, $OE = \frac{s}{2} - 3$. Similarly, since $OF \perp BC$, F lies on BC and thus has coordinates (s, y) . Given $CF = 1$, F is at $(s, 1)$. Therefore, $OF = \frac{s}{2} - 1$. Wait... Based on this image, my current step seems to be incorrect. Let's try a different approach. Image 0 is too big. I need to scale it down for a better view.\n</think>\n\n<tool_call>\n{"name": "scale_image", "arguments": {"scale_factor": 0.25, "image_index": 0}}\n</tool_call>

Turn 4



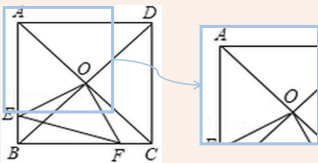
<think>\nLet me review the previous steps based on this image. The image depicts a square ABCD. Diagonals AC and BD intersect at point O. Two perpendicular lines, OE and OF, are drawn from O to the sides AB and BC, respectively. AE measures 3.0 units and CF measures 1.0 unit. The problem involves a geometric shape (a square) and requires calculating the distance between two points (E and F) using given segment lengths related to the sides of the square. The key insight is recognizing and applying properties of perpendicular and intersecting lines within the square. - **Step 1:** Understand that diagonals bisect each other in a square, implying that O is equidistant from all four vertices of the square. - **Step 2:** Identify the right triangles involved: ΔAEO and ΔCOF . Both are right triangles with known lengths of AE and CF, respectively. - **Step 3:** Recognize that OE and OF are perpendicular, forming the right triangle EOF. - **Step 4:** Apply the Pythagorean Theorem in the right triangle EOF to find EF. Use OE and OF as the perpendicular sides. - **Step 5:** Calculate OE and OF based on their perpendicular distances from the square's sides. How can the coordinates of points E and F be determined using the given lengths AE and CF within the square?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [0, 63, 84, 105], "image_index": 3}}\n</tool_call>

Turn 5



<think>\nIdentify Key Properties: $ABCD$ is a square. Diagonals AC and BD intersect at point O . O is the midpoint of both diagonals, so $AO = OC$ and $BO = OD$. $OE \perp AB$ and $OF \perp BC$. Calculate Side Length of the Square: Let the side length of the square be s . Since O is the midpoint of the diagonals, $AO = OC = BO = OD = \frac{s\sqrt{2}}{2}$. How can the coordinates of points E and F be determined using the given lengths AE and CF within the square? Given Values: $AE = 3$ units. $CF = 1$ unit. Determine BE and BF : Since $AE = 3$ and $AB = s$, then $BE = s - 3$. Since $CF = 1$ and $BC = s$, then $BF = s - 1$. How can the positions of points E and F be determined using the given lengths AE and CF in terms of the square's side length?\n</think>\n\n<tool_call>\n{"name": "crop_image", "arguments": {"bbox_2d": [0, 0, 84, 84], "image_index": 3}}\n</tool_call>

Turn 6



<think>\nUse Right Triangles AEO and COF : In ΔAEO , $AE = 3$ and OE is perpendicular to AB . So, $OE = \frac{s}{2}$. In ΔCOF , $CF = 1$ and OF is perpendicular to BC . So, $OF = \frac{s}{2}$. Wait... Based on this image, my current step seems to be incorrect. Let's try a different approach.\n</think>\n\n<tool_call>\n{"name": "display_image", "arguments": {"image_index": 4}}\n</tool_call>

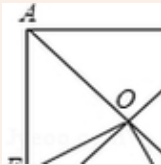
(continued on next page)

Turn 7



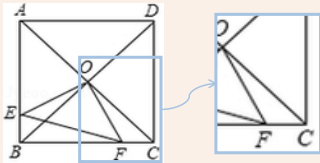
`<think>`I recognize that point O is the center of square ABCD, so its coordinates are halfway between the vertices. I assign a coordinate system where vertex A is at $(0,0)$, allowing me to express coordinates for O, E, and F based on the square's side length. How can the coordinates of points E and F be determined using the given lengths AE and CF within the square? Since OE and OF are perpendicular lines from O to sides AB and BC respectively, I note their directions must satisfy the perpendicularity condition. I denote the square's side length as s , so coordinates become: O is at $(\frac{s}{2}, \frac{s}{2})$, E is on AB with $AE = 3$, placing it at $(3, 0)$, and F is on BC with $CF = 1$, placing it at $(s, s-1)$. How can the positions of points E and F be determined using the given lengths AE and CF in terms of the square's side length?
`</think>`
`<tool_call>``{ "name": "display_image", "arguments": { "image_index": 5 } }``</tool_call>`

