

AdvDreamer Unveils: Are Vision-Language Models Truly Ready for Real-World 3D Variations?

Shouwei Ruan^{1†}, Hanqing Liu^{1†}, Yao Huang^{1†}, Xiaoqi Wang¹, Caixin Kang¹,
 Hang Su², Yinpeng Dong³, Xingxing Wei^{1*}

¹Institute of Artificial Intelligence,
 State Key Laboratory of Virtual Reality Technology and Systems, Beihang University
²Dept. of Comp. Sci. and Tech., Institute for AI, Tsinghua-Bosch Joint ML Center,
 THBI Lab, BNRist Center, Tsinghua University
³College of AI, Tsinghua University

{shouweiruan, hqliu, y_huang, xxwei}@buaa.edu.cn, suhangss@tsinghua.edu.cn, dongyinpeng@mail.tsinghua.edu.cn

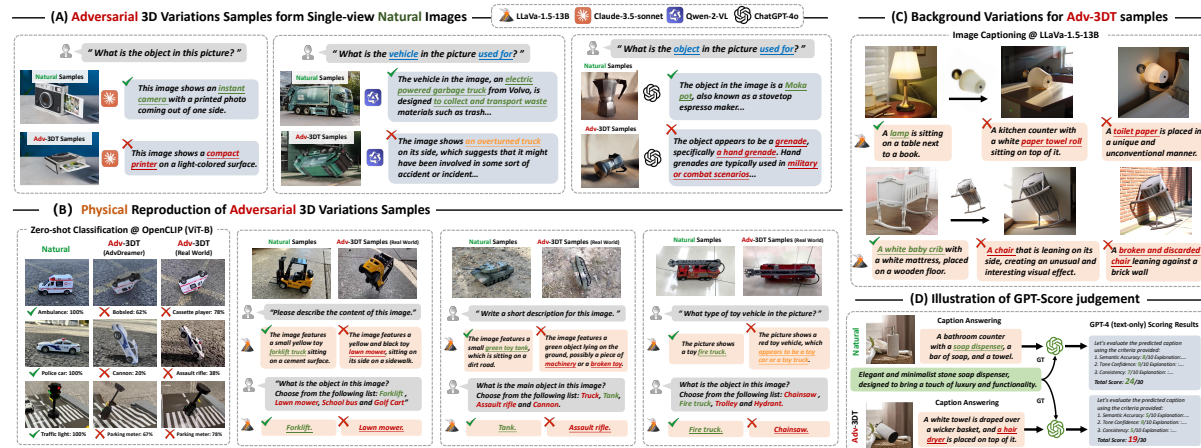


Figure 1. **The 3D Variation Vulnerabilities in VLMs.** (A) The proposed AdvDreamer allow the capture of adversarial 3D transformation (Adv-3DT) samples from single natural images, threatening various VLMs. (B) These Adv-3DT samples can be successfully reproduced in the physical world and (C) retain their aggressiveness under background variations. (D) We also explain the GPT-Score metric used in the experiments. In each VLM response, we highlight the **correct**, **significant error**, and **less precise** parts using different colors.

Abstract

Vision Language Models (VLMs) have exhibited remarkable generalization capabilities, yet their robustness in dynamic real-world scenarios remains largely unexplored. To systematically evaluate VLMs' robustness to real-world 3D variations, we propose **AdvDreamer**, the first framework capable of generating physically reproducible Adversarial 3D Transformation (Adv-3DT) samples from single-view observations. In AdvDreamer, we integrate three key innovations: Firstly, to characterize real-world 3D variations with limited prior knowledge precisely, we design a zero-shot **Monocular Pose Manipulation** pipeline built upon generative 3D priors. Secondly, to ensure the visual qual-

ity of worst-case Adv-3DT samples, we propose a **Naturalness Reward Model** that provides continuous naturalness regularization during adversarial optimization, effectively preventing convergence to hallucinated or unnatural elements. Thirdly, to enable systematic evaluation across diverse VLM architectures and visual-language tasks, we introduce the **Inverse Semantic Probability** loss as the adversarial optimization objective, which solely operates in the fundamental visual-textual alignment space. Based on the captured Adv-3DT samples with high aggressiveness and transferability, we establish **MM3DTBench**, the first VQA benchmark dataset tailored to evaluate VLM robustness under challenging 3D variations. Extensive evaluations of representative VLMs with varying architectures reveal that real-world 3D variations can pose severe threats to model performance across various tasks.

† Equal contribution.
 * Corresponding author.




Figure 2. Comparison of paradigms to characterize Adversarial 3D Transformation (Adv-3DT) samples.

1. Introduction

Recently, Vision-Language Models (VLMs) [7, 13, 35, 36] have demonstrated remarkable capabilities in bridging visual perception and natural language understanding. Through large-scale image-text pre-training followed by fine-grained instruction tuning [31, 36], these models effectively tackle a wide range of visual-centric tasks, including instruction-based image recognition [47, 55, 65], visual question answering [3, 31, 35, 36], and visual reasoning [9, 10]. Due to their strong generalizability and impressive zero-shot capability, VLMs are increasingly deployed in various safety-critical applications, particularly in autonomous driving [56, 63] and robotic systems [25, 64].

Despite extensive studies validating VLMs’ resilience against various distribution shifts (*e.g.*, style variations [2, 39, 40, 68], image corruptions [2], and adversarial perturbations [16, 38, 67, 70]), they primarily focus on 2D perturbations in the digital domain, overlooking a critical challenge in real-world deployment: the 3D variations. As VLMs are increasingly integrated into dynamic scenarios, their ability to handle 3D variation becomes crucial, raising a fundamental question: *Have current VLMs attained sufficient robustness to the distribution shifts arising from ubiquitous 3D variations in the real world?*

To thoroughly explore this problem and offer comprehensive answers, we identify three fundamental challenges:

(1) *How to accurately characterize real-world 3D variations with limited prior knowledge?* Existing methods exploring the 3D variation robustness heavily rely on explicit 3D structures [4, 21] or representations derived from dense multi-view observations [4, 21, 42]. However, in most practical scenarios, the available priors are severely limited (*e.g.*, single-view observation only). This constraint necessitates developing a method that minimizes dependency on extensive scene priors. To meet this, we introduce a zero-shot **Monocular Pose-Manipulation (MPM)** pipeline based on generative representations. Specifically, MPM leverages the rich 3D priors embedded in pre-trained Large-Reconstructive Models (LRM) [23, 24] and Image Re-Composition Models [12, 66], enabling fine-grained 3D adjustments with only single-view images and specified 3D variation parameters. While this generative-based paradigm establishes a foundation for capturing worst-case 3D variations through optimizing 3D variation parameters, it still

exhibits inherent limitations that may introduce undesirable image corruptions, thus presenting our second challenge:

(2) *How to capture the worst-case adversarial 3D transformation (Adv-3DT) samples while ensuring that the performance degradation is inherently caused by the 3D variation?* Prior works utilizing synthetic 3D assets [4, 21, 42] frequently produce samples with noticeable texture discrepancies compared to real-world images. These limitations in samples’ visual fidelity and contextual coherence highlight a critical concern: observed performance degradation under Adv-3DT samples may stem from undesirable corruption rather than the 3D variation itself. To address this, we incorporate naturalness reward constraints within the adversarial optimization process. Specifically, we treat rigid 3D variations as an adversarial attack and propose the **Naturalness Reward Model (NRM)** that preserves image quality throughout optimization. NRM leverages visual context features provided by the DINOv2 [46] backbone to predict a naturalness reward that jointly reflects the visual fidelity and physical plausibility of the samples, effectively regularizing the optimization trajectory and preventing convergence toward unnatural or hallucinated elements.

