# RETHINKING TRANSFERABLE TARGETED ATTACKS ON MULTIMODAL LARGE LANGUAGE MODELS VIA CAUSAL INTERVENTION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Multimodal Large Language Models (MLLMs) have demonstrated impressive capabilities in vision-language understanding, yet their robustness against targeted adversarial manipulation remains largely unexplored. Existing approaches on transferable targeted attacks mainly optimize target-related feature representations but overlook which visual components truly drive the transfer toward the target semantics. To address this limitation, we introduce a causal-guided framework that leverages target-aware counterfactual intervention to uncover the visual factors contributing to adversarial transfer. To achieve efficient estimation, we develop a gradient-based approximation strategy that evaluates the causal influence of visual components. The identified target-relevant components are then incorporated into a causal-guided feature optimization objective to improve alignment with the desired semantic representation. Extensive experiments across various MLLMs demonstrate that our approach consistently enhances targeted transferability over existing attack methods, validating the effectiveness of causal guidance for generating transferable adversarial examples.

## 1 INTRODUCTION

Multimodal Large Language Models (MLLMs) have recently achieved remarkable progress by combining visual perception with language reasoning through the integration of visual encoders and large language models (Cui et al., 2024; Li et al., 2023; Wang et al., 2024a). Their powerful multimodal reasoning abilities have enabled a wide range of applications, including visual assistance (Wu et al., 2023), content understanding (Cui et al., 2024), and content interaction (Wang et al., 2024a). Despite their impressive capabilities, recent studies (Zhao et al., 2023; Zhang et al., 2025a; Xie et al., 2025) have revealed that MLLMs remain vulnerable to adversarial perturbations, where carefully crafted inputs can fool the model into predicting the wrong output. Among different attack settings, transferable targeted adversarial attacks are particularly challenging, as they aim to generate adversarial examples that can mislead unseen black-box models while consistently driving them toward a predefined target response (Fu et al., 2025; Yan et al., 2026; Li et al., 2026b).

Existing targeted transfer attacks (Xie et al., 2025; Zhang et al., 2025a; Wang et al., 2025b) on MLLMs mainly improve transferability by aligning the representations of adversarial and target images. M-Attack (Li et al., 2026b) optimizes global feature alignment between the adversarial example and the target image to guide the model toward the desired response. More recently, FOA-Attack (Jia et al., 2026) further explores fine-grained feature alignment by incorporating local patch-level representations, aiming to exploit the rich visual information. However, such optimization objectives provide limited insight into which visual components actually drive the target response. The vision encoder in MLLMs represents an image as a sequence of visual tokens, which capture different semantic information and contribute differently to the final prediction (Zhang et al., 2025c; Liu et al., 2026). Treating all tokens equally may overlook the most important ones, leading to less effective optimization and reducing targeted transferability.

A natural question arises: *which visual components should be modified to most effectively induce the target response?* Existing approaches (Chefer et al., 2021; Jiang et al., 2025; Liu et al., 2026) attempt to estimate visual importance using attention maps, activation magnitudes, or gradient-based



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

sensitivity. Nevertheless, these measurements mainly capture observational correlations or local responsiveness between intermediate features and model outputs. They cannot reveal whether actively changing a specific visual component toward the target representation will actually improve the target response. For targeted adversarial attacks, understanding such intervention effects is crucial, as the desired objective is not merely identifying important regions, but discovering regions whose modification can causally promote target transfer.

To address this limitation, we introduce a counterfactual causal perspective into transferable targeted attacks on MLLMs. In addition to focusing on token alignment, we also intervene on intermediate visual representations to enhance the transferability. Specifically, we view visual token importance estimation from a target-oriented counterfactual perspective, where each visual token is intervened by replacing it with the corresponding representation from the target image. The resulting change in target alignment is then used to quantify its contribution. Since explicitly evaluating every token intervention requires substantial computational cost, we further derive a gradient-based approximation that estimates the intervention effect of all tokens. Based on the estimated causal scores, we identify target-relevant visual tokens and introduce a causal-guided representation alignment objective to selectively optimize their movement toward the target representation.

Our approach provides a new perspective for targeted transferable attacks by shifting the focus from image representation matching to intervention-based identification of target-driving visual components. Extensive experiments demonstrate that the proposed method achieves improved transferability over existing approaches, validating the effectiveness of counterfactual token attribution for generating target-oriented adversarial examples. The main contributions of this work are summarized as follows:

- We introduce a counterfactual causal formulation for transferable targeted attacks on MLLMs, which effectively identifies visual tokens that contribute to target transfer through target-oriented interventions.
- We develop an efficient gradient-based approximation to estimate token-level counterfactual effects, avoiding the computational overhead of individual forward interventions for each token.
- We further introduce a causal-guided representation alignment strategy that focuses optimization on high-contribution visual tokens, leading to improved targeted transferability across unseen black-box MLLMs.

## 2 Related Work

Transferable Adversarial Attacks on MLLMs. Multimodal Large Language Models (MLLMs) have recently demonstrated impressive understanding capabilities, but their vulnerability to adversarial perturbations has attracted increasing attention (Zhao et al., 2023; Wang et al., 2024b; Hu et al., 2025). Existing methods (Hu et al., 2025; Wang et al., 2021; Huang et al., 2025) mainly improve transferability by optimizing the semantic alignment between adversarial examples and target images (Xie et al., 2025). Early studies (Zhao et al., 2023; Wang et al., 2024b) on attacks against MLLMs mainly explored image-level representation matching. AttackVLM (Zhao et al., 2023) enhances targeted transferability by aligning image features between the source and target. Building on this idea, subsequent studies (Liu et al., 2025; Li et al., 2026a) explore more effective feature alignment strategies to further improve transferability. CWA (Chen et al., 2024) introduced sharpness-aware optimization to improve the robustness of adversarial perturbations, while SSA-CWA (Chen et al., 2024; Long et al., 2022) incorporated spectrum-based augmentation to enhance the transferable attack. More recently, M-Attack (Li et al., 2026b) improves transferability by exploiting local patch-level semantic information. FOA-Attack (Jia et al., 2026) uses local patch tokens and optimal transport for fine-grained target alignment, with dynamic ensemble weighting to improve targeted transferability (Chen et al., 2024; Hu et al., 2025). Existing methods align representations but overlook which visual tokens drive target transfer. We identify these tokens via counterfactual intervention and use them to guide adversarial optimization.

**Counterfactual Causal Inference.** Counterfactual causal inference aims to quantify the causal effect of an alternative intervention by estimating the outcome that would have occurred under a counterfactual condition (Rubin, 1974; Pearl, 2009; Pearl et al., 2016; Peters et al., 2017). Struc-

tural Causal Models (SCMs) provide a formal framework for estimating causal effects while accounting for potential confounding, enabling model predictions to be analyzed beyond observed correlations (Pearl et al., 2016; Schölkopf et al., 2021; Arjovsky et al., 2019). In computer vision, counterfactual reasoning has been widely adopted to analyze how model predictions change under controlled interventions on visual evidence. Rather than relying on passive correlations in the training distribution, these methods intervene on objects, regions, attributes, or latent factors to reveal the causal contribution of visual evidence. This perspective has been used to mitigate spurious correlations, improve visual explanations, and learn more disentangled representations (Wang et al., 2020; Qi et al., 2020; Goyal et al., 2019; Chang et al., 2018; Yang et al., 2021). These studies suggest that counterfactual analysis provides a principled mechanism for identifying causally relevant visual evidence. In multimodal foundation models, counterfactual reasoning has been explored for visual evidence estimation and the mitigation of spurious cross-modal correlations (Zhang et al., 2025b; Gao & Dong, 2025). Related causal formulations have also been applied to retrieval augmentation, prompt optimization, and hallucination mitigation (Wang et al., 2025a; Li et al., 2025). While these works demonstrate the potential of causal reasoning in multimodal understanding, it is still unexplored to utilize causal token intervention to enhance transferability.

