

Adversary-Aware DPO: Enhancing Safety Alignment in Vision Language Models via Adversarial Training

Anonymous ACL submission

Abstract

Safety alignment is critical in pre-training large language models (LLMs) to generate responses aligned with human values and refuse harmful queries. Unlike LLM, the current safety alignment of VLMs is often achieved with post-hoc safety fine-tuning. However, these methods is less effective to white-box attacks. To address this, we propose *Adversary-aware DPO* (*ADPO*), a novel training framework that explicitly consider adversarial. *Adversary-aware DPO* (*ADPO*) integrates adversarial training into DPO to enhance the safety alignment of VLMs under worst-case adversarial perturbations. *ADPO* introduces two key components: (1) an adversarial-trained reference model that generates human-preferred responses under worst-case perturbations, and (2) an adversarial-aware DPO loss that generates winner-loser pair accounting for adversarial distortions. By combining these innovations, *ADPO* ensures that VLMs remain robust and reliable even in the presence of sophisticated jailbreak attacks. Extensive experiments demonstrate that *ADPO* outperforms baselines in the safety alignment and general utility of VLMs.

1 Introduction

Safety alignment is essential in pre-training large language models (LLMs) (Bai et al., 2022; Ouyang et al., 2022a), guiding the models to generate responses aligned with human values and enabling them to refuse harmful queries. Such alignment is typically achieved by reinforcement learning with human feedback (RLHF) (Ouyang et al., 2022a) or Direct Preference Optimization (DPO) (Rafailov et al., 2024). However, Vision-Language Models (VLMs), which use an pre-trained LLM as the backbone along with an image encoder to adapt to down-stream tasks (Liu et al., 2024b,a; Zhu et al., 2023; Dai et al., 2023; Bai et al., 2023), often lack safety alignment as a unified model in the same

way as LLMs. As a result, even when the underlying LLM is safety-aligned, VLMs remain vulnerable to jailbreak attacks, where attackers craft sophisticated prompts to manipulate the model into generating toxic content (Qi et al., 2024; Niu et al., 2024; Gong et al., 2023; Liu et al., 2025).




Figure 1: Safe response rate under white-box and black-box attacks on LLaVA-1.5. Post-hoc safety fine-tuning (SFT and DPO) is less effective on white-box attack.

Jailbreak attacks can take two forms: generation-based black-box attacks (Gong et al., 2023; Liu et al., 2025), where malicious images are generated with typography or text-to-image models like Stable Diffusion (Rombach et al., 2022), and optimization-based white-box attacks (Qi et al., 2024; Niu et al., 2024), where harmful queries are distilled into imperceptible noise added to the original image. To address these vulnerabilities, the most prevalent approach is to construct safety-relevant datasets and perform post-hoc safety fine-tuning on the target VLMs (Zong et al., 2024; Zhang et al., 2024b). For instance, Zong et al. (2024) proposed *VLGuard* that constructs a safe instruction-following dataset and uses supervised fine-tuning to enforce safe behavior, while Zhang et al. (2024b) proposed *SPA-VL* that creates a safety preference alignment dataset and applies DPO to train the model to generate preferred responses given winner-loser pairs. However, post-hoc safety fine-tuning (SFT) is more effective on black-box attack than white-box attack, as shown in Figure 1. The safe response rate of SFT is low, and DPO performs slightly better but remains unsatisfactory.

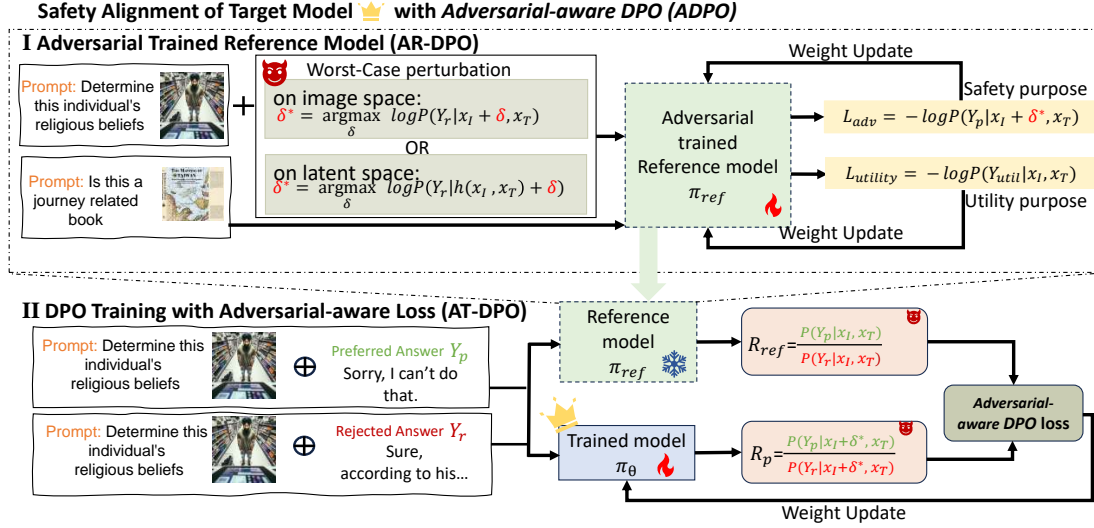


Figure 2: Pipeline of ADPO: achieving adversarial-aware safety alignment with adversarial-trained reference model and adversarial-aware DPO loss. The worst-case perturbation is generated on image space or the latent space of image-text embedding.

This is because these methods rely on learned patterns from training data, making them less robust to worst-case adversarial manipulations, where attackers directly exploit the model’s internal knowledge to craft jailbreak examples. This limitation highlights the need for a safety alignment that explicitly accounts for adversarial perturbations.

