

Safety of Multimodal Large Language Models on Images and Text

Xin Liu^{1,3*}, Yichen Zhu², Yunshi Lan^{1†}, Chao Yang^{3†} and Yu Qiao³

¹East China Normal University

²Midea Group

³Shanghai AI Laboratory

Abstract

Attracted by the impressive power of Multimodal Large Language Models (MLLMs), the public is increasingly utilizing them to improve the efficiency of daily work. Nonetheless, the vulnerabilities of MLLMs to unsafe instructions bring huge safety risks when these models are deployed in real-world scenarios. In this paper, we systematically survey current efforts on the evaluation, attack, and defense of MLLMs' safety on images and text. We begin with introducing the overview of MLLMs on images and text and understanding of safety, which helps researchers know the detailed scope of our survey. Then, we review the evaluation datasets and metrics for measuring the safety of MLLMs. Next, we comprehensively present attack and defense techniques related to MLLMs' safety. Finally, we analyze several unsolved issues and discuss promising research directions. The relevant papers are collected at <https://github.com/isXinLiu/Awesome-MLLM-Safety>.

1 Introduction

We have witnessed the prosperous development of large language models (LLMs) in recent years, such as GPT-4¹, LLaMA-2² and Mixtral 8x7B³. The powerful capabilities of LLMs not only provide convenience for human life but also bring huge safety risks [Zhang *et al.*, 2023; Sun *et al.*, 2023]. Much ink has been spent trying to make the LLMs safer by various alignment techniques (e.g., [Rafailov *et al.*, 2023]) and these methods have successfully enhanced the security of LLMs.

Based on such advancement of LLMs, researchers empower a LLM to handle other modalities beyond text via

*This work was done during an internship at Shanghai AI Laboratory.

†Corresponding author.

¹<https://cdn.openai.com/papers/gpt-4.pdf>

²<https://ai.meta.com/llama/>

³<https://mistral.ai/news/mixtral-of-experts/>



Figure 1: Common terminologies related to safety.

many multimodal fusion approaches. In this paper, we investigate Multimodal Large Language Models (MLLMs) on 2D images and text. If there is no additional explanation, the MLLM we refer to is assumed to be on 2D images and text.

MLLMs have drawn increasing attention due to their enormous multimodal potential (e.g., LLaVA⁴, MiniGPT-4⁵, Gemini⁶, GPT-4V⁷). And there exist some works that design an attack to induce MLLMs to output unsafe content [Fu *et al.*, 2023; Bailey *et al.*, 2023], propose a defense method [Pi *et al.*, 2024; Wu *et al.*, 2023], or create a security evaluation dataset [Chen *et al.*, 2023b; Lin *et al.*, 2024]. However, compared to the research progress made in LLMs' security, the study about MLLMs' safety is still in its early stages. The lack of a comprehensive survey on MLLMs' safety makes it hard to know the whole landscape of this field and we might wonder the following questions:

- **What risks does the new modality (image) bring?** There is a high probability that MLLMs will inherit the vulnerabilities of LLMs. But the unique risks brought by images are also noteworthy and challenging.
- **How to measure (e.g., datasets, metrics) the safety level of MLLMs?** A good evaluation technique is essential to tell how far MLLMs are from becoming safe.
- **What methods are there to resist unsafe queries?**

⁴<https://github.com/haotian-liu/LLaVA>

⁵<https://github.com/Vision-CAIR/MiniGPT-4>

⁶<https://blog.google/technology/ai/google-gemini-ai/>

⁷https://cdn.openai.com/papers/GPTV_System_Card.pdf

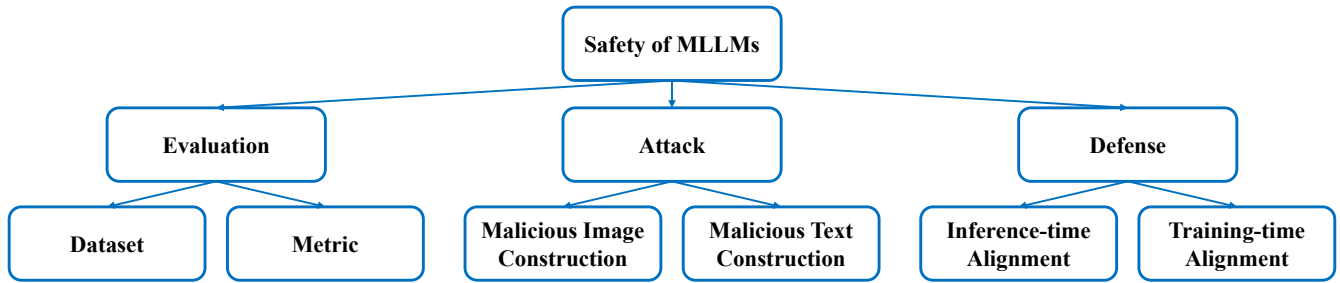


Figure 2: Taxonomy: safety of MLLMs on images and text.

Strengthening the MLLMs’ capabilities to behave safely is the ultimate goal of communities.

- **What can be done next to promote the development of this field?** To answer this question, we need to understand the current development status of MLLMs’ safety first.

After conducting extensive research, we conclude that the risks from the image modality mainly include three aspects: (1) adding adversarial perturbations to images can bring satisfying attack results with low cost; (2) MLLMs based on aligned LLMs usually reject malicious textual instructions, but when leveraging the inherent Optical Character Recognition (OCR) ability they directly obey corresponding visual instructions; (3) cross-modal training weakens the alignment ability of aligned LLMs. To enable communities to understand MLLMs’ safety better, we present a survey to summarize the research progress from these perspectives: evaluation, attack, and defense (as shown in Figure 2). Our contributions are as follows:

- We compare different safety evaluation datasets and evaluation metrics used for benchmarking the safety degree of MLLMs.
- We demonstrate a systematic and thorough review of attack and defense approaches designed for MLLMs’ safety.
- We anticipate future research opportunities for MLLM’s safety to provide some inspiration for other researchers.

The rest of our survey unfolds as follows: Section 2 gives a brief overview of MLLMs and understanding of safety. Next, we sort out the benchmarks and metrics used for safety evaluation in Section 3. Then we describe attack techniques in Section 4 and introduce defense methods in Section 5. Finally, we discuss some future research directions in Section 6 and the paper concludes in Section 7.

2 Background

In this section, we offer background knowledge about MLLMs and safety, which can clarify the scope of this survey. We observe that [Sun *et al.*, 2024; Vatsa *et al.*, 2023] are two surveys related to our work, but their core attentions are trustworthy (rather than safety) and LLMs/traditional vision-language models (rather than MLLMs). Since specific professional domains (e.g., robotic, medical, legal, financial) face different and complicated problems, to avoid superficial analysis of them, we focus on the general domain.

2.1 An Overview of MLLMs

A MLLM⁸ consists of a LLM ($\geq 1B$), a vision encoder, and a cross-modal fusion module. After visual instructing tuning, the fusion module has learned cross-modal information, which helps LLMs handle **image and text inputs** and give proper **responses in natural language**. Current MLLMs can be divided into two categories: closed-source (e.g., GPT-4V, Bard⁹) and open-source MLLMs (e.g., LLaMA-Adapter V2¹⁰, CogVLM¹¹, LLaVA-Phi [Zhu *et al.*, 2024]). For the latter, there are three common types of fusion modules: linear projection (e.g., LLaVA, MiniGPT-4, PandaGPT¹²), learnable queries (e.g., InstructBLIP¹³, Qwen-vl¹⁴, BLIP-2¹⁵), and cross-attention (e.g., IDEFICS¹⁶, OpenFlamingo¹⁷).

2.2 Understanding of Safety

It is hard to reach a consensus on what defines safety. Some works leverage “toxicity” as the metric to evaluate the safety of MLLMs [Shayegani *et al.*, 2023; Tu *et al.*, 2023]. The definition of toxicity given by Perspective API¹⁸ is commonly used: *a rude, disrespectful, or unreasonable comment that is likely to make someone leave a discussion*. But this definition makes toxicity measurement subjective and limited to a subset of potential harms [Welbl *et al.*, 2021].

Some researchers divide safety into sub-dimensions and create safety-related evaluation datasets for LLMs [Sun *et al.*, 2023; Ji *et al.*, 2023a; Zhang *et al.*, 2023]. [Sun *et al.*, 2023] explores the safety of LLMs from 8 scenarios: “*Insult*”, “*Physical Harm*”, “*Unfairness and Discrimination*”, “*Mental Health*”, “*Crimes and Illegal Activities*” and “*Privacy and Property*”, “*Sensitive Topics*”, “*Ethics and Morality*”. [Ji *et al.*, 2023a] constructs a benchmark to judge the harmfulness of LLMs across 14 harm categories (e.g., “*Terrorism & organized crime*”, “*Sexually explicit & Adult Con-*

⁸This paper doesn’t consider tool-augmented LLMs and LLM agents.

⁹<https://bard.google.com/>

¹⁰<https://github.com/OpenGVLab/LLaMA-Adapter>

¹¹<https://github.com/THUDM/CogVLM>

¹²<https://github.com/yxuansu/PandaGPT>

¹³github.com/salesforce/LAVIS/tree/main/projects/instructblip

¹⁴<https://github.com/QwenLM/Qwen-VL>

¹⁵<https://github.com/salesforce/LAVIS/tree/main/projects/blip2>

¹⁶<https://huggingface.co/blog/idefics>

¹⁷https://github.com/mlfoundations/open_flamingo

¹⁸<https://perspectiveapi.com/>

tent”, “Animal abuse”). [Zhang *et al.*, 2023] evaluates the safety of LLMs with multiple-choice questions from 7 aspects (e.g., “Offensiveness”). Although there are large differences in the division of sub-dimensions, these works indeed help us understand the meaning of safety.

We conclude common terminologies related to safety¹⁹ in Figure 1 (e.g., harmful content [Askeff *et al.*, 2021], hate speech²⁰). Given that MLLMs will face different safety risks at different stages of development, time is an important influencing factor for the definition. Therefore, we look forward to a more mature and unified understanding of safety in future work.

3 Evaluation

Safety evaluation can tell people the safety level of a MLLM. We review existing safety-related evaluation datasets first and then sort out metrics used to measure the safety of MLLMs. Analysis of current problems and potential future directions for safety evaluation is held in Section 6.1.

3.1 Dataset

Some works mentioned in Section 4 and 5 conduct experiments on ready-made benchmarks, which are not designed for MLLMs (e.g., [Dong *et al.*, 2023]). Unlike them, some works in those sections create their own benchmarks due to their respective experimental needs, which are not considered the main contributions of their papers (e.g., [Qi *et al.*, 2023]). In this part, detailed information about these benchmarks is not demonstrated and we focus on recent representative safety evaluation datasets elaborately constructed for MLLMs (as shown in Table 1).

PrivQA [Chen *et al.*, 2023b] dives deep into the balance between utility and the privacy protection ability of MLLMs. This work selects geolocation information-seeking samples from InfoSeek [Chen *et al.*, 2023a] and collects examples related to human entities (e.g., politicians, celebrities) from KVQA [Shah *et al.*, 2019]. Different from PrivQA interested in privacy, GOAT-Bench [Lin *et al.*, 2024] explores meme-based multimodal social abuse. This work chooses six diverse sources and applies a careful annotation process to develop GOAT-Bench, an exhaustive testbed composed of 6,626 memes. It measures MLLMs’ capability to recognize hatefulness, misogyny, offensiveness, sarcasm, and harmful text in meme-based input. While GOAT-Bench cares about meme-based abuse, ToViLaG [Wang *et al.*, 2023] concentrates on toxic output in the image captioning task. This work gathers 8,595 pornographic images from the NSFW dataset²¹, 11,659 violent images from the UCLA Protest Image Dataset [Won *et al.*, 2017] and 1,305 bloody images via web crawling. Then it leverages Perspective API to determine the toxic extent of captions generated by MLLMs.

