

Point-RFT: Improving Multimodal Reasoning with Visually Grounded Reinforcement Finetuning

Minheng Ni^{1,2 †}, Zhengyuan Yang^{3 †}, Linjie Li³, Chung-Ching Lin³, Kevin Lin³,
Wangmeng Zuo^{2 ☒}, Lijuan Wang^{3 ☒}
¹Hong Kong Polytechnic University ²Harbin Institute of Technology ³Microsoft

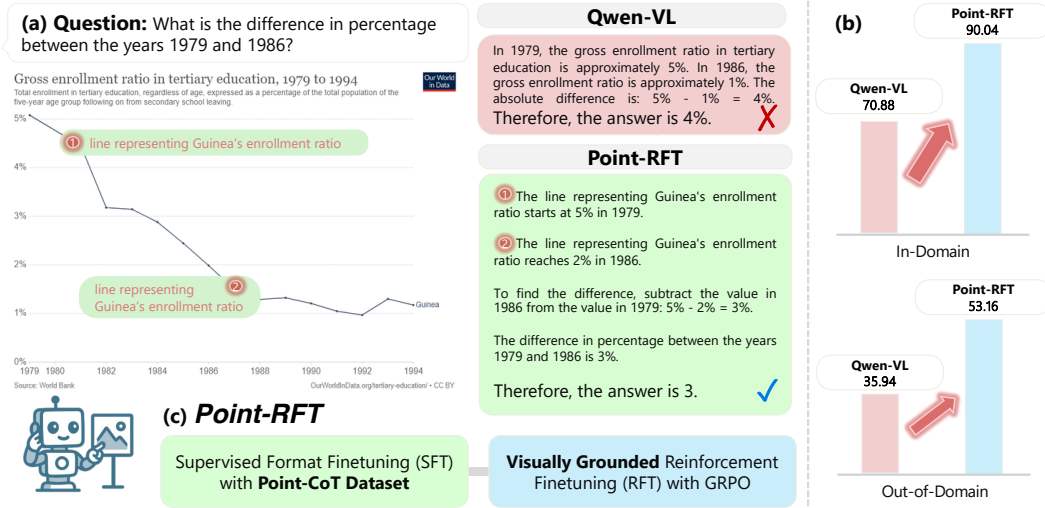


Figure 1: **Point-RFT improves multimodal reinforcement finetuning with visually grounded CoT.** (1) We first construct a Point-CoT dataset for Supervised Format Finetuning (SFT). It allows the model to generate step-by-step reasoning traces explicitly linked to visual pointing, mitigating hallucinations and enhancing perception-reasoning alignment. (2) Reinforcement Finetuning (RFT) with GRPO: Optimizes answer correctness and grounded rationale coherence by rewarding localized visual-textual reasoning paths. Our model, code, and dataset can be found at this link.

Abstract

Recent advances in large language models have significantly improved textual reasoning through the effective use of Chain-of-Thought (CoT) and reinforcement learning. However, extending these successes to vision-language tasks remains challenging due to inherent limitations in text-only CoT, such as visual hallucinations and insufficient multimodal integration. In this paper, we introduce Point-RFT, a multimodal reasoning framework explicitly designed to leverage visually grounded CoT reasoning for visual document understanding. Our approach consists of two stages: First, we conduct format finetuning using a curated dataset of 71K diverse visual reasoning problems, each annotated with detailed, step-by-step rationales explicitly grounded to corresponding visual elements. Second, we employ reinforcement finetuning targeting visual document understanding. On ChartQA, our approach improves accuracy from 70.88% (format-finetuned baseline) to 90.04%, surpassing the 83.92% accuracy achieved by reinforcement finetuning relying solely on text-based CoT. The result shows that our grounded CoT is more effective for multimodal reasoning compared with the text-only CoT. Moreover, Point-RFT exhibits superior generalization capability across several out-of-domain visual document reasoning benchmarks, including CharXiv, PlotQA, IconQA, TabMWP, *etc.*, and highlights its potential in complex real-world scenarios.

1 Introduction

Large reasoning models (o1b, 2024; Guo et al., 2025) have shown impressive capabilities in solving complex text-based problems, such as STEM reasoning (Lightman et al., 2023; Rein et al., 2024) and coding tasks (Jain et al., 2024b; Quan et al., 2025). Central to their success is the large language model’s ability to improve responses by expanding the textual Chain-of-Thought (CoT) (Wei et al., 2022), also known as scaling inference-time compute (Brown et al., 2024; Snell et al., 2024). Reinforcement learning (Mnih et al., 2015, 2016; Schulman et al., 2017; Shao et al., 2024) further optimizes these models by maximizing the rewards associated with these generated longer, more comprehensive, and accurate responses. Despite these successes in textual reasoning, adapting them to vision-language tasks remains challenging, with unique open challenges such as imperfect visual perception (Fu et al., 2024) as well as integrating reasoning across both visual and text modalities (Hao et al., 2025).

Inspired by the introduction of visual-centric tokens in vision-language models (Yang et al., 2022a; Jain et al., 2024a; Dong et al., 2023; Yang et al., 2022b; Wang et al., 2022b) to improve visual perception, we present Point-RFT, with the insight that the effective CoT for multimodal reasoning should contain grounded visual tokens. To achieve this, we first perform format finetuning across diverse problem sets to enable Point-RFT to generate grounded CoT rationales. Subsequently, the format-finetuned model undergoes reinforcement finetuning with outcome reward on specialized visual document datasets, allowing it to explore the effective usage of grounded CoT in solving visual document questions.

To curate our format finetuning dataset, we establish a pipeline that integrates step-by-step text rationales generated by a large multimodal model (Hurst et al., 2024) with a grounding model (Deitke et al., 2024) that links relevant reasoning steps explicitly to visual points. The resulting dataset comprises 71K images covering a wide range of question types (Guo et al., 2024). This diversity in the dataset design ensures that the model generalizes its learned formatting strategies broadly, allowing it to explore effective solutions for various problems during subsequent reinforcement finetuning. After the format finetuning for grounded CoT, we conducted reinforcement finetuning (Shao et al., 2024) specifically on the ChartQA dataset (Masry et al., 2022). Experimental results demonstrate notable performance improvements: accuracy increased from 87.21% to 90.04% compared to the Supervised finetuning (SFT) baseline and from 83.92% to 90.04% compared to reinforcement learning using text-only CoT. These results show the effectiveness of our Point-RFT and the grounded CoT approach for enhancing multimodal reasoning capabilities.