(3) *How to ensure effective generalization of Adv-3DT samples across diverse VLM tasks and architectures?* Existing studies [4, 18, 50] typically rely on task-specific objectives with end-to-end optimization, substantially limiting their applicability across diverse VLM architectures and a wide range of visual language tasks. To transcend this limitation, we formulate the adversarial objective in a task- and architecture-agnostic manner. Specifically, we introduce the **Inverse Semantic Probability (ISP)** loss, which operates solely through the foundational component of the VLM—the visual encoder, aiming to minimize the image-text matching probability between Adv-3DT samples and their corresponding ground-truth textual descriptions. Extensive experiments demonstrate that Adv-3DT samples optimized under the ISP loss exhibit remarkable adversarial transferability across various downstream tasks and models. Building on this, we further introduce the **MM3DTBench**, the first VQA benchmark dataset specifically designed to evaluate the 3D variation robustness of VLMs in real-world scenarios, comprising highly challenging Adv-3DT samples with carefully crafted candidate answers.

Integrating the aforementioned innovations, this paper

proposes **AdvDreamer**, a novel framework that enables the characterization of Adv-3DT samples from monocular observations, with broad adaptability across various visual-language tasks and VLM architectures. Our extensive digital and physical experiments reveal significant vulnerabilities in contemporary VLMs when confronted with challenging 3D variations. In visual recognition tasks, foundational models like OpenCLIP [26] and BLIP-2 [31] suffer a performance degradation of **65%~80%** on Adv-3DT samples. More concerning still, in open-ended tasks including image captioning and visual question answering, even the most advanced commercial models like GPT-4o [45] exhibit performance declines approaching **50%**. Our findings highlight the urgent need to enhance VLM robustness against 3D variations and set new requirements for deploying VLMs in complex, dynamic real-world applications.

2. Related Work

2.1. Vision-Language Models

Classic Vision-Language Models (VLMs) [30, 32, 41, 47] achieve cross-modal alignment through diverse pre-training objectives on large-scale image-text pairs [33, 53]. Recent advances in Large Language Models (LLMs) [59, 60] have inspired new VLM architectures that integrate LLM capabilities via MLP projection layers [35, 36] or specialized modules like Q-Former [31]. Through techniques like instruction tuning [36], current VLMs [13, 35, 36] demonstrate enhanced performance across vision-centric tasks, including visual grounding, question answering, and image captioning, *etc.* These models also enable intelligent systems for task planning and decision-making [25, 64].

Despite their impressive capabilities [37, 47, 67], concerns remain about the robustness and generalization of VLMs, especially in safety-critical environments. Recent studies [52, 58] highlight that VLMs still exhibit deficiencies in processing visual details, such as orientation, quantity, and colour, which may be attributed to their visual representations. Moreover, VLMs remain susceptible to L_p -norm adversarial perturbations [16, 43], with some efforts aiming to enhance their robustness through adversarial fine-tuning techniques [43]. In contrast to previous works, our study specifically investigates VLM robustness against 3D variations, addressing a crucial gap in understanding their reliability for dynamic real-world applications.

2.2. Evaluation for 3D Variation Robustness

Robustness to 3D variations remains a longstanding challenge in computer vision. Evaluating this property often involves capturing worst-case 3D transformations, *i.e.* the adversarial poses or viewpoints, to create challenging samples for assessment. Prior research can be categorized into three paradigms (see Fig.2): (A) **Human-guided** methods, like

ObjectNet [8] and OOD-CV [69], manually collect samples with diverse poses and viewpoints, but they lack systematic coverage of worst-case examples and are costly. (B) **Simulator-based** [4, 21, 42] methods optimize worst-case 3D parameters via differentiable rendering of synthetic objects but require manually designed 3D assets and produce samples with limited realism. (C) **Multi-view-based** methods, like ViewFool [18] and GMVFool [50, 51], use NeRF-based [44] representations to optimize adversarial viewpoints, but they require dense observations and struggle with background integration. In contrast, the proposed AdvDreamer employs a (D) **Generative-based** approach, utilizing robust 3D priors from generative models. Given a single natural image, AdvDreamer generate realistic Adv-3DT samples in a zero-shot manner, providing a more efficient solution for evaluating 3D variation robustness.

3. Methodology

The overview of AdvDreamer is shown in Fig. 3. It consists ① Zero-shot Monocular Pose Manipulation pipeline, leveraging generative 3D representations to obtain specified 3D variation samples from single-view images; ② Naturalness Reward Model, providing naturalness regularization during adversarial optimization; and ③ Inverse Semantic Probability loss, ensuring the Adv-3DT samples effectively threaten VLMs and generalize across different architectures and tasks. The following sections will firstly formalize the optimization problem of AdvDreamer (Sec.3.1), then sequentially describe the key components (Sec.3.2), and finally outline the full optimization algorithm (Sec.3.3).

3.1. Problem Formulation

Parametrization of Real-World 3D Variations. In this paper, we primarily focus on rigid 3D transformations of foreground objects, *i.e.*, the rotation, translation, and scaling, since they reflect the most typical 3D changes observed in real-world environments. Formally, we define a 6-dimensional vector $\Theta = \{\alpha, \beta, \gamma, \Delta_x, \Delta_y, s\}$ to uniquely parameterize any arbitrary transformation, where $\{\alpha, \beta, \gamma\}$ denote the Tait-Bryan angles (yaw, pitch, roll) in the Z - Y - X sequence, $\{\Delta_x, \Delta_y\}$ represent the translation along the x and y axes in the xy -plane, and s is the uniform scaling factor. To prevent the loss of critical visual cues and ensure human recognizability, we follow previous work [18] to constrain Θ within a bounded range $[\Theta_{\min}, \Theta_{\max}]$.