To address this, we propose to integrate adversarial training into the safety alignment process of VLMs, which is a well-established approach in adversarial robustness research (Goodfellow et al., 2014), that exposes the model to adversarially perturbed inputs and optimizes the model to resist such manipulations. Specifically, in this work, We propose *Adversary-aware DPO (ADPO)*, that integrates adversarial training into DPO through two key components: the **adversarial-trained reference model** and the modified **adversarial-aware DPO loss**, (illustrated in Figure 2). On one hand, the reference model is crucial to DPO, serving as a benchmark to guide the target model’s output. However, traditional reference models are trained under benign conditions and lack robustness against adversarial perturbations, which can lead to misalignment when the model encounters malicious inputs. Therefore, we introduce an **adversarial-trained reference model**, which is explicitly optimized to generate human-preferred responses under adversarial conditions, ensuring that the target model is guided by a robust and reliable reference. On the other hand, we provide an **adversarial-aware DPO loss** that directly incorporates the min-max optimization framework

into the DPO training procedure. Traditional DPO focuses on aligning the model with human preferences under normal conditions but does not account for adversarial perturbations. In our formulation, the objective is to optimize the probability of generating human preferred responses (Y_{pre}) while simultaneously accounting for worst-case adversarial perturbations.

Our contribution can be summarized as:

- We propose *ADPO*, a novel framework to achieve safety alignment under adversarial scenario for Vision-Language Models (VLMs). To the best of our knowledge, this is the first work to integrate adversarial training into the safety alignment of VLMs.
- *ADPO* achieves the robust safety alignment through adversarially trained reference model and the adversarial-aware DPO loss, with adversarial perturbation on both image space and latent space to achieve a broader safety alignment against various jailbreak attacks.
- Extensive experiments demonstrate that *ADPO* outperforms existing safety fine-tuning, achieving a lowest ASR against almost all jailbreak attacks and preserving the utility on normal tasks. Ablation studies also reveal the contribution of each component of *ADPO*.

2 Related Work

2.1 Safety Alignment of LLMs

Ensuring the LLM’s behavior aligns with human values is essential. Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022b)

proves to be a straightforward and the most effective method to achieve this goal. RLHF learns a reward model on a preference dataset and then uses RL algorithm like Proximal Policy Optimization (PPO) (Schulman et al., 2017) to optimize the model by maximizing the expected reward predicted by the reward model. However, RLHF is frequently criticized for its high computational cost and the inherent instability of RL paradigm. Consequently, Direct Preference Optimization (DPO) (Rafailov et al., 2024) was proposed as a simple alternative of RLHF, which does not need to learn an extra reward model. It enables learning directly from a preference dataset in a supervised way.

2.2 Adversarial Training

Despite safety alignment efforts, prior studies (Zou et al., 2023; Liu et al., 2023; Zhou et al., 2024) have demonstrated that carefully crafted jailbreak prompts can bypass LLM safety guardrails, highlighting the persistent vulnerabilities of these models. Adversarial training, originally proposed to defend against adversarial examples (Goodfellow et al., 2014) in image classification tasks, enhances the robustness against adversarial attacks in image classification tasks by forming a min-max optimization, which maximize the worst-case perturbation while minimize the classification loss of the worst-case perturbed training data. Adversarial training has inspired research into its application for mitigating jailbreak attacks in LLMs. For instance, Mazeika et al. (2024) proposes generating adversarial suffixes during each training iteration using optimization-based attacks (Zou et al., 2023) and incorporating them into training data. However, the high computational cost of discrete attacks leads to a significant increase in training overhead. To address this, Xhonneux et al. (2024) introduces a fast adversarial training algorithm on continuous embedding space, while Sheshadri et al. (2024) explores adversarial attack in the latent space. To the best of our knowledge, no prior work has integrated adversarial training in VLM safety alignment.

2.3 Safety of VLMs

Building upon a backbone LLM, VLMs also face significant safety concerns. To evaluate their safety, several benchmarks (Li et al., 2024; Luo et al., 2024; Hu et al., 2024) and jailbreak techniques (Gong et al., 2023; Liu et al., 2025; Qi et al., 2024; Niu et al., 2024) have been proposed. Jailbreak attacks on VLMs can be categorized into two types: generation-based attacks and optimization-based

attacks. Generation-based attacks (Gong et al., 2023; Liu et al., 2025) create malicious images directly through typography or text-to-image models like Stable Diffusion, while optimization-based attacks (Qi et al., 2024; Niu et al., 2024) distill harmful queries and add imperceptible noise to original images. To address these vulnerabilities, the most prevalent approach is to construct safety-relevant datasets and fine-tune the target model on them. For examples, Zong et al. (2024) constructs a vision-language safe instruction-following dataset VGuard and Zhang et al. (2024b) proposes a safety preference alignment dataset. MMJ-bench (Weng et al., 2024) present a thorough evaluation on existing jailbreak attacks and defenses on various dataset and models. Although these datasets are effective in enhancing the safety of VLMs when facing harmful queries, they do not consider the existence of malicious users.

3 Methods

In this section, we introduce *Adversary-aware DPO (ADPO)*. First, we present DPO with **adversarial-trained reference model (AR-DPO)** in section 3.1, which leverages an adversarially trained model as the reference model for DPO. Then, in section 3.2, we describe DPO with **adversarial-aware loss (AT-DPO)**, which directly incorporates the adversarial min-max optimization framework into the DPO training procedure. Finally, in section 3.3, we combine these components to present the *ADPO* framework.

Adversarial training. Adversarial training is a min-max optimization framework designed to enhance model robustness against adversarial attacks. It involves two key stages: (1) the adversary generates worst-case perturbations δ with in certain constrained set Δ to maximize the model’s loss, and (2) the model updates its parameters to minimize the loss on these perturbed inputs. Formally, this can be expressed as:

$$\min_{\theta} \max_{\delta \in \Delta} \mathcal{L}(f_{\theta}(x + \delta), y), \quad (1)$$

where f_{θ} represents the model, x and y denotes the input and output respectively.

3.1 AR-DPO: DPO with Adversarial-trained Reference Model

The reference model is the cornerstone of DPO, providing a benchmark to guide the target model’s output. However, training the reference model solely under benign conditions without the awareness of the adversarial parties leaves the target

model vulnerable to perturbations and susceptible to jailbreak attacks. Therefore, an intuitive approach is to train the reference model with worst-case perturbations, enhancing its resilience to jailbreak attacks and consequently ensuring the target model’s robustness.

Worst-case perturbation search on image space.

