

# S2H-DPO : Hardness-Aware Preference Optimization for Vision–Language Models

Anonymous ACL submission

## Abstract

Vision-Language Models (VLMs) have demonstrated remarkable progress in single-image understanding, yet effective reasoning across multiple images remains challenging. We identify a critical capability gap in existing multi-image alignment approaches: current methods focus primarily on localized reasoning with pre-specified image indices (“Look at Image 3 and...”), bypassing the essential skills of global visual search and autonomous cross-image comparison. To address this limitation, we introduce a Simple-to-Hard (S2H) framework that systematically constructs multi-image preference data across three hierarchical reasoning levels requiring an increasing level of capabilities: (1) single-image localized reasoning, (2) multi-image localized comparison, and (3) global visual search. Unlike prior work that relies on model-specific attributes, such as hallucinations or attention heuristics, to generate preference pairs, our approach leverages prompt-driven complexity to create chosen/rejected pairs that are applicable across different models. Through extensive evaluations on LLaVA-v1.5 and Qwen2.5-VL models, we show that our diverse multi-image reasoning data significantly enhances multi-image reasoning performance, yielding average improvements of 5.27% and 2.49% over baseline methods across benchmarks. Importantly, our approach maintains strong single-image reasoning performance while simultaneously strengthening multi-image understanding capabilities, thus advancing the state of the art for holistic visual preference alignment.

## 1 Introduction

Vision–Language Models (VLMs) (Li et al., 2025; Bai et al., 2023; Lin et al., 2024) have recently made rapid progress in understanding visual inputs and solving reasoning-intensive tasks. Yet reasoning across multiple images remains a challenging frontier: the model must not only interpret each

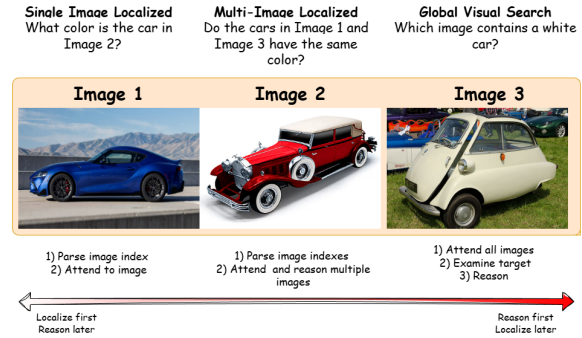


Figure 1: Multi-image reasoning skill hierarchy: From localized reasoning ( $L1$ : "What is the color of the car in Image 2?") to global visual search ( $L3$ : "Which image contains a white car?"). Each level requires strictly more capabilities than the previous.

image in isolation, but also align, compare, and integrate evidence across images. While proprietary systems such as GPT-4o (OpenAI, 2024) demonstrate strong multi-image capabilities, open-source VLMs (Li et al., 2025; Zhang et al., 2024) still struggle to reliably aggregate information when visual context is distributed across several images. The difficulty compounds as the number of images grows, increasing both the search space and the need for consistent cross-image correspondence. Effective multi-image reasoning therefore hinges on two core abilities: (1) localizing where to look across multiple images, and (2) composing information from multiple regions into a coherent conclusion. Existing work primarily addresses this via multi-image pre-training (Awadalla et al., 2023) and supervised fine-tuning (SFT) (Jiang et al., 2024; Chen et al., 2024b; Liu et al., 2024c), with only a few studies (Liu et al., 2025) approaching the problem from the perspective of visual preference alignment. In particular, (Liu et al., 2025) observes that simply incorporating multi-image data during SFT is not sufficient to improve multi-image reasoning—and can, in some cases, even degrade performance on single-image tasks—thereby moti-

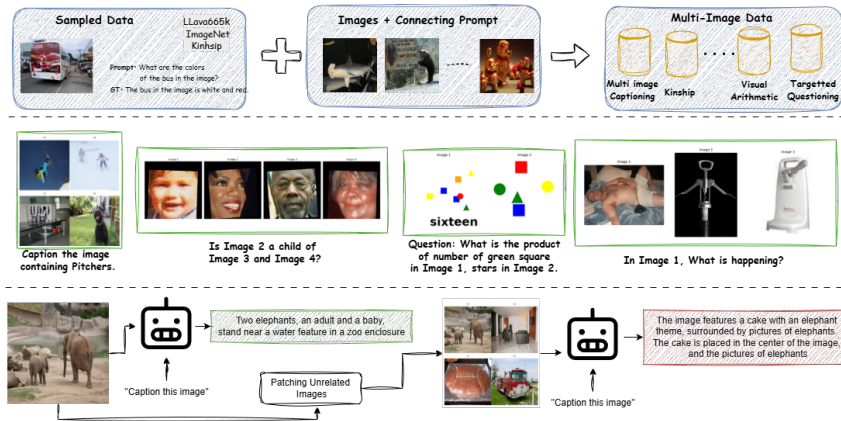


Figure 2: **Overview of our method.** We transform single-image data sources (e.g., LLaVA 665k, ImageNet) into multi-image datasets across visual modalities with progressively increasing cognitive load. Our synthetic transformation pipeline creates task hierarchies that advance from basic single-image understanding to complex multi-image reasoning requiring visual comparison, spatial reasoning, and cross-modal integration.

vating Direct Preference Optimization (DPO)-style preference alignment for multi-image reasoning. However, a key challenge is constructing multi-image preference pairs because (1) each example must include a carefully selected set of multiple images and corresponding questions, and (2) each example requires well-curated chosen and rejected text response. These requirements make large-scale human annotation prohibitively expensive, motivating synthetic data approaches for generating multi-image preference data.

