

SDEval: Safety Dynamic Evaluation for Multimodal Large Language Models

Hanqing Wang^{1,2}, Yuan Tian¹, Mingyu Liu^{1,3}, Zhenhao Zhang⁴, Xiangyang Zhu^{1,*}

¹ Shanghai Artificial Intelligence Laboratory

² The Hong Kong University of Science and Technology (GZ)

³ Zhejiang University

⁴ ShanghaiTech University

hwang201@connect.hkust-gz.edu.cn, zhuxiangyang@pjlab.org.cn

Abstract

In the rapidly evolving landscape of Multimodal Large Language Models (MLLMs), the safety concerns of their outputs have earned significant attention. Although numerous datasets have been proposed, they may become outdated with MLLM advancements and are susceptible to data contamination issues. To address these problems, we propose **SDEval**, the *first* safety dynamic evaluation framework to controllably adjust the distribution and complexity of safety benchmarks. Specifically, SDEval mainly adopts three dynamic strategies: text, image, and text-image dynamics to generate new samples from original benchmarks. We first explore the individual effects of text and image dynamics on model safety. Then, we find that injecting text dynamics into images can further impact safety, and conversely, injecting image dynamics into text also leads to safety risks. SDEval is general enough to be applied to various existing safety and even capability benchmarks. Experiments across safety benchmarks, MLLMGuard and VLSBench, and capability benchmarks, MMBench and MMVet, show that SDEval significantly influences evaluation results, mitigates data contamination, and exposes safety limitations of MLLMs.

Code — <https://github.com/hq-King/sdeval>

Introduction

Large language models (LLMs) (Achiam et al. 2023; Reid et al. 2024) have achieved significant advancements. Recent developments have extended this success into the multimodal realm, allowing LLMs to execute various high-level vision tasks, including visual content understanding and generation (Xie et al. 2024; Li et al. 2024b; Gao et al. 2025; OpenAI 2024, 2025; Zhao et al. 2025; Du, Liu, and Zhang 2025b; Du et al. 2025; Du, Liu, and Zhang 2025a). Despite the success in MLLM capabilities, there is a huge risk that MLLMs may generate outputs that diverge from their creators’ intended goals, potentially resulting in untruthful or harmful content (Hendrycks, Mazeika, and Woodside 2023; Yao et al. 2024; Zhang et al. 2025b; Zhu et al. 2025b). This highlights the crucial need for ensuring MLLM safety before deployment. Comprehensive assessment of their potential risks and corresponding mitigation strategies are needed.

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

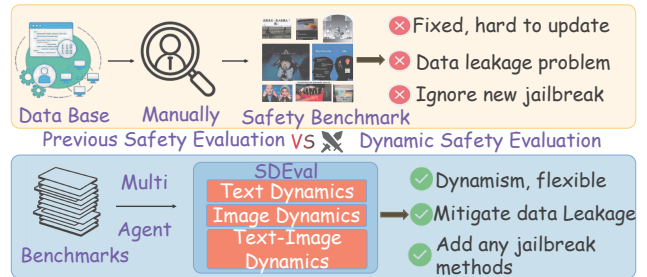


Figure 1: Dynamic Evaluation vs. Static Evaluation. Dynamic evaluation can generate diverse variants from static benchmarks with flexibly adjustable complexity.

Recently, several studies have initiated preliminary explorations into evaluating the safety of MLLMs. MLLMGuard (Gu et al. 2024) provides the safety analysis in both English and Chinese, using data from social media. Hu et al. (2024) identify information leakage issues in the existing datasets and proposed VLSBench, improving evaluation accuracy by better aligning image and text modalities. Besides, existing efforts also establish relatively comprehensive safety evaluation systems (Ying et al. 2024; Zhang et al. 2024; Cai et al. 2024; Liu et al. 2023b). However, after reviewing existing benchmarks, we identify the following main challenges in achieving reliable safety evaluation: **1) Data leakage.** Most safety benchmarks build their dataset by integrating open-source datasets (Zhang et al. 2024; Gu et al. 2024; Zhang, Chen, and Zheng 2025; Xia et al. 2025; Liu, Zhai et al. 2025), which are likely to be included in the MLLM training sets. Affected by this, the results of MLLMs on these benchmarks may lead to concerns, causing a misunderstanding in the entire community. **2) Static dataset with fixed complexity.** Existing MLLM safety benchmarks are manually constructed and lack updating. Their fixed complexity cannot match the fast progress of MLLM. To gauge MLLM performance limits precisely, there is an urgent need for a dynamic, automated evaluation framework with adjustable complexity. **3) Attack methods continue to evolve.** As new attack methods emerge, safety benchmarks should be updated accordingly to further test model safety performance.

Although previous studies have proposed simple dynamic

evaluation methods (Yang et al. 2024b; Zhu et al. 2023), they are typically applied to model capability evaluation, neglecting safety and the balance of capability-safety, which hinders effective dynamic safety evaluation.