Beyond the performance improvements on ChartQA, we note that Point-RFT demonstrates several compelling properties for multimodal reasoning: **(1)** Improved generalization capabilities compared to supervised finetuning and reinforcement learning with text-only CoT, with superior performance on various out-of-domain tasks including ChartXiv (Wang et al., 2024), PlotQA (Methani et al., 2020), IconQA (Lu et al., 2021), TabMWP (Lu et al., 2022b), *etc.* **(2)** Enhanced pointing accuracy not only contributes directly to higher-quality final answers but also provides increased interpretability. This capability allows the disentanglement of perception errors from reasoning errors, facilitating more precise diagnostics and targeted improvements.

Our contributions are summarized as follows.

- We present, Point-RFT, the first multimodal reasoning framework explicitly designed for visually grounded reinforcement fine-tuning. We empirically show that “visually grounded CoT” is more effective for multimodal reasoning compared with text-only thoughts.
- We curate a 71K-example dataset where every reasoning step is aligned with point-level visual references, enabling supervised format finetuning that teaches the model to “think while pointing.” We then perform RFT with grounded CoT.
- On ChartQA, Point-RFT boosts accuracy from 70.88 % to 90.04 %. It also achieves the best average score over five out-of-domain benchmarks, demonstrating superior generalization ability and interpretability.

Point-CoT

Question: Compare the levels of satisfaction with democracy across different countries and regions. Discuss how cultural, political, and economic contexts influence these levels of satisfaction.



GT Answer: Satisfaction with democracy can vary significantly across different countries and regions, influenced by a multitude of cultural, political, and economic factors. For instance, countries with a long history of stable democratic governance and strong institutions tend to have higher levels of satisfaction. In contrast, regions experiencing political turmoil, corruption, or economic instability often report lower satisfaction levels. Cultural factors, such as the value placed on individual freedoms and social equality, also play a role. Additionally, economic prosperity and equitable wealth distribution can positively impact satisfaction with democracy, while economic inequality and poverty can lead to disillusionment. Understanding these variations can provide insights into how to improve democratic processes and enhance citizen satisfaction globally.

Figure 2: **Visualization of Point-CoT dataset.** Point-CoT dataset integrates the reasoning process of answering questions with point grounding, creating a novel form of multimodal CoT.

2 Related Work

Large Reasoning Language Models. Recent studies (o1b, 2024; Guo et al., 2025) have demonstrated that reinforcement learning (Mnih et al., 2015, 2016; Schulman et al., 2017; Shao et al., 2024) can significantly enhance the reasoning capabilities of language models (Lightman et al., 2023; Rein et al., 2024; Jain et al., 2024b; Quan et al., 2025). The approach involves the use of chain-of-thought (CoT) reasoning to systematically generate high-quality responses, coupled with outcome-based reward optimization, enabling models to explore and identify effective solutions. However, a purely textual CoT may be sub-optimal for vision-centric (Fu et al., 2024; Rahmanzadehgervi et al., 2024; Hua et al., 2024; Tong et al., 2024) and vision-language tasks (Yu et al., 2023, 2024b; Hao et al., 2025). In this work, we propose and investigate a grounded CoT approach to improve multimodal reasoning.

Vision-Language Models. Advancements in vision-language modeling (VLM) (Radford et al., 2021; Yuan et al., 2021; Wang et al., 2022c,a; Yu et al., 2022; Li et al., 2024; Hu et al., 2022) have significantly enhanced the joint understanding of visual and textual inputs, benefiting various downstream tasks such as image captioning (Chen et al., 2015) and visual question answering (Goyal et al., 2017). Recently, contemporary VLMs (Alayrac et al., 2022; OpenAI, 2023a; Yang et al., 2023; Liu et al., 2023; Wang et al., 2023; Chen et al., 2024; Bai et al., 2023; Team, 2023) have integrated multimodal modeling techniques with large language models (LLMs) (OpenAI, 2023b; Grattafiori et al., 2024; Yang et al., 2024), resulting in enhanced capabilities including instruction-following, in-context learning, and improved zero-shot generalization. Nevertheless, the application of reinforcement learning to enhance multimodal reasoning (Xiyao et al., 2024; Zhou et al., 2024; Yu et al., 2024a; Sun et al., 2023; Wang et al., 2025a) remains relatively underexplored due to the unique challenges associated with extending outcome reward-based reinforcement learning approaches from purely textual to multimodal.

Multimodal Chain-of-Thought. Multimodal chain-of-thought reasoning Wang et al. (2025b) has recently attracted intense attention. Early works focus on text-only thoughts for multimodal problems Lu et al. (2022a); Zhang et al. (2024); Xu et al. (2024), showing that explicit step-by-step reasoning boosts visual-question-answering accuracy. More recent studies extend this idea to true multimodal thoughts Rose et al. (2023); Wu et al. (2024); Hu et al. (2024); Fu et al. (2025); Ma et al. (2024), including having models directly generating visualizations, as well as calling vision tools that generate crops and other augmented views. Despite this progress, all existing multimodal chain-of-thought approaches remain predominantly prompt-driven, rather than end-to-end RFT, as explored in this study.

3 Method

Our core hypothesis is that reinforcement finetuning with grounded chain-of-thought (CoT) reasoning enhances visual-language reasoning capabilities. To investigate this, we propose a two-stage training framework: (1) supervised fine-tuning (SFT) with format-finetuning dataset, to support grounded CoT, followed by (2) GRPO-based RL optimization with outcome reward signals. Section 3.1 introduces the format finetuning data collection pipeline, and Section 3.2 details the model training approach.

Table 1: **Statistics of dataset.** The final dataset comprises 71K images with diverse question type. This mixture ensures coverage of both free-form and structured visual reasoning scenarios.

