
Adversarial Attacks against Closed-Source MLLMs via Feature Optimal Alignment

Xiaojun Jia¹, Sensen Gao², Simeng Qin^{1†}, Tianyu Pang³, Chao Du³,
Yihao Huang^{1†}, Xinfeng Li¹, Yiming Li¹, Bo Li⁴, Yang Liu¹

¹Nanyang Technological University, Singapore ²MBZUAI, United Arab Emirates

³Sea AI Lab, Singapore ⁴University of Illinois Urbana-Champaign, USA

{jiaxiaojunq, sensen.gao2002, qinsimeng670}@gmail.com;
{tianyupang3, duchao, lxfmakeit, liyiming.tech}@gmail.com;
lbo@illinois.edu; yangliu@ntu.edu.sg;

Abstract

Multimodal large language models (MLLMs) remain vulnerable to transferable adversarial examples. While existing methods typically achieve targeted attacks by aligning global features—such as CLIP’s [CLS] token—between adversarial and target samples, they often overlook the rich local information encoded in patch tokens. This leads to suboptimal alignment and limited transferability, particularly for closed-source models. To address this limitation, we propose a targeted transferable adversarial attack method based on feature optimal alignment, called FOA-Attack, to improve adversarial transfer capability. Specifically, at the global level, we introduce a global feature loss based on cosine similarity to align the coarse-grained features of adversarial samples with those of target samples. At the local level, given the rich local representations within Transformers, we leverage clustering techniques to extract compact local patterns to alleviate redundant local features. We then formulate local feature alignment between adversarial and target samples as an optimal transport (OT) problem and propose a local clustering optimal transport loss to refine fine-grained feature alignment. Additionally, we propose a dynamic ensemble model weighting strategy to adaptively balance the influence of multiple models during adversarial example generation, thereby further improving transferability. Extensive experiments across various models demonstrate the superiority of the proposed method, outperforming state-of-the-art methods, especially in transferring to closed-source MLLMs. The code is released at <https://github.com/jiaxiaojunQAQ/FOA-Attack>.

1 Introduction

Recent advancements in Large Language Models (LLMs) [49, 45, 3, 9, 1, 52, 53] have showcased extraordinary capabilities in language comprehension, reasoning, and generation. Capitalizing on the potent capabilities of Large Language Models (LLMs), a series of works [2, 30, 36, 63, 10] have attempted to seamlessly integrate visual input into LLMs, paving the way for the development of Multimodal Large Language Models (MLLMs). Commonly, these methods adopt pre-trained vision encoders, such as Contrastive Language Image Pre-training (CLIP) [47], to extract features from images and subsequently align them with language embeddings. MLLMs have achieved remarkable performance in vision-related tasks, including visual reasoning [34, 24], image captioning [32, 48], visual question answering [41, 29], etc. Beyond open-source advancements, commercial closed-source MLLMs such as GPT-4o, Claude-3.7, and Gemini-2.0 are widely adopted.

[†]Correspondence to Simeng Qin and Yihao Huang.




Figure 1: Targeted adversarial examples generated by FOA-Attack, with responses from commercial MLLMs to the prompt “Describe this image”.

Although large-scale foundation models have achieved remarkable successes, the security problems [15, 46, 65, 39, 27, 26, 25] associated with them are equally alarming and represent an ongoing challenge that remains unresolved. Recent works [13, 62, 17, 38, 44] have indicated that MLLMs are vulnerable to adversarial examples [19], as they inherit the adversarial vulnerability of vision encoders. The existence of adversarial examples poses significant security and safety risks to the real-world deployment of large-scale foundation models. Recently, some studies [4, 7, 50, 62] have delved into the adversarial robustness of MLLMs and have found that existing MLLMs remain vulnerable to adversarial attacks. Adversarial attacks on MLLMs are broadly classified as untargeted or targeted. Untargeted attacks aim to induce incorrect output, while targeted attacks force specific outputs. Adversarial transferability—the ability of adversarial examples to generalize across models—is critical for both types, especially in black-box settings where the target model is inaccessible. Targeted black-box attacks are particularly challenging [5, 64, 58]. Previous works integrate multiple pre-trained image encoders (*e.g.*, CLIP) to generate adversarial examples, which can significantly improve adversarial transferability. Notably, adversarial examples generated using open-source CLIP models can successfully carry out targeted attacks against closed-source commercial MLLMs. However, they achieve the limit improvement of adversarial transferability. Specifically, existing methods typically generate adversarial examples by minimizing contrastive loss between the global features of adversarial examples and target samples, where global features are often represented by the [CLS] token in open-source image encoders such as CLIP. While this strategy can produce semantically aligned adversarial samples in the feature space of the source model, it largely ignores the rich local features encoded by patch tokens. These local features contain fine-grained spatial and semantic details essential for comprehensive understanding in vision-language tasks. Neglecting them leads to weak alignment at the local level, resulting in adversarial perturbations that are less generalizable and highly dependent on the specific characteristics of the source model. Consequently, the generated adversarial examples tend to overfit the surrogate models and exhibit poor transferability to other models, especially commercial closed-source MLLMs.

To alleviate these issues, we propose FOA-Attack, a targeted transferable adversarial attack method based on optimal alignment of global and local features. Specifically, at the global level, we propose to adopt a coarse-grained feature alignment loss based on cosine similarity, encouraging the global features (*e.g.*, [CLS] tokens) of the adversarial example to align closely with those of the target sample. At the local level, previous works [14] indicate that the [CLS] token in the Transformer architecture represents global features, while other tokens represent local patch features. To fully extract the information from the target image, we use local features to generate adversarial samples. Although local features are rich, they are also redundant. We employ clustering techniques to distill compact and discriminative local patterns; that is, we use the features of the cluster centers to represent the characteristics of each cluster. We then formulate the alignment of these local features as an optimal transport (OT) problem and propose a local clustering OT loss to achieve fine-grained alignment between adversarial and target samples. Moreover, to further improve adversarial transferability, we propose a dynamic ensemble model weighting strategy that adaptively balances the weights of multiple models during adversarial example generation. Specifically, we generate adversarial samples using multiple CLIP image encoders, treating enhancement of feature similarity to the target sample across different encoders as separate tasks. The convergence of each objective

can be indicated by the rate at which its loss decreases—faster loss reduction implies a higher learning speed. Consequently, a higher learning speed results in a lower weight assigned to that objective. Extensive experiments demonstrate that the proposed FOA-Attack consistently outperforms state-of-the-art targeted adversarial attack methods, achieving superior transferability against both open-source and closed-source MLLMs. As shown in Fig. 1, the proposed FOA-Attack generates adversarial examples with superior transferability. Our main contributions are as follows:

- We propose FOA-Attack, a targeted transferable attack framework that jointly aligns global and local features, effectively guiding adversarial examples toward the target feature distribution and enhancing transferability.
- At the global level, we propose a cosine similarity-based global feature loss to align coarse-grained representations, while at the local level, we extract compact patch-level features via clustering and formulate their alignment as an optimal transport (OT) problem. Subsequently, we propose a local clustering OT loss for fine-grained alignment.
- We propose a dynamic ensemble model weighting strategy that adaptively balances multiple image encoders based on their convergence rates, substantially boosting the transferability of adversarial examples.
- Extensive experiments across various models are conducted to demonstrate that FOA-Attack consistently outperforms state-of-the-art methods, achieving remarkable performance even against closed-source MLLMs.

2 Related work

2.1 Multimodal large language models

Large language models (LLMs) have demonstrated remarkable performance in Natural Language Processing (NLP). Leveraging the impressive capabilities of LLMs, several studies have explored their integration with visual inputs, enabling strong performance across applications such as multimodal dialogue systems [2, 59, 1], visual question answering [54, 60, 23], etc. This integration marks a pivotal step toward the evolution of Multimodal Large Language Models (MLLMs). Existing studies achieve the integration of textual and visual modes through different strategies. Specifically, some studies focus on utilizing learnable queries to extract visual information and then adopt LLMs to generate text information based on the extracted visual features, such as Flamingo [2], BLIP-2 [30]. Some works propose to adopt several projection layers to align the visual features with text embeddings, such as PandaGPT [51], LLaVA [36, 37]. In addition, some works [16] propose to use some lightweight adapters to perform fine-tuning for performance improvement. Moreover, several studies [31, 42] have expanded the scope of research to include video inputs, utilizing the extensive capabilities of LLMs for enhanced video understanding tasks.

2.2 Adversarial attacks

Previous adversarial attack methods have primarily focused on image classification tasks. They usually utilize model gradients to generate adversarial examples, such as FGSM [18], PGD [43], C&W [6]. These studies have shown that deep neural networks are easily fooled by adversarial examples. Some studies [20, 55, 57] have demonstrated that MLLMs not only inherit the advantages of vision modules but also their vulnerabilities to adversarial examples. Adversarial attacks for MLLMs can be categorized as untargeted attacks and targeted attacks. Untargeted attacks aim to induce MLLMs to produce incorrect textual outputs, whereas targeted attacks aim to force specific, predetermined outputs. A series of recent works has paid more attention to the transferability of adversarial attacks, particularly in targeted scenarios. Adversarial transferability refers to the ability of adversarial examples generated on surrogate models to successfully attack unseen models. In particular, Zhao et al. [62] propose AttackVLM, involving generating targeted adversarial examples using pre-trained models like CLIP [47] and BLIP [30], and then transferring these examples to other VLMs such as MiniGPT-4 [63], LLaVA. They have demonstrated that image-to-image feature matching can improve adversarial transferability more effectively than image-to-text feature matching, a finding that has inspired subsequent research. Chen et al. [8] propose the Common Weakness Attack (CWA), a method that enhances the transferability of adversarial examples by targeting shared vulnerabilities among ensemble surrogate models. Subsequently, Dong et al. [13] propose




Figure 2: **Overview of the proposed FOA-Attack.** (a) The proposed feature optimal alignment loss which includes the coarse-grained feature loss and the fine-grained feature loss. (b) The proposed dynamic ensemble model weighting strategy.

the SSA-CWA method, which combines Spectrum Simulation Attack [40] (SSA) and Common Weakness Attack (CWA) to enhance the transferability of adversarial examples against closed-source commercial MLLMs like Google’s Bard. Guo et al. [22] propose AdvDiffVLM, a diffusion-based framework that integrates Adaptive Ensemble Gradient Estimation (AEGE) and GradCAM-guided Mask Generation (GCMG) to efficiently generate targeted and transferable adversarial examples for MLLMs. Zhang et al. [61] propose AnyAttack, a self-supervised framework, which trains a noise generator on the large-scale LAION-400M dataset using contrastive learning, to generate targeted adversarial examples for MLLMs without labels. Li et al. [33] propose the M-Attack method, which uses random cropping and resizing during optimization, to significantly improve the transferability of adversarial examples against MLLMs.