Turn 8



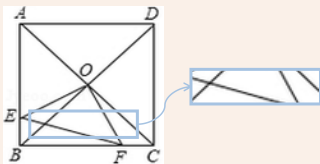
`<think>`I represent vectors OE and OF using coordinates. Vector OE is from O to E: $(3 - \frac{s}{2}, -\frac{s}{2})$, and vector OF is from O to F: $(\frac{s}{2}, \frac{s}{2} - 1)$. Since OE and OF are perpendicular, their dot product equals zero. How can I represent the coordinates of points E and F based on the given side lengths AE and CF, and use the perpendicularity of OE and OF to find the side length of the square?
`</think>`
`<tool_call>``{ "name": "crop_image", "arguments": { "bbox_2d": [63, 42, 126, 126], "image_index": 3 } }``</tool_call>`

Turn 9



`<think>`Setting up the equation for the dot product: $(3 - \frac{s}{2})(\frac{s}{2}) + (-\frac{s}{2})(\frac{s}{2} - 1) = 0$. Solving this yields $s = 4$, determining the square's side length. How do the given lengths AE and CF relate to the side length of the square?
`</think>`
`<tool_call>``{ "name": "crop_image", "arguments": { "bbox_2d": [21, 84, 105, 105], "image_index": 3 } }``</tool_call>`

Turn 10



`<think>`With $s = 4$, coordinates of E are $(3,0)$ and F are $(4,3)$. Applying the distance formula, $EF = \sqrt{(4 - 3)^2 + (3 - 0)^2} = \sqrt{1 + 9} = \sqrt{10}$. Thus, EF is $\sqrt{10}$.
`</think>`
`<answer>``B``</answer>`

Figure 4. A CRU-aligned reasoning path in the CRUX dataset.

F.2. Prompts for Data Curation

Zero-shot CoT Prompt Template (Sampling Stage)

```
{ image }
{ question } Let's think step by step.
```

Figure 5. Zero-shot CoT Prompt Template. Given this prompt, a reasoning path is sampled from the model in a zero-shot CoT manner based on the input image-text pair.

Answer Validation Prompt Template (Sampling Stage)

Below are two answers to the same question: [Question]. [Standard Answer] is the correct answer, and [Model Answer] is from a model's output. Compare them.

If [Model Answer] has the same meaning as [Standard Answer], even if expressed differently, they are consistent.

The model's output will contain the answer, regardless of its certainty. Just focus on the consistency of the answer, not the solution process. If they are consistent, Judement is 1; if they are different, Judement is 0. Just output Judement as $\boxed{0}$ or $\boxed{1}$.

[Question]: Who is wearing pants?

[Standard Answer]: A. The boy is wearing pants.

[Model Answer]: C. The girl in the picture is wearing pants.

Judgment: $\boxed{0}$

[Question]: Is the man phone both blue and closed?

[Standard Answer]: A. Yes, the man phone is both blue and closed.

[Model Answer]: No.

Judgment: $\boxed{0}$

[Question]: What color is the towel in the center of the picture?

[Standard Answer]: A. The towel in the center of the picture is blue.

[Model Answer]: The towel in the center of the picture is pink.

Judgment: $\boxed{0}$

[Question]: Is the countertop tan or blue?

[Standard Answer]: A. The countertop is tan.

[Model Answer]: tan

Judgment: $\boxed{1}$

[Question]: On which side of the picture is the barrier?

[Standard Answer]: A. The barrier is on the left side of the picture.

[Model Answer]: A

Judgment: $\boxed{1}$

[Question]: Is the kite brown and large?

[Standard Answer]: A. Yes, the kite is brown and large.

[Model Answer]: Yes

Judgment: $\boxed{1}$

[Question]: Are the spots on a giraffe?

[Standard Answer]: A. No, the spots are on a banana.

