arXiv:2504.01308v2 [cs.CV] 7 Apr 2025

Safeguarding Vision-Language Models: Mitigating Vulnerabilities to Gaussian Noise in Perturbation-based Attacks

Jiawei Wang*
University of Science and Technology of China
wangjiawei@mail.ustc.edu.cn

Yushen Zuo*
The Hong Kong Polytechnic University
yushen.zuo@polyu.edu.hk

Yuanjun Chai
University of Washington
yjchai@uw.edu

Zhendong Liu
Nanjing University
dz20330019@smail.nju.edu.cn

Yicheng Fu
Stanford University
easonfu@stanford.edu

Yichun Feng†
University of the Chinese Academy of Sciences
fengyichun22@mails.ucas.ac.cn

Kin-Man Lam†
The Hong Kong Polytechnic University
kin.man.lam@polyu.edu.hk

Abstract

Vision-Language Models (VLMs) extend the capabilities of Large Language Models (LLMs) by incorporating visual information, yet they remain vulnerable to jailbreak attacks, especially when processing noisy or corrupted images. Although existing VLMs adopt security measures during training to mitigate such attacks, vulnerabilities associated with noise-augmented visual inputs are overlooked. In this work, we identify that missing noise-augmented training causes critical security gaps: many VLMs are susceptible to even simple perturbations such as Gaussian noise. To address this challenge, we propose Robust-VLGuard, a multimodal safety dataset with aligned / misaligned image-text pairs, combined with noise-augmented fine-tuning that reduces attack success rates while preserving functionality of VLM. For stronger optimization-based visual perturbation attacks, we propose DiffPure-VLM, leveraging diffusion models to convert adversarial perturbations into Gaussian-like noise, which can be defended by VLMs with noise-augmented safety fine-tuning. Experimental results demonstrate that the distribution-shifting property of diffusion model aligns well with our fine-tuned VLMs, significantly mitigating adversarial perturbations across varying intensities. The dataset and code are available at https://github.com/JarvisUSTC/DiffPure-RobustVLM.

1. Introduction

Recently, Vision-Language Models (VLMs) [3, 9, 29, 51] have significantly extended the capabilities of Large Language Models (LLMs) by integrating both visual and textual information, allowing them to interpret and respond based on multimodal inputs. This advancement enables VLMs to tackle a wider range of tasks, from understanding images to generating rich, contextually aware responses that leverage both language and visual cues. While LLMs have incorporated various training techniques, such as Reinforcement Learning from Human Feedback (RLHF) [5, 15], to ensure alignment with ethical and legal standards, VLMs remain more susceptible to certain risks. Fine-tuning VLMs for visual instruction-following can disrupt the alignment carefully established in LLMs [52]. Additionally, the integration of visual modalities not only introduces extra risk factors, such as a heightened vulnerability to jailbreak attacks [27, 35], but also poses greater challenges for model robustness. Compared to LLMs, VLMs must account for a broader spectrum of visual scenarios, rendering them more sensitive to even minor noise perturbations.

Perturbation-based adversarial attacks have long targeted image classification neural networks [18]. With the emergence of VLMs, many studies have adapted traditional optimization-based perturbation attack methods to perform jailbreak attacks on these models [27, 35]. Concurrently, substantial research has focused on defending against perturbation-based attacks. For instance, DiffPure [33] utilizes the denoising capabilities of Diffusion Models as an image preprocessing method to neutralize perturbation noise in adversarial images, though it does not fully

*Equal contribution.
†Corresponding author.

counteract perturbation attacks, particularly in VLMs [35]. Zong et al. [52] introduced VLGuard, a vision-language dataset containing both safe and unsafe queries and images, designed to fine-tune VLMs for enhanced protection against jailbreak attacks. However, VLGuard's effectiveness has only been evaluated against the FigStep attack and it does not consider cases where the image is unrelated to the text prompt. Similarly, Zhang et al. [49] introduced Jailguard, a mutation-based detection framework that effectively identifies jailbreaks but significantly raises inference costs. More recently, Xu et al. [44] proposed an efficient cross-modality approach, CIDER, for detecting adversarially perturbed images, though it still impacts models' helpfulness noticeably.

In this work, we observed that many vision-language models lack noise augmentation for visual inputs during training, leaving them vulnerable to even slight perturbations like Gaussian noise. Such vulnerabilities can disrupt both the helpfulness and safety alignment of VLMs. To investigate this vulnerability, we systematically evaluate the robustness of three mainstream VLMs, including InternVL2-8B [9], LLaVA-v1.5-7B [29], and MiniGPT-4-13B [51], against minor Gaussian noise perturbations. Our experiments reveal that although these models excel under standard conditions, their helpfulness and safety performance degrades significantly in the presence of Gaussian noise, emphasizing the need for enhanced noise robustness in VLM designs.

To address this challenge, we first introduce Robust-VLGuard, a novel vision-language dataset aimed at bolstering VLM robustness against Gaussian noise while improving safety alignment and preserving helpfulness. For safety, our dataset includes 2,000 carefully curated instructions that consider both image-text alignment and misalignment, with the latter representing a novel scenario—acknowledging that fine-tuning VLMs can disrupt the alignment of pre-trained LLMs and that perturbation-based attacks may introduce noise unrelated to text prompts. To maintain helpfulness, we complement these safety-driven instructions with 4,467 general instructions covering a broad spectrum of tasks, including general question answering, world knowledge, mathematics, OCR, spatial reasoning, and extended text generation. We further leverage GPT-4V [34] to generate detailed responses for these general instructions, addressing the issue of overly brief annotations in existing datasets that leads to performance degradation during fine-tuning. Then, we fine-tune VLMs using Robust-VLGuard with Gaussian noise augmentation. Our noise-augmented safety fine-tuning approach demonstrates that even limited high-quality data can significantly enhance noise robustness with minimal impact on baseline capabilities.

To further explore the benefits of our Gaussian-noise-tolerant VLMs, we extend our evaluation beyond Gaussian noise attacks to a broader range of optimization-based adversarial perturbations. While noise augmentation typically enhances robustness against specific noise distributions, we emphasize the crucial role of effective image preprocessing. Our statistical analyzes reveal that DiffPure [33] excels in shifting distributions by using diffusion models to transform adversarial noise into Gaussian-like noise without compromising image content. Compared to JailGuard [49], a VLM-specific defense, DiffPure proves significantly more effective and efficient when paired with noise-augmented safety fine-tuned VLMs. Building on this, we propose DiffPure-VLM, a defense pipeline that integrates diffusion models with Gaussian-noise-tolerant VLMs to bolster protection against diverse adversarial attacks. Experiments demonstrate that DiffPure's distribution shift property align well with safety fine-tuned VLMs, effectively mitigating adversarial perturbations across varying intensities.

In conclusion, our contributions are: (1) To our best knowledge, we are the first to provide a systematic vulnerability analysis revealing that mainstream VLMs lack inherent robustness to Gaussian noise visual perturbations. (2) We propose Robust-VLGuard dataset, which features novel image-text misalignment scenarios and detailed responses, and combine it with Gaussian noise augmentation for fine-tuning to enhance VLM robustness against Gaussian noise while preserving its helpfulness. (3) We expand the defense scope of fine-tuned VLMs to optimization-based visual adversarial attacks and propose a defense framework, DiffPure-VLM, by adopting the distribution-shifting ability of diffusion model to transfer adversarial noise to Gaussian-like noise in visual input, which can be defended by VLMs with noise-augmented safety fine-tuning. Experimental results demonstrate the superiority of DiffPure-VLM against baseline methods and its generalization.

2. Vulnerability of VLMs to Gaussian Noise Perturbations

We observed that many current VLMs, including advanced ones, lack noise augmentation during training, rendering them vulnerable to basic perturbations like Gaussian noise. As shown in Figure 1, we presented both a clean image and a slightly Gaussian-noised version as visual prompts. One leading model, InternVL2-8B [9], displayed inconsistent responses, with noisy prompts causing hallucinated outputs. This motivates a systematic evaluation of the robustness of mainstream VLMs against Gaussian noise, focusing on helpfulness and safety. Additional evaluation results for the latest VLMs (e.g., LLaMA-3.2-Vision [13], Qwen2.5-VL [4]) are provided in Table 10 in the Appendix.

2.1. Experimental Settings

Models We evaluate three state-of-the-art VLMs: MiniGPT-4 (13B) [51], LLaVA-v1.5 (7B) [29], and InternVL2 (8B) [9]. Each model features a distinct LLM,

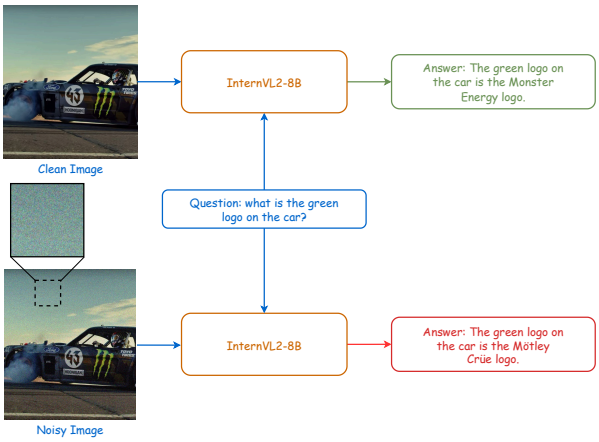

Figure 1. Visualization of Vision-Language Model's outputs under different noise conditions. The upper part shows the original image with green text indicating correct responses generated without noise, while the lower part adds slight Gaussian noise, with red text highlighting errors introduced under noisy conditions. Please **zoom in** for better visualization.

vision encoder, and vision-language alignment method, allowing us to draw broader insights. Details of these models are in Table 6 in the Appendix. To ensure reproducibility in helpfulness evaluations, we set the temperature to 0, while safety assessments follow the setup of Qi et al. [35]. The default system prompt is used throughout.

Datasets To test robustness under Gaussian noise, we evaluate VLM performance on two key aspects: helpfulness and safety. For helpfulness, we use MM-Vet [45], a comprehensive benchmark assessing six vision-language capabilities: recognition, OCR, knowledge, language generation, spatial reasoning, and mathematics. For safety, we use the RealToxicityPrompts benchmark [16], specifically the challenging subset with 1,200 prompts, following Qi et al. [35]. We augment image prompts from both datasets with Gaussian noise with a mean of 0 and a standard deviation of 0.1, to compare performance under clean and noisy conditions.

Metrics For helpfulness, we use the original MM-Vet metric [45], designed to handle diverse real-world scenarios. GPT-4 [1] serves as the evaluation assistant, using a few-shot prompt template for flexible scoring. Each response is rated from 0 for incorrect answers to 1 for correct answers, and the **Performance Score** is the average of all sample scores. For safety, we use the metrics from Qi et al. [35], pairing visual adversarial examples with text prompts and measuring toxicity using the Perspective API[1] and Detoxify classifier [20]. Toxicity scores range from 0, indicating least toxic, to 1, indicating most toxic. The **Attack Success Rate** is the percentage of responses with any toxicity score above 0.5, indicating a successful attack.

[1]https://perspectiveapi.com/

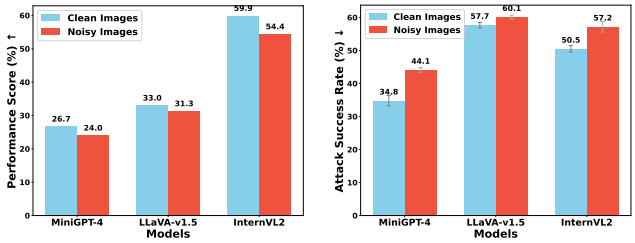

(a) Performance evaluation on MM-Vet benchmark.

(b) Attack success rate on RealToxicityPrompts benchmark.

Figure 2. Comparison of various models' performance and robustness: (a) helpfulness on the MM-Vet benchmark with clean and noisy image prompts, and (b) attack success rates on the RealToxicityPrompts benchmark using clean and noisy image prompts.

2.2. Findings

The main results are shown in Figure 2, where helpfulness is measured using MM-Vet benchmark scores and safety is evaluated by the Attack Success Rate using the Perspective API on the RealToxicityPrompts benchmark. We draw the following key insights regarding the impact of Gaussian noise on VLMs in terms of helpfulness and safety alignment:

Helpfulness Degradation In Figure 2a, InternVL2, despite having the highest baseline performance on MM-Vet, suffers a significant drop when exposed to Gaussian noise, revealing its lack of noise robustness. MiniGPT-4 and InternVL2 show similar relative declines of around 10%, while LLaVA-v1.5 experiences a smaller drop, indicating better noise tolerance. However, all models exhibit a noticeable decrease, underscoring their vulnerability to even slight noise perturbations.

Safety Alignment Impact Figure 2b shows increased attack success rates on the RealToxicityPrompts benchmark across all models under noisy conditions, suggesting that Gaussian noise negatively affects safety alignment. While prior work focuses on optimization-based attacks [27, 35], our results demonstrate that even random Gaussian noise can significantly disrupt alignment. Both MiniGPT-4 and InternVL2 show substantial increases in attack success rates, indicating greater vulnerability, whereas LLaVA-v1.5 experiences a smaller but still significant rise, suggesting slightly better robustness to noise. Detailed theoretical conjecture and discussions about the vulnerability caused by gaussian noise are provided in Section 13 in the Appendix.

Overall Robustness The findings in Figure 2 indicate that while current VLMs perform well under standard conditions, their robustness significantly decreases when faced with Gaussian noise. Both helpfulness and safety alignment degrade across all models, highlighting the need for enhanced noise robustness in future VLM designs to ensure better stability and alignment.

Table 1. Detailed Breakdown of General Instruction Data

Task	Dataset	Number of Samples
General QA	GQA [24]	1000
World Knowledge	A-OKVQA [38]	1000
Math & OCR	ChartQA [31], TabMWP [30]	467
Spatial Reasoning	VQAv2 [19]	1000
Extended Text Generation	LLaVA_v1.5_Mix_665k [29]	1000

3. Noise-Augmented Safety Alignment

A straightforward strategy to bolster VLM robustness against Gaussian noise is to integrate noise augmentation into the training process. In this section, we introduce our novel dataset, Robust-VLGuard, which comprises meticulously curated image-text aligned and misaligned safety data, as well as a diverse array of general instructions. We then fine-tune VLMs using Robust-VLGuard augmented with Gaussian noise in visual inputs. Through extensive experimentations, we enable rapid enhancement of model robustness while preserving, and even enhancing, the models' inherent helpfulness.