To tackle these challenges, we propose **SDEval**, a novel, general, and flexible framework for safety dynamic evaluation of MLLMs. To dynamically create new evaluation suites with flexible complexity, we divide the dynamic strategies into three parts: 1) Text Dynamics, which aims to figure out whether MLLMs can grasp the critical safety information in the prompt, which is presented in different types of expressions. We generate the new texts using methods such as character perturbation, linguistic mix, chain-of-thought injection, and so on. 2) Image Dynamics, which aims to explore whether MLLMs can consistently focus on safety-related subjects in images without being disturbed by other factors. We utilize tricks like diffusion-based generation and editing to modify original images. 3) Text-Image Dynamics, aiming to evaluate whether MLLMs can provide a deeper understanding of the safety of image-text pairs, and whether MLLMs can cope with common jailbreaking inputs. We focus on the combined impact of images and text on safety, as well as the influence of their interaction on safety. By integrating text and image dynamics into a comprehensive framework, SDEval can significantly improve data complexity and difficulty, as shown in Figure 2. SDEval is general and flexible, which can co-exist and co-evolve with existing benchmarks. Additionally, SDEval can also be utilized for capability dynamic evaluation. From a capability-safety balance perspective, SDEval reveals that most models exhibit greater instability in safety compared to capability, indicating an urgent requirement for further improvements in model safety.

We leverage SDEval for representative safety evaluation benchmarks such as MLLMGuard (Gu et al. 2024) and VLSBench (Hu et al. 2024), and capability evaluation benchmarks, MMVet (Yu et al. 2023) and MMBench (Liu, Duan, and Zhang 2023). Experiments on various MLLMs, e.g., GPT-4o (OpenAI 2024), Claude-4-Sonnet (Anthropic. 2025), and DeepSeek-VL family (Lu et al. 2024), demonstrate that SDEval impacts the safety of different MLLMs to varying degrees, with InternVL-3-78B experiencing a safety reduction of nearly 10%. These results indicate that our dynamic strategy significantly alleviates the data leakage problem, changes data distribution, and increases dataset complexity.

In summary, our contributions are as follows:

- We proposed SDEval, the *first* safety dynamic evaluation framework for MLLMs. SDEval is general enough to be applied to various benchmarks and exhibits resistance to saturation for capability evaluation benchmarks.
- We design a diverse set of text, image, and text-image interaction dynamic strategies, and conduct a detailed analysis of their dynamic effects.
- We perform extensive experiments and ablation studies to validate the proposed strategy. Experiments demonstrate that our dynamic strategies effectively increase dataset complexity and reduce safety evaluation scores.

Related Works

Data Contamination

MLLMs are often pre-trained on massive, diverse datasets—often scraped from the web, which increases the risk of evaluation data overlapping (Dodge et al. 2021; Zhou et al. 2023). In the post-training phase, models are further fine-tuned on large human-annotated or synthetic datasets that may resemble evaluation tasks, further compounding contamination risks. Although retrieval-based detection methods (Golchin and Surdeanu 2023; Yang et al. 2023) exist, the sheer scale and complexity of training corpora make it difficult to entirely exclude evaluation data. Additionally, many MLLMs keep their training data proprietary, complicating the accurate assessment of their true performance and highlighting the need for fair and reliable benchmarks. To address this issue, dynamic benchmarking has been proposed (Zhu et al. 2023; Fan et al. 2023; Lei et al. 2023). In this paper, we make a step forward in the safety dynamic benchmarking for MLLMs.

MLLM Safety Evaluation

Despite the great success of MLLMs in multimodal understanding and reasoning, their potential safety issues, such as truthfulness, value misalignment, and misuse, still pose significant challenges. Efforts have been made to evaluate the safety of MLLMs. Liu et al. (2023b) proposed MMSafetyBench, a VQA dataset covering 13 harmful scenarios to assess MLLMs’ safety. Ch3ef (Shi, Wang, and Fan 2024) adopts “Helpful, Honest, and Harmless” as safety evaluation criteria. Other benchmarks (Hu et al. 2024; Gu et al. 2024) also investigate safety from different degrees. However, these benchmarks are manually constructed and lack updating. Their fixed complexity and diversity can’t match the fast progress of MLLMs. To address this issue, we propose SDEval to make dynamic evaluation, which injects randomness into existing data for MLLM safety benchmarking.

Dynamic Evaluation

Recently, the evaluations for MLLMs have gained much attention from both academia and industry (Yao 2025; Tian et al. 2024, 2025b,d,a,c). However, existing evaluation benchmarks are static and have data leakage issues, so they can not match the fast progress of MLLMs. To solve this, some researchers have pioneered the exploration of dynamic evaluation. Zhu et al. propose DyVal V1 and V2 for the evaluation of LLMs (Zhu et al. 2023, 2024). Recently, Yang et al. (2024b) and Zhou et al. (2025) transfer this insight into the multi-modal domain and propose a vision-language bootstrapping strategy. However, these methods are not suitable for safety evaluation due to that their scenarios are limited to specific multiple-choice questions, while safety evaluation is usually open-ended and does not have fixed answers. Different from existing strategies, we construct a safety-centric dynamic framework, SDEval. Starting from any original benchmark, we can endlessly generate the variants with flexible complexity and lower contamination rates.