Source	Samples	Ratio	Turn	Points
ChartQA	29533	41.5%	3.97	2.02
DVQA	25269	35.5%	8.78	5.79
PlotQA	14055	19.8%	8.71	5.89
CLEVR (MathV360K)	1322	1.9%	4.61	3.43
TallyQA	931	1.3%	2.93	2.01

3.1 Point-CoT: Format-Finetune Dataset

To enable models to explicitly understand and interact with content in images during multimodal reasoning, we propose the Point-CoT dataset. As shown in Figure 2, the Point-CoT dataset integrates the reasoning process of answering questions with point grounding, creating a novel form of multimodal chain-of-thought (CoT).

Our data generation pipeline addresses two critical challenges in visual-language reasoning: 1) synthesizing logically coherent reasoning chains with explicit visual grounding and 2) ensuring consistency between textual rationales and spatial referring. We design a hybrid pipeline that combines GPT-4o and Molmo-7B, leveraging their complementary strengths: GPT-4o excels at generating human-like reasoning patterns, while Molmo-7B specializes in fine-grained visual localization due to its pretraining on large-scaled geometric annotation tasks. Figure 3 illustrates the overall dataset generation pipeline.

Step 1: CoT Generation with GPT-4o. For each image-question-answer triplet from Mammoth-VL datasets, GPT-4o generates comprehensive, step-by-step reasoning traces with text tokens only. Each step may explicitly reference visual elements (*e.g.*, chart components) crucial for deriving the correct and convincing answer. For the referenced visual elements, the model will be required to generate their descriptions, *i.e.*, `<point>...</point>` and the number of objects, `<point_result>...</point_result>`.

Step 2: Points Grounding Generation with Molmo-7B. We then use Molmo-7B to annotate coordinates of points $\{x_0, y_0, x_1, y_1, x_2, y_2, \dots\}$ for visual elements mentioned in each reasoning step. Coordinates follow raw image resolution rather than normalized scales, and points are formatted as `<points x0="..." y0="..." ...>description</points>` preceding each reasoning step. After point generation, cross-validation on counts of points between GPT-4o rationales, *i.e.*, number in `<point_result>...</point_result>` and Molmo-7B annotations will be conducted, and we will discard points with mismatches after eight retries.

The validation process improves the integrity and accuracy of the data, enhancing the quality of the Chain-of-Thought (CoT) reasoning from two perspectives. First, by using Molmo-7B for visual grounding, hallucinated elements in GPT-4o’s reasoning are filtered out because they will fail the grounding process. Second, if Molmo-7B fails to ground or provides incorrect annotations, the number of visual elements mentioned in GPT-4o’s text (*e.g.*, “three bars in the chart”) will not match the number of instances grounded by Molmo-7B. Thus, if a model’s visual perception is accurate, the resulting CoT reasoning will also be accurate. This rigorous process reduces annotation errors compared to single-model approaches.

The final dataset comprises 71K images with diverse question types (counting, comparison, arithmetic) from the subset of Mammoth-VL. This mixture ensures coverage of both free-form and structured visual reasoning scenarios.

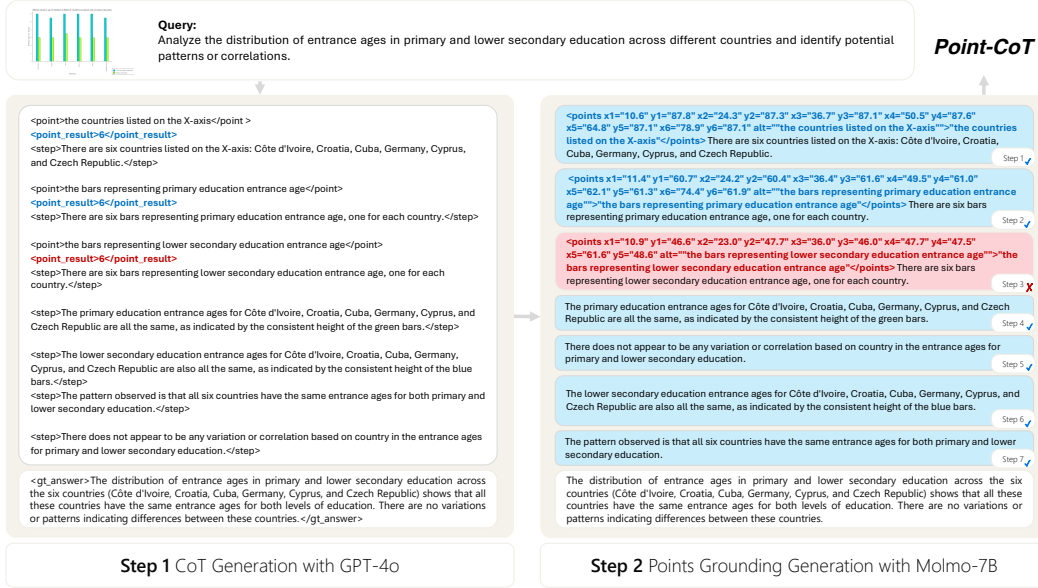


Figure 3: **Overall dataset generation pipeline.** The whole construction process combining LLM reasoning (GPT-4o) and geometric grounding (Molmo-7B). The pipeline ensures spatial-textual consistency through cross-validation, producing our Point-CoT dataset.

3.2 Model Training

Supervised Format Finetuning (SFT) with Point-CoT. We fine-tune a Qwen2.5-VL as the base model using our format finetuning dataset. Each training instance adheres to the following format:

```
<think>
<points ...>...</points>...
<points ...>...</points>...
...
</think>
<answer>...</answer>
```

The template design serves three purposes: 1) explicitly separating visual grounding from textual reasoning, 2) enabling deterministic parsing of coordinate references during evaluation, and 3) forcing the model to attend to spatial features before generating answers. Our ablation studies indicate that interleaving point annotations with textual content significantly enhance answer accuracy, suggesting that the model benefits from the alignment between spatial referring and thinking in grounded CoT.