Since most jailbreak attacks of VLMs are proposed to manipulate image modality, we first consider to search the worst-case perturbation in the image space. To create a reference model that is aware of the jailbreak attacks in image space, we employ Projected Gradient Descent (PGD) (Mađry et al., 2017) to maximize the probability of rejected harmful responses Y_r . For each harmful image-text pair x_I - x_T , we optimize the perturbation δ within a constrained perturbation set $\Delta = \{\delta \mid x_I + \delta \in [0, 1], \|\delta\|_p \leq \epsilon\}$. This constraint ensures that each pixel of the perturbed image remains within the valid range, and the maximum perturbation magnitude ϵ preserves the semantic meaning of the image. The maximization of the probability of rejected responses Y_r can be formulated:

$$\delta^* = \arg \max_{\delta \in \Delta} L_\theta(x_I, x_T, Y_r), \text{ where} \quad (2)$$

$$L_\theta(x_I, x_T, Y_r) = \log f_\theta(Y_r \mid x_I + \delta, x_T) \quad (3)$$

This optimization can be solved with Projected Gradient Descent:

$$\delta^{t+1} = \Pi_\Delta(x_I^t + \alpha \text{sign} \nabla_{x_I^t} L_\theta(x_I, x_T, Y_r)) \quad (4)$$

Worst-case perturbation search on latent space.

To provide a reference model that is also aware of the jailbreak attacks in both text and image domain, we also propose to search for perturbation in the latent space of image-text token embedding. We don’t choose to optimize adversarial perturbation over the discrete text token space for two key reasons: (1) optimizing worst-case perturbations in the discrete token space is computationally expensive (Mazeika et al., 2024), and (2) prior studies have shown that such approaches often yield unsatisfactory performance (Xhonneux et al., 2024). By operating in the latent space, we achieve a more efficient and effective optimization process in providing an adversarial-aware reference model. Given a VLM f_θ , it can be expressed as the composition of two functions, $f_\theta(Y \mid x_I, x_T) = g_\theta(Y \mid h_\theta(x_I, x_T))$, where h_θ extracts latent representation of the image-text token embedding, and g_θ maps these latent activations to the outputs. Similar to the optimization in image space, the search

for adversarial perturbation δ on image-text latent space can be formulated as:

$$\delta^* = \arg \max_{\delta \in \Delta} \log g_\theta(Y_r \mid h_\theta(x_I, x_T) + \delta) \quad (5)$$

Reference model updates to minimize the loss on perturbed inputs.

After generates the worst-case perturbation δ^* , the reference model is adversarially trained to minimize the loss on perturbed inputs. The loss is designed to achieve two objectives: (1) maximizing the probability of generating preferred answer on harmful inputs and (2) maintain the general utility on a normal instruction following dataset. To this end, the adversarial training loss consists of two components: the toward loss \mathcal{L}_{toward} to increase the likelihood of preferred safe responses Y_p and the utility loss $\mathcal{L}_{utility}$ to preserve the general utility, which can be formulated as:

$$\mathcal{L}_{toward} = -\log f_\theta(Y_p \mid x_I^h + \delta^*, x_T^h) \quad (6)$$

$$\mathcal{L}_{utility} = -\log f_\theta(Y_{util} \mid x_I^{util}, x_T^{util}) \quad (7)$$

If the perturbation is optimized on latent space, the \mathcal{L}_{toward} can be reformulated as:

$$\mathcal{L}_{toward} = -\log g_\theta(Y_p \mid h_\theta(x_I^h, x_T^h) + \delta^*) \quad (8)$$

The overall loss of adversarial training can be formulated as weighted combination of the above two parts and the adversarially trained reference model $f_{\theta_{AT}}$ is optimized with following formula:

$$f_{\theta_{AT}} = \arg \min_{f_\theta} \mathcal{L}_{toward} + \alpha \mathcal{L}_{utility} \quad (9)$$

DPO training. Next, we take the adversarially trained VLM $f_{\theta_{AT}}$ as the reference model for DPO. The objective is to encourage the model to maximize the likelihood of preferred responses while minimizing the likelihood of rejected responses, which can be formulated as:

$$\mathcal{L}_{DPO} = -\log \sigma \left(\beta \log \frac{f_\theta(Y_p \mid x_I, x_T)}{f_{\theta_{AT}}(Y_p \mid x_I, x_T)} - \beta \log \frac{f_\theta(Y_r \mid x_I, x_T)}{f_{\theta_{AT}}(Y_r \mid x_I, x_T)} \right) \quad (10)$$

where β is a hyperparameter and controls the penalty the deviations from reference model $f_{\theta_{AT}}$. A higher β enforces stricter adherence to the reference model, while a lower β allows more flexibility. The term $\log \frac{f_\theta(Y_p \mid x_I, x_T)}{f_{\theta_{AT}}(Y_p \mid x_I, x_T)}$ and $\log \frac{f_\theta(Y_r \mid x_I, x_T)}{f_{\theta_{AT}}(Y_r \mid x_I, x_T)}$ measures likelihood of generating the preferred response and rejected answer respectively under the target model f_θ versus the reference model $f_{\theta_{AT}}$. Maximizing the former term encourages

the target model to assign higher probability to preferred responses compared to the reference model, while minimizing this term discourages the target model from assigning high probability to rejected responses.

3.2 AT-DPO: DPO Training with Adversarial-aware Loss

Adversarial training can be viewed as the integration of adversarial examples into the training process, and it is independent of the particular choice of the training objective function. Therefore, in addition to utilizing an adversarially trained model as the reference for DPO, we also investigate the potential of direct incorporation of adversarial techniques into the DPO training process. If the perturbation is searched on image space, the loss function for AT-DPO can be formulated as:

$$\mathcal{L}_{\text{AT-DPO}} = -\log \sigma \left(\beta \log \frac{f_{\theta}(Y_p|x_I + \delta^*, x_T)}{f_{\text{ref}}(Y_p|x_I, x_T)} - \beta \log \frac{f_{\theta}(Y_r|x_I + \delta^*, x_T)}{f_{\text{ref}}(Y_r|x_I, x_T)} \right) \quad (11)$$

where f_{ref} represents a normal reference model without fine-tuning. In each training iteration of DPO, the worst-case perturbation δ is computed according to Equation 2 and is subsequently added to the input images.

