

Why Vision Language Models Struggle with Visual Arithmetic? Towards Enhanced Chart and Geometry Understanding

Kung-Hsiang Huang¹ Can Qin¹ Haoyi Qiu²
Philippe Laban¹ Shafiq Joty¹ Caiming Xiong¹ Chien-Sheng Wu¹

¹Salesforce AI Research ²UCLA

¹{kh.huang, cqin, sjoty, cxiong, wu.jason}@salesforce.com

Abstract

Vision Language Models (VLMs) have achieved remarkable progress in multimodal tasks, yet they often struggle with visual arithmetic, seemingly simple capabilities like object counting or length comparison, which are essential for relevant complex tasks like chart understanding and geometric reasoning. In this work, we first investigate the root causes of this deficiency through a suite of probing tasks focusing on basic visual arithmetic. Our analysis reveals that while pre-trained vision encoders typically capture sufficient information, the text decoder often fails to decode it correctly for arithmetic reasoning. To address this, we propose COGALIGN, a novel post-training strategy inspired by Piaget’s theory of cognitive development. COGALIGN trains VLMs to recognize invariant properties under visual transformations. We demonstrate that this approach significantly improves the performance of three diverse VLMs on our proposed probing tasks. Furthermore, COGALIGN enhances performance by an average of 4.6% on CHOCOLATE and 2.9% on MATH-VISION, outperforming or matching supervised fine-tuning methods while requiring only 60% less training data. These results highlight the effectiveness and generalizability of COGALIGN in improving fundamental visual arithmetic capabilities and their transfer to downstream tasks. ¹

1 Introduction

In recent years, vision language models (VLMs) have rapidly advanced, demonstrating remarkable capabilities in integrating and processing multimodal information (Liu et al., 2023; Dai et al., 2023; Chen et al., 2024; Xue et al., 2024). These models have found extensive applications across various domains, ranging from visual common-sense reasoning to sophisticated tasks like web

agents (Xu et al., 2024; Zhang et al., 2024a; Xie et al., 2024; Lin et al., 2024). By leveraging both visual and textual data, VLMs promise a nuanced understanding that surpasses what can be achieved by analyzing them individually.

Despite these advancements, current VLMs exhibit noticeable deficiencies in performing fundamental *visual arithmetic*: these models struggle with seemingly simple tasks like accurately counting objects, comparing lengths, assessing angles, and evaluating relative sizes or areas (Rahmanzadehgervi et al., 2024; Wang et al., 2024c; Huang et al., 2024a; Ullman, 2024; Wei et al., 2024; Kamoj et al., 2024). These shortcomings are particularly evident in complex tasks such as chart understanding (Huang et al., 2024c) and geometric problem-solving (Gao et al., 2025).

In this study, we first delve into the root causes of VLMs’ difficulties with visual arithmetic, exploring several hypotheses to elucidate why VLMs often fail when faced with such challenges (§2). We propose a suite of probing tasks, focusing on basic visual arithmetic such as length comparison, to answer this question. Our analysis reveals that pre-trained vision encoders coupled with a simple linear classifier perform poorly on these probing tasks, indicating that a single linear layer is insufficient to decode the complex visual representations for arithmetic reasoning. However, when we fine-tune the text decoder of a VLM on these tasks, performance significantly improves. This suggests **the bottleneck lies in the decoder’s ability to effectively process and utilize the visual information, rather than in the visual representation itself.**

To tackle these challenges, we propose a novel post-training strategy, COGALIGN, designed to improve the performance of VLMs in visual arithmetic tasks (§3). Drawing inspiration from Piaget’s theory of cognitive development (Piaget, 1952), our method focuses on enhancing VLMs’ understanding of *conservation* (recognizing that certain prop-

¹COGALIGN data has been released at: <https://github.com/SalesforceAIResearch/CogAlign>.

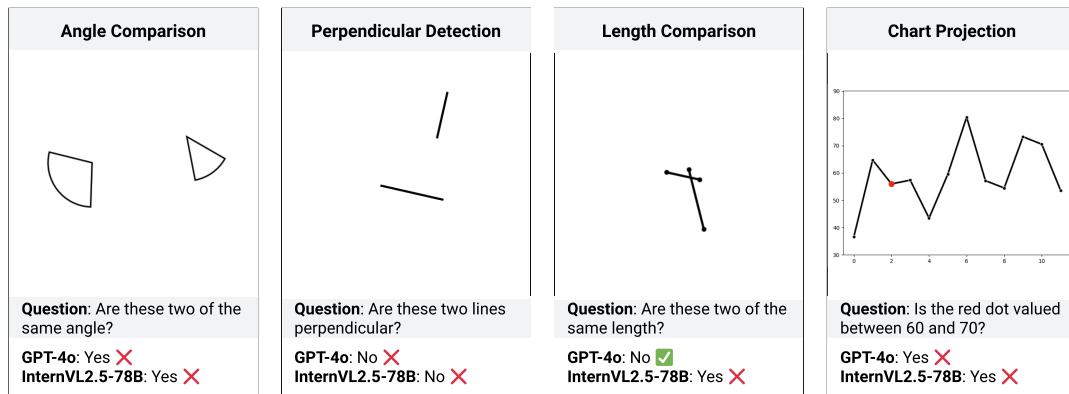


Figure 1: Examples of probing tasks designed to assess visual arithmetic abilities. Each task presents a visual input and a question requiring comparison or evaluation of geometric properties. At the bottom of each task, we see that even top-performing VLMs like GPT-4o and InternVL2.5-78B struggle with these seemingly simple tasks.

erties remain unchanged despite transformations) and *decentration* (considering multiple aspects simultaneously). We train VLMs using synthetically generated image pairs that demonstrate transformations, enabling them to compare and evaluate based on specific properties like length, angle, and quantity. By employing Direct Preference Optimization (DPO) (Rafailov et al., 2023), the model learns from both positive and negative examples, offering a richer learning signal than traditional Supervised Fine-Tuning (SFT). Our experiments show that COGALIGN significantly enhances performance across three VLMs of different scales and architectures on the proposed probing tasks.

Furthermore, we evaluate COGALIGN on two downstream benchmarks: CHOCOLATE (Huang et al., 2024c) for chart understanding, and MATH-VISION (Wang et al., 2024a) for geometric problem-solving (§4). Our results demonstrate the effectiveness of COGALIGN in enhancing performance on these complex tasks. On average, COGALIGN boosts performance by 4.6% and 2.9% on CHOCOLATE and MATH-VISION respectively, demonstrating that improving fundamental visual arithmetic capabilities translates to improved performance on downstream tasks. Notably, COGALIGN outperforms or achieves comparable performance to SFT methods while requiring 60% less training data, even though COGALIGN does not involve direct optimization for specific tasks. This showcases its strong generalizability and highlights its potential of focusing on foundational skills to unlock broader capabilities in VLMs.