3 Methodology

Previous works show ensemble-based adversarial examples exhibit better transferability than single-model ones; thus, we employ a dynamic ensemble framework in this work. As shown in Fig. 2, the proposed FOA-Attack incorporates a feature optimal alignment loss and a dynamic ensemble weighting strategy to jointly enhance adversarial transferability across different foundation models.

3.1 Preliminary

Given an ensemble of image encoders from vision-language pre-training models $\mathcal{F} = \{f_{\theta_1}, f_{\theta_2}, \dots, f_{\theta_t}\}$, where each image encoder $f : \mathbb{R}^D \rightarrow \mathbb{R}^F$ outputs the image features for an input $\mathbf{x} \in \mathbb{R}^D$. Given a natural image $\mathbf{x}_{nat}$ and a target image $\mathbf{x}_{tar}$, the goal of the transfer-based attack is to generate an adversarial example $\mathbf{x}_{adv}$ whose features are as close as possible to those of the target image. It can be formulated as a constrained optimization problem:

$$\min_{\mathbf{x}_{adv}} \sum_{i=1}^t [\mathcal{L}(f_{\theta_i}(\mathbf{x}_{adv}), f_{\theta_i}(\mathbf{x}_{tar}))], \quad \text{s.t. } \|\mathbf{x}_{adv} - \mathbf{x}_{nat}\|_{\infty} \leq \epsilon, \quad (1)$$

where $\mathcal{L}$ represents the loss function, ϵ represents the maximum perturbation strength, and the adversarial examples are generated under the ℓ_{∞} norm.

3.2 The proposed coarse-grained feature optimal alignment

Given an image encoder (e.g., CLIP) f_{θ} , we extract the coarse-grained global features (e.g., [CLS] token) of the adversarial example $\mathbf{x}_{adv}$ as $\mathbf{X} = f_{\theta}(\mathbf{x}_{adv}) \in \mathbb{R}^{1 \times d}$, where d is the feature dimension. Similarly, the coarse-grained global feature of the target image is extracted as $\mathbf{Y} = f_{\theta}(\mathbf{x}_{tar}) \in \mathbb{R}^{1 \times d}$. To promote the adversarial example to align with the semantics of the target image at a global level, we minimize the negative cosine similarity between their coarse-grained features as the optimization

objective. The loss function can be defined as:

$$\mathcal{L}_{coa} = 1 - \cos(\mathbf{X}, \mathbf{Y}) = 1 - \frac{\langle \mathbf{X}, \mathbf{Y} \rangle}{\|\mathbf{X}\| \cdot \|\mathbf{Y}\|}, \quad (2)$$

where $\langle \mathbf{X}, \mathbf{Y} \rangle$ is the inner product and $\|\cdot\|$ is the ℓ_2 norm.

3.3 The proposed fine-grained feature optimal alignment

Given an image encoder (e.g., CLIP) f_θ , we extract the fine-grained local features (e.g., patch tokens) of the adversarial example and the target image. They can be defined as:

$$\mathbf{X}_{loc} = f_\theta^{loc}(\mathbf{x}_{adv}) \in \mathbb{R}^{m \times d}, \quad \mathbf{Y}_{loc} = f_\theta^{loc}(\mathbf{x}_{tar}) \in \mathbb{R}^{m \times d} \quad (3)$$

where $\mathbf{X}_{loc}$ and $\mathbf{Y}_{loc}$ represent the local features of the adversarial sample and the target image respectively, f_θ^{loc} represents the image features extracted from patch tokens of the image encoder, and m represents the number of patch or local features. Since local features contain fine-grained image information as well as more redundant image information, to reduce redundancy and retain discriminative information from the local features, we apply K-means clustering on $\mathbf{X}_{loc}$ and $\mathbf{Y}_{loc}$ to obtain representative cluster centers. Formally, we define:

$$\mathbf{X}_{clu} = \text{KMeans}(\mathbf{X}_{loc}, n) \in \mathbb{R}^{n \times d}, \quad \mathbf{Y}_{clu} = \text{KMeans}(\mathbf{Y}_{loc}, n) \in \mathbb{R}^{n \times d}, \quad (4)$$

where $\mathbf{X}_{clu}$ and $\mathbf{Y}_{clu}$ denote the n cluster centers obtained from the local features of the adversarial and target images, respectively. Each cluster center summarizes a semantically coherent region in the original image feature space, thus providing a more compact and informative representation for alignment. In our modeling of fine-grained local feature loss, we have drawn inspiration from the theory of optimal transport [56]. This theory was proposed by Villani with the objective of achieving the transportation of goods at minimal cost. In our study, we model the local features of the adversarial example and the target image as two separate distributions. Our goal is to identify the most efficient transportation scheme to more appropriately match the features of the target image onto the adversarial example, which can facilitate the transition between the two distributions. Let $\mu = \{\mathbf{X}_{clu}^a\}_{a=1}^n$ represent the distribution of clustering local features in the adversarial example, where n is the number of clustering local features, and $\mathbf{X}_{clu}^a$ denotes the a -th clustering local feature. Similarly, let $\nu = \{\mathbf{Y}_{clu}^b\}_{b=1}^n$ represent the distribution of clustering local features in the target image, with $\mathbf{Y}_{clu}^b$ representing the b -th clustering local feature. The cost function $c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b)$ defines the cost of transporting a feature from $\mathbf{X}_{clu}^a$ in the adversarial example to $\mathbf{Y}_{clu}^b$ in the target image. Hence, the optimization problem is formulated as:

$$\min \sum_{a=1}^n \sum_{b=1}^n c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b) \cdot \pi_{ab}, \quad \text{s.t.} \quad \forall a, \sum_{b=1}^n \pi_{ab} = 1; \quad \forall b, \sum_{a=1}^n \pi_{ab} = 1; \quad \forall a, b, \pi_{ab} \geq 0,$$

where the matrix π represents the transport plan between the features of the adversarial examples and target images. Each element π_{ab} of this matrix indicates the proportion of the a -th feature from the adversarial example that is assigned to the b -th feature in the target image. The constraints ensure the alignment of local features in accordance with μ and ν . The cost function is commonly computed using the negative cosine similarity as below:

$$c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b) = 1 - \langle \mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b \rangle, \quad (5)$$

The Sinkhorn algorithm [11] is employed to solve this optimal transport problem. Let $C_{ab} = c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b)$ be the cost of transporting the a -th local feature of the adversarial example to the b -th local feature of the target image. Local feature loss begins by defining the cost matrix:

$$C_{ab} = c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b), \quad \forall a, b \quad (6)$$

Then iteratively update u and v :

$$u_a = \frac{1}{n} \left(\sum_b \exp \left(-\frac{C_{ab}}{\lambda} \right) v_b \right)^{-1}, \quad v_b = \frac{1}{n} \left(\sum_a \exp \left(-\frac{C_{ab}}{\lambda} \right) u_a \right)^{-1}, \quad (7)$$

where $\lambda > 0$ is the regularization parameter (default: $\lambda = 0.1$). The transport plan is:

$$\pi_{ab} = u_a \exp\left(-\frac{C_{ab}}{\lambda}\right) v_b. \quad (8)$$

Finally, the local feature loss is:

$$\mathcal{L}_{fin} = \sum_{a,b} C_{ab} \cdot \pi_{ab}. \quad (9)$$

Finally, the total loss of FOA-Attack for the image encoder f_θ can be defined as:

$$\mathcal{L}_\theta = \mathcal{L}_{coa} + \eta \cdot \mathcal{L}_{fin}, \quad (10)$$

where η is the weighting factor that balances the local loss component. To handle varying local feature complexity, we adopt a progressive strategy that increases the number of cluster centers if the attack fails. In this paper, the number of centers is set to 3 and 5.