While recent work—most notably MIA-DPO (Liu et al., 2025)—proposes promising strategies for generating synthetic preference pairs, its data curation remains fundamentally limited in capturing the skills required for robust multi-image reasoning. In particular, it relies on questions that pre-specify where the model should look (e.g., “What is the color of the car in Image 2?” given multiple images). By explicitly indicating which image(s) to examine, this formulation sidesteps key competencies in which the model must autonomously determine relevant visual evidence, compare images, and compose information across multiple views. Therefore, we argue that, beyond simply reducing the cost of large-scale data generation, it is crucial to explicitly define the capabilities needed for multi-image reasoning and to ensure that data curation targets them. In this regard, let us consider the progression of multi-image reasoning capabilities:

- **Level 1 (Single-Image Localized):** "What is the color of the car in Image 2?" — Reason about one pre-specified image.
- **Level 2 (Multi-Image Localized):** "Do the cars

in Image 1 and Image 3 have same color?" — Reason and Compare across multiple pre-specified images, unlike Level 1 which reduces to single-image reasoning.

- **Level 3 (Global Visual Search):** "Which images contain a white car?" — Reason and Search all images to locate relevant regions and carefully aggregate information, unlike Level 1 and Level 2, where regions are pre-specified.

Figure 1 illustrates these capabilities, where each level demands strictly more from the model than the previous. **Critically, MIA-DPO (Liu et al., 2025) trains only on Level 1.** We argue this is not merely a difference in question format, but an important gap: different formulations induce qualitatively different reasoning patterns with increasing difficulty. In particular, at Level 1 the model reasons about a single, pre-specified image (“look at Image 2”); at Level 2 it must compare multiple pre-specified images (“look at Image 1 and Image 3”); and at Level 3 it must autonomously determine where to look and then reason over the relevant regions (e.g., “look for the image with a white car”). Thus, we ask the question: **Can preference tuning on diverse multi-image reasoning tasks achieve broader generalization than index-based localization training alone?**

More precisely, we automatically construct a diverse synthetic training set containing preference pairs spanning Levels 1-3, allowing models to learn multi-image reasoning capabilities simultaneously. We then perform a **comprehensive multi-image DPO approach that jointly trains models across all reasoning levels.** Through systematic abla-

tions, we demonstrate that: (1) joint training on mixed-difficulty data enables robust multi-image reasoning, (2) models trained on diverse question types, including global search, generalize better than those trained only on localized questions.

Experimental results highlight a significant capability gap between our approach and the current state-of-the-art. While MIA-DPO achieves 44.2% accuracy on Level 1–2 benchmarks, its performance declines to 42.9% on Level 3 tasks (Table 1). Conversely, our method—which leverages training samples across all reasoning levels—improves these figures to 47.9% and 45.04%, respectively. A critical limitation of MIA-DPO is its reliance on Level 1 examples, which hinge on model-specific intrinsic hallucinations to generate chosen-rejected pairs. This necessitates the generation of a new dataset for every new model architecture. By introducing higher hardness levels through prompt-driven complexity, our method alleviates this dependency. Furthermore, our approach preserves single-image reasoning capabilities, matching the performance of the pre-finetuned baseline. These findings demonstrate that our method substantially enhances multi-image reasoning across all semantic levels without compromising foundational single-image performance.

In summary, our main contributions are:

- We identify a “capability gap” in existing multi-image alignment: current methods train only on localized reasoning with pre-specified images (Level 1), missing the global and partial visual search capabilities (Levels 2-3) required for real-world multi-image understanding.
- We develop a systematic technique for generating high-quality rejected answers from correct-answer-only datasets, enabling comprehensive multi-image preference optimization across all reasoning levels without manual annotation.
- Through extensive experiments, we demonstrate that our comprehensive multi-image DPO approach substantially outperforms existing methods while preserving single-image reasoning capabilities.

## 2 Proposed Method

### 2.1 Preliminaries

In this section, we present the concept of visual preference alignment and use the Direct Preference Optimization (DPO) method as a representative example.

**Visual Preference Alignment:** Preference alignment focuses on aligning a model’s outputs with a set of preferences. Common methods include Reinforcement Learning from Human Feedback (RLHF) (Yu et al., 2024a) and Reinforcement Learning from AI Feedback (RLAIF) (Yu et al., 2024b). Consider a dataset  $D$  where each sample consists of an input prompt  $x$ , a preferred response  $y_w$ , and a rejected response  $y_l$ . Formally, we represent the dataset as  $D = \{x, y_w, y_l\}$ . The input  $x$  can be an interleaved sequence of images  $v$  and text  $t$ .

When a VLM processes  $x$  to generate an output  $y$ , a reward  $r(x, y)$  is assigned by a reward model  $r$ , which assigns higher scores to preferred outputs and lower scores to rejected ones. Visual preference alignment aims to maximize this reward:

$\max_{\theta} \mathbb{E}_{x \sim D, y \sim \pi_{\theta}(y|x)} [r(x, y)]$ , where  $\theta$ ,  $\pi_{\theta}$  and  $\pi_{\theta}(y|x)$  denote the parameters, policy, and output distribution of VLM, respectively.