Our main contributions are as follows:

- We conduct an in-depth analysis to uncover the root causes of VLMs’ underperformance in tasks that involve visual arithmetic.

- We develop COGALIGN, a post-training strategy designed to enhance VLMs’ abilities in understanding performing visual arithmetic.
- Extensive experiments on three VLMs show that COGALIGN significantly improves performance in chart comprehension and geometric problem-solving, highlighting its generalizability.

2 Why Vision Language Models Struggle with Visual Arithmetic?

As suggested in previous studies, VLMs struggle with visual arithmetic (Rahmanzadehgervi et al., 2024; Wang et al., 2024c), leading to poor performance in tasks involving such capabilities such as chart understanding (Huang et al., 2024c) and geometric problem-solving (Gao et al., 2025). In this section, we aim to understand the root causes behind such phenomenon. We first propose a suite of probing tasks we design to facilitate our analysis (§2.1) and then illustrate the various analyses we conduct to validate our hypotheses (§2.2).

2.1 Probing Tasks

We propose four probing tasks for assessing visual arithmetic capabilities, motivated by the fundamental operations needed to interpret visual data quantitatively. For a VLM to successfully understand a chart, for example, it must be able to compare lengths of bars or lines, discern relationships indicated by line slopes, and projecting points onto axes. An overview of the probing tasks are shown in Figure 1. All four tasks are discriminative and can be considered binary classification tasks. Below, we illustrate these tasks in details.

Angle Comparison asks models to determine whether the angle of two wedges are the same.

Vision Encoder	Angle Comparison	Perpendicular Detection	Length Comparison	Chart Projection
LLaVA-v1.5-proj	89.7	87.3	74.4	61.3
CLIP ViT-L/14	95.8	92.2	81.0	60.0
SigLIP-SO400M/14	98.5	92.1	89.9	74.5
InternViT-300M-V2.5	88.2	85.5	70.8	59.0
DINOv2-Large	96.2	94.7	81.8	57.3
Random guessing	50.0	50.0	50.0	50.0

Table 1: Accuracy (%) of different vision encoders with a linear classifier on the test set of each probing task. We conduct feature probing experiments by freezing the vision encoder and only fine-tuning the linear layer for binary classification. LLaVA-v1.5-proj refers to the representations obtained from the projection layer of LLaVA-v1.5.

This requires the model to differentiate and measure angular magnitude, a seemingly more complex operation that tests the model’s grasp of spatial relationships and angular geometry. This task assesses the model’s capacity to interpret rotational dimensions and engage in deeper analytical processing to distinguish subtle differences in angle, thereby evaluating the core geometric understanding of the model in angular perception.

Perpendicular Detection challenges models to determine if two given lines are perpendicular to each other. Building upon the concept of angles, this task requires a deeper understanding of specific angular relationships, where perpendicularity implies a 90° angle. While Angle Comparison focuses on general angle differentiation, Perpendicular Detection assesses a model’s ability to recognize this specific geometric configuration.

Length Comparison asks models whether two lines with arbitrary slopes are of the same lengths. In addition to basic spatial reasoning, this task requires models to consider trigonometric relationships between the lines, demanding higher-level understanding of equivalence regardless of orientation. The variability in slopes necessitates an advanced ability to rotate or translate lines, challenging the model’s proficiency in geometric reasoning beyond simple horizontal and vertical comparisons.

Chart Projection challenges the model to determine if the value of a red dot on a black line chart lies between 60 and 70. As the most complex task, this task integrates key aspects of the preceding tasks. It requires spatial reasoning to project the dot’s position onto the y-axis, similar to Angle and Perpendicular Detection. It then involves comparing the projected value’s magnitude against the specified range, akin to Length Comparison.

2.2 Probing Analysis

The research question we aim to answer is: *Do visual representations from pre-trained vision encoders contain enough information to perform vi-*

sual arithmetic tasks? To answer this question, we conduct experiments by feeding the outputs from various encoders into a linear classifier to perform binary classification on the probing tasks. For each task, we randomly generate 12,000 images programmatically with a train:development:test split of 10:1:1. Each split has a balanced portion of positive and negative labels. We test a wide range of vision encoder, including CLIP ViT-L/14 (Radford et al., 2021), SigLIP-SO400M/14 (Zhai et al., 2023), InternViT-300M-V2.5 (Chen et al., 2024), and DINOv2-Large (Oquab et al., 2024). We also evaluate the features produced by the projection layer of LLaVA-v1.5 (Liu et al., 2023). Each model (i.e., the single classifier) was trained for 200 epochs and the checkpoint that achieves the highest performance on the development set is selected.

The results are presented in Table 1. Overall, we observe that fixed visual representations, when paired with a single linear layer, yield reasonable performance on simpler tasks such as Simple Length Comparison and Angle Comparison. However, they struggle significantly with more complex tasks like Length Comparison and Chart Projection. Therefore, we conclude that **pre-trained vision encoders do not convey sufficient information through their fixed visual representations for a linear classifier to succeed at visual arithmetic**. This may be attributed to two potential reasons: (1) the visual representations genuinely lack the information necessary for visual arithmetic tasks, or (2) a linear layer lacks the capacity to effectively leverage the visual features provided.

To further investigate the underlying cause of this limitation, we perform additional experiments by fine-tuning the LLM-based text decoder component of LLaVA-v1.5, while keeping its vision encoder frozen. In VLMs, visual representations are concatenated with text representations in the decoder. Unlike a linear layer, the text decoder can process textual queries as inputs, offering an opportunity to understand the effect of textual clues

VLM	Query Type	Fine-tuned?	Length Comparison	Chart Projection
	-	✗	50.0	51.3
LLaVA-v1.5-7B	ORIGINAL	✓	95.4	98.9
	EMPTY	✓	95.2	98.1
	IRRELEVANT	✓	95.8	97.8

Table 2: Accuracy (%) of fine-tuned LLaVA-v1.5-7B with different queries on the test set of each probing task. We conduct experiments by freezing the vision encoder and only fine-tune the LLM decoder for binary classification. The top row displays the zero-shot performance.

provided by input queries. We test with three different queries: an **ORIGINAL** query reflecting the task as shown in Figure 1, an **EMPTY** query which is a blank string, and an **IRRELEVANT** query such as “*My name is John?*”. Additionally, we evaluate LLaVA-v1.5 in a zero-shot setting for comparisons. Given our previous observations, these experiments focus exclusively on the Length Comparison and Chart Projection tasks.