3.4 The proposed dynamic ensemble model weighting strategy

Building upon prior work, we generate adversarial examples using ensemble losses from multiple models to enhance adversarial transferability, computed as:

$$\mathcal{L} = \sum_{i=1}^t W_i \cdot \mathcal{L}_{\theta_i}, \quad (11)$$

where $\mathcal{L}_{\theta_i}$ represents the loss generated on the i -th image encoder and W_i represents the corresponding weight coefficient. Previous studies typically set all weights W_i at 1.0 without investigating the impact of varying W_i values on adversarial transferability, leading to limited improvements. Due to inconsistent vulnerabilities in different models, assigning uniform weights can cause optimization to favor certain losses. This often results in adversarial examples that are effective only on specific models, thereby reducing adversarial transferability. To further boost adversarial transferability, we propose a dynamic ensemble model weighting strategy to adaptively balance the weights of multiple models for adversarial example generation. Specifically, we generate adversarial examples using multiple CLIP image encoders, where improving the feature alignment between the adversarial and target samples on each encoder is treated as an independent optimization task. To balance these tasks, we monitor the convergence behavior of each objective by measuring the rate of loss reduction. A faster decrease in loss indicates a higher effective learning speed, suggesting that the task is easier to optimize. Hence, we assign a lower weight to objectives with higher learning speeds, ensuring that the optimization does not overemphasize the easily aligned tasks while neglecting others. At step $\mathbb{T}$, the learning speed is calculated by the loss ratio between steps $\mathbb{T}$ and $\mathbb{T} - 1$:

$$S_i(\mathbb{T}) = \frac{\mathcal{L}_{\theta_i}^{\mathbb{T}}(f_{\theta_i}(\mathbf{x}_{adv}), f_{\theta_i}(\mathbf{x}_{tar}))}{\mathcal{L}_{\theta_i}^{\mathbb{T}-1}(f_{\theta_i}(\mathbf{x}_{adv}), f_{\theta_i}(\mathbf{x}_{tar}))}, \quad (12)$$

where $\mathcal{L}_{\theta_i}$ is calculated by using Eq. (10) and $S_i(\mathbb{T})$ represents the learning speed of the adversarial example generation on the i -th model. The weight parameters in Eq. (11) can be calculated by:

$$W_i = W_{\text{init}} \times t \times \frac{\exp(S_i(\mathbb{T})/T)}{\sum_{j=1}^t \exp(S_j(\mathbb{T})/T)}, \quad (13)$$

where W_{init} denotes the initial setting of each W_i , consistent with the M-Attack configuration of 1.0. Multiplying by the number of surrogate models t scales the weights to fluctuate around 1.0, thereby refining the initialization. The temperature coefficient T further adjusts the relative differences between task weights. A detailed description of the algorithm is provided in the Appendix A.

4 Experiment

4.1 Settings

Datasets. Following previous works [13, 33], we use 1,000 clean images of size $224 \times 224 \times 3$ from the NIPS 2017 Adversarial Attacks and Defenses Competition dataset¹. Additionally, we randomly select 1,000 images from the MSCOCO validation set [35] as target images.

¹<https://nips.cc/Conferences/2017/CompetitionTrack>

Table 1: Performance of ASR (%) and AvgSim on different open-source MLLMs.

Method	Model	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	4.9	0.08	9.7	0.14	31.4	0.31	27.7	0.28	8.2	0.16	2.3	0.07
	B/32	8.7	0.12	13.3	0.17	11.3	0.14	9.5	0.12	8.4	0.15	1.7	0.05
	Laion	14.0	0.17	26.1	0.27	46.3	0.42	47.1	0.42	15.7	0.23	11.6	0.16
AdvDiffVLM [22]	Ensemble	2.1	0.01	2.5	0.01	1.5	0.01	1.6	0.01	0.7	0.00	0.8	0.01
SSA-CWA [13]	Ensemble	0.9	0.03	0.7	0.03	1.1	0.03	1.2	0.03	7.6	0.15	0.9	0.03
AnyAttack [61]	Ensemble	13.7	0.16	21.6	0.24	37.5	0.35	38.4	0.37	10.2	0.17	8.3	0.15
M-Attack [33]	Ensemble	38.6	0.35	52.6	0.46	68.3	0.56	67.1	0.56	23.0	0.29	21.3	0.25
FOA-Attack (Ours)	Ensemble	52.4	0.45	70.7	0.58	79.6	0.65	78.9	0.66	38.1	0.41	35.3	0.35


Implementation Settings. Following [33], we adopt three CLIP variants, which include ViT-B/16, ViT-B/32, and ViT-g-14-laion2B-s12B-b42K, as surrogate models to generate adversarial examples. The perturbation budget ϵ is set to $16/255$ under the norm ℓ_∞ . The attack step size is set to $1/255$. The number of attack iterations is set to 300. We evaluate the transferability of adversarial examples across fourteen MLLMs, including six open-source models (Qwen2.5-VL-3B/7B, LLaVa-1.5/1.6-7B, Gemma-3-4B/12B), five closed-source models (Claude-3.5/3.7, GPT-4o/4.1, Gemini-2.0), and three reasoning-oriented closed-source models (GPT-o3, Claude-3.7-thinking, Gemini-2.0-flash-thinking-exp). The text prompt of these models is set to “Describe this image.” All experiments are run on an Ubuntu system using an NVIDIA A100 Tensor Core GPU with 80GB of RAM.

Competitive Methods. We compare the proposed FOA-Attack with five advanced targeted and transfer-based adversarial attack methods for MLLMs: AttackVLM [62], SSA-CWA [13], AdvDiffVLM [22], AnyAttack [61], and M-Attack [33].

Evaluation metrics. Following [33], we adopt the widely used LLM-as-a-judge framework. Specifically, we use the same target MLLM to generate captions for both adversarial examples and target images, then assess their similarity using GPTScore. An attack is considered successful if the similarity score exceeds 0.5², which means that the adversarial example and the target image have the same subject. Additional results under varied thresholds are provided in the Appendix B. We report the attack success rate (ASR) and the average similarity score (AvgSim). For reproducibility, we include detailed evaluation prompts in the Appendix C.

4.2 Hyper-parameter Selection

We have two hyper-parameters in the proposed method: the temperature coefficient T and the weighting factor η . To study their effects, we conduct hyper-parameter selection experiments. As shown in Fig. 3 (a), setting $T = 1.0$ achieves the best trade-off between ASR and AvgSim, particularly on GPT-4o. While the ASR on Claude-3.5 shows minor variation, the performance on GPT-4o is more sensitive to T , with $T = 1.0$ leading to optimal semantic alignment. In Fig. 3 (b), we find that $\eta = 0.2$ consistently delivers the best performance on both models. A larger η overemphasizes the fine-grained loss, which slightly harms overall alignment. Therefore, we set $T = 1.0$ and $\eta = 0.2$ as the default values in our experiments.

Figure 3: (a) Impact of the temperature coefficient T ; (b) Impact of the weighting factor η .

4.3 Comparisons results

Comparisons with different attack methods. We compare our proposed FOA-Attack with several existing adversarial attack baselines, including AttackVLM, AdvDiffVLM, SSA-CWA, AnyAttack, and M-Attack, across both open-source and closed-source MLLMs. As shown in Table 1, on open-source models such as Qwen, LLaVa, and Gemma series, FOA-Attack consistently outperforms all baselines by a large margin. Specifically, it achieves an average ASR of 70.7% and 79.6% on Qwen2.5-VL-7B and LLaVa-1.5-7B, respectively, significantly surpassing the prior strongest method, M-Attack (52.6% and 68.3%). Moreover, FOA-Attack achieves the highest AvgSim scores

²This work adopts a stricter success threshold than the 0.3 used in M-Attack [33].

Table 2: Performance of ASR (%) and AvgSim on different closed-source MLLMs.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	0.1	0.02	0.2	0.03	16.2	0.21	17.5	0.22	7.0	0.12
	B/32	4.8	0.08	7.3	0.11	5.3	0.10	6.4	0.11	2.6	0.06
	Laion	0.3	0.02	1.2	0.03	39.7	0.38	42.4	0.39	28.9	0.30
AdvDiffVLM [22]	Ensemble	0.8	0.01	1.1	0.01	2.3	0.01	2.5	0.01	1.6	0.01
SSA-CWA [13]	Ensemble	0.4	0.02	0.4	0.03	0.5	0.03	0.2	0.02	0.4	0.02
AnyAttack [61]	Ensemble	4.6	0.09	4.3	0.08	8.2	0.15	7.3	0.13	6.1	0.12
M-Attack [33]	Ensemble	6.0	0.10	8.9	0.12	60.3	0.50	60.8	0.51	44.8	0.41
FOA-Attack (Ours)	Ensemble	11.9	0.16	15.8	0.18	75.1	0.59	77.3	0.62	53.4	0.50

across all models, indicating a better semantic alignment between adversarial and target captions. Table 2 further demonstrates the superiority of FOA-Attack on closed-source commercial MLLMs, including Claude-3, GPT-4, and Gemini-2.0. Notably, FOA-Attack yields 75.1% and 77.3% ASR on GPT-4o and GPT-4.1, outperforming M-Attack by 14.8% and 16.5%, respectively. On Gemini-2.0, FOA-Attack achieves a remarkable 53.4% ASR and 0.50 AvgSim, while other baselines perform poorly with ASRs below 8%. These results validate the effectiveness of our method across a wide range of both open- and closed-source MLLMs. We also provide the computation-time normalized comparison between M-Attack and FOA-Attack in the Appendix D. FOA-Attack results against defenses are in the Appendix E.

Comparisons on reasoning MLLMs. We further evaluate our FOA-Attack on 100 randomly selected images with reasoning-enhanced closed-source MLLMs, including GPT-o3, Claude-3.7-thinking, and Gemini-2.0-flash-thinking-exp, as shown in Table 3. Compared to the strong baseline M-Attack, our method consistently achieves higher ASR and AvgSim across all models. Specifically, on GPT-o3, FOA-Attack achieves an ASR of 81.0% and an AvgSim of 0.63, outperforming M-Attack by 14.0% and 0.09, respectively. Similarly, on Gemini-2.0-flash-thinking-exp, FOA-Attack improves ASR from 49.0% to 57.0% and AvgSim from 0.43 to 0.51. Even for the highly robust Claude-3.7-thinking model, our method raises ASR from 10.0% to 16.0%, along with a slight improvement in AvgSim. These results demonstrate that FOA-Attack remains highly effective even against reasoning-enhanced MLLMs, which are typically assumed to be more robust due to their advanced alignment and reasoning capabilities. However, our findings reveal that these models exhibit comparable or even weaker resistance to adversarial inputs than their non-reasoning MLLMs. This may stem from their reliance on textual reasoning, while shared visual encoders remain vulnerable to visual perturbations.