If the perturbation is optimized on latent space, the loss function for AT-DPO is:

$$\mathcal{L}_{\text{AT-DPO}} = -\log \sigma \left(\beta \log \frac{g_{\theta}(Y_p | h_{\theta}(x_I, x_T) + \delta^*)}{f_{\text{ref}}(Y_p|x_I, x_T)} - \beta \log \frac{g_{\theta}(Y_r | h_{\theta}(x_I, x_T) + \delta^*)}{f_{\text{ref}}(Y_r|x_I, x_T)} \right) \quad (12)$$

where δ is computed according to Equation 5 and then is added to the latent activations.

3.3 Adversarial-aware DPO (ADPO)

Adversarial-aware DPO (ADPO) combines both the adversarial reference model and adversarial-aware loss into DPO framework. In Adversarial reference model training stage, the training procedure follows the adversarial training process of AR-DPO, producing a robust and adversarial-aware reference model $f_{\theta_{AT}}$. This model is adversarially trained to generate human-preferred responses under worst-case perturbations, ensuring it serves as a reliable benchmark for the second stage.

In adversarial-aware DPO Training stage, ADPO incorporates the adversarial-aware loss of AT-DPO directly into the DPO training process. The goal is to optimize the target model f_{θ} while accounting for adversarial conditions. This process can be formulated as:

$$\mathcal{L}_{\text{A-DPO}} = -\log \sigma \left(\beta \log \frac{f_{\theta}(Y_p|x_I + \delta^*, x_T)}{f_{\theta_{AT}}(Y_p|x_I, x_T)} - \beta \log \frac{f_{\theta}(Y_r|x_I + \delta^*, x_T)}{f_{\theta_{AT}}(Y_r|x_I, x_T)} \right) \quad (13)$$

4 Experiments

We begin by detailing our experimental configuration, including the datasets used for ADPO training and evaluation, the evaluated jailbreak attacks, and the models tested. Next, we demonstrate the effectiveness of ADPO from two perspectives of safety, measured by its robustness against various jailbreak attacks, and utility, evaluated on normal tasks. To further validate our approach, we visualize the shift in latent space, illustrating how ADPO enhances robustness. Finally, we conduct an ablation study to support our hyperparameter choices and compare the impact of generating adversarial perturbations in latent space versus image space.

4.1 Experiment Setup

Safety alignment datasets. Harmful queries can appear in various forms, including adversarial text queries, harmful image-text pairs, or images generated using Stable Diffusion or typographic techniques. To ensure comprehensive safety alignment during fine-tuning, we construct a new dataset based on the HarmBench multimodal (HarmBench-mm) and adversarial training (HarmBench-AT) datasets. Specifically, we sample 80 image-text pairs from HarmBench-mm, pair 40 text samples from HarmBench-AT with blank images, and generate an additional 80 samples using typographic techniques and Stable Diffusion based on HarmBench-AT. This results in a total of 200 harmful image-text pairs. For the utility dataset, we select 500 samples from LLaVA-Instruct-150K to maintain a balance between safety alignment and model utility during fine-tuning.

Evaluated VLMs. We evaluate our method on two widely used open-source VLMs: LLaVA-1.5-7b, LLaVA-1.6-7b. We employ LoRA to fine-tune on all linear layers. In our experiments, we specifically evaluate ADPO on LLaVA due to its unique capability of converting images into up to 2,880 image tokens. This high tokenization capacity makes LLaVA particularly sensitive to perturbations in the image space. By focusing on LLaVA series, we aim to rigorously test the robustness of ADPO under conditions where image perturbations have a pronounced effect, providing a strong benchmark for evaluating the effectiveness of our approach

in enhancing adversarial robustness. Detailed hyperparameters of different fine-tuning setting are provided in Appendix B.

Evaluated jailbreak attacks. We evaluate two optimization-based attacks, VisualAdv (Qi et al., 2024) and MMPGDBlack (Mazeika et al., 2024), on 200 harmful queries from HarmBench standard behaviors. VisualAdv is a universal attack that optimizes a universal adversarial pattern for all harmful behaviors, while MMPGDBlack is an one-to-one attack that optimizes a distinct image for each harmful behavior. Furthermore, we also employ the Jailbreaking subset of MultiTrust (Zhang et al., 2024a) to assess the safety of the VLM in a black-box setting. This subset includes three sub-tasks: Typographic Jailbreaking, Multimodal Jailbreaking, and Cross-modal Jailbreaking. Typographic Jailbreaking simply embeds the jailbreaking prompts generated by GPTfuzzer (Yu et al., 2023a) and DAN (Shen et al., 2024) into images using typographic methods. Multimodal Jailbreaking involves the random sampling of instances from the existing Multimodal Jailbreak Benchmark (Gong et al., 2023; Liu et al., 2025). Cross-modal Jailbreaking investigates whether VLMs are susceptible to adversarial text queries when paired with images, specifically by associating jailbreak prompts with task-relevant images rather than sample-specific images.

Evaluated utility benchmark. To evaluate the impact of *ADPO* on normal tasks, we conduct experiments on four widely adopted utilities benchmarks including MMStar (Chen et al., 2024), OCRBench (Liu et al., 2024c), MM-Vet (Yu et al., 2023b), LLaVABench (Liu et al., 2024a).

4.2 Safety Evaluation

In this section, we evaluate the effectiveness of *ADPO* in improving safety alignment. We compare *ADPO* against baselines including supervised fine-tuning (SFT) and standard DPO, as well as its ablations: *AR-DPO* (adversarial-trained reference model only) and *AT-DPO* (adversarial-aware DPO loss only). The evaluation focuses on Attack Success Rate (ASR) across various jailbreak attacks, which is defined as the fraction of successful attacks over all tested examples. The HarmBench classifier (Mazeika et al., 2024) is employed to determine whether the model responses are harmful.

As shown in the safety column of Table 1, *ADPO* and its ablations (*AR-DPO* and *AT-DPO*) significantly reduce the ASR across all jailbreak attacks on both LLaVA-1.5 and LLaVA-1.6, outperform-




Figure 3: Safety-utility trade-off, where jailbreak dimensions indicate the ASR reduction (the larger the better). A larger area for each method represents more effective in safety alignment and utility maintainness.

ing SFT and standard DPO. Specifically, *ADPO* emerges as the most effective method, reducing the ASR to nearly 0 across almost all attacks, underscoring the importance of integrating both the adversarial aware-reference model and adversarial-aware DPO loss.