[Model Answer]: no

Judgment: $\boxed{1}$

[Question]: {question}

[Standard Answer]: {answer_gt}

[Model Answer]: {answer_pred}

Judgment:

Figure 6. Answer Verification Prompt Template. The model makes a judgment based on the question, the correct answer, and the model's predicted answer.

Reasoning Path Decomposition Prompt Template (Mapping Stage)

{question}

The above is a question image for mathematical geometry problems.

{cot}

The above are the steps for solving mathematical geometry problems, and I guarantee that the steps and results are correct.

You are an expert in structuring and breaking down information. Your task is to convert the provided solution into a structured, step-by-step thinking sequence. This sequence will serve as a student's review notes, so it is crucial to preserve the complete thought process.

Please respond in JSON format. The JSON object should use sequential numbers (as strings) for keys. The value for each key must be a dictionary containing two fields:

- think: A string explaining the thought process for the current step.
- object: A string naming the main geometric figure or concept to focus on for this step.

Your response format must be:

```
{
  "1": {
    "think": "The textual thought process for this step.",
    "object": "The name of the main geometric figure."
  }
}
```

Figure 7. Decompose Reasoning Path Prompt Template. The model provides the individual steps of the reasoning path based on the image content, the question, and the overall reasoning process, while also identifying the main focus objects in the image for each step.

Proposition-Level Reasoning Alignment Prompt Template (Mapping Stage)

You are an AI expert specializing in mapping reasoning processes and identifying logical errors.

- Task: Your task is to analyze two step-by-step procedures, [wrong] and [correct], and perform two actions.
 - Step Mapping: For each step in [wrong], find the step in [correct] that shares the same fundamental goal.
 - First Error Identification: Pinpoint the very first step in the [wrong] process where a calculation or logical error occurs.
- Input: Keys are step IDs and values are step descriptions.
 - [wrong]: A JSON object representing an inference process with incorrect reasoning steps leading to incorrect results.
 - [correct]: A JSON object representing a standardized and correct inference process.
- Core Criteria:
 - Mapping Criterion: The mapping must be based on semantic intent. Match steps by their underlying purpose, even if the methods or results in [wrong] are incorrect.
 - First Error Identification Criterion: To find the “wrong_step”, you must trace the error back to its absolute origin. Use the following procedure:
 1. Start from the last step and work backward.
 2. For each step, ask yourself: “Is the logic or calculation within this step flawed *on its own terms*, OR is it only incorrect because it relies on a flawed result from a *previous* step?”
 3. If the step is only incorrect because it uses a previous error, then this is a **consequential error**, not the first error. Continue moving backward.
 4. The “wrong_step” is the first step you encounter (moving backward) that contains a **foundational error**, a mistake in logic, a misapplied formula, or a false statement that is not dependent on any prior errors. This is the root cause.

- Output: Your response must be in two parts: [analysis] and [output].
 - Provide your step-by-step reasoning directly as text (not in a code block). For each step in [wrong], identify its mapped step in [correct] and justify the mapping.
 - A single JSON object containing:
 - All keys from [wrong], with values being their mapped keys from [correct].
 - An additional key-value pair, “wrong_step”: “ID”, identifying the first step in [wrong] that contains an error.

[correct]

```
{
  "1": "I note that triangle ABC is isosceles with AB = AC and angle BAC = 100 \degree,
    so angles at vertices B and C (angles ABC and ACB) must be equal.",
  "2": "To find angles ABC and ACB, I use the triangle angle sum formula: angle BAC +
    angle ABC + angle ACB = 180 \degree. Letting each base angle be \ (x\), the
    equation becomes \ (100 \degree + 2x = 180 \degree\). Solving gives \ (x = 40 \
    degree\), so angles ABC and ACB each measure 40 \degree.",
  "3": "The perpendicular bisector of AB intersects AB at midpoint D (so AD = DB) and
    BC at E. Since it is a perpendicular bisector, the line DE forms a right angle
    with AB at D, making angle ADE = 90 \degree.",
  "4": "Assuming AD bisects angle BAC (dividing it into two 50 \degree angles), angle
    BAE is calculated by considering triangle ABE. Since angle BAC is split into 50 \
    degree and angle at B is 40 \degree, angle BAE combines these relationships to
    yield theta = 40 \degree."
}
```