To avoid overfitting on the dataset  $D$ , a KL-divergence loss  $D_{\text{KL}}$  is incorporated by preference alignment approaches to regularize the difference between the model’s policy  $\pi_{\theta}(y|x)$  and a reference model’s policy  $\pi_{\text{ref}}(y|x)$ :

$$\max_{\theta} [\mathbb{E}_{x \sim D, y \sim \pi_{\theta}(y|x)} [r(x, y)] - \beta \cdot D_{\text{KL}}(\pi_{\theta}(y|x) \parallel \pi_{\text{ref}}(y|x))], \quad (1)$$

where, the hyper-parameter  $\beta$  controls the influence of KL-divergence on the optimization objective. Note that the reference model is the model’s state prior to preference alignment.

**Direct Preference Optimization (DPO):** To optimize the preference alignment objective in Eq. (1), we can use either an online reward model (e.g., PPO (Schulman et al., 2017)) or pre-computed off-line chosen/rejected pairs (e.g., DPO (Rafailov et al., 2024)). Given its simplicity, DPO has been widely adopted in previous visual alignment works (Zhao et al., 2023; Zhou et al., 2024). Eq. (1) can be reformulated as the loss function of DPO:  $L_{\text{DPO}}(\pi_{\theta}; \pi_{\text{ref}}) =$

$$-\mathbb{E}_{(x, y_w, y_l) \sim D} \left[ \log \sigma \left( \beta \log \frac{\pi_{\theta}(y_w|x)}{\pi_{\text{ref}}(y_w|x)} - \beta \log \frac{\pi_{\theta}(y_l|x)}{\pi_{\text{ref}}(y_l|x)} \right) \right], \quad (2)$$

where,  $\sigma(\cdot)$  denotes the sigmoid function. As shown in Eq. (2), DPO-based alignment methods concentrate on designing input prompts  $x$  and pairing each with a preferred output  $y_w$  and a rejected output  $y_l$ .



Figure 3: Simple-to-Hard (S2H) DPO Data Format. **Level 1:** The query directly points to a specific image (e.g., “Caption the first image”), and rejected pairs provide responses completely unrelated to the target image. **Level 2:** The query references multiple images to enforce multi-image reasoning (e.g., “Compare the first and third images”). **Level 3:** The query is open-ended and requires the model to examine all images before identifying which one satisfies the semantic constraint (e.g., “Caption the image containing a peacock”).

## 2.2 Simple-to-hard (S2H) Tasks

We propose three Simple-to-Hard (S2H) probe tasks with increasing levels of complexity, defined by the cognitive load required for localization, cross-image grounding, and reasoning. Our tasks therefore escalate from basic single-image query to complex multi-step reasoning that demands joint understanding of all images. We illustrate the overall process in Figure 2. Below, we explain these tasks in detail.

### 2.2.1 Single Image Localized (Level 1)

In the simplest case, following prior work (Liu et al., 2025), we convert single-image datasets to multi-image format by appending unrelated distractor images to existing VQA samples. The chosen response corresponds to the gold answer for the target image, while the rejected response is an incorrect answer generated by the model. Critically, this approach does not establish explicit relationships between images—the task simply requires the model to ignore irrelevant visual context. The rejected responses primarily arise from model hallucinations when processing noisy multi-image inputs, rather than from meaningful cross-image reasoning errors. While this provides a baseline for multi-image handling, it does not test genuine multi-image reasoning capabilities.

### 2.2.2 Multi-Image Localized (Level 2)

In this category, we extend Level 1 tasks to a multi-image setting. Rather than analyzing a single image

in isolation, we expand the context by requiring the model to localize features across multiple images and respond to prompts that necessitate explicit cross-image connections. We introduce two distinct tasks in this category: (i) Kinship Recognition, where the objective is to determine familial or social relationships within a set of facial images; and (ii) Visual Arithmetic, which requires the model to perform mathematical reasoning—such as counting or logical operations—on shapes, numbers, and objects distributed across several images. We select Kinship Recognition and Visual Arithmetic as Level-2 tasks because they capture complementary aspects of multi-image localized reasoning. Kinship Recognition requires asymmetric relational inference: although multiple images are presented, the answer pertains primarily on a single target image in relation to others, demanding cross-image relational grounding. Visual Arithmetic, in contrast, requires symmetric compositional aggregation, where each image contributes equally and the model must combine information across all images to compute the answer. Together, these sub-skills capture the core challenges of L2 multi-image reasoning.

*Chosen/Rejected pair generation.* For Kinship Recognition, we leverage preexisting datasets (Wang et al., 2017; Robinson et al., 2016) to generate chosen and rejected pairs. The process is straightforward: we first read the labels and relationships of the individuals. If the relationship is correct, we generate a chosen

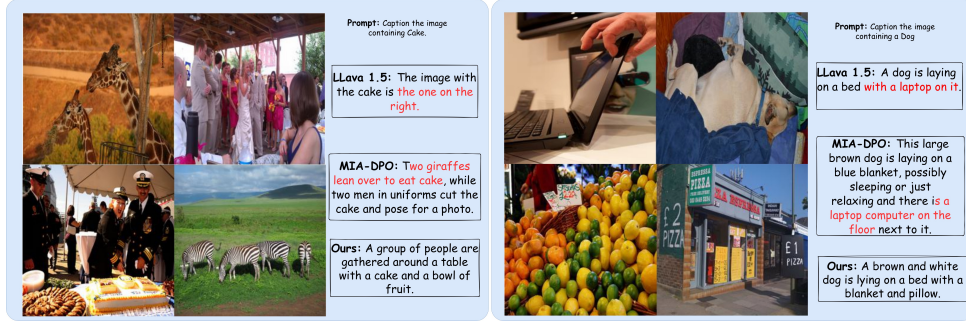


Figure 4: Outputs of our method on Global Visual Search. Our approach effectively isolates the queried concept and produces accurate responses to the prompt.