Let θ be model's parameter and T be the length of the text. The loss function is:

$$\mathcal{L}(\theta) := -\mathbb{E}_{(x,y) \sim \mathcal{D}_{\text{Point-CoT}}} \sum_{t=1}^T \log P(y_t | x, y_{<t}; \theta), \quad (1)$$

where (x, y) is the query and target response with thinking in dataset $\mathcal{D}_{\text{Point-CoT}}$.

Reinforcement Finetuning (RFT) with GRPO. Despite obtaining point-grounded reasoning capability, initial SFT may cause performance drops in OOD domains, suggesting overfitting to synthetic reasoning patterns. To address this, we apply Reinforcement Learning (RL) to further optimize it.

We use the Group-wise Relative Policy Optimization (GRPO) algorithm (Shao et al., 2024) as the optimization method, which extends PPO by introducing task-specific advantage normalization. We implement GRPO with dual rewards to tackle the tension between structural and semantic correctness:

- **Format Reward (R_f)** measures structural adherence to the `<think> ... </think>` `<answer> ... </answer>` template via regex matching. It employs the grammar-level regex parser that checks 1) thinking process integrity and 2) whether the answer can be extracted from the answer tag.

- **Accuracy Reward** (R_a) computes answer correctness against ground-truth answer using

$$R_a := \mathbb{I}(\hat{y} = y). \quad (2)$$

The combined reward $\hat{R} := R_f + R_a$ is optimized on 15K ChartQA vanilla training data by:

$$\mathcal{J}(\theta) := \mathbb{E}_{(x,y) \sim \mathcal{D}_{\text{ChartQA}}} \frac{\pi_\theta}{\pi_{\theta'}} \hat{R} - \beta \cdot D_{\text{KL}}(\pi_\theta \| \pi_{\text{SFT}}), \quad (3)$$

where β is a hyper-parameter of KL divergence, and π_{SFT} , π_θ , and $\pi_{\theta'}$ is the model after SFT, the optimized model and the model before GRPO optimization.

4 Experiment

4.1 Implementation Details

All models use AdamW optimizer with 5×10^{-5} learning rate, and 512 batch size for SFT and 56 batch size for RL. Training converges in 500 steps for SFT and 100 RL steps with $\beta = 0.00$. We use soft matching for numeric answers (tolerating $\pm 5\%$ relative error) and exact matching for other responses in GRPO. We implement our two-stage training pipeline using PyTorch with $8 \times \text{A100}$ GPUs based on Easy-R1 (Zheng et al., 2025).

4.2 Experimental Setup

We evaluate our approach on six multimodal reasoning benchmark datasets spanning diverse domains:

- **ChartQA** (Masry et al., 2022): A chart understanding task requiring extraction and reasoning over chart elements. We use the official test split for evaluation.
- **CharXiv** (Wang et al., 2024): Scientific charts from academic papers with complex visual encodings. We use the official validation split for evaluation.
- **PlotQA** (Methani et al., 2020): Statistical plots testing trend analysis and numerical reasoning. We randomly sampled 2,000 examples as the test split.
- **IconQA** (Lu et al., 2021): Iconographic reasoning requiring symbolic interpretation. We randomly sampled 2,000 examples as the test split.
- **TabMWP** (Lu et al., 2022b): A task involving tabular math problems that combine table parsing and arithmetic reasoning. We use the official 1,000-sample mini-test split.
- **Counting** (Li et al., 2023): Based on the SuperCLEVR dataset, we sampled 200 examples for evaluation.

We compare the following approaches: (1) **Base-RFT**: A baseline model trained with supervised finetuning on Point-CoT without explicit point annotations, followed by reinforcement learning finetuning with GRPO on ChartQA. This serves as our primary baseline. (2) **Point-SFT**: A model trained with supervised finetuning on Point-CoT with explicit point annotations. (3) **Point-RFT**: A model based on Point-SFT that undergoes reinforcement learning finetuning with GRPO on ChartQA. (4) **Base**: The backbone model without any finetuning is also included in the tables as a reference.

For the base model and Base-RFT, we only require the model to output the reasoning process and the final answer. For Point-SFT and Point-RFT, we additionally require the model to output a reasoning process with explicit point annotations, along with the final answer.

To evaluate all models, we extract the content from the answer tag in their output and compare it to the ground truth to calculate the answer accuracy (Overall) as the primary evaluation metric. In ablation studies, we further analyze the format compliance (Format) and the accuracy of answers within compliant formats (Inner). Note that all models are evaluated with CoT.

4.3 Quantitative Results

4.3.1 Main Results

Table 3 presents our main findings across in-domain and out-of-domain test sets. Our Point-RFT achieves state-of-the-art performance on the primary ChartQA benchmark, significantly outperform-

Table 2: **Overall results among different datasets.** Comparison of different training variants on in-domain (ChartQA) and out-of-domain test sets. Point-RFT achieves state-of-the-art in-domain performance (90.04%) and demonstrates strongest OOD generalization (53.16% average), significantly outperforming baselines. The improvement is statistically significant with $p < 0.01$ under t -test.

Method	Backbone	Setting		In-Domain ChartQA	Out-of-Domain					Avg.
		Point	RL		CharXiv	PlotQA	IconQA	TabMWP	Counting	
Base	Qwen2.5-VL-7B	-	-	70.88	26.50	17.80	53.40	61.00	21.00	35.94
Base-RFT	Qwen2.5-VL-7B	-	✓	83.92	25.20	14.85	48.55	61.40	19.00	33.80
Point-SFT		✓	-	87.21	28.00	20.34	55.25	60.30	70.00	46.78
Point-RFT		✓	✓	90.04	36.20	20.40	59.80	70.90	78.50	53.16

Table 3: **Overall results among different models.** The improvement is statistically significant with $p < 0.01$ under t -test.