The fine-tuned LLaVA results are shown in Table 2. We have the following observations. First, **existing VLMs do struggle with challenging visual arithmetic when used in zero-shot manners, achieving less than 90% and 75% on the two more challenging probing tasks, even with extensive in-domain training.** The finding is consistent with prior studies (Rahmanzadehgervi et al., 2024; Wang et al., 2024c) and highlights the validity and complexity of our proposed probing tasks. Second, **VLMs fine-tuned on in-domain data perform reasonably well in visual arithmetic.** LLaVA-v1.5-7B is able to achieve an accuracy of above 95% on both Length Comparison and Chart Projection tasks. Third, **the high performance of fine-tuned VLMs on in-domain data is due to the larger capacity of an LLM.** Comparing the three different queries, we see the performance on Length Comparison and Chart Projection does not vary too much. This means that a fine-tuned LLaVA-v1.5-7B performs well even when the query provides no clue or irrelevant information about the given tasks. Combining this observation with our findings in Table 1, we learn that fine-tuned VLMs perform well because the LLM-based text decoder have larger capacity than a linear layer rather than leveraging the semantics of the input query.

2.3 Visual Representation Inspection

While the analysis conducted in §2.2 reveals important findings, one may argue that the fine-tuning success might be caused by the fine-tuned model exploiting patterns in binary tasks. Motivated by

the visual representation analysis of Rudman et al. (2025), we conduct additional experiments focusing on visual representations’ ability to differentiate isolated geometric properties. We generated datasets of images, each containing a single angle (sampled in 10-degree increments from 10° to 100°) or a single line (sampled in 0.1 length increments from 0.1 to 1.0). We created 100 images per class, resulting in 1,000 images each for the angle and line datasets. We then extracted visual embeddings for these images using the pre-trained CLIP and SigLIP vision encoders. To assess whether embeddings for different angles or lengths occupy distinct regions in the feature space (a prerequisite for successful differentiation or classification), we performed t-SNE visualization (van der Maaten and Hinton, 2008) to understand cluster separation. The results are shown in Figure 2.

We have the following observations. First, there are some clear clusters for both CLIP and SigLIP embeddings, particularly for smaller angles and shorter lines. However, the separation is less distinct for longer lines and larger angles, especially with CLIP embeddings. Second, the separation for angles is more obvious than that for lines, aligning with our findings in Table 3 that Length Comparison is more challenging than Angle Comparison. Third, SigLIP embeddings demonstrate superior clustering for both angles and lines. This is also consistent with our Table 3 results where SigLIP outperforms CLIP across the proposed probing tasks.

Based on these observations and our findings in §2.2, we conclude that **while visual representations from VLMs do encode information for visual arithmetic, it cannot be effectively decoded without further fine-tuning**, which may in turn affect their zero-shot performance on downstream tasks like chart understanding and that **fine-tuning success is not merely exploiting patterns in a binary task.**

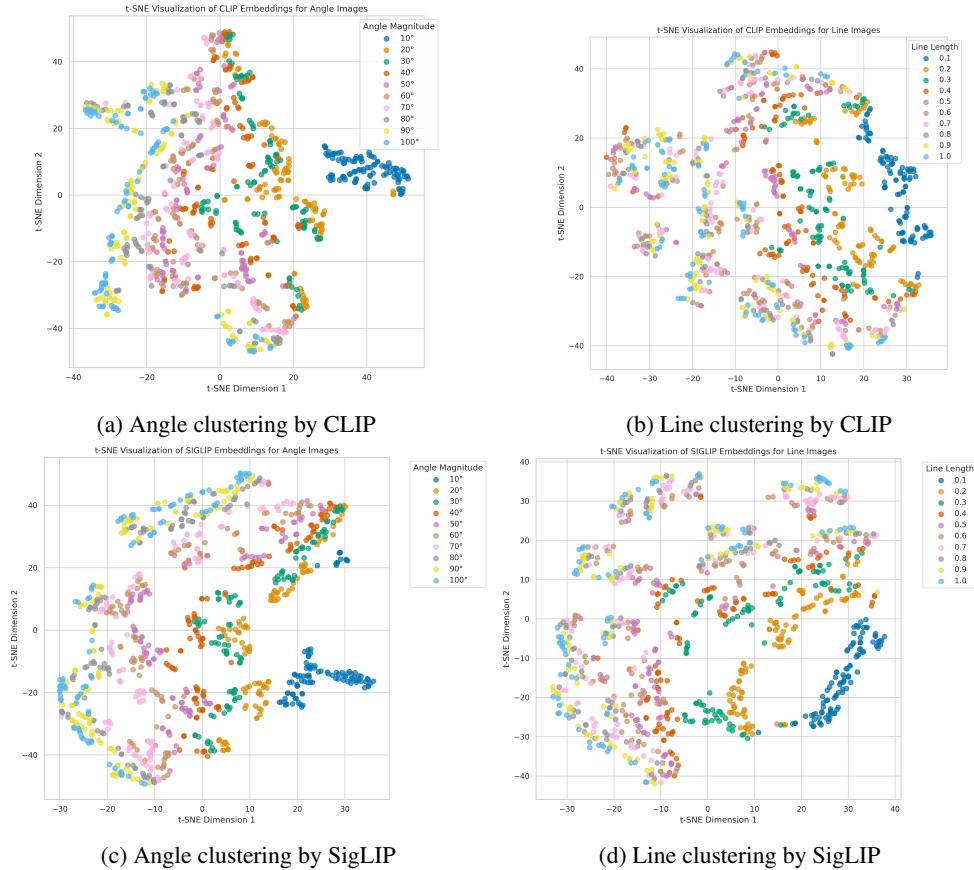


Figure 2: t-SNE visualizations of clustering by CLIP and SigLIP for angles and lines.