Table 3: Performance of ASR (%) and AvgSim on reasoning-enhanced closed-source MLLMs.

Method	Model	GPT-o3		Claude-3.7-thinking		Gemini-2.0-flash-thinking-exp	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
M-Attack [33]	Ensemble	67.0	0.54	10.0	0.15	49.0	0.43
FOA-Attack (Ours)	Ensemble	81.0	0.63	16.0	0.18	57.0	0.51

4.4 Ablation study

To understand the contribution of each component in FOA-Attack, we conduct an ablation study on 100 randomly selected images. As shown in Table 4, we systematically remove three core modules from FOA-Attack: global alignment, local alignment, and dynamic loss weighting. Removing global alignment results in a noticeable drop in performance, with ASR decreasing from 81.0% to 78.0% on GPT-4o and from 16.0% to 14.0% on Claude-3.5. It indicates the importance of aligning coarse-grained features for effective adversarial transferability. Excluding local alignment leads to a more significant degradation, especially in AvgSim, indicating that fine-grained feature alignment is essential for preserving semantic consistency between the adversarial and target samples. ASR on GPT-4o drops to 76.0%, and AvgSim decreases from 0.62 to 0.58. Lastly, removing dynamic loss weighting also reduces performance (e.g., 81.0% $\rightarrow$ 79.0% ASR on GPT-4o), showing that adaptively balancing optimization objectives also contributes to improving adversarial transferability. We further evaluate the role of K-Means clustering in local feature alignment by removing it and directly aligning all patch tokens. As shown in Table 5, removing K-Means leads to a clear drop in both attack success rate and feature similarity.

Table 4: Ablation study of our FOA-Attack.

Method	Claude-3.5		GPT-4o	
	ASR	AvgSim	ASR	AvgSim
M-Attack	10.0	0.13	73.0	0.56
FOA-Attack (Ours)	16.0	0.18	81.0	0.62
w/o global alignment	14.0	0.17	78.0	0.60
w/o local alignment	12.0	0.15	76.0	0.58
w/o dynamic loss weighting	13.0	0.17	79.0	0.61

Table 5: Ablation study on the necessity of the K-Means clustering step.

Method	GPT-4o ASR (%)	AvgSim	Claude-3.5 ASR (%)	AvgSim	Gemini-2.0 ASR (%)	AvgSim
FOA-Attack (w/o K-Means)	77.0	0.59	12.0	0.16	50.0	0.42
FOA-Attack (w/ K-Means)	81.0	0.62	16.0	0.18	56.0	0.46




Figure 4: Visualization of adversarial images and perturbation.

4.5 Performance analysis

Keyword matching rate (KMR). Previous work manually assigned three semantic keywords to each image and introduced three success thresholds— KMR_α (at least one matched), KMR_β (at least two matched), and KMR_γ (all three matched)—to evaluate attack transferability under different semantic matching levels. Following their setting, we compare the proposed method with previous works on 100 randomly selected images. As shown in Table 6, FOA-Attack consistently outperforms all baselines across different models (GPT-4o, Gemini-2.0, and Claude-3.5) and all keyword matching thresholds (KMR_α , KMR_β , KMR_γ), demonstrating superior targeted transferability. Notably, it achieves 92.0% on KMR_α and significantly higher scores on stricter metrics (76.0% KMR_β , 27.0% KMR_γ) on GPT-4o. Even on the more robust Claude-3.5, FOA-Attack achieves the best performance with 37.0% KMR_α . These results highlight the effectiveness of our FOA-Attack in enhancing adversarial transferability.

Table 6: Keyword Matching Rate (KMR) comparison across different models and attack methods.

Method	Model	GPT-4o			Gemini-2.0			Claude-3.5		
		KMR_α	KMR_β	KMR_γ	KMR_α	KMR_β	KMR_γ	KMR_α	KMR_β	KMR_γ
AttackVLM [62]	B/16	9.0	4.0	0.0	7.0	2.0	0.0	6.0	3.0	0.0
	B/32	8.0	2.0	0.0	7.0	2.0	0.0	4.0	1.0	0.0
	Laion	7.0	4.0	0.0	7.0	2.0	0.0	5.0	2.0	0.0
AdvDiffVLM [22]	Ensemble	2.0	0.0	0.0	2.0	0.0	0.0	2.0	0.0	0.0
SSA-CWA [13]	Ensemble	11.0	6.0	0.0	5.0	2.0	0.0	7.0	3.0	0.0
AnyAttack [61]	Ensemble	44.0	20.0	4.0	46.0	21.0	5.0	25.0	10.0	2.0
M-Attack [33]	Ensemble	82.0	54.0	13.0	75.0	53.0	11.0	31.0	18.0	3.0
FOA-Attack (Ours)	Ensemble	92.0	76.0	27.0	88.0	69.0	24.0	37.0	23.0	5.0

Sample visualization. Fig. 4 shows adversarial images and perturbations from different methods. Our method preserves image quality with minimal visible artifacts, while baselines such as AnyAttack and M-Attack introduce more noticeable noise. The perturbation maps on the right reveal that our method produces more structured and semantically aligned patterns, indicating stronger feature-level alignment and better adversarial transferability. Commercial MLLM responses are in the Appendix F.

Impact of more cluster centers. To enhance transferability, we adopt a progressive strategy that increases the number of cluster centers upon attack failure. We conduct experiments on 100 randomly selected images to explore the impact of more cluster centers. As shown in Table 7, incorporating more centers consistently im-

Table 7: Performance with varying cluster centers.

Method	Time (mins)	Claude-3.5		GPT-4o	
		ASR	AvgSim	ASR	AvgSim
M-Attack [33]	90	10.0	0.13	73.0	0.56
FOA-Attack ([3])	113	12.0	0.14	76.0	0.58
FOA-Attack ([3,5])	217	16.0	0.18	81.0	0.62
FOA-Attack ([3,5,8])	315	17.0	0.20	83.0	0.63
FOA-Attack ([3,5,8,10])	410	18.0	0.21	84.0	0.64

proves ASR and AvgSim, but also leads to higher time cost. To strike a balance between effectiveness and efficiency, we adopt the ([3,5]) setting in our main experiments. Moreover, in practice, the Sinkhorn algorithm needs minimal overhead compared to the dominant cost of gradient-based optimization in adversarial attacks. Since it only involves matrix-vector multiplications and element-wise operations, its computational cost is negligible relative to backpropagation. Specifically, FOA-Attack takes about 1.13 minutes to generate an adversarial example, whereas the Sinkhorn algorithm accounts for only around 0.015 minutes.

Table 8: Performance of attack transferability on more surrogate models.

Method	GPT-4o		Claude-3.5		Gemini	
	ASR (%)	AvgSim	ASR (%)	AvgSim	ASR (%)	AvgSim
M-Attack [33]	73.0	0.56	10.0	0.13	48.0	0.40
FOA-Attack	81.0	0.62	16.0	0.18	56.0	0.46
M-Attack (+SigLIP) [33]	73.0	0.56	11.0	0.13	50.0	0.41
FOA-Attack (+SigLIP)	82.0	0.62	19.0	0.20	60.0	0.50
M-Attack (+MetaCLIP) [33]	74.0	0.58	13.0	0.15	51.0	0.42
FOA-Attack (+MetaCLIP)	85.0	0.64	24.0	0.24	60.0	0.51
M-Attack (+SigLIP+MetaCLIP) [33]	75.0	0.58	13.0	0.15	53.0	0.43
FOA-Attack (+SigLIP+MetaCLIP)	86.0	0.64	26.0	0.25	62.0	0.53

Performance on more surrogate models. we additionally explore the effect of incorporating more diverse models into the ensemble. Specifically, we include SigLIP and MetaCLIP to construct a stronger ensemble. We conduct experiments on 100 randomly selected images. The results are in Table 8. It can be observed that expanding the ensemble to include more diverse models (SigLIP, MetaCLIP) improves transferability across all target models. Notably, our proposed FOA-Attack consistently outperforms M-Attack in all settings, confirming that the effectiveness of FOA-Attack is not solely dependent on the backbone ensemble size or diversity, but benefits from our proposed optimization strategies.

Table 9: Comparison of CLIP similarities between successful and failed samples.

CLIP Similarity	Qwen2.5-VL-3B	Qwen2.5-VL-7B	LLaVa-1.5-7B	LLaVa-1.6-7B	Gemma-3-4B	Gemma-3-12B
Success	0.4719	0.4742	0.4714	0.4711	0.4742	0.4767
Fail	0.4656	0.4560	0.4588	0.4607	0.4656	0.4646
CLIP Similarity	Claude-3.5	Claude-3.7	GPT-4o	GPT-4.1	Gemini-2.0	—
Success	0.4890	0.4762	0.4713	0.4714	0.4739	—
Fail	0.4662	0.4674	0.4610	0.4600	0.4604	—

Impact of Source-Target Similarity on ASR. To investigate whether visual similarity between source and target images affects attack success, we compute CLIP similarities for successful and failed samples across various models. Table 9 shows the average CLIP similarity scores for each case. On average, successful samples exhibit slightly higher CLIP similarity ($\Delta \approx +0.009$), suggesting that visual similarity between source and target images can marginally benefit attack success. However, the overall difference is small, indicating that our method does not rely heavily on image similarity and remains robust across diverse natural images.