Method	Backbone	Setting		In-Domain ChartQA	CharXiv	PlotQA	Out-of-Domain		Counting	Avg.
		CoT	Point				IconQA	TabMWP		
Base	Qwen2-VL-7B	✓	-	55.40	25.20	10.50	49.15	61.00	68.00	42.77
Point-RFT	Qwen2-VL-7B	✓	✓	86.16	28.30	18.80	46.90	61.60	68.50	44.82
Base	Qwen2.5-VL-3B	✓	-	79.28	27.50	10.80	44.75	61.00	40.50	36.91
Point-RFT	Qwen2.5-VL-3B	✓	✓	86.24	28.80	17.90	51.00	62.20	60.50	44.08
Base	Qwen2.5-VL-7B	✓	-	70.88	26.50	17.80	53.40	61.00	21.00	35.94
Point-RFT	Qwen2.5-VL-7B	✓	✓	90.04	36.20	20.40	59.80	70.90	78.50	53.16

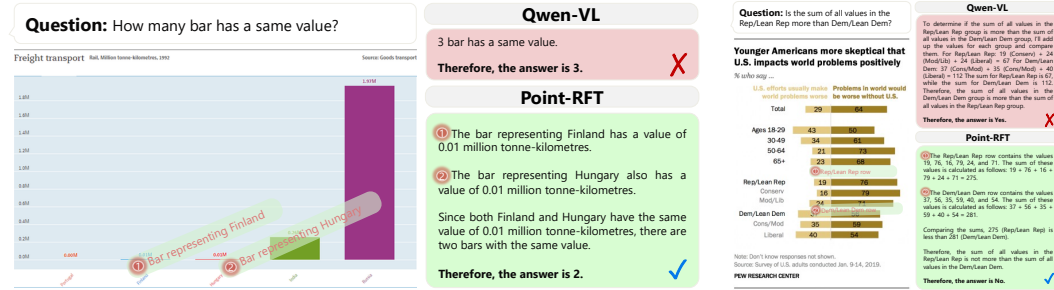


Figure 4: **Cases of In-domain Chart.** These cases highlight the limitations of pure text reasoning in analyzing complex visual elements. Point-RFT successfully reasons by integrating the visual content.

ing the Point-SFT model and text-only reinforcement learning variants, Base-RFT. More importantly, we observe substantial improvements in out-of-domain generalization.

In-Domain Performance. On ChartQA, Point-RFT achieves an accuracy of 90.04%, showing a significant improvement over Point-SFT (87.21%), which demonstrates the importance of RL in helping the model learn point-based reasoning. Compared with the text-only reinforcement learning variant (83.92%), the absolute improvement validates that visual grounding significantly enhances reasoning capability through better perception-reasoning alignment. Our method also outperforms the base instruction-tuned Qwen model (70.88%). Since the base model does not perform CoT reasoning or visual localization, Point-RFT can learn advanced visual reasoning patterns.

Out-of-Domain Generalization. Point-RFT achieves an average accuracy of 53.16% across five out-of-domain benchmarks, surpassing the text-only reinforcement learning variant by 19.36% absolute. Notably, the gains on PlotQA (+6.55%) and IconQA (+11.25%) suggest that explicit visual references facilitate the transfer of reasoning patterns to novel visual layouts. Furthermore, similar to the in-domain results, Point-RFT significantly outperforms Point-SFT, demonstrating the critical role of RL in acquiring this unique capability. Additionally, we observe that the OOD performance is comparable to the backbone model that undergoes large-scale multi-task instruction finetuning, further validating the generalization capability of Point-RFT.

For extra comparisons, please refer to the **Supplementary Material**.

4.3.2 Ablation Studies

Pointing Format. Table 4 compares the XML and JSON syntax formats and the impact of point indexing schemes on performance. XML syntax achieves an accuracy of 56.72%, while JSON

Table 4: **Ablation studies of pointing format.** Different coordinate formats and indexing schemes highly influence performance.

Method	Type		ChartQA		
	Format	Point	Overall	Inner	Format
Base	-	-	79.28	82.17	96.48
Point-SFT	JSON	✓	49.36	71.88	69.56
	XML	Random	31.32	75.82	39.72
	XML	✓	56.72	76.20	74.44

Table 5: **Ablation studies of point grounding.** Removing visual grounding reduces performance.

Method	Training			ChartQA		
	Point	SFT	RL	Overall	Inner	Format
Base	-	-	-	79.28	82.17	96.48
Base-SFT	-	✓	-	27.60	75.82	36.40
Point-SFT	✓	✓	-	56.72	76.20	74.44
Base-RFT	-	✓	✓	77.16	79.61	96.92
Point-RFT	✓	✓	✓	86.24	86.52	99.68

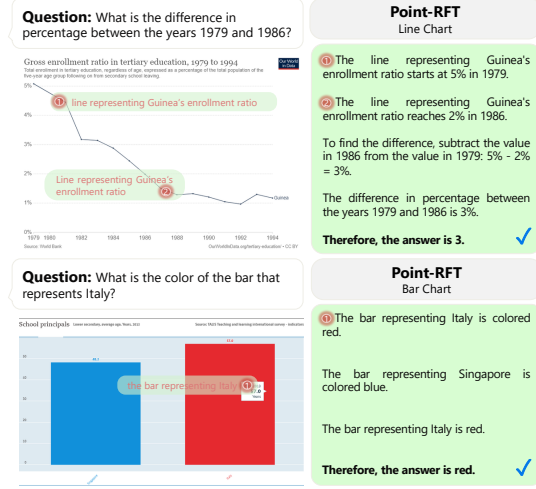


Figure 5: **Case of OOD Chart.** Point-RFT successfully transfers coordinate referencing skills learned from bar charts.



Figure 6: **Cases of different RL steps.** As the steps increased, accurate localization led to correct answers.

achieves 49.36%. This suggests that for more complex reasoning, the XML format better aids the model in generating accurate outputs. Surprisingly, the indexing scheme (0-based vs. 1-based) shows a significant difference (56.72% vs. 31.32%). Interestingly, Qwen’s default learned indexing starts from 1. We hypothesize that introducing appropriate variation helps the model better learn the ability to perform pointing during reasoning.