3 COGALIGN

To address the challenges VLMs face in performing visual arithmetic, we propose a novel post-training method inspired by Piaget’s theory of cognitive development (Piaget, 1952), which outlines four stages: Sensorimotor, Preoperational, Concrete Operational, and Formal Operational. Each stage represents a different ability to process information and solve problems, culminating in abstract reasoning. The Concrete Operational Stage is particularly relevant. At this stage, children develop (1) *conservation*, understanding that certain properties like length remains constant despite changes in appearance, and (2) *decentration*, the ability to consider multiple aspects of a situation at once. These skills are essential for VLMs to perform visual arithmetic accurately, recognizing invariant properties such as length or angle across transformations. Current VLM training paradigms often neglect these cognitive processes, resulting in models that struggle to maintain key properties during visual transformations and to integrate multiple visual features effectively. While pre-trained visual encoders use losses that encourage some invariance to transformations, the integration of vision and language representa-

tions in decoders often lacks explicit enforcement of conservation and decentration principles, leading to models that capture visual features but fail to reason about them effectively.

To address these issues, we present a post-training method, Cognitive Alignment (COGALIGN), aimed at enhancing VLMs’ understanding of *conservation* and *decentration*. Our approach explicitly trains VLMs to recognize invariant properties like length, angle, and count across different visual transformations. We achieve this by presenting the model with pairs of figures and associated queries designed to highlight these properties. The queries prompt the model to compare and contrast the figures, focusing on whether a specific property is different or same despite variations in appearance. This approach encourages the model to develop a stronger understanding of geometric concepts and move beyond superficial visual comparisons. Furthermore, we leverage DPO (Rafailov et al., 2023) for training, rather than SFT. DPO allows the model to learn from both positive and negative examples within the preference framework, providing a richer learning signal compared to SFT. By strengthening these cognitive capacities within VLMs, our goal

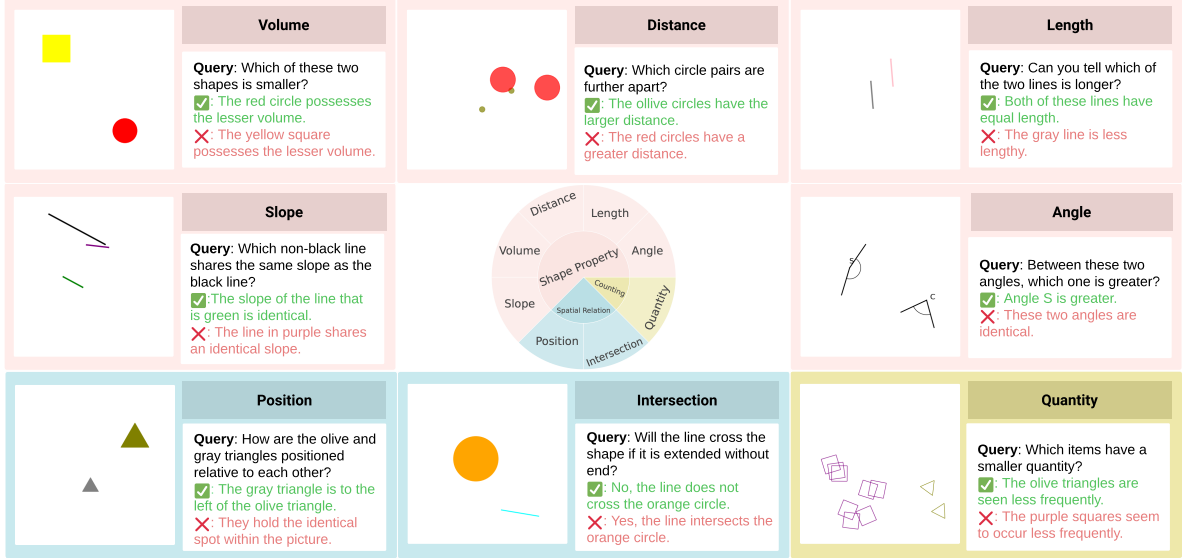


Figure 3: Example training data for COGALIGN. Each example consists of a visual input, a query prompting comparison of a specific property (i.e. angle, length, distance, and etc), a positive response consistent with the visual input, and a negative response that contradicts it.

is to improve their performance on tasks involving visual arithmetic. The subsequent subsections detail the specific training procedure employed (§3.1) and the automated construction of our training data (§3.2).

3.1 DPO Training Objective

Our goal is to train a model with parameters θ that learns *conservation* and *decentration* from contrasting responses by maximizing the conditional probability of positive responses over their negative counterparts. Concretely, the DPO training data consists of preference pairs, each containing a user query Q , an input image I , a positive response R_p and a negative response R_n . The entire set of DPO training data can be represented as $\mathcal{D} = \left\{ (Q, I, R_p, R_n)^{(i)} \right\}_{i=1}^{|\mathcal{D}|}$. The objective function \mathcal{L}_{DPO} that DPO minimizes is:

$$\mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(Q, I, R_p, R_n) \sim \mathcal{D}} \left[\log \sigma(r_\Delta) \right],$$

$$r_\Delta = \beta \log \frac{\pi_\theta(R_p | Q, I)}{\pi_{\text{ref}}(R_p | Q, I)} - \beta \log \frac{\pi_\theta(R_n | Q, I)}{\pi_{\text{ref}}(R_n | Q, I)},$$

where σ is the sigmoid function, π_θ is the parameterized policy under training, π_{ref} is the initial frozen policy, and β is a hyper-parameter that controls the deviation from π_{ref} .

3.2 Training Data Synthesis

To effectively train VLMs on the principles of *conservation* and *decentration*, we require a training

dataset designed to highlight these concepts. This section details our automated process for synthesizing training data, encompassing visual generation and tailored query-response construction. We draw inspiration from our probing tasks but adapt the format to better suit the DPO training procedure. By plotting two shapes within a single image, we allow the model to directly compare invariant properties like length and angle across various transformations. We devise eight fundamental tasks, each of which aim to enhance VLMs’ different abilities to reason about visual arithmetic operations: understanding *angle*, *length*, *distance*, *quantity*, *volume*, *position*, *slope*, and *intersection*. An overview of these tasks is shown in Figure 3

To automate data synthesis, we present a data generation pipeline. First, we programmatically generate images using Python, allowing precise control over each figure’s properties, such as lengths and positions. Next, we create query-response pairs for these images based on predefined templates (see Table 7). These queries are designed to prompt VLMs to make comparisons, identify similarities/dissimilarities, and reason about geometric properties. Given the known ground truth, positive responses are generated by accurately populating placeholders in the templates, while negative responses are created with incorrect values. For instance, a positive query for the first sub-figure in Figure 3 might be, “The angle S is larger.”, and a negative query might be, “The angle C is larger.”