5 Conclusion

In this work, we propose FOA-Attack, a targeted transferable adversarial attack framework that jointly aligns global and local features to improve transferability against both open- and closed-source MLLMs. Our method incorporates a global cosine similarity loss, a local clustering optimal transport loss, and a dynamic ensemble weighting strategy to comprehensively enhance adversarial transferability. Extensive experiments across various models demonstrate that the proposed FOA-Attack significantly outperforms existing state-of-the-art attack methods in both attack success rate and semantic similarity, especially on closed-source commercial and reasoning-enhanced MLLMs. These results reveal persistent vulnerabilities in MLLMs and highlight the importance of fine-grained feature alignment in designing transferable adversarial attacks. Further discussion, including limitations and broader impacts, is provided in the Appendix G.

Acknowledgments and Disclosure of Funding

This work is supported by the National Research Foundation, Singapore, and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG4-GC-2023-008-1B); by the National Research Foundation Singapore and the Cyber Security Agency under the National Cybersecurity R&D Programme (NCRP25-P04-TAICeN); and by the Prime Minister's Office, Singapore under the Campus for Research Excellence and Technological Enterprise (CREATE) Programme. Any opinions, findings and conclusions, or recommendations expressed in these materials are those of the author(s) and do not reflect the views of the National Research Foundation, Singapore, Cyber Security Agency of Singapore, Singapore.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35: 23716–23736, 2022.
- [3] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [4] Luke Bailey, Euan Ong, Stuart Russell, and Scott Emmons. Image hijacks: Adversarial images can control generative models at runtime. *arXiv preprint arXiv:2309.00236*, 2023.
- [5] Junyoung Byun, Seungju Cho, Myung-Joon Kwon, Hee-Seon Kim, and Changick Kim. Improving the transferability of targeted adversarial examples through object-based diverse input. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15244–15253, 2022.
- [6] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. Ieee, 2017.
- [7] Nicholas Carlini, Milad Nasr, Christopher A Choquette-Choo, Matthew Jagielski, Irena Gao, Anas Awadalla, Pang Wei Koh, Daphne Ippolito, Katherine Lee, Florian Tramèr, et al. Are aligned neural networks adversarially aligned? *arXiv preprint arXiv:2306.15447*, 2023.
- [8] Huanran Chen, Yichi Zhang, Yinpeng Dong, Xiao Yang, Hang Su, and Jun Zhu. Rethinking model ensemble in transfer-based adversarial attacks. *arXiv preprint arXiv:2303.09105*, 2023.
- [9] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113, 2023.
- [10] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, Yang Zhou, Kaizhao Liang, Jintai Chen, Juanwu Lu, Zichong Yang, Kuei-Da Liao, et al. A survey on multimodal large language models for autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 958–979, 2024.
- [11] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [12] Gavin Weiguang Ding, Luyu Wang, and Xiaomeng Jin. AdverTorch v0.1: An adversarial robustness toolbox based on pytorch. *arXiv preprint arXiv:1902.07623*, 2019.
- [13] Yinpeng Dong, Huanran Chen, Jiawei Chen, Zhengwei Fang, Xiao Yang, Yichi Zhang, Yu Tian, Hang Su, and Jun Zhu. How robust is google's bard to adversarial image attacks? *arXiv preprint arXiv:2309.11751*, 2023.

- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [15] Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022.
- [16] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023.
- [17] Sensen Gao, Xiaojun Jia, Xuhong Ren, Ivor Tsang, and Qing Guo. Boosting transferability in vision-language attacks via diversification along the intersection region of adversarial trajectory. *arXiv preprint arXiv:2403.12445*, 2024.
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [19] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6572>.
- [20] Xiangming Gu, Xiaosen Zheng, Tianyu Pang, Chao Du, Qian Liu, Ye Wang, Jing Jiang, and Min Lin. Agent smith: A single image can jailbreak one million multimodal llm agents exponentially fast. In *International Conference on Machine Learning*, 2024.
- [21] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Van Der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017.
- [22] Qi Guo, Shanmin Pang, Xiaojun Jia, Yang Liu, and Qing Guo. Efficient generation of targeted and transferable adversarial examples for vision-language models via diffusion models. *IEEE Transactions on Information Forensics and Security*, 2024.
- [23] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023.
- [24] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025.
- [25] Yihao Huang, Felix Juefei-Xu, Qing Guo, Jie Zhang, Yutong Wu, Ming Hu, Tianlin Li, Geguang Pu, and Yang Liu. Personalization as a shortcut for few-shot backdoor attack against text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [26] Yihao Huang, Le Liang, Tianlin Li, Xiaojun Jia, Run Wang, Weikai Miao, Geguang Pu, and Yang Liu. Perception-guided jailbreak against text-to-image models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.
- [27] Yihao Huang, Chong Wang, Xiaojun Jia, Qing Guo, Felix Juefei-Xu, Jian Zhang, Yang Liu, and Geguang Pu. Efficient universal goal hijacking with semantics-guided prompt organization. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, pages 5796–5816, July 2025.
- [28] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Hassan Foroosh. Comdefend: An efficient image compression model to defend adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6084–6092, 2019.

- [29] Jiayi Kuang, Ying Shen, Jingyou Xie, Haohao Luo, Zhe Xu, Ronghao Li, Yinghui Li, Xianfeng Cheng, Xika Lin, and Yu Han. Natural language understanding and inference with mllm in visual question answering: A survey. *ACM Computing Surveys*, 57(8):1–36, 2025.
- [30] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [31] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023.
- [32] Wei Li, Hehe Fan, Yongkang Wong, Yi Yang, and Mohan Kankanhalli. Improving context understanding in multimodal large language models via multimodal composition learning. In *Forty-first International Conference on Machine Learning*, 2024.
- [33] Zhaoyi Li, Xiaohan Zhao, Dong-Dong Wu, Jiacheng Cui, and Zhiqiang Shen. A frustratingly simple yet highly effective attack baseline: Over 90% success rate against the strong black-box models of gpt-4.5/4o/o1. *arXiv preprint arXiv:2503.10635*, 2025.
- [34] Zhiyuan Li, Dongnan Liu, Chaoyi Zhang, Heng Wang, Tengfei Xue, and Weidong Cai. Enhancing advanced visual reasoning ability of large language models. *arXiv preprint arXiv:2409.13980*, 2024.
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part v 13*, pages 740–755. Springer, 2014.
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [37] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024.
- [38] Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. Safety of multimodal large language models on images and text. *arXiv preprint arXiv:2402.00357*, 2024.
- [39] Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*, 2023.
- [40] Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xianglong Liu, Jian Zhang, and Jingkuan Song. Frequency domain model augmentation for adversarial attack. In *European conference on computer vision*, pages 549–566. Springer, 2022.
- [41] Duc-Tuan Luu, Viet-Tuan Le, and Duc Minh Vo. Questioning, answering, and captioning for zero-shot detailed image caption. In *Proceedings of the Asian Conference on Computer Vision*, pages 242–259, 2024.
- [42] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023.
- [43] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [44] Hefei Mei, Zirui Wang, Shen You, Minjing Dong, and Chang Xu. Veattack: Downstream-agnostic vision encoder attack against large vision language models. *arXiv preprint arXiv:2505.17440*, 2025.

- [45] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [46] Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. Red teaming language models with language models. *arXiv preprint arXiv:2202.03286*, 2022.
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [48] Sara Sarto, Marcella Cornia, and Rita Cucchiara. Image captioning evaluation in the age of multimodal llms: Challenges and future perspectives. *arXiv preprint arXiv:2503.14604*, 2025.
- [49] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- [50] Christian Schlarman and Matthias Hein. On the adversarial robustness of multi-modal foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3677–3685, 2023.
- [51] Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.
- [52] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothee Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [53] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [54] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021.
- [55] Haoqin Tu, Chenhang Cui, Zijun Wang, Yiyang Zhou, Bingchen Zhao, Junlin Han, Wangchunshu Zhou, Huaxiu Yao, and Cihang Xie. How many unicorns are in this image? a safety evaluation benchmark for vision llms. *arXiv preprint arXiv:2311.16101*, 2023.
- [56] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [57] Xuguang Wang, Zhenlan Ji, Pingchuan Ma, Zongjie Li, and Shuai Wang. Instructta: Instruction-tuned targeted attack for large vision-language models. *arXiv preprint arXiv:2312.01886*, 2023.
- [58] Phoenix Neale Williams and Ke Li. Black-box sparse adversarial attack via multi-objective optimisation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12291–12301, 2023.
- [59] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023.
- [60] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3081–3089, 2022.
- [61] Jiaming Zhang, Junhong Ye, Xingjun Ma, Yige Li, Yunfan Yang, Jitao Sang, and Dit-Yan Yeung. Anyattack: Towards large-scale self-supervised generation of targeted adversarial examples for vision-language models. *arXiv preprint arXiv:2410.05346*, 2024.