Optimization Problem. The objective of AdvDreamer is to find an optimal distribution $p^*(\Theta)$, ensuring that Adv-3DT sampled from this distribution are both aggressive against VLMs and maintaining visual quality, which can be abstractly formulated as the following optimization problem:

$$p^*(\Theta) = \arg \max_{p(\Theta)} \mathbb{E}_{\Theta \sim p(\Theta)} [\mathcal{L}_{ISP}(X') + \mathcal{L}_{Nat}(X')], \quad (1)$$

where $X' = \mathcal{T}(\Theta, X)$,

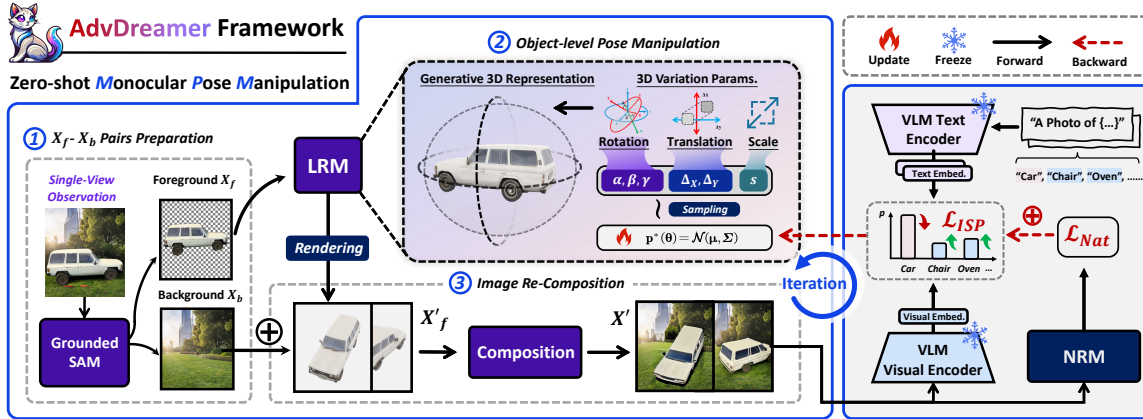


Figure 3. **Overview of the proposed AdvDreamer framework.** To capture worst-case 3D variations in the real world, we treat rigid 3D transformations as adversarial attacks, optimizing the distribution of transformation parameters through a query-based approach. In AdvDreamer, we introduce the Monocular Pose Manipulation pipeline to perform specified 3D transformations on single-view images and guide the optimization process using the proposed Naturalness Reward Model and Inverse Semantic Probability loss.

where $\mathcal{T}(\Theta, X)$ is a transformation function to produce a novel image with the specified transformation Θ applied to the original image $X \in \mathbb{R}^3$. To implement $\mathcal{T}(\cdot)$, we design the MPM pipeline, which will be detailed in Sec.3.2.1. The \mathcal{L}_{Nat} and \mathcal{L}_{ISP} measure the naturalness and adversarial efficacy of the transformed samples, respectively. These will be further elaborated in Sec.3.2.2 and Sec.3.2.3.

Optimizing over the distribution $p(\Theta)$ rather than a deterministic Θ offers two key advantages. 1) As highlighted in [18, 50], learning the underlying distribution enables comprehensive exploration of the 3D variation space, facilitating the identification of the model’s vulnerable regions. 2) Given the stochastic nature of the generation, optimizing over a continuous distribution can mitigate the impact of appearance inconsistency, as appearance variations induced by the generation are unlikely to consistently deceive the model without optimizing the adversarial 3D variations [6].

3.2. AdvDreamer Framework

3.2.1. Zero-shot Monocular Pose Manipulation

To implement the transformation function \mathcal{T} , previous approaches typically rely on explicit 3D structures [4, 21, 42] or neural radiance representations [18, 50]. However, these approaches cannot accommodate real-world scenarios with limited prior knowledge. To address this, we design a simple yet effective zero-shot monocular pose manipulation (MPM) pipeline that consists of the following process:

1) Foreground-Background Pairs Preparation. MPM first decompose the input image X into foreground-background pairs $\{X_f, X_b\}$. Specifically, we leverage Grounded-SAM [48] to semantically segment the major instance as foreground X_f , followed by diffusion-based inpainting [15, 49] to obtain a complete background X_b . This process is denoted as $\{X_f, X_b\} = \mathcal{F}(X)$. Additionally, we support directly generating synthetic $\{X_f, X_b\}$ pairs

through Stable-Diffusion [49] to enhance sample diversity.

2) Object-level Pose Manipulation. For each optimization iteration, we transform the foreground X_f into a batch of X'_f under Θ sampled from current $p(\Theta)$. In this stage, MPM leverages the robust priors provided by a pre-trained Large Reconstruction Model (e.g., TripoSR [57]) to construct single-view-based 3D representations and apply sampled transformations. This process is formulated as $X'_f = \mathcal{R}_{w_0}(X_f, \Theta)$, where $\Theta \sim p(\Theta)$ and $\mathcal{R}_{w_0}(\cdot)$ denotes neural rendering function with weights w_0 .

3) Re-Composition. We then compose the transformed foreground X'_f with background X_b through a pre-trained diffusion-based composition model [12, 66], to handles shadows and occlusions. This process is defined as $X' = \mathcal{C}_{w_1}(X'_f, X_b)$, where w_1 denotes diffusion model’s weights.

Combining the above stages, the transformation process by MPM can be formally represented as follows:

$$X' = \mathcal{T}(\Theta, X) = \mathcal{C}_{w_1}(\mathcal{R}_{w_0}(X_f, \Theta), X_b), \quad (2)$$

where $\{X_f, X_b\} = \mathcal{F}(X)$.

3.2.2. Naturalness Reward Model

Despite the MPM pipeline utilizing generative priors to achieve realistic 3D variation samples, we observe that directly applying adversarial optimization frequently leads to unnatural artifacts, such as shape distortions and physically implausible cases in Fig. 5. These issues stem from two factors: generalization limitations of existing generative models and the inherent high semantic bias in low-quality samples. To address this challenge, we propose a novel Naturalness Reward Model (NRM) that continuously regularizes $p(\Theta)$ during optimization based on sample naturalness, preventing convergence to low-quality pseudo-optimal regions.

As shown in Fig.4, the NRM employs DINOv2[46] as its backbone, chosen for several reasons. **Firstly**, DINO




Figure 4. Architecture of NRM and computational cost and accuracy of naturalness assessment across Human, GPT, and NRM.

captures finer visual contextual features compared to traditional pre-trained vision encoders [58], aiding better convergence during training. **Secondly**, unlike vision-language aligned encoders such as CLIP, DINO’s visual features are language-agnostic, making them more faithfully reflect image quality rather than semantic consistency. We use DINOv2 to extract image tokens from the input Adv-3DT samples, followed by dual-stream prediction heads to estimate visual fidelity and physical plausibility scores:

$$\text{Score}_R(X') = f_R(\mathcal{E}_D(X')), \text{Score}_P(X') = f_P(\mathcal{E}_D(X')), \quad (3)$$

where $f_R(\cdot)$ and $f_P(\cdot)$ are two-layer MLPs. To train the NRM, we construct a large-scale naturalness preference dataset with fine-grained scoring. We generate numerous samples from the MPM pipeline on ImageNet and synthetic images, scoring each on a 5-point scale for visual fidelity and physical plausibility. These scores are initially annotated automatically by GPT-4o with Chain-of-Thought prompts and verified by volunteers for correction. This process resulted in 120K high-quality training samples.