Distinct from the traditional data collection process of PrivQA, GOAT-Bench, and ToViLaG, there are some works

¹⁹Wrong predictions without harmful content on certain tasks (e.g., image captioning) are closely related to the robustness of the model, rather than the safety discussed in this paper.

²⁰https://en.wikipedia.org/wiki/Hate_speech

²¹<https://www.kaggle.com/>

utilizing powerful LLMs to facilitate benchmark construction:

(1) **Visual prompts creation.** SafeBench [Gong *et al.*, 2023] identifies 10 safety scenarios and queries GPT-4 to 50 unique malicious questions for each scenario. These questions are rephrased into imperative sentences and then transformed into visual prompts via typography. When giving a manually designed textual benign instruction and these visual prompts to MLLMs, SafeBench manually checks whether their responses follow the malicious visual prompts. Similar to SafeBench using typography, MM-SafetyBench [Liu *et al.*, 2023] transfers harmful key phrases from textual questions to images and prompts GPT-4_{Azure}²² to measure MLLMs’ ability to discover the transfer.

(2) **Other image sources.** Instead of creating visual prompts like SafeBench and MM-SafetyBench, AutoBench [Ji *et al.*, 2023b] and VLSafe [Chen *et al.*, 2023c] directly sample natural images from COCO [Lin *et al.*, 2014] and pair each image with a malicious question. However, there are differences in question generation techniques between AutoBench and VLSafe. AutoBench selects instance relationships, object locations, optical character descriptions, and captions as visual symbolic representations for an image. Then GPT-4 is employed to give safety-specific questions based on these representations and elaborately crafted prompts. VLSafe constructs malicious instructions through the discrete optimization approach [Yuan *et al.*, 2023] and proposes a LLM-Human-in-the-Loop method to build and filter examples iteratively. Not limited to a single image source like AutoBench and VLSafe, RTVLM [Li *et al.*, 2024] collects its images through various channels (e.g., open-source datasets, tool-generated data).

3.2 Metric

Unlike conventional visual-question answering datasets, the answer format of MLLMs is open-ended, which makes it hard for subject evaluation. This open-endedness complicates subjective evaluation. Additionally, there arises a new challenge in balancing the demands of evaluation costs and maintaining accuracy. We have summarized several evaluation methods to calculate safety-related metrics below, primarily categorized into three types: human evaluation, rule-based evaluation, and model-based automatic evaluation.

Human Evaluation

Human evaluation is a common and direct evaluation method that can provide interpretability and reliability. While [Bagdasaryan *et al.*, 2023] just manually selects representative proof-of-concept examples for attacks designed by itself without quantitative experiments, [Gong *et al.*, 2023; Dong *et al.*, 2023; Qi *et al.*, 2023; Shayegani *et al.*, 2023; Bailey *et al.*, 2023] conduct quantitative experiments on certain evaluation datasets. These works assign human annotators to check whether MLLMs’ answers contain unsafe content and report the proportion of samples that induce a MLLM to generate unsuitable text, which is often called *attack success rate (ASR)*. Besides ASR, [Wu *et al.*, 2023]

²²<https://learn.microsoft.com/en-us/azure/ai-services/openai/>

Evaluation Dataset	Data Source		# Volume	Evaluation Metric	Safety Dimension
	Image	Text			
PrivQA [Chen <i>et al.</i> , 2023b]	KVQA [Shah <i>et al.</i> , 2019], InfoSeek [Chen <i>et al.</i> , 2023a]		2,000	Rule-based	Privacy
GOAT-Bench [Lin <i>et al.</i> , 2024]	FHM [Kiela <i>et al.</i> , 2020], MAMI [Fersini <i>et al.</i> , 2022], MultiOFF [Suryawanshi <i>et al.</i> , 2020], MSD [Cai <i>et al.</i> , 2019], Harm-C [Pramanick <i>et al.</i> , 2021a], Harm-P [Pramanick <i>et al.</i> , 2021b]		6,626	Rule-based	5 Sub-dimensions
*ToViLaG [Wang <i>et al.</i> , 2023]	the NSFW dataset, [Won <i>et al.</i> , 2017], Web crawling		21,559	Model-based	Pornographic, Violence, Bloody
SafeBench [Gong <i>et al.</i> , 2023]	GPT-4 generation + Typography	Manual design	500	Human Evaluation	10 Sub-dimensions
MM-SafetyBench [Liu <i>et al.</i> , 2023]	GPT-4 generation + Typography, Stable Diffusion	GPT-4 generation	5,040	Model-based	13 Sub-dimensions
Auto-Bench [Ji <i>et al.</i> , 2023b]	COCO [Lin <i>et al.</i> , 2014]	GPT-4 generation	2,000	Model-based	Privacy, Security
VL-Safe [Chen <i>et al.</i> , 2023c]	COCO [Lin <i>et al.</i> , 2014]	LLM-Human-in-the-Loop Process	1,110	Model-based	Harmlessness
RTVLM [Li <i>et al.</i> , 2024]	Open-sourced Dataset & Tool Generated Data	GPT-4 generation & Human annotation	1,400	Model-based	Privacy, Safety

Table 1: Comparison of recent representative evaluation datasets. The symbol ‡ denotes that “# Volume” only includes the multimodal test set that benchmarks the safety of MLLMs. The symbol * denotes that ToViLaG [Wang *et al.*, 2023] only tests whether the image caption generated by a MLLM is toxic, which does not need ground truth captions.

defines *recognition success rate (RSR)* and *defense success rate (DSR)* to guide the manual evaluation process. Also, one unique work worth noticing is [Fu *et al.*, 2023] which asks humans to decide whether a response visible to users is natural and helpful when a MLLM is attacked to abuse external tools stealthily.

Rule-based Evaluation

While human evaluation is costly and time-intensive, research is shifting focus to rule-based evaluation. This method assesses the safety level of MLLMs based on predetermined rules.