deterministic caption (e.g., ‘Yes, Image 3 is...’) and a corresponding rejected caption (e.g., ‘No, Image 3 is not...’). Conversely, if the relationship is incorrect, we generate a chosen–rejected pair by flipping the responses accordingly. For Visual Arithmetic, images are generated synthetically using deterministic procedures, ensuring that each image contributes clearly to the final answer while maintaining controlled, reproducible inputs for multi-image reasoning. For Visual Arithmetic, images are generated synthetically by placing random shapes and objects in each image while keeping track of their counts. To create a question, we randomly sample a subset of images, select some objects, and choose a random mathematical operator (e.g., addition, subtraction) to form a problem whose answer depends on aggregating information across the selected images. This deterministic procedure guarantees correct chosen/reject pair generation.

### 2.2.3 Global Visual Search (Level 3)

This category inverts the computational ordering: the model must first perform global reasoning across the entire image set before localizing the relevant image. Consider the task of captioning an image that satisfies a semantic criterion, e.g., “Caption the image containing a peacock”. Unlike previous levels, localization cannot precede reasoning. Instead, the model must evaluate all images against the query constraint, identify the matching instance, and only then generate the caption. This formulation tests the model’s ability to integrate cross-image search with subsequent generative tasks.

#### Single-to-Multi-Image Prompt Construction.

Figure 3 (bottom) illustrates our procedure for converting single-image captions into multi-image chosen-rejected pairs. We construct training data

from existing annotated datasets through a two-stage process designed to distinguish targeted search from general summarization. *Chosen pair generation:* We first select a target concept from ImageNet (Deng et al., 2009) and sample a target image  $\mathcal{I}_{\text{target}}$  belonging to that concept. We then prompt a vision-language model to generate both (1) a detailed caption and (2) the primary object present in the image (e.g., peacock). To form a multi-image set, we pair  $\mathcal{I}_{\text{target}}$  with  $N - 1$  distractor images randomly sampled from *different* concept classes, ensuring that the same primary object does not appear in more than one image. The chosen response is then assigned as the generated caption. *Rejected pair generation:* For the same multi-image set, we prompt the captioning model *without* specifying a target object, asking it to simply “caption the images.” This yields a rejected response that typically either aggregates information across multiple images or provides a generic description, failing to isolate the target. Importantly, these rejected responses are not randomly unrelated—they represent plausible but incorrect responses, which provides a better learning signal to model compared to trivial negatives. We use Qwen2.5-VL-32B (Bai et al., 2025) as the caption generator in this work. *Quality filtering:* To ensure meaningful contrast between chosen and rejected pairs, we compute semantic similarity between the two responses using text encoders (CLIP, MPNet). Pairs with similarity exceeding threshold  $\tau$  (top quartile) are discarded, as high similarity indicates the rejected response inadvertently captured the correct behavior. This filtering ensures the contrastive pairs provide a clear learning signal for distinguishing conditional visual search from unconstrained multi-image description.

### 2.3 Implementation Details

Our S2H-DPO is model-agnostic and can be directly applied to a wide range of LLMs without requiring the generation of additional training datasets, in contrast to prior approaches. We evaluate our method on two widely used MLLMs: LLaVA-v1.5-7B (Liu et al., 2024a) and Qwen2.5-VL-7B (Bai et al., 2025). All models are trained for three epochs with a learning rate of  $5 \times 10^{-5}$ , and the temperature parameter  $\beta$  is set to 0.1. We create 20K samples for each level.

Table 1: **Main results on multi-image benchmarks.** We compare our S2H-DPO with other DPO algorithms across five multi-image benchmarks. Our method yields consistent improvements over both the classic LLaVA-v1.5 and the recent Qwen2.5-VL-7B. In contrast, single-image DPO methods perform poorly on multi-image benchmarks.

Models	Parameter	BLINK	Mantis	NLVR2	Average
GPT-4V (Achiam et al., 2023)	-	51.1	62.7	88.8	67.53
LLaVA-v1.6 (Li et al., 2025)	7B	39.6	45.6	58.9	48.03
Qwen-VL-Chat (Bai et al., 2023)	7B	31.2	39.2	58.7	43.03
VideoLLaVA (Lin et al., 2024)	7B	38.9	35.9	56.5	43.77
Fuyu (Bavishi et al., 2023)	8B	36.6	27.2	51.1	38.3
Idefics2 (Laurençon et al., 2024)	8B	45.2	48.9	86.9	60.33
InstructBLIP (Dai et al., 2023)	13B	42.2	45.6	60.3	49.37
CogVLM (Wang et al., 2024)	17B	41.5	45.2	58.6	48.43
Emu2-Chat (Sun et al., 2024)	37B	36.2	37.8	58.2	44.07
LLaVA-v1.5 (Liu et al., 2024a)	7B	37.1	41.9	52.1	43.7
+ LLaVA-RLHF (Sun et al., 2023)	7B	40.8	30.4	51.8	41.0
+ HA-DPO (Zhao et al., 2023)	7B	38.6	34.6	51.6	41.6
+ POVID (Zhou et al., 2024)	7B	19.9	37.8	21.4	26.37
+ MIA-DPO (Liu et al., 2025)	7B	42.9	44.2	54.2	47.1
+ S2H-DPO (Ours)	7B	<b>43.40</b>	<b>46.54</b>	<b>55.59</b>	<b>48.97</b>
$\Delta$	-	<b>+6.30</b>	<b>+6.03</b>	<b>+3.49</b>	<b>+5.27</b>
Qwen2.5-VL (Bai et al., 2025)	7B	54.29	68.66	74.28	65.74
+ MIA-DPO (Liu et al., 2025)	7B	41.28	59.45	74.18	58.30
+ S2H-DPO (Ours)	7B	<b>55.85</b>	<b>74.19</b>	<b>74.67</b>	<b>68.24</b>
$\Delta$	-	<b>+1.56</b>	<b>+5.53</b>	<b>+0.39</b>	<b>+2.49</b>