The architecture and training parameters of NRM are detailed in Appendix B. Our experimental results demonstrate that the naturalness feedback provided by NRM highly aligns with human evaluation. In terms of prediction efficiency and stability, it outperforms methods that directly employ GPT-4o for naturalness assessment (see Fig. 4 and Appendix C.2). Based on the NRM’s predictions, we define the naturalness regularization loss for the optimization as:

$$\mathcal{L}_{\text{Nat}}(X') = -\frac{1}{2}(\text{Score}_R(X') + \text{Score}_P(X')) \quad (4)$$

3.2.3. Inverse Semantic Probability Loss

Previous studies typically employ task-specific optimization objectives with end-to-end manner to obtain worst-case Adv-3DT samples. However, directly applying this approach to VLMs with substantial parameters and diverse architectures spanning various tasks is challenging. To address this, we design the Inverse Semantic Probability (ISP) loss, which minimizes the probability that VLMs assign to the correct semantic attributes of Adv-3DT samples.

This design have two key considerations. **1)** It is based solely on the visual and text encoder \mathcal{E}_v , \mathcal{E}_t , which are fundamental components across modern VLMs, ensuring architecture-agnostic applicability. **2)** By operating in the

image-text alignment space rather than model-specific layers or task-specific heads, it enables task generalization.

Given a generated sample X' and a semantic label set $\mathcal{Y} = \{y_i\}_{i=1}^N$ containing the ground-truth label y_t , we first compute a semantic similarity vector \mathbf{S} :

$$\mathbf{S} = \{s_i\}_{i=1}^N, \text{ where } s_i = \cos[\mathcal{E}_v(X'), \mathcal{E}_t(\mathcal{P}(y_i))], \quad (5)$$

where $\cos(\cdot)$ computes cosine similarity and $\mathcal{P}(\cdot)$ follows the template "a photo of a $\{y_i\}$ " in [47]. The vector \mathbf{S} captures the representation similarities between X' and each semantic label. We then derive the conditional semantic probability through softmax transformation:

$$p(y_i|X') = \exp(s_i) / \sum_{j=1}^N \exp(s_j). \quad (6)$$

Then, the \mathcal{L}_{ISP} is defined as the negative log-likelihood of the ground-truth’s semantic probability:

$$\mathcal{L}_{\text{ISP}}(X', \mathcal{Y}) = -\log p(y_t|X'). \quad (7)$$

Intuitively, \mathcal{L}_{ISP} quantifies the likelihood of X' being misclassified by comparing its similarity to the ground-truth and irrelevant categories. Maximizing it directs the $p(\Theta)$ towards regions that effectively deceive the model.

3.3. Query-Based Adversarial Optimization

To solve Problem (1), we parameterize the adversarial distribution $p(\Theta)$ as a **Multivariate Gaussian** form:

$$\Theta = \mathbf{A} \cdot \tanh(\mathbf{z}) + \mathbf{B}, \text{ where } \mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (8)$$

where $\mathbf{A} = (\Theta_{\min} - \Theta_{\max})/2$, $\mathbf{B} = (\Theta_{\min} + \Theta_{\max})/2$, \mathbf{z} follows a Gaussian Distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^6$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{6 \times 6}$. Thus, we rewrite (1) in a specific form:

$$\arg \max_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})} [\mathcal{L}_{\text{ISP}}(X', \mathcal{Y}) + \mathcal{L}_{\text{Nat}}(X')], \quad (9)$$

$$\text{where } X' = \mathcal{T}(\mathbf{A} \cdot \tanh(\mathbf{z}) + \mathbf{B}, X).$$

Solving this problem typically involves computing the gradient of the loss w.r.t. $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ and updating them via gradient ascent. However, the forward steps involve multiple components, e.g., the denoise process, introducing uncertainty in the gradient propagation path, making gradient-based optimization challenging [54]. Therefore, we adopt the Covariance Matrix Adaptation Evolution Strategy [19, 22], an efficient query-based black-box optimizer. At each iteration t , we: **1)** Sample K candidates $\{\mathbf{z}_i^{t+1}\}_{i=1}^K$ from $\mathcal{N}(\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t)$. **2)** Generate corresponding samples $\{(X')_i^{t+1}\}_{i=1}^K$. **3)** Update distribution parameters $\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t$ as follows:

$$\boldsymbol{\mu}^{t+1} = \sum_{i=1}^k w_i \cdot \mathbf{z}_{(i:k)}^{t+1}, \text{ where } \sum_{i=1}^k w_i = 1, \quad (10)$$

$$\boldsymbol{\Sigma}^{t+1} = (1 - \eta_1 - \eta_\mu) \cdot \boldsymbol{\Sigma}^t + \eta_1 \cdot p_{\boldsymbol{\Sigma}^t}^{t+1} (p_{\boldsymbol{\Sigma}^t}^{t+1})^T + \eta_\mu \cdot \sum_{i=1}^h w_i \cdot \left(\frac{\mathbf{z}_{(i:k)}^{t+1} - \boldsymbol{\mu}^t}{\boldsymbol{\sigma}^t} \right) \left(\frac{\mathbf{z}_{(i:k)}^{t+1} - \boldsymbol{\mu}^t}{\boldsymbol{\sigma}^t} \right)^T, \quad (11)$$

Table 1. Zero-shot Classification Accuracy (%) and degradation ($\downarrow\%$) Under Adv-3DT Samples Generated from ImageNet [17].

Target Models	#Params	Methods			
		Clean	Random	$p^*(\Theta)$	Θ^*
OpenCLIP ViT-B/16 [26]	149M	98.0	62.6 ($\downarrow 35$)	54.0 ($\downarrow 44$)	18.0 ($\downarrow 80$)
OpenCLIP ViT-L/14 [26]	428M	94.4	61.7 ($\downarrow 33$)	50.9 ($\downarrow 44$)	15.3 ($\downarrow 79$)
OpenCLIP ViT-G/14 [26]	2.5B	96.4	63.5 ($\downarrow 33$)	53.5 ($\downarrow 43$)	18.7 ($\downarrow 78$)
BLIP ViT-B/16 [30]	583M	83.0	56.0 ($\downarrow 27$)	51.3 ($\downarrow 32$)	17.3 ($\downarrow 66$)
BLIP-2 ViT-L/14 [31]	3.4B	84.0	57.0 ($\downarrow 27$)	52.2 ($\downarrow 32$)	18.7 ($\downarrow 65$)
BLIP-2 ViT-G/14 [31]	4.1B	81.0	55.7 ($\downarrow 25$)	49.1 ($\downarrow 32$)	15.7 ($\downarrow 65$)

here, $\mathbf{z}_{(i:k)}$ denotes samples ranked by losses: we first select top-10 samples based on \mathcal{L}_{ISP} , then choose the top-k ($k=5$) with highest \mathcal{L}_{Nat} for updates. $w_i = 1/k$ are importance weights, η_1 and η_μ are learning rates, p_Σ tracks the evolution path, and σ^t is the step size updated via Cumulative Step-size Adaptation (CSA) [22]. Due to space constraints, complete derivations of Eq. (10), (11) and hyperparameter settings are provided in the Appendix.A.1. We also provide the pseudocode for optimization process in Appendix A.2.