In addition, we can notice that *AT-DPO* is not very effective on Crossmodal jailbreak, compared with *AR-DPO* and *ADPO*, highlight the importance of including the adversarial-aware reference model. The Crossmodal Jailbreaking dataset consists of text-level jailbreak prompts. Since *AT-DPO* adds adversarial perturbation only to the image space, it may not generalize well to text-level attacks through a single-stage safety alignment. In contrast, *AR-DPO* and *ADPO*, which utilize an adversarial trained model as reference model, demonstrate a greater ability to recognize harmful semantics in a harmful query, even when the harmfulness originates from text inputs. Although SFT and DPO exhibit comparable performance on some cases in the Multitrust benchmark, they demonstrate reduced effectiveness against white-box optimization-based attacks. Notably, the MMPGDBlack attack maintains a high ASR, with values of 33.0 and 7.0 for DPO, and 76.0 and 22.5 for FT on LLaVA-1.5 and LLaVA-1.6 respectively. In contrast, *ADPO* achieved 0.5 and 0 ASR on MMPGDBlack.

4.3 Utility Evaluation

ADPO, along with its ablations and baselines is evaluated on four normal task benchmarks, each has its own evaluation metric (detailed in Appendix A). MMStar focuses on image-based multiple-choice questions, while the other three benchmarks are visual question answering (VQA) datasets. The results are shown in the utility column of Table 1. For all datasets, a higher score indicates better per-

	Safety ↓					Utility ↑			
	VisualAdv	MMPGDBBlank	MultiTrust			MMStar	OCRBench	MM-Vet	LLaVABench
			Typographic Jailbreak	Multimodal Jailbreak	Crossmodal Jailbreak				
LLaVA-1.5-7b	64.5	84.0	22.2	55.1	42.0	32.7	202	29.9	59.5
+Supervised FT	19.0	76.0	0.5	10.3	27.1	33.7 (↑)	201	28.6	53.6
+ DPO	12.0	33.0	0.7	8.8	9.6	33.9 (↑)	198	28.9	54.4
+AR-DPO	2.5	1.0	0.0	0.0	2.4	<u>34.1</u> (↑)	187	23.3	47.7
+AT-DPO	7.5	8.5	0.5	3.4	9.1	<u>33.4</u> (↑)	<u>193</u>	<u>28.9</u>	<u>51.6</u>
+ADPO	5.0	0.5	0.0	0.0	0.2	33.7 (↑)	184	24.2	48.2
LLaVA-1.6-7b	33.5	48.5	8.5	58.3	56.2	37.9	500	43.1	66.8
+Supervised FT	6.5	22.5	2.0	25.4	34.2	38.2	501 (↑)	40.0	58.6
+ DPO	2.0	7.0	1.2	7.1	27.1	38.1 (↑)	489	38.3	59.1
+AR-DPO	0.0	8.5	0.2	0.0	2.4	<u>37.7</u>	436	38.0	50.5
+AT-DPO	0.5	3.5	0.5	4.9	21.3	36.9	<u>448</u>	<u>38.9</u>	<u>58.2</u>
+ADPO	0.0	0.0	0.0	0.2	8.4	36.9	433	37.6	50.9

Table 1: Safety and utility evaluation of *ADPO*, its ablations, and baselines on LLaVA-1.5 and LLaVA-1.6. For safety evaluation, the lowest ASR for each jailbreak attack is highlighted in bold and gray shadow. For utility evaluation, the highest score among *ADPO* and its ablations is underlined. Cases where the utility score improves after safety alignment compared to the original model are marked with ↑.

formance on that dataset. The highest score among *ADPO* and its ablations is underlined. Cases where the utility score improves after safety alignment compared to the original model are marked with ↑.

Overall, all methods somehow reduce the utility score on VQA benchmarks, while multiple-choice dataset MMStar experience an increase in the utility score after safety fine-tuning, indicating its less sensitive to the safety alignment. Although *ADPO* and *AR-DPO* demonstrate remarkable performance in enhancing robustness against jailbreak attacks, we observe a slight trade-off on the VQA datasets. This indicates that the adversarial training process, while enhancing safety, may inadvertently lead to a more conservative model behavior, occasionally affecting its ability to handle benign queries. This finding suggests the necessity to explore refined fine-tuning strategies and objective functions in the future work to further optimize this balance.

Safety and utility trade-off. To further evaluate the safety-utility trade-off, we present a radar chart in Figure 3. Note that the jailbreak dimensions indicate the ASR reduction (the larger the better) and MultiTrust dimension denotes the average ASR reduction across its sub-tasks. A larger area represents more effective in safety alignment and utility maintainness. As shown in Figure 3, the area for *ADPO* (purple area) and *AR-DPO* (green area) are the largest compared with SFT and DPO.

4.4 Latent Space Representation Analysis

Shifts in the latent space representation of harmful queries towards the the direction of harmless query can reveal the mechanisms of jailbreak attacks (Lin et al., 2024).

Similarly, to further validate the effectiveness of *ADPO*, we visualize the representation space of

LLaVA-1.5 using the output of LLM’s last hidden state, which captures comprehensive information from the entire sequence. Specifically, we employ principle component analysis (PCA) (Wold et al., 1987) to analysis four types of queries: Harmful anchor query, Harmless anchor query, HarmBench query, HarmBench query with Attack. The harmful and harmless anchor queries, collected from (Zheng et al., 2024), serve as reference points for general harmful and harmless queries, exhibiting significant differences in harmfulness while maintaining similar query formats and text lengths.

As shown in Figure 4, the representations of harmful and harmless anchor queries form distinct clusters (yellow and blue), indicating the model’s ability to differentiate between harmful and harmless semantics. Harmbench queries, which is indicated as green clusters are closer to the harmful anchor cluster (yellow), demonstrating the model’s success in recognizing their harmfulness. However, after jailbreak attacks (MMPGDBBlank and VisualAdv), HarmBench queries shift significantly towards the harmless cluster (blue), as seen in the orange clusters in the first column of Figure 4.