- [62] Yunqing Zhao, Tianyu Pang, Chao Du, Xiao Yang, Chongxuan Li, Ngai-Man Cheung, and Min Lin. On evaluating adversarial robustness of large vision-language models. *arXiv preprint arXiv:2305.16934*, 2023.
- [63] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- [64] Hegui Zhu, Xiaoyan Sui, Yuchen Ren, Yanmeng Jia, and Libo Zhang. Boosting transferability of targeted adversarial examples with non-robust feature alignment. *Expert Systems with Applications*, 227:120248, 2023.
- [65] Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Neil Zhenqiang Gong, Yue Zhang, et al. Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. *arXiv preprint arXiv:2306.04528*, 2023.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We have introduced our contributions and scope in the abstract and introduction sections.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: In the conclusion section, we discussed the limitations of the proposed method.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[No\]](#)

Justification: The formulas in the paper are simple to understand, and there is no theoretical conclusion.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have provided the loss calculation in our method. The dataset and model's sources are clearly identified. Detailed parameters for the experimental setup are also provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have submitted the code for review, and it is included anonymously in a single zip file.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide detailed information on the experimental setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We do not report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We introduce the details in setup of experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The paper is discussing on an AI security problem and involve the Deception & Harassment part. We have add an "Impact Statement" paragraph at the end of the submitted manuscript.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have added an "Impact Statement" paragraph at the end of the submitted manuscript to discuss the broader impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not release the model or dataset.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use the open-sourced dataset and give the url.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: There is no new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer:[Yes]

Justification: In this work, LLM was used to assist with language polishing and clarity improvement of the manuscript. The scientific content, experimental design, and analysis were entirely conducted by the authors.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A A Detailed Description of Our FOA-Attack

Following the M-Attack [33], we propose a targeted transferable adversarial attack method based on feature optimal alignment, called FOA-Attack. The detailed description of the proposed FOA-Attack is shown in Algorithm 1.

Algorithm 1: FOA-Attack

Input: clean image $\mathbf{x}_{\text{nat}}$, target image $\mathbf{x}_{\text{tar}}$, perturbation budget ϵ , iterations n , loss function $\mathcal{L}$, surrogate model ensemble $\mathcal{F} = \{f_{\theta_1}, f_{\theta_2}, \dots, f_{\theta_t}\}$, image processing $\mathcal{T}$, step size α

Output: adversarial image $\mathbf{x}_{\text{adv}}$

```

1 Initialize:  $\mathbf{x}_{\text{adv}}^0 = \mathbf{x}_{\text{nat}} + \delta_0$  (i.e.,  $\delta_0 = 0$ ); // Initialize adversarial image  $\mathbf{x}_{\text{adv}}$ 
2 for  $\mathbb{T} = 0$  to  $n - 1$  do
3    $\hat{\mathbf{x}}_i^a = \mathcal{T}(\mathbf{x}_{\text{adv}}^i)$ ,  $\hat{\mathbf{x}}^t = \mathcal{T}(\mathbf{x}_{\text{tar}})$ ;
4   ; // Perform random crop
5   for  $j = 1$  to  $t$  do
6      $\mathcal{L}_{coa} = 1 - \frac{\langle f_{\theta_j}(\hat{\mathbf{x}}_i^a), f_{\theta_j}(\hat{\mathbf{x}}^t) \rangle}{\|f_{\theta_j}(\hat{\mathbf{x}}_i^a)\| \cdot \|f_{\theta_j}(\hat{\mathbf{x}}^t)\|}$ ,
7      $\mathbf{X}_{loc} = f_{\theta_j}^{loc}(\mathbf{x}_{\text{adv}})$ ,  $\mathbf{Y}_{loc} = f_{\theta_j}^{loc}(\mathbf{x}_{\text{tar}})$ ,
8      $\mathbf{X}_{clu} = \text{KMeans}(\mathbf{X}_{loc}, n)$ ,  $\mathbf{Y}_{clu} = \text{KMeans}(\mathbf{Y}_{loc}, n)$ ,
9      $C_{ab} = c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b)$ ,  $\forall a, b$   $c(\mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b) = 1 - \langle \mathbf{X}_{clu}^a, \mathbf{Y}_{clu}^b \rangle$ ,
10     $u_a = \frac{1}{n} (\sum_b \exp(-\frac{C_{ab}}{\lambda}) v_b)^{-1}$ ,  $v_b = \frac{1}{n} (\sum_a \exp(-\frac{C_{ab}}{\lambda}) u_a)^{-1}$ ,
11     $\pi_{ab} = u_a \exp(-\frac{C_{ab}}{\lambda}) v_b$ ,
12     $\mathcal{L}_{fin} = \sum_{a,b} C_{ab} \cdot \pi_{ab}$ 
13     $\mathcal{L}_{\theta_j}^{\mathbb{T}} = \mathcal{L}_{coa} + \eta \cdot \mathcal{L}_{fin}$ ,
14    if  $\mathbb{T} == 0$  then
15       $S_j(\mathbb{T}) = 1$ ,
16    else
17       $S_j(\mathbb{T}) = \frac{\mathcal{L}_{\theta_j}^{\mathbb{T}}}{\mathcal{L}_{\theta_j}^{\mathbb{T}-1}}$ ,
18     $W_{\text{init}} = 1$ 
19    for  $j = 1$  to  $t$  do
20       $W_j = W_{\text{init}} \times t \times \frac{\exp(S_j(\mathbb{T})/T)}{\sum_{j=1}^t \exp(S_j(\mathbb{T})/T)}$ ,
21     $g_i = \frac{1}{m} \nabla_{\hat{\mathbf{x}}_i^a} \sum_{j=1}^m W_j \cdot \mathcal{L}_{\theta_j}$ ;
22     $\delta_{i+1} = \text{Clip}(\delta_i + \alpha \cdot \text{sign}(g_i), -\epsilon, \epsilon)$ ;
23     $\hat{\mathbf{x}}_{i+1}^a = \hat{\mathbf{x}}_i^a + \delta_{i+1}$ ;
24     $\mathbf{x}_{\text{adv}}^{i+1} = \hat{\mathbf{x}}_{i+1}^a$ 
25 return  $\hat{\mathbf{x}}_n^a$ 

```

Table 10: Performance (threshold is 0.3) of ASR (%) and AvgSim on different open-source MLLMs.

Method	Model	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	14.6	0.08	26.5	0.14	57.3	0.31	49.8	0.28	36.1	0.16	13.9	0.07
	B/32	22.4	0.12	31.6	0.17	27.3	0.14	23.1	0.12	35.0	0.15	9.1	0.05
	Laion	32.8	0.17	48.7	0.27	70.2	0.42	68.2	0.42	50.3	0.23	33.8	0.16
AdvDiffVLM [22]	Ensemble	2.7	0.01	3.1	0.01	1.9	0.01	2.1	0.01	0.9	0.00	1.2	0.01
SSA-CWA [13]	Ensemble	4.8	0.03	5.3	0.03	3.9	0.03	4.9	0.03	38.0	0.15	6.0	0.03
AnyAttack [61]	Ensemble	34.7	0.16	41.9	0.24	56.3	0.35	59.2	0.37	36.5	0.17	28.6	0.15
M-Attack [33]	Ensemble	63.3	0.35	80.2	0.46	89.8	0.56	87.4	0.56	64.3	0.29	50.3	0.25
FOA-Attack (Ours)	Ensemble	77.4	0.45	91.1	0.58	95.3	0.65	93.0	0.66	80.5	0.41	67.6	0.35

Table 11: Performance (threshold is 0.3) of ASR (%) and AvgSim on different closed-source MLLMs.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	2.4	0.02	4.1	0.03	40.8	0.21	42.6	0.22	23.5	0.12
	B/32	14.8	0.08	20.5	0.11	20.1	0.10	21.9	0.11	9.9	0.06
	Laion	3.5	0.02	4.9	0.03	69.9	0.38	71.8	0.39	55.8	0.30
AdvDiffVLM [22]	Ensemble	1.1	0.01	1.4	0.01	3.2	0.01	2.9	0.01	2.0	0.01
SSA-CWA [13]	Ensemble	3.2	0.02	3.7	0.03	3.8	0.03	3.0	0.02	4.0	0.02
AnyAttack [61]	Ensemble	19.1	0.09	18.7	0.08	40.8	0.15	39.5	0.13	31.1	0.12
M-Attack [33]	Ensemble	17.9	0.10	23.8	0.12	86.8	0.50	89.1	0.51	75.5	0.41
FOA-Attack (Ours)	Ensemble	28.4	0.16	36.4	0.18	94.8	0.59	95.6	0.62	86.7	0.50

B More Comparison Results under Varied Thresholds

We further evaluate the performance of FOA-Attack at the threshold of 0.3. As shown in Table 10, FOA-Attack consistently achieves superior adversarial success rates (ASR) and average semantic similarity (AvgSim) on open-source MLLMs, such as 95.3% ASR and 0.66 AvgSim on LLaVa-1.6-7B, significantly outperforming baseline ensemble attacks. Similarly, Table 11 highlights FOA-Attack’s strong transferability to closed-source models under the 0.3 threshold, achieving notably high performance (e.g., 95.6% ASR and 0.62 AvgSim on GPT-4.1), confirming its effectiveness and semantic alignment across diverse evaluation scenarios.

Table 12: Performance (threshold is 0.7) of ASR (%) and AvgSim on different open-source MLLMs.