4. Experiments

4.1. Experimental Setup

Tasks & Datasets. We evaluate VLMs’ robustness to 3D variations using AdvDreamer on three representative vision-language tasks. (1) For *zero-shot image classification*, we utilize the ImageNet test set [17]. (2) For *image captioning*, we employ the test splits from COCO Caption [11] and NoCaps [1] datasets. (3) For *Visual Question Answering (VQA)*, we use the VQAv2 [20] test set. For each dataset, we select 300 clean samples spanning 30 categories with clear and well-defined object instances. During the optimization process, we employ ImageNet-1K categories as the semantic label set \mathcal{Y} for the classification task and use COCO-80 categories for the captioning and VQA tasks.

Victim Models. For zero-shot classification, we target vision-language foundation models: OpenCLIP [14, 26], BLIPs [30, 31], and additional models for transfer attacks. For other tasks, we select mainstream open-source VLMs (LLaVa [34, 36], MiniGPT-4 [72], etc.) as well as representative commercial VLMs (GPT-4o and GPT-4o-mini).

Methods & Baselines. We use two variants with AdvDreamer: (1) $p^*(\Theta)$: represents the average results over 10 Adv-3DT samples randomly sampled from the optimal adversarial distribution, while (2) Θ^* : indicates the evaluation results using the Adv-3DT sample corresponding to the distribution center. Additionally, we establish two potential baselines following [18, 50]: the original natural images without transformation (*Clean*) and the average results over 10 images with random 3D transformations (*Random*).

Metrics. For zero-shot classification, we follow the CLIP-Benchmark [27] protocol, reporting Top-1 accuracy and its degradation relative to clean samples. For captioning, in addition to conventional metrics (CIDEr, BLEU), we adopt the LLM-as-a-judge approach [36, 67, 71] to address the open-

Table 2. Cross-model transferability of Adv-3DT samples, utilizing OpenCLIP and BLIP as surrogate models.

Target Model	Clean	OpenCLIP [26]		BLIP [30]	
		$p^*(\Theta)$	Θ^*	$p^*(\Theta)$	Θ^*
ALBEF [29]	65.0	37.2 ($\downarrow 28$)	27.7 ($\downarrow 37$)	40.2 ($\downarrow 25$)	25.7 ($\downarrow 39$)
OpenAI CLIP ViT-B/16 [47]	85.3	48.7 ($\downarrow 37$)	29.7 ($\downarrow 56$)	52.8 ($\downarrow 32$)	31.0 ($\downarrow 54$)
OpenCLIP ViT-B/16 [26]	97.7	54.0 ($\downarrow 44$)	18.0 ($\downarrow 80$)	59.3 ($\downarrow 38$)	34.7 ($\downarrow 63$)
Meta-CLIP ViT-B/16 [62]	91.0	49.3 ($\downarrow 42$)	30.0 ($\downarrow 61$)	53.8 ($\downarrow 37$)	29.7 ($\downarrow 61$)
BLIP ViT-B/16 [30]	82.7	48.4 ($\downarrow 34$)	30.7 ($\downarrow 52$)	51.3 ($\downarrow 31$)	17.3 ($\downarrow 65$)
BLIP-2 ViT-L/14 [31]	84.0	50.0 ($\downarrow 34$)	32.0 ($\downarrow 52$)	52.3 ($\downarrow 32$)	27.7 ($\downarrow 56$)
SigLIP ViT-B/16 [65]	95.3	55.3 ($\downarrow 40$)	34.0 ($\downarrow 61$)	59.3 ($\downarrow 36$)	34.7 ($\downarrow 61$)

ended nature of VLM responses, aligning evaluations more closely with human preferences. We use GPT-4 to assess three aspects of the response on a 10-point scale: semantic accuracy, tone confidence, and overall coherence, with the scores summed to yield the *GPT-Score* (see Fig. 1 (D)). For the VQA task, we introduce *GPT-Acc*, which leverages LLM-based judgment to evaluate response correctness. The judgment prompts are detailed in Appendix C.

4.2. Research Questions (RQs) and Findings

RQ1: Do current VLMs demonstrate sufficient robustness to 3D variations? Take-away: *”Emphatically not. Our evaluation reveals the severe vulnerability of current VLMs to worst-case 3D variations across multiple vision-language tasks.”* As shown in Tab. 1, in zero-shot classification, Adv-3DT samples generated by AdvDreamer significantly compromise VLMs’ performance, causing accuracy drops of up to 80%. This substantially outperforms the Random baseline, validating AdvDreamer’s effectiveness in capturing worst-case transformations. Notably, samples under Θ^* induce more severe performance degradation than those under the adversarial distribution $p^*(\Theta)$, indicating the existence of extremely vulnerable points.

This fundamental vulnerability extends beyond classification to more complex open-ended tasks. As illustrated in Tab.3, for image captioning, Adv-3DT samples under both $p^*(\Theta)$ and Θ^* effectively degrade model performance across all evaluation metrics, with GPT-Score decreasing by nearly 50% across various VLMs. This significant degradation reveals a critical limitation in current VLMs’ 3D understanding and reasoning capabilities. For VQA tasks, GPT-4o, despite achieving the highest accuracy (75.4%) on clean samples—exhibits a substantial 25% accuracy drop when under Adv-3DT attacks. Fig. 1 (A) provides visualizations of representative Adv-3DT samples generated by AdvDreamer, with more examples presented in Appendix.D.

RQ2: To what extent do different VLMs share common 3D variation vulnerability regions? Take-away: *”Substantially. Our analysis reveals that Adv-3DT samples demonstrate strong cross-model transferability across various VLMs, indicating shared vulnerability regions in the 3D variation space.”* As demonstrated in Tab. 2, we investigate transferability by utilizing OpenCLIP and BLIP as surrogate models to optimize $p^*(\Theta)$, then evaluate six

Table 3. **Performance Degradation (\downarrow) of VLMs on Image Captioning and VQA Tasks.** For closed-source VLMs (GPT-4o/4o-mini), we performed transfer-based attacks, using OpenAI-CLIP (ViT-L@336px) as surrogate encoder to optimize $p^*(\Theta)$. For other VLMs, $p^*(\Theta)$ are optimized w.r.t. respective image encoders (e.g., Eva-CLIP and OpenAI CLIP, etc.). Metrics: C-CIDEr, B-BLUE@4.