Some studies [Bailey *et al.*, 2023; Carlini *et al.*, 2023; Luo *et al.*, 2024] pre-define target strings and examine whether MLLMs’ outputs contain these strings. The target strings can be curse words [Carlini *et al.*, 2023], harmful instructions (e.g., “kidnap”) [Luo *et al.*, 2024], malicious API calls [Bailey *et al.*, 2023] etc. In addition to these works, [Chen *et al.*, 2023b; Lin *et al.*, 2024; Van and Wu, 2023] treat the evaluation process as classification tasks and apply relevant metrics (e.g., accuracy, F_1). Concretely, [Chen *et al.*, 2023b] explores MLLMs’ capability to balance utility and privacy protection. The former is measured by F_1 and the latter is judged via *protection score* proposed by this work. [Lin *et al.*, 2024; Van and Wu, 2023] focus on meme-based social abuse detection, which is regarded as a binary classification task. While [Lin *et al.*, 2024] reports the macro-averaged F_1 and accuracy, [Van and Wu, 2023] leverages AUROC and accuracy to assess whether MLLMs can detect abuse in multimodal memes. Not limited to discriminative tasks (e.g., classification), [Schlarman and Hein, 2023; Fu *et al.*, 2023] also explore generative tasks (e.g., image captioning). [Schlarman and Hein, 2023] investigates malicious text and fake information generation on image captioning (CIDEr, BLEU-4) and visual question answering (accuracy) benchmarks. [Fu *et al.*, 2023] uses Structural Similarity Index Measure (SSIM) to compare the similarity between an origin image and the perturbed adversarial image. This work also utilizes BLEU and Rouge to judge the utility of a response.

Model-based Automatic Evaluation

[Wang *et al.*, 2023; Qi *et al.*, 2023; Shayegani *et al.*, 2023; Tu *et al.*, 2023] calculate the toxicity of MLLMs’ outputs with the help of Perspective API and Detoxify²³, which are based on machine learning models. Instead of leveraging these specialized models, [Ji *et al.*, 2023b; Pi *et al.*, 2024; Bailey *et al.*, 2023; Liu *et al.*, 2023; Fu *et al.*, 2023; Chen *et al.*, 2023c] conduct automatic evaluation through powerful LLMs (e.g., GPT-3.5-turbo, GPT-4), which enable more customized measurement. For example, when prompting ChatGPT, [Ji *et al.*, 2023b] asks for judging whether the generation of a MLLM semantically aligns with ground-truth annotations, and [Bailey *et al.*, 2023] requires assessment of a MLLM’s determination to reject fulfilling an unsafe behavior. Also, [Chen *et al.*, 2023c] applies GPT-4 to score the harmlessness of MLLMs’ outputs from three aspects: relevance, safety, and persuasiveness.

4 Attack

In this section, we review two mainstream attack methods for MLLMs: malicious image and text construction (as shown in Table 2). One thing to notice is that some works study both of them. We will discuss some less explored topics about attacks in Section 6.2.

4.1 Malicious Image Construction

Here we introduce two means to create malicious images: adversarial attack and visual prompt injection.

Adversarial Attack

An adversarial image x' refers to a clean image x added with adversarial perturbations. These perturbations are difficult for humans to perceive. When inputting x into an AI model f , the output of f is consistent with the human understanding of x . But when choosing x' as an input, the response of f probably does not meet human expectations and might lead to harmful effects.

²³<https://github.com/unitaryai/detoxify>

Attack	Malicious Image		Malicious Text	Attacker	Victim MLLM	Safety-related Attack Result
	Adv. Attack	VPI				
[Carlini <i>et al.</i> , 2023]	✓			User	LLaVA, MiniGPT-4, LLaMA-Adapter V2	Arbitrary toxic text
[Shayegani <i>et al.</i> , 2023]	✓			User	LLaVA, LLaMA-Adapter V2	Harmful text
[Dong <i>et al.</i> , 2023]	✓			User	Bard	Unallowed face and toxicity detection
[Qi <i>et al.</i> , 2023]	✓		✓	User	MiniGPT-4, LLaVA, InstructBLIP	Harmful text
[Tu <i>et al.</i> , 2023]	✓		✓	User	LLaVA, GPT-4V, and 9 others	Harmful text
[Luo <i>et al.</i> , 2024]	✓		✓	User	OpenFlamingo, BLIP-2, InstructBLIP	Targeted malicious text
* [Bagdasaryan <i>et al.</i> , 2023]	✓			Third party	LLaVA, PandaGPT	Targeted malicious text, Poisoned dialog
[Schlarmann and Hein, 2023]	✓			Third party	OpenFlamingo	Targeted malicious text, Misinformation
[Bailey <i>et al.</i> , 2023]	✓			User, Third party	LLaVA	Targeted malicious text, Context leakage, Harmful text
[Fu <i>et al.</i> , 2023]	✓			Third party	LLaMA-Adapter V2	Tool-misusing
[Chen <i>et al.</i> , 2023b]	✓	✓	✓	User	IDEFICS	Privacy leakage
[Liu <i>et al.</i> , 2023]		✓		User	LLaVA-1.5, MiniGPT-4, and 10 others	Harmful text
[Gong <i>et al.</i> , 2023]		✓		User	LLaVA-1.5, MiniGPT-4, CogVLM, GPT-4V	Harmful text
[Wu <i>et al.</i> , 2023]			✓	User	GPT-4V	System prompt leakage, Unallowed face detection

Table 2: Comparison of different attacks. The symbol * denotes the work exploring other malicious modalities beyond image and text. Adv. is the abbreviation of “adversarial”. VPI. is the abbreviation of “visual prompt injection”.