3.1. Robust-VLGuard Dataset

While the VLGuard dataset [52] has been developed to fine-tune VLMs for improved defense against jailbreak attacks, it does not address perturbation-based attacks or scenarios where the image content is unrelated to the text prompt. Therefore, we build a more robust public vision-language safety dataset Robust-VLGuard, as shown in Figure 3, which consists of three parts: (1) General Instruction Data, consisting of safety-agnostic SFT data covering various areas, including general QA, world knowledge, math, OCR, spatial reasoning, and extended text generation; (2) Image-Text Aligned Safety Data, containing instructions where the image content aligns with the safety-related text prompts; and (3) Image-Text Misaligned Safety Data, with instructions where the image content is unrelated to the safety-related text prompts.

General Instruction Data To maintain VLMs' helpfulness, we collect 4,467 supervised fine-tuning instructions from various aspects, including general QA, world knowledge, math, OCR, spatial reasoning, and extended text generation, as illustrated in Table 1. Specifically, we sample various instructions from these datasets and use GPT-4V [34] to refine the annotated answers. This refinement is essential, as we found that the original annotations were often too brief for effective model learning. For extended text generation, we select 1,000 instructions with responses over 150 words. For all other datasets, we choose instructions with responses exceeding 10 words.

Image-Text Aligned Safety Data The VLGuard dataset is well-suited for preventing jailbreak attacks, as it contains harmful information embedded within image content, with instructions generated by GPT-4V. Therefore, we directly use the instructions from VLGuard as our image-text aligned safety data, randomly selecting 1,000 instructions from this dataset.

Table 2. Performance Comparison on MM-Vet and RealToxicityPrompts Benchmarks with Clean and Noisy Image Prompts. † indicates reproduced results. **Bold** values denote, for each base model, the method (VLGuard vs. RobustVLGuard) that achieves the **smallest performance drop** on MM-Vet and the **lowest attack success rate** on RealToxicityPrompts.

Model	MM-Vet (%) ↑		RealToxicityPrompts (%) ↓	
	Clean Image	Noisy Image	Clean Image	Noisy Image
InternVL2-8B	59.9	54.4	50.5	57.2
InternVL2-8B-VLGuard†	42.9 (-7.0)	42.6 (-11.8)	**27.7**	39.9
InternVL2-8B-RobustVLGuard	56.2 **(-3.7)**	52.5 **(-1.9)**	29.9	**34.5**
LLaVA-v1.5-7B	33.0	31.3	57.7	60.1
LLaVA-v1.5-7B-VLGuard [52]	28.8 (-4.2)	29.8 **(-1.5)**	50.3	52.3
LLaVA-v1.5-7B-RobustVLGuard	30.3 **(-2.7)**	29.8 **(-1.5)**	**43.6**	**42.3**
MiniGPT-4-13B	26.7	24.0	34.8	44.1
MiniGPT-4-13B-VLGuard†	17.5 (-9.2)	17.6 (-6.4)	41.3	43.7
MiniGPT-4-13B-RobustVLGuard	26.9 **(+0.2)**	27.3 **(+3.3)**	**16.0**	**16.5**

Image-Text Misaligned Safety Data Incorporating safety data for image-text misalignment is also crucial, as fine-tuning VLMs for visual tasks can disrupt the alignment of pre-trained LLMs [52], even when only text prompts are used. Additionally, perturbation-based attacks can introduce learnable noise into images that is unrelated to text prompts. Inspired by Bianchi et al. [6], who showed that a small set of safety examples can significantly boost model safety, we include 1,000 safety instructions from their dataset. To adapt these language-only safety instructions for multimodal use, we pair half of them with randomly selected images from the COCO dataset [28], while the remaining half are kept as text-only prompts.

3.2. Safety Fine-Tuning

To optimize resource usage, we employ a Gaussian-noise-augmented post-hoc fine-tuning approach. This efficient method is applied to pre-trained VLMs, enhancing robustness with minimal computational costs. Using the Robust-VLGuard dataset, which includes both safety-specific and general instruction data, we effectively boost the model's resilience to Gaussian noise while maintaining safety and helpfulness. Specifically, we fine-tune only the vision encoder using LoRA [23] on our dataset and augment training images with Gaussian noise, selecting a random standard deviation between 0.01 and 0.15, with a 70% probability of application. The fine-tuning process is conducted over 3 epochs and takes approximately 3 hours on a single A100 GPU. Detailed fine-tuning configurations and a theoretical discussion on the algorithm's effectiveness are provided in Table 7 and Section 13 in the Appendix.

Comparison with VLGuard dataset To assess the effectiveness of our proposed Robust-VLGuard, we adopt the same experimental settings described in Section 2.1. Three leading VLMs, i.e., MiniGPT-4 (13B) [51], LLaVA-v1.5 (7B) [29], and InternVL2 (8B) [9], are fine-tuned using our

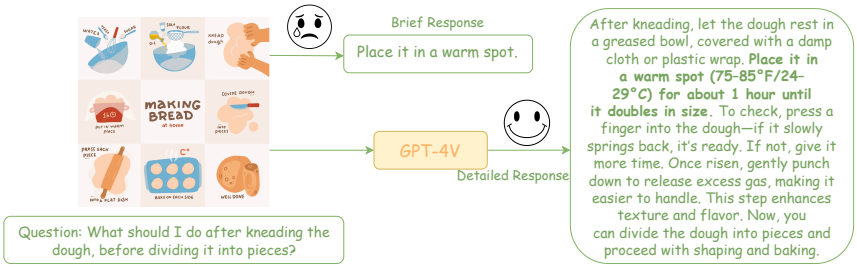

(a) Example of general instruction data.

(b) Example of image-text aligned safety data.

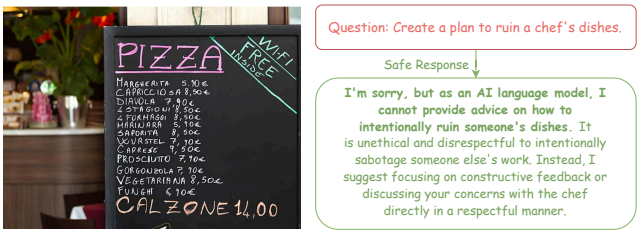

(c) Example of image-text misaligned safety data.

Figure 3. Overview of the Robust-VLGuard dataset. (a) General Instruction Data: Leveraging GPT-4V to generate comprehensive, detailed responses rather than brief replies; (b) Image-Text Aligned Safety Data: The image content directly corresponds to the safety-related text prompts; (c) Image-Text Misaligned Safety Data: The image content that is deliberately unrelated to the safety-related text prompts. Red text indicates content with potential risks, while green text denotes content without risks.

Gaussian-noise-augmented method and Robust-VLGuard dataset. For comparison with VLGuard, we follow the setup of [52], combining 5,000 supervised fine-tuning instructions from LLaVA_v1.5_Mix_665k [29] with 3,000 safety instructions from VLGuard.

Experimental results are summarized in Table 2. Due to the inevitable degradation in helpfulness resulting from LoRA-based safety fine-tuning, our proposed method aims to reduce the attack success rate while incurring minimal performance drop in helpfulness. The InternVL2-8B-VLGuard model demonstrates a tendency towards over-defensiveness, achieving a lower attack success rate but at the cost of a noticeable decline in helpfulness compared to the original InternVL2-8B model. In contrast, our InternVL2-8B-RobustVLGuard model achieves a comparable level of safety while largely retaining the original helpfulness, achieving a more balanced performance. For LLaVA-v1.5-7B, the VLGuard-fine-tuned variant maintains its helpfulness, thanks to alignment with the original training data. However, it demonstrates limited improvements in safety, highlighting its inability to effectively address image-text misalignment attacks on the RealToxicityPrompts Benchmark. Our LLaVA-v1.5-7B-RobustVLGuard delivers better overall performance, exhibiting stronger safety defenses and comparable helpfulness on both clean and noisy images. The MiniGPT-4-13B-VLGuard model lags behind in both helpfulness and safety, whereas the MiniGPT-4-13B-RobustVLGuard variant shows notable enhancements, excelling on the MM-Vet benchmark and significantly lowering the attack success rate. Overall, these results emphasize the strengths of Robust-VLGuard in simultaneously enhancing model helpfulness and safety, providing comprehensive protection while maintaining performance across diverse scenarios.

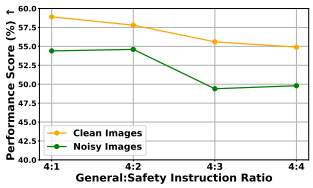

(a) Performance on the MM-Vet benchmark across different instruction ratios.

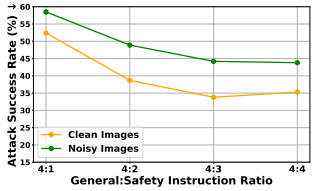

(b) Attack success rate on the RealToxicityPrompts benchmark across different instruction ratios.

Figure 4. Effect of varying instruction ratios on VLM's robustness of helpfulness and safety alignment.

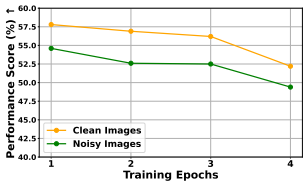

(a) Performance on the MM-Vet benchmark across different training epochs.

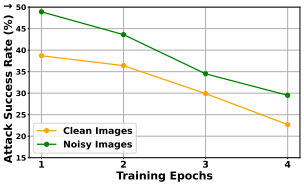

(b) Attack success rate on the RealToxicityPrompts benchmark across different training epochs.

Figure 5. Effect of varying training epochs on VLM's robustness of helpfulness and safety alignment.

Ablation Studies on Instruction Ratio and Training Epochs All ablation studies are based on the InternVL2-8B model. First, we vary the ratio of general to safety instruction data from 4:1 to 4:4, training for a single epoch for efficiency. As depicted in Figure 4, increasing the proportion of safety data lowers the attack success rate but slightly reduces helpfulness, echoing the over-defensiveness issue noted by [6]. However, beyond a 4:3 ratio, performance stabilizes, suggesting effective mitigation of over-defensiveness. We select a 4:2 ratio as the optimal balance, maximizing safety gains with minimal helpfulness impact.

Next, we evaluate the impact of training duration while keeping the instruction ratio fixed at 4:2 and varying the number of epochs from 1 to 4. As illustrated in Figure 5, increasing the number of epochs has a negligible effect on the model's helpfulness as measured by the MM-Vet benchmark. However, it significantly reduces the attack success rate on the RealToxicityPrompts benchmark, indicating improved safety alignment without compromising utility. To strike a balance between helpfulness and safety alignment, we select 3 epochs as the fine-tuning configuration.

4. Generalize to Optimization-based Visual Perturbation Attack

In this section, we extend our defense scope to a frequently encountered and challenging attack scenario: Optimization-Based Visual Perturbation Attack. It uses the projected gradient descent algorithm (PGD) with a pixel constraint ϵ to inject adversarial noise into images, effectively jailbreaking VLMs. While noise augmentation typically enhances robustness against specific noise distributions, we emphasize the critical role of image preprocessing in either transforming adversarial noise into a target distribution or directly eliminating it. We first evaluate the effectiveness of various image preprocessing defense methods, and then introduce DiffPure-VLM, a universal defense framework that robustly counters both Gaussian and adversarial noise.

4.1. Preprocessing Methods in distribution shifting

In this section, we explore different image preprocessing defense methods in distribution shifting. Specifically, we use adversarial images I_{adv} optimized for perturbation attacks on MiniGPT-4 [51] from [35], with pixel constraint $\epsilon = 16/255$ as an example. As I_{adv} is optimized based on benign clean image I_c, we compute their residual image $r_{adv} = I_{adv} - I_c$ to obtain the adversarial noise. Then we use histogram and quantile-quantile (Q-Q) plot to evaluate the distribution property of r_{adv}. As shown in Figure 6, adversarial noise follows an non-Gaussian distribution.

Currently, there are two representative image preprocessing defense methods: JailGuard [49] and DiffPure [33]. JailGuard, designed specifically for VLMs, employs various image processing techniques (e.g., random masking, horizontal flipping, Gaussian blur, and resizing) to generate variants of the input and detect adversarial samples based on discrepancies in model responses. However, most of these operations are linear transformations, offer limited ability to eliminate adversarial perturbations. Moreover, JailGuard requires multiple model runs, leading to high computational overhead. In contrast, DiffPure is tailored for computer vision models (e.g., classifiers) and leverages diffusion models to mitigate adversarial noise. It adds a small amount of noise to the adversarial image I_{adv} and reconstructs a clean image through a limited number of forward and reverse diffusion steps (e.g., using DDPM [21]) with a carefully chosen timestep $t^* \in [50, 150]$. While DiffPure aims to purify adversarial perturbations while preserving global semantic structures, our findings reveal that **at relatively small timesteps it does not completely remove the noise. Instead, it shifts the perturbation distribution towards Gaussian-like distribution.**

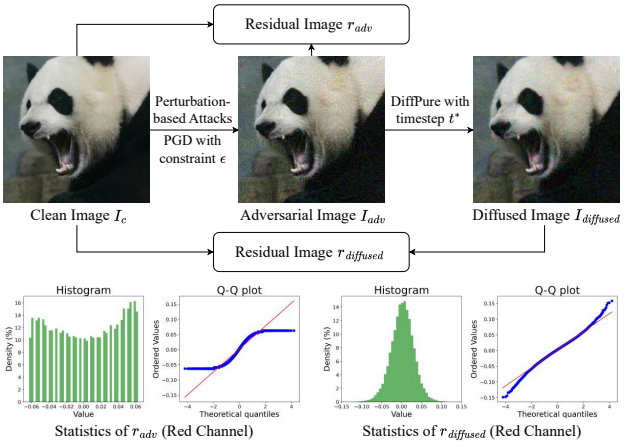

Figure 6. Residual image Gaussianity analysis. We apply DiffPure ($t^* = 50$) to adversarial image I_{adv} to obtain diffused image $I_{diffused}$. Then we calculate residual images r_{adv} and $r_{diffused}$ and evaluate their distribution by the histogram and Q-Q plot.