### 3 Related Works in Preference Alignment

Aligning Large Language Models with human preferences has become a critical requirement for their safe and reliable deployment in real-world applications, particularly under stringent ethical and safety constraints. Preference alignment methods can be broadly divided into two categories. The first involves feedback-driven alignment, which leverages either human-annotated preferences (Rafailov et al., 2024; Bai et al., 2022) or AI-generated feedback (Yu et al., 2024b). The second category focuses on prompt-based guidance, where carefully designed instructions are used to steer model behavior without modifying model parameters (Chen et al., 2023).

Recent studies have begun to investigate preference alignment in the context of vision-language large models (VLLMs). Sun et al. (Sun et al., 2023) introduced LLaVA-RLHF, which utilizes

Table 2: **Main results on single-image benchmarks.** We compare S2H-DPO with other DPO approaches across two key single-image benchmarks. S2H-DPO maintains strong proficiency in single-image tasks.

Models	Parameter	MMStar	POPE	Average
LLaVA-v1.6 (Li et al., 2025)	7B	37.6	70.3	53.95
Qwen-VL-Chat (Bai et al., 2023)	7B	34.5	74.9	54.7
Idefics2 (Laurençon et al., 2024)	8B	49.5	86.2	67.85
OpenFlamingo (Awadalla et al., 2023)	9B	36.9	52.6	44.75
InstructBLIP (Dai et al., 2023)	13B	32.7	86.1	59.4
CogVLM (Wang et al., 2024)	17B	39.9	88.0	63.95
Emu2-Chat (Sun et al., 2024)	37B	40.7	88.0	64.35
LLaVA-v1.5 (Liu et al., 2024a)	7B	32.9	85.9	59.4
+ LLaVA-RLHF (Sun et al., 2023)	7B	31.6	80.8	56.2
+ HA-DPO (Zhao et al., 2023)	7B	33.5	84.3	58.9
+ POVID (Zhou et al., 2024)	7B	36.2	86.3	<b>61.25</b>
+ MIA-DPO (Liu et al., 2025)	7B	32.9	<b>87.2</b>	60.05
+ S2H-DPO (Ours)	7B	<b>33.62</b>	<b>85.70</b>	<b>59.66</b>
Qwen2.5-VL (Bai et al., 2025)	7B	62.24	83.13	72.69
+ S2H-DPO (Ours)	7B	<b>62.47</b>	<b>83.59</b>	<b>73.03</b>

human-annotated preference data to reduce hallucinations in LLaVA. Li et al. (Li et al., 2024) proposed a preference distillation framework that transfers alignment signals into VLLMs, improving visual grounding and response relevance. Similarly, Yu et al. (Yu et al., 2024a) collected fine-grained human preferences in the form of segment-level corrections to hallucinated content and optimized model behavior using dense supervision. HA-DPO (Zhao et al., 2023) addresses alignment by leveraging GPT-4’s API to create the required DPO data, though it incurs substantial API expenses. POVID (Zhou et al., 2024) induces hallucinations using blurred images GPT-4 to generate DPO data. Recently, MIA-DPO (Liu et al., 2025) automates the generation of chosen-rejected pairs by exploiting the model’s intrinsic hallucinations, but it still necessitates creating a fresh set of DPO data for each model.

## 4 Experiments

### 4.1 Benchmarks

To assess our method, we evaluate its performance across two primary categories of benchmarks. First, we examine multi-image reasoning using three multi-image benchmarks BLINK (Fu et al., 2024), Mantis (Jiang et al., 2024) and NLVR2 (Suhr et al., 2019). Second, to ensure our approach maintains robust single-image capabilities, we test on two single-image benchmarks, including MMStar (Chen et al., 2024a) and POPE (Li et al., 2023). This extensive evaluation demonstrates our method’s versatility and confirms significant performance gains in multi-image scenarios without sacrificing single-image proficiency.

Table 3: Effect of curriculum training strategies. Comparison of simple-to-hard (S2H) curriculum training against flat training. Flat training outperforms incremental curriculum training, as the latter leads to myopic optimization that biases the model to look for localization cues in the prompt.