# 3 PRELIMINARIES

## 3.1 TRANSFERABLE TARGETED ADVERSARIAL ATTACK

Given a source image $x_s$, a target image $x_t$, and a surrogate visual encoder $f$, the goal of a targeted transfer attack is to construct an adversarial example $x_{\text{adv}}$ on the surrogate model that can fool the unseen black-box models into predicting content that resembles that of the target image. Let $f(x)$ denote the feature representation of an input image $x$ extracted by the surrogate model. The objective of the targeted transferable adversarial attack is to minimize the feature distance between the adversarial image and the target image, which can be defined as:

$$\mathcal{L}_{\text{bas}} = 1 - \cos\left(f(x_{\text{adv}}), f(x_t)\right). \tag{1}$$

Accordingly, the generation of the adversarial example can be formulated as

$$x^*_{\text{adv}} = \arg\min_{x_{\text{adv}}} \mathcal{L}_{\text{bas}}, \qquad |x_{\text{adv}} - x_s|_\infty \leq \epsilon, \tag{2}$$

where $\epsilon$ is the maximum allowable perturbation budget.

## 3.2 COUNTERFACTUAL CAUSAL ANALYSIS

Counterfactual causal analysis focuses on how the system outcome changes when a variable is actively modified. We review three fundamental concepts relevant to counterfactual causal analysis: Structural Causal Model, causal interventions, and counterfactuals.

**Structural Causal Model.** A Structural Causal Model (SCM) describes a causal system through a set of variables and their corresponding structural mechanisms. Formally, the causal system is represented as

$$\mathcal{C} = (\mathcal{E}, \mathcal{Z}, \mathcal{G}, p(\mathcal{E})), \tag{3}$$

where $\mathcal{E}$ denotes the exogenous variables capturing external factors and uncertainties, $\mathcal{Z} = \{Z_1, \ldots, Z_n\}$ denotes the endogenous variables, and $\mathcal{G}$ represents the structural functions that specify the causal generation process of each endogenous variable. The distribution $p(\mathcal{E})$ characterizes the uncertainty of the exogenous variables.

**Causal Intervention.** In causal analysis, an intervention goes beyond observing a variable at a particular value. Instead, it deliberately changes the state of the variable and examines how this change influences downstream variables. For a variable $Z_i$, we denote $\text{do}(Z_i = \tilde{z}_i)$ as an intervention that actively sets it to a new state $\tilde{z}_i$. Under this operation, the original mechanism that determines $Z_i$ from its upstream variables is replaced by $Z_i := \tilde{z}_i$, while the remaining structural relationships are kept unchanged. The effect of this intervention then propagates through the original causal structure and is ultimately reflected in the outcome variable.

**Counterfactual Effect.** A counterfactual evaluates how the outcome of a factual instance would change under an alternative intervention. Let $\mathbf{z}$ denote the factual state with outcome $R(\mathbf{z})$. If only

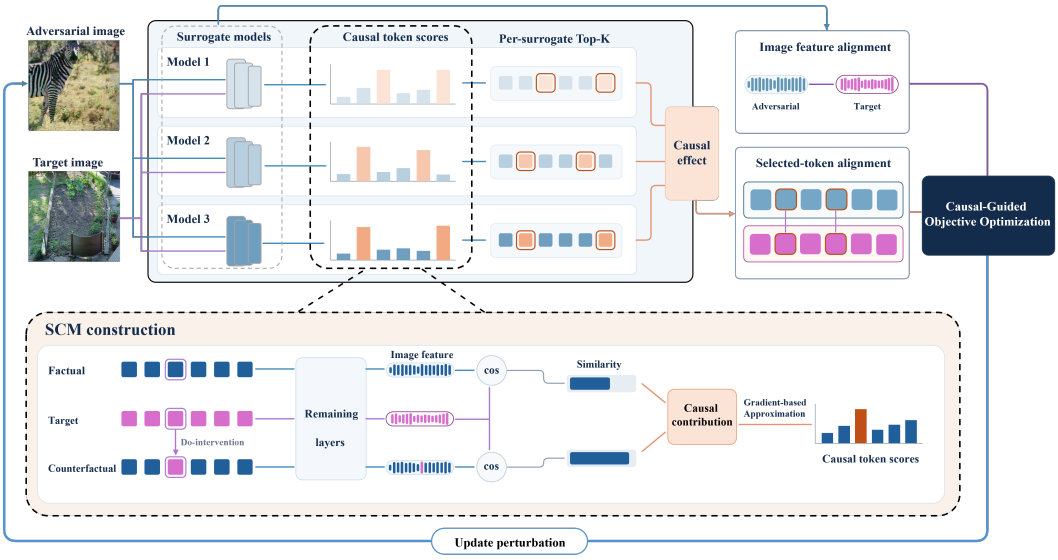


Figure 1: Overview of the proposed method. Multiple surrogate models estimate model-specific causal token scores and select Top-$K$ visual tokens. The causal tokens are combined with image-level feature alignment to form the causal-guided object optimization. The bottom panel illustrates SCM-based token intervention and gradient-based approximation for efficient causal contribution.

$Z_i$ is changed to $\tilde{z}_i$ while all other conditions remain fixed, the counterfactual outcome is $R_{Z_i \leftarrow \tilde{z}_i}(\mathbf{z})$. The corresponding instance-level effect is defined as $\mathcal{E}_i = R_{Z_i \leftarrow \tilde{z}_i}(\mathbf{z}) - R(\mathbf{z})$. Here, $\mathcal{E}_i$ measures the contribution of $Z_i$ to the target outcome: a positive value indicates a promoting effect, while a negative value indicates a suppressive effect.

## 4 METHOD

Existing targeted transfer attacks optimize representation alignment but overlook which visual tokens actually contribute to this process. To identify the tokens that truly drive target alignment, as shown in Figure 1, we introduce a counterfactual causal attribution strategy, which measures the effect of modifying each token toward the target representation. The identified tokens are then used to guide the optimization process to enhance transferability.

### 4.1 CAUSAL MODELING OF TARGET-ORIENTED INTERVENTION

Given a surrogate model $f$ containing $L$ Transformer layers, its forward propagation process can be written as:

$$H^0(x) = \mathrm{Embed}(x), H^l(x) = F^l\left(H^{l-1}(x)\right), \qquad l = 1, \ldots, L, \tag{4}$$

where $F^l(\cdot)$ denotes the mapping of the $l$-th Transformer layer, and $H^l(x)$ denotes the visual token of input image $x$ at layer $l$. The $l$-th embedding contains one CLS token and $N$ patch tokens and can be written as

$$H^l(x) = \left[h^l_{\mathrm{cls}}(x), h^l_1(x), \ldots, h^l_N(x)\right], \tag{5}$$

where $h^l_i(x) \in \mathbb{R}^D$ denotes the $i$-th patch token, and $D$ is the hidden feature dimension. The final image representation is presented as:

$$z(x) = \left[H^L(x)\right]_{\mathrm{cls}}. \tag{6}$$

The above forward process exhibits a layer-wise dependency structure, where each layer is computed from the output of the preceding layer, ultimately producing the final model output. Therefore, intermediate visual tokens can be treated as internal variables in a structural causal model, and their influence on the final representation can be analyzed through controlled interventions. By changing

the state of a specific token and observing how this change propagates through the model and affects the final target representation, we perform token-level interventions.

Given the $\ell$-th layer as the intervention layer, we obtain the corresponding embeddings of the adversarial image $x_{\text{adv}}$ and target image $x_t$:

$$H^{\ell}_{\text{adv}} = [h^{\ell}_{\text{cls,adv}}, h^{\ell}_{1,\text{adv}}, \ldots, h^{\ell}_{N,\text{adv}}], H^{\ell}_{\text{t}} = [h^{\ell}_{\text{cls,t}}, h^{\ell}_{1,\text{t}}, \ldots, h^{\ell}_{N,\text{t}}]. \tag{7}$$

where $h^{\ell}_{\text{i,adv}}$ denotes the $i$-th patch token of the adversarial image at the selected layer, and $h^{\ell}_{\text{i,t}}$ denotes the $i$-th patch token at the corresponding layer of the target image. $H^{\ell}_{adv}$ is treated as the factual state, while $H^{\ell}_{t}$ provides a reference state. For the $i$-th patch token, we consider the following target-oriented intervention:

$$\text{do}\left(h^{\ell}_{i,\text{adv}} = h^{\ell}_{i,\text{t}}\right), \tag{8}$$

which replaces the $i$-th token of the adversarial image with the corresponding token representation from the target image while keeping all other tokens at this layer unchanged. After applying this intervention, the modified intermediate state becomes

$$H^{\ell,(i\leftarrow t)}_{\text{adv}} = [h^{\ell}_{\text{cls,adv}}, h^{\ell}_{1,\text{adv}}, \ldots, h^{\ell}_{i-1,\text{adv}}, h^{\ell}_{i,\text{t}}, h^{\ell}_{i+1,\text{adv}}, \ldots, h^{\ell}_{N,\text{adv}}]. \tag{9}$$

This operation neither modifies the model parameters nor changes the current states of the other tokens. Instead, it actively changes only one local variable in the intermediate representation space. The intervened state $H^{\ell,(i\leftarrow t)}_{\text{adv}}$ is then propagated through the remaining Transformer layers:

$$\widetilde{H}^{l} = F^{l}\left(\widetilde{H}^{l-1}\right), \widetilde{H}^{\ell} = H^{\ell,(i\leftarrow t)}_{\text{adv}}, l = \ell+1, \ldots, L. \tag{10}$$

The representation under this intervention is

$$z^{(i\leftarrow t)}_{\text{adv}} = \left[\widetilde{H}^{L}\right]_{\text{cls}}. \tag{11}$$

Through the intervention process, we obtain the importance of each visual token by measuring its causal contribution to target representation alignment. Importantly, the intervention is directional and consistent with the target transfer objective rather than being an arbitrary feature replacement. Each token is moved toward its corresponding target representation, and the resulting change in target alignment reflects its contribution.

### 4.2 Gradient-Based Efficient Causal Approximation

We quantify the counterfactual causal effect of each token by replacing a single patch token and measuring the change in the target outcome. However, directly applying such interventions to all tokens requires an individual forward propagation for each candidate token after the intervention layer. This exhaustive evaluation is computationally expensive due to the large number of tokens and becomes particularly prohibitive in iterative adversarial optimization.

To efficiently estimate the counterfactual effect, we approximate the target-oriented intervention as a directional perturbation in the intermediate representation space. For the $i$-th patch token, the corresponding displacement toward the target representation is defined as:

$$\Delta h^{\ell}_{i} = h^{\ell}_{i,t} - h^{\ell}_{i,adv}. \tag{12}$$

Accordingly, the counterfactual intervention that replaces $h^{\ell}_{i,\text{adv}}$ with $h^{\ell}_{i,\text{t}}$ is equivalently formulated as moving the current token representation along the target-oriented direction $\Delta h^{\ell}_{i}$.

To avoid explicitly constructing a counterfactual state for each token, we approximate the intervention effect using a first-order Taylor expansion. Specifically, we define the outcome function $Y_{\ell}(\cdot)$ as the downstream objective after the intervention layer $\ell$, which maps an intermediate representation to the final target alignment score:

$$Y_{\ell}(H^{\ell}) = \mathcal{S}\left(z(H^{\ell}), z_{\text{t}}\right), \tag{13}$$

where $z(H^{\ell})$ denotes the final representation obtained by propagating $H^{\ell}$ through the remaining Transformer layers, and $\mathcal{S}(\cdot,\cdot)$ represents the target similarity function.

We approximate the counterfactual effect using a first-order Taylor expansion around the factual state $H^{\ell}_{adv}$. Specifically, the intervened outcome can be expressed as:

$$Y_\ell\left(H^{\ell,(i\leftarrow t)}_{\text{adv}}\right) = Y_\ell\left(H^{\ell}_{\text{adv}}\right) + \left(\frac{\partial Y_\ell}{\partial h^{\ell}_{i,\text{adv}}}\right)^{\top} \Delta h^{\ell}_{i} \quad + \mathcal{O}\left(\|\Delta h^{\ell}_{i}\|_2^2\right). \tag{14}$$

By ignoring the higher-order terms, the counterfactual causal effect of the $i$-th token can be approximated as:

$$\widehat{\mathcal{E}}_i = \left(\frac{\partial Y_\ell}{\partial h^{\ell}_{i,\text{adv}}}\right)^{\top} \left(h^{\ell}_{i,t} - h^{\ell}_{i,\text{adv}}\right). \tag{15}$$

A large positive value indicates that moving the token toward its target representation is more likely to increase target similarity, whereas a negative value suggests that such a change may be detrimental to the current target transfer objective. Notably, this formulation is different from conventional token importance estimation based on gradient norms. Gradient magnitude mainly reflects how sensitive the model output is to a particular token, but it does not account for the direction of the change.

After obtaining the approximate causal scores for all tokens, we select the $K$ patch tokens with the highest estimated scores:

$$\mathcal{S} = \text{TopK}\left(\left\{\widehat{\mathcal{E}}_i\right\}_{i=1}^{N}, K\right). \tag{16}$$

It should be noted that the causal contribution of each token changes dynamically as the adversarial perturbation is updated. As the current adversarial representation evolves, both the local gradient and the target displacement direction change accordingly. Therefore, we periodically recompute the approximate causal scores and update the selected token set at a predefined interval. The causal scores are used only to determine the optimization targets in the subsequent stage. After the Top-$(K)$ selection is completed, gradient propagation through the scoring procedure is stopped, thereby avoiding additional higher-order gradient computation during attack optimization.