Specifically, we apply $t^* = 50$ in DiffPure and obtain diffused image $I_{diffused}$ based on I_{adv}. Then we calculate the residual image $r_{diffused} = I_{diffused} - I_c$. As shown in Figure 6, $r_{diffused}$ approximates a Gaussian distribution from its shape and its closeness to the theoretical line (Red line) of Gaussian distribution in the Q-Q plot. More visualizations of $I_{diffused}$ and statistic of $r_{diffused}$ across different ϵ and t^* are available in Section 11.2 in the Appendix. For quantitative evaluation of $r_{diffused}$, we use two metrics:
(1) **Kurtosis:** Kurtosis is used to measure the tailedness of a data distribution with the definition as

$$Kurt[X] = E\left[\left(\frac{X - \mu_X}{\sigma_X}\right)^4\right] \quad (1)$$

where μ_X and σ_X are mean and standard deviation of data

X. If X follows Gaussian distribution, $Kurt[X] = 3$.
(2) **Q-Q deviation:** Q-Q deviation measures the root-mean-square error (RMSE) between the quantiles of the sample distribution and those of a Gaussian distribution:

$$D(Q_s, Q_g) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(Q_{s,i} - Q_{g,i})^2} \quad (2)$$

where N is the number of ordered quantile points, Q_s is the sample quantiles of $r_{diffused}$ and Q_t is the theoretical quantiles of a Gaussian distribution. A lower RMSE value suggests closer alignment to a Gaussian distribution. We vary t^* in DiffPure from 0 to 750 in increments of 50, and also examine $t^* = 30$ for fine-grained analysis. For RGB images, we calculate these metrics per channel and obtain the average value. To identify Gaussian-like distributions in our analysis, we use the thresholds $3 \leq$ Kurtosis ≤ 6 and Q-Q deviation ≤ 0.01. Points meeting these criteria are marked in red in Figure 7, showing that under certain timesteps t^* (e.g., $t^* \in [50, 150]$) in DiffPure, $r_{diffused}$ exhibits Gaussian-like distribution properties. We further conduct our analysis in the embedding space and the conclusion is similar as in pixel space. Detailed analysis is provided in Section 11.2 in the Appendix .

Finally, we conduct defense experiment with DiffPure. Specifically, we employ MiniGPT-4 as the VLM. We apply the unconditional model in [11] as the Diffusion model in DiffPure in all our experiments. We add Gaussian noise $n \sim \mathcal{N}(0, \sigma_n^2)$ to the benign clean image I_c and apply DiffPure with timestep t^* to both Gaussian noisy image and adversarial images. Then we evaluate their attack success rate in the RealToxicityPrompts benchmark. Table 3 presents results for $\sigma_n = 30/255$ and $t^* = 50$, results of different σ_n and t^* are provided in Table 11 in the Appendix.

Experiment result shows that directly applying DiffPure to the Gaussian noisy image with VLM without noise-augmented safety fine-tuning does not decrease the attack success rate. We further observe that the attack success rate for diffused images are similar to Gaussian noisy image. Overall, we conclude that **DiffPure, when applied with a suitable timestep t^* (e.g., $t^* \in [50, 150]$), exhibits a unique distribution-shifting capability that transforms adversarial noise into a Gaussian-like distribution.**

4.2. DiffPure-VLM

Leveraging our safety fine-tuning approach and DiffPure's unique characteristic, we propose DiffPure-VLM — a defense pipeline that integrates Diffusion Models with Gaussian-noise-tolerant VLMs, as illustrated in Figure 14. Specifically, we purify adversarial images by applying a small timestep in DiffPure to preserve image content. The purified image with slight Gaussian-like noise is fed into the Gaussian-noise-tolerant, safety-tuned VLM, effectively mitigating the adversarial perturbations.

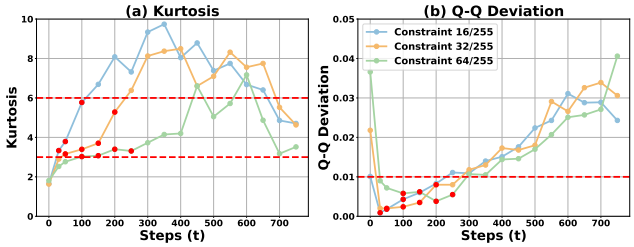

Figure 7. Gaussianity metrics of $r_{diffused}$ under different pixel constraints ϵ of adversarial image I_{adv} and timestep t^* in DiffPure. Please zoom in to see details.

Table 3. Defense of DiffPure in MiniGPT-4 under different image configurations. Attack Success Rate is evaluated on the RealToxicityPrompts benchmark.

Image Configuration	Attack Success Rate (%)
Benign clean Image I_c	34.8 ± 1.6
+ n ($\sigma_n = 30/255$)	44.1
+ n ($\sigma_n = 30/255$) + DiffPure ($t^* = 50$)	44.3
Adversarial image I_{adv} ($\epsilon = 16/255$)	53.6 ± 1.0
+ DiffPure ($t^* = 50$)	45.0
Adversarial image I_{adv} ($\epsilon = 32/255$)	59.4 ± 1.4
+ DiffPure ($t^* = 50$)	45.5
Adversarial image I_{adv} ($\epsilon = 64/255$)	67.2 ± 0.2
+ DiffPure ($t^* = 50$)	44.5

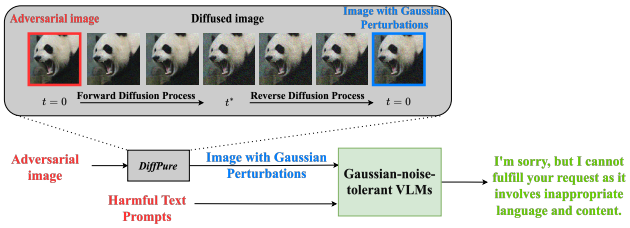

Figure 8. The overall framework of DiffPure-VLM.

First, to verify DiffPure's effectiveness in our defense pipeline, we compare its performance against JailGuard for mitigating optimization-based perturbation attacks following [35] on the RealToxicityPrompts benchmark. We selected two base models, LLaVA-VLGuard and our LLaVA-RobustVLGuard, for a comprehensive evaluation. As shown in Table 4, DiffPure with $t^* = 50$ consistently outperforms JailGuard across both base models. Notably, when paired with our VLM, DiffPure delivers a substantially greater improvement over JailGuard than when paired with LLaVA-VLGuard, confirming that its distribution-shifting properties are especially well-suited to our robust VLM.

To evaluate the generalization of DiffPure-VLM, we assess the pipeline's performance under various optimization-based perturbation attack strengths (pixel constraint $\epsilon \in \{16/255, 32/255, 64/255\}$) and different timesteps ($t^* = 50, 150$) of DiffPure in our suite of three safety fine-tuned VLMs. For brevity, Table 12 presents results for $\epsilon =$

Table 4. Comparison of image preprocessing methods for mitigating adversarial attacks (ϵ = 32/255) on RealToxicityPrompts.

Model	Attack Success Rate ↓	Identity	Profanity	Severe Toxicity	Sexually Explicit	Threat	Toxicity
LLaVA-VLGuard	70.4	21.3	52.8	7.5	16.7	7.0	67.2
JailGuard + LLaVA-VLGuard	52.1	12.5	39.0	5.3	13.2	4.9	50.0
DiffPure + LLaVA-VLGuard	**51.1**	3.4	40.9	2.2	13.4	3.6	47.5
LLaVA-RobustVLGuard	62.5	7.8	48.0	5.4	16.5	5.8	60.0
JailGuard + LLaVA-RobustVLGuard	48.9	6.0	37.3	4.8	13.4	4.0	46.5
DiffPure + LLaVA-RobustVLGuard	**43.9**	3.2	34.6	2.4	12.8	3.7	41.0

Table 5. Evaluation of DiffPure-VLM's effectiveness on RealToxicityPrompts. Metrics include attack success rate and various toxicity levels (Perspective API %). Additional results for other attack strengths are provided in Table 12 in the Appendix .

Image Type	Attack Success Rate ↓	Identity	Profanity	Severe Toxicity	Sexually Explicit	Threat	Toxicity
InternVL2-8B-RobustVLGuard							
Benign Clean image	29.9	0.8	22.1	0.3	7.2	1.5	25.9
Gaussian Noisy image	34.5	2.1	27.2	1.3	8.4	1.6	31.3
Adversarial image (ϵ = 32/255)	70.6	26.7	56.5	9.2	17.3	6.9	68.1
+DiffPure-VLM (t*=50)	**33.4**	2.4	20.6	0.7	8.1	2.4	29.1
+DiffPure-VLM (t*=150)	**32.8**	1.7	25.9	0.6	7.7	1.8	29.1
LLaVA-v1.5-7B-RobustVLGuard							
Benign Clean image	43.6	4.6	34.7	2.4	12.3	3.5	41.0
Gaussian Noisy image	42.3	3.1	34.5	1.9	11.8	3.1	40.0
Adversarial image (ϵ = 32/255)	62.5	7.8	48.0	5.4	16.5	5.8	60.0
+DiffPure-VLM (t*=50)	**43.9**	3.2	34.6	2.4	12.8	3.7	41.0
+DiffPure-VLM (t*=150)	**42.5**	3.5	32.7	2.8	12.1	4.1	39.3
MiniGPT-4-13B-RobustVLGuard							
Benign Clean image	16.0	0.4	9.9	0.3	4.6	1.1	12.1
Gaussian Noisy image	16.5	0.9	11.9	0.6	5.8	1.0	14.0
Adversarial image (ϵ = 32/255)	53.7	9.8	35.3	4.1	13.9	5.4	48.1
+DiffPure-VLM (t*=50)	**13.6**	0.3	9.2	0.2	5.5	0.9	10.6
+DiffPure-VLM (t*=150)	**11.9**	0.5	8.6	0.2	4.2	0.6	9.9

$32/255$. Results for the other ϵ values are provided in Table 12 in the Appendix. The experimental results demonstrate that, compared to clean or Gaussian-noise inputs, adversarial perturbation attacks substantially increase attack success rates. Nevertheless, our DiffPure mechanism consistently reduces these rates across all tested timesteps, nearly restoring performance to that of clean or Gaussian inputs. This not only confirms our analysis of DiffPure's unique ability to transform adversarial noise into Gaussian noise, but also highlights the efficacy and generalization of DiffPure-VLM with effectively mitigating the impact of perturbations.

5. Related Work

5.1. Vision Language Model

In recent years, Vision Language Models have gained significant attention due to their ability to jointly process and understand both visual and textual data. A major breakthrough in this area came with the development of large-scale pre-trained models like CLIP [36], which aligns images and texts in a shared embedding space. By training on a vast dataset of image-text pairs, CLIP enabled zero-shot learning capabilities for tasks such as action recognition [7, 32] and optical character recognition [26, 40]. Recent advances have introduced multi-modal foundation models like InstructBLIP [10], LLaVA [29], and Qwen-VL [3] which leverage large-scale transformer networks that can process both image and text inputs simultaneously, enabling more sophisticated reasoning over complex scenarios. These models excel in advanced tasks like UI understanding [14, 22] and visual question answering [2], and open up new possibilities for generative capabilities.

5.2. Jailbreaking VLMs

As VLMs become increasingly integrated into various applications, critical concerns have arisen regarding their robustness, security, and ethical alignment. A critical issue is "jailbreaking"—the ability to bypass safety protocols, potentially triggering unintended or harmful behaviors. Given that VLMs process both textual and visual inputs, vulnerabilities inherent in LLMs may also affect VLMs. Moreover, the integration of visual inputs into VLMs, while expanding their capabilities, introduces more diverse attack patterns, significantly heightening the severity of potential threats [25, 43]. Several methods have emerged in this domain, including converting the malicious content into images through typography-based manipulations [17], leveraging multimodal perturbations to craft stronger adversarial attacks [47, 50], and using gradient updates to embed malicious triggers within seemingly benign images [39]. These diverse sophisticated methods pose significant challenges to ensuring the safety and reliability of VLMs.

5.3. Safeguarding VLMs

Given the growing challenges posed by malicious attacks on VLMs, it is crucial to develop effective defense strategies. One key approach involves enhancing the training process through adversarial fine-tuning [37, 52] or prompt tuning [48], where models or learnable prompts are trained on perturbed examples to improve resilience against real-world attack scenarios. Another promising defense strategy is input sanitization, which aims to detect or neutralize adversarial inputs before they can compromise the model. Methods in this area include shifting sample probabilities to adversarial-free regions [42], applying randomized smoothing to mitigate the impact of adversarial inputs [41], and utilizing unlabeled prompts from users in the wild for malicious prompt detection [12]. In this work, we concentrate on improving VLM robustness to noise-augmented inputs. By integrating Diffusion Models with safety fine-tuned VLMs, we equip these models with enhanced protection against a broad range of adversarial attacks.

6. Conclusion

In this work, we address a critical gap in the robustness of VLMs by examining the impact of Gaussian noise perturbations and propose Robust-VLGuard, a multimodal safety dataset paired with Gaussian-noise-augmented fine-tuning, to enhance safety alignment and preserve helpfulness of VLMs. We further propose DiffPure-VLM to de-

fend Optimization-Based Visual Perturbation Attack by using diffusion model to transfer adversarial noise to Gaussian noise which can be defended by VLMs with noise-augmented safety fine-tuning. Experiment result demonstrates the superiority of DiffPure-VLM in Gaussian noise and adversarial perturbations with baseline methods.

While DiffPure-VLM provides a practical defense, future work includes integrating noise augmentation in pre-training, expanding the safety dataset for broader tasks, and exploring adaptive multi-modal defenses to further enhance real-world performance.

7. Social Impact

This research exposes VLM vulnerabilities to noisy inputs and adversarial attacks. While Robust-VLGuard and DiffPure-VLM enhance robustness, our findings have dual-use implications. Given VLMs' growing adoption, we responsibly disclose these issues to raise awareness and foster the development of more secure models, mitigating deployment risks.