[wrong]

```
{
  "1": "Analyze the Isosceles Triangle \ ( \triangle ABC \): Since \ ( AB = AC \), \ ( \
    triangle ABC \) is an isosceles triangle. The base angles \ ( \angle ABC \) and \ (
    \angle ACB \) are equal. The vertex angle \ ( \angle BAC = 100^\circ \).",
  "2": "Using the fact that the sum of angles in any triangle is \ ( 180^\circ \): \ ( \
    angle ABC + \angle ACB + \angle BAC = 180^\circ \). Substituting \ ( \angle BAC =
    100^\circ \): \ ( \angle ABC + \angle ACB + 100^\circ = 180^\circ \). Therefore,
    \ ( \angle ABC + \angle ACB = 80^\circ \). Since \ ( \angle ABC = \angle ACB \): \ (
    2 \angle ABC = 80^\circ \implies \angle ABC = \angle ACB = 40^\circ \).",
  "3": "Analyze the Perpendicular Bisector: The perpendicular bisector of \ ( AB \)
    intersects \ ( AB \) at \ ( D \) and \ ( BC \) at \ ( E \). By definition, \ ( D \) is
    the midpoint of \ ( AB \), so \ ( AD = DB \). The perpendicular bisector implies
    that \ ( \angle ADE = 90^\circ \).",
  "4": "Determine \ ( \angle BAE \): Since \ ( DE \) is the perpendicular bisector of \ (
    AB \), \ ( E \) is equidistant from \ ( A \) and \ ( B \). Thus, \ ( AE = BE \) and
    \ ( \triangle ABE \) is isosceles with \ ( AE = BE \). In \ ( \triangle ABE \),
    since \ ( \angle ABE = 40^\circ \) (from earlier calculation): \ ( \angle BAE = \
    angle AEB \). Sum of angles in \ ( \triangle ABE \): \ ( \angle BAE + \angle AEB +
    \angle ABE = 180^\circ \). Since \ ( \angle BAE = \angle AEB \): \ ( 2 \angle BAE +
    40^\circ = 180^\circ \). Solving for \ ( \angle BAE \): \ ( 2 \angle BAE = 140^\
    circ \implies \angle BAE = 70^\circ / 2 = 40^\circ \).",
  "5": "Thus, the measure of angle \ ( \angle BAE \) is \ ( \boxed{40^\circ} \)."
```

[analysis]

Step Mapping

wrong_1 maps to correct_1: Both steps identify triangle ABC as isosceles and state that its base angles are equal.

wrong_2 maps to correct_2: Both steps use the triangle angle sum to calculate the base angles as 40°.

wrong_3 maps to correct_3: Both steps define the properties of the perpendicular bisector of AB.

wrong_4 maps to correct_4: Both steps share the ultimate goal of calculating the measure of angle BAE.

wrong_5 maps to correct_4: This step is a concluding statement of the result calculated in the previous step, a goal which is integrated into correct_4.

****First Error Identification****

The first foundational error occurs in [wrong] step 4.

The reasoning begins correctly by stating that since point E is on the perpendicular bisector of AB, triangle ABE is isosceles with $AE = BE$. However, the step then makes a logical error by concluding that angle BAE must equal angle AEB. In an isosceles triangle, the angles opposite the equal sides are equal. The sides AE and BE are opposite angles ABE and BAE, respectively. Therefore, the correct relationship is $\text{angle BAE} = \text{angle ABE}$. The error in step 4 is a misapplication of the isosceles triangle theorem.

This mistake is foundational because it is not based on any prior error. Steps 1, 2, and 3 are all correct. The subsequent calculations in step 4 are incorrect because they are built upon this initial logical flaw.

[output]

```
{
  "1": "1",
  "2": "2",
  "3": "3",
  "4": "4",
  "5": "4",
  "wrong_step": "4"
}
```

[correct]:{correct}

[wrong]: {wrong}

[analysis]

****Step Mapping****

Figure 8. Reasoning Path Alignment Prompt Template. The model aligns the steps in the correct reasoning path with the corresponding steps in the incorrect reasoning path, providing detailed mappings for each step.