### 4.3 Causal-Guided Objective Optimization

To improve transferability across different surrogate models, we introduce a causal token alignment constraint that explicitly guides high-contribution visual tokens toward the target representation. Let $h^{\ell}_{m,i,\text{adv}}$ and $h^{\ell}_{m,i,t}$ denote the intermediate-layer token representations of the adversarial and target images from the $m$-th surrogate model, respectively. The causal loss can be presented as:

$$\mathcal{L}^{m}_{\text{cas}} = 1 - \frac{1}{|\mathcal{S}_m|} \sum_{i\in\mathcal{S}_m} \frac{(h^{\ell}_{m,i,\text{adv}})^{\top} h^{\ell}_{m,i,t}}{\|h^{\ell}_{m,i,\text{adv}}\|_2 \|h^{\ell}_{m,i,t}\|_2}, \qquad |\mathcal{S}_m| \neq 0. \tag{17}$$

Given $M$ surrogate models $\{f_m\}_{m=1}^{M}$, the losses from multiple surrogate models are formulated as

$$\bar{\mathcal{L}}_{\text{cas}} = \frac{1}{M} \sum_{m=1}^{M} \mathcal{L}^{m}_{\text{cas}}. \tag{18}$$

Meanwhile, following (Jia et al., 2026), we employ the same dynamic weighting strategy for aggregating the basic target losses from multiple surrogate models (see A.5 for details).

$$\bar{\mathcal{L}}_{\text{bas}} = \sum_{m=1}^{M} \omega_m \mathcal{L}^{m}_{bas}, \qquad \sum_{m=1}^{M} \omega_m = M. \tag{19}$$

The final attack objective combines the basic target optimization and causal token alignment:

$$\mathcal{L}_{\text{attack}} = \bar{\mathcal{L}}_{\text{bas}} + \lambda \bar{\mathcal{L}}_{\text{cas}}, \tag{20}$$

where $\lambda$ balances the basic target attack objective and the causal token representation constraint.

Table 1: Transferable attack results of various methods on open-source MLLMs.

| | InternVL3-8B | | GLM-4.1V-9B | | Molmo-7B-D | | Phi-4-MM | |
|---|---|---|---|---|---|---|---|---|
| Method | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim |
| AdvDiffVLM | 0.12 | 0.140 | 0.11 | 0.218 | 0.08 | 0.193 | 0.09 | 0.200 |
| SSA-CWA | 0.01 | 0.111 | 0.00 | 0.114 | 0.00 | 0.106 | 0.00 | 0.102 |
| AnyAttack | 0.21 | 0.277 | 0.32 | 0.347 | 0.25 | 0.325 | 0.29 | 0.344 |
| M-Attack | 0.67 | 0.537 | 0.61 | 0.529 | 0.65 | 0.540 | 0.57 | 0.513 |
| FOA-Attack | 0.70 | 0.566 | 0.66 | 0.561 | 0.64 | 0.553 | 0.63 | 0.544 |
| Ours | **0.80** | **0.617** | **0.74** | **0.621** | **0.80** | **0.619** | **0.67** | **0.587** |

# 5 EXPERIMENTS

## 5.1 EXPERIMENTAL SETUP

**Datasets.** We conduct our experiments on the NIPS 2017 Adversarial Attacks and Defenses Competition dataset (K et al., 2017). Following (Jia et al., 2026), we randomly select 100 images for the evaluation. To further validate the reliability of our results on a larger test set, we additionally report results on 1,000 images in the supplementary material. To ensure that the target image is different from the source image, we randomly select them from the MSCOCO dataset (Lin et al., 2014). Meanwhile, we constrain adversarial perturbations to an $\ell_\infty$ budget of $16/255$ to limit pixel-level changes.

**Models.** Following previous works (Jia et al., 2026; Li et al., 2026b), we leverage three CLIP variants as surrogate models to generate adversarial examples, including ViT-B/16, ViT-B/32 and ViT-g-14-laion2B-s12B-b42K (ViT-g/14). For the black-box models, we consider both open-source and closed-source models. The open-source models include InternVL3-8B (Zhu et al., 2025), GLM-4.1V-9B (Hong et al., 2025), Molmo-7B-D (Deitke et al., 2025), and Phi-4-multimodal-instruct (Abouelenin et al., 2025), while the closed-source models include GPT-5 (Singh et al., 2025), GPT-5.6-Terra (OpenAI, 2026), Claude Sonnet 5 (Anthropic, 2026), and Gemini-2.5-Flash (Comanici et al., 2025).

**Comparison Methods.** We compare our method against several representative targeted transfer attack methods for MLLMs, namely AdvDiffVLM (Guo et al., 2024), SSA-CWA (Dong et al., 2023), AnyAttack (Zhang et al., 2025a), M-Attack (Li et al., 2026b), and FOA-Attack (Jia et al., 2026).

**Evaluation Metrics.** Following prior work on targeted transfer attacks against MLLMs (Li et al., 2026b), we evaluate attack performance using Attack Success Rate (ASR) and Average Similarity (AvgSim) (Jia et al., 2026). We adopt an LLM-as-a-judge framework (Li et al., 2026b; Jia et al., 2026) by feeding both adversarial and target images into the same target MLLM and using GPTScore to measure semantic similarity between the generated responses and the target response. ASR denotes the percentage of samples with GPTScore greater than or equal to 0.5, while AvgSim measures the average semantic similarity (see A.2 for details).

## 5.2 COMPARISON WITH STATE-OF-THE-ART METHODS

To evaluate the effectiveness of our proposed method, we compare it with five state-of-the-art transferable targeted attack methods. All methods are evaluated under the same black-box transfer setting on both open-source and closed-source target models.

**Transferable Performance on Open-source Models.** As shown in Table 1, our method achieves the best performance across all four open-source MLLMs in terms of ASR and AvgSim. Compared with FOA-Attack, our method improves ASR from 0.70 to 0.80, from 0.66 to 0.74, from 0.64 to 0.80, from 0.63 to 0.67 on InternVL3-8B, GLM-4.1V-9B, Molmo-7B-D, and Phi-4-MM, respectively, with corresponding AvgSim gains of 0.051, 0.060, 0.066, and 0.043. These improvements support the effectiveness of the proposed token attribution strategy in guiding adversarial optimization.

**Transferable Performance on Closed-source Models.** Table 2 shows that the performance of transferable targeted attacks varies across different closed-source MLLMs. Compared with FOA-

Table 2: Transferable attack results of various methods on closed-source MLLMs.

| | GPT-5 | | GPT-5.6-Terra | | Claude Sonnet 5 | | Gemini-2.5-Flash | |
|---|---|---|---|---|---|---|---|---|
| Method | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim |
| AdvDiffVLM | 0.09 | 0.126 | 0.02 | 0.154 | 0.06 | 0.183 | 0.06 | 0.082 |
| SSA-CWA | 0.01 | 0.017 | 0.00 | 0.086 | 0.00 | 0.064 | 0.00 | 0.006 |
| AnyAttack | 0.11 | 0.210 | 0.11 | 0.268 | 0.03 | 0.231 | 0.04 | 0.208 |
| M-Attack | 0.65 | 0.525 | 0.12 | 0.229 | 0.49 | 0.459 | 0.40 | 0.372 |
| FOA-Attack | 0.72 | 0.582 | 0.16 | 0.275 | 0.59 | 0.506 | 0.56 | 0.456 |
| Ours | **0.83** | **0.645** | **0.26** | **0.327** | **0.69** | **0.556** | **0.61** | **0.488** |

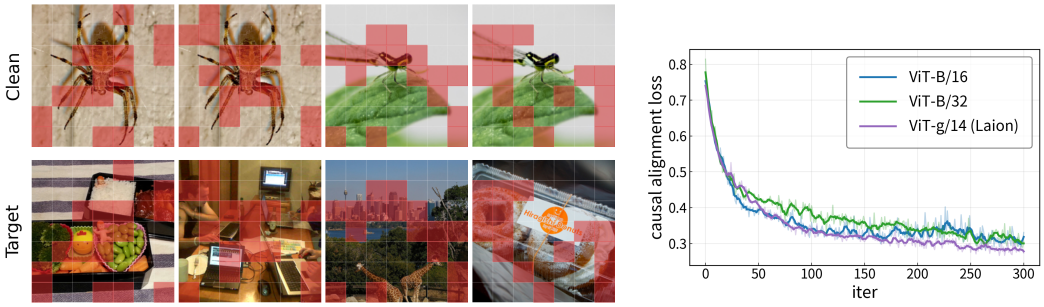


(a) Visualization of causal tokens under different target images.