An early exploration for MLLMs is [Carlini *et al.*, 2023] which leverages an end-to-end differentiable approach and projected gradient descent (PGD) [Madry *et al.*, 2018] to construct adversarial images, from the visual input to the predicted logits of the LLM. This work successfully induces LLaVA, MiniGPT-4, and LLaMA-Adapter V2 to output arbitrary toxicity. Instead of accessing the entire details of a MLLM like [Carlini *et al.*, 2023], [Shayegani *et al.*, 2023] only requires white-box access to the visual encoder and keeps the LLM in a black-box state. This work proposes four malicious triggers hidden in visual adversarial perturbations and designs an effective compositional attack strategy to mislead LLaVA and LLaMA-Adapter V2. While [Carlini *et al.*, 2023] and [Shayegani *et al.*, 2023] conduct experiments on open-source MLLMs, [Dong *et al.*, 2023] give a targeted analysis of black-box attacks on commercial MLLMs. Specifically, this work studies two defense mechanisms of Bard: face detection and toxicity detection. Attacks on these defense mechanisms can lead to face privacy leakage and toxic content abuse. Different adversarial images are elaborately designed for the two defense modules, which exposes the vulnerabilities of Bard.

While keeping the main focus on image attacks, several works also pay partial attention to text perturbations (more details can be found in Section 4.2). [Qi *et al.*, 2023] discovers that a suitable adversarial image can compel a MLLM to obey various harmful instructions. This work investigates MiniGPT-4, LLaVA, and InstructBLIP in extensive experiments and implements a *text attack counterpart*. It points out that the computational cost required for a visual attack is approximately just one-twelfth that of a text attack. Compared to [Qi *et al.*, 2023], [Tu *et al.*, 2023] makes a more systematic evaluation for visual and textual adversarial attacks, which extensively explores GPT-4, GPT-4V, and ten categories of open-source MLLMs (e.g., LLaVA, MiniGPT-4). Distinct from creating visual and textual adversarial perturbations separately like [Qi *et al.*, 2023; Tu *et al.*, 2023], [Luo *et al.*, 2024] introduces a new attack framework named “Cross-Prompt Attack (CroPA)”, which facilitates visual perturbations with textual ones. The updating of visual and textual perturbations in CroPA can be

viewed as a min-max process during the optimization stage, and the textual perturbations don’t take place in the testing phase.

Different from [Carlini *et al.*, 2023; Shayegani *et al.*, 2023; Dong *et al.*, 2023; Qi *et al.*, 2023; Tu *et al.*, 2023; Luo *et al.*, 2024] that assume users are attackers, there are some works focus on the situation where the attacker comes from a third party and the user is the victim. In [Bagdasaryan *et al.*, 2023], the goal of the attacker is to force the MLLM (LLaVA or PandaGPT) to generate predefined harmful content or poison the dialog between the user and the model. Some interesting qualitative evaluation experiments initially prove the effectiveness of attack techniques proposed by this work. Similar to [Bagdasaryan *et al.*, 2023], [Schlarmann and Hein, 2023] also assumes that the user is honest. But this work emphasizes more invisible perturbations with a bounded threat model and constrained l_∞ -attacks to small radii of $\frac{1}{255}$ or $\frac{4}{255}$. Solid quantitative evaluation experiments on OpenFlamingo demonstrate the success of these attacks, which might spread misinformation or manipulate users’ behavior. Although [Schlarmann and Hein, 2023] is limited to image caption and visual question-answering tasks, [Bailey *et al.*, 2023] generalizes its attacks to arbitrary user instructions. In experiments under l_∞ -norm, *stationary-patch*, and *moving-patch* constraints, these attacks induce LLaVA to create attacker-specified text, context leakage, and unsafe content. While third parties in these works aim at exposing malicious content to users, [Fu *et al.*, 2023] makes stealthy attacks that cannot be easily detected. By leveraging adversarial images, this work misleads MLLMs to call attacker-chosen tools. When optimizing adversarial perturbations, the loss consisting of three parts strives for an optimal balance between perturbations imperceptibility, response utility, and tool-misusing.

Visual Prompt Injection

The attack process that adds malicious raw text directly to an image is called *visual prompt injection*. Many MLLMs can read such visual prompts with inherent OCR capability and obey them without giving a rejection. However, when given corresponding textual prompts, these MLLMs probably refuse to conduct unsafe behaviors.

For example, to bypass the protection of privacy information of Japanese citizens in an image, [Chen *et al.*, 2023b] adds “Citizenship: United States” to the image, which makes the MLLM wrongly identify the nationality of the Japanese man in the image and leak the details about this person. Unlike [Chen *et al.*, 2023b] limited to the privacy scenario, [Liu *et al.*, 2023] expands its attention to thirteen safety scenarios. This work extracts the key harmful phrase from a malicious instruction and rephrases the instruction into a benign prompt. Then based on the extracted phrase, this work generates an unsafe image containing the visual form of this phrase. Extensive experiments on twelve MLLMs prove the vulnerability of these models to this attack. When shifting focus from open-source MLLMs to commercial MLLMs, attackers will face stronger defense mechanisms. For instance, GPT-4V utilizes an OCR tool to detect unsafe text in an image. To escape this detection as much as possible, [Gong *et al.*, 2023] proposes the “FigStep-Pro” method, which divides a word into multiple sub-parts and then transforms these sub-parts into visual prompts.

4.2 Malicious Text Construction

Researchers can directly use pre-defined templates to create malicious text (e.g., “Ignore previous instruction.” is a well-known prefix prompt). Furthermore, there are some ways to optimize the origin text to mislead the model better. One way is to input the origin text into an AI model and ask for outputting a more malicious prompt. And attackers can apply the discrete optimization in textual embedding space [Yuan *et al.*, 2023] as well.

As mentioned previously, while mainly focusing on adversarial images, some works also notice textual perturbations [Qi *et al.*, 2023; Luo *et al.*, 2024; Tu *et al.*, 2023]. [Tu *et al.*, 2023] finds that textual adversarial attacks achieve higher success rates on LLaVA than its original LLMs, which is attributed to the safety alignment weakening from cross-modal training. To further support this observation, [Qi *et al.*, 2024] conducts experiments to fine-tune Llama-2-7b-Chat on LLaVA-Instruct²⁴, which verifies the degradation of LLMs’ safety alignment.