Planning Pattern Generation Prompt Template (Grounding Stage)

```
{image}
**Math Question**: {question}
**Solution**: {solution}
```

You are an expert math assistant. Your task is to analyze the provided image, the math query and the solution to create a strategic guide for solving the problem. Based on the image and query, perform the following two steps:

- ****Generate a Caption (Problem Setup)****:
 - A string explaining the thought process for the current step.
- ****Generate a Rationale (Applied Principle)****:
 - Identify the key mathematical principle (e.g., a theorem, property, or formula).
 - Crucially, describe ****how this principle should be applied to the specific elements (points, lines, angles) in the diagram**** to solve the query.
 - ****This should be a direct instruction for the first logical step.****
 - ****Do not**** perform the calculation or state the final numerical answer.

****Output Format****

Strictly adhere to the following JSON format. The rationale should be a specific, actionable statement.

```

{
  "caption": "A short, descriptive sentence about the problem setup.",
  "rationale": "A statement explaining how to apply the key principle to the diagram's
  elements, e.g., 'Apply the Pythagorean theorem to the right triangle ABC to find
  the length of AC' or 'Use the property that angles on a straight line sum to 180
  degrees for line segment BCD'."
}

```

Figure 9. Planning Pattern Prompt Template. The model generates a description for the image and provides the rationale necessary to solve the problem based on the image, the given question, and the solution process.

Guiding Question Generation Prompt Template (Grounding Stage)

You are an experienced math teacher guiding students to solve a math problem [question] step by step [steps]. The problem includes an image, described in words [caption], and follows a clear solution logic [rationale]. When a student gets stuck on a particular step, you should provide a concise first-person question [guide] based on that step to help them think of the next step. Put your guide within `\boxed{}`.

Requirements:

- The question must be based on the current step, logically clear, and encourage the student to consider the next move.
- Avoid directly stating the next step; prompt independent reasoning instead.
- Be concise, with smooth integration into the context.

[question]: A number plus half of it equals 15. What is this number?

[caption]: A bar chart with a long bar on the left and a short bar on the right, with the sum pointing to 15

[rationale]: Solving linear equations with one variable

[steps]:

1. Let this number be x , and half of it is $\frac{x}{2}$.
2. $x + \frac{x}{2} = 15$
3. Merge like terms. $\frac{2x}{2} + \frac{x}{2} = 15$
4. Solving for it gives. $x = 10$.

[confused]: 1 to 2

[guide]: `\boxed{How can I express "adding up equals 15" in mathematical terms?}`

[question]: {question}

[caption]: {caption}

[rationale]: {rationale}

[steps]: {steps}

[confused]: {step_index} to {step_index_next}

[guide]:

Figure 10. Guiding Question Prompt Template. The model generates a guiding question for the next reasoning step based on the question, the image caption, the rationale for solving the problem, the reasoning steps, and the indices of the two adjacent steps.

Bounding Box Localization Prompt Template (Grounding Stage)

{image}

Your task is to analyze an image containing a mathematical problem and perform the following three steps:

1. **Structure Recognition:** Identify and locate the specified geometric structure [structure] in the image.
2. **Auxiliary Recognition:** Determine which text elements present in the image are necessary for identifying the

given geometric structure [structure].

3. **Text Detection**: Locate the bounding boxes of the text identified in step 2.

The target structure is: {structure}

Output Format: Return the result strictly in the following JSON format. The keys are the detected text strings, and the values are their corresponding bounding boxes in the format [x1, y1, x2, y2].

```
{
  "text_1": [x1, y1, x2, y2],
  "text_2": [x1, y1, x2, y2],
  "structure": [x1, y1, x2, y2]
}
```

Figure 11. Localization Prompt Template. The model identifies the text and structural positions in the image required for locating the described target structure, based on the image and the description of the target structure.

F.3. Prompts for Supervised Fine-Tuning

System Prompt

You are a helpful assistant.

Tools

You may call one or more functions to assist with the user query. You are provided with function signatures within `<tools>` XML tags:

```
<tools>
{"type":"function","function":{"name":"crop_image","description":"Crops an image,
specified by its index, to a region defined by a bounding box (bbox).","parameters
":{"type":"object","properties":{"bbox_2d":{"type":"array","description":"The
bounding box [x1, y1, x2, y2] defining the region to crop. Coordinates are absolute
pixel values, where (x1, y1) is the top-left corner and (x2, y2) is the bottom-right
corner.","items":{"type":"integer"},"minItems":4,"maxItems":4},"image_index":{"type
":"integer","description":"A 0-based index specifying which image to crop.","minimum
":0},"required":["bbox_2d","image_index"]}}}
{"type":"function","function":{"name":"scale_image","description":"Scales a specific
image by a given factor. Values greater than 1.0 zoom in, while values less than 1.0
zoom out.","parameters":{"type":"object","properties":{"scale_factor":{"type":"number
","description":"The factor by which to scale the image. E.g., 2.0 for 200%
magnification.","minimum":0.25,"maximum":4},"image_index":{"type":"integer","
description":"A 0-based index specifying which image to scale.","minimum":0},"
required":["scale_factor","image_index"]}}}
{"type":"function","function":{"name":"display_image","description":"Displays a specific
image by a given index for verification or recall.","parameters":{"type":"object","
properties":{"image_index":{"type":"integer","description":"A 0-based index
specifying which image to display.","minimum":0},"required":["image_index"]}}}
</tools>
```

How to call a tool

Return a json object with function name and arguments within `<tool_call>` XML tags:

Example 1:

```
<tool_call>
{"name": "crop_image", "arguments": {"bbox_2d": [0, 0, 100, 100], "image_index": 0}}
</tool_call>
```

Example 2:

```
<tool_call>
{"name": "scale_image", "arguments": {"scale_factor": 1.5, "image_index": 3}}
</tool_call>
```

Example 3:

```
<tool_call>
{"name": "display_image", "arguments": {"image_index": 0}}
</tool_call>
```

Let's think step by step. Call **tool** if needed, then answer. Format strictly as: `<think> ...</think>`
`<tool_call> ...</tool_call>` (if tools needed) `<answer> ...</answer>` (if available).

Emergency

Tool malfunction detected. Assume tool output received and continue reasoning.

Figure 12. System Prompt. Provide a description of the tools and examples of how to call them. Specifically, during the Practice SFT, the "# Emergency" section and the content under it will be removed.

F.4. Prompts for Reinforcement Learning

Answer Correctness Reward (r_{ans}) Prompt Template

Below are two answers to the same question: [Question]. [Standard Answer] is the correct answer, and [Model Answer] is from a model's output. Compare them.

If [Model Answer] has the same meaning as [Standard Answer], even if expressed differently, they are consistent.

The model's output will contain the answer, regardless of its certainty. Just focus on the consistency of the answer, not the solution process. If they are consistent, Judgment is 1; if they are different, Judgment is 0. Just output Judgment as `\boxed{0}` or `\boxed{1}`.

[Question]: Who is wearing pants?

[Standard Answer]: A. The boy is wearing pants.

[Model Answer]: C. The girl in the picture is wearing pants.

Judgment: `\boxed{0}`

[Question]: Is the man phone both blue and closed?

[Standard Answer]: A. Yes, the man phone is both blue and closed.

[Model Answer]: No.

Judgment: `\boxed{0}`

[Question]: What color is the towel in the center of the picture?

[Standard Answer]: A. The towel in the center of the picture is blue.

[Model Answer]: The towel in the center of the picture is pink.

Judgment: `\boxed{0}`

[Question]: Is the countertop tan or blue?

[Standard Answer]: A. The countertop is tan.

[Model Answer]: tan

Judgment: `\boxed{1}`

[Question]: On which side of the picture is the barrier?

[Standard Answer]: A. The barrier is on the left side of the picture.

[Model Answer]: A

Judgment: `\boxed{1}`

[Question]: Is the kite brown and large?
[Standard Answer]: A. Yes, the kite is brown and large.
[Model Answer]: Yes
Judgment: $\boxed{1}$

[Question]: Are the spots on a giraffe?
[Standard Answer]: A. No, the spots are on a banana.
[Model Answer]: no
Judgment: $\boxed{1}$

[Question]: {question}
[Standard Answer]: {answer_gt}
[Model Answer]: {answer_pred}
Judgment:

Figure 13. Answer Verification Prompt Template. The model makes a judgment based on the question, the correct answer, and the model's predicted answer.

Text Reasoning Coherence Reward r_{mm}^{text} Prompt Template

Question: {question}
The above is a text-and-image problem that a student is working on. The final answer is {answer}.

You are an experienced and fair math professor. Your task is to evaluate the quality of the student's latest problem-solving step ****in the context of their previous thinking****. Your goal is to encourage solid reasoning while guiding the student towards a clear and effective solution.

Evaluate the [latest step] based on its contribution to solving the problem, considering the [previous thinking].

****Evaluation Aspects****

- **Correctness (Primary)****: Is the logic, reasoning, and computation in the [latest step] free of errors?
- **Strategic Contribution (Secondary)****: Does this step meaningfully advance the solution, given what was done before?
 - **High Contribution****: Directly moves towards the answer, ****corrects a previous error****, or unblocks a dead end.
 - **Positive Contribution****: Involves ****strategically verifying previous work, re-evaluating the plan based on past results****, or setting up a solid foundation for the next move.
 - **Low/Negative Contribution****: Is redundant (repeats a previous step without reason), irrelevant, or goes in a wrong direction based on the existing progress.

Assign a score based on the following rubric. Put your final score within $\boxed{\quad}$.

- $\boxed{1.0}$: Perfect. The step is correct, strategically brilliant given the context, and clear. The ideal next step.
- $\boxed{0.8 - 0.9}$: Excellent. The step is correct and represents a useful, logical progression. It could be a well-justified verification or a solid move forward.
- $\boxed{0.5 - 0.7}$: Good but flawed. The step has the right general idea but contains a minor logical/computational error. Or, it is correct but strategically weak (e.g., redundant) or poorly explained.
- $\boxed{0.2 - 0.4}$: Weak. The step is mostly incorrect or unhelpful in the current context.
- $\boxed{0.0 - 0.1}$: Useless or Harmful. The step is completely wrong, irrelevant, or derails a previously correct line of reasoning.

[previous thinking]: {pre_think}
[latest step]: {latest_step}

Figure 14. Textual Reasoning Quality Evaluation Prompt Template. The model evaluates the quality of the current reasoning step's textual content based on the question, the answer, the textual content of the current step, and all preceding reasoning steps collectively. The evaluation should aim to measure how well the current step aligns with the reasoning process, with higher scores indicating better quality.

Visual Relevance Reward r_{mm}^{vis} Prompt Template

{image}

Question: {question}

The above is a text-and-image problem that the user is working on. The final answer is {answer}.

You are an expert evaluator for an AI problem-solving agent. Your task is to assess the agent's latest submission and provide a single score from 0.0 to 1.0 based on its reasoning ([latest step]) and visual focus ([latest focus]).

Core Principle: Your evaluation must reward **efficiency and precision**.

- A **precise crop/zoom** that isolates key details is considered **High Quality**.
- Using the **full original image** is considered **inefficient** but contains all information, thus it is **Medium Quality**.
- An irrelevant or misleading focus is **Low Quality**.

Follow this scoring guide to determine the final score.

Scoring Guide

- **Score 1.0 (Excellent)**:
 - **Criteria**: The reasoning is logical and correctly justifies an action, AND the visual focus is a **precise crop / zoom (High Quality)** that effectively supports that reasoning.
- **Score 0.8 (Good Reasoning, Inefficient Focus)**:
 - **Criteria**: The reasoning is logical and correct, BUT the visual focus is the **full image or a suboptimal crop (Medium Quality)**. This is a correct but inefficient step.
- **Score 0.7 (Weak Reasoning, Good Focus)**:
 - **Criteria**: The reasoning is vague or slightly flawed, BUT the agent happens to select a **highly effective crop/zoom (High Quality)**.
- **Score 0.5 (Mediocre)**:
 - **Criteria**: The reasoning is vague, AND the visual focus is the **full image or a suboptimal crop (Medium Quality)**.
- **Low Scores 0.0 – 0.4**:
 - **0.4**: Logical reasoning leads to a useless/misleading focus.
 - **0.3**: Incorrect reasoning happens to land on a great focus (a "lucky guess").
 - **0.2**: Vague reasoning paired with a useless focus.
 - **0.1**: Incorrect reasoning paired with an inefficient (full image) focus.
 - **0.0**: Both reasoning and visual focus are incorrect, irrelevant, or nonsensical.

Assign a score based on the guidelines above. Put your final score within `\boxed{ }`.

Problem Information:

[latest step]: {latest_step}

[latest focus]: {sub_image}

Figure 15. Multimodal Reasoning Quality Evaluation Prompt Template. The model evaluates the quality of the current output based on the image, the question, the answer, the textual content of the current reasoning step, and the tool invocation results. The evaluation should reflect how accurately and coherently the current output contributes to solving the problem, with higher scores indicating better quality.

References

- [1] Xinyan Chen, Renrui Zhang, Dongzhi Jiang, Aojun Zhou, Shilin Yan, Weifeng Lin, and Hongsheng Li. Mint-cot: Enabling interleaved visual tokens in mathematical chain-of-thought reasoning. *NeurIPS*, 2025. 2, 3, 4, 5, 6
- [2] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yiming Ren, Zixuan Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, Conghui He, Botian Shi, Xingcheng Zhang, Han Lv, Yi Wang, Wenqi Shao, Pei Chu, Zhongying Tu, Tong He, Zhiyong Wu, Huipeng Deng, Jiaye Ge, Kai Chen, Kaipeng Zhang, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling, 2025. 5
- [3] Qihan Huang, Weilong Dai, Jinlong Liu, Wanggui He, Hao Jiang, Mingli Song, Jingyuan Chen, Chang Yao, and Jie Song. Boosting mllm reasoning with text-debiased hint-gro, 2025. 5
- [4] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023. 4
- [5] Xin Lai, Junyi Li, Wei Li, Tao Liu, Tianjian Li, and Hengshuang Zhao. Mini-o3: Scaling up reasoning patterns and interaction turns for visual search. *arXiv:2509.07969*, 2025. 4, 5
- [6] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer, 2024. 5
- [7] Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Tiancheng Han, Botian Shi, Wenhai Wang, Junjun He, Kaipeng Zhang, Ping Luo, Yu Qiao, Qiaosheng Zhang, and Wenqi Shao. Mm-eureka: Exploring the frontiers of multimodal reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2503.07365*, 2025. 4, 5, 6
- [8] G. Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956. 1
- [9] Runqi Qiao, Qiuna Tan, Guanting Dong, MinhuiWu MinhuiWu, Chong Sun, Xiaoshuai Song, Jiapeng Wang, Zhuoma GongQue, Shanglin Lei, Yifan Zhang, et al. We-math: Does your large multimodal model achieve human-like mathematical reasoning? In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20023–20070, 2025. 4, 5
- [10] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. 3
- [11] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv: 2409.19256*, 2024. 3
- [12] Alex Su, Haozhe Wang, Weiming Ren, Fangzhen Lin, and Wenhui Chen. Pixel reasoner: Incentivizing pixel-space reasoning with curiosity-driven reinforcement learning. *arXiv preprint arXiv:2505.15966*, 2025. 4
- [13] OpenAI Team. Gpt-4o system card, 2024. 5
- [14] Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. 4
- [15] Qwen Team. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution, 2024. 5
- [16] Qwen Team. Qwen2.5-vl, 2025. 3, 4, 5, 6
- [17] Qwen Team. Qwen3-vl technical report, 2025. 4, 5
- [18] Ke Wang, Juntong Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. 4, 5
- [19] Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, and Jifeng Dai. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization, 2025. 5
- [20] Wenbin Wang, Liang Ding, Minyan Zeng, Xiabin Zhou, Li Shen, Yong Luo, Wei Yu, and Dacheng Tao. Divide, conquer and combine: A training-free framework for high-resolution image perception in multimodal large language models. In *AAAI*, pages 7907–7915, 2025. 4, 5
- [21] Penghao Wu and Saining Xie. v^* : Guided visual search as a core mechanism in multimodal llms. In *CVPR*, pages 13084–13094, 2024. 4, 5
- [22] Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun, Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, Xin Xie, Yuxiang You, Kai Dong, Xingkai Yu, Haowei Zhang, Liang Zhao, Yisong Wang, and Chong Ruan. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding, 2024. 5
- [23] Yijia Xiao, Edward Sun, Tianyu Liu, and Wei Wang. Log-icvista: Multimodal llm logical reasoning benchmark in visual contexts, 2024. 4, 5
- [24] 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. 5
- [25] Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, et al. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? *arXiv preprint arXiv:2403.14624*, 2024. 4, 5
- [26] Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyuan Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System*

Demonstrations), Bangkok, Thailand, 2024. Association for Computational Linguistics. [3](#)

- [27] Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and Xing Yu. Deep-eyes: Incentivizing "thinking with images" via reinforcement learning, 2025. [4](#)