(b) Evolution of causal alignment loss.

Figure 2: This figure illustrates the target dependence and alignment of causal tokens, where (a) shows the visualization of high-contribution tokens for fixed clean images paired with different targets, and (b) presents the causal alignment loss evaluated on the currently selected Top-(K) tokens in each of three surrogate models.

Attack, our method achieves the most notable improvement on GPT-5, increasing ASR from 0.72 to 0.83 and AvgSim from 0.582 to 0.645. Our method also yields clear gains on GPT-5.6-Terra and Claude Sonnet 5, improving ASR from 0.16 to 0.26 and from 0.59 to 0.69, respectively, while simultaneously increasing AvgSim from 0.275 to 0.327 and from 0.506 to 0.556. On Gemini-2.5-Flash, our method improves ASR from 0.56 to 0.61 and achieves the highest AvgSim of 0.488. These results demonstrate that our approach consistently improves targeted transferability across diverse open-source and closed-source MLLMs, highlighting its generalization capability under black-box settings.

### 5.3 Counterfactual Causal Effect Analysis

**Causal Token Identification.** We evaluate whether different token importance estimation strategies can effectively identify visual tokens that contribute to transferability. We compare random selection, attention-based, activation-based, and gradient-based strategies. As shown in Table 3, our proposed counterfactual causal strategy consistently achieves the best performance, improving both ASR and AvgSim over other selection approaches. These results demonstrate that causal intervention provides a more accurate estimation of target-relevant visual tokens, leading to more effective targeted transfer attacks.

Table 3: The performance of different token selection strategies.

| Metric | Random | Attention | Activation | Gradient Norm | Ours |
|---|---|---|---|---|---|
| ASR ↑ | 0.74 | 0.68 | 0.71 | 0.73 | **0.83** |
| AvgSim ↑ | 0.618 | 0.556 | 0.582 | 0.621 | **0.645** |

**Visualization of Causal Tokens under Different Targets.** For each clean image, we vary the target image and estimate the causal contributions of visual tokens. As shown in Figure 2a, the spatial distribution of high-contribution tokens changes across targets, even when the clean image

remains fixed. This observation illustrates the target dependence of causal token importance and motivates selecting tokens according to their contributions to the specific target representation.

**Causal Token Alignment Analysis.** We fix the same clean image and select multiple target images to analyze the target dependency of causal tokens. Figure 2b shows that the causal alignment loss consistently decreases across different surrogate models, demonstrating that our strategy effectively guides optimization toward target-relevant representations.

### 5.4 Ablation Studies

**Dynamic Causal Token Selection.** To investigate the effectiveness of dynamically updating causal token contributions during optimization, we conduct an ablation study on different token update strategies using GPT-5 as the black-box target model. The same target model is used in all subsequent ablation studies, including those in the appendix. Specifically, we compare our dynamic update strategy with a static selection strategy, where the important tokens are identified only once and kept fixed throughout the optimization process. For dynamic updating, we introduce an update interval $R$, which denotes the number of optimization iterations between two consecutive causal token re-estimation steps. We evaluate different update intervals, with $R$ set to 10, 5, and 1, where a smaller $R$ indicates more frequent causal token updates. As shown in Table 4, the static selection strategy achieves an ASR of 71% and an AvgSim of 0.571, while dynamically updating causal tokens consistently improves the attack performance. In particular, reducing the update interval from 10 to 1 gradually increases both ASR and AvgSim, with $R = 1$ achieving the best performance. This indicates that continuously adapting the selected tokens according to their updated causal contributions provides more accurate optimization guidance and better aligns adversarial examples with the target representation.

**Causal Loss Weight.** The final attack objective combines the basic target representation alignment with the causal token alignment term, where the hyperparameter $\lambda$ balances the contribution of causal token alignment during optimization. To investigate the effect of this parameter on attack performance, we vary $\lambda$ while keeping other settings unchanged and compare the results. As shown in Table 5, when $\lambda$ increases from 0 to 5, both ASR and AvgSim consistently improve, achieving the best performance at $\lambda = 5$ with an ASR of 0.83 and an AvgSim of 0.645. However, further increasing $\lambda$ to 10 leads to a performance drop, indicating that excessively emphasizing causal token alignment may weaken the contribution of the overall target representation alignment. These results demonstrate that a proper balance between target representation and causal token alignment leads to more effective target-oriented optimization and improved transferability.

Table 4: Ablation study of different causal token update strategies.

| Strategy | ASR ↑ | AvgSim ↑ |
|---|---|---|
| Static Selection | 0.71 | 0.571 |
| $R = 10$ | 0.76 | 0.610 |
| $R = 5$ | 0.79 | 0.620 |
| $R = 1$ | 0.83 | 0.645 |

Table 5: Ablation study of the causal alignment weight $\lambda$.

| $\lambda$ | ASR ↑ | AvgSim ↑ |
|---|---|---|
| 0 | 0.65 | 0.525 |
| 1 | 0.76 | 0.600 |
| 5 | 0.83 | 0.645 |
| 10 | 0.75 | 0.591 |

## 6 Conclusion

In this work, we propose a causal-guided framework for transferable targeted attacks on MLLMs. Through target-oriented counterfactual interventions, we estimate the contributions of individual visual tokens to target representation alignment. We further derive a gradient-based approximation to efficiently estimate these contributions without explicitly evaluating each token intervention. Based on the estimated causal scores, our method dynamically selects high-contribution tokens and aligns them with their target counterparts, alongside a global feature alignment objective. Experiments on diverse open- and closed-source MLLMs demonstrate improved targeted transferability over existing methods, while ablation studies support the effectiveness of causal token selection and dynamic updates.

## AI USE STATEMENT

We used AI tools to assist with language editing, grammar checking, and checks for consistency in wording and reported numerical values. The authors reviewed all AI-assisted suggestions and take responsibility for the accuracy, integrity, and final content of this paper.

## REFERENCES


Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, et al. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*, 2025.

Anthropic. Introducing claude sonnet 5. Anthropic, June 2026. URL `https://www.anthropic.com/news/claude-sonnet-5`.

Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.

Chun-Hao Chang, Elliot Creager, Anna Goldenberg, , and David Duvenaud. Explaining image classifiers by adaptive dropout and generative in-filling. *arXiv preprint arXiv:1807.08024*, 2018.

Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 782–791. IEEE, 2021.

Huanran Chen, Yichi Zhang, Yinpeng Dong, Xiao Yang, Hang Su, and Jun Zhu. Rethinking model ensemble in transfer-based adversarial attacks. In *International Conference on Learning Representations*, volume 2024, pp. 19199–19225, 2024.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

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 *2024 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pp. 958–979. IEEE, 2024.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 91–104. IEEE, 2025.

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.

Tingchao Fu, Jinhong Zhang, Fanxiao Li, Ping Wei, Xianglong Zeng, and Wei Zhou. Multimodal alignment augmentation transferable attack on vision-language pre-training models. *Pattern Recognition Letters*, 191:131–137, 2025.