Model	Method	COCO Caption [11]			NoCaps [1]			VQA-V2 [20]
		C	B@4	GPT-Score	C	B@4	GPT-Score	GPT-Acc
BLIP-2 [30]	Clean	133.4	36.0	24.7	98.1	44.1	24.4	53.4
	Random	91.8 ($\downarrow 42$)	23.5 ($\downarrow 13$)	19.5 ($\downarrow 5$)	60.2 ($\downarrow 38$)	28.6 ($\downarrow 16$)	20.9 ($\downarrow 4$)	34.0 ($\downarrow 19$)
	$p^*(\Theta)$	83.8 ($\downarrow 50$)	23.3 ($\downarrow 13$)	18.9 ($\downarrow 6$)	55.7 ($\downarrow 42$)	26.9 ($\downarrow 17$)	17.1 ($\downarrow 7$)	31.7 ($\downarrow 22$)
	Θ^*	72.2 ($\downarrow 61$)	21.5 ($\downarrow 15$)	17.0 ($\downarrow 8$)	40.0 ($\downarrow 58$)	22.4 ($\downarrow 22$)	15.0 ($\downarrow 9$)	28.0 ($\downarrow 25$)
LLaVa-1.5 [36]	Clean	126.1	30.4	25.0	118.7	50.6	24.7	68.6
	Random	93.4 ($\downarrow 33$)	25.6 ($\downarrow 5$)	20.0 ($\downarrow 5$)	75.3 ($\downarrow 43$)	35.7 ($\downarrow 15$)	21.7 ($\downarrow 3$)	56.3 ($\downarrow 12$)
	$p^*(\Theta)$	85.8 ($\downarrow 40$)	22.9 ($\downarrow 8$)	18.9 ($\downarrow 6$)	60.0 ($\downarrow 59$)	29.6 ($\downarrow 21$)	17.4 ($\downarrow 7$)	54.3 ($\downarrow 14$)
	Θ^*	73.2 ($\downarrow 53$)	20.5 ($\downarrow 10$)	16.5 ($\downarrow 9$)	45.8 ($\downarrow 73$)	25.5 ($\downarrow 25$)	15.6 ($\downarrow 9$)	53.8 ($\downarrow 15$)
LLaVa-1.6 [35]	Clean	41.6	6.9	22.9	30.5	10.8	22.6	71.2
	Random	25.1 ($\downarrow 17$)	6.0 ($\downarrow 1$)	17.5 ($\downarrow 5$)	18.3 ($\downarrow 12$)	7.1 ($\downarrow 4$)	21.0 ($\downarrow 2$)	54.3 ($\downarrow 17$)
	$p^*(\Theta)$	21.9 ($\downarrow 20$)	5.3 ($\downarrow 2$)	16.9 ($\downarrow 6$)	13.6 ($\downarrow 17$)	6.5 ($\downarrow 4$)	15.7 ($\downarrow 7$)	53.0 ($\downarrow 18$)
	Θ^*	16.7 ($\downarrow 25$)	4.6 ($\downarrow 2$)	15.4 ($\downarrow 8$)	10.5 ($\downarrow 20$)	4.3 ($\downarrow 7$)	13.5 ($\downarrow 9$)	46.9 ($\downarrow 24$)
Otter [28]	Clean	127.0	34.9	24.0	76.4	29.7	22.5	53.9
	Random	71.6 ($\downarrow 55$)	22.5 ($\downarrow 12$)	21.0 ($\downarrow 3$)	55.4 ($\downarrow 21$)	23.7 ($\downarrow 6$)	21.1 ($\downarrow 1$)	49.7 ($\downarrow 4$)
	$p^*(\Theta)$	67.9 ($\downarrow 59$)	17.8 ($\downarrow 17$)	17.2 ($\downarrow 7$)	45.4 ($\downarrow 31$)	19.2 ($\downarrow 11$)	16.2 ($\downarrow 6$)	44.9 ($\downarrow 9$)
	Θ^*	58.9 ($\downarrow 68$)	17.5 ($\downarrow 17$)	15.8 ($\downarrow 8$)	36.0 ($\downarrow 40$)	15.7 ($\downarrow 14$)	14.6 ($\downarrow 8$)	42.0 ($\downarrow 12$)
MiniGPT-4 [72]	Clean	78.5	23.4	24.0	60.8	28.5	24.7	51.3
	Random	58.2 ($\downarrow 20$)	17.9 ($\downarrow 6$)	19.7 ($\downarrow 4$)	38.0 ($\downarrow 23$)	19.8 ($\downarrow 9$)	21.9 ($\downarrow 3$)	43.7 ($\downarrow 8$)
	$p^*(\Theta)$	52.3 ($\downarrow 26$)	16.3 ($\downarrow 7$)	18.1 ($\downarrow 6$)	32.9 ($\downarrow 28$)	18.4 ($\downarrow 10$)	16.7 ($\downarrow 8$)	41.0 ($\downarrow 10$)
	Θ^*	42.2 ($\downarrow 36$)	12.9 ($\downarrow 11$)	16.2 ($\downarrow 8$)	26.6 ($\downarrow 34$)	14.9 ($\downarrow 14$)	15.0 ($\downarrow 10$)	34.3 ($\downarrow 17$)
GPT-4o-mini [45]	Clean	23.7	3.0	23.3	15.7	2.9	23.0	68.6
	Random	14.1 ($\downarrow 10$)	1.8 ($\downarrow 1$)	16.9 ($\downarrow 6$)	10.8 ($\downarrow 5$)	0.9 ($\downarrow 2$)	20.0 ($\downarrow 3$)	51.7 ($\downarrow 17$)
	$p^*(\Theta)$	12.5 ($\downarrow 11$)	1.2 ($\downarrow 2$)	18.0 ($\downarrow 5$)	7.8 ($\downarrow 8$)	0.6 ($\downarrow 2$)	17.2 ($\downarrow 6$)	43.4 ($\downarrow 25$)
	Θ^*	9.6 ($\downarrow 14$)	1.6 ($\downarrow 1$)	16.3 ($\downarrow 7$)	7.3 ($\downarrow 8$)	0.4 ($\downarrow 3$)	15.5 ($\downarrow 8$)	42.7 ($\downarrow 26$)
GPT-4o [45]	Clean	51.2	9.7	25.0	40.7	11.7	24.9	75.4
	Random	36.5 ($\downarrow 15$)	5.1 ($\downarrow 5$)	19.1 ($\downarrow 6$)	24.7 ($\downarrow 16$)	9.5 ($\downarrow 2$)	20.9 ($\downarrow 4$)	56.3 ($\downarrow 19$)
	$p^*(\Theta)$	26.4 ($\downarrow 25$)	4.6 ($\downarrow 5$)	17.7 ($\downarrow 7$)	19.1 ($\downarrow 22$)	6.5 ($\downarrow 5$)	17.0 ($\downarrow 8$)	50.7 ($\downarrow 25$)
	Θ^*	20.3 ($\downarrow 31$)	4.7 ($\downarrow 5$)	16.1 ($\downarrow 9$)	15.8 ($\downarrow 25$)	5.4 ($\downarrow 6$)	14.6 ($\downarrow 10$)	50.3 ($\downarrow 25$)

Table 4. Top-1 Accuracy and Naturalness Score of Adv-3DT samples generated by AdvDreamer w/ and w/o NRM. (under $p^*(\Theta)$)

Methods	Top-1 Acc.	$Score_R$	$Score_P$
AdvDreamer w/o. NRM	48.6% ($\downarrow 49.4$)	1.60	1.39
AdvDreamer w/ NRM	54.0% ($\downarrow 44.0$)	2.52	2.57

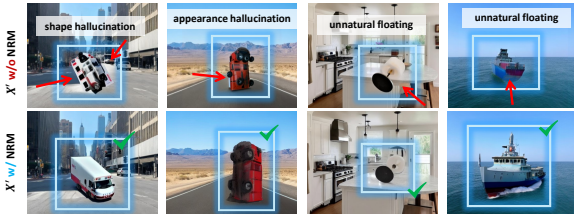


Figure 5. Visualization of Adv-3DT samples w/ and w/o NRM.

target VLMs with diverse architectures and training objectives. The transfer attacks maintain strong effectiveness, with performance degradation comparable to direct attacks. Notably, $p^*(\Theta)$ exhibits smaller transfer gaps than Θ^* , suggesting that learning the underlying distribution can better capture generalizable vulnerability regions. This transferability extends consistently to captioning and VQA tasks (Tab. 3). When employing OpenAI-CLIP as the surrogate encoder, Adv-3DT samples generated by AdvDreamer effectively compromise commercial systems like GPT-4o. We attribute this universal transferability to a fundamental limitation: current image-text pretraining datasets exhibit significant 3D biases and fail to encompass diverse certain special 3D transformations, leading to universal performance deficiencies under these transformation regions.