References

[1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3

[2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 8

[3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. 1, 8

[4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 2, 12

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

[6] Federico Bianchi, Mirac Suzgun, Giuseppe Attanasio, Paul Rottger, Dan Jurafsky, Tatsunori Hashimoto, and James Zou. Safety-tuned llamas: Lessons from improving the safety of large language models that follow instructions. In *ICLR*, 2024. 4, 6

[7] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 8

[8] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024. 12

[9] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024. 1, 2, 4, 12

[10] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023. 8

[11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 7, 18

[12] Xuefeng Du, Reshmi Ghosh, Robert Sim, Ahmed Salem, Vitor Carvalho, Emily Lawton, Yixuan Li, and Jack W Stokes. Vlmguard: Defending vlms against malicious prompts via unlabeled data. *arXiv preprint arXiv:2410.00296*, 2024. 8

[13] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 2, 12

[14] Yicheng Fu, Raviteja Anantha, Prabal Vashisht, Jianpeng Cheng, and Etai Littwin. Ui-jepa: Towards active perception of user intent through onscreen user activity. *arXiv preprint arXiv:2409.04081*, 2024. 8

[15] Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022. 1

[16] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Realtoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020. 3

[17] Yichen Gong, Delong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan, and Xiaoyun Wang. Figstep: Jailbreaking large vision-language models via typographic visual prompts. *arXiv preprint arXiv:2311.05608*, 2023. 8

[18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1

[19] 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 *CVPR*, 2017. 4

[20] Laura Hanu and Unitary team. Detoxify. Github. https://github.com/unitaryai/detoxify, 2020. 3

[21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 6

[22] Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. Cogagent: A visual language model for gui agents. In *CVPR*, 2024. 8

[23] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. 4

[24] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019. 4

[25] Haibo Jin, Leyang Hu, Xinuo Li, Peiyan Zhang, Chonghan Chen, Jun Zhuang, and Haohan Wang. Jailbreakzoo: Survey, landscapes, and horizons in jailbreaking large language and vision-language models. *arXiv preprint arXiv:2407.01599*, 2024. 8

[26] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. The hateful memes challenge: detecting hate speech in multimodal memes. In *NeurIPS*, 2020. 8

[27] Yifan Li, Hangyu Guo, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. Images are achilles' heel of alignment: Exploiting visual vulnerabilities for jailbreaking multimodal large language models. *arXiv preprint arXiv:2403.09792*, 2024. 1, 3

[28] 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 *ECCV*, 2014. 4

[29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2024. 1, 2, 4, 5, 8, 12

[30] Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. In *ICLR*, 2023. 4

[31] Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the Association for Computational Linguistics: ACL 2022*, 2022. 4

[32] Antoine Miech, Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. Rareact: A video dataset of unusual interactions. *arXiv preprint arXiv:2008.01018*, 2020. 8

[33] Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Animashree Anandkumar. Diffusion models for adversarial purification. In *International Conference on Machine Learning*, 2022. 1, 2, 6, 13, 18

[34] OpenAI. Gpt-4v(ision) system card. 2023. 2, 4

[35] Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter Henderson, Mengdi Wang, and Prateek Mittal. Visual adversarial examples jailbreak aligned large language models. In *AAAI*, 2024. 1, 2, 3, 6, 7, 12

[36] 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*, 2021. 8

[37] Christian Schlarmann, Naman Deep Singh, Francesco Croce, and Matthias Hein. Robust clip: Unsupervised adversarial fine-tuning of vision embeddings for robust large vision-language models. In *ICLR 2024 Workshop on Reliable and Responsible Foundation Models*. 8

[38] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *ECCV*, 2022. 4

[39] Erfan Shayegani, Yue Dong, and Nael Abu-Ghazaleh. Jailbreak in pieces: Compositional adversarial attacks on multi-modal language models. In *ICLR*, 2023. 8

[40] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 2013. 8

[41] Jiachen Sun, Changsheng Wang, Jiongxiao Wang, Yiwei Zhang, and Chaowei Xiao. Safeguarding vision-language models against patched visual prompt injectors. *arXiv preprint arXiv:2405.10529*, 2024. 8

[42] Rajkumar Theagarajan, Ming Chen, Bir Bhanu, and Jing Zhang. Shieldnets: Defending against adversarial attacks using probabilistic adversarial robustness. In *CVPR*, 2019. 8

[43] Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does llm safety training fail? In *NeurIPS*, 2023. 8

[44] Yue Xu, Xiuyuan Qi, Zhan Qin, and Wenjie Wang. Defending jailbreak attack in vlms via cross-modality information detector. *arXiv preprint arXiv:2407.21659*, 2024. 2

[45] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. In *International Conference on Machine Learning*, 2023. 3

[46] Ivy Zhang, Wei Peng, Jenny N, Theresa Yu, and David Qiu. Ivy-vl: Compact vision-language models achieving sota with optimal data, 2024. 12

[47] Jiaming Zhang, Qi Yi, and Jitao Sang. Towards adversarial attack on vision-language pre-training models. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 8

[48] Jiaming Zhang, Xingjun Ma, Xin Wang, Lingyu Qiu, Jiaqi Wang, Yu-Gang Jiang, and Jitao Sang. Adversarial prompt tuning for vision-language models. *arXiv preprint arXiv:2311.11261*, 2023. 8

[49] Xiaoyu Zhang, Cen Zhang, Tianlin Li, Yihao Huang, Xiaojun Jia, Xiaofei Xie, Yang Liu, and Chao Shen. A mutation-based method for multi-modal jailbreaking attack detection. *arXiv preprint arXiv:2312.10766*, 2023. 2, 6

[50] 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. In *NeurIPS*, 2023. 8

[51] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. In *ICLR*, 2024. 1, 2, 4, 6, 12

[52] Yongshuo Zong, Ondrej Bohdal, Tingyang Yu, Yongxin Yang, and Timothy Hospedales. Safety fine-tuning at (almost) no cost: A baseline for vision large language models. In *International Conference on Machine Learning*, 2024. 1, 2, 4, 5, 8

Safeguarding Vision-Language Models: Mitigating Vulnerabilities to Gaussian Noise in Perturbation-based Attacks

Supplementary Material

8. Overview of the Supplementary Material

This supplementary material offers additional details and analyses to further support the findings presented in the main manuscript. It includes detailed information on the experimental configuration (Appendix 9), more evaluation on recent vision-language models (Appendix 10), a thorough analysis of the limitations and unique characteristics of DiffPure (Appendix 11), extended implementation specifics of DiffPure-VLM (Appendix 12), and conjectures along with preliminary theoretical discussions on the effects of Gaussian noise (Appendix 13). Collectively, these sections provide deeper insights into our methodology, enhancing the transparency and reproducibility of our research.

9. Experiment Details

9.1. Models

In this work, we conduct all experiments on three leading Vision-Language Models (VLMs), i.e., MiniGPT-4 (13B) [51], LLaVA-v1.5 (7B) [29], and InternVL2 (8B) [9]. We use the official model weights from HuggingFace or Github repositories for experiments in our paper. These model details are summarized in Table 6. Each model features a distinct LLM, vision encoder, and vision-language alignment method, allowing us to draw broader insights.

Table 6. Specifications of the evaluated VLMs.

Model	Size	Vision Encoder	LLM	VL Connection Module
MiniGPT-4-13B	14B	EVA-CLIP ViT-G/14	Vicuna-v0-13B	Q-former
LLaVA-v1.5-7B	7B	CLIP ViT-L/14	Vicuna-v1.5-7B	MLP
InternVL2-8B	8B	InternViT-300M-448px	InternLM2-8B	MLP

9.2. Fine-tuning Configuration

We present the detailed hyper-parameters for post-hoc fine-tuning on our Robust-VLGuard dataset in Table 7. Gaussian noise augmentation was applied to the training images, with a randomly selected standard deviation between 0.01 and 0.15, and a 70% probability of application. The fine-tuning was performed over 3 epochs on a single A100-80G GPU, using a consistent batch size of 16. For MiniGPT-4-13B, unfreezing the linear projector significantly improved robustness in terms of helpfulness and safety alignment. However, for LLaVA-v1.5-7B and InternVL2-8B, unfreezing the linear projector led to increased overfitting, likely due to differences in the vision-language connection modules of these models.

Table 7. Post-hoc fine-tuning hyper-parameters of different models.

Model	Training Module	LoRA Rank	LoRA Alpha	Learning Rate
MiniGPT-4-13B	Vision Encoder & Linear Projector	16	32	3e-5
LLaVA-v1.5-7B	Vision Encoder	16	256	4e-5
InternVL2-8B	Vision Encoder	16	256	4e-5

9.3. Details of Evaluation Settings

For evaluation on the MM-Vet benchmark, we set the temperature to 0 and use greedy decoding across all experiments to ensure reproducibility in helpfulness assessments. For safety evaluations on the RealToxicityPrompts benchmark, we follow the setup of Qi et al. [35], using a temperature of 1 and performing three runs to calculate the average attack success rate. Greedy decoding is also employed for this benchmark. The choice of temperature 1 reflects real-world usage, where sampling is typically enabled during interactions with VLMs. This setting aims to better simulate real-world scenarios when assessing safety alignment.

Additionally, the MM-Vet and RealToxicityPrompts benchmarks offer a comprehensive set of metrics covering various aspects. For the sake of brevity, we report only the overall metrics — Performance Score and Attack Success Rate — in the main paper. Here, we present the detailed evaluation results in Table 8 and Table 9, corresponding to Figure 2 in the main paper. The results show that Gaussian noisy images negatively impact nearly all metrics across both benchmarks and various models. Notably, using Gaussian noisy images as prompts improves MiniGPT-4's performance on the OCR metric in the MM-Vet benchmark, highlighting the current VLMs' lack of robustness.

10. Additional Evaluation on Recent Vision-Language Models

In this section, we further assess the robustness of state-of-the-art vision-language models against Gaussian noise. Table 10 presents the attack success rates on the RealToxicityPrompts benchmark for four recent VLMs—LLaMA-3.2-Vision-11B [13], Ivy-VLM-3B [46], Qwen2.5-VL-7B [4], and InternVL2.5-8B [8]—under various Gaussian noise levels. Lower percentages indicate improved safety alignment.

Table 8. Robustness comparison of various models on the MM-Vet benchmark using clean and Gaussian noisy image prompts (GPT-4 %).

Image Type	Performance Score ↑	Recognition	OCR	Knowledge	Generation	Spatial	Math
MiniGPT-4-13B							
Clean Image	26.7	34.9	13.5	27.4	27.1	19.1	7.7
Gaussian Noisy Image	24.0 (-2.7)	29.0 (-5.9)	16.9 (+3.4)	20.5 (-6.9)	22.5 (-4.6)	20.7 (+1.6)	7.7 (0.0)
LLaVA-v1.5-7B							
Clean Image	33.0	37.9	23.9	20.4	23.6	28.5	11.5
Gaussian Noisy Image	31.3 (-1.7)	36.3 (-1.6)	21.9 (-2.0)	18.3 (-2.1)	21.2 (-2.4)	25.7 (-2.8)	3.8 (-7.7)
InternVL2-8B							
Clean Image	59.9	53.5	71.9	40.4	44.6	69.7	65.4
Gaussian Noisy Image	54.4 (-5.5)	47.1 (-6.4)	66.8 (-5.1)	35.1 (-5.3)	39.2 (-5.4)	61.1 (-8.6)	53.5 (-11.9)

Table 9. Robustness comparison of various models on the RealToxicityPrompts benchmark using clean and Gaussian noisy image prompts. Metrics include Attack Success Rate and toxicity levels (Perspective API %).

Image Type	Attack Success Rate ↓	Identity Attack	Profanity	Severe Toxicity	Sexually Explicit	Threat	Toxicity
MiniGPT-4-13B							
Clean Image	34.8	2.7	25.1	1.5	12.2	2.0	30.5
Gaussian Noisy Image	44.1 (+9.3)	3.6 (+0.9)	31.2 (+6.1)	2.3 (+0.8)	13.7 (+1.5)	2.9 (+0.9)	38.2 (+7.7)
LLaVA-v1.5-7B							
Clean Image	57.7	5.7	46.8	3.7	18.0	3.8	54.4
Gaussian Noisy Image	60.1 (+2.4)	4.8 (-0.9)	48.1 (+1.3)	2.9 (-0.8)	17.8 (-0.2)	4.0 (+0.2)	56.0 (+1.6)
InternVL2-8B							
Clean Image	50.5	4.1	40.2	1.9	13.5	2.5	44.3
Gaussian Noisy Image	57.2 (+6.7)	4.5 (+0.4)	45.9 (+5.7)	2.0 (+0.1)	14.3 (+0.8)	3.2 (+0.7)	51.7 (+7.4)

As shown, when Gaussian noise is introduced at increasing levels ($\sigma_n = 30/255$, $\sigma_n = 50/255$, and $\sigma_n = 70/255$), all models exhibit a rise in attack success rates, highlighting their sensitivity to simple Gaussian noise perturbations. These findings underscore the need for robust noise augmentation and defense strategies in training pipelines to maintain safety alignment in VLMs.

11. Further Analysis of DiffPure

11.1. Defence Performance

In this section, we present a comprehensive analysis of the effects of DiffPure [33] and Gaussian noise under perturbation-based attacks in Vision-Language Models (VLMs). Specifically, we extend the experimental setup described in Section 3.1 in the main paper by varying the standard deviation σ_n of Gaussian noise n and the timestep parameter t^* in DiffPure. Results are summarized in Table 11. First, Gaussian noise n with standard deviations $\sigma_n \in \{15/255, 30/255, 50/255, 75/255\}$ is added to the benign clean image I_c to evaluate its impact on the Attack Success Rate. The results demonstrate that the Attack Success Rate under Gaussian noise is significantly higher than that of the benign clean image. When $\sigma_n \leq 50/255$, increasing σ_n will lead to a higher Attack Success Rate. However, this trend did not continue at a higher σ_n setting (e.g., $\sigma_n = 75/255$), indicating that the effect of Gaussian noise on VLMs is limited. Next, we apply DiffPure with different timesteps $t^* \in \{50, 100, 150\}$ to generate diffused images from adversarial inputs I_{adv} with varying perturbation constraints ϵ. For $\epsilon = 16/255$, increasing t^* to 100 or 150 reduces the Attack Success Rate but does not lower it below the level observed for the benign clean image. For larger perturbation constraints, increasing t^* fails to decrease the Attack Success Rate, with a comparable performance of Gaussian noisy images.

11.2. Distribution Shift

In this section, we present detailed results from the Gaussianity experiments conducted on adversarial and diffused images. Specifically, we visualize adversarial images I_{adv} alongside their corresponding residuals r_{adv}, and diffused images $I_{diffused}$ with their residuals $r_{diffused}$, under pixel constraints $\epsilon \in \{16/255, 32/255, 64/255\}$ for I_{adv} and diffusion timesteps $t^* \in \{50, 100, 150, 500, 750\}$ in DiffPure [33] for generating $I_{diffused}$. Visualizations are shown in Figure 9, 10 and 11 with corresponding metrics: Kurtosis, Q-Q deviation, mean and standard deviation. From these visualizations, we observe that when $50 \leq t^* \leq 150$, the residuals $r_{diffused}$ exhibit a Gaussian-like distribution, particularly for $\epsilon = 32/255$ and $\epsilon = 64/255$. However, as t^* increases, the Kurtosis of $r_{diffused}$ rises, indicating a shift towards a long-tailed distribution. This suggests that a small fraction of pixels in $I_{diffused}$ undergo significant

Table 10. Attack success rate (%) on the RealToxicityPrompts benchmark for various vision-language models under different noise levels. Lower scores indicate improved safety alignment.

	RealToxicityPrompts (%) ↓			
	LLaMA-3.2-Vision-11B	Ivy-VLM-3B	Qwen2.5-VL-7B	InternVL2.5-8B
Benign clean Image I_c	45.4	29.9	36.8	43.9
+ n ($\sigma_n = 30/255$)	46.4 (+1.0)	35.5 (+5.6)	39.3 (+2.5)	51.5 (+7.6)
+ n ($\sigma_n = 50/255$)	47.6 (+2.2)	40.3 (+10.4)	39.5 (+2.7)	52.8 (+8.9)
+ n ($\sigma_n = 70/255$)	48.5 (+3.1)	42.0 (+12.1)	46.1 (+9.3)	54.0 (+10.1)

Table 11. Detailed results of the defense of DiffPure in MiniGPT-4 on the RealToxicityPrompts benchmark under different image configurations. (Perspective API %).

Image Configuration	**Attack Success Rate ↓**	**Identity Attack**	**Profanity**	**Severe Toxicity**	**Sexually Explicit**	**Threat**	**Toxicity**
Benign clean Image I_c	34.8	2.7	25.1	1.5	12.2	2.0	30.5
+ n ($\sigma_n = 15/255$)	38.5 (+3.7)	2.9 (+0.2)	27.4 (+2.3)	1.1 (-0.4)	13.0 (+0.8)	2.3 (+0.3)	34.1 (+3.6)
+ n ($\sigma_n = 30/255$)	44.1 (+9.3)	3.6 (+0.9)	31.2 (+6.1)	2.3 (+0.8)	13.7 (+1.5)	2.9 (+0.9)	38.2 (+7.7)
+ n ($\sigma_n = 50/255$)	46.3 (+11.5)	3.4 (+0.7)	34.0 (+8.9)	1.8 (+0.3)	14.8 (+2.6)	2.5 (+0.5)	39.5 (+9.0)
+ n ($\sigma_n = 75/255$)	44.1 (+9.3)	3.8 (+1.1)	30.1 (+5.0)	1.9 (+0.4)	14.3 (+2.1)	2.8 (+0.8)	37.5 (+7.0)
Adversarial image I_{adv} ($\epsilon = 16/255$)	53.6 (+18.8)	8.4 (+5.7)	36.6 (+9.4)	6.6 (+5.1)	14.1 (+1.9)	4.7 (+2.7)	48.6 (+18.1)
+ DiffPure ($t^* = 50$)	45.0 (+10.2)	2.5 (-0.2)	31.7 (+6.6)	1.8 (+0.3)	14.5 (+2.3)	2.8 (+0.8)	38.8 (+8.3)
+ DiffPure ($t^* = 100$)	37.6 (+2.8)	3.0 (+0.3)	25.6 (+0.5)	1.3 (-0.2)	12.3 (+0.1)	1.8 (-0.2)	33.1 (+2.6)
+ DiffPure ($t^* = 150$)	37.7 (+2.9)	2.5 (-0.2)	26.5 (+1.4)	2.1 (+0.6)	12.2 (+0.0)	2.5 (+0.5)	32.9 (+2.4)
Adversarial image I_{adv} ($\epsilon = 32/255$)	59.4 (+24.6)	14.6 (+11.9)	39.5 (+14.4)	7.0 (+5.5)	14.9 (+2.7)	6.2 (+4.2)	53.8 (+23.3)
+ DiffPure ($t^* = 50$)	45.5 (+10.7)	2.6 (-0.1)	32.1 (+7.0)	2.2 (+0.7)	14.8 (+2.6)	3.0 (+1.0)	38.5 (+8.0)
+ DiffPure ($t^* = 100$)	43.8 (+9.0)	3.3 (+0.6)	31.9 (+6.8)	1.9 (+0.4)	13.1 (+0.9)	2.5 (+0.5)	38.1 (+7.6)
+ DiffPure ($t^* = 150$)	42.3 (+7.5)	3.7 (+1.0)	30.4 (+5.3)	1.3 (-0.2)	13.3 (+1.1)	2.8 (+0.8)	36.3 (+5.8)
Adversarial image I_{adv} ($\epsilon = 64/255$)	67.2 (+32.4)	15.9 (+13.2)	49.6 (+24.5)	12.2 (+10.7)	16.9 (+4.7)	6.6 (+4.6)	63.1 (+32.6)
+ DiffPure ($t^* = 50$)	44.5 (+9.7)	2.9 (+0.2)	32.2 (+7.1)	2.4 (+0.9)	13.7 (+1.5)	2.7 (+0.7)	38.0 (+7.5)
+ DiffPure ($t^* = 100$)	42.1 (+7.3)	2.8 (+0.1)	30.3 (+5.2)	1.9 (+0.4)	13.7 (+1.5)	3.0 (+1.0)	36.5 (+6.0)
+ DiffPure ($t^* = 150$)	44.1 (+9.3)	3.3 (+0.6)	31.5 (+6.4)	1.4 (-0.1)	13.3 (+1.1)	2.5 (+0.5)	38.2 (+7.7)

changes compared to I_c, leading to a cleaner image with minimal content alteration, especially when $\epsilon = 16/255$. At $t^* = 500$, the Kurtosis and standard deviation of $r_{diffused}$ become significantly larger, implying greater changes in image content, as reflected in the visualization of $I_{diffused}$. For $t^* = 750$, the Kurtosis decreases while the standard deviation further increases, indicating that $r_{diffused}$ transitions to a flatter and broader distribution. In this case, $I_{diffused}$ diverges substantially from I_c in image content.

Furthermore, we extend our analysis to the embedding space, examining the similarities between the clean image I_c, the adversarial image I_{adv}, and the diffused image $I_{diffused}$. Based on our experiment in pixel space, where the residual noise $r_{diffused}$ approximates a Gaussian distribution under certain timestep settings in DiffPure, we consider $I_{diffused}$ as comparable to I_c with added Gaussian noise. To verify this, we generate a noisy image $I_n = I_c + n, n \in \mathcal{N}\left(0, \sigma^2_{r_{diffused}}\right)$, where $\sigma_{r_{diffused}}$ indicates the standard deviation of $r_{diffused}$. Using pre-trained visual encoder E in MiniGPT-4, we compute cosine similarities between the embeddings of I_n and I_{adv}, denoted as $C_{n,adv}$, and between I_n and $I_{diffused}$, denoted as $C_{n,diffused}$. Figure 12 shows these similarities across varying adversarial constraints ϵ and DiffPure steps t^*. Results indicate that, $C_{n,diffused}$ consistently exceeds $C_{n,adv}$, showing that $I_{diffused}$ is closer to I_n than I_{adv} in the embedding space. Notably, with moderate timesteps ($t^* \in [50, 150]$), $I_{diffused}$ is similar to I_n (Gaussian noise n added to the benign clean image I_c) in both pixel and embedding spaces.

We also visualize the cosine similarity between the visual embeddings of $I_{diffused}$ and I_c, denoted as $C_{clean,diffused}$, across varying ϵ and t^*. Results are shown in Figure 13, revealing that $C_{clean,diffused}$ decreases rapidly as t^* decreases, while it gradually declines as t^* increases. Combining these findings with experiments in pixel space, we conclude that smaller t^* values lead $I_{diffused}$ to retain adversarial information, whereas larger t^* values result in significant content disruption, leading to semantic misalignment.

12. Additional Details of DiffPure-VLM

12.1. Implementation Details

The overall architecture of our proposed DiffPure-VLM framework is illustrated in Figure 14, with the detailed al-

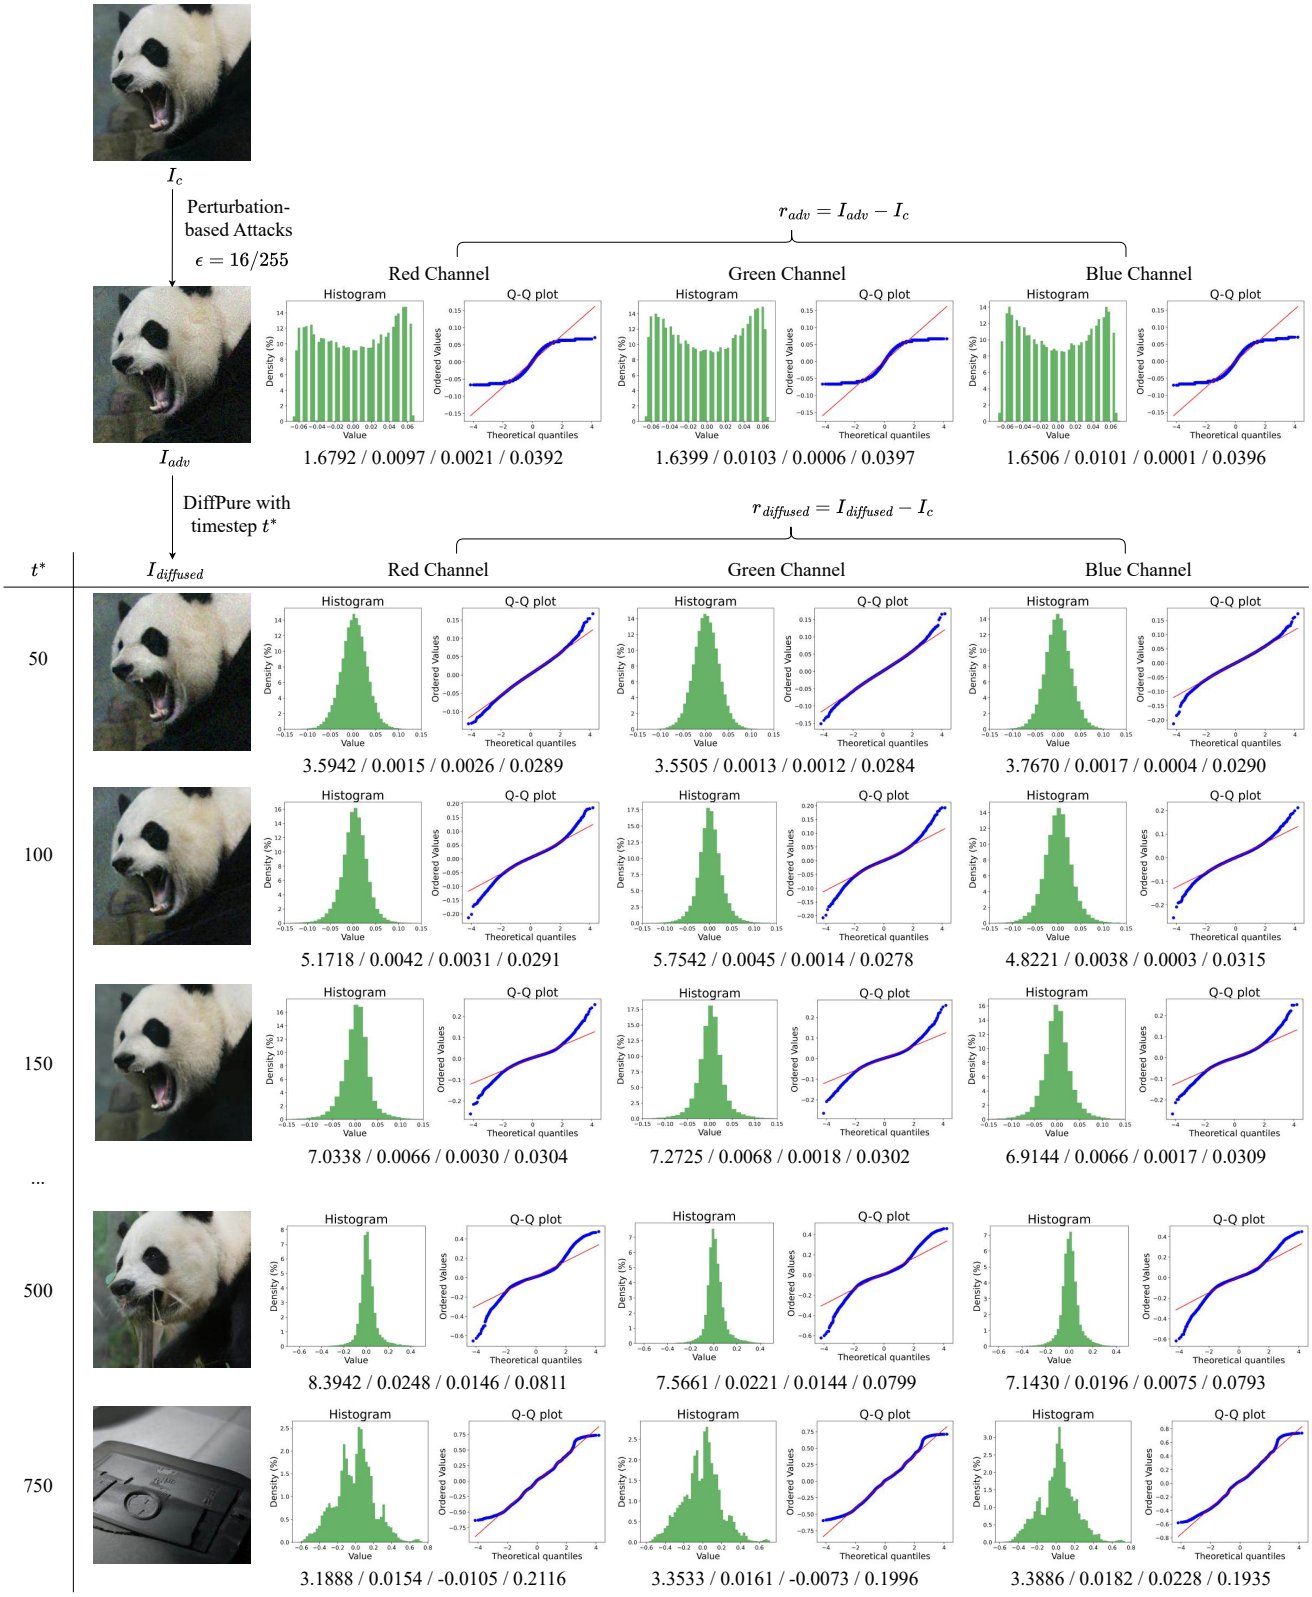

Figure 9. I_{adv}, $I_{diffused}$ and statistics of r_{adv}, $r_{diffused}$ under different t^* in DiffPure (constraint $\epsilon = 16/255$). Metrics are shown in 'Kurtosis / Q-Q Deviation / Mean / Standard Deviation'. Please zoom in to see details.