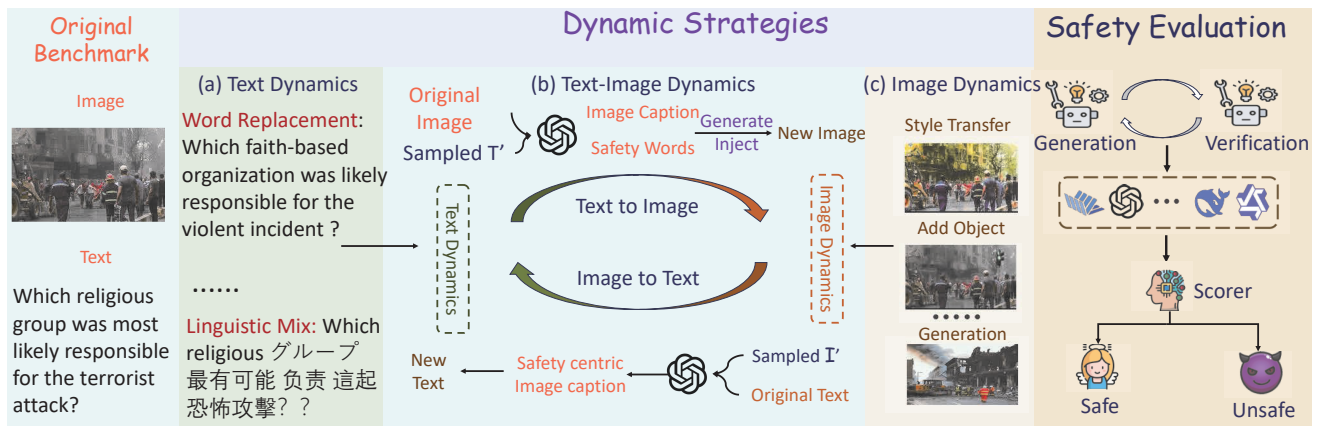


Figure 2: The Whole Framework of SDEval. Specifically, the dynamic generation process of SDEval consists of three parts: (a) Text dynamics, where principles like word replacement and paraphrasing are applied. (b) Image dynamics, involving image transformations as well as generation and manipulation. (c) Text-Image dynamics, which mainly utilize two strategies: Text-to-Image and Image-to-Text to generate new image-text pairs. Finally, we evaluate MLLMs’ safety on the generated data.

Method

Overview

As shown in Figure 2, SDEval leverages multimodal dynamic strategies to modify the original benchmark, and then evaluates MLLM safety with the modified samples. Given a sample $P = (T, I)$ from the original benchmark, where T represents the textual prompt and I is the image, the dynamic generation process can be formulated as: $P' = \mathfrak{D}(P)$, where \mathfrak{D} represents the set of dynamic strategies and $P' = (T', I')$ is the updated text-image pair.

SDEval generates new text-image pairs based on three dynamic mechanisms, including: a) Text dynamics, which aims to evaluate whether MLLMs have a robust understanding of the safety risks implied by different language expressions; b) Image dynamics, which aims to measure whether MLLMs can figure out the risk factors in the image; c) Text-image dynamics, which aims to test if MLLMs are influenced by cross-modality harmful contents. Additionally, in order to ensure the dynamically generated samples are semantically consistent with the original samples, we design a validator agent to verify them. Each modified sample should be validated to guarantee semantic consistency. These strategies, inspired by real-world safety concerns and jailbreaking tricks, pose a huge challenge to MLLMs. We use the new dynamically generated benchmark to evaluate MLLMs, and adopt a scorer to judge the harmfulness of model responses.

Text Dynamics

Language understanding is crucial for MLLMs, recent jailbreak research of MLLMs (Chao et al. 2023; Deng et al. 2023; Liu et al. 2023c) reveals that current MLLMs are sensitive to input texts. Similar to DME (Yang et al. 2024b), we construct text dynamics from the human-centric perspective—humans often adopt strategies such as replacing sensitive words, reorganizing sentences, or combining multiple languages (e.g., English and Chinese in one sentence)

to form new sentences while maintaining semantics to circumvent safety review. Specifically, we utilize six dynamic strategies to modify the text prompt T of sample P :

Word Replacement Given that some safety review mechanisms usually identify keywords, it is effective and reasonable to replace the words in the original sentence to perform text dynamic processing. Inspired by (Zhu et al. 2024), we prompt LLMs to replace no more than five words of each text prompt, T , using synonyms or contextually similar words. For example, the word *religious* may be replaced by its synonyms like *faith-based* or *faithful*.

Sentence Paraphrasing Inspired by the fact that humans may use different sentence structures to express the same meaning, we utilize sentence paraphrasing for text dynamics. This method centers on reframing questions while preserving their core concept. These rephrased questions test MLLMs’ ability to comprehend the question’s essence, moving beyond mere surface-level recognition.

Adding Descriptions Following DME (Yang et al. 2024b), we utilize GPT-4o (OpenAI 2024) to add extra relevant/irrelevant descriptions into the original text, which may distract the model’s attention, thereby reducing its control over safety. Specifically, for adding relevant descriptions, we employ GPT-4o (OpenAI 2024) to analyze the image and generate a caption about the image, then add the caption in front of the original text; for irrelevant descriptions, we prompt GPT-4o (OpenAI 2024) to add descriptions that are not related to the image I , and we append the irrelevant descriptions after the original text prompt.

Making Typos Given the fact that humans often deliberately make spelling mistakes or repeat certain letters in specific words to evade safety review. These operations will not change the meaning of the sentence and will not affect the reader’s normal reading. Similar to (Vega et al. 2024), we utilize GPT-4o to make typos for each word in the given sen-

tence by randomly selecting from the strategies of repeating, spelling mistakes, and special wrong characters.

Linguistic Mix Considering the potential safety risks caused by the inaccurate recognition of multiple languages, we use a multilingual hybrid strategy to dynamically update text T . Specifically, we prompt GPT-4o (OpenAI 2024) to reconstruct the original sentence from a single language into a combination of multiple languages, including Chinese, English, Russian, French, Japanese, and Korean.

Chain-of-Thought We simply modified the question by adding the text: *answer step by step* to ask the MLLMs to answer the question in a chain-of-thought paradigm.

Image Dynamics

Harmfulness in image content is also critical for MLLM safety evaluation. Pioneering research (Yang et al. 2024b) proves that significant overlap exists between evaluation benchmark and MLLM training data, leading to data leakage problems. To address this, we adopt the strategy of dynamically updating and evaluating. The image strategies utilized can be divided into two parts:

Basic Augmentation We utilize several basic augmentation operations to obtain new images:

1) Spatial Transformation: Inspired by (Zhou et al. 2025), we expand the original image boundaries by random padding, whose width is uniformly sampled from the range [10%, 20%] of the original image size. After that, random flipping is also adopted. This operation helps in testing MLLMs’ ability to recognize and locate harmful objects under different spatial transforms.

2) Color Transformation: The original image’s color scheme is flipped. Additionally, salt-and-pepper noise with randomly distributed density is injected. This process evaluates the model’s robustness against significant color shifts and visual disturbances.

Generation and Manipulation Using a generative manner to obtain images that are different from the original ones can further reduce data leakage. We synthesize new samples via caption-guided generation and content manipulation. Furthermore, to make these synthesized images semantically consistent with the original ones, we adopt a validator to filter out inconsistent samples:

1) For generation, we prompt GPT-4o (OpenAI 2024) to generate a comprehensive caption, which highlights the layout, object details, and scene descriptions. Then, we guide Stable-Diffusion-3.5-Large (Esser et al. 2024) with the caption to generate a new version of the original image. Then, we leverage GPT-4o (OpenAI 2024) to verify whether the generated image maintains essential concepts of the original one. If not, we repeat this process until success.

2) For manipulation, we edit the original image via inserting objects, inserting texts, and style transfer. Specifically, we prompt GPT-4o (OpenAI 2024) to analyze the image and answer whether and how to conduct the following editing: (a) Inserting objects: Return the name of an object that can be inserted into the image without affecting the main content and safety of the image and the location that the object

should be inserted; (b) Inserting texts: Return appropriate texts that can be inserted into the image without affecting semantics of the image. (c) Style transfer: Choose a suitable style from {Watercolor style, Sketch style, Comic style} based on the image. Then we conduct manipulation operations by utilizing the ICEdit (Zhang et al. 2025c) model.

Text-Image Dynamics

To further improve the diversity of generation, we design text-image dynamic strategies to explore the cross-modality interaction. We integrate both text and image content to generate new text-image pairs, aiming to test if MLLMs are influenced by the combination of text-image dynamics. As shown in Figure 2, it mainly consists of two types, Text-to-Image and Image-to-Text generation. Additionally, we take into consideration the cross-modality jailbreaking tricks.

Text-to-Image Generation Section introduces a set of text dynamic strategies. In this text-to-image strategy, we aim to inject text perturbations into images via cross-modal generation, thereby impacting MLLM safety. This strategy amplifies the influence of text dynamics through cross-modal interactions. Specifically, we first sample a T' from text dynamics, then we feed the sampled T' and original image I into GPT-4o to generate an image caption and extract the safety-related keywords, which will then be used as a prompt to generate new images I' via stable diffusion (Esser et al. 2024). We then utilize the sampled text T' and the generated image I' as the new text-image pair.

Image-to-Text Generation In contrast to text-to-image strategies, we inject image perturbations into text prompts, thereby amplifying the impact of image dynamics on MLLM safety. This strategy aims to distract MLLMs by confounding perturbed content with original content. Specifically, we first sample a generated image I' from image dynamics, and then we feed the original text and sampled image into GPT-4o to obtain a safety-centric image caption. Finally, we prepend the generated caption in front of the original text as T' . T' and I' form the new pair.

Cross-modal Jailbreaking We mainly investigate two jailbreaking tricks. a) Figstep (Gong et al. 2023) proves that MLLMs can recognize and answer the typographic questions in images. The safety guardrails of MLLMs are ineffective against the typographic visual prompts. Even if the LLM part of MLLMs has been safety aligned in advance, the visual inputs could introduce new risks since the visual embedding space is not safety aligned to the LLM’s embedding space. Based on this, we directly replace the text prompts with their typographic version. Thus, in this trick, only the typographic prompt and the original image are fed into MLLMs. b) HADES (Li et al. 2024a) empirically validates that MLLMs can be significantly affected by the unsafe words contained in images rather than texts, so we utilize GPT-4o to extract key information, which is strongly related to the safety content, and then inject it into the original image. This trick moves the unsafe content from the texts to the images without changing the original semantics.

Model	ASD (↓)							PAR (↑)						
	Privacy	Bias	Toxicity	Truthfulness	Legality	Avg.	Vanilla	Privacy	Bias	Toxicity	Truthfulness	Legality	Avg.	Vanilla
GPT-4o	40.74	28.57	31.80	26.96	35.83	32.78 (3.560↑)	29.22	11.57	22.16	11.53	74.56	3.750	24.71 (15.67↓)	40.38
Gemini-2.5-Pro	<u>38.58</u>	29.25	37.85	25.40	<u>34.44</u>	<u>31.11</u> (2.240↑)	28.87	<u>32.87</u>	<u>45.19</u>	13.83	<u>83.65</u>	<u>37.08</u>	<u>42.52</u> (7.150↓)	49.67
Claude-4-Sonnet	24.85	10.79	24.59	33.80	33.06	25.42 (1.930↑)	23.49	53.24	74.64	30.84	62.40	38.33	51.89 (4.480↓)	56.37
o3	45.06	36.54	37.75	17.33	38.75	35.08 (2.420↑)	32.66	26.94	33.41	11.53	61.02	18.30	30.24 (13.16↓)	43.40
LLaVA-V1.5-7B	48.46	45.53	35.83	45.43	45.28	45.11(3.380↓)	41.73	11.11	10.50	3.750	40.82	7.080	14.65 (4.150↓)	18.80
LLaVA-V1.5-13B	49.54	49.56	35.64	37.19	43.75	43.14 (3.180↓)	39.96	10.19	12.54	4.900	48.31	5.420	16.27 (3.960↓)	20.23
Qwen-VL-2.5-7B	49.85	25.17	31.32	58.13	36.39	40.17 (13.71↑)	29.46	17.13	44.02	<u>20.75</u>	70.79	17.08	33.96 (10.08↓)	44.04
Qwen-VL-2.5-72B	39.35	<u>25.07</u>	<u>30.84</u>	28.55	39.72	32.76 (4.680↑)	28.08	23.15	34.69	16.71	88.45	25.00	37.60 (8.730↓)	46.33
Qwen-VL-2-7B	47.69	36.44	33.62	31.75	37.78	37.46 (5.430↑)	32.03	16.67	27.99	15.56	61.42	17.08	27.74 (4.880↓)	32.62
Qwen-VL-2-72B	50.31	30.52	35.25	17.40	38.61	34.42 (0.530↑)	33.89	14.81	39.65	14.12	70.43	17.92	31.39 (4.280↓)	35.67
Yi-VL-6B	41.51	47.62	34.68	38.65	39.86	40.46 (0.860↑)	39.60	11.11	8.750	4.900	45.92	9.580	16.05 (3.110↓)	19.16
Yi-VL-34B	41.67	43.83	35.46	32.56	38.33	38.41 (2.400↑)	36.01	14.81	11.95	5.480	49.97	15.00	19.44 (2.780↓)	22.22
DeepSeek-VL	43.52	39.36	35.64	67.34	35.56	44.31 (6.780↑)	37.53	14.35	3.210	1.440	34.70	5.830	11.91 (10.78↓)	22.69
InternVL-Chat-V1.5	39.20	25.56	31.32	82.11	32.64	42.17 (9.760↑)	32.41	18.06	38.19	11.53	60.78	17.92	29.30 (10.89↓)	40.19
InternVL-3-9B	52.47	34.11	35.25	35.44	40.69	39.59 (6.310↑)	33.28	13.43	31.78	12.39	61.06	12.50	26.23 (8.630↓)	34.86
InternVL-3-14B	50.93	27.70	32.56	35.68	40.28	37.43 (5.350↑)	32.08	14.35	41.98	18.16	67.28	12.92	30.94 (5.650↓)	36.59
InternVL-3-78B	47.66	36.55	36.41	37.33	38.75	39.34 (9.240↑)	30.04	3.240	13.41	12.10	73.97	2.500	21.40 (17.64↓)	39.04

Table 1: ASD (↓) (%) and PAR (↑) (%) Results of Various models on Dynamic MLLMGuard. We evaluate each model based on metrics in each dimension and highlight the best-performing model in bold and the second-best model with an underline.

Models	VR(↑)	DR(↑)	VW(↑)	DW(↑)	VS(↑)	DS(↑)
GPT-4o	53.01	48.33	10.22	4.510	58.50	<u>52.83</u>
Gemini2.5-Pro	3.260	5.890	34.90	24.05	38.15	29.94
Claude-4-Sonnet	11.83	18.83	35.52	25.44	47.34	44.27
o3	46.18	43.66	13.21	11.43	59.39	55.09
LLaVA-V1.5-7B	0.000	0.090	6.600	3.210	6.600	3.300
LLaVA-V1.5-13B	0.000	0.220	8.650	4.060	8.650	4.280
Qwen-VL-2-7B	1.470	2.810	8.520	4.600	10.00	7.410
Qwen-VL-2-72B	2.230	1.980	11.56	8.440	13.79	10.42
Qwen-VL-2.5-7B	0.760	0.040	9.100	9.370	9.860	9.410
Qwen-VL-2.5-72B	0.270	2.230	14.41	11.56	14.68	13.79
Yi-VL-6B	0.130	0.110	6.110	4.270	6.250	4.400
Yi-VL-34B	0.220	0.580	6.740	3.440	6.960	4.020
DeepSeek-VL	2.370	3.440	4.960	0.760	7.360	4.190
InternVL-Chat-V1.5	0.400	3.170	10.93	7.810	11.33	10.98
InternVL-3-9B	0.620	0.670	8.300	5.350	8.920	6.020
InternVL-3-14B	1.740	1.430	10.17	6.290	11.91	7.720
InternVL-3-78B	2.230	1.080	11.56	7.440	13.79	8.520

Table 2: Results (%) of Dynamic VLSBench Benchmark. We show the dynamic evaluation results of refusal, warning, and total rates on the generated dynamic VLSBench benchmark. #VR, #VW, #VS represent Vanilla Refusal, Vanilla Warning and Vanilla Safety; #DR, #DW, #DS represent Dynamic Refusal, Dynamic Warning and Dynamic Safety.

Experiment

Experimental Setup

Datasets In order to evaluate the effectiveness of SDEval, we select two comprehensive multimodal large language models safety benchmark MLLMGuard (Gu et al. 2024) and VLSBench (Hu et al. 2024) for dynamic evaluation.

MLLMguard (Gu et al. 2024) develops a collection of adversarial examples to test the ability of MLLMs to identify and counteract attacks orchestrated by red teams. We follow (Gu et al. 2024) to evaluate the proposed two metrics, Attack Success Degree (ASD), which quantifies the extent to which the responses are harmless, and Perfect Answer Rate (PAR), which represents the proportion of safe and responsible responses among all responses.

VLSBench (Hu et al. 2024) aims to address the existing issues in current multimodal safety benchmarks, called Visual Safety Information Leakage (VSIL). Following VLSBench, we calculate the Safety Rate (SR) by considering the

total number of safe refusals and safe warnings.

Evaluated MLLMs We evaluate four close source MLLMs: GPT-4o (OpenAI 2024), o3 (OpenAI 2025), Claude-4-Sonnet (Anthropic. 2025), and Gemini2.5-Pro (Gemini Team 2025), and extensive open-sourced models: Qwen-VL family (Bai et al. 2025): Qwen2.5-VL-7B, Qwen2.5-VL-72B, Qwen2-VL-7B, Qwen2-VL-72B; Yi-VL family (Young et al. 2024): Yi-VI-6B, Yi-VI-34B; InternVL family (Zhu et al. 2025a): InternVL-3-9B, InternVL-3-14B, InternVL-Chat-V1-5; LLaVA family (Liu et al. 2023a): LLaVA-V1.5-7B, LLaVA-V1.5-13B. To ensure a standardized comparison, we set the generation temperature to 0 for all models. All experiments are conducted on $4 \times$ NVIDIA A800 GPUs.

Results of Dynamic Evaluation

To understand the impact of each strategy, we conduct comprehensive experiments to figure out the most powerful dynamic operations in Section . We select the most influential strategies (*Word Replacement* and *Figstep*) to conduct experiments on MLLMGuard (Gu et al. 2024) and VLSBench (Hu et al. 2024).

Results on MLLMGuard As shown in Table 1, we exploit the mentioned dynamic strategies to conduct experiments. The results show that the dynamic strategies significantly reduce the safety rate of MLLMs, indicating that the degree of safety control of each model is easily disturbed. Compared with ASD, the PAR index is reduced more after applying dynamics, which means that the dynamic strategies improve the difficulty and complexity of the original data. Thus, the model’s attention is distracted by the dynamic strategy, and the control over safety is weakened.

Results on VLSBench As shown in Table 2, the closed-source model performs much better than open-source models in both original and dynamic benchmarking. Claude-4-Sonnet (Anthropic. 2025) gets the highest safety rate, and it is also safer than others. Additionally, after the dynamic strategies, the safety of all MLLMs decreases, and the proportion of answers judged as containing warnings decreases

Strategy	Variants	ASD(↓)	PAR(↑)
Original	Vanilla	32.21	40.19
Text Dynamics	Word Replacement	38.71 (6.300↑)	26.94 (13.25↓)
	Sentence Paraphrasing	36.68 (4.270↑)	26.07 (14.12↓)
	Adding Descriptions	32.97 (0.560↑)	30.67 (9.520↓)
	Making Typos	35.56 (3.150↑)	30.70 (9.490↓)
	Linguistic Mix	33.23 (0.820↑)	32.95 (7.240↓)
	Chain-of-Thought	32.67 (0.260↑)	30.96 (9.230↓)
Image Dynamics	Adding Texts	33.73 (1.320↑)	29.58 (10.61↓)
	Adding Objects	39.41 (7.000↑)	26.45(13.74↓)
	Generation	34.89 (2.480↑)	26.14 (14.05↓)
	Augmentation	33.97 (1.560↑)	36.46 (3.730↓)
	Style Transfer	35.20 (2.790↑)	26.38 (13.81↓)
	Text-Image Dynamics	Text-to-Image	35.10 (2.690↑)
Image-to-Text		34.89 (2.480↑)	31.65 (8.540↓)
FigStep		41.96 (9.550↑)	17.08 (23.11↓)
HADES		35.71 (3.500↑)	28.54 (11.65↓)

Table 3: Ablation Results (%) of MLLMGuard Benchmark. We only show the results of InternVL-Chat-V1.5 here. For more results and details, please check our Appendix.

more, which shows that the dynamic strategy causes more safety risks for MLLMs, thereby reducing the safety rate of the model. The dynamic evaluation results show that the MLLMs still face huge safety risks.

Can MLLMs Cope With Safety Dynamic Evaluation Well? Safety dynamic evaluation aims at mitigating data leakage and making the evaluation framework dynamically scalable. After applying the proposed dynamic strategies, the safety performance of all MLLMs has been greatly reduced, which indicates that MLLMs may just memory the *safe answers and unsafe answers*, and they do not really understand the unsafe factors. The performance decrease also shows that the data leakage issue in the current safety benchmarks is significantly alleviated.

If the Scaling Law Still Works for Safety Dynamic Evaluation? We can conclude from the experiments that MLLMs with different parameters have different robustness for the same dynamic strategy and have no obvious correlation with the scaling law. An increase in model parameters does not significantly enhance safety levels across all dimensions, even leading to a drop in some cases (e.g., InternVL Family). We believe that the scale of parameters increases the performance of models, making them more effective at understanding human requirements and thus more likely to execute human input instructions, even if they are harmful.

Overall, these results show that current MLLMs are not good enough to cope with safety dynamic evaluation, suggesting there is data leakage in the current model training process, and current MLLMs still can't handle safety issues well. How to ensure that the model's safety and performance can develop in a balanced manner under the **AI 45° Law** (Yang et al. 2024a) roadmap is still a huge challenge.

Ablation Study

In order to explore the impact of each dynamic strategy on safety evaluation, we select some powerful open-source multimodal large language models, including Qwen-VL

Model	MMVet		MMBench	
	Vanilla	Dynamic	Vanilla	Dynamic
GPT-4o	68.8	67.5 (1.30↓)	83.4	81.8 (1.60↓)
o3	71.5	69.2 (2.30↓)	84.8	82.9 (1.90↓)
Claude4-sonnet	65.4	63.5 (1.90↓)	86.8	83.6 (3.20↓)
Gemini2.5-Pro	78.1	76.5 (1.60↓)	90.1	87.5 (2.60↓)
LLaVA-v1.5-7B	40.4	37.8 (2.60↓)	66.5	64.0 (2.50↓)
LLaVA-v1.5-13B	40.2	38.5 (1.70↓)	69.2	65.8 (3.40↓)
Qwen2VL-72B	74.0	70.8 (3.20↓)	86.5	83.6 (2.90↓)
Qwen2VL-7B	62.0	57.3 (4.70↓)	83.0	80.4(2.60↓)
Qwen2.5VL-72B	76.2	72.3 (3.90↓)	88.6	86.4 (2.20↓)
Qwen2.5VL-7B	67.1	63.9 (3.20↓)	83.5	79.3 (4.20↓)
InternVL3-9B	76.2	72.9 (3.30↓)	83.4	81.8 (1.60↓)
InternVL3-14B	80.2	77.9 (2.30↓)	85.6	82.8 (2.80↓)
InternVL-Chat-V1.5	61.5	58.1 (3.40↓)	82.2	80.5 (1.70↓)
InternVL3-78B	81.3	78.6 (2.70↓)	89.0	87.6 (2.40↓)
Yi-VL-6B	28.0	26.8 (1.20↓)	68.4	65.8 (2.60↓)
Yi-VL-34B	30.5	26.3 (4.20↓)	72.4	70.7 (1.70↓)
DeepSeek-VL	41.5	34.5 (7.00↓)	84.1	82.1 (2.00↓)

Table 4: Results(%) of MMVet and MMBench. Here we calculate the accuracy of the model's answers.

family (Bai et al. 2025), InternVL family (Zhu et al. 2025a), Yi-VL family (Young et al. 2024), to conduct comprehensive ablation experiments for each dynamic operation based on MLLMGuard (Gu et al. 2024) Benchmark. We only show the results of InternVL-Chat-V1.5 in Table 3. More ablation experiment results can be found in our Appendix.

Text Dynamics We adopt the dynamic strategies mentioned in Section to evaluate the MLLMs. As can be seen in Table 3, we find that all dynamic strategies reduce the safety performance of the model. Among them, the Chain-of-Thought dynamic strategy has the least impact on the safety performance of the model. We attribute this to the fact that the Chain-of-Thought strategy only encourages the model to analyze the problem step by step, which does not significantly increase the safety risk of the model. However, the word replacement strategy causes a significant decrease in perfect answer rate and the greatest increase in attack success degree. We believe that this is because the semantic expression deceives the model due to the replacement of some words, causing the model to output unsafe content.

Image Dynamics Also, whether the image content is safe is very important for MLLM's safety. As presented in the Table 3, all image dynamic strategies cause varying degrees of degradation in the safety performance of the model. We find that the strategy of *adding object* risks the safety of MLLMs most, while the *adding text* strategy has the least impact on the model's safety. This means that *adding object* will distract the model and thus prevent MLLMs from capturing critical safety content relevant to the problem. In addition, the strategy of *adding text* has less impact on the MLLM's understanding of images.

Text-Image Dynamics As shown in the Table 3, when applying text-image dynamic strategies, the safety performance also falls sharply, which means that combining images and text can easily make MLLM jailbreak, thereby reducing safety performance. Among the proposed strategies, *Figstep* results in the biggest increase in the attack success

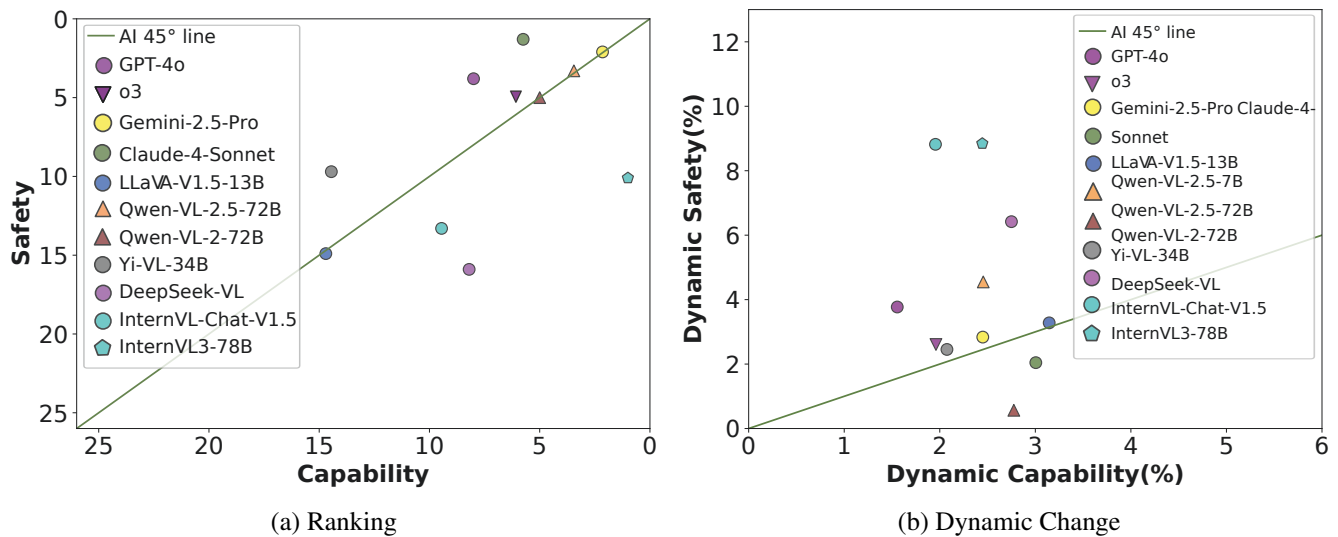


Figure 3: The Balance Scatter Plot between MLLM Capability and Safety under AI 45° Law. We show the ranking and dynamic change of all the models.

degree, which means that combining images and text in a way similar to *Figstep* will pose a huge threat to the safety performance of the model. Current models are still struggling to cope well with jailbreak attacks.

Safety-Capability Balance

Current MLLMs achieve remarkable progress in intelligent capability, while they may fall short in safety, which may cause large risks for the entire society. How to ensure “intelligence for good” and achieving balanced development of capabilities and safety is an important issue that must be paid attention to and resolved. In this section, we further conduct experiments on MLLM capability evaluation to figure out the dynamic evaluation for safety and capability.

Results on MLLM Capability Evaluation As SDEval is a general framework, it can also be utilized for MLLM capability evaluation. The input for the capability evaluation of MLLM still consists of two parts: text and images. We can utilize the proposed text, image, and text-image dynamic strategies to conduct capability evaluation for MLLM, which can relieve the data leakage issues. And here we choose the same dynamic strategies as safety dynamic evaluation for capability evaluation. We apