

Mitigating Visual Forgetting via Take-along Visual Conditioning for Multi-modal Long CoT Reasoning

Hai-Long Sun^{1,2} *, Zhun Sun^{4,5} *, Houwen Peng³, Han-Jia Ye^{1,2} †

¹School of Artificial Intelligence, Nanjing University

²National Key Laboratory for Novel Software Technology, Nanjing University

³Tencent ⁴Center for Language AI Research, Tohoku University

⁵RIKEN Center for Advanced Intelligence Project

Abstract

Recent advancements in Large Language Models (LLMs) have demonstrated enhanced reasoning capabilities, evolving from Chain-of-Thought (CoT) prompting to advanced, product-oriented solutions like OpenAI o1. During our re-implementation of this model, we noticed that in multimodal tasks requiring visual input (*e.g.*, geometry problems), Multimodal LLMs (MLLMs) struggle to maintain focus on the visual information, in other words, MLLMs suffer from a gradual decline in attention to visual information as reasoning progresses, causing text-over-relied outputs. To investigate this, we ablate image inputs during long-chain reasoning. Concretely, we truncate the reasoning process midway, then re-complete the reasoning process with the input image removed. We observe only a $\sim 2\%$ accuracy drop on MathVista’s test-hard subset, revealing the model’s textual outputs dominate the following reasoning process. Motivated by this, we propose Take-along Visual Conditioning (TVC), a strategy that shifts image input to critical reasoning stages and compresses redundant visual tokens via dynamic pruning. This methodology helps the model retain attention to the visual components throughout the reasoning. Our approach achieves state-of-the-art performance on average across five mathematical reasoning benchmarks (+3.4 points vs previous sota), demonstrating the effectiveness of TVC in enhancing multimodal reasoning systems. The project page is available at <https://sun-hailong.github.io/projects/TVC>.

1 Introduction

Large Language Models (LLMs) have achieved significant advancements in natural language processing (NLP), particularly in the area of reasoning. These models have evolved from simple

*Work done during their internship/employment at Tencent Hunyuan

†Corresponding author: yehj@lamda.nju.edu.cn

prompt-based Chain-of-Thought (CoT) (Wei et al., 2022) techniques to sophisticated product-oriented solutions like OpenAI’s o1 (OpenAI, 2024b), DeepSeek-R1 (DeepSeek-AI, 2024), and Qwen-QVQ (QwenTeam, 2024), demonstrating iterative reasoning capabilities for complex multi-step tasks, which enables them to handle tasks that require multi-step thinking, logic, and knowledge integration. Recently, several works also extended the CoT reasoning process to MLLMs settings through data-centric innovations. For instance, Math-LLaVA (Shi et al., 2024) pioneers domain-specific training with the MathV360K dataset, while MAMmoTH-VL (Guo et al., 2024) scales up multimodal CoT data generation.

While such progress is notable in text-based domains, extending these advancements to multimodal scenarios presents unique challenges that transcend traditional language model boundaries. Reasoning in MLLMs requires fused understanding across different modalities, for example, in geometric reasoning tasks the model should interpret and reason about images alongside text. Therefore, the model’s ability to integrate and maintain focus on both types of information is critical. Unlike text-only LLMs that reinforce problem context through linguistic recurrence, MLLMs struggle to sustain visual attention across reasoning steps. That is, as the length of the context increases, the model is more inclined to conduct the next step of reasoning based on the previously outputted text rather than the information of the image itself, which eventually leads to the continuation of the wrong text reasoning process and degraded model performance. We term this phenomenon as visual forgetting.

In this work, we conduct a diagnostic analysis of the visual forgetting effect within a long-chain reasoning system. The system processes multimodal Q&A tasks through a series of interconnected reasoning steps. We demonstrate significantly reduced attentional allocation to visual inputs during multi-

stage reasoning processes. More formally, our analysis: 1) truncates the reasoning process midway and removes the image embeddings; 2) regenerates subsequent reasoning trajectories; 3) evaluates the reasoning outcomes of the pre/post-ablation inference trajectories. Intuitively, the performance gap between normal reasoning and diagnostic reasoning reveals the model’s dependency on generated text over original visual evidence. Our results (See Section 2.1) on the MathVista-Hard datasets show that, removing the image midway through the reasoning only causes an insignificant performance degradation (*i.e.* $\sim 2\%$), indicating that the model completes the reasoning process based primarily on its output text. More importantly, we also observe that the model’s dependency on the visual evidence diminishes over time since the early removal of the image inputs could hurt accuracy by $\sim 20\%$. This suggests model’s reasoning employs both visual and textual information in the early stage, then over-relying on text history which limits full visual reasoning potential.

Motivated by this, we propose a novel strategy to mitigate the visual forgetting effect and maintain visual attention throughout the reasoning process. Our methodology compresses and shifts the image input to the later stages of the reasoning process, ensuring the model integrates sufficient visual evidence into its reasoning. This approach results in improved performance and achieves state-of-the-art results on average across five mathematical reasoning benchmarks (*i.e.*, $+3.4\%$ vs previous sota). Our findings highlight the effectiveness of this strategy in enhancing the performance of multimodal reasoning systems, providing a robust solution to the problem of visual forgetting in long-chain reasoning tasks.

2 Take-along Visual Conditioning: Sustaining Visual Evidence for Multi-modal Long CoT Reasoning

In this section, we first discuss our motivation and observations of the visual forgetting phenomenon in MLLM reasoning systems (Section 2.1). Then, we propose the Take-alone Visual Conditioning (TVC) approach to mitigate visual forgetting and enhance the model’s long-chain reasoning capabilities (Section 2.2).

2.1 Capturing the Visual Forgetting

Text-based reasoning systems often repeat key terms or phrases (*e.g.*, restating “triangle ABC” or




Figure 1: **The visual forgetting phenomenon by removing the image at different reasoning stages.** It shows that by the midpoint of the reasoning process, the model becomes less dependent on the image, causing text-over-relied outputs.

“equation (1)”) to keep the problem’s context clear and focused. By restating these critical details, they create strong connections between each step of the reasoning process and the original problem’s requirements. This repetition ensures the logic stays on track, and consistent to follow.

$$\mathcal{C}_{\text{MLLM}} = f(V, T_1, \dots, T_n) \quad (1)$$

On the other hand, MLLMs struggle with this approach due to their design. As formalized in Equation (1), visual inputs V are confined to the initial processing stages. Unlike text, these visual evidence aren’t revisited or reinforced later in the reasoning process. Because there’s no built-in way to keep visual information “active” or relevant throughout the task, the system’s ability to connect visual details with text or logic weakens over time, leading to a progressive *visual attention decay*. The model is more likely to reason with previously outputted text and becomes particularly problematic in visual reasoning tasks that require continuous validation of spatial relationships.

We conduct two analytic analyses using the QVQ-72B-Preview model (QwenTeam, 2024) to capture this visual forgetting phenomenon quantitatively and qualitatively. On one hand, we remove the visual inputs at eight different stages to observe their impact. On the other hand, we depict the attention matrix to directly track the attention decay of the visual evidence over time.

Progressive Image Removing. To assess the extent to which the reasoning process depends on

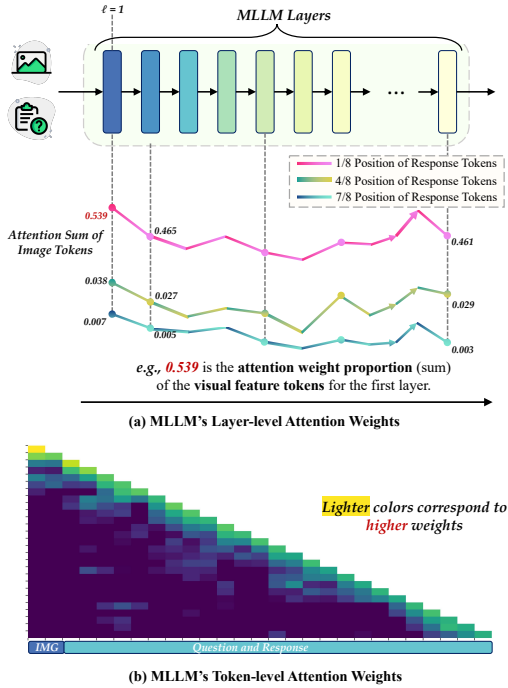


Figure 2: **Illustration of layer-level and token-level attention weights.** (a) The layer-level attention weights of image tokens across different response token positions. (b) The token-level attention weights at the middle layer. It shows that the model’s attention to the image gradually decreases during the reasoning process.

previously generated text, we first perform a normal reasoning process, then reset the KV cache at various stages of the reasoning process. This effectively removed image tokens and forced subsequent steps to rely solely on text-based information. Specifically, the reasoning process was divided into $K = 8$ intervals based on output token counts regardless of the length of the normal reasoning process, with visual input progressively masked by resetting the KV cache at different cut-off positions $k \in \{0, 1, \dots, K-1\}$. In other words, the first $k/8$ part of the normal reasoning process is now employed as a prompt, and the model now re-complete the reasoning process without image inputs. Furthermore, we discovered that for some questions (30.9% of the MathVista dataset), the model could answer correctly using only the text-based prompt. Consequently, we excluded these questions and designated the remaining dataset as the *MathVista-Hard* dataset.

As illustrated in Figure 1, we observe that the performance is 40.9 at $k = 4$ and 43.1 for the normal reasoning. This minimal 2.2% gap suggests that the model overly relies on text outputs during the later reasoning stages rather than on the image.

This indicates that once the model completes the half stages of reasoning, it becomes overly reliant on its own generated CoT. Over-reliance on this self-generated reasoning worsens the model’s tendency to forget visual evidence over time, which can result in hallucinations during lengthy, multi-step reasoning tasks. This result also reveals a near-exponential forgetting effect in performance as k increases, which we formalize as:

$$\mathcal{R}(k) = \mathcal{R}_{\text{full}} - \Delta_{\text{visual}}(k), \quad \Delta_{\text{visual}}(k) \propto e^{-k}, \quad (2)$$

where $\mathcal{R}_{\text{full}}$ represents the full multimodal reasoning performance, and $\Delta_{\text{visual}}(k)$ quantifies the performance degradation caused by visual masking at position k .

Visual Attention Decay. To more intuitively observe the changes in visual feature attention across different stages of reasoning, we investigate the layer-level attention weights at various response token positions and the token-level attention weights at layer 16. Given that MLLM’s attention weights reflect the focus on tokens and influence the decision-making process, we first analyze the attention weights at each layer of the MLLM. Specifically, for each layer, we calculate the proportion of attention weights on all image tokens.

As shown in Figure 2(a), we observe that at the 1/8 position of the reasoning process, the model effectively focuses on the visual inputs. However, as reasoning progresses, despite fluctuations in attention to visual tokens across layers, the model’s overall attention to visual evidence gradually decreases, leading to visual forgetting. Next, following the methodology of FastV (Chen et al., 2024a), we analyze the attention maps for several cases and find that the model predominantly focuses on previously generated text tokens rather than the input image. After approximately 20% tokens, the existence of image inputs on attention maps diminishes significantly, as illustrated in Figure 2(b). Both of the observations indicate a phenomenon of visual memory degradation, revealing the model’s limitations in maintaining consistent attention to visual inputs throughout the reasoning process.

2.2 Take-along Visual Conditioning

In this section, we introduce our solution to tackle this problem in detail. We propose Take-along Visual Conditioning (TVC), a dynamic image retention mechanism that re-introduces visual inputs at strategic intervals throughout the reasoning pro-

cess. TVC mitigates visual attention decay by periodically reaffirming visual information, akin to human problem-solving behaviors where individuals frequently refer back to visual inputs. Our approach enhances the model’s ability to incorporate visual information continuously, improving its long-chain reasoning capacity by ensuring that visual evidence is revisited during critical decision-making moments.

The TVC method consists of two key stages: training and testing. In the training stage, we introduce Dynamic Visual Reaffirmation (DVR), which guides the model through iterative reinforcement of visual evidence during long reasoning chains. In the testing phase, we present Periodic Visual Calibration (PVC), where visual reactivation is periodically triggered at self-reflection intervals. To prevent the model from forgetting previous text-based reasoning steps due to an excessive number of image tokens, we adopt image compression through adaptive pooling to reduce the image token size while preserving spatial semantics. This dual-modality engagement mechanism ensures consistent interaction between textual reasoning and visual evidence. We present an illustration of the TVC system in Figure 3.

Dynamic Visual Reaffirmation. Our dynamic visual reaffirmation training strategy combines two key components: 1) *Data Curation*: We curate long-chain reasoning data using the pipeline described in Section 3, sourced from high-quality academic datasets (e.g., MathV360K, Geo170K, and LLaVA-OneVision). This process yields a high-quality dataset optimized for long-chain reasoning training. 2) *Visual Content Injection*: While the curated data ensures correctness, the QVQ model inherently lacks the ability to iteratively reference visual inputs during reasoning. Therefore, we manually re-inject the visual content (visual embeddings and bridging prompt) to triggers visual re-activation at predefined self-reflection intervals. Specifically, given the initial multimodal input $\mathcal{M}_0 = (V, T_0)$, DVR performs visual reactivation at self-reflection intervals $\{r_1, \dots, r_m\}$:

$$\mathcal{M}_i = \left(V, [T_{\text{prev}}; \underbrace{\text{Prompt}}_{\text{Re-activation}}; T_{\text{new}}] \right) \text{ at step } r_i \quad (3)$$

where T_{prev} represents the previous reasoning steps and T_{new} denotes the new reasoning steps that are based on prior reasoning and reintroduce focus on

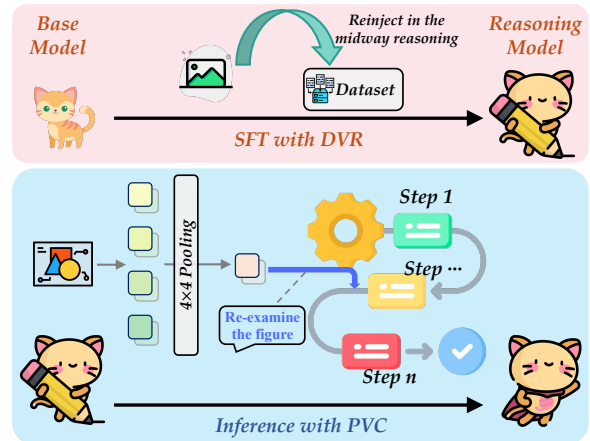


Figure 3: **Overview of TVC System Design.** We enable the model to have take-along visual conditioning capabilities through two stages: training and inference.

visual information. The bridging prompt is employed to hint the existence of the image, e.g., Let me see the image again. To improve efficiency, our initial implementation adopts midpoint reactivation ($m = 1, r_1 = 0.5L$ for L -step chains).

During the self-reflection phase of the reasoning steps, we randomly select m positions to inject visual content. Specifically, we inject reactivation prompts and regenerate visual captions using the model’s intrinsic capabilities. This forces joint attention to both textual reasoning and visual evidence. By leveraging the model’s intrinsic image captioning capability, we continuously reinforce visual information throughout the reasoning process. This ensures that the model incorporates image evidence during reflection, rather than relying solely on textual reasoning.

Periodic Visual Calibration. Calibrating visual attention plays a crucial role in enhancing long-chain reasoning capabilities. After training our model, we further design the periodic visual calibration process. Specifically, we coordinate operations during reactivation as follows: 1) *Token Compression*: We first compress visual tokens using average pooling to prevent text-based reasoning from forgetting visual information. 2) *Visual Cache Reset*: We then prepend an instruction (bridging prompt) to reintroduce the image and re-inject image tokens by resetting the KV cache of the generation process.

We also provide an example of how PVC is implemented in the case study section (Section 4.5). PVC both improves reasoning efficiency and prevents the model from forgetting previous reasoning steps due to an overload of visual tokens.

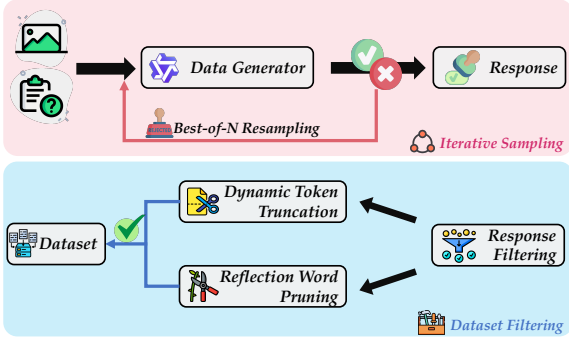


Figure 4: **Data Generation Pipeline of TVC.** We use iterative distillation to collect long-chain reasoning data, followed by a comprehensive response filtering process to ensure high-quality reasoning.

3 Data-Centric Implementation of Multimodal Reasoning System

In this section, we briefly describe our implementation of the multimodal reasoning system through a scalable curated data generation pipeline.

3.1 Long-Chain Reasoning Data Collection

Prior research (Qin et al., 2024; Jiang et al., 2024) has identified two dominant paradigms for constructing long-chain reasoning systems: (1) explicit search-based methods, which utilize structures such as Monte Carlo Tree Search (MCTS) combined with specially trained reward models to guide the search process toward optimal solutions, and (2) instruction distillation approaches, which fine-tune models on curated long chain-of-thought (CoT) datasets. To efficiently develop an MLLM with long-chain reasoning capabilities, we adopt the distillation paradigm. In this section, we describe the distillation process and present the data generation pipeline aimed at enhancing the reasoning capability of MLLM.

Our distillation pipeline begins with aggregating publicly available reasoning datasets (e.g., MathV360K (Shi et al., 2024) and Geo170K (Gao et al., 2023)). Through an iterative process of distillation and response filtering, we progressively enhance the model’s long-chain reasoning capabilities. Specifically, the teacher model, QVQ-72B-Preview, generates long CoT reasoning data, while the student model, Qwen2-VL, undergoes supervised fine-tuning on the filtered CoT data to enhance its reasoning performance.

3.2 Iterative Distillation with Reject Sampling

After obtaining the long-chain reasoning responses, we employ an assessment pipeline to ensure data quality. Specifically, we introduce a dual-temperature sampling mechanism to optimize data quality through variance exploitation.

Deterministic Initial Sampling. For the first-stage data generation, we employ temperature $\tau = 0$ to obtain the QVQ’s most confident reasoning paths:

$$\mathcal{D}_{\text{init}} = \{(x, y_{\text{CoT}}) \sim \mathcal{P}_{\text{QVQ}}(\cdot | x; \tau = 0)\}, \quad (4)$$

where y_{CoT} represents the response generated by QVQ. This ensures that the model generates the most probable reasoning path for each input. Subsequently, we use these high-confidence responses as a foundation for further refinement in later stages of data generation.

Answer-Centric Reject Sampling. To ensure the quality of the data, we employ an LLM-as-a-Judge approach for answer-centric reject sampling. Specifically, we use a strong LLM (e.g., Qwen2.5-72B-Instruct) as the verifier and prompt the model to output a ‘yes’ or ‘no’ in JSON format, indicating whether the long-chain reasoning answer is correct, i.e., $\mathcal{Y}_{\text{valid}} = \{y_{\text{CoT}} | \text{Match}(y_{\text{CoT}}, y_{\text{gt}})\}$. After this process, we obtain $\sim 200\text{K}$ high-quality samples.

Best-of-N Error Correction. While initial sampling with temperature $\tau = 0$ yields a set of data, there are substantial residual errors ($\mathcal{D}_{\text{error}}$) where QVQ’s responses (y_{CoT}) unmatched with ground truth answers (55.8% of $\mathcal{D}_{\text{init}}$). To fully leverage the available open-source data, we implement a contrastive regeneration strategy using $\tau = 1$:

$$\mathcal{D}_{\text{corrected}} = \bigcup_{x \in \mathcal{D}_{\text{error}}} \left\{ \arg \min_{y^i \in \mathcal{Y}_{\text{valid}}} \text{Length}(y^i) \right\}_{i=1}^8, \quad (5)$$

where $\mathcal{Y}_{\text{valid}}$ denotes the subset of 8 regenerated responses (at $\tau = 1$) that satisfy answer correctness. For cases with multiple valid responses, we prioritize minimal reasoning token length to enhance conciseness and efficiency:

$$\text{Length}(y) = \sum_{t=1}^T \mathbb{I}(w_t \notin \{\text{[PAD]}, \text{[SEP]}\}) \quad (6)$$

This two-stage sampling achieves partial error recovery while reducing the average token count to minimize meaningless self-reflection, effectively balancing correctness and efficiency. To further enhance the quality of the dataset, we also perform

dynamic token truncation and reflection word pruning. This approach helps reduce the ineffective reflection parts in the dataset, thereby mitigating the issue of the model being unable to complete its answers. More details are in Section B.

4 Experiments

We conduct comprehensive experiments across multiple vision-language benchmarks to demonstrate the effectiveness of our method. Section 4.1 provides implementation details for TVC. In Section 4.3, we present key results on visual reasoning tasks, supplemented with findings on general image understanding. Section 4.4 details ablation studies to evaluate the influence of critical design decisions. Finally, Section 4.5 presents visualizations and case studies to illustrate the method’s operational characteristics and insights derived from it.

4.1 Training Recipe

We integrate the TVC system with MLLMs of varying scales to demonstrate the effectiveness and generalizability of our approach. Initial implementation with Qwen2-VL-7B-Instruct confirmed the method’s validity. To further evaluate scalability and establish robust baselines against state-of-the-art MLLMs, we expanded the approach to a 72B model. Prior to training TVC, we follow the long-chain reasoning pipeline described earlier. We use the LLaMA-Factory (Zheng et al., 2024) framework, with a learning rate of $2e-5$, a batch size of 256, and 5 training epochs. During optimization, only the LLM parameters and cross-modal connector are trained, while the visual encoder remains frozen. The training process requires 10 hours on a 64×H20 GPU setup for the 7B model and approximately 4 days for the 72B model. Additional details are presented in Table 4.

4.2 Evaluation Setup

We conduct a comprehensive experimental analysis across various visual reasoning benchmarks that require advanced visual reasoning skills. To ensure a well-rounded evaluation, we select several widely recognized and representative benchmarks, including MathVista (Lu et al., 2024a), MathVerse (Zhang et al., 2024a), MathVision (Wang et al., 2024a), Dynamath (Zou et al., 2024), and OlympiadBench (He et al., 2024). MathVista consists of 6,141 examples that require fine-grained, deep visual understanding and compositional reasoning. MathVerse contains 2,612 multi-subject

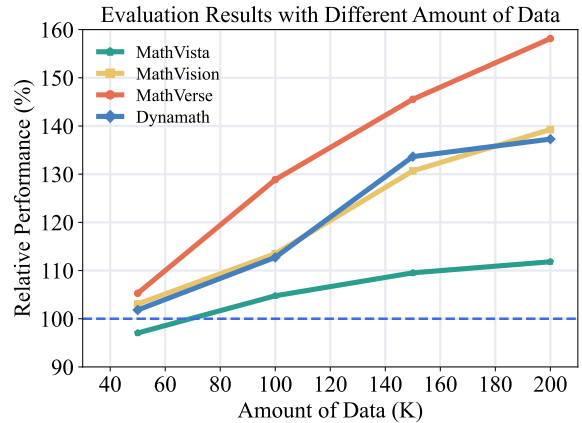


Figure 5: **Ablations on the amount of training data.** TVC benefits from data scaling, continually improving the reasoning capabilities.

math problems from a variety of sources. MathVision includes 3,040 high-quality mathematical problems sourced from established mathematics competitions. OlympiadBench features 8,476 bilingual multimodal problems tailored to Olympic-level mathematics and physics competitions. These benchmarks evaluate problem-solving abilities in mathematics, and following standard practice, we use GPT-4o-mini as the evaluator. Following the VLMEvalKit guidelines, we exclude the text-only split from MathVerse and the theorem-proof sections from OlympiadBench. For a fair comparison, we conduct evaluations using the testmini sets of MathVerse and MathVista. Fast evaluation is made possible through the use of the VLMEvalKit (Duan et al., 2024) and vLLM (Kwon et al., 2023).

4.3 Main Results on Visual Reasoning

The results in Table 1 demonstrate the effectiveness and generalizability of TVC across multiple visual reasoning benchmarks. Notably, our model achieves a 16.7% improvement over QVQ-72B-Preview on MathVision and a 17.6% gain on MathVerse, highlighting enhanced reasoning capabilities. Unlike conventional datasets where textual descriptions may include implicit visual cues, MathVerse is an all-around visual math benchmark specifically designed for equitable, in-depth evaluation of MLLMs. The significant gains on MathVerse underscore the significance of TVC, given the benchmark’s unique design principles. Furthermore, the TVC-7B model, despite its smaller size compared to counterparts, achieves competitive performance, even outperforming leading MLLMs in multiple cases. This demonstrates the robustness of our

Table 1: **Results on Visual Reasoning Tasks.** We conduct evaluation experiments across 6 benchmarks, covering both general reasoning and task-specific reasoning assessments. TVC exhibits notable effectiveness and generalizability when applied to Qwen2-VL, surpassing other state-of-the-art MLLMs by a large margin.

Model	Size	MathVista	MathVision	MathVerse	Dynamath	OlympiadBench	Average
MiniCPM-V-2.6 (Yadav et al., 2025)	8B	60.8	18.4	17.6	9.8	-	-
VITA-1.5 (Fu et al., 2025a)	8B	66.2	19.5	23.4	9.6	-	-
LLaVA-COT (Xu et al., 2024)	11B	52.5	19.9	22.6	7.8	-	-
Qwen2-VL (Wang et al., 2024b)	7B	60.9	16.3	24.6	11.0	3.2	23.2
InternVL2.5 (Chen et al., 2024b)	8B	64.5	17.0	22.8	9.4	0.1	22.8
POINTS1.5 (Liu et al., 2024b)	8B	66.4	22.0	26.6	14.2	-	-
Ovis1.6-Gemma2 (Lu et al., 2024b)	27B	70.2	20.6	37.8	17.0	-	-
InternVL2.5-COT (Chen et al., 2024b)	78B	71.4	32.5	40.1	28.5	-	-
LLaVA-OneVision (Li et al., 2024)	72B	67.1	25.3	27.2	15.6	-	-
Qwen2-VL (Wang et al., 2024b)	72B	69.7	26.6	36.2	20.0	10.3	32.6
QVQ-72B-preview (QwenTeam, 2024)	72B	71.4	35.9	41.5	30.7	20.4	40.0
TVC	7B	68.1	22.7	38.9	15.1	9.8	30.9
TVC	72B	72.2	41.9	48.8	30.0	24.3	43.4

Table 2: **Ablations on the TVC System.** TVC enhances reasoning capabilities, showing significant improvements on both general and task-specific reasoning benchmarks.