Mengyu Gao and Qiulei Dong. Causality-guided prompt learning for vision-language models via visual granulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1141–1151, 2025.

Yash Goyal, Ziyan Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee. Counterfactual visual explanations. In *International Conference on Machine Learning*, pp. 2376–2384. PMLR, 2019.

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*, 20:1333–1348, 2024.

Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, et al. Glm-4.5 v and glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*, 2025.

Kai Hu, Weichen Yu, Li Zhang, Alexander Robey, Andy Zou, Chengming Xu, Haoqi Hu, and Matt Fredrikson. Transferable adversarial attacks on black-box vision-language models. *arXiv preprint arXiv:2505.01050*, 2025.

Hanxun Huang, Sarah Erfani, Yige Li, Xingjun Ma, and James Bailey. X-transfer attacks: Towards super transferable adversarial attacks on clip. In *International Conference on Machine Learning (ICML)*, 2025.

Xiaojun Jia, Sensen Gao, Simeng Qin, Tianyu Pang, Chao Du, Yihao Huang, Xinfeng Li, Yiming Li, Bo Li, and Yang Liu. Adversarial attacks against closed-source mllms via feature optimal alignment. *Advances in Neural Information Processing Systems*, 38:57993–58028, 2026.

Yutao Jiang, Qiong Wu, Wenhao Lin, Wei Yu, and Yiyi Zhou. What kind of visual tokens do we need? training-free visual token pruning for multi-modal large language models from the perspective of graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 4075–4083, 2025.

Alex K, Ben Hamner, and Ian Goodfellow. Nips 2017: Defense against adversarial attack. `https://www.kaggle.com/competitions/nips-2017-defense-against-adversarial-attack`, 2017. Kaggle.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.

Li Li, Jiashu Qu, Linxin Song, Yuxiao Zhou, Yuehan Qin, Tiankai Yang, and Yue Zhao. Treble counterfactual vlms: A causal approach to hallucination. Association for Computational Linguistics, 2025.

Yang Li, Jia-Li Yin, Luojun Lin, and Wei Lin. Transform to transfer: Boosting adversarial attack transferability on vision-language pre-training models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 30141–30150, 2026a.

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. *Advances in Neural Information Processing Systems*, 38:45018–45051, 2026b.

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 *European conference on computer vision*, pp. 740–755. Springer, 2014.

Chengxin Liu, Wonseok Choi, Chenshuang Zhang, and Tae-Hyun Oh. Aligning what vision-language models see and perceive with adaptive information flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24706–24715, June 2026.

Yunqi Liu, Xue Ouyang, and Xiaohui Cui. Gleam: Enhanced transferable adversarial attacks for vision-language pre-training models via global-local transformations. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1665–1674. IEEE, 2025.

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*, pp. 549–566. Springer, 2022.

OpenAI. Gpt-5.6 terra. OpenAI, July 2026. URL `https://openai.com/index/gpt-5-6/`.

Judea Pearl. *Causality*. Cambridge university press, 2009.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.

Jonas Peters, Dominik Janzing, and Bernhard Scholkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.

Jiaxin Qi, Yulei Niu, Jianqiang Huang, and Hanwang Zhang. Two causal principles for improving visual dialog. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10860–10869, 2020.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*, 2025.

Nengbo Wang, Xiaotian Han, Jagdip Singh, Jing Ma, and Vipin Chaudhary. Causalrag: Integrating causal graphs into retrieval-augmented generation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 22680–22693, 2025a.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.

Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense r-cnn. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10760–10770, 2020.

Yubo Wang, Chaohu Liu, Yanqiu Qu, Haoyu Cao, Deqiang Jiang, and Linli Xu. Break the visual perception: Adversarial attacks targeting encoded visual tokens of large vision-language models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 1072–1081, 2024b.

Zhibo Wang, Hengchang Guo, Zhifei Zhang, Wenxin Liu, Zhan Qin, and Kui Ren. Feature importance-aware transferable adversarial attacks. In *2021 IEEE/CVF international conference on computer vision (ICCV)*, pp. 7619–7628. IEEE, 2021.

Zhuoran Wang, Jiawei Zhao, Liang Zhang, Zheng Wang, and Yining Hu. Targeted and transferable adversarial attacks on large vision-language models via two-stage region-specific perturbations. In *2025 IEEE 24th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pp. 282–289. IEEE, 2025b.

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.

Peng Xie, Yequan Bie, Jianda Mao, Yangqiu Song, Yang Wang, Hao Chen, and Kani Chen. Chain of attack: On the robustness of vision-language models against transfer-based adversarial attacks. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14679–14689. IEEE, 2025.

Hai Yan, Haijian Ma, Xiaowen Cai, Daizong Liu, Zenghui Yuan, Xiaoye Qu, Jianfeng Dong, Runwei Guan, Xiang Fang, Hongyang He, et al. Fit the distribution: Cross-image/prompt adversarial attacks on multimodal large language models. *Advances in Neural Information Processing Systems*, 38:75204–75247, 2026.

Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9593–9602, 2021.

Jiaming Zhang, Junhong Ye, Xingjun Ma, Yige Li, Yunfan Yang, Yunhao Chen, Jitao Sang, and Dit-Yan Yeung. Anyattack: Towards large-scale self-supervised adversarial attacks on vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 19900–19909, 2025a.

Jusheng Zhang, Kaitong Cai, Yijia Fan, Jian Wang, and Keze Wang. Cf-vlm: counterfactual vision-language fine-tuning. *arXiv preprint arXiv:2506.17267*, 2025b.

Qizhe Zhang, Aosong Cheng, Ming Lu, Renrui Zhang, Zhiyong Zhuo, Jiajun Cao, Shaobo Guo, Qi She, and Shanghang Zhang. Beyond text-visual attention: Exploiting visual cues for effective token pruning in vlms. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 20857–20867. IEEE, 2025c.

Yunqing Zhao, Tianyu Pang, Chao Du, Xiao Yang, Chongxuan Li, Ngai-Man Man Cheung, and Min Lin. On evaluating adversarial robustness of large vision-language models. *Advances in Neural Information Processing Systems*, 36:54111–54138, 2023.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

# A Appendix

## A.1 Results on Large-scale Dataset

To further evaluate our method on a larger test set, we compare it with FOA-Attack using 1,000 test samples across four closed-source MLLMs. As shown in Table 6, our method consistently achieves higher ASR and AvgSim on all evaluated models. On GPT-5, ASR increases from 0.682 to 0.757, while AvgSim improves from 0.5526 to 0.6036. Similar improvements are observed on the other models, with ASR increasing from 0.179 to 0.230 on GPT-5.6-Terra, from 0.574 to 0.633 on Claude Sonnet 5, and from 0.488 to 0.517 on Gemini 2.5 Flash. These results show that the performance advantage over FOA-Attack persists on the larger test set, demonstrating the effectiveness of causal-guided token alignment for targeted adversarial transfer.