Method	Model	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	2.0	0.08	5.3	0.14	17.9	0.31	16.6	0.28	3.9	0.16	0.7	0.07
	B/32	4.6	0.12	6.6	0.17	6.5	0.14	4.8	0.12	3.8	0.15	0.4	0.05
	Laion	8.0	0.17	15.7	0.27	31.2	0.42	32.8	0.42	8.1	0.23	4.1	0.16
AdvDiffVLM [22]	Ensemble	0.2	0.01	0.4	0.01	0.3	0.01	0.5	0.01	0.2	0.00	0.2	0.01
SSA-CWA [13]	Ensemble	0.3	0.03	0.5	0.03	0.5	0.03	0.2	0.03	3.0	0.15	0.1	0.03
AnyAttack [61]	Ensemble	11.6	0.16	17.3	0.24	26.7	0.35	23.2	0.37	5.8	0.17	6.4	0.15
M-Attack [33]	Ensemble	22.7	0.35	35.4	0.46	47.4	0.56	48.0	0.56	11.1	0.29	12.3	0.25
FOA-Attack (Ours)	Ensemble	35.2	0.45	53.1	0.58	62.5	0.65	63.6	0.66	23.2	0.41	19.6	0.35

Table 13: Performance (threshold is 0.7) of ASR (%) and AvgSim on different closed-source MLLMs.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	0.0	0.02	0.1	0.03	7.8	0.21	8.2	0.22	3.4	0.12
	B/32	2.4	0.08	3.3	0.11	3.0	0.10	3.0	0.11	0.9	0.06
	Laion	0.2	0.02	0.7	0.03	25.5	0.38	26.0	0.39	15.9	0.30
AdvDiffVLM [22]	Ensemble	0.1	0.01	0.2	0.01	0.5	0.01	0.4	0.01	0.2	0.01
SSA-CWA [13]	Ensemble	0.1	0.02	0.0	0.03	0.4	0.03	0.2	0.02	0.1	0.02
AnyAttack [61]	Ensemble	1.5	0.09	1.3	0.08	1.8	0.15	1.7	0.13	0.8	0.12
M-Attack [33]	Ensemble	3.3	0.10	4.4	0.12	38.8	0.50	39.8	0.51	26.6	0.41
FOA-Attack (Ours)	Ensemble	6.3	0.16	9.6	0.18	57.9	0.59	58.9	0.62	41.5	0.50

Continuing with the threshold set to 0.7, Table 12 shows FOA-Attack maintains its lead among open-source MLLMs, achieving significantly higher ASR and AvgSim, such as 62.5% ASR and 0.66 AvgSim on LLaVa-1.6-7B, notably surpassing all baseline ensemble methods. Similarly, results in Table 13 indicate that FOA-Attack retains effectiveness against challenging closed-source models even at the higher threshold, notably achieving 58.9% ASR and 0.62 AvgSim on GPT-4.1, reinforcing its strong adversarial transferability and semantic alignment in stringent attack scenarios.

Continuing with the threshold set to 0.8, Table 14 illustrates FOA-Attack’s superior transferability across open-source MLLMs, achieving notably high ASR and AvgSim (e.g., 44.1% ASR, 0.65

Table 14: Performance (threshold is 0.8) of ASR (%) and AvgSim on different open-source MLLMs.

Method	Model	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	1.2	0.08	2.7	0.14	8.7	0.31	10.1	0.28	3.4	0.16	0.2	0.07
	B/32	2.3	0.12	3.0	0.17	3.4	0.14	2.6	0.12	3.5	0.15	0.4	0.05
	Laion	4.1	0.17	8.6	0.27	19.1	0.42	23.2	0.42	6.0	0.23	2.0	0.16
AdvDiffVLM [22]	Ensemble	0.1	0.01	0.1	0.01	0.1	0.01	0.1	0.01	0.1	0.00	0.0	0.01
SSA-CWA [13]	Ensemble	0.2	0.03	0.1	0.03	0.3	0.03	0.1	0.03	2.6	0.15	0.0	0.03
AnyAttack [61]	Ensemble	4.6	0.16	7.3	0.24	11.9	0.35	13.4	0.37	2.8	0.17	2.2	0.15
M-Attack [33]	Ensemble	12.0	0.35	19.6	0.46	32.2	0.56	33.7	0.56	6.8	0.29	6.5	0.25
FOA-Attack (Ours)	Ensemble	20.2	0.45	34.2	0.58	44.1	0.65	47.6	0.66	14.2	0.41	11.1	0.35

Table 15: Performance (threshold is 0.8) of ASR (%) and AvgSim on different closed-source MLLMs.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	0.0	0.02	0.0	0.03	4.3	0.21	4.3	0.22	1.7	0.12
	B/32	1.1	0.08	1.5	0.11	1.3	0.10	1.5	0.11	0.3	0.06
	Laion	0.0	0.02	0.1	0.03	14.6	0.38	13.0	0.39	7.7	0.30
AdvDiffVLM [22]	Ensemble	0.0	0.01	0.0	0.01	0.2	0.01	0.1	0.01	0.1	0.01
SSA-CWA [13]	Ensemble	0.0	0.02	0.0	0.03	0.1	0.03	0.2	0.02	0.1	0.02
AnyAttack [61]	Ensemble	0.5	0.09	0.4	0.08	0.6	0.15	0.7	0.13	0.1	0.12
M-Attack [33]	Ensemble	1.6	0.10	1.7	0.12	23.6	0.50	23.0	0.51	14.7	0.41
FOA-Attack (Ours)	Ensemble	4.5	0.16	5.1	0.18	37.2	0.59	37.1	0.62	25.4	0.50

Table 16: Performance (threshold is 0.9) of ASR (%) and AvgSim on different open-source MLLMs.

Method	Model	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	0.3	0.08	0.6	0.14	3.8	0.31	4.2	0.28	2.7	0.16	0.0	0.07
	B/32	0.6	0.12	0.5	0.17	0.8	0.14	1.3	0.12	2.9	0.15	0.0	0.05
	Laion	1.1	0.17	2.1	0.27	6.6	0.42	10.2	0.42	3.3	0.23	0.2	0.16
AdvDiffVLM [22]	Ensemble	0.0	0.01	0.0	0.01	0.1	0.01	0.0	0.01	0.1	0.00	0.0	0.01
SSA-CWA [13]	Ensemble	0.1	0.03	0.0	0.03	0.2	0.03	0.0	0.03	2.3	0.15	0.0	0.03
AnyAttack [61]	Ensemble	1.3	0.16	1.7	0.24	5.2	0.35	6.4	0.37	0.9	0.17	0.3	0.15
M-Attack [33]	Ensemble	4.0	0.35	5.8	0.46	13.2	0.56	18.1	0.56	2.9	0.29	1.1	0.25
FOA-Attack (Ours)	Ensemble	5.6	0.45	10.8	0.58	22.4	0.65	27.2	0.66	6.5	0.41	2.8	0.35

Table 17: Performance (threshold is 0.9) of ASR (%) and AvgSim on different closed-source MLLMs.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
AttackVLM [62]	B/16	0.0	0.02	0.0	0.03	0.8	0.21	0.7	0.22	0.2	0.12
	B/32	0.1	0.08	0.2	0.11	0.1	0.10	0.1	0.11	0.1	0.06
	Laion	0.0	0.02	0.1	0.03	2.2	0.38	2.7	0.39	1.2	0.30
AdvDiffVLM [22]	Ensemble	0.0	0.01	0.0	0.01	0.1	0.01	0.0	0.01	0.1	0.01
SSA-CWA [13]	Ensemble	0.0	0.02	0.0	0.03	0.0	0.03	0.0	0.02	0.0	0.02
AnyAttack [61]	Ensemble	0.0	0.09	0.1	0.08	0.0	0.15	0.0	0.13	0.0	0.12
M-Attack [33]	Ensemble	0.1	0.10	0.1	0.12	4.7	0.50	6.3	0.51	2.1	0.41
FOA-Attack (Ours)	Ensemble	0.7	0.16	0.4	0.18	11.2	0.59	12.1	0.62	4.9	0.50

AvgSim on LLaVa-1.5-7B), substantially surpassing baseline methods. Similarly, in Table 15, FOA-Attack retains significant effectiveness against closed-source models even at this challenging threshold, notably reaching 37.2% ASR on GPT-4o and 37.1% ASR on GPT-4.1, while maintaining high AvgSim scores, reinforcing its exceptional adversarial transfer capability.

With an even stricter threshold of 0.9, Tables 16 and 17 show FOA-Attack still effectively maintains its superior adversarial transferability. In Table 16, FOA-Attack outperforms baseline ensemble attacks on open-source MLLMs, notably achieving 27.2% ASR and 0.66 AvgSim on LLaVa-1.6-7B. In the closed-source scenario (Table 17), FOA-Attack demonstrates notable effectiveness, particularly on GPT-4o and GPT-4.1 (11.2% and 12.1% ASR, respectively), continuing to exhibit strong semantic alignment (AvgSim ≥ 0.59). These results confirm FOA-Attack’s remarkable transferability even under highly stringent evaluation conditions.

Evaluation Prompt

Rate the semantic similarity between the following two texts on a scale from 0 to 1.

Criteria for similarity measurement:

- Main Subject Consistency:** If both descriptions refer to the same key subject or object (e.g., a person, food, an event), they should receive a higher similarity score.
- Relevant Description:** If the descriptions are related to the same context or topic, they should also contribute to a higher similarity score.
- Ignore Fine-Grained Details:** Do not penalize differences in phrasing, sentence structure, or minor variations in detail. Focus on whether both descriptions fundamentally describe the same thing.
- Partial Matches:** If one description contains extra information but does not contradict the other, they should still have a high similarity score.
- Similarity Score Range:**
 - 1.0:** Nearly identical in meaning.
 - 0.8-0.9:** Same subject, with highly related descriptions.
 - 0.7-0.8:** Same subject, core meaning aligned, even if some details differ.
 - 0.5-0.7:** Same subject but different perspectives or missing details.
 - 0.3-0.5:** Related but not highly similar (same general theme but different descriptions).
 - 0.0-0.2:** Completely different subjects or unrelated meanings.

Text 1: {input_text1}
Text 2: {input_text2}

Output only a single number between 0 and 1. Do not include any explanation or additional text.

Figure 5: Evaluation prompt template.