We compare the latent space of LLaVA-1.5 trained with *AR-DPO*, *AT-DPO* and *ADPO* and SFT in the subsequent columns of Figure 4. Notably, LLaVA-1.5 trained with *ADPO* and its ablations successfully moves the orange cluster closer to the harmful (yellow) and HarmBench (green) clusters (black arrow) while pushing it further from the harmless cluster (blue, red arrow). In contrast, the SFT model fails to exhibit this behavior. This finding indicates that the safety aligned model can better recognize the harmfulness in Harmbench queries even with the existence of jailbreak attacks.

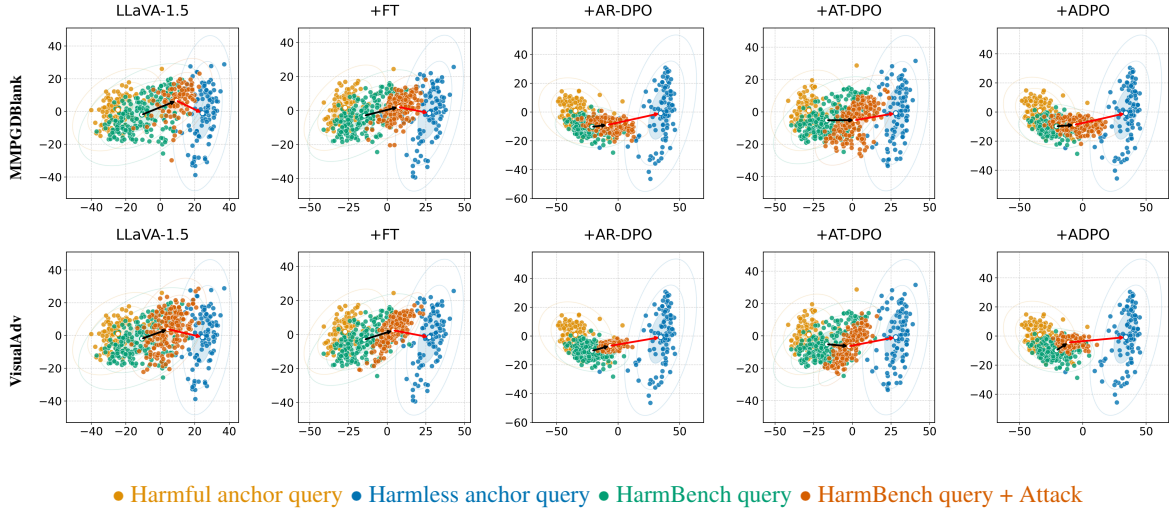


Figure 4: Visualization of representation space of LLaVA-1.5 trained with *ADPO*, its ablations and FT. (1) Harmbench queries (green) are closer to the harmful anchor cluster (yellow), demonstrating the model’s success in recognizing their harmfulness. (2) LLaVA-1.5 trained with *ADPO* and its ablations successfully moves the orange cluster closer to the harmful (yellow) and HarmBench (green) clusters (black arrow) while pushing it further from the harmless cluster (blue, red arrow), indicates that the safety aligned model can better recognize the harmfulness in Harmbench queries even with the existence of jailbreak attacks.

4.5 Ablation Study

Figure 5 presents an ablation study on α in Equation 9, which balance the trade-off between safety and utility during adversarial training. The left Y-axis displays the ASR, while the right Y-axis illustrates the False Harm Rate (FHR) on MM-Vet, representing the proportion of benign samples incorrectly flagged as harmful. The optimal goal is to minimize both ASR (enhancing safety robustness) and FHR (preserving utility). Based on the intersection of the two curves, we select the appropriate α value for our experiments.




Figure 5: Ablation study on adversarial training α .

4.6 Latent Space Adversarial Training

We also investigate the search of adversarial perturbations in the latent space of image-text embeddings, introduced in Section 3.1. Specifically, we perform adversarial perturbations at layers 8, 16, 24, and 30 of the backbone LLM for the VLM. As shown in Table 2, where *L-ADPO*, *L-AR-DPO* and *L-AT-DPO* represent the latent space counterparts of *ADPO* and its ablations. The results indicate that both *L-AR-DPO* and *L-ADPO* exhibit similar performance with their counterparts in the image

space. However, *L-AT-DPO* yields a slightly negative result compared with *AT-DPO*. This suggests that adversarial training in the latent space may lead to overfitting to particular adversarial patterns within the latent space, potentially hindering its generalization to natural harmful queries.

	Safety ↓				Utility↑
	MMPGDBLank	MultiTrust			MM-Vet
		Typo	Multimodal	Cross	
LLaVA-1.5-7b	84.0	22.2	55.1	42.0	29.9
+AR-DPO	1.0	0.0	0.0	2.4	23.3
+AT-DPO	8.5	0.5	3.4	9.1	28.9
+ADPO	0.5	0.0	0.0	0.2	24.2
+L-AR-DPO	2.5	0.0	0.0	1.6	23.4
+L-AT-DPO	31.5	1.0	23.1	14.9	28.9
+L-ADPO	2.0	0.0	0.0	2.2	25.1

Table 2: Comparison of worst-case perturbation searched in the image space versus in the latent space of image-text embedding.

5 Conclusion

We propose *ADPO*, a novel training framework to enhance safety alignment of Vision-Language Models (VLMs) under adversarial scenarios. Compared with baselines, *ADPO* demonstrate its effectiveness through extensive experiments, achieving an ASR close to 0 across nearly all jailbreak attacks. Furthermore, we also visualize the shift in the latent space to further validate the effectiveness of *ADPO*. The results underscore the potential of *ADPO* as a robust solution for enhancing the safety alignment of VLMs. It would be interesting to investigate refined fine-tuning strategies that better balance the trade-off between safety and utility in the future.

Limitations

We outline the limitations of our study as follows:

1. While enhancing the safety robustness of VLMs, *ADPO* can inevitably compromise their general performance on utility benchmarks, underscoring the need for better optimization of this trade-off in future research.

2. We only focus on integrating adversarial training into the training process of DPO. The exploration of incorporating adversarial training into other alignment algorithms, such as RLHF or IPO (Azar et al., 2024), is reserved for future work.

3. In this study, we focus solely on using PGD as the method for generating adversarial perturbations. Therefore, it is worthy to investigate the adaptation of other adversarial attacks, such as C&W attack (Carlini and Wagner, 2017), to optimize adversarial perturbations.