Table 6: Transferable attack results of different methods on closed-source MLLMs using 1,000 test samples. The best results are highlighted in bold.

| | GPT-5 | | GPT-5.6-Terra | | Claude Sonnet 5 | | Gemini 2.5 Flash | |
|---|---|---|---|---|---|---|---|---|
| Method | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim | ASR | AvgSim |
| FOA-Attack | 0.682 | 0.5526 | 0.179 | 0.2169 | 0.574 | 0.5058 | 0.488 | 0.4236 |
| Ours | **0.757** | **0.6036** | **0.230** | **0.2509** | **0.633** | **0.5326** | **0.517** | **0.4545** |

## A.2 Details of Evaluation

**Evaluation protocol.** We evaluate targeted transfer using an LLM-based semantic similarity score. Specifically, for each image pair, the evaluated MLLM separately generates descriptions of the adversarial image and the corresponding target image using the same captioning prompt: *"Describe this image in one concise sentence, no longer than 20 words."* The resulting descriptions are then provided to GPT-4o, which directly assigns a similarity score between 0 and 1. The judge receives only the two descriptions, without access to the images.

**Scoring prompt.** Figure 3 presents the complete scoring prompt, where `text1` and `text2` are replaced with the target-image and adversarial-image descriptions, respectively. The scoring rubric emphasizes agreement in the main subject and context while allowing differences in wording and minor details. An attack is considered successful when the similarity score is greater than or equal to 0.5. We report ASR and AvgSim as defined in the main text.

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

**Criteria for similarity measurement:**

1. **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.
2. **Relevant Description**: If the descriptions are related to the same context or topic, they should also contribute to a higher similarity score.
3. **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.**
4. **Partial Matches:** If one description contains extra information but does not contradict the other, they should still have a high similarity score.
5. **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: {text1}

Text 2: {text2}

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

Figure 3: The details of the scoring prompt.

### A.3 Additional Results

**Computational Efficiency.** We also present the computational costs of counterfactual computation and the proposed gradient approximation method, with the results reported in Table 7. The vanilla counterfactual approach requires constructing a separate intervention state for each visual token and performing an additional forward pass after the intervention layer, resulting in substantial computational overhead. In contrast, our method requires only a single backward propagation from the current adversarial example to obtain gradients for all tokens, followed by lightweight vector operations to estimate their causal scores. As shown in Table 7, the proposed gradient approximation reduces the scoring time from 5.60 s to 0.05 s and decreases GPU memory consumption from 934 MB to 27 MB, achieving a $112\times$ speedup over the vanilla counterfactual computation. These results demonstrate that our approximation provides an efficient and practical solution for causal token estimation during adversarial optimization.

Table 7: Computational efficiency comparison between explicit token Intervention and the proposed gradient-based method.

| Method | Scoring Time (s) ↓ | GPU Memory (MB) ↓ | Speedup ↑ |
|---|---|---|---|
| Explicit token Intervention | 5.60 | 934 | $1.0\times$ |
| Gradient Approximation (Ours) | 0.05 | 27 | $112.0\times$ |

### A.4 Additional Ablation Studies

**Token Ratio.** Our method selects only visual tokens with high causal contributions for target representation optimization. Therefore, the number of selected tokens $K$ is an important factor affecting attack performance. To analyze the influence of different selection scales, we vary the proportion of tokens involved in optimization while keeping the remaining attack settings unchanged, with results presented in Table 8. When too few tokens are selected, they may provide insufficient target semantic information. In contrast, selecting too many tokens may introduce tokens with weak causal contributions or even negative effects on target transfer, thereby weakening the advantage of selective optimization. Therefore, we evaluate different token selection ratios and compare the results to analyze the sensitivity of our method to $K$.

Table 8: Ablation results of different token selection ratios.

| Token Ratio | ASR ↑ | AvgSim ↑ |
|---|---|---|
| 0 | 0.65 | 0.525 |
| $1/16$ | 0.75 | 0.605 |
| $1/4$ | 0.83 | 0.645 |
| $9/16$ | 0.76 | 0.614 |
| 1 | 0.67 | 0.558 |

**Intervention Layer Position.** Our method performs target-oriented token intervention in the intermediate representation space of the image encoder. Therefore, the intervention layer position may affect the estimated intervention effects. Different layers capture distinct representation characteristics, with shallow layers focusing on low-level visual patterns and deeper layers encoding more semantic information. To investigate the impact of intervention depth on attack performance, we perform causal token analysis at multiple layers with different relative depths. The results in Table 9 show that the intervention depth has a significant impact on attack performance. The layer at $0.35L$ achieves the best results with an ASR of 0.83 and an AvgSim of 0.645. Both shallower and deeper layers lead to performance degradation, suggesting that effective causal token identification requires a proper balance between visual details and semantic information.

Table 9: Ablation results of different intervention layer positions.

| Relative Depth | ASR ↑ | AvgSim ↑ |
|---|---|---|
| $0.20L$ | 0.73 | 0.573 |
| $0.35L$ | 0.83 | 0.645 |
| $0.50L$ | 0.78 | 0.625 |
| $0.65L$ | 0.75 | 0.595 |
| $0.80L$ | 0.65 | 0.537 |

### A.5 Dynamic Ensemble Weighting

For image-level feature alignment, we adopt the dynamic ensemble weighting strategy of (Jia et al., 2026) to account for differences in optimization progress across surrogate models. Let $\mathcal{L}_m^{\mathrm{bas},(k)}$ denote the image-level alignment loss of the $m$-th surrogate model at iteration $k$. Its relative change between consecutive iterations is measured by

$$r_m^{(k)} = \frac{\mathcal{L}_m^{\mathrm{bas},(k)}}{\mathcal{L}_m^{\mathrm{bas},(k-1)}}, \qquad k \geq 1. \tag{21}$$

A larger ratio indicates less relative progress in reducing the loss. The model weights are therefore computed as

$$\omega_m^{(k)} = M \, \frac{\exp\left(r_m^{(k)}/\tau\right)}{\sum_{j=1}^{M} \exp\left(r_j^{(k)}/\tau\right)}, \tag{22}$$

where $\tau > 0$ is a temperature parameter controlling the differences between model weights. This rule assigns larger weights to models with slower relative loss reduction. The factor $M$ ensures that the average model weight remains one. At initialization, we set $r_m^{(0)} = 1$, yielding $\omega_m^{(0)} = 1$ for all models.

The resulting image-level alignment objective is

$$\mathcal{L}_{\text{bas}}^{(k)} = \sum_{m=1}^{M} \omega_m^{(k)} \mathcal{L}_m^{\text{bas},(k)}. \tag{23}$$

This weighting is applied to the image-level alignment term, while the causal token alignment losses are averaged uniformly across surrogate models, as defined in Eq. 18.

### A.6 Visualization Results

Figure 4 presents clean images, their corresponding target images, and adversarial examples generated by different methods, together with the associated perturbations. In the displayed examples, the adversarial images produced by our method preserve the overall appearance and recognizable content of the clean images. The perturbation visualizations further illustrate differences in the spatial patterns of the modifications introduced by each method. These examples provide a qualitative comparison of visual fidelity, complementing the quantitative attack evaluation.

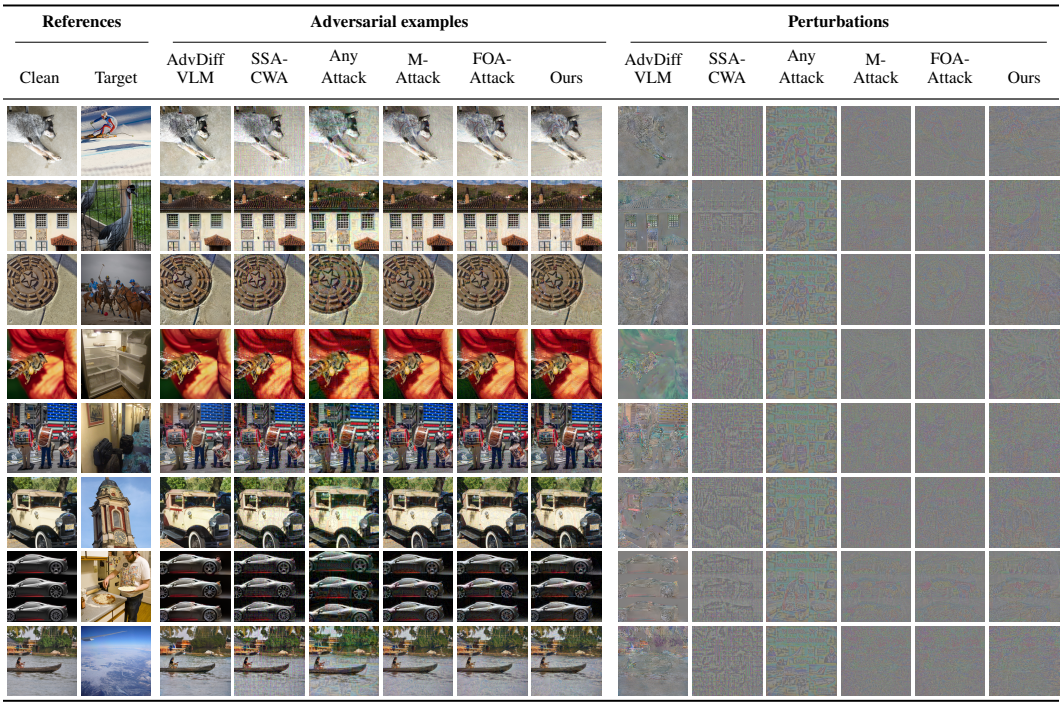


Figure 4: Visual comparison of adversarial examples and perturbations. Each row corresponds to one sample. Clean and target images are followed by adversarial examples in the left method block and corresponding perturbations in the right method block. Both blocks use the same method order. Our method is indicated by bold column headers.

### A.7 Qualitative Examples of Targeted Transfer

We provide additional qualitative examples to complement the quantitative evaluation of our method on closed-source MLLMs. As shown in Figure 5, adversarial images generated by our method can induce responses aligned with the intended target semantics across different models. Although the responses vary in wording and descriptive details, they exhibit target-related semantic content. These examples provide a qualitative illustration of targeted transfer and help characterize model behavior under the proposed attack.

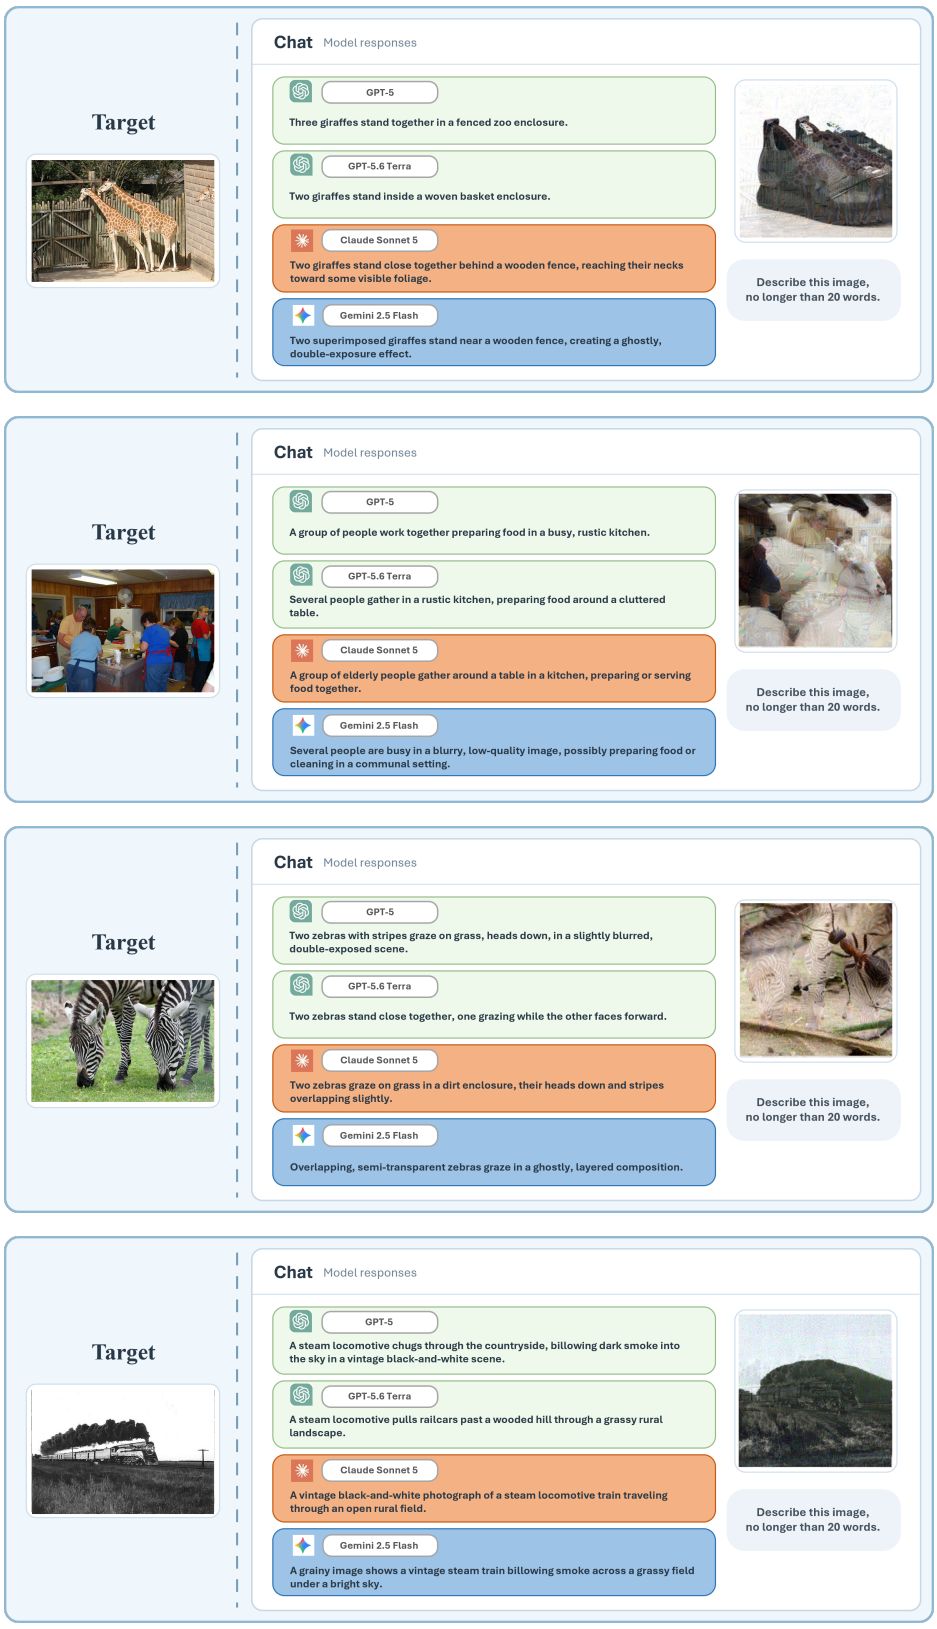


Figure 5: Qualitative example of targeted transfer to closed-source MLLMs.