Model	Angle Comparison	Perpendicular Detection	Length Comparison	Chart Projection
LLaVA-OV-0.5B	51.8	50.5	52.5	50.7
+COGALIGN (Ours)	79.8 (+54.0%)	51.5 (+2.0%)	53.4 (+1.7%)	53.2 (+4.9%)
InternVL-2.5-MPO-1B	51.8	49.6	52.3	60.0
+COGALIGN (Ours)	52.1 (+0.6%)	50.7 (+2.0%)	52.6 (+0.6%)	66.9 (+11.5%)
InternVL-2.5-MPO-4B	60.6	54.9	56.3	84.0
+COGALIGN (Ours)	72.3 (+19.3%)	56.4 (+2.7%)	60.0 (+6.6%)	86.3 (+2.7%)

Table 3: Accuracy (%) of different VLMs on our proposed probing tasks. All models produce output in a zero-shot fashion without fine-tuning on the tasks.

To ensure diversity, we use an LLM² to create multiple variations of each query and response, following the approach of Huang et al. (2024b). We synthesize a total of 64,000 training instances for DPO, with a balanced splits of each task.

3.3 Effectiveness on the Probing Tasks

To assess the effectiveness of COGALIGN on our proposed probing tasks, we trained three VLMs with varying scales and architectures: LLaVA-OV-0.5B (Li et al., 2024), InternVL-2.5-MPO-1B (Wang et al., 2024b), and InternVL-2.5-MPO-4B using COGALIGN, as described in §3.1 and §3.2, for one epoch. The results are presented in Table 3.

We observe that COGALIGN demonstrably improves performance across all three models across all probing tasks. More significant gains are observed on simpler tasks of Angle Comparison, likely due to their similarity with the DPO training instances. For instance, LLaVA-OV-0.5B sees a substantial 54.0% improvement on Angle Comparison after training with COGALIGN. This highlights the effectiveness of our approach in enhancing the core visual arithmetic capabilities that are crucial for these tasks. Interestingly, while the gain in angle-related tasks like Angle Comparison was substantial, the performance increases for Perpendicularity Detection were more modest (e.g., a 2.0% improvement for LLaVA-OV-0.5B). This suggests that certain geometric properties, such as perpendicularity, may pose greater challenges.

Overall, these findings collectively demonstrate the effectiveness of COGALIGN in enhancing visual arithmetic capabilities across various VLMs.

4 Generalizability of COGALIGN

Now that we have demonstrated the advantage of COGALIGN on our probing tasks, we ask: *does the improvement on simple visual arithmetic tasks transfer to more complex tasks?* To answer this question, we explore whether COGALIGN en-

hances model performance in chart understanding and geometric problem-solving. In the following subsections, we detail the experimental setup (§4.1) and present our findings (§4.2).

4.1 Experimental Setups

Benchmarks We evaluate the effectiveness of our method on two tasks relevant to visual arithmetic: chart understanding and geometry problem-solving. For chart understanding, we utilize the CHOCOLATE dataset (Huang et al., 2024c), which tests a model’s capability to determine whether a given caption is factually consistent with its corresponding chart.³ CHOCOLATE comprises three splits: LVLM, LLM, and FT, each generated by models of varying architectures and scales. Each CHOCOLATE instance is annotated with a binary label $\mathcal{L} \in \{\text{consistent}, \text{inconsistent}\}$. The dataset includes a total of 1,187 chart-caption pairs. For geometry problem-solving, we assess performance using the test set of the MATH-VISION dataset (Wang et al., 2024a), which comprises 3,040 questions spanning 16 mathematical disciplines. We concentrate on the eight disciplines related to geometry: analytic geometry (ANAG), combinatorial geometry (COMBG), descriptive geometry (DESCG), solid geometry (SOLG), transformation geometry (TRANSG), and three metric geometry branches - angle, area, and length. For evaluations, we employ AUC score for CHOCOLATE and accuracy for MATH-VISION, in alignment with Huang et al. (2024c) and Wang et al. (2024a). Detailed dataset statistics for these benchmarks are provided in Appendix A.

Models and Baselines To assess the efficacy of COGALIGN compared to methods that directly optimize model capabilities towards specific tasks, we consider a chart supervised fine-tuning dataset: CHARTGEMMA160K (Masry et al., 2024), as well as one geometric problem-solving dataset:

³We decided against using other common datasets like ChartQA (Masry et al., 2022) due to their training data already being included in some VLMs, such as LLaVA-OneVision.

²gpt-4o is used for paraphrasing.

Model	CHOCOLATE				MATH-VISION								
	LVLN	LLM	Ft	AVG	ANAG	COMBG	DESCG	ANGLE	AREA	LEN	SOLG	TRANSG	AVG
GPT-4o	74.7	69.6	68.5	70.9	35.7	31.2	28.8	38.2	34.8	36.7	31.1	21.4	32.2
LLaVA-Next-Vicuna-7B	50.0	50.3	50.1	50.1	9.5	14.9	18.3	12.7	15.6	13.4	7.0	17.3	15.6
Qwen2-VL-7B-Instruct	57.3	64.0	71.1	65.4	17.9	17.9	24.0	19.7	22.4	21.4	16.0	25.0	19.5
ChartGemma-3B	51.8	54.2	53.7	53.2	11.9	13.6	14.4	9.8	11.2	10.0	10.6	13.7	11.9
G-LLaVA-13B	50.0	50.0	50.0	50.0	14.3	15.9	22.1	19.1	20.0	21.2	15.6	16.1	18.0
LLaVA-OV-0.5B	56.6	50.4	57.8	54.9	16.7	17.2	22.1	17.3	13.6	18.9	11.1	16.7	16.7
+CHARTGEMMA160K	50.2	50.0	50.3	50.2	11.9	15.6	16.3	17.9	12.4	14.6	10.2	19.0	14.7
+GEO170K	50.8	48.7	50.5	50.0	16.7	16.9	13.5	17.3	14.6	16.9	10.7	17.3	15.5
+COGALIGN (Ours)	56.7	64.7	57.7	59.7	15.5	17.5	19.2	17.3	13.8	17.8	11.5	16.1	16.1
InternVL-2.5-MPO-1B	53.2	60.2	65.0	59.5	16.7	18.2	22.1	26.0	15.8	18.0	10.7	17.3	18.1
+CHARTGEMMA160K	54.5	61.8	60.9	59.1	20.2	16.6	24.0	26.6	19.2	16.3	12.3	21.4	19.6
+GEO170K	54.6	62.1	60.9	59.2	19.0	16.6	32.7	24.9	20.0	15.4	13.1	18.5	20.0
+COGALIGN (Ours)	59.7	60.1	64.6	61.5	16.7	16.2	31.7	25.4	17.0	17.4	13.1	20.8	19.7
InternVL-2.5-MPO-4B	60.3	67.2	75.9	67.8	28.6	23.1	22.1	32.4	22.6	24.9	19.7	17.9	23.8
+CHARTGEMMA160K	62.1	66.0	76.2	68.1	23.8	18.5	18.3	16.8	24.4	23.4	10.7	17.9	22.3
+GEO170K	60.0	65.5	64.3	59.9	32.1	18.8	23.7	31.2	23.6	23.8	15.6	21.4	23.8
+COGALIGN (Ours)	61.2	68.6	76.8	68.9	27.4	19.8	24.0	32.9	22.6	26.9	17.2	25.6	24.6

Table 4: Performance (%) on the CHOCOLATE and MATH-VISION datasets.