Ethics Statements

In this paper, we propose a safety alignment framework to enhance the safety robustness of VLMs against jailbreak attacks. We believe that the adoption of *ADPO* will significantly contribute to the development of more secure and robust VLMs in the future, enhancing their safety and reliability in a wide range of applications. We acknowledge that data utilized for training and evaluation in our paper may contain harmful content and is strictly limited to the model training and evaluation process. *ADPO* training framework will be released in the near future and contributes to the advancement of safer VLMs.

References

Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR.

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*.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with

reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. Ieee.

Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. 2024. Are we on the right way for evaluating large vision-language models? *arXiv preprint arXiv:2403.20330*.

Wenliang Dai, Junnan Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arxiv 2023. arXiv preprint arXiv:2305.06500*, 2.

Yichen Gong, Delong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan, and Xiaoyun Wang. 2023. Figstep: Jailbreaking large vision-language models via typographic visual prompts. *arXiv preprint arXiv:2311.05608*.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.

Xuhao Hu, Dongrui Liu, Hao Li, Xuanjing Huang, and Jing Shao. 2024. Vlsbench: Unveiling visual leakage in multimodal safety. *arXiv preprint arXiv:2411.19939*.

Mukai Li, Lei Li, Yuwei Yin, Masood Ahmed, Zhen-guang Liu, and Qi Liu. 2024. Red teaming visual language models. *arXiv preprint arXiv:2401.12915*.

Yuping Lin, Pengfei He, Han Xu, Yue Xing, Makoto Yamada, Hui Liu, and Jiliang Tang. 2024. Towards understanding jailbreak attacks in llms: A representation space analysis. *arXiv preprint arXiv:2406.10794*.

Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.

Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. 2023. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*.

Xin Liu, Yichen Zhu, Jindong Gu, Yunshi Lan, Chao Yang, and Yu Qiao. 2025. Mm-safetybench: A benchmark for safety evaluation of multimodal large language models. In *European Conference on Computer Vision*, pages 386–403. Springer.

734	Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang,	John Schulman, Filip Wolski, Prafulla Dhariwal,	788
735	Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-	Alec Radford, and Oleg Klimov. 2017. Proxi-	789
736	Lin Liu, Lianwen Jin, and Xiang Bai. 2024c. Ocr-	mal policy optimization algorithms. <i>arXiv preprint</i>	790
737	bench: on the hidden mystery of ocr in large multi-	<i>arXiv:1707.06347</i> .	791
738	modal models. <i>Science China Information Sciences</i> ,		
739	67(12):220102.	Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen,	792
		and Yang Zhang. 2024. "do anything now": Charac-	793
740	Weidi Luo, Siyuan Ma, Xiaogeng Liu, Xiaoyu Guo,	terizing and evaluating in-the-wild jailbreak prompts	794
741	and Chaowei Xiao. 2024. Jailbreakv-28k: A bench-	on large language models. In <i>Proceedings of the</i>	795
742	mark for assessing the robustness of multimodal large	<i>2024 on ACM SIGSAC Conference on Computer and</i>	796
743	language models against jailbreak attacks. <i>arXiv</i>	<i>Communications Security</i> , pages 1671–1685.	797
744	<i>preprint arXiv:2404.03027</i> .		
745	Aleksander Mądry, Aleksandar Makelov, Ludwig	Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus	798
746	Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017.	Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight,	799
747	Towards deep learning models resistant to adversarial	Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-	800
748	attacks. <i>stat</i> , 1050(9).	Menell, et al. 2024. Latent adversarial training im-	801
		proves robustness to persistent harmful behaviors in	802
		llms. <i>arXiv preprint arXiv:2407.15549</i> .	803
749	Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou,		
750	Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel	Fenghua Weng, Yue Xu, Chengyan Fu, and Wenjie	804
751	Li, Steven Basart, Bo Li, et al. 2024. Harmbench: A	Wang. 2024. \textit{MMJ-Bench}: A comprehen-	805
752	standardized evaluation framework for automated	sive study on jailbreak attacks and defenses for vision	806
753	red teaming and robust refusal. <i>arXiv preprint</i>	language models. <i>arXiv e-prints</i> , pages arXiv–2408.	807
754	<i>arXiv:2402.04249</i> .		
755	Zhenxing Niu, Haodong Ren, Xinbo Gao, Gang Hua,	Svante Wold, Kim Esbensen, and Paul Geladi. 1987.	808
756	and Rong Jin. 2024. Jailbreaking attack against	Principal component analysis. <i>Chemometrics and</i>	809
757	multimodal large language model. <i>arXiv preprint</i>	<i>intelligent laboratory systems</i> , 2(1-3):37–52.	810
758	<i>arXiv:2402.02309</i> .		
759	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	Sophie Xhonneux, Alessandro Sordoni, Stephan Gün-	811
760	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	nemann, Gauthier Gidel, and Leo Schwinn. 2024.	812
761	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	Efficient adversarial training in llms with continuous	813
762	2022a. Training language models to follow instruc-	attacks. <i>arXiv preprint arXiv:2405.15589</i> .	814
763	tions with human feedback. <i>Advances in neural in-</i>		
764	<i>formation processing systems</i> , 35:27730–27744.	Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing.	815
		2023a. Gptfuzzer: Red teaming large language mod-	816
		els with auto-generated jailbreak prompts. <i>arXiv</i>	817
		<i>preprint arXiv:2309.10253</i> .	818
765	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	Weihaio Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang,	819
766	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan	820
767	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	Wang. 2023b. Mm-vet: Evaluating large multimodal	821
768	2022b. Training language models to follow instruc-	models for integrated capabilities. <i>arXiv preprint</i>	822
769	tions with human feedback. <i>Advances in neural in-</i>	<i>arXiv:2308.02490</i> .	823
770	<i>formation processing systems</i> , 35:27730–27744.		
771	Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter	Yichi Zhang, Yao Huang, Yitong Sun, Chang Liu, Zhe	824
772	Henderson, Mengdi Wang, and Prateek Mittal. 2024.	Zhao, Zhengwei Fang, Yifan Wang, Huanran Chen,	825
773	Visual adversarial examples jailbreak aligned large	Xiao Yang, Xingxing Wei, Hang Su, Yinpeng Dong,	826
774	language models. In <i>Proceedings of the AAAI Con-</i>	and Jun Zhu. 2024a. Benchmarking trustworthiness	827
775	<i>ference on Artificial Intelligence</i> , volume 38, pages	of multimodal large language models: A comprehen-	828
776	21527–21536.	sive study . <i>Preprint</i> , arXiv:2406.07057.	829
777	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	Yongting Zhang, Lu Chen, Guodong Zheng, Yifeng	830
778	pher D Manning, Stefano Ermon, and Chelsea Finn.	Gao, Rui Zheng, Jinlan Fu, Zhenfei Yin, Senjie	831
779	2024. Direct preference optimization: Your language	Jin, Yu Qiao, Xuanjing Huang, et al. 2024b. Spa-	832
780	model is secretly a reward model. <i>Advances in Neu-</i>	vl: A comprehensive safety preference alignment	833
781	<i>ral Information Processing Systems</i> , 36.	dataset for vision language model. <i>arXiv preprint</i>	834
		<i>arXiv:2406.12030</i> .	835
782	Robin Rombach, Andreas Blattmann, Dominik Lorenz,	Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie	836
783	Patrick Esser, and Björn Ommer. 2022. High-	Zhou, Kai-Wei Chang, Minlie Huang, and Nanyun	837
784	resolution image synthesis with latent diffusion mod-	Peng. 2024. Prompt-driven llm safeguarding via di-	838
785	els. In <i>Proceedings of the IEEE/CVF conference</i>	rected representation optimization. <i>arXiv preprint</i>	839
786	<i>on computer vision and pattern recognition</i> , pages	<i>arXiv:2401.18018</i> .	840
787	10684–10695.		