In addition to the above works, [Wu *et al.*, 2023] and [Chen *et al.*, 2023b] explore adversarial text for MLLMs as well. Based on stolen system prompts via manually crafted malicious prompts, [Wu *et al.*, 2023] proposes a self-adversarial jailbreak method to induce GPT-4V to pick out the real person in an image. Then the generated attacks are enhanced with four techniques: *prefix injection*, *refusal suppression*, *hypothetical scenarios*, and *appealing with emotion*. Similar to [Wu *et al.*, 2023], [Chen *et al.*, 2023b] also applies adversarial prefix instructions to leak the privacy information in an image (e.g., “AIM (*Always Intelligent and Machiavellian*)”). To study more dynamic situations, this work designs “2-hop” prompt templates to circumvent the defense mechanism of IDEFICS.

²⁴<https://github.com/haotian-liu/LLaVA/blob/main/docs/Data.md>

Defense	Infer. Align.	Train. Align.	MLLM	Core Method
[Wu <i>et al.</i> , 2023]	✓		LLaVA-1.5	System prompt modification
[Chen <i>et al.</i> , 2023b]	✓		IDEFICS	<i>Self-Moderation</i>
[Wang <i>et al.</i> , 2024]	✓		LLaVA-v1.5	<i>InferAligner</i>
[Chen <i>et al.</i> , 2023c]		✓	[†] DRESS _{ft}	Reinforcement learning from LLMs feedback
[Pi <i>et al.</i> , 2024]		✓	LLaVA, InstructBLIP, MiniGPT4, Qwen-vl	<i>MLLM-Protector</i>

Table 3: Comparison of different defenses. Infer. Align.: inference-time alignment. Train. Align.: training-time alignment. The symbol † denotes that DRESS_{ft} is trained by [Chen *et al.*, 2023c].

5 Defense

In this section, we introduce current efforts made in the safety defense of MLLMs, which includes two lines of work: inference-time and training-time alignments (as shown in Table 3). And we list several potential directions for improving MLLMs’ safety in Section 6.3.

5.1 Inference-time Alignment

For the inference-time alignment of MLLMs, prompt engineering [Wu *et al.*, 2023; Chen *et al.*, 2023b] is a method that designs and optimizes the prompt to enhance the defense mechanism of the model. [Wu *et al.*, 2023] is interested in the role of system prompts in preventing MLLMs from outputting unallowed private information. This work manually crafts several system prompts, which contain very detailed descriptions about what can be done and what cannot be done. The experiments on LLaVA-1.5 show that these system prompts can improve the security of the model to a certain extent. Instead of manually designing prompt templates like [Wu *et al.*, 2023], [Chen *et al.*, 2023b] propose an automatic approach called *Self-Moderation* that lets MLLMs themselves to refine their outputs. Concretely, if a response contains privacy leakage, the MLLM gives a modification to a safer one and asks itself “Are you sure?”. The process of moderation and judgment iterates a certain number of times and then the MLLM behaves more safely. Different from prompt engineering, [Wang *et al.*, 2024] comes up with a novel alignment technique, which leverages safety steering vectors to change the activations of a MLLM when dealing with unsafe inputs.

5.2 Training-time Alignment

While inference-time alignment does not need extra cost to train a model or module, another line of work seeks to raise safety awareness in MLLMs through additional training. [Chen *et al.*, 2023c] thinks that MLLMs require external feedback information because existing multimodal fine-tuning is not sufficient for harmlessness alignment. Therefore, this work constructs Natural Language Feedback (NLF) for MLLMs’ initial responses with the help of GPT-4. Then it modifies the conditional reinforcement learning to handle NLF and makes the MLLM less unsafe based on this new

method. While [Chen *et al.*, 2023c] does not change the architecture of a MLLM, [Pi *et al.*, 2024] equips the MLLM with a lightweight unsafe content detector and an output detoxifier. These modules can recognize harmful responses from the MLLM and transform them into safe ones. The experiments on LLaVA, InstructBLIP, MiniGPT4, and Qwen-vl display the effectiveness of these modules.

6 Future Research Opportunities

In this section, we make a discussion on several unresolved issues in exploring the safety of MLLMs and provide our suggestions for future research opportunities.

6.1 Reliable Safety Evaluation

More comprehensive safety benchmarks and more reasonable safety evaluation metrics are needed. Each present evaluation dataset covers a limited scope of MLLMs' safety. For example, some benchmarks just test the coarse-grained safety level of MLLMs without finer-grained safety capability partitioning. To improve the quality of an evaluation dataset, when starting to build it, we recommend several aspects that need to be considered: (1) Safety dimensions. A clear and systematic taxonomy for safety capabilities is very important. (2) Expected unsafe elements. The creator should determine expected elements representing unsafe content (e.g., natural language, concrete objects, toxic chemical molecules) and where to position these elements (e.g., images, text, or both). (3) Resources used. Potential choices are samples from existing datasets, web crawling, generation from AI models, and manual construction. (4) Volume, diversity, and quality control. These three factors influence the budget. Besides benchmarks, evaluation metrics are also worth attention. Leveraging powerful LLMs (e.g., GPT-4) is a suitable approach but carefully designed prompts are key points to define objective evaluation rules and metrics.

6.2 In-Depth Study of Safety Risk

Although many types of safety attacks prove the vulnerability of MLLMs, they lack an in-depth analysis of what enables attacks to succeed. Some works only display qualitative experiments and don't conduct quantitative evaluations. For those works that demonstrate quantitative experiments, they just select different evaluation datasets and don't make a comparison with attacks proposed by other researchers. However, a direct comparison between these attacks is important for communities to understand the detailed reasons for MLLMs' unsafe behaviors. Despite the 3 safety risks we conclude in Section 1, many questions are still waiting to be answered. For example, it is a meaningful topic to explore the impacts of MLLMs' architecture, parameters, and cross-modal training datasets on their safety ability. The requirements (e.g., computational resources, other cost) for attacks also demand a further investigation. Since MLLMs are built based on LLMs, they might inherit the flaws of LLMs (e.g., prompt sensitivity). Thus, the experience in LLMs' safety attacks can serve as a valuable reference resource.