Curriculum	BLINK	MANTIS	NLVR2	Mean
LLaVA-v1.5	37.1	41.9	52.1	43.70
+L1 (Liu et al., 2025)	42.9	44.2	54.2	47.10
L1 → L2	40.2	39.6	54.7	44.83
L2 flat	44.1	44.7	55.6	48.13
L1 → L3	37.03	41.0	54.7	44.24
L3 flat	43.8	42.8	55.2	47.27
L1 → (L2 ∪ L3)	40.17	41.0	55.5	45.56
(L2 ∪ L3) flat	43.2	41.4	55.0	46.53
L1 → L2 → L3	39.96	41.4	55.3	45.55

## 4.2 Baseline Methods

We benchmark our method against several state-of-the-art preference optimization strategies. LLaVA-RLHF (Sun et al., 2023) serves as the standard RLHF baseline, utilizing human and GPT-generated feedback. HA-DPO (Zhao et al., 2023) focuses on error correction by leveraging GPT-4V as an expert annotator to rectify model hallucinations. In contrast, POVID (Zhou et al., 2024) adopts an adversarial approach, pairing distorted images with synthetic hallucinations to stress-test modality alignment. Finally, we compare against MIA-DPO (Liu et al., 2025), which automates the generation of preference data by exploiting the model’s internal hallucinations in multi-image contexts.

## 4.3 Results on Multi-Image Benchmarks

**Results on LLaVA-1.5:** We report the performance of our method across several multi-image benchmarks in Table 1. Our approach yields significant gains, specifically improving by 6.30%, 6.03%, and 3.49% across three key datasets. Notably, on the complex BLINK benchmark—which requires specialized domain knowledge—S2H-DPO outperforms the LLaVA-v1.5 baseline by 6.30%. Furthermore, evaluations on the Mantis and NLVR2 benchmarks show improvements of 6.03% and 3.49%, respectively. These results clearly highlight the effectiveness of our proposed method in enhancing the model’s capacity for visual synthesis and reasoning within multi-image contexts. We illustrate these results in Figure 4.

**Comparison with Preference Optimization baselines:** As shown in Table 1, S2H-DPO consistently outperforms established preference optimization methods. While standard LLaVA-RLHF and HA-DPO show marginal or even negative gains on

specific multi-image tasks, our method achieves substantial improvements across all benchmarks. Notably, compared to the closest baseline, MIA-DPO, our approach yields an additional boost of 0.50%, 2.34%, 1.39%. These results indicate that our method provides a more robust signal for multi-image alignment than existing DPO or RLHF variants.

**Results on QwenVL2.5-7B:** As shown in Table 1, our method consistently improves QwenVL2.5-7B across all three multi-image benchmarks, yielding gains of 1.56%, 5.53%, and 0.39%, with an average improvement of 2.49%. In contrast to preference-optimization baselines such as MIA-DPO, which degrades performance on the first two benchmarks, S2H-DPO achieves stable and substantial improvements across all settings, demonstrating its effectiveness for multi-image alignment and reasoning.

## 4.4 Results on Single-Image Benchmarks

While S2H-DPO is primarily designed to enhance performance in multi-image reasoning tasks, it is crucial that the optimized model maintains effectiveness on its core capability: single-image reasoning. We evaluate S2H-DPO on established single-image benchmarks to verify this property. As shown in Table 2, S2H-DPO consistently outperforms the LLaVA-v1.5 baseline and existing preference-alignment methods, including LLaVA-RLHF and HA-DPO, across all evaluated single-image benchmarks. These results demonstrate that S2H-DPO not only substantially improves multi-image reasoning performance but also preserves strong single-image capabilities. This dual proficiency makes S2H-DPO well-suited for training robust vision-language models capable of handling both single- and multi-image scenarios in real-world deployments.

## 5 Ablation Studies

**Ablation on curriculum training and data mixture:** We study the effect of curriculum learning strategies and data mixture on multimodal reasoning performance. Specifically, we compare incremental curriculum-based training—where models are exposed to tasks in increasing order of complexity—against flat training baselines that directly optimize on the target task distribution using different data mixtures. As summarized in Table 3, we denote incremental setups as  $A \rightarrow B$ , indicating training on task set  $B$  after pretraining on




Figure 5: Effect of the number of distractors (2–5) on L3 task performance across BLINK, MANTIS, and NLVR2, showing peak accuracy at 3 distractors.

A. Across all mixtures, flat training consistently outperforms curriculum-based approaches. For example, L2 flat training attains a mean accuracy of 48.13%, 3.3% improvement over the L1→L2 curriculum (44.83%). This improvement corresponds to gains of 3.9% and 5.1% on BLINK and MANTIS, respectively. This trend generalizes across task combinations: L3 flat training improves upon L1→L3 by 3.0% (47.27% vs. 44.24%), while (L2∪L3) flat training yields a 1.0% accuracy increase over L1→(L2∪L3) (46.53% vs. 45.56%). Even the most gradual curriculum, L1→L2→L3, achieves only 45.55% mean accuracy, underperforming all flat training variants. **These results suggest that, for multimodal reasoning tasks, direct exposure to the target task distribution is more effective than gradual adaptation through curriculum hierarchies.** This is because curricula based training on simpler tasks induce myopic reasoning behavior, causing models to over-rely on localized or explicitly referenced visual cues present in the prompt, rather than developing a holistic, global understanding of the visual scene.

**Ablation on SFT vs DPO:** We compare the effectiveness of supervised fine-tuning (SFT) and direct preference optimization (DPO) for enhancing multimodal reasoning capabilities. We simply convert the DPO data to SFT data by discarding the rejected samples. As shown in Table 4, both approaches substantially improve over the LLaVA-v1.5 baseline (43.70% mean accuracy), with SFT achieving 46.92% and DPO reaching 48.97% mean accuracy. DPO demonstrates consistent advantages across all benchmarks, outperforming SFT by 2.88% on BLINK (43.40% vs. 40.52%), 2.78% on MANTIS (47.93% vs. 45.15%), and 0.41% on NLVR2 (55.50% vs. 55.09%). These results indicate that preference-based optimization provides stronger learning signals for multimodal reasoning tasks

Table 4: DPO vs. SFT on LLaVA-v1.5. Preference-based optimization (DPO) provides stronger guidance than standard supervised fine-tuning (SFT), leading to improved performance across benchmarks.