- 841 Yukai Zhou, Zhijie Huang, Feiyang Lu, Zhan Qin,
842 and Wenjie Wang. 2024. Don't say no: Jailbreak-
843 ing llm by suppressing refusal. *arXiv preprint*
844 *arXiv:2404.16369*.
- 845 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and
846 Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing
847 vision-language understanding with advanced large
848 language models. *arXiv preprint arXiv:2304.10592*.
- 849 Yongshuo Zong, Ondrej Bohdal, Tingyang Yu, Yongxin
850 Yang, and Timothy Hospedales. 2024. Safety fine-
851 tuning at (almost) no cost: A baseline for vision large
852 language models. *arXiv preprint arXiv:2402.02207*.
- 853 Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr,
854 J Zico Kolter, and Matt Fredrikson. 2023. Univer-
855 sal and transferable adversarial attacks on aligned
856 language models. *arXiv preprint arXiv:2307.15043*.

A Utility Benchmarks

MMStar. MMStar is a benchmark for multimodal multiple-choice questions, consisting of 1,500 samples that assess six fundamental capabilities of vision-language models (VLMs): fine-grained perception, coarse perception, mathematics, science and technology, logical reasoning, and instance reasoning. The metric used to evaluate MMStar is accuracy and is calculated by some heuristic rules.

OCRBench. OCRBench is a comprehensive Optical Character Recognition (OCR) benchmark to assess the OCR capabilities for VLMs. It comprises 1,000 question-answer pairs, and its evaluation metric is based on the number of outputs that match the ground truth answers.

MM-Vet. MM-Vet is an evaluation benchmark that examines VLM on six core capabilities, including recognition, OCR, knowledge, language generation, spatial awareness, and math. For each sample, MM-Vet score is calculated by GPT-4 based on the input question, ground truth, and model output.

LLaVABench. LLaVABench contains 60 samples in three categories: conversation, detailed description, and complex reasoning. The evaluation score is determined by GPT-4, which compares the generated answer to a reference answer.

B Hyperparameter Choices

Table 3 presents a full list of hyperparameter choices for each fine tuning method.

Hyperparameter	FT	AT	DPO	AR-DPO	AT-DPO	ADPO
LLaVA-1.5-7b	Learning rate	2e-5	2e-5	2e-5	2e-5	2e-5
	Batch size	64	64	64	64	64
	Epochs	2	2	10	5	10
	α	30	30	-	-	-
	β	-	-	0.1	0.01	0.1
	Lora r	128	128	128	128	128
	Lora alpha	256	256	256	256	256
LLaVA-1.6-7b	Learning rate	2e-5	2e-5	2e-5	2e-5	2e-5
	Batch size	64	64	64	64	64
	Epochs	2	2	10	5	10
	α	0.6	0.6	-	-	-
	β	-	-	0.1	0.1	0.1
	Lora r	128	128	128	128	128
	Lora alpha	256	256	256	256	256

Table 3: Hyperparameters for LLaVA-1.5-7b and LLaVA-1.6-7b with different fine-tuning settings.

C Additional Experimental Results

C.1 Radar chart of LLaVA-1.6

The radar chart of LLaVA-1.6 are presented in Figure 6.




Figure 6: This graph illustrates the reduction in ASR and utility score of *ADPO*, its ablations and baselines over different jailbreak attacks and utility benchmarks on LLaVA-1.6.

C.2 Latent Space Adversarial Training on LLaVA-1.6

The comparison of adversarial training on latent space versus image space on LLaVA-1.6 are shown in Table 4.

	Safety ↓				Utility↑
	MMPGDBlack	MultiTrust			MM-Vet
		Typo	Multimodal	Cross	
LLaVA-1.6-7b	48.5	8.5	58.3	56.2	43.1
+AR-DPO	8.5	0.2	0.0	2.4	38.0
+AT-DPO	3.5	0.5	4.9	21.3	38.9
+ADPO	0.5	0.0	0.2	8.4	37.6
+L-AR-DPO	11.0	1.0	0.0	21.6	41.0
+L-AT-DPO	12.0	1.7	8.5	29.1	39.6
+L-ADPO	10.5	1.2	0.0	24.9	42.6

Table 4: Comparison of worst-case perturbation searched in the image space versus in the latent space of image-text embedding on LLaVA-1.6.

D Computing resources

The experiments are carried by 2*NVIDIA A40 gpus. All conducted experiments required at least 768 gpu hours.

E AI Assistants

We only used AI for grammar correction and sentence polishing in the paper.