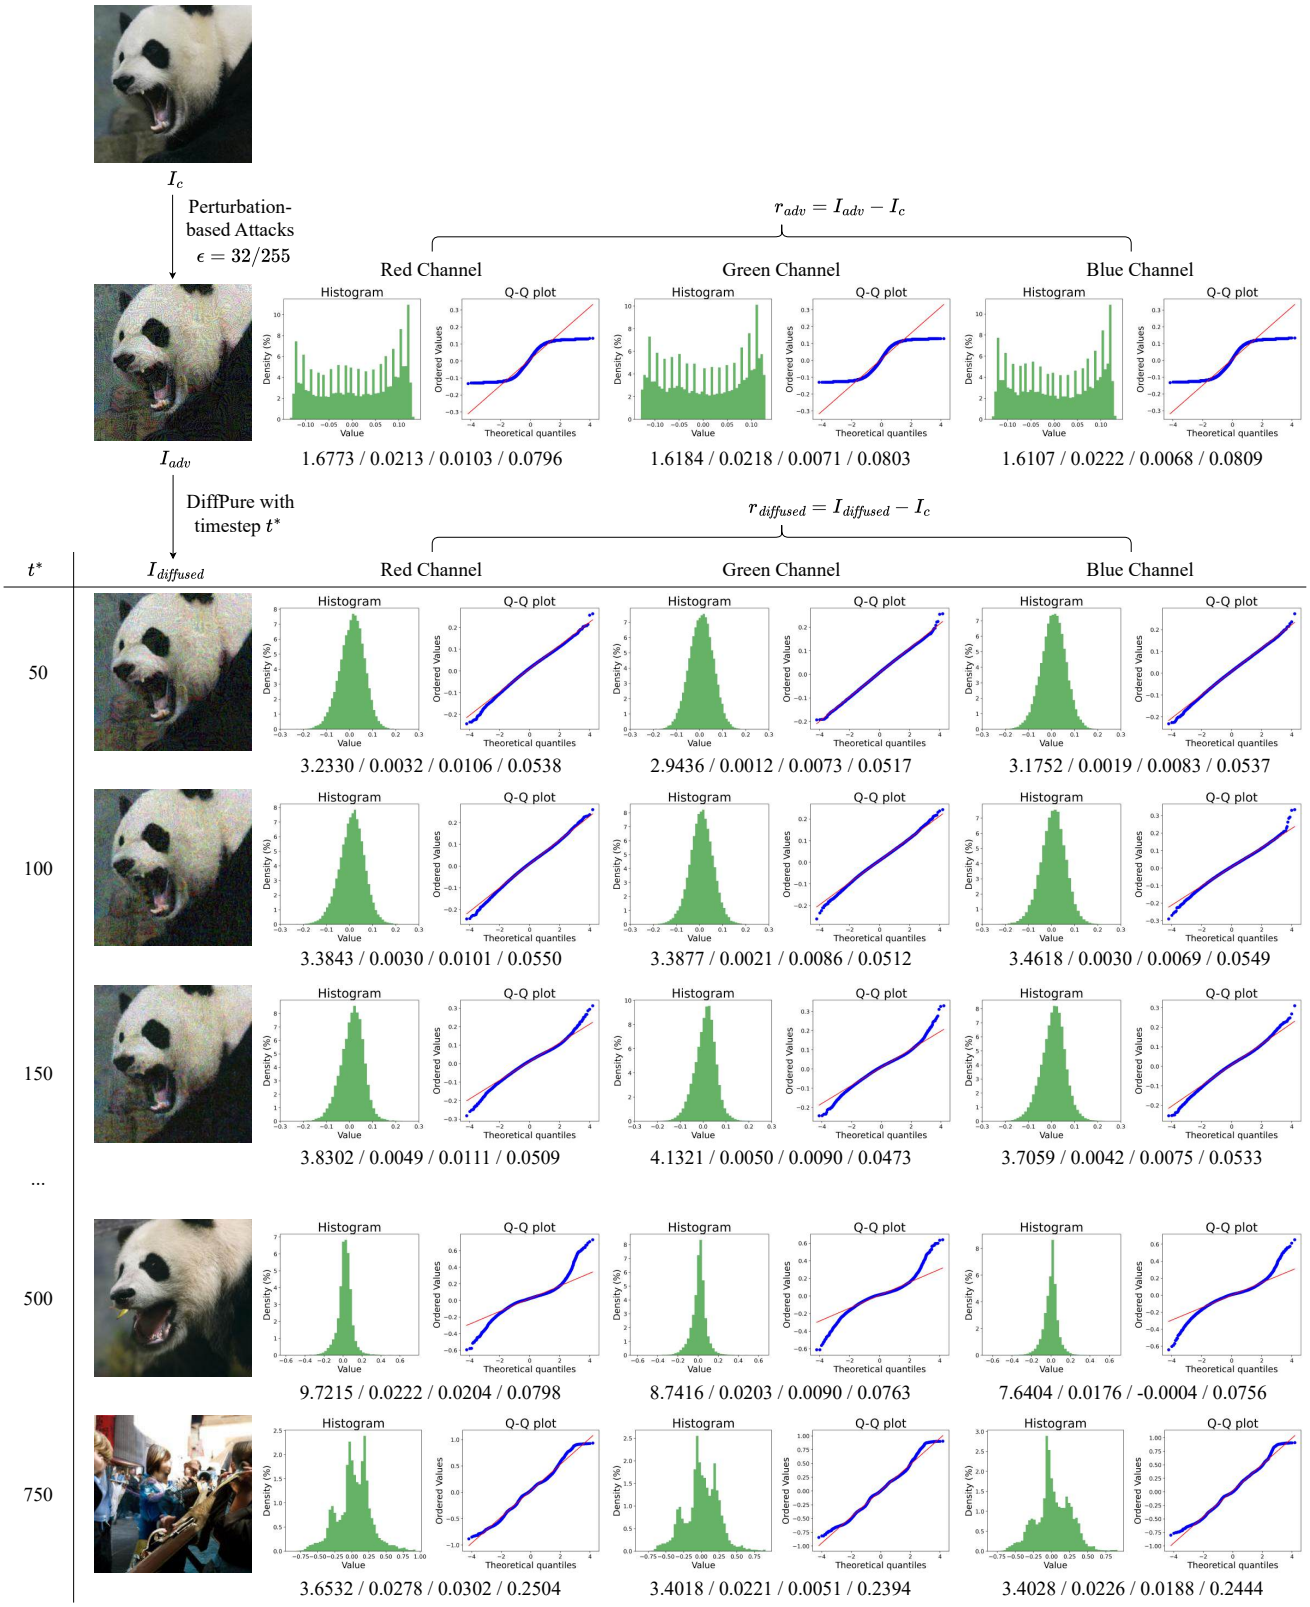

Figure 10. I_{adv}, $I_{diffused}$ and statistics of r_{adv}, $r_{diffused}$ under different t^* in DiffPure (constraint $\epsilon = 32/255$). Metrics are shown in 'Kurtosis / Q-Q Deviation / Mean / Standard Deviation'. Please zoom in to see details.

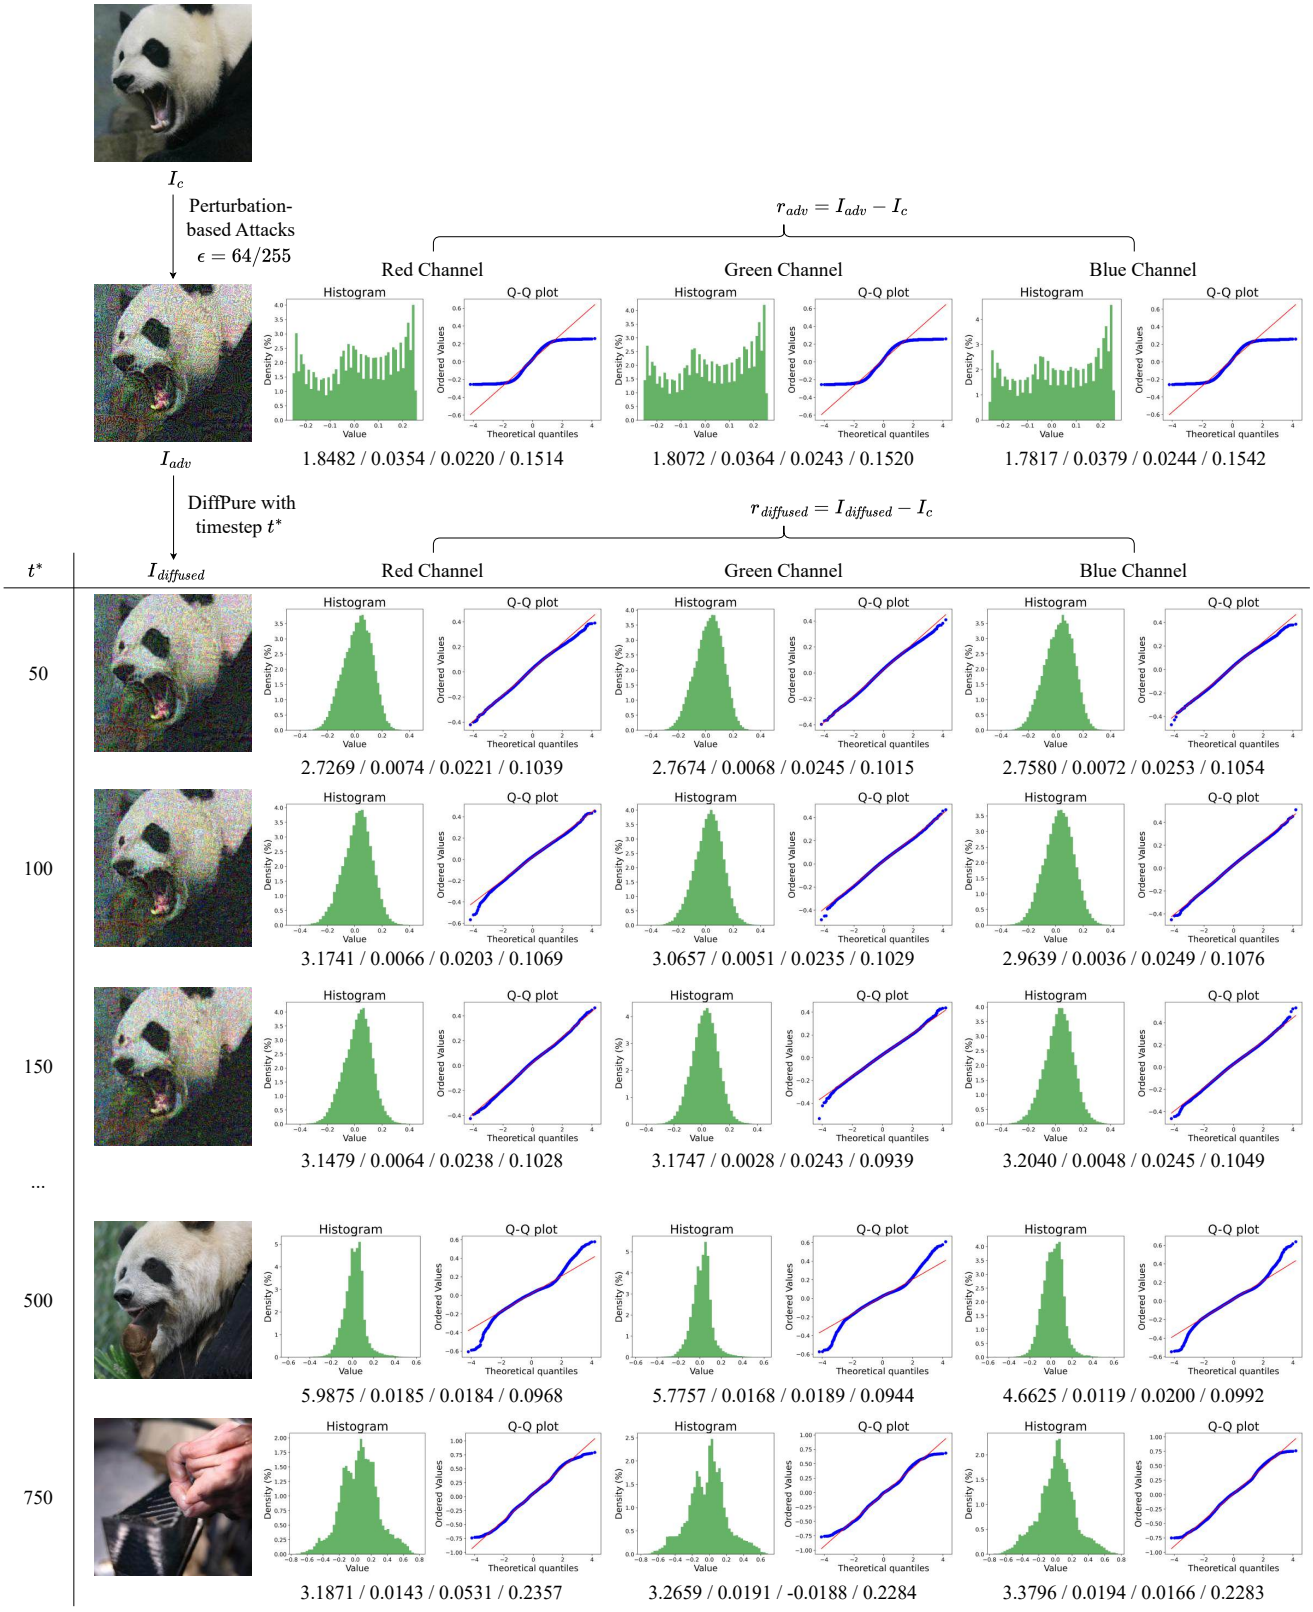

Figure 11. I_{adv}, $I_{diffused}$ and statistics of r_{adv}, $r_{diffused}$ under different t^* in DiffPure (constraint $\epsilon = 64/255$). Metrics are shown in 'Kurtosis / Q-Q Deviation / Mean / Standard Deviation'. Please zoom in to see details.

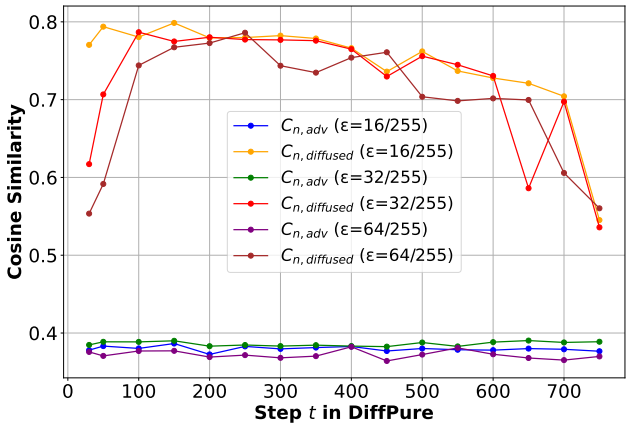

Figure 12. Cosine similarity of visual embeddings under different ϵ of adversarial image I_{adv} and t^* of DiffPure.