GEO170K (Gao et al., 2025). We use the above methods to train three open-source VLMs for one epoch: InternVL2.5-1B-MPO (Wang et al., 2024b), InternVL2.5-4B-MPO (Wang et al., 2024b), and LLaVA-OV-0.5B (Li et al., 2024). We also compare performance of two VLMs instruction-tuned specifically for chart understanding and geometric problem-solving: ChartGemma-3B (Masry et al., 2024) and G-LLaVA-13B (Gao et al., 2025). Experimental details can be found in Appendix B.

4.2 Results

The results for experiments on CHOCOLATE and MATH-VISION are shown in Table 4. We find that **COGALIGN is effective in enhancing chart understanding and geometric problem-solving capabilities of VLMs even though COGALIGN was not specifically optimized for these two tasks.** On average, COGALIGN boosts the performance by 4.6% and 2.9% on the CHOCOLATE and MATH-VISION datasets, respectively. This shows that patching fundamental capabilities such as visual arithmetic of VLMs can enhance their capabilities in tasks involving such abilities.

More importantly, we find that **COGALIGN demonstrates better generalizability compared to supervised fine-tuning VLMs using task-specific data.** For instance, when comparing the InternVL-2.5-MPO-1B variants, COGALIGN achieves an average score of 61.5% on CHOCOLATE, outperforming both the CHARTGEMMA160K (59.1%) and GEO170K (59.2%) variants. Similarly, on the MATH-VISION dataset, while the GEO170K variant shows competitive performance, COGALIGN achieves a comparable average performance across all geometry subtasks, indicating a broader improvement. Notably, COGALIGN requires only 60% less training data com-

pared to these two baseline methods.

The results suggest that COGALIGN offers a valuable approach to enhancing VLMs by improving their fundamental visual arithmetic capabilities. It exhibits strong generalizability across different tasks and base models, often outperforming or achieving comparable performance to task-specific fine-tuning methods without being explicitly trained on the target datasets. This highlights the potential of focusing on foundational skills to unlock broader capabilities in VLMs.

4.3 Discussions

Impact of learning from contrasting examples

We investigate the impact of learning from contrasting examples versus solely positive examples by comparing DPO (the default COGALIGN setting) and SFT training method (using only the positive response). Figure 4 presents the results. We observe that the SFT approaches can lead to much worse performance (e.g. LLaVA-OV-0.5B), while the DPO approach improves performance over the original models more consistently. This suggests that learning from contrasting examples provides a richer learning signal compared to traditional supervised learning, leading to better performance.

Impact on general VLM benchmarks

To assess the impact of COGALIGN on general VLM capabilities, we compare the performance on two additional benchmarks: MME (Fu et al., 2023) and MMMU (Yue et al., 2024). The results are presented in Figure 5. Overall, COGALIGN consistently improves performance across most settings (five out of six), indicating that its benefits extend beyond the specific probing tasks and generalize to other multimodal reasoning challenges. This suggests that **COGALIGN enhances visual**

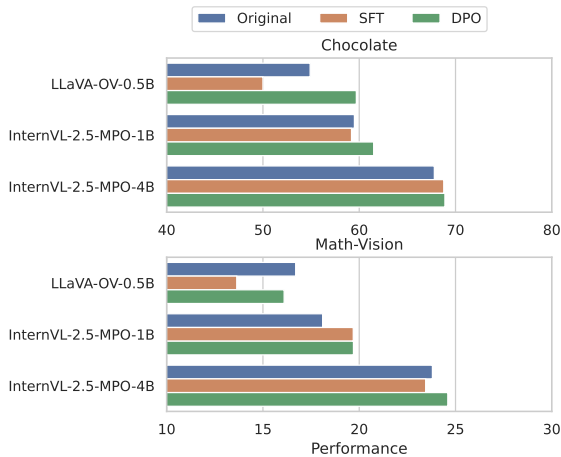


Figure 4: Performance comparison when training different models with SFT and DPO.

arithmetic capabilities without compromising performance on general tasks.

5 Related Works

5.1 Vision Language Models

Vision language models (VLMs) are multimodal models that learn to generate text outputs based on both visual and textual inputs. The development of large-scale VLMs has demonstrated impressive zero-shot capabilities, enabling them to perform well with a variety of image types, such as documents and web pages (Liu et al., 2023; Dai et al., 2023; OpenAI, 2023; Google, 2023; Anthropic, 2023). These VLMs generally consist of three major components: a vision encoder, such as CLIP (Radford et al., 2021) or SigLIP (Zhai et al., 2023), which processes visual inputs; a language model that handles textual inputs and generates text tokens; and a projector layer that connects the image and text modalities. Typically, VLMs are trained using image captioning data and instruction-tuning datasets. Recently, several post-training strategies have been suggested to enhance VLM capabilities in areas like conversational interaction (Xiong et al., 2024) and reasoning (Wang et al., 2024b). In this work, we propose a new post-training strategy, COGALIGN, for improving VLMs’ proficiency in understanding visual arithmetic operations.

5.2 Shortcomings of Vision Language Models

While Vision-Language Models (VLMs) demonstrate impressive performance across a range of tasks, several studies have highlighted their limitations by examining various aspects such as architectures (McKinzie et al., 2024; Karamcheti et al., 2024; Tong et al., 2024; Shi et al., 2025),




Figure 5: Performance on two general VLM benchmarks: MME and MMMU with or without COGALIGN.

training methods (Laurençon et al., 2024), and data considerations (Udandarao et al., 2024; Gadre et al., 2024; Zhang et al., 2024b; Wei et al., 2024). Some research indicates that VLMs struggle with specific tasks, including basic geometric understanding (Gao et al., 2025; Ullman, 2024) and chart comprehension (Huang et al., 2024c), and are prone to hallucinations (Qiu et al., 2024). Our study seeks to uncover the root causes behind these challenges, especially those that involve visual arithmetic operations, and proposes solutions to address these shortcomings.