Training	BLINK	MANTIS	NLVR2	Mean
LLaVA (Liu et al., 2024b)	37.1	41.9	52.1	43.70
+SFT	40.52	45.15	55.09	46.92
+DPO	43.40	47.93	55.50	48.97

compared to standard supervised learning.

**Ablation on number of distractors.** We study the effect of distractor count in the L3 task by training separate models on 20K samples with 2–5 distractors. Figure 5 shows the resulting accuracy trends. Performance improves as distractors increase, peaks at 3, and declines beyond that. This suggests that while more distractors enrich supervision and encourage robust reasoning, too many introduce noise. Overall, 3 distractors provide the optimal balance between contextual diversity and reasoning difficulty across benchmarks.

**Does S2H training enhance L3 ability?** To evaluate this, we construct L3-style questions using the ImageNet validation set. For each question, we randomly sample a target image belonging to an object  $c$ , and select  $N \in \{2, 3, 4\}$  distractor images from different classes. The model is then asked: “Which of the following images contains object  $c$ ?” On 500 such questions, our fine-tuned model achieves an accuracy of **31.05%**, outperforming baseline LLaVA-v1.5 by **6.94%**. This improvement indicates that the proposed S2H training tasks enhance higher-order visual reasoning and leads to stronger L3 capability.

## 6 Conclusion

We identify and address a critical capability gap in multi-image reasoning within Vision-Language Models. While prior methods focused heavily on localized indexing, our Simple-to-Hard (S2H) framework demonstrates that multi-image reasoning requires a hierarchical approach—advancing from basic isolation to complex global visual search and cross-image composition. By systematically generating synthetic preference pairs across these cognitive levels, we move beyond model-specific hallucinations toward a more robust, prompt-driven alignment strategy. Our results show that this comprehensive training not only significantly improves performance on complex multi-image benchmarks but also preserves the foundational single-image reasoning capabilities.

## 7 Limitations

Despite its effectiveness, our S2H-DPO framework has few limitations. First, although preference pairs are generated without human annotation, the process still relies on existing datasets and a strong captioning model, whose biases may propagate into the training data. Second, our experiments focus on multi-image settings with a limited number of images (capped at 6 images) ; scaling to much larger image sets may introduce additional challenges in attention and efficiency that we do not explore. Finally, while S2H-DPO captures key forms of multi-image reasoning, it does not cover all possible reasoning patterns, such as temporal or causal reasoning across images, which we leave for future work.

## References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, and 1 others. 2023. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-VL: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, Augustus Odena, Arushi Somani, and Sagnak Tasirlar. 2023. Introducing our multimodal models. *Adept Blog*.

Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. 2024a. Are we on the right way for evaluating large vision-language models? In *Proceedings of the 38th International*

*Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc. 654  
655  
656

Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Bin Lin, Zhenyu Tang, Li Yuan, Yu Qiao, Dahua Lin, Feng Zhao, and Jiaqi Wang. 2024b. Sharegpt4video: improving video understanding and generation with better captions. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc. 657  
658  
659  
660  
661  
662  
663  
664  
665

Zhipeng Chen, Kun Zhou, Zheng Zhang, Beichen Gong, Xin Zhao, and Ji-Rong Wen. 2023. ChatCoT: Tool-augmented chain-of-thought reasoning on chat-based large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 666  
667  
668  
669  
670

Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: towards general-purpose vision-language models with instruction tuning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. 671  
672  
673  
674  
675  
676  
677

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. 678  
679  
680  
681  
682

Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. 2024. Blink: Multimodal large language models can see but not perceive. In *18th European Conference on Computer Vision, 2024*. 683  
684  
685  
686  
687  
688

Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max W.F. Ku, Qian Liu, and Wenhui Chen. 2024. Mantis: Interleaved multi-image instruction tuning. *Transactions on Machine Learning Research*, 2024. 689  
690  
691  
692

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*. 693  
694  
695  
696  
697

Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun MA, and Chunyuan Li. 2025. LLaVA-neXT-interleave: Tackling multi-image, video, and 3d in large multimodal models. In *The Thirteenth International Conference on Learning Representations*. 698  
699  
700  
701  
702  
703

Lei Li, Zihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou Wang, Lingpeng Kong, and Qi Liu. 2024. VLFeedback: A large-scale AI feedback dataset for large vision-language models alignment. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6227–6246. 704  
705  
706  
707  
708  
709  
710