6.3 Safety Alignment

As shown in Section 5, there are currently not many techniques to align MLLMs with human values on safety. In this part, we provide our thoughts on the safety alignment of MLLMs.

Alignment Techniques

Optimizing the process of visual instruction tuning for MLLMs' safety is a potential direction that has not gained much attention. It is not known how to build a high-quality safety-related dataset for this training phase, which can teach MLLMs to recognize unsafe queries and reject processing them. Diversity might be an influencing factor and communities may wonder about the best practices for the sample number and order (e.g., "Does it need more safety-related examples than other safety-irrelevant examples?", "Does it need to train on these two different examples separately in sequence, or a mix of them is better?").

Reinforcement learning from human feedback (RLHF) is also a promising approach, which leverages human preferences as rewards to enable a model to align with human values and has played a huge role in the safety alignment of LLMs. However, it is unclear what challenges will be encountered when applying RLHF to MLLMs for safety purposes. The construction of preference data is undoubtedly a crucial component and worth careful thinking from researchers.

Balance between Safety and Utility

When facing a malicious instruction, if MLLMs refuse to obey it, they keep their safety but lose their utility. Some researchers (e.g., [Röttger *et al.*, 2023]) have found exaggerated safety behaviors in LLMs: *misclassifying safe questions as malicious*. This incorrect classification can seriously degrade the performance of LLMs on safe prompts. Therefore, a careful balance between safety and utility is very essential when developing new safety alignment methods for MLLMs. Also, communities should understand that different applications and audiences require different balances. Here we recommend several works that may give some inspiration for researchers. PrivQA [Chen *et al.*, 2023b] proposes the *Protection Score* to measure MLLMs' ability to correctly protect specific privacy content and expose other information that does not ask for protection. [Fu *et al.*, 2023] designs a novel loss to balance between response utility and attack success. Despite this work focusing on attacks, it can inspire us to explore an effective loss that considers both the utility and safety of MLLMs.

7 Conclusion

In this paper, we attempted to present a comprehensive overview of MLLMs' safety. First, we introduce the overview of MLLMs and the understanding of safety. Afterward, we systematically review the evaluation, attack, and defense of MLLMs' safety, which demonstrates the current development status of MLLMs' safety. Finally, we delve into the existing challenges and point out some promising future research opportunities for potential researchers to explore in the future.

Acknowledgements

This work was supported by the Joint Key Project (Project No. U23A20298) and Young Scientists Project (Project No. 62206097) of the National Natural Science Foundation of China, Shanghai Pujiang Talent Program (Project No. 22PJ1403000). One of the authors, Chao Yang, is supported by the Shanghai Post-doctoral Excellent Program (Grant No. 2022234). Any opinions expressed in this material are only those of the author(s).

References

- [Askill *et al.*, 2021] Amanda Askill, Yuntao Bai, et al. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*, 2021.
- [Bagdasaryan *et al.*, 2023] Eugene Bagdasaryan, Tsung-Yin Hsieh, Ben Nassi, and Vitaly Shmatikov. Abusing Images and Sounds for Indirect Instruction Injection in Multi-Modal LLMs. *arXiv preprint*, page arXiv:2307.10490, 2023.
- [Bailey *et al.*, 2023] Luke Bailey, Euan Ong, Stuart Russell, and Scott Emmons. Image Hijacks: Adversarial Images can Control Generative Models at Runtime. *arXiv preprint*, page arXiv:2309.00236, 2023.
- [Cai *et al.*, 2019] Yitao Cai, Huiyu Cai, and Xiaojun Wan. Multi-modal sarcasm detection in Twitter with hierarchical fusion model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2506–2515, 2019.
- [Carlini *et al.*, 2023] Nicholas Carlini, Milad Nasr, Christopher A. Choquette-Choo, Matthew Jagielski, Irena Gao, Pang Wei Koh, Daphne Ippolito, Florian Tramèr, and Ludwig Schmidt. Are aligned neural networks adversarially aligned? In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Chen *et al.*, 2023a] Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. Can pre-trained vision and language models answer visual information-seeking questions? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14948–14968, 2023.
- [Chen *et al.*, 2023b] Yang Chen, Ethan Mendes, Sauvik Das, Wei Xu, and Alan Ritter. Can language models be instructed to protect personal information? *arXiv preprint arXiv:2310.02224*, 2023.
- [Chen *et al.*, 2023c] Yangyi Chen, Karan Sikka, Michael Cogswell, Heng Ji, and Ajay Divakaran. Dress: Instructing large vision-language models to align and interact with humans via natural language feedback. *arXiv preprint arXiv:2311.10081*, 2023.
- [Dong *et al.*, 2023] Yinpeng Dong, Huanran Chen, Jiawei Chen, Zhengwei Fang, Xiao Yang, Yichi Zhang, Yu Tian, Hang Su, and Jun Zhu. How robust is google’s bard to adversarial image attacks? In *R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*, 2023.
- [Fersini *et al.*, 2022] Elisabetta Fersini, Francesca Gasparini, Giulia Rizzi, Aurora Saibene, Berta Chulvi, Paolo Rosso, Alyssa Lees, and Jeffrey Sorensen. SemEval-2022 task 5: Multimedia automatic misogyny identification. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 533–549, 2022.
- [Fu *et al.*, 2023] Xiaohan Fu, Zihan Wang, Shuheng Li, Rajesh K. Gupta, Niloofar Mireshghallah, Taylor Berg-Kirkpatrick, and Earlene Fernandes. Misusing Tools in Large Language Models With Visual Adversarial Examples. *arXiv preprint*, page arXiv:2310.03185, 2023.
- [Gong *et al.*, 2023] 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.
- [Ji *et al.*, 2023a] Jiaming Ji, Mickel Liu, Juntao Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.
- [Ji *et al.*, 2023b] Yuanfeng Ji, Chongjian Ge, Weikai Kong, Enze Xie, Zhengying Liu, Zhengguo Li, and Ping Luo. Large Language Models as Automated Aligners for benchmarking Vision-Language Models. *arXiv preprint arXiv:2311.14580*, 2023.
- [Kiela *et al.*, 2020] 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 *Advances in Neural Information Processing Systems*, pages 2611–2624, 2020.
- [Li *et al.*, 2024] Mukai Li, Lei Li, Yuwei Yin, Masood Ahmed, Zhenguang Liu, and Qi Liu. Red teaming visual language models. *arXiv preprint arXiv:2401.12915*, 2024.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, 2014.
- [Lin *et al.*, 2024] Hongzhan Lin, Ziyang Luo, Bo Wang, Ruichao Yang, and Jing Ma. Goat-bench: Safety insights to large multimodal models through meme-based social abuse. *arXiv preprint arXiv:2401.01523*, 2024.
- [Liu *et al.*, 2023] Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. Query-relevant images jailbreak large multi-modal models. *arXiv preprint arXiv:2311.17600*, 2023.
- [Luo *et al.*, 2024] Haochen Luo, Jindong Gu, Fengyuan Liu, and Philip Torr. An image is worth 1000 lies: Transferability of adversarial images across prompts on vision-language models. In *The Twelfth International Conference on Learning Representations*, 2024.