6 Conclusion

This study investigates the challenges faced by VLMs in performing visual arithmetic, revealing that while visual encoders often capture necessary information, text decoders struggle to effectively utilize it. We introduce COGALIGN, a novel post-training strategy inspired by Piaget’s theory of cognitive development, focusing on enhancing VLMs’ understanding of conservation and decentration through DPO training. Our evaluations show that COGALIGN not only enhances VLMs’ understanding of visual arithmetic, but also improves their performance in chart understanding and geometric problem-solving through experiments on the CHOCOLATE and MATH-VISION datasets, showcasing its effectiveness and generalizability across various models and tasks. Notably, COGALIGN often outperforms or achieves comparable results to task-specific supervised fine-tuning methods without direct training on the target domain, highlighting the potential of bolstering foundational cognitive skills for broader VLM capabilities. Future work could explore how COGALIGN impacts other multimodal tasks beyond charts and geometry, potentially leading to a more unified approach in VLM training where generalizability is prioritized.

7 Limitations

Probing Tasks While the probing tasks we have proposed provide valuable insights into the visual arithmetic capabilities of VLMs, it is important to acknowledge that they may not encompass all possible dimensions of visual reasoning. Our choice to limit the scope of these tasks was intentional, as they serve as initial, simple tests to determine whether VLMs exhibit failure in fundamental aspects of visual arithmetic. These tasks allow us to iterate different experiments in a controlled and efficient manner, providing clear, actionable insights without the complexity that more comprehensive tasks might introduce. However, there is potential to explore additional tasks that involve more complex interactions of basic geometric properties. For instance, tasks requiring the model to simultaneously assess both length and angle, or combinations of length and area, could be valuable for understanding the compositionality of these atomic tasks.

Training Data Synthesis The training data synthesis method of COGALIGN is not only scalable but also effectively enhances the visual arithmetic capabilities of VLMs. Our approach serves as a proof-of-concept, demonstrating the potential of automated data generation for improving models' understanding of basic geometric properties. To further enrich the training data, we could consider utilizing additional configurations for each task. For instance, in generating positive and negative responses, we could leverage LLMs to produce rationales based on the specific configuration of each figure. By including explanations or justifications for why a particular geometric property holds or does not hold, we could foster deeper understanding within the VLMs.

References

- Anthropic. 2023. [Claude: A new ai assistant](#). *Anthropic*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024. 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.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [InstructBLIP: Towards general-purpose vision-language models with instruction tuning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. 2023. [MME: A comprehensive evaluation benchmark for multimodal large language models](#). *CoRR*, abs/2306.13394.
- Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, et al. 2024. Datacomp: In search of the next generation of multimodal datasets. *Advances in Neural Information Processing Systems*, 36.
- Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wan-jun Zhong, Yufei Wang, Lanqing HONG, Jianhua Han, Hang Xu, Zhenguo Li, and Lingpeng Kong. 2025. [G-LLaVA: Solving geometric problem with multi-modal large language model](#). In *The Thirteenth International Conference on Learning Representations*.
- Google. 2023. [Gemini: Advancements in ai models](#).
- Kung-Hsiang Huang, Hou Pong Chan, Yi R. Fung, Haoyi Qiu, Mingyang Zhou, Shafiq Joty, Shih-Fu Chang, and Heng Ji. 2024a. From pixels to insights: A survey on automatic chart understanding in the era of large foundation models. *IEEE Transactions on Knowledge and Data Engineering*.
- Kung-Hsiang Huang, Akshara Prabhakar, Sidharth Dhawan, Yixin Mao, Huan Wang, Silvio Savarese, Caiming Xiong, Philippe Laban, and Chien-Sheng Wu. 2024b. [Crmarena: Understanding the capacity of llm agents to perform professional crm tasks in realistic environments](#). *Preprint*, arXiv:2411.02305.
- Kung-Hsiang Huang, Mingyang Zhou, Hou Pong Chan, Yi Fung, Zhenhailong Wang, Lingyu Zhang, Shih-Fu Chang, and Heng Ji. 2024c. [Do LVLMs understand charts? analyzing and correcting factual errors in chart captioning](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 730–749, Bangkok, Thailand. Association for Computational Linguistics.
- Ryo Kamoi, Yusen Zhang, Sarkar Snigdha Sarathi Das, Ranran Haoran Zhang, and Rui Zhang. 2024. [Visionyqa: Large vision language models still struggle with visual perception of geometric information](#).
- Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa Sadigh. 2024. [Prismatic vlms: investigating the design space of visually-conditioned language models](#). In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Hugo Laurençon, Leo Tronchon, Matthieu Cord, and Victor Sanh. 2024. [What matters when building](#)

- vision-language models? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Weixian Lei, Lijuan Wang, and Mike Zheng Shou. 2024. Showui: One vision-language-action model for gui visual agent. *arXiv preprint arXiv:2411.17465*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. **Visual instruction tuning**. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. **ChartQA: A benchmark for question answering about charts with visual and logical reasoning**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279, Dublin, Ireland. Association for Computational Linguistics.
- Ahmed Masry, Megh Thakkar, Aayush Bajaj, Aaryaman Kartha, Enamul Hoque, and Shafiq Joty. 2024. **Chart-gemma: Visual instruction-tuning for chart reasoning in the wild**. *Preprint*, arXiv:2407.04172.
- Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, Anton Belyi, Haotian Zhang, Karanjeet Singh, Doug Kang, Ankur Jain, Hongyu Hè, Max Schwarzer, Tom Gunter, Xiang Kong, Aonan Zhang, Jianyu Wang, Chong Wang, Nan Du, Tao Lei, Sam Wiseman, Guoli Yin, Mark Lee, Zirui Wang, Ruoming Pang, Peter Grasch, Alexander Toshev, and Yinfei Yang. 2024. Mm1: Methods, analysis & insights from multimodal llm pre-training. *ArXiv*, abs/2403.09611.
- OpenAI. 2023. **Gpt-4 technical report**.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. **DINOv2: Learning robust visual features without supervision**. *Transactions on Machine Learning Research*. Featured Certification.
- John Piaget. 1952. The origins of intelligence in children. *International University*.
- Haoyi Qiu, Wenbo Hu, Zi-Yi Dou, and Nanyun Peng. 2024. Valor-eval: Holistic coverage and faithfulness evaluation of large vision-language models. *arXiv preprint arXiv:2404.13874*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. **Learning transferable visual models from natural language supervision**. In *International Conference on Machine Learning*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. **Direct preference optimization: Your language model is secretly a reward model**. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Pooyan Rahmazadehgervi, Logan Bolton, Mohammad Reza Taesiri, and Anh Totti Nguyen. 2024. Vision language models are blind. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 18–34.
- William Rudman, Michal Golovanevsky, Amir Bar, Vedant Palit, Yann LeCun, Carsten Eickhoff, and Ritambhara Singh. 2025. Forgotten polygons: Multimodal large language models are shape-blind. *arXiv preprint arXiv:2502.15969*.
- Baifeng Shi, Ziyang Wu, Maolin Mao, Xin Wang, and Trevor Darrell. 2025. When do we not need larger vision models? In *European Conference on Computer Vision*, pages 444–462. Springer.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9568–9578.
- Vishaal Udandarao, Ameya Prabhu, Adhiraj Ghosh, Yash Sharma, Philip Torr, Adel Bibi, Samuel Albanie, and Matthias Bethge. 2024. No "zero-shot" without exponential data: Pretraining concept frequency determines multimodal model performance. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Tomer Ullman. 2024. The illusion-illusion: Vision language models see illusions where there are none. *arXiv preprint arXiv:2412.18613*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. **Visualizing data using t-sne**. *Journal of Machine Learning Research*, 9(Nov):2579–2605.
- Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. 2024a. **Measuring multimodal mathematical reasoning with MATH-vision dataset**. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, and Jifeng Dai. 2024b. Enhancing the reasoning ability of multimodal large