Explicit Grounding. We find that removing visual pointing (Table 5) results in severe performance degradation: a 28.5% absolute drop in accuracy for SFT (from 56.72% to 27.60%) and a 3.88% drop for RL (from 81.04% to 77.16%). This validates our hypothesis that explicit visual references are crucial for both initial format learning and subsequent reinforcement training. The 98.32% format compliance rate indicates that structured pointing stabilizes RL optimization, reducing hallucinated reasoning steps. This is critical for real-world applications, where predictable output formats facilitate downstream integration.

For more ablation studies, please refer to the **Supplementary Material**.

4.4 Qualitative Results

4.4.1 Chart Reasoning

In-domain Chart To investigate how visually grounded reasoning impacts the model’s inference, we selected two groups of cases to observe the differences between our reasoning and our baseline model Base-RFT in Figure 4. We observe that our model successfully acquires the ability to perform step-by-step reasoning and locate corresponding visual elements, enabling Point-RFT to arrive at the correct answers in these two high-information-density examples. This demonstrates the strengths of Point-RFT, particularly its systematic exploration of visual evidence.

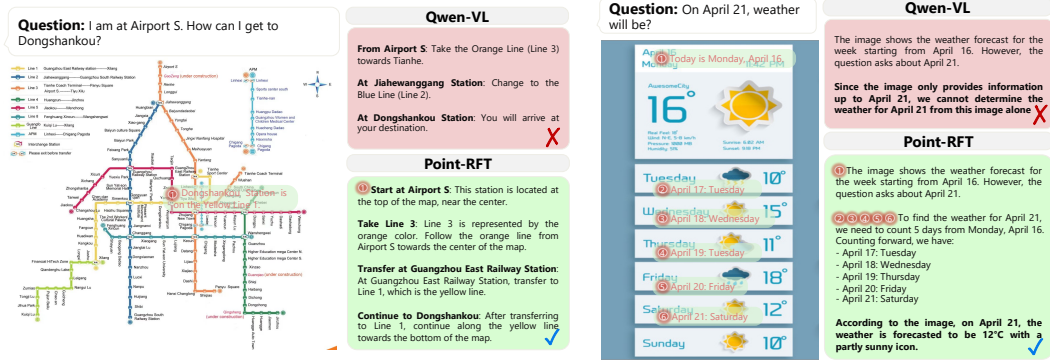


Figure 7: **Cases of generalization.** Point-RFT can perform reasoning on more complex graphics, showcasing its generalization capability.

Out-of-domain Chart Figure 5 illustrates how CoT-based explicit grounding facilitates OOD adaptation. When faced with chart types lacking explicit visual elements (*e.g.*, line charts), Point-RFT successfully transfers coordinate referencing skills learned from bar charts. Despite consistent descriptions of visual elements across steps, the model correctly identified different locations. Meanwhile, the model was able to accurately locate bars containing special shapes and text. This suggests that explicit grounding helps abstract visual reasoning from domain-specific details.

4.4.2 Training Steps

Figure 6 tracks the reasoning quality across different reinforcement learning iterations. In the early stages (10 steps), identified points gradually deviated from their correct positions in the chart, leading to reasoning errors. By the final model (100 steps), the correct positions were successfully located, demonstrating concise and evidence-based reasoning chains. Additionally, we observed that more steps led to more diverse but equally correct reasoning steps in textual form. This evolution underscores the critical role of reinforcement learning in extracting effective grounded reasoning strategies from the exploration space.

4.4.3 Generalization

Not only can Point-RFT operate effectively in abstract charts, we observe its equal efficacy in real-world scenarios. As illustrated in Figure 7 through two representative examples: In the first case, point-grounded reasoning significantly enhances the model’s analytical capabilities when processing complex traffic diagrams, enabling it to recognize fine-grained visual elements and perform logical reasoning. In the second example, the pointing mechanism facilitates explicit expression and reasoning about latent content, allowing accurate interpretation of scenarios that would otherwise be incomprehensible to the original model.

For further analysis on dataset, please refer to the **Supplementary Material**.

5 Conclusion

We have presented Point-RFT, a visually grounded Chain-of-Thought framework designed to enhance multimodal document reasoning. Our results demonstrate that structured visual grounding significantly improves accuracy and interpretability in multimodal reasoning, improving the ChartQA accuracy from 70.88% to 90.04%. Among various compared baselines, Point-RFT shows superior generalization capabilities across diverse, out-of-domain benchmarks such as CharXiv, PlotQA, and TabMWP. Additionally, we provide a visually grounded CoT dataset containing 71K samples, facilitating easy format finetuning for future studies on multimodal reasoning. Our findings highlighted the potential on disentanglement of perception errors from reasoning errors, facilitating more precise diagnostics and targeted improvements for large multimodal models.

For statements of broader impact and limitations, please refer to the **Supplementary Material**.

Acknowledgment

This work was supported by National Key R&D Program of China under Grant No. 2022YFA1004100.

References

- Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, 2024.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*, 2024.
- Liang Chen, Lei Li, Haozhe Zhao, Yifan Song, and Vinci. R1-v: Reinforcing super generalization ability in vision language models with less than \$3. <https://github.com/Deep-Agent/R1-V>, 2025. Version v0.2 released 3 Feb 2025.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198, 2024.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, et al. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. *arXiv preprint arXiv:2409.17146*, 2024.
- Yihe Deng, Hritik Bansal, Fan Yin, Nanyun Peng, Wei Wang, and Kai-Wei Chang. Openvlthinker: An early exploration to complex vision-language reasoning via iterative self-improvement, 2025.
- Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023.
- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not perceive. In *European Conference on Computer Vision*, pages 148–166. Springer, 2024.
- Xingyu Fu, Minqian Liu, Zhengyuan Yang, John Corring, Yijuan Lu, Jianwei Yang, Dan Roth, Dinei Florencio, and Cha Zhang. Refocus: Visual editing as a chain of thought for structured image understanding. *arXiv preprint arXiv:2501.05452*, 2025.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Jarvis Guo, Tuney Zheng, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Yizhi Li, Graham Neubig, Wenhui Chen, and Xiang Yue. Mammoth-vl: Eliciting multimodal reasoning with instruction tuning at scale. *arXiv preprint arXiv:2412.05237*, 2024.

- Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and Yu Cheng. Can mllms reason in multimodality? emma: An enhanced multimodal reasoning benchmark. *arXiv preprint arXiv:2501.05444*, 2025.
- Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17980–17989, 2022.
- Yushi Hu, Weijia Shi, Xingyu Fu, Dan Roth, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, and Ranjay Krishna. Visual sketchpad: Sketching as a visual chain of thought for multimodal language models. *arXiv preprint arXiv:2406.09403*, 2024.
- Hang Hua, Yunlong Tang, Ziyun Zeng, Liangliang Cao, Zhengyuan Yang, Hangfeng He, Chenliang Xu, and Jiebo Luo. Mmcomposition: Revisiting the compositionality of pre-trained vision-language models. *arXiv preprint arXiv:2410.09733*, 2024.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models, 2025.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Jitesh Jain, Zhengyuan Yang, Humphrey Shi, Jianfeng Gao, and Jianwei Yang. Ola-vlm: Elevating visual perception in multimodal llms with auxiliary embedding distillation. *arXiv preprint arXiv:2412.09585*, 2024a.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024b.
- Chunyu Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, Jianfeng Gao, et al. Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and Trends® in Computer Graphics and Vision*, 16(1-2):1–214, 2024.
- Zhuowan Li, Xingrui Wang, Elias Stengel-Eskin, Adam Kortylewski, Wufei Ma, Benjamin Van Durme, and Alan L Yuille. Super-clevr: A virtual benchmark to diagnose domain robustness in visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14963–14973, 2023.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.
- Haotian Liu, Chunyu Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning. *arXiv preprint arXiv:2110.13214*, 2021.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022a.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv preprint arXiv:2209.14610*, 2022b.
- Zixian Ma, Jianguo Zhang, Zhiwei Liu, Jieyu Zhang, Juntao Tan, Manli Shu, Juan Carlos Niebles, Shelby Heinecke, Huan Wang, Caiming Xiong, et al. Taco: Learning multi-modal action models with synthetic chains-of-thought-and-action. *arXiv preprint arXiv:2412.05479*, 2024.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Botian Shi, Wenhai Wang, Junjun He, Kaipeng Zhang, et al. Mm-eureka: Exploring visual aha moment with rule-based large-scale reinforcement learning. *arXiv preprint arXiv:2503.07365*, 2025.
- Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536, 2020.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PmLR, 2016.
- OpenAI. Gpt-4v(ision) system card. 2023a.
- OpenAI. Gpt-4 technical report, 2023b.
- Yingzhe Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, Qipeng Zhu, Kai Yang, Xingzhong Xu, Xin Geng, and Xu Yang. Lmm-r1: Empowering 3b llms with strong reasoning abilities through two-stage rule-based rl. *arXiv preprint arXiv:2503.07536*, 2025.
- Shanghaoran Quan, Jiayi Yang, Bowen Yu, Bo Zheng, Dayiheng Liu, An Yang, Xuancheng Ren, Bofei Gao, Yibo Miao, Yunlong Feng, et al. Codeelo: Benchmarking competition-level code generation of llms with human-comparable elo ratings. *arXiv preprint arXiv:2501.01257*, 2025.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- Pooyan Rahmanzadehgervi, Logan Bolton, Mohammad Reza Taesiri, and Anh Totti Nguyen. Vision language models are blind. In *Proceedings of the Asian Conference on Computer Vision*, pages 18–34, 2024.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
- Daniel Rose, Vaishnavi Himakunthala, Andy Ouyang, Ryan He, Alex Mei, Yujie Lu, Michael Saxon, Chinmay Sonar, Diba Mirza, and William Yang Wang. Visual chain of thought: bridging logical gaps with multimodal infillings. *arXiv preprint arXiv:2305.02317*, 2023.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*, 2023.
- Google Gemini Team. Gemini: A family of highly capable multimodal models, 2023.
- Peter Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Adithya Jairam Vedagiri IYER, Sai Charitha Akula, Shusheng Yang, Jihan Yang, Manoj Middepogu, Ziteng Wang, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *Advances in Neural Information Processing Systems*, 37:87310–87356, 2024.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022a.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340. PMLR, 2022b.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023.
- Xiyao Wang, Zhengyuan Yang, Chao Feng, Hongjin Lu, Linjie Li, Chung-Ching Lin, Kevin Lin, Furong Huang, and Lijuan Wang. Sota with less: Mcts-guided sample selection for data-efficient visual reasoning self-improvement. *arXiv preprint arXiv:2504.07934*, 2025a.

- Yaoting Wang, Shengqiong Wu, Yuecheng Zhang, Shuicheng Yan, Ziwei Liu, Jiebo Luo, and Hao Fei. Multi-modal chain-of-thought reasoning: A comprehensive survey. *arXiv preprint arXiv:2503.12605*, 2025b.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *ICLR*, 2022c.
- Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang, Xindi Wu, Haotian Liu, Sadhika Malladi, et al. Charxiv: Charting gaps in realistic chart understanding in multimodal llms. *Advances in Neural Information Processing Systems*, 37:113569–113697, 2024.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
- Wenshan Wu, Shaoguang Mao, Yadong Zhang, Yan Xia, Li Dong, Lei Cui, and Furu Wei. Mind’s eye of llms: Visualization-of-thought elicits spatial reasoning in large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Wang Xiyao, Yang Zhengyuan, Li Linjie, Lu Hongjin, Xu Yuancheng, Lin Chung-Ching Lin, Lin Kevin, Huang Furong, and Wang Lijuan. Scaling inference-time search with vision value model for improved visual comprehension. *arXiv preprint arXiv:2412.03704*, 2024.
- Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*, 2024.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Faisal Ahmed, Zicheng Liu, Yumao Lu, and Lijuan Wang. Unitab: Unifying text and box outputs for grounded vision-language modeling. In *European Conference on Computer Vision*, pages 521–539. Springer, 2022a.
- Zhengyuan Yang, Jianfeng Wang, Zhe Gan, Linjie Li, Kevin Lin, Chenfei Wu, Nan Duan, Zicheng Liu, Ce Liu, Michael Zeng, et al. Reco: Region-controlled text-to-image generation. *arXiv preprint arXiv:2211.15518*, 2022b.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of Imms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 2023.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwan He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. RLhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13807–13816, 2024a.
- Weihaoyu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities, 2023.
- Weihaoyu, Zhengyuan Yang, Lingfeng Ren, Linjie Li, Jianfeng Wang, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, and Xinchao Wang. Mm-vet v2: A challenging benchmark to evaluate large multimodal models for integrated capabilities. *arXiv preprint arXiv:2408.00765*, 2024b.
- Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal chain-of-thought reasoning in language models. *Trans. Machine Learning Research*, 2024.
- Yaowei Zheng, Junting Lu, Shenzhi Wang, Zhangchi Feng, Dongdong Kuang, and Yuwen Xiong. Easyr1: An efficient, scalable, multi-modality rl training framework. <https://github.com/hiyouga/EasyR1>, 2025.
- Yiyang Zhou, Chenhang Cui, Rafael Rafailov, Chelsea Finn, and Huaxiu Yao. Aligning modalities in vision large language models via preference fine-tuning. *arXiv preprint arXiv:2402.11411*, 2024.