C Detailed Evaluation Prompt

Following M-Attack [33], we adopt the same way to evaluate the adversarial performance. Below is the detailed evaluation prompt used to assess semantic similarity between textual inputs: **ASR**: the “{input_text_1}” and “{input_text_2}” are used as placeholders for text inputs. The evaluation prompt template is shown in Fig. 5.

D Computation-time normalized comparison between M-Attack and FOA-Attack

To ensure that the performance improvements of FOA-Attack are not merely due to longer runtime, we conduct a computation-time normalized evaluation against the M-Attack baseline. Specifically, we increase M-Attack’s number of iterations to match the total computational budget of FOA-Attack (217 minutes). M-Attack (700 iterations) and M-Attack (800 iterations) require approximately 210 and 240 minutes, respectively, providing a fair basis for comparison. The results are summarized in Table 18.

Table 18: Computation-time normalized comparison between M-Attack and FOA-Attack.

Method	GPT-4o ASR (%)	AvgSim	Claude-3.5 ASR (%)	AvgSim	Gemini-2.0 ASR (%)	AvgSim	Compute Time (min)
M-Attack [33]	73.0	0.56	10.0	0.13	48.0	0.40	90
M-Attack [33] (700 iters)	76.0	0.57	9.0	0.11	52.0	0.44	210
M-Attack [33] (800 iters)	78.0	0.59	5.0	0.09	51.0	0.43	240
FOA-Attack (ours)	81.0	0.62	16.0	0.18	56.0	0.46	217

As shown in Table 18, FOA-Attack consistently outperforms both compute-matched M-Attack variants across all metrics. Compared to M-Attack (700 iters), FOA-Attack achieves +5%, +7%, and +4% higher ASR on GPT-4o, Claude-3.5, and Gemini, along with notably higher feature similarity (e.g., 0.62 vs. 0.57 on GPT-4o, 0.18 vs. 0.11 on Claude-3.5). Even when compared with M-Attack (800 iters), which consumes slightly more time, FOA-Attack maintains +3% higher ASR on GPT-4o, +11% on Claude-3.5, and +5% on Gemini. These findings demonstrate that FOA-Attack’s performance gains primarily stem from its methodological design rather than additional computational resources.

E Comparison Results on Series of Defense Methods

We evaluate the attack performance of FOA-Attack against a series of defense methods, including smoothing-based defenses [12] (Gaussian, Medium, and Average), JPEG compression [21], and Comdefend [28]. The experimental results on both open-source and closed-source MLLMs are shown in Table 19 and Table 20. Across all defenses, FOA-Attack consistently outperforms M-Attack in both ASR and AvgSim. On open-source models, FOA-Attack maintains a strong ASR (e.g., 25.0% vs. 13.0% under Comdefend on Qwen2.5-VL-7B), while preserving semantic alignment. On closed-source models, the advantage is even more evident. Under Comdefend, our FOA-Attack achieves 61.0% ASR on GPT-4o and 55.0% on GPT-4.1, while M-Attack drops below 10%. Even under JPEG, FOA-Attack maintains over 50% ASR with stable AvgSim values. These results indicate that the proposed FOA-Attack achieves superior adversarial transferability and resilience across diverse defense strategies.

Table 19: Attack performance of adversarial images against open-source Multimodal Large Language Models (MLLMs) after defense processing.

Defense	Method	Qwen2.5-VL-3B		Qwen2.5-VL-7B		LLaVa-1.5-7B		LLaVa-1.6-7B		Gemma-3-4B		Gemma-3-12B	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
Gaussian	M-Attack [33]	14.0	0.18	27.0	0.29	50.0	0.48	48.0	0.47	17.0	0.25	14.0	0.17
	FOA-Attack (Ours)	27.0	0.27	50.0	0.42	67.0	0.60	65.0	0.58	29.0	0.35	22.0	0.27
Medium	M-Attack [33]	17.0	0.21	35.0	0.33	44.0	0.41	41.0	0.39	13.0	0.18	6.0	0.10
	FOA-Attack (Ours)	36.0	0.31	60.0	0.45	62.0	0.54	60.0	0.53	18.0	0.25	9.0	0.16
Average	M-Attack [33]	9.0	0.14	20.0	0.23	38.0	0.36	36.0	0.36	11.0	0.18	8.0	0.12
	FOA-Attack (Ours)	22.0	0.24	38.0	0.35	57.0	0.51	56.0	0.51	28.0	0.33	11.0	0.17
JPEG	M-Attack [33]	13.0	0.20	35.0	0.35	60.0	0.51	59.0	0.50	29.0	0.34	22.0	0.27
	FOA-Attack (Ours)	29.0	0.32	58.0	0.49	77.0	0.63	77.0	0.62	50.0	0.44	44.0	0.42
Comdefend	M-Attack [33]	10.0	0.13	27.0	0.27	48.0	0.42	46.0	0.41	14.0	0.22	12.0	0.17
	FOA-Attack (Ours)	25.0	0.28	49.0	0.46	65.0	0.54	63.0	0.54	33.0	0.36	22.0	0.29

Table 20: Attack performance of adversarial images against closed-source Multimodal Large Language Models (MLLMs) after defense processing.

Method	Model	Claude-3.5		Claude-3.7		GPT-4o		GPT-4.1		Gemini-2.0	
		ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim	ASR	AvgSim
Gaussian	M-Attack [33]	2.0	0.04	5.0	0.06	57.0	0.45	53.0	0.44	29.0	0.29
	FOA-Attack (Ours)	3.0	0.06	6.0	0.07	72.0	0.57	71.0	0.57	50.0	0.42
Medium	M-Attack [33]	3.0	0.04	4.0	0.06	39.0	0.37	40.0	0.38	23.0	0.24
	FOA-Attack (Ours)	4.0	0.07	6.0	0.09	59.0	0.48	63.0	0.50	41.0	0.37
Average	M-Attack [33]	2.0	0.04	1.0	0.03	38.0	0.37	39.0	0.36	19.0	0.22
	FOA-Attack (Ours)	5.0	0.06	3.0	0.06	59.0	0.48	62.0	0.50	36.0	0.34
JPEG	M-Attack [33]	9.0	0.12	14.0	0.17	60.0	0.48	52.0	0.45	36.0	0.35
	FOA-Attack (Ours)	14.0	0.20	22.0	0.24	75.0	0.59	78.0	0.59	58.0	0.49
Comdefend	M-Attack [33]	2.0	0.04	5.0	0.08	35.0	0.35	37.0	0.37	22.0	0.25
	FOA-Attack (Ours)	6.0	0.07	11.0	0.15	61.0	0.49	63.0	0.51	38.0	0.39

F Commercial MLLM Response

To further validate the efficacy of FOA-Attack, we provide real-world interaction results indicating that adversarial examples can guide advanced commercial closed-source MLLMs, which include GPT-4o, GPT-o3, GPT-4.1, GPT-4.5, Claude-3.5-Sonnet, Claude-3.7-Sonnet, Gemini-2.0-Flash, and Gemini-2.5-Flash, to generate descriptions semantically aligned with the specified target images. Specifically, Fig. 6 to 13 correspond to the attack results on each of these models in order: Fig. 6 shows GPT-4o, Fig. 7 shows GPT-o3, Fig. 9 shows GPT-4.1, Fig. 8 shows GPT-4.5, Fig. 10 shows Claude-3.5-Sonnet, Fig. 11 shows Claude-3.7-Sonnet, Fig. 12 shows Gemini-2.0-Flash, and Fig. 13 shows Gemini-2.5-Flash. The consistent attack success across all models highlights the high transferability of the proposed FOA-Attack.

G Limitations and Impact Statement

Limitations. Although the proposed method demonstrates excellent performance in transferring target adversarial examples, it introduces additional computations, such as local OT loss, which

decrease the efficiency of generating adversarial examples. Enhancing the efficiency of these attacks will be a key focus of our future research.

Impact Statement. This paper proposes a method for targeting transferrable adversarial attacks on MLLMs using targeted multi-modal alignment. The proposed method, like previous adversarial attack methods, investigates adversarial examples in order to identify adversarial vulnerabilities in MLLMs. This effort aims to guide future research into improving MLLMs against adversarial attacks and developing more effective defense approaches. Furthermore, the victim MLLMs employed in this study are open-source models with publicly available weights. The research on adversarial examples will help shape the landscape of AI security.




Figure 6: Example responses from the commercial MLLM-GPT-4o to targeted attacks generated by our method.




Figure 7: Example responses from the commercial MLLM-GPT-o3 to targeted attacks generated by our method.






Figure 8: Example responses from the commercial MLLM-GPT-4.5 to targeted attacks generated by our method.

Target




Description 4.1 ▾



Describe this image, no longer than 25 words.

A group of zebras with overlapping stripes, some grazing on grass, creating a visually complex pattern of black and white lines.

Target




Description 4.1 ▾



Describe this image, no longer than 25 words.

An elephant standing in a shaded forest area, lifting its trunk while surrounded by foliage and wooden fencing in the background.

Target



Description 4.1 ▾



Describe this image, no longer than 25 words.

Several small birds are perched on a cylindrical bird feeder, pecking at the seeds, with greenery visible in the background.

Figure 9: Example responses from the commercial MLLM-GPT-4.1 to targeted attacks generated by our method.




Figure 10: Example responses from the commercial MLLM-Claude-3.5-Sonnet to targeted attacks generated by our method.




Figure 11: Example responses from the commercial MLLM-Claude-3.7-Sonnet to targeted attacks generated by our method.




Figure 12: Example responses from the commercial MLLM-Gemini-2.0-Flash to targeted attacks generated by our method.




Figure 13: Example responses from the commercial MLLM-Gemini-2.5-Flash to targeted attacks generated by our method.