- language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*.
- Zhenhailong Wang, Joy Hsu, Xingyao Wang, Kuan-Hao Huang, Manling Li, Jiajun Wu, and Heng Ji. 2024c. Visually descriptive language model for vector graphics reasoning. *arXiv preprint arXiv:2404.06479*.
- Haoran Wei, Youyang Yin, Yumeng Li, Jia Wang, Liang Zhao, Jianjian Sun, Zheng Ge, and Xiangyu Zhang. 2024. Slow perception: Let's perceive geometric figures step-by-step. *arXiv preprint arXiv:2412.20631*.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. 2024. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*.
- Tianyi Xiong, Bo Li, Dong Guo, Huizhuo Yuan, Quanquan Gu, and Chunyuan Li. 2024. [Llava-onevision-chat: Improving chat with preference learning](#).
- Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. 2024. [Llava-cot: Let vision language models reason step-by-step](#). *Preprint*, arXiv:2411.10440.
- Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. 2024. xgen-mm (blip-3): A family of open large multimodal models. *arXiv preprint arXiv:2408.08872*.
- Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. 2024. [MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 9556–9567. IEEE.
- 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 (ICCV)*, pages 11975–11986.
- Ruohong Zhang, Bowen Zhang, Yanghao Li, Haotian Zhang, Zhiqing Sun, Zhe Gan, Yinfei Yang, Ruoming Pang, and Yiming Yang. 2024a. Improve vision language model chain-of-thought reasoning. *arXiv preprint arXiv:2410.16198*.
- Yuhui Zhang, Alyssa Unell, Xiaohan Wang, Dhruva Ghosh, Yuchang Su, Ludwig Schmidt, and Serena Yeung-Levy. 2024b. [Why are visually-grounded language models bad at image classification?](#) In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

A Dataset statistics

Table 5 and Table 6 show the dataset statistics for the CHOCOLATE and MATH-VISION datasets, respectively.

B Training Settings

For all models and all training approaches, we set all other hyper-parameters according to the guidelines described in their corresponding GitHub repository. These settings are described in Table 8.

Statistic	Number
Total questions	3,040
- multiple-choice questions	1,532 (50.4%)
- Free-form questions	1,508 (49.6%)
- Questions in the testmini set	304 (10.0%)

Table 6: Statistics of the MATH-VISION dataset.

	# Factual	# Non-factual	# Total
Sentence	2,561	2,762	5,323
Caption	213	974	1,187

Table 5: Statistics of the CHOCOLATE dataset. A sentence is considered factual if and only if it does not contain any factual error. A caption is considered factual if all its sentences are factual.

Task	Query Templates
Angle	<ol style="list-style-type: none"> 1. The angle with the [COLOR] color is larger. 2. The angle X is larger. 3. The angle with the [COLOR] color is smaller. 4. The angle X is smaller. 5. These two angles are the same.
Length	<ol style="list-style-type: none"> 1. The line with the [COLOR] color is longer. 2. The line X is longer. 3. The line with the [COLOR] color is shorter. 4. The line X is shorter. 5. These two lines are the same length.
Distance	<ol style="list-style-type: none"> 1. The pair of circles with the [COLOR] color has the longer distance. 2. The pair of circles with the [COLOR] color has the smaller distance. 3. These two pair of circles have the same distance.
Quantity	<ol style="list-style-type: none"> 1. The [COLOR] [SHAPE] appears more times. 2. The [COLOR] [SHAPE] appears less times. 3. The [COLOR-A] [SHAPE-A] and [COLOR-B] [SHAPE-B] appear the same number of times.
Volume	<ol style="list-style-type: none"> 1. The [COLOR] [SHAPE] has the larger volume. 2. The [COLOR] [SHAPE] has the smaller volume. 3. These two shapes have the same volume.
Slope	<ol style="list-style-type: none"> 1. The line with the [COLOR] has the same slope. 2. Both lines have the same slope as the black line. 3. Neither line has the same slope as the black line.
Position	<ol style="list-style-type: none"> 1. The [COLOR-A] [SHAPE-A] is [POSITION] of [COLOR-B] [SHAPE-B]. 2. They occupy the exact same position in the image. 3. The [COLOR-A] [SHAPE-A] is [WRONG-POSITION] of [COLOR-B] [SHAPE-B].
Intersection	<ol style="list-style-type: none"> 1. Yes, the line does intersect the [COLOR] [SHAPE]. 2. No, the line does not intersect the [COLOR] [SHAPE].

Table 7: The full set of query templates used for query generation.

Model	Training Method	Batch Size	Learning Rate
LLaVA-OV-0.5B	DPO	128	5e-7
	SFT	16	2e-6
InternVL-2.5-MPO-1B	DPO	256	1e-6
	SFT	16	4e-5
InternVL-2.5-MPO-4B	DPO	256	1e-6
	SFT	16	4e-5

Table 8: Experimental details for different training approaches. All models are trained for one epoch for fair comparisons.