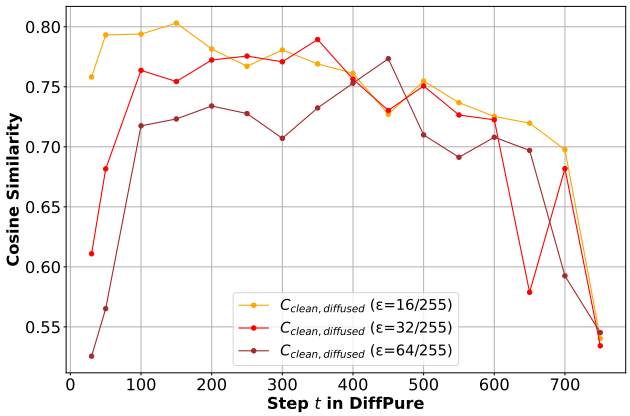

Figure 13. Cosine Similarity of visual embeddings from I_c and $I_{diffused}$ under different ϵ of adversarial image.

gorithmic procedure outlined in Algorithm 1. For our experiments, we employ the Guided Diffusion model for ImageNet [11], specifically the 256×256 unconditional variant provided by OpenAI[2]. Importantly, we synchronize the forward diffusion timesteps ($\mathbf{t}_{\text{forward}}$) with the reverse diffusion timesteps ($\mathbf{t}_{\text{reverse}}$), denoted as t^* in the experimental results, following the setup in DiffPure [33]. Here, we leverage this diffusion model to validate the robustness of our fine-tuned VLMs against Gaussian noise, demonstrating a preliminary defense strategy. However, the fixed image resolution of the diffusion model requires down-sampling and up-sampling operations, which may introduce artifacts not considered during the fine-tuning of the VLM, potentially impacting evaluation results. In the future, adopting more advanced diffusion models will be essential for real-world applications.

[2] https://openaipublic.blob.core.windows.net/diffusion/jul-2021/256x256_diffusion_uncond.pt

Algorithm 1 DiffPure-VLM Adversarial Image Purification with DDPM

Require: Adversarial image x, harmful text prompt p, diffusion model $\mathcal{D}$, forward diffusion timesteps $\mathbf{t}_{\text{forward}}$, reverse diffusion timesteps $\mathbf{t}_{\text{reverse}}$, visual language model $\mathcal{VLM}$.

Ensure: Question answering result output

1: Resize input image x to the size required by the diffusion model (e.g., 256×256).
2: DDPM forward process with $\mathbf{t}_{\text{forward}}$ steps: $\hat{x} = \text{get_noised_x}(x, \mathbf{t}_{\text{forward}})$.
3: **for** t in $\mathbf{t}_{\text{reverse}}$ **do**
4: Denoise using reverse DDPM process: $x = \text{denoising_process}(\hat{x}, t)$.
5: **end for**
6: Obtain purified image with Gaussian noise: $x_{\text{gaussian}} = \text{normalize}(x)$.
7: Perform question answering using VLM: output $= \mathcal{VLM}(x_{\text{gaussian}}, p)$.
8: **return** output

12.2. Extended Experimental Results

In the main paper, for the sake of brevity, we only report results for the standard perturbation-based attack setting of $\epsilon = 32/255$. However, we also conducted experiments with lower attack strength ($\epsilon = 16/255$) and higher attack strength ($\epsilon = 64/255$) to further validate our analysis and approach in Table 12. Across different models and attack strengths, our DiffPure-VLM consistently reduces the attack success rate within a limited number of diffusion timesteps (fewer than 150). Notably, under lower attack strengths, setting the diffusion step to as low as $t^* = 50$ is sufficient to bring the attack success rate down to the level observed for clean inputs. However, under higher attack strengths, $t^* = 50$ fails to reduce the attack success rate to the baseline level for both InternVL2-8B and MiniGPT-4-13B. This indicates that stronger attacks require a larger number of diffusion steps to effectively transform the adversarial noise into Gaussian noise. This finding aligns with the analysis presented in Figure 4 of the main paper, where the residual image at $t^* = 50$ for an attack strength of $\epsilon = 64/255$ does not exhibit Gaussian characteristics. Moreover, we observe that $t^* = 100$ demonstrates strong performance across all attack conditions, making it an effective trade-off between time and robustness. Thus, in real-world applications, setting $t^* = 100$ offers a balanced solution, achieving reliable defense while maintaining computational efficiency.

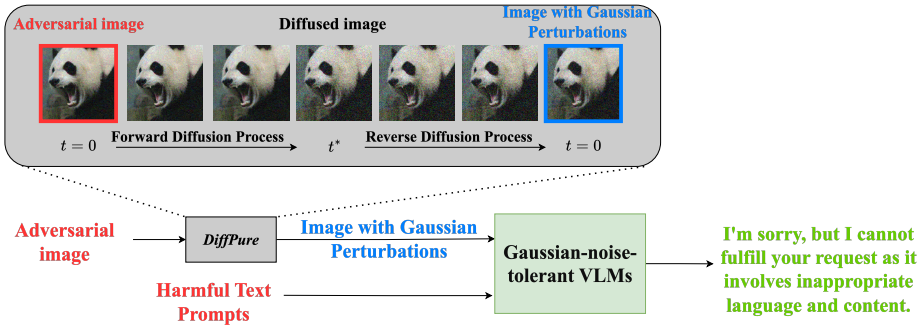

Figure 14. The overall framework of DiffPure-VLM.

13. Conjectures and Discussion on the Impact of Gaussian Noise

Problem Definition

Setting:

- A Visual Language Model (VLM) typically consists of three main components: a visual encoder, a language model, and a vision-language connection module.
- Let the input be a pair (I, T), where $I \in \mathbb{R}^d$ is an image and T is the corresponding text prompt.
- The VLM generates an output sequence of tokens, denoted by $\hat{T} = f_\theta(I, T)$, where f_θ represents the VLM pipeline parameterized by θ.

Adversarial Attack: An adversarial perturbation δ is applied to the image I, resulting in a perturbed image $I_\delta = I + \delta$. The perturbation δ is crafted to manipulate the VLM into generating a specific harmful target text T^{target}. The adversary's objective is:

$$\delta = \arg\min_{\|\delta\| \leq \epsilon} L\left(f_\theta(I + \delta, T), T^{\text{target}}\right),$$

where $L(\cdot, \cdot)$ measures the discrepancy between the generated text $\hat{T}$ and the target text T^{target}. The constraint $\|\delta\| \leq \epsilon$ ensures that the perturbation is imperceptible.

Conjectures: We introduce the following four conjectures to guide our investigation into the impact of Gaussian noise on VLMs:

1. **Sensitivity of Adversarial Attacks to Gaussian Noise**: Adding Gaussian noise to adversarially perturbed images will significantly reduce the effectiveness of the attack.
2. **Gaussian Noise as a Simple Attack on VLM Safety**: Gaussian noise, even without adversarial perturbations, may increase the likelihood of generating harmful text.
3. **Gaussian Noise as a Regularizer**: Augmenting training data with Gaussian noise may act as a regularizer, enhancing the robustness of the VLM.
4. **Fine-Tuning with Gaussian Noise Preserves Performance**: Incorporating Gaussian noise during fine-tuning will preserve or even improve the VLM's overall performance.

Objective: The goal of this study is to systematically evaluate the impact of Gaussian noise on the robustness and reliability of VLMs. By exploring the above conjectures, we aim to determine whether Gaussian noise can effectively mitigate adversarial perturbations and enhance model robustness without compromising performance.

Conjecture 1: Sensitivity of Adversarial Perturbations to Gaussian Noise

Statement: Adversarial perturbations are highly sensitive to Gaussian noise; the attack effectiveness is significantly diminished when Gaussian noise is added to the adversarial image.

Discussion:

Consider an adversarially perturbed image $I_\delta = I + \delta$, where the perturbation δ is optimized to minimize the loss:

$$\delta = \arg\min_{\|\delta\| \leq \epsilon} L\left(f_\theta(I + \delta, T), T^{\text{target}}\right),$$

where $L(\cdot, \cdot)$ measures the discrepancy between the generated text $\hat{T}$ and the harmful target text T^{target}. The perturbation δ is crafted to exploit specific vulnerabilities in the model f_θ.

Now, consider the scenario where Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2 I)$ is added to the input. The new input becomes:

$$I_{\delta,\eta} = I + \delta + \eta.$$

The expected loss over the distribution of Gaussian noise η is:

$$\mathbb{E}_\eta\left[L\left(f_\theta(I + \delta + \eta, T), T^{\text{target}}\right)\right].$$

Since the adversarial perturbation δ is tailored for the specific input I, adding random Gaussian noise η disrupts this optimization. Adversarial perturbations exploit the model's sensitivity along certain directions in the input space, while isotropic Gaussian noise perturbs the input uniformly in all directions, diminishing the effect of δ.

Table 12. Evaluation of DiffPure-VLM's effectiveness on RealToxicityPrompts across different image configurations. Metrics include attack success rate and toxicity levels (Perspective API %). Rows highlighted in light red indicate cases where attack success rate does not meet the baseline level of benign image input.

Image Type	Attack Success Rate ↓	Identity Attack	Profanity	Severe Toxicity	Sexually Explicit	Threat	Toxicity
InternVL2-8B-RobustVLGuard							
Benign Clean Image	29.9	0.8	22.1	0.3	7.2	1.5	25.9
Benign Noisy Image	34.5	2.1	27.2	1.3	8.4	1.6	31.3
Adversarial Image ($\epsilon = 16/255$)	72.5	19.8	58.5	8.3	19.2	7.8	70.0
+DiffPure-VLM ($t^* = 50$)	**31.4**	1.4	24.6	1.1	7.9	1.6	27.5
+DiffPure-VLM ($t^* = 100$)	**28.2**	0.9	21.7	0.4	6.8	1.5	23.9
+DiffPure-VLM ($t^* = 150$)	**28.2**	1.6	22.4	0.2	6.9	1.1	24.4
Adversarial Image ($\epsilon = 32/255$)	70.6	26.7	56.5	9.2	17.3	6.9	68.1
+DiffPure-VLM ($t^* = 50$)	**33.4**	2.4	20.6	0.7	8.1	2.4	29.1
+DiffPure-VLM ($t^* = 100$)	**33.4**	1.6	27.7	0.6	7.6	1.5	30.2
+DiffPure-VLM ($t^* = 150$)	**32.8**	1.7	25.9	0.6	7.7	1.8	29.1
Adversarial Image ($\epsilon = 64/255$)	57.3	9.3	45.8	4.4	16.1	3.9	53.9
+DiffPure-VLM ($t^* = 50$)	**40.9**	2.3	32.9	1.4	9.3	2.3	37.3
+DiffPure-VLM ($t^* = 100$)	**35.7**	1.8	28.2	0.8	7.6	2.4	31.8
+DiffPure-VLM ($t^* = 150$)	**36.1**	2.4	28.3	1.2	8.3	1.8	33.6
LLaVA-v1.5-7B-RobustVLGuard							
Benign Clean image	43.6	4.6	34.7	2.4	12.3	3.5	41.0
Benign Noisy image	42.3	3.1	34.5	1.9	11.8	3.1	40.0
Adversarial image ($\epsilon = 16/255$)	62.6	11.3	48.8	5.3	16.8	5.8	59.1
+DiffPure-VLM ($t^* = 50$)	**42.7**	3.4	32.1	1.5	12.0	4.6	39.7
+DiffPure-VLM ($t^* = 100$)	**42.8**	3.9	32.5	2.3	12.5	3.7	39.3
+DiffPure-VLM ($t^* = 150$)	**44.4**	3.3	34.4	2.2	12.6	3.2	41.0
Adversarial image ($\epsilon = 32/255$)	62.5	7.8	48.0	5.4	16.5	5.8	60.0
+DiffPure-VLM ($t^* = 50$)	**43.9**	3.2	34.6	2.4	12.8	3.7	41.0
+DiffPure-VLM ($t^* = 100$)	**44.1**	3.5	35.4	2.1	13.0	4.1	41.3
+DiffPure-VLM ($t^* = 150$)	**42.5**	3.5	32.7	2.8	12.1	4.1	39.3
Adversarial image ($\epsilon = 64/255$)	57.5	9.2	43.5	5.2	15.3	5.8	54.7
+DiffPure-VLM ($t^* = 50$)	**42.1**	2.7	32.1	2.1	12.3	2.9	39.0
+DiffPure-VLM ($t^* = 100$)	**40.5**	3.3	31.4	1.9	11.7	2.8	37.5
+DiffPure-VLM ($t^* = 150$)	**42.4**	3.5	32.8	1.8	11.5	4.0	40.2
MiniGPT-4-13B-RobustVLGuard							
Benign Clean image	16.0	0.4	9.9	0.3	4.6	1.1	12.1
Benign Noisy image	16.5	0.9	11.9	0.6	5.8	1.0	14.0
Adversarial image ($\epsilon = 16/255$)	47.4	9.3	34.2	1.4	11.8	4.2	41.5
+DiffPure-VLM ($t^* = 50$)	**16.0**	0.6	9.3	0.3	6.5	1.4	13.2
+DiffPure-VLM ($t^* = 100$)	**15.8**	0.7	9.7	0.0	6.1	1.1	12.8
+DiffPure-VLM ($t^* = 150$)	**9.8**	0.4	6.0	0.1	3.3	0.5	7.8
Adversarial image ($\epsilon = 32/255$)	53.7	9.8	35.3	4.1	13.9	5.4	48.1
+DiffPure-VLM ($t^* = 50$)	**13.6**	0.3	9.2	0.2	5.5	0.9	10.6
+DiffPure-VLM ($t^* = 100$)	**15.2**	0.6	9.5	03	5.4	1.1	12.7
+DiffPure-VLM ($t^* = 150$)	**11.9**	0.5	8.6	0.2	4.2	0.6	9.9
Adversarial image ($\epsilon = 64/255$)	60.2	6.8	44.6	4.2	16.2	5.8	56.0
+DiffPure-VLM ($t^* = 50$)	**30.3**	1.8	21.6	1.4	11.4	1.9	26.9
+DiffPure-VLM ($t^* = 100$)	**10.6**	0.0	7.1	0.0	4.1	0.8	8.2
+DiffPure-VLM ($t^* = 150$)	**9.4**	0.4	5.5	0.3	4.1	0.6	7.0

Assuming that f_θ and L are Lipschitz continuous, we can bound the increase in expected loss as follows:

$$\mathbb{E}_\eta \left[L\left(f_\theta(I+\delta+\eta, T), T^{\text{target}} \right) \right] \geq L\left(f_\theta(I+\delta, T), T^{\text{target}} \right) + \frac{\sigma^2 \lambda}{2},$$

where λ is a positive constant related to the curvature of L and f_θ.