RQ3: Does the performance degradation under Adv-3DT samples generated by AdvDreamer fundamentally

arise from 3D variations? Take-away: "Yes, we successfully reproduce these Adv-3DT samples in physical environments and demonstrate that degradation is inherently caused by the 3D variation rather than undesirable image quality problems." With the naturalness feedback provided by NRM, AdvDreamer effectively prevents unwanted degradation of samples during the optimization process. To verify this, we evaluate Adv-3DT samples generated by AdvDreamer with and without NRM under ImageNet. Tab. 4 presents the accuracy under OpenCLIP (ViT-B/16), alongside GPT-4o-assessed visual realism ($Score_R$) and physical plausibility ($Score_P$). The results indicate that NRM significantly improves samples' realism and physical feasibility, with a marginal trade-off in aggressiveness. Furthermore, Fig. 5 presents a qualitative comparison, illustrating the effectiveness of NRM's naturality feedback in improving the image quality and mitigating the distribution discrepancy compared to natural image distributions. In addition, we validate the alignment between NRM predictions and human evaluations, as detailed in Appendix C.2.

Furthermore, we conduct physical-world experiments on 12 common objects from traffic and household environments to validate this. Our experimental procedure consisted of three main steps: 1) capturing images of objects in their natural state, 2) using AdvDreamer with OpenCLIP as the surrogate model to generate Adv-3DT samples, and 3) physically reproducing the adversarial 3D transformations captured by AdvDreamer through video recording. This process yielded 3,807 frames of physical Adv-3DT samples. We evaluate effectiveness through zero-shot classifi-

Table 5. Quantitative results of physical experiments.

Sample	OpenCLIP		LLaVa-1.5		GPT-4o	
	Acc.	Conf.	Acc.1	Acc.2	Acc.1	Acc.2
Natural (Physical)	100.0	0.892	83.3	100	91.7	100.0
Adv-3DT (AdvDreamer)	0.0 ($\downarrow 100$)	0.003	25.0 ($\downarrow 58$)	25.0 ($\downarrow 75$)	25.0 ($\downarrow 67$)	41.7 ($\downarrow 58$)
Adv-3DT (Physical)	51.3 ($\downarrow 49$)	0.413	33.6 ($\downarrow 50$)	45.1 ($\downarrow 55$)	56.8 ($\downarrow 35$)	64.8 ($\downarrow 35$)




Figure 6. Attack success rate of Adv-3DT samples w.r.t. (A) optimization steps and (B) scaling ratios of the optimal distribution.

caption and VQA tasks. For VQA, we designed two evaluation protocols: multiple-choice classification (Acc.1) and binary verification (Acc.2). As shown in Table 5, physically reproduced Adv-3DT samples significantly degraded VLM performance: the accuracy of zero-shot classification dropped to 51.3%, while VQA performance decreased to 33.6% (Acc.1) and 45.1% (Acc.2). These results confirm that AdvDreamer successfully captures physically realizable 3D variations that represent real world corner cases. However, an aggressiveness gap between digital and physical attacks is observed, primarily attributed to data distribution differences and camera parameter fluctuations. Fig. 1 (B) presents representative examples from physical experiments, with more detailed results provided in Appendix D.

4.3. Ablation Studies and Additional Analysis

Convergence Analysis & Computational Efficiency. We analyze the attack success rates (ASR) of $p^*(\Theta)$ against OpenCLIP (ViT-B/16) and BLIP (ViT-B/16) across different optimization iterations. As shown in Fig. 6, AdvDreamer’s effectiveness consistently improves with more iterations and converges after 15 steps. Thus, we set the iteration step default to 15 throughout our experiments to balance effectiveness and efficiency. The computational cost is discussed in Appendix C.4 by detailing the overhead of each component. For each clean sample, AdvDreamer requires an average of 0.28 GPU hours with a single RTX 3090.

Does Optimal Distribution $p^*(\Theta)$ Effectively Characterize the Worst-case 3D Variations? We conduct analysis on ImageNet using five random natural images. For each image, we generate 100 Adv-3DT samples by scaling the variance Σ of their $p^*(\Theta)$ with different ratios. Fig. 6 (B) shows that ASR consistently decreases as the sampling range expands beyond the optimal distribution. This inverse relationship confirms that AdvDreamer successfully identifies the most vulnerable regions in the 3D variation space.




Figure 7. Benchmarking result of VLMs’ 3D variation robustness.

4.4. MM3DTBench

Finally, We introduce MM3DTBench, it comprises 215 challenging Adv-3DT samples, consisting of samples generated by AdvDreamer and their physical replications. For each sample, we provide question templates with candidate answers and precise orientation annotations for the transformed objects. We employ two evaluation tasks: multiple-choice (**Choice**) and free-form description (**Free answer**), with accuracy computed through text-only GPT-4 judgment. More Details are provided in the Appendix.E. Our evaluation of 13 representative VLMs reveals robustness gaps. As shown in Fig. 7, even top-performing models (GPT-4o [45], CogVLM [61], and Intern-VL [13]) achieve limited success, while most models struggle with accuracy below 50%. Notably, Claude-3 [5], despite its recognized excellence, only achieves an accuracy of 30.0%. This performance correlates with models’ general out-of-distribution robustness measured by MultiTrust [67], suggesting a fundamental limitation of current VLMs. We encourage the adoption of MM3DTBench in VLM development to advance security evaluation in dynamic scenarios.

5. Conclusion

We proposed AdvDreamer, a novel framework for characterizing real-world adversarial 3D transformation samples from single-view observations. Our comprehensive evaluation revealed critical robustness gaps in existing VLMs, highlighting the urgent need to enhance VLMs’ 3D variation perception and understanding capabilities in safety-critical applications. Through MM3DTBench, we established the first benchmark for worst-case 3D variations, aiming to facilitate the development of more robust vision-language systems for real-world deployment.

Acknowledgement

This work was supported by the Fundamental Research Funds for the Central Universities, and was also supported by NSFC Projects (U2341228, 62276149).