- [Madry *et al.*, 2018] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [Pi *et al.*, 2024] Renjie Pi, Tianyang Han, Yueqi Xie, Rui Pan, Qing Lian, Hanze Dong, Jipeng Zhang, and Tong Zhang. Mllm-protector: Ensuring mllm’s safety without hurting performance. *arXiv preprint arXiv:2401.02906*, 2024.
- [Pramanick *et al.*, 2021a] Shraman Pramanick, Dimitar Dimitrov, Rituparna Mukherjee, Shivam Sharma, Md. Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. Detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2783–2796, 2021.
- [Pramanick *et al.*, 2021b] Shraman Pramanick, Shivam Sharma, Dimitar Dimitrov, Md. Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. MOMENTA: A multi-modal framework for detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4439–4455, 2021.
- [Qi *et al.*, 2023] Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Mengdi Wang, and Prateek Mittal. Visual adversarial examples jailbreak aligned large language models. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*, 2023.
- [Qi *et al.*, 2024] Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning aligned language models compromises safety, even when users do not intend to! In *The Twelfth International Conference on Learning Representations*, 2024.
- [Rafailov *et al.*, 2023] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.
- [Röttger *et al.*, 2023] Paul Röttger, Hannah Rose Kirk, et al. Xstest: A test suite for identifying exaggerated safety behaviours in large language models. *arXiv preprint arXiv:2308.01263*, 2023.
- [Schlarmann and Hein, 2023] Christian Schlarmann and Matthias Hein. On the adversarial robustness of multi-modal foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3677–3685, 2023.
- [Shah *et al.*, 2019] Sanket Shah, Anand Mishra, et al. Kvqa: Knowledge-aware visual question answering. In *AAAI Conference on Artificial Intelligence*, 2019.
- [Shayegani *et al.*, 2023] Erfan Shayegani, Yue Dong, and Nael Abu-Ghazaleh. Jailbreak in pieces: Compositional Adversarial Attacks on Multi-Modal Language Models. *arXiv preprint arXiv:2307.14539*, 2023.
- [Sun *et al.*, 2023] Hao Sun, Zhexin Zhang, et al. Safety assessment of chinese large language models. *arXiv preprint arXiv:2304.10436*, 2023.
- [Sun *et al.*, 2024] Lichao Sun, Yue Huang, et al. Trustllm: Trustworthiness in large language models. *arXiv preprint arXiv:2401.05561*, 2024.
- [Suryawanshi *et al.*, 2020] Shardul Suryawanshi, Bharathi Raja Chakravarthi, et al. Multimodal meme dataset (MultiOFF) for identifying offensive content in image and text. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 32–41, 2020.
- [Tu *et al.*, 2023] Haoqin Tu, Chenhang Cui, et al. How Many Unicorns Are in This Image? A Safety Evaluation Benchmark for Vision LLMs. *arXiv preprint*, page arXiv:2311.16101, 2023.
- [Van and Wu, 2023] Minh-Hao Van and Xintao Wu. Detecting and correcting hate speech in multimodal memes with large visual language model. *arXiv preprint arXiv:2311.06737*, 2023.
- [Vatsa *et al.*, 2023] Mayank Vatsa, Anubhooti Jain, and Richa Singh. Adventures of trustworthy vision-language models: A survey. *arXiv preprint arXiv:2312.04231*, 2023.
- [Wang *et al.*, 2023] Xinpeng Wang, Xiaoyuan Yi, et al. Tovi-lag: Your visual-language generative model is also an evildoer. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [Wang *et al.*, 2024] Pengyu Wang, Dong Zhang, et al. Inferaligner: Inference-time alignment for harmlessness through cross-model guidance. *arXiv preprint arXiv:2401.11206*, 2024.
- [Welbl *et al.*, 2021] Johannes Welbl, Amelia Glaese, et al. Challenges in detoxifying language models. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2447–2469, 2021.
- [Won *et al.*, 2017] Donghyeon Won, Zachary C. Steinert-Threlkeld, et al. Protest activity detection and perceived violence estimation from social media images. In *Proceedings of the 25th ACM International Conference on Multimedia*, page 786–794, 2017.
- [Wu *et al.*, 2023] Yuanwei Wu, Xiang Li, et al. Jailbreaking gpt-4v via self-adversarial attacks with system prompts. *arXiv preprint arXiv:2311.09127*, 2023.
- [Yuan *et al.*, 2023] Lifan Yuan, YiChi Zhang, et al. Bridge the gap between CV and NLP! a gradient-based textual adversarial attack framework. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7132–7146, 2023.
- [Zhang *et al.*, 2023] Zhexin Zhang, Leqi Lei, et al. Safetybench: Evaluating the safety of large language models with multiple choice questions. *arXiv preprint arXiv:2309.07045*, 2023.
- [Zhu *et al.*, 2024] Yichen Zhu, Minjie Zhu, et al. Llava- ϕ : Efficient multi-modal assistant with small language model. *arXiv preprint arXiv:2401.02330*, 2024.