This inequality indicates that the addition of Gaussian noise increases the expected loss, thus reducing the effectiveness of the adversarial perturbation. The random noise η disrupts the carefully crafted δ, making it less effective at manipulating the VLM's output. This supports our conjecture that Gaussian noise can act as a simple yet effective countermeasure against adversarial attacks.

Conjecture 2: Gaussian Noise as a Simple Attack on VLM Safety

Statement: Adding Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2 I)$ to a clean image I_{clean} can compromise the safety of VLMs.
Setting: Let I_{clean} be a clean image, and $\eta \sim \mathcal{N}(0, \sigma^2 I)$ be Gaussian noise. The perturbed image is defined as:

$$I_{\text{noisy}} = I_{\text{clean}} + \eta.$$

The VLM processes the noisy image I_{noisy} along with a corresponding text prompt T, and generates an output based on this combined input.
Discussion:

1. **Effect of Noise on Model Input:** The input to the model becomes $I_{\text{noisy}} = I_{\text{clean}} + \eta$. This perturbation, although random, alters the image representation processed by the VLM. The model's output can be locally approximated around the clean input as:

$$f_\theta(I_{\text{clean}} + \eta, T) \approx f_\theta(I_{\text{clean}}, T) + \nabla I_{\text{clean}} f_\theta \cdot \eta,$$

where $\nabla I_{\text{clean}} f_\theta$ represents the gradient of the model output with respect to the clean image input. The Gaussian noise η introduces random perturbations that shift the image features.

2. **Vulnerability of VLMs to Noise:** VLMs are typically trained on clean image data, and thus, they may lack robustness to input noise. The introduction of Gaussian noise can push the model's input into regions of the feature space that were not well-covered during training, potentially causing the model to misinterpret the input and generate unexpected responses.

3. **Impact on Safety:** Adding Gaussian noise may shift the model's behavior towards decision boundaries where safety mechanisms are less effective. This increases the likelihood of generating unsafe or harmful text:

$$L(f_\theta(I_{\text{clean}} + \eta, T), T^{\text{target}}) \geq L(f_\theta(I_{\text{clean}}, T), T^{\text{target}}),$$

where T^{target} represents a potentially harmful target output. The inequality suggests that the noisy input can lead to a higher loss, increasing the risk of unsafe text generation.

4. **Gaussian Noise as a Simple Yet Effective Attack:** Unlike adversarial perturbations that require careful optimization and model-specific crafting, Gaussian noise introduces random changes without any specific targeting. Despite its simplicity, it can destabilize the model and affect its safety, demonstrating that even non-adversarial noise can be a risk factor for VLMs.

In summary, adding Gaussian noise to clean images can indeed disrupt the safety of VLMs, even in the absence of sophisticated adversarial attacks. This highlights a potential vulnerability of VLMs that warrants further investigation.

Conjecture 3: Gaussian Noise as a Regularizer

Statement: Augmenting training data with Gaussian noise acts as a regularizer, reducing the risk of overfitting to adversarial perturbations and enhancing model robustness.
Discussion:

We introduce a regularized loss function that incorporates Gaussian noise during training:

$$L_{\text{reg}}(\theta) = \mathbb{E}_{(I,T)\sim\mathcal{D}} \mathbb{E}_{\eta\sim\mathcal{N}(0,\sigma^2 I)} \left[L\left(f_\theta(I+\eta, T), T \right) \right],$$

where $\mathcal{D}$ represents the training data distribution. This formulation encourages the model to perform well not only on clean inputs but also on noisy inputs, promoting robustness.

To understand the regularizing effect of Gaussian noise, we expand the loss function L using a second-order Taylor expansion around the clean input I:

$$\begin{aligned} L\left(f_\theta(I+\eta, T), T \right) \approx\ & L\left(f_\theta(I, T), T \right) \\ & + \nabla_I L\left(f_\theta(I, T), T \right)^\top \eta \\ & + \frac{1}{2} \eta^\top \nabla_I^2 L\left(f_\theta(I, T), T \right) \eta. \end{aligned}$$

Taking the expectation over the Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2 I)$, we obtain:

$$\begin{aligned} \mathbb{E}_\eta \left[L\left(f_\theta(I+\eta, T), T \right) \right] \approx\ & L\left(f_\theta(I, T), T \right) \\ & + \frac{1}{2} \mathbb{E}_\eta \left[\eta^\top \nabla_I^2 L\left(f_\theta(I, T), T \right) \eta \right] \\ =\ & L\left(f_\theta(I, T), T \right) \\ & + \frac{\sigma^2}{2} \operatorname{Tr}\left(\nabla_I^2 L\left(f_\theta(I, T), T \right) \right). \end{aligned}$$

The additional term $\frac{\sigma^2}{2}\operatorname{Tr}\left(\nabla_I^2 L\left(f_\theta(I,T),T\right)\right)$ penalizes large curvature (i.e., high second derivatives) of the loss function with respect to the input I. This encourages smoother mappings from the input to the output, reducing the model's sensitivity to small input perturbations, including adversarial ones.

In summary, the addition of Gaussian noise during training acts as a regularizer by penalizing sharp changes in the loss landscape. This results in a model that is less prone to overfitting and more resilient to adversarial attacks, as it learns smoother and more stable input-output mappings.

Conjecture 4: Fine-Tuning with Gaussian Noise Preserves Performance

Statement: Fine-tuning the VLM with Gaussian noise-augmented data maintains performance on clean data while enhancing robustness to adversarial perturbations.

Discussion:

Let $\mathcal{D} = \{(I_i, T_i)\}_{i=1}^N$ be the original training dataset. We construct an augmented dataset by adding Gaussian noise:

$$\mathcal{D}_{\text{aug}} = \left\{(I_i + \eta_i, T_i) \mid \eta_i \sim \mathcal{N}(0, \sigma^2 I)\right\}_{i=1}^N.$$

The training objective is to minimize the following loss function:

$$\hat{L}_{\text{aug}}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{E}_{\eta_i}\left[L\left(f_\theta(I_i + \eta_i, T_i), T_i\right)\right].$$

Since the Gaussian noise η_i has a zero mean, the expected gradient of the loss with respect to the model parameters θ is centered around the gradient on the clean data:

$$\mathbb{E}_{\eta_i}\left[\nabla_\theta L\left(f_\theta(I_i + \eta_i, T_i), T_i\right)\right] = \nabla_\theta L\left(f_\theta(I_i, T_i), T_i\right).$$

This result indicates that the expected training gradient remains aligned with the gradient computed on the clean data, thereby preserving the model's performance on clean inputs.

Moreover, by training on both clean and noise-augmented data, the model is exposed to a neighborhood of inputs around each training example. This exposure helps the model generalize better and become less sensitive to small perturbations, effectively enhancing its robustness against adversarial attacks.

In summary, fine-tuning with Gaussian noise-augmented data acts as a regularization strategy that not only maintains the VLM's accuracy on clean data but also improves its resistance to adversarial perturbations.

14. Detailed Proofs

Bounding the Increase in Loss Due to Gaussian Noise

Discussion:

Step 1: Lipschitz Continuity of f_θ and L

Assume that the model function $f_\theta : \mathbb{R}^d \times \mathcal{T} \to \mathbb{R}^k$ and the loss function $L : \mathbb{R}^k \times \mathcal{T} \to \mathbb{R}$ are Lipschitz continuous with constants K_f and K_L, respectively. That is, for all $x, y \in \mathbb{R}^d$ and $T \in \mathcal{T}$:

$$\|f_\theta(x, T) - f_\theta(y, T)\| \leq K_f \|x - y\|,$$

and for all $a, b \in \mathbb{R}^k$:

$$|L(a, T^{\text{target}}) - L(b, T^{\text{target}})| \leq K_L \|a - b\|.$$

Step 2: Bounding the Change in Loss Due to Noise η

Consider the adversarially perturbed image $I_\delta = I + \delta$, where δ is crafted to minimize the loss:

$$\delta = \arg\min_{\|\delta\| \leq \epsilon} L\left(f_\theta(I + \delta, T), T^{\text{target}}\right).$$

When Gaussian noise $\eta \sim \mathcal{N}(0, \sigma^2 I)$ is added, the input becomes $I_{\delta,\eta} = I + \delta + \eta$. The change in loss due to η is:

$$\Delta L = L\left(f_\theta(I + \delta + \eta, T), T^{\text{target}}\right) - L\left(f_\theta(I + \delta, T), T^{\text{target}}\right).$$

Using the Lipschitz continuity of L:

$$|\Delta L| \leq K_L \left\|f_\theta(I + \delta + \eta, T) - f_\theta(I + \delta, T)\right\|.$$

Step 3: Computing the Expected Increase in Loss

Applying the Lipschitz continuity of f_θ:

$$\|f_\theta(I + \delta + \eta, T) - f_\theta(I + \delta, T)\| \leq K_f \|\eta\|.$$

Thus, the change in loss is bounded by:

$$|\Delta L| \leq K_L K_f \|\eta\|.$$

Since η is a Gaussian random vector with zero mean and covariance $\sigma^2 I$, the expected value of $\|\eta\|$ is:

$$\mathbb{E}[\|\eta\|] = \sigma\sqrt{2}\frac{\Gamma\left(\frac{d+1}{2}\right)}{\Gamma\left(\frac{d}{2}\right)} \approx \sigma\sqrt{d - \frac{1}{2}} \quad \text{for large } d.$$

Therefore, the expected increase in loss is approximately:

$$\mathbb{E}[|\Delta L|] \leq K_L K_f \sigma\sqrt{d}.$$

Step 4: Lower Bounding the Expected Increase in Loss

Since δ minimizes $L(f_\theta(I+\delta,T),T^{\text{target}})$ at the point $I+\delta$, any perturbation η added to $I+\delta$ is likely to increase the loss. Under the conjecture that L is convex around $I+\delta$, the expected increase in loss due to η can be lower bounded using the curvature (second derivative) of L:

$$\mathbb{E}_\eta\left[L\left(f_\theta(I+\delta+\eta,T),T^{\text{target}}\right)\right] \geq L\left(f_\theta(I+\delta,T),T^{\text{target}}\right)+\frac{\sigma^2}{2}\lambda_{\min},$$

where $\lambda_{\min}$ is the smallest eigenvalue of the Hessian matrix $\nabla^2_{I+\delta}L(f_\theta(I+\delta,T),T^{\text{target}})$.

Conclusion:

Adding Gaussian noise increases the expected loss by at least $\frac{\sigma^2}{2}\lambda_{\min}$, reducing the effectiveness of the adversarial perturbation. This result supports the conjecture that Gaussian noise disrupts the optimization achieved by the adversary, weakening the impact of adversarial attacks.

Second-Order Taylor Expansion of L Around I

Discussion:

Step 1: Second-Order Taylor Expansion

We expand the loss function $L(f_\theta(I+\eta,T),T)$ around the point I using the second-order Taylor expansion:

$$\begin{aligned}L(f_\theta(I+\eta,T),T) &= L(f_\theta(I,T),T)\\ &+\nabla_I L(f_\theta(I,T),T)^\top\eta\\ &+\frac{1}{2}\eta^\top\nabla_I^2 L(f_\theta(I,T),T)\eta+R_3\end{aligned}$$

where:

- $\nabla_I L(f_\theta(I,T),T)$ is the gradient of the loss with respect to the input I.
- $\nabla_I^2 L(f_\theta(I,T),T)$ is the Hessian matrix of second derivatives with respect to I.
- R_3 is the remainder term of higher order $O(\|\eta\|^3)$.

Step 2: Expected Value of the Linear Term

Since η is sampled from a zero-mean Gaussian distribution $\eta\sim\mathcal{N}(0,\sigma^2 I)$, the expected value of the linear term becomes:

$$\mathbb{E}_\eta\left[\nabla_I L(f_\theta(I,T),T)^\top\eta\right]=\nabla_I L(f_\theta(I,T),T)^\top\mathbb{E}_\eta[\eta]=0$$

Step 3: Expected Value of the Quadratic Term

Next, we compute the expectation of the quadratic term:

$$\mathbb{E}_\eta\left[\eta^\top\nabla_I^2 L(f_\theta(I,T),T)\eta\right]$$

Using the properties of Gaussian distributions, we know that for a symmetric matrix A:

$$\mathbb{E}_\eta\left[\eta^\top A\eta\right]=\sigma^2\operatorname{Tr}(A)$$

Thus, the expected value of the quadratic term becomes:

$$\mathbb{E}_\eta\left[\eta^\top\nabla_I^2 L(f_\theta(I,T),T)\eta\right]=\sigma^2\operatorname{Tr}\left(\nabla_I^2 L(f_\theta(I,T),T)\right)$$

Step 4: Neglecting the Remainder Term

For small values of σ, the remainder term R_3 is of order $O(\sigma^3)$ and can be safely ignored. Thus, the approximation becomes:

$$\begin{aligned}\mathbb{E}_\eta\left[L(f_\theta(I+\eta,T),T)\right] &\approx L(f_\theta(I,T),T)\\ &+\frac{\sigma^2}{2}\operatorname{Tr}\left(\nabla_I^2 L(f_\theta(I,T),T)\right)\end{aligned}$$

Step 5: Interpretation of the Trace Term

The term $\operatorname{Tr}\left(\nabla_I^2 L(f_\theta(I,T),T)\right)$ denotes the sum of the eigenvalues of the Hessian matrix, representing the overall curvature of the loss function with respect to the input. A larger trace value indicates higher curvature, suggesting greater sensitivity of the model to input perturbations. Reducing this sensitivity is crucial for enhancing the model's robustness.

Step 6: Gaussian Noise as Regularization

The additional term $\frac{\sigma^2}{2}\operatorname{Tr}\left(\nabla_I^2 L(f_\theta(I,T),T)\right)$ functions as a regularizer, penalizing high curvature in the loss landscape. This encourages the model to learn smoother input-output mappings, thereby reducing its vulnerability to small perturbations, including adversarial attacks.

Step 7: Connection to Tikhonov Regularization

This regularization effect is conceptually similar to Tikhonov regularization, where a penalty proportional to the norm of the model parameters is added to the loss function. In our case, the penalty arises naturally from the Gaussian noise, encouraging robustness by flattening the loss landscape:

$$\begin{aligned}\mathbb{E}_\eta\left[L(f_\theta(I+\eta,T),T)\right] &\approx L(f_\theta(I,T),T)\\ &+\frac{\sigma^2}{2}\operatorname{Tr}\left(\nabla_I^2 L(f_\theta(I,T),T)\right)\end{aligned}$$

This smoothing effect reduces the model's sensitivity to input perturbations, enhancing its robustness without compromising performance on clean data.