711	Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang,	Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu,	767
712	Wayne Xin Zhao, and Ji-Rong Wen. 2023. Eval-	Chunyuan Li, Yikang Shen, Chuang Gan, Liang-	768
713	uating object hallucination in large vision-language	Yan Gui, Yu-Xiong Wang, Yiming Yang, and 1	769
714	models. In <i>The 2023 Conference on Empirical Meth-</i>	others. 2023. Aligning large multimodal mod-	770
715	<i>ods in Natural Language Processing</i> .	els with factually augmented rlhf. <i>arXiv preprint</i>	771
		<i>arXiv:2309.14525</i> .	772
716	Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning,	Shuyang Wang, Joseph P Robinson, and Yun Fu. 2017.	773
717	Peng Jin, and Li Yuan. 2024. Video-LLaVA: Learn-	Kinship verification on families in the wild with	774
718	ing united visual representation by alignment before	marginalized denoising metric learning. In <i>Auto-</i>	775
719	projection. In <i>Proceedings of the 2024 Conference on</i>	<i>matic Face and Gesture Recognition (FG), 2017 12th</i>	776
720	<i>Empirical Methods in Natural Language Processing</i> .	<i>IEEE International Conference and Workshops</i> .	777
721	Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae	Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi	778
722	Lee. 2024a. Improved baselines with visual instruc-	Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang,	779
723	tion tuning. In <i>Proceedings of Computer Vision and</i>	Lei Zhao, Song XiXuan, Jiazheng Xu, Keqin Chen,	780
724	<i>Pattern Recognition (CVPR)</i> .	Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie	781
725	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae	Tang. 2024. <i>CogVLM: Visual expert for pretrained</i>	782
726	Lee. 2024b. Visual instruction tuning. <i>Advances in</i>	<i>language models</i> . In <i>The Thirty-eighth Annual Con-</i>	783
727	<i>neural information processing systems</i> , 36.	<i>ference on Neural Information Processing Systems</i> .	784
728	Ziyu Liu, Tao Chu, Yuhang Zang, Xilin Wei, Xiaoyi	Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng	785
729	Dong, Pan Zhang, Zijian Liang, Yuanjun Xiong,	Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao	786
730	Yu Qiao, Dahua Lin, and 1 others. 2024c. Mmdu: A	Zheng, Maosong Sun, and 1 others. 2024a. RIHF-V:	787
731	multi-turn multi-image dialog understanding bench-	Towards trustworthy mllms via behavior alignment	788
732	mark and instruction-tuning dataset for lvlms. <i>arXiv</i>	from fine-grained correctional human feedback. In	789
733	<i>preprint arXiv:2406.11833</i> .	<i>Proceedings of Computer Vision and Pattern Recog-</i>	790
734	Ziyu Liu, Yuhang Zang, Xiaoyi Dong, Pan Zhang,	<i>nition (CVPR)</i> .	791
735	Yuhang Cao, Haodong Duan, Conghui He, Yuanjun	Tianyu Yu, Haoye Zhang, Yuan Yao, Yunkai Dang,	792
736	Xiong, Dahua Lin, and Jiaqi Wang. 2025. <i>MIA-DPO:</i>	Da Chen, Xiaoman Lu, Ganqu Cui, Taiwen He,	793
737	<i>Multi-image augmented direct preference optimiza-</i>	Zhiyuan Liu, Tat-Seng Chua, and 1 others. 2024b.	794
738	<i>tion for large vision-language models</i> . In <i>The Thir-</i>	RLAIF-V: Aligning mllms through open-source ai	795
739	<i>teenth International Conference on Learning Repre-</i>	feedback for super gpt-4v trustworthiness. <i>arXiv</i>	796
740	<i>sentations</i> .	<i>preprint arXiv:2405.17220</i> .	797
741	OpenAI. 2024. <i>Hello gpt-4o</i> .	Pan Zhang, Xiaoyi Dong, Yuhang Zang, Yuhang Cao,	798
742	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	Rui Qian, Lin Chen, Qipeng Guo, Haodong Duan,	799
743	pher D Manning, Stefano Ermon, and Chelsea Finn.	Bin Wang, Linke Ouyang, and 1 others. 2024.	800
744	2024. Direct preference optimization: Your language	Internlm-Xcomposer-2.5: A versatile large vision	801
745	model is secretly a reward model. In <i>NeurIPS</i> .	language model supporting long-contextual input and	802
746	Joseph P. Robinson, Ming Shao, Yue Wu, and Yun Fu.	output. <i>arXiv preprint arXiv:2407.03320</i> .	803
747	2016. Families in the wild (fiw): Large-scale kinship	Zhiyuan Zhao, Bin Wang, Linke Ouyang, Xiaoyi Dong,	804
748	image database and benchmarks. In <i>Proceedings</i>	Jiaqi Wang, and Conghui He. 2023. Beyond hallu-	805
749	<i>of the 2016 ACM on Multimedia Conference</i> , pages	cinations: Enhancing lvlms through hallucination-	806
750	242–246. ACM.	aware direct preference optimization. <i>arXiv preprint</i>	807
751	John Schulman, Filip Wolski, Prafulla Dhariwal,	<i>arXiv:2311.16839</i> .	808
752	Alec Radford, and Oleg Klimov. 2017. Proxi-	Yiyang Zhou, Chenhang Cui, Rafael Rafailov, Chelsea	809
753	mal policy optimization algorithms. <i>arXiv preprint</i>	Finn, and Huaxiu Yao. 2024. Aligning modalities	810
754	<i>arXiv:1707.06347</i> .	in vision large language models via preference fine-	811
755	Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang,	tuning. <i>arXiv preprint arXiv:2402.11411</i> .	812
756	Huajun Bai, and Yoav Artzi. 2019. A corpus for		
757	reasoning about natural language grounded in pho-		
758	tographs. In <i>Proceedings of the 57th Annual Meet-</i>		
759	<i>ing of the Association for Computational Linguistics,</i>		
760	<i>2019</i> .		
761	Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang,		
762	Qiying Yu, Yueze Wang, Yongming Rao, Jingjing		
763	Liu, Tiejun Huang, and Xinlong Wang. 2024. Gener-		
764	ative multimodal models are in-context learners. In		
765	<i>Proceedings of Computer Vision and Pattern Recog-</i>		
766	<i>nition (CVPR)</i> .		