A Quality of Point-CoT

To verify the quality of Point-CoT, we invited 10 volunteers and GPT-4o itself to validate 1,000 randomly selected samples, calculating the proportion of high-quality samples. We define high-quality samples as those that do not contain reasoning errors and have essentially correct points. As shown in Table 6, we found that the proportion of high-quality samples in Point-CoT approaches 95%, demonstrating the effectiveness of our Point-CoT.

Table 6: Statistics of dataset.

Evaluation	Samples	Ratio
Human	1000	94.5%
Machine	1000	96.8%

B Comparison of Models

Table 7 extends the main paper’s Table 2 by comparing two additional baselines, Molmo-7B and Qwen2.5-VL-7B evaluated without the CoT prompting. Superficially, Qwen2.5-VL-7B appears to regress when verbose CoT reasoning is enabled. This drop, however, stems primarily from the model’s in-domain short-answer post-training, not from an intrinsic weakness of long-form reasoning—a phenomenon echoed in other RFT studies Deng et al. (2025); Wang et al. (2025a); Huang et al. (2025); Meng et al. (2025); Peng et al. (2025); Chen et al. (2025); Zheng et al. (2025). Nonetheless, our Point-RFT, which couples grounded point-level reasoning with CoT, consistently outperforms both versions of Qwen2.5-VL-7B as well as Molmo-7B. These results confirm that integrating point-based reasoning substantially boosts multimodal reasoning performance, validating the design of both the Point-CoT dataset and the Point-RFT model.

Table 7: Overall results among different datasets.

Method	Setting CoT	In-Domain ChartQA	Out-of-Domain					
			CharXiv	PlotQA	IconQA	TabMWP	Counting	Avg.
Qwen2.5-VL-7B		87.30	34.40	22.30	59.10	56.40	15.00	39.06
Qwen2.5-VL-7B	✓	70.88	26.50	17.80	53.40	61.00	21.00	35.94
Molmo-7B	✓	20.48	7.60	3.95	16.50	20.70	0.00	9.75
Point-RFT	✓	90.04	36.20	20.40	59.80	70.90	78.50	53.16

C Training Steps

As shown in Table 8, SFT requires careful balancing of training steps, reaching peak accuracy at 500 steps (56.72%). Training beyond this point leads to overfitting and performance degradation. For RL, performance monotonically improves with increasing steps, peaking at 100 steps (81.04%). This indicates that extended reward shaping continually refines the reasoning patterns grounded in visual pointing.

D Training Dataset

Table 9 reveals that mixing SFT datasets harms performance (81.04% drops to 75.20%). This suggests that targeted format learning on high-quality grounded CoT data is more effective than training on broader but noisier datasets. Notably, Point-RFT even outperforms models trained on pure ChartQA data for both SFT and RL, demonstrating its ability to generalize.

E Broader Impact

We present Point-RFT, a multimodal reasoning framework that bridges textual and visual reasoning through visually grounded Chain-of-Thought (CoT). By explicitly anchoring rationales to visual

elements, our approach mitigates hallucinations and enhances multimodal integration, advancing the reliability of AI systems in real-world document analysis. This innovation holds significant potential for human-AI collaboration. Point-RFT paves the way for trustworthy AI assistants capable of seamless multimodal reasoning across diverse domains.

Table 8: Ablation studies of training steps.

Method	Steps		ChartQA		
	SFT	RL	Overall	Inner	Format
Base	-	-	79.28	82.17	96.48
Point-SFT	300	-	50.00	71.88	69.56
	500	-	56.72	76.20	74.44
	600	-	45.56	75.48	60.36
	700	-	37.00	74.71	49.52
	1000	-	20.44	75.45	25.68
Point-RFT	500	5	80.92	82.14	98.52
	500	10	83.36	83.66	99.64
	500	20	84.72	84.96	99.72
	500	50	85.48	85.58	99.88
	500	100	86.24	86.52	99.68

F Limitation

Despite the promising results, our approach has several limitations. First, the two-stage training pipeline (format finetuning and reinforcement finetuning) requires substantial computational resources, particularly when scaling to larger datasets. Moreover, like most autoregressive models, the sequential reasoning process in Point-RFT introduces significant inference latency and resource consumption, especially for complex visual documents. These challenges highlight the need for more efficient training paradigms and inference acceleration techniques. In future work, we will explore lightweight training strategies and parallelizable decoding mechanisms to address these limitations while maintaining reasoning accuracy.

Table 9: Ablation studies of training dataset.

Method	Dataset		ChartQA		
	SFT	RL	Overall	Inner	Format
Base	-	-	79.28	82.17	96.48
Base-SFT	ChartQA	-	3.88	74.62	5.20
Base-RFT	ChartQA	ChartQA	80.92	81.51	99.28
Point-RFT	Point-CoT (w/o ChartQA)	ChartQA	81.04	82.42	98.32
Point-RFT	Point-CoT	ChartQA	86.24	86.52	99.68