Method	MathVista	MathVision	MathVerse	Avg
Baseline	60.9	16.3	24.6	33.9
Vanilla - Direct SFT	63.5	19.8	31.6	38.3
TVC w/o PVC	66.7	21.8	35.6	41.4
TVC w/o DVR	66.2	22.3	34.7	41.0
TVC Full	68.1	22.7	38.9	43.2

methodology even with more compact model configurations. Beyond task-specific visual reasoning, we extend our evaluation to general reasoning benchmarks (*e.g.*, MathVista), where TVC consistently delivers strong performance, achieving a 3.6% improvement over the original Qwen2-VL-72B model. These results emphasize TVC’s ability to excel in tasks requiring both perception and reasoning. Collectively, the findings indicate that TVC not only advances performance in specialized visual reasoning tasks but also offers substantial benefits in broader application scenarios.

4.4 Further Analysis

In this section, we conduct comprehensive experiments to evaluate the design choices of TVC, emphasizing the key contributions of our approach. We also present a case study to further illustrate the qualitative effectiveness of TVC.

Effectiveness of TVC system. To evaluate the effectiveness of the TVC system, we conduct comprehensive ablation experiments on various components using the Qwen2-VL-7B as the Baseline. We begin by performing supervised fine-tuning on the

Table 3: **Ablations on Token Compression.**

Method	MathVista	MathVision	MathVerse	Avg
TVC Baseline	68.3	21.5	39.6	43.1
+ 2x2 Avg Pooling	67.8	22.9	38.3	43.0
+ 4x4 Avg Pooling	68.1	22.7	38.9	43.2

Qwen2-VL-7B model with the data from Section 3, referred to as Vanilla - Direct SFT. Furthermore, we apply the DVR training approach outlined in Section 2.2, which increases the focus on the visual information in the training data, enabling the model to implicitly learn visual conditioning capabilities. Additionally, during the testing phase, we experiment with resetting the visual KV cache midway through the reasoning process, and after compressing the visual tokens, we add them to the end of the reasoning steps. This strategy allows the model to further observe the image content during its thought process, mitigating the visual forgetting and suppressing hallucinations.

As demonstrated in Table 2, the results highlight that the TVC system is crucial for enhancing visual reasoning capabilities. When using only the SFT training data without the DVR strategy in TVC (Vanilla - Direct SFT), improvements in visual reasoning tasks are limited, as the model lacks the ability to reflect on visual information. Furthermore, models trained with the DVR strategy alone still produce sub-optimal results, underscoring the importance of a comprehensive approach that integrates PVC and DVR. The contributions of PVC and DVR are relatively balanced in enhancing the reasoning capabilities. To further validate the effectiveness of our visual token compression, we

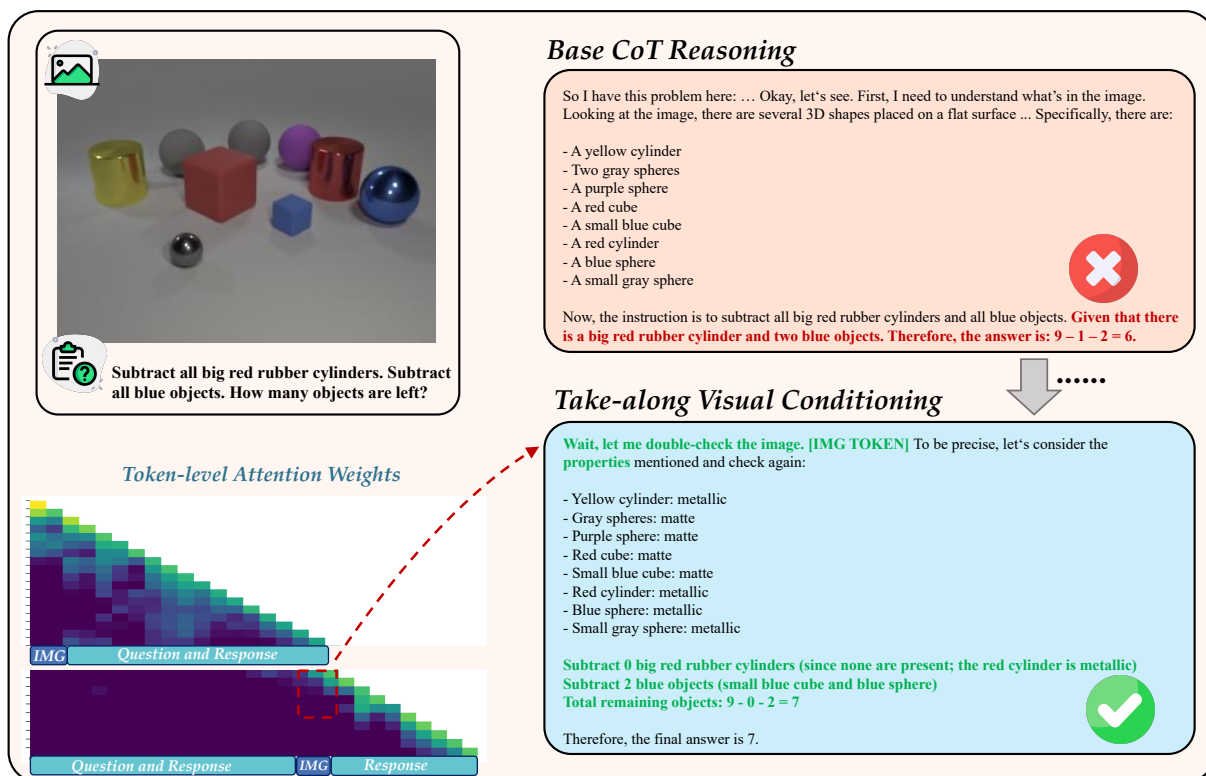


Figure 6: **Case Study of TVC.** TVC effectively re-examines the image during the reflection process to correct mistakes, guiding the model to the correct answer.

conduct experiments with different pooling methods. As shown in Table 3, the TVC Baseline represents the method without image compression. We observe that the use of pooling methods has little impact on the model’s capabilities. Utilizing 4x4 average pooling for compression not only enhances the model’s inference efficiency but also achieves a slight performance improvement.

Data Scaling Law. To validate the scalability of our approach, we also conduct a data scaling law experiment for training. As shown in Figure 5, we compare models across various data size: 50K, 100K, 150K, and 200K samples, and present the performance relative to the base model (Qwen2-VL-7B). Our findings show that as the dataset size increases, the model’s performance continues to improve. However, it is evident that with a small dataset, the model’s reasoning ability cannot reach an optimal level solely through SFT. Therefore, data scaling proves highly effective during SFT training, and the results indicate that TVC benefits from increased data.

To prevent underfitting from becoming a performance bottleneck, we increase the number of training epochs with smaller datasets, which further improves model performance. This indicates that in-

sufficient data or epochs lead to undertrained models, making it difficult to learn long-chain reasoning. Increasing both the data and training epochs can effectively alleviate this issue.

4.5 Case Study

We present a comprehensive case study in Figure 6 to illustrate the improvements of our TVC approach. Specifically, we provide an example that highlights advancements in the reasoning process. In this example, the model is tasked with carefully observing the objects within an image and, after eliminating certain shapes and attributes, providing the count of the remaining objects. During base CoT reasoning, the model fails to check the attributes of each object and only focuses on their shapes, leading to an incorrect final answer. Subsequently, the model learns to re-focus on the image and carefully describe the attributes of each object in detail. This process allows the model to identify the issue in its previous reasoning and provide the correct answer.

5 Conclusion

In this paper, we introduce Take-along Visual Conditioning (TVC), a novel strategy designed to en-

hance the reasoning capabilities of MLLMs by addressing the challenge of visual attention degradation during long-chain reasoning. By dynamically shifting the image input to critical stages of reasoning and compressing redundant visual tokens, we ensure that the model maintains focus on the visual information throughout the process. Our extensive evaluation on several mathematical reasoning benchmarks demonstrates the effectiveness of TVC in improving multimodal reasoning, providing a robust approach to equip MLLMs with better visual grounding for complex tasks.

Acknowledgments

This work is partially supported by National Key R&D Program of China (2024YFE0202800), NSFC (62376118), Key Program of Jiangsu Science Foundation (BK20243012), CCF-Tencent Rhino-Bird Open Research Fund (RAGR20240101), AMED (JP25wm0625405), and Collaborative Innovation Center of Novel Software Technology and Industrialization.

Limitations

Despite advancements, our method may still exhibit several limitations. First, for highly complex reasoning tasks requiring sophisticated analytical capabilities, simply increasing visual revisits proves insufficient. In contrast, it is crucial to enhance the model’s inherent reasoning capacity. Second, our method assumes the availability of delayed visual processing, making it potentially unsuitable for real-time applications requiring instantaneous visual feedback, such as robotic navigation or time-sensitive decision-making scenarios.

References

- Meta AI. 2024. Build the future of ai with meta llama 3. Technical report, Meta AI.
- anthropic. 2024. Introducing the next generation of claude. Technical report, anthropic.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv:2309.16609*.
- Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. 2024a. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. *Preprint, arXiv:2403.06764*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024b. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Zhong Muyan, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2023. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv:2312.14238*.
- DeepSeek-AI. 2024. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. Technical report, DeepSeek.
- Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkang Yang, Winston Hu, Yongming Rao, and Ziwei Liu. 2024. Insight-v: Exploring long-chain visual reasoning with multimodal large language models. *arXiv preprint arXiv:2411.14432*.
- Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. 2024. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. *arXiv preprint arXiv:2407.11691*.
- Chaoyou Fu, Haojia Lin, Xiong Wang, Yi-Fan Zhang, Yunhang Shen, Xiaoyu Liu, Yangze Li, Zuwei Long, Heting Gao, Ke Li, et al. 2025a. Vita-1.5: Towards gpt-4o level real-time vision and speech interaction. *arXiv preprint arXiv:2501.01957*.
- Jiale Fu, Yuchu Jiang, Junkai Chen, Jiaming Fan, Xin Geng, and Xu Yang. 2025b. Speculative ensemble: Fast large language model ensemble via speculation. *arXiv preprint arXiv:2502.01662*.
- Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wanjun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, et al. 2023. G-llava: Solving geometric problem with multi-modal large language model. *arXiv preprint arXiv:2312.11370*.
- Jarvis Guo, Tuney Zheng, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Yizhi Li, Graham Neubig, Wenhui Chen, and Xiang Yue. 2024. Mammoth-vl: Eliciting multimodal reasoning with instruction tuning at scale. *arXiv preprint arXiv:2412.05237*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. 2024. Olympiad-bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Yushi Hu, Weijia Shi, Xingyu Fu, Dan Roth, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, and Ranjay Krishna. 2024. Visual sketchpad: Sketching as a visual chain of thought for multimodal language models. *arXiv preprint arXiv:2406.09403*.

- Jinhao Jiang, Zhipeng Chen, Yingqian Min, Jie Chen, Xiaoxue Cheng, Jiapeng Wang, Yiru Tang, Haoxiang Sun, Jia Deng, Wayne Xin Zhao, et al. 2024. Technical report: Enhancing llm reasoning with reward-guided tree search. *arXiv preprint arXiv:2411.11694*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv:2306.00890*.
- Zongzhao Li, Jiacheng Cen, Bing Su, Wenbing Huang, Tingyang Xu, Yu Rong, and Deli Zhao. 2025. Large language-geometry model: When llm meets equivariance. *arXiv preprint arXiv:2502.11149*.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024a. Llava-next: Improved reasoning, ocr, and world knowledge.
- Yuan Liu, Le Tian, Xiao Zhou, Xinyu Gao, Kavio Yu, Yang Yu, and Jie Zhou. 2024b. Points1. 5: Building a vision-language model towards real world applications. *arXiv preprint arXiv:2412.08443*.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024a. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *ICLR*.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *Advances in Neural Information Processing Systems*, 36:43447–43478.
- Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. 2024b. Ovis: Structural embedding alignment for multimodal large language model. *arXiv preprint arXiv:2405.20797*.
- Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, et al. 2024. Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv:2403.09611*.
- OpenAI. 2024a. Gpt-4o: Hello gpt-4o. Technical report, OpenAI.
- OpenAI. 2024b. Learning to reason with llms. Technical report, OpenAI.
- Yingzhe Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, Qipeng Zhu, Kai Yang, Xingzhong Xu, Xin Geng, and Xu Yang. 2025. Lmm-rl: Empowering 3b llms with strong reasoning abilities through two-stage rule-based rl. *arXiv preprint arXiv:2503.07536*.
- Yiwei Qin, Xuefeng Li, Haoyang Zou, Yixiu Liu, Shijie Xia, Zhen Huang, Yixin Ye, Weizhe Yuan, Hector Liu, Yuanzhi Li, et al. 2024. O1 replication journey: A strategic progress report—part 1. *arXiv preprint arXiv:2410.18982*.
- QwenTeam. 2024. Qvq: To see the world with wisdom. Technical report, Alibaba.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv:2403.05530*.
- Wenhao Shi, Zhiqiang Hu, Yi Bin, Junhua Liu, Yang Yang, See-Kiong Ng, Lidong Bing, and Roy Ka-Wei Lee. 2024. Math-llava: Bootstrapping mathematical reasoning for multimodal large language models. *arXiv preprint arXiv:2406.17294*.
- Hai-Long Sun, Da-Wei Zhou, Yang Li, Shiyin Lu, Chao Yi, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, De-Chuan Zhan, et al. 2025a. Parrot: Multilingual visual instruction tuning. In *ICML*.
- Hai-Long Sun, Da-Wei Zhou, De-Chuan Zhan, and Han-Jia Ye. 2025b. Pilot: A pre-trained model-based continual learning toolbox. *SCIENCE CHINA Information Sciences*, 68(4):147101.
- Hai-Long Sun, Da-Wei Zhou, Hanbin Zhao, Le Gan, De-Chuan Zhan, and Han-Jia Ye. 2025c. Mos: Model surgery for pre-trained model-based class-incremental learning. In *AAAI*, volume 39, pages 20699–20707.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.

- Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. 2024a. Measuring multimodal mathematical reasoning with math-vision dataset. *Preprint*, arXiv:2402.14804.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Yibo Wang, Yuanyu Wan, and Lijun Zhang. 2025. Revisiting projection-free online learning with time-varying constraints. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, pages 21339–21347.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv:2201.11903*.
- Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. 2024. Llava-o1: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*.
- Ankit Yadav, Lingqiao Liu, and Yuankai Qi. 2025. Exploring primitive visual measurement understanding and the role of output format in learning in vision-language models. *arXiv preprint arXiv:2501.15144*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, et al. 2024a. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? *arXiv preprint arXiv:2403.14624*.
- Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Ziyu Guo, Shicheng Li, Yichi Zhang, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Bin Wei, et al. 2024b. Mavis: Mathematical visual instruction tuning with an automatic data engine. *arXiv preprint arXiv:2407.08739*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. 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. Association for Computational Linguistics.
- Da-Wei Zhou, Hai-Long Sun, Jingyi Ning, Han-Jia Ye, and De-Chuan Zhan. 2024a. Continual learning with pre-trained models: A survey. In *IJCAI*, pages 8363–8371.
- Da-Wei Zhou, Hai-Long Sun, Han-Jia Ye, and De-Chuan Zhan. 2024b. Expandable subspace ensemble for pre-trained model-based class-incremental learning. In *CVPR*, pages 23554–23564.
- Wenwen Zhuang, Xin Huang, Xiantao Zhang, and Jin Zeng. 2024. Math-puma: Progressive upward multimodal alignment to enhance mathematical reasoning. *arXiv preprint arXiv:2408.08640*.
- Chengke Zou, Xingang Guo, Rui Yang, Junyu Zhang, Bin Hu, and Huan Zhang. 2024. Dynamath: A dynamic visual benchmark for evaluating mathematical reasoning robustness of vision language models. *arXiv preprint arXiv:2411.00836*.

A Related Work

Multimodal Large Language Models. Multimodal Large Language Models (MLLMs) (Li et al., 2023; Liu et al., 2024a; Sun et al., 2025a; Wang et al., 2024b; Lu et al., 2024b; McKinzie et al., 2024; Sun et al., 2025c,b; Fu et al., 2025b; Dong et al., 2024) integrate vision encoders (Radford et al., 2021; Zhai et al., 2023) with LLMs (AI, 2024; Bai et al., 2023), endowing them with robust capabilities across a wide range of domains. These include general visual understanding (OpenAI, 2024a; Li et al., 2024), mathematical reasoning (Shi et al., 2024; Gao et al., 2023), and answering college-level questions (Chen et al., 2023), demonstrating their versatility in real-world tasks. The rapid advancements in open-source models have also spurred the development of proprietary models, such as GPT-4o (OpenAI, 2024a), Gemini (Team et al., 2023; Reid et al., 2024), Qwen2-VL-MAX (Wang et al., 2024b), and Claude3 (anthropic, 2024). These models have demonstrated remarkable performance in both evaluation benchmarks and practical applications, solidifying their position at the forefront of AI research and deployment.

Reasoning with MLLMs. Recent advancements in MLLMs have significantly enhanced performance in reasoning tasks across both text and multimodal scenarios (OpenAI, 2024b; DeepSeek-AI, 2024; QwenTeam, 2024; Peng et al., 2025). Current methods primarily rely on CoT (Wei et al., 2022) to train MLLMs for step-by-step reasoning. Data-driven approaches include MathLLaVA (Shi et al., 2024), which introduced the MathV360K dataset, and MAMmoTH-VL (Guo et al., 2024), which curates a large-scale multimodal CoT dataset in a scalable manner. Another line of research explores vision-text alignment. MAVIS (Zhang et al., 2024b) fine-tunes a math-specific vision encoder with curated caption data, while Math-PUMA (Zhuang et al., 2024) leverages the Kullback-Leibler (KL) divergence of next-token prediction distributions for modality alignment. In a different paradigm, MLLMs act as coordinators, utilizing external tools such as LLMs, web search engines, and computer programs for complex reasoning. Chameleon (Lu et al., 2023) orchestrates tool-call sequences, and Visual Sketchpad (Hu et al., 2024) enables models to generate visual sketches to aid reasoning.

B More Details of Reasoning Dataset

In this section, we provide a detailed description of dynamic token truncation and reflection word pruning in the process of constructing the reasoning dataset. We also provide detailed information about the training data in Table 5.

B.1 Dynamic Token Truncation

To further improve the dataset quality, we analyze the distribution of token lengths after the answer-centric rejection sampling. We find that many samples are close to the maximum token limit, and manual checks show that these long reasoning chains often have problems—such as logical errors, mistakes in multi-step calculations, and reliance on shortcuts that don’t work in general cases (*e.g.*, substituting specific values). Motivated by the correlation between extreme token lengths and compromised solution quality, we implement adaptive truncation thresholds to keep the answers within the 200-8000 token range. This dynamic filtering not only eliminates the invalid cases (overly verbose or terse responses) but also enhances the overall quality of the data. The final length distribution matches how human experts solve problems and keeps the important reasoning steps intact.

B.2 Reflection Word Pruning

Our analysis reveals a critical flaw in distilled reasoning chains: excessive metacognitive loops caused by uncontrolled reflection markers (*e.g.*, ‘Alternatively,’ ‘Wait’), which led to performance degradation through infinite loops or ungrounded speculation. Term frequency analysis of reflection density shows a heavy-tailed distribution—95% of samples contained fewer than 10 reflection markers per chain, while 1% exhibited over 50 markers, with this extreme group strongly correlating to hallucination rates. To address this, we introduce a reflection token quota system that automatically prunes samples exceeding 25 reflection markers while retaining core reasoning logic using semantic-aware span detection. As shown in Figure 4, this approach significantly reduced infinite-loop instances in validation tasks while improving answer accuracy. The refined reflection pattern mirrors expert human problem-solving strategies, wherein targeted self-correction enhances, rather than disrupts, the continuity of the reasoning process.

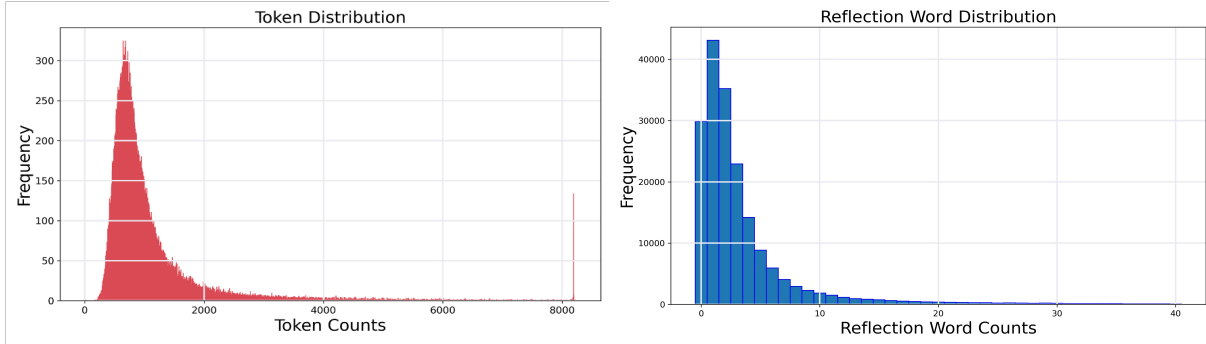


Figure 7: The token and reflection word distribution of the long-chain reasoning dataset.

Table 4: The detailed training hyperparameters.

Config	SFT
Deepspeed	Zero3
Epoch	5
Warmup Ratio	0.1
Max Grad Norm	1.0
Optimizer	AdamW
Learning rate	2e-5
Learning rate scheduler	Cosine
Text max length	8192
Batch size per GPU	1
Gradient Accumulation Steps	4
GPU	64×H20-96G
Precision	Bf16

C Discussion

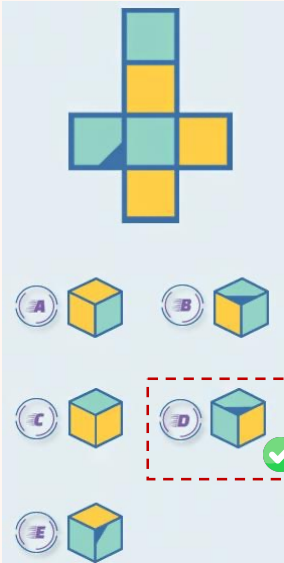
In this work, we investigate the phenomenon of visual information forgetting in MLLMs during long-chain reasoning. Through comprehensive analysis and experiments, we show that as reasoning chains progressively lengthen, models exhibit a gradual deterioration in retaining visual inputs, ultimately undermining their multimodal reasoning capabilities and exacerbating hallucination issues. To address this challenge, we propose a take-along visual conditioning mechanism that enables models to dynamically revisit visual inputs during reasoning steps, thereby enhancing content fidelity throughout the inference process.

However, as illustrated in Figure 8, we acknowledge several limitations. First, for highly complex reasoning tasks requiring sophisticated analytical capabilities, simply increasing visual revisits proves insufficient. In contrast, it is crucial to enhance the model’s inherent reasoning capacity. Second, our method assumes the availability of delayed visual processing, making it potentially unsuitable for real-time applications requiring instantaneous visual feedback, such as robotic navigation or time-sensitive decision-making scenarios.

Table 5: Details on the TVC’s training data, which is derived from publicly available datasets.

Datasets	Samples
MathV360K (Shi et al., 2024)	221K
Geo170K (Gao et al., 2023)	22K
LLaVA-OneVision (Li et al., 2024)	97K
Cambrian-1 (Tong et al., 2024)	1K

Our work represents an initial exploration into mitigating visual forgetting in extended multimodal reasoning chains. We envision future research directions including: (1) Developing hybrid architectures that synergistically enhance both visual retention and intrinsic reasoning capabilities; (2) Investigating adaptive attention mechanisms for real-time multimodal applications; (3) Exploring curriculum learning strategies to progressively strengthen long-chain reasoning capacities. We hope this foundational study will inspire further advances in understanding and improving multimodal reasoning systems for complex real-world applications.



Which of the cubes is not identical to the unfolded net?
 Choices: (A) A (B) B (C) C (D) D (E) E

Let's see here. I've got this problem where there's an unfolded net of a cube, and then there are several cubes shown below with letters A through E. The question is asking which one of these cubes is not identical to the unfolded net. Okay, so I need to compare each cube to the net and see if they match or not.

.....

Let me label the squares for clarity:

1. Front face (F)
2. Back face (B)
3. Top face (T)
4. Bottom face (D)
5. Left face (L)
6. Right face (R)

.....

Let's look at each cube option:

A) This is a cube with letters on its faces. Wait, no, it's a cube with colors, similar to the others. Wait, actually, all cubes have combinations of teal and yellow faces. I need to carefully observe the color placement on each cube.

Option A:

- Front face: teal
- Top face: yellow
- Right face: teal
- Left face: yellow
- Back face: teal
- Bottom face: yellow

.....

Option E:

- Front: teal
- Top: yellow
- Right: teal
- Left: yellow
- Back: teal
- Bottom: yellow

Comparing to the net: F(t), T(y), R(t), L(y), B(y), D(t)

So, according to the net, B should be yellow, but in Option E, it's teal. Therefore, Option E differs from the net in the back face color. Hence, Option E is not identical to the unfolded net.

Therefore, the answer is E.

Figure 8: Qualitative Results of TVC.