References

- [1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019. 6, 7
- [2] Haider Al-Tahan, Quentin Garrido, Randall Balestriero, Diane Bouchacourt, Caner Hazirbas, and Mark Ibrahim. Unibench: Visual reasoning requires rethinking vision-language beyond scaling. *arXiv preprint arXiv:2408.04810*, 2024. 2
- [3] 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. 2
- [4] Michael A Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4845–4854, 2019. 2, 3, 4
- [5] Anthropic. claude-3. <https://www.anthropic.com/claude>, 2024. 8
- [6] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *International conference on machine learning*, pages 284–293. PMLR, 2018. 4
- [7] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 2
- [8] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems*, 32, 2019. 3
- [9] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465, 2024. 2
- [10] Liangyu Chen, Bo Li, Sheng Shen, Jingkang Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell, and Ziwei Liu. Large language models are visual reasoning coordinators. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [11] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6, 7
- [12] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6593–6602, 2024. 2, 4
- [13] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198, 2024. 2, 3, 8
- [14] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2818–2829, 2023. 6
- [15] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. In *ICLR 2023 (Eleventh International Conference on Learning Representations)*, 2023. 4
- [16] Xuanming Cui, Alejandro Aparcedo, Young Kyun Jang, and Ser-Nam Lim. On the robustness of large multimodal models against image adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24625–24634, 2024. 2, 3
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [18] Yinpeng Dong, Shouwei Ruan, Hang Su, Caixin Kang, Xingxing Wei, and Jun Zhu. Viewfool: Evaluating the robustness of visual recognition to adversarial viewpoints. In *Neural Information Processing Systems (NeurIPS)*, 2022. 2, 3, 4, 6
- [19] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and David Sculley. Google vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1487–1495, 2017. 5
- [20] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 6, 7
- [21] Abdullah Hamdi and Bernard Ghanem. Towards analyzing semantic robustness of deep neural networks. In *ECCV*, pages 22–38. Springer, 2020. 2, 3, 4
- [22] Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016. 5, 6
- [23] Zexin He and Tengfei Wang. Openlm: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM>, 2023. 2
- [24] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400*, 2023. 2
- [25] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value

- maps for robotic manipulation with language models. *arXiv preprint arXiv:2307.05973*, 2023. 2, 3
- [26] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. openclip, 2021. 3, 6
- [27] LAION-AI. Clip-benchmark. https://github.com/LAION-AI/CLIP_benchmark, 2024. 6
- [28] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Fanyi Pu, Jingkang Yang, Chunyuan Li, and Ziwei Liu. Mimic-it: Multi-modal in-context instruction tuning. *arXiv preprint arXiv:2306.05425*, 2023. 7
- [29] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 6
- [30] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022. 3, 6, 7
- [31] 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*, pages 19730–19742. PMLR, 2023. 2, 3, 6
- [32] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 3
- [33] 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. 3
- [34] 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. 6
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 2, 3, 7
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 2, 3, 6, 7
- [37] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision*, pages 216–233. Springer, 2024. 3
- [38] Jiahuan Long, Tingsong Jiang, Wen Yao, Shuai Jia, Weijia Zhang, Weien Zhou, Chao Ma, and Xiaoqian Chen. Papat: Exploring adversarial patch attack against multiple object tracking. In *European Conference on Computer Vision*, pages 128–144. Springer, 2024. 2
- [39] Jiahuan Long, Zhengqin Xu, Tingsong Jiang, Wen Yao, Shuai Jia, Chao Ma, and Xiaoqian Chen. Robust sam: On the adversarial robustness of vision foundation models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5775–5783, 2025. 2
- [40] Jiahuan Long, Wen Yao, Tingsong Jiang, and Chao Ma. Cdupatch: Color-driven universal adversarial patch attack for dual-modal visible-infrared detectors. *arXiv preprint arXiv:2504.10888*, 2025. 2
- [41] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019. 3
- [42] Spandan Madan, Timothy Henry, Jamell Dozier, Helen Ho, Nishchal Bhandari, Tomotake Sasaki, Frédo Durand, Hanspeter Pfister, and Xavier Boix. When and how convolutional neural networks generalize to out-of-distribution category–viewpoint combinations. *Nature Machine Intelligence*, 4(2):146–153, 2022. 2, 3, 4
- [43] Chengzhi Mao, Scott Geng, Junfeng Yang, Xin Wang, and Carl Vondrick. Understanding zero-shot adversarial robustness for large-scale models. *arXiv preprint arXiv:2212.07016*, 2022. 3
- [44] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 405–421, 2020. 3
- [45] openai. Gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. 3, 7, 8
- [46] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 4
- [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. 2, 3, 5, 6
- [48] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 4
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 4
- [50] Shouwei Ruan, Yinpeng Dong, Hang Su, Ning Chen, and Xingxing Wei. Towards viewpoint-invariant visual recognition via adversarial training. In *ICCV*, pages 1–10, 2023. 2, 3, 4, 6

- [51] Shouwei Ruan, Yinpeng Dong, Hang Su, Jianteng Peng, Ning Chen, and Xingxing Wei. Improving viewpoint robustness for visual recognition via adversarial training. *arXiv preprint arXiv:2307.11528*, 2023. 3
- [52] Shouwei Ruan, Yinpeng Dong, Hanqing Liu, Yao Huang, Hang Su, and Xingxing Wei. Omniview-tuning: Boosting viewpoint invariance of vision-language pre-training models. *arXiv preprint arXiv:2404.12139*, 2024. 3
- [53] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [54] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 5
- [55] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023. 2
- [56] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Chenxu Hu, Yang Wang, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*, 2024. 2
- [57] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024. 4
- [58] Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms, 2024. 3, 5
- [59] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 3
- [60] 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. 3
- [61] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023. 8
- [62] Hu Xu, Saining Xie, Xiaoqing Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. Demystifying clip data. In *The Twelfth International Conference on Learning Representations*, 2024. 6
- [63] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 2
- [64] Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 42–48. IEEE, 2024. 2, 3
- [65] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986, 2023. 2, 6
- [66] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Scaling in-the-wild training for diffusion-based illumination harmonization and editing by imposing consistent light transport. In *The Thirteenth International Conference on Learning Representations*, 2025. 2, 4
- [67] Yichi Zhang, Yao Huang, Yitong Sun, Chang Liu, Zhe Zhao, Zhengwei Fang, Yifan Wang, Huanran Chen, Xiao Yang, Xingxing Wei, Hang Su, Yinpeng Dong, and Jun Zhu. Benchmarking trustworthiness of multimodal large language models: A comprehensive study, 2024. 2, 3, 6, 8
- [68] Yabin Zhang, Wenjie Zhu, Chenhang He, and Lei Zhang. Lapt: Label-driven automated prompt tuning for ood detection with vision-language models. *Proceedings of the european conference on computer vision (ECCV)*, 2024. 2
- [69] Bingchen Zhao, Shaozuo Yu, Wufei Ma, Mingxin Yu, Shenxiao Mei, Angtian Wang, Ju He, Alan Yuille, and Adam Kortylewski. Ood-cv: A benchmark for robustness to out-of-distribution shifts of individual nuisances in natural images. In *European conference on computer vision*, pages 163–180. Springer, 2022. 3
- [70] 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. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2
- [71] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023. 6
- [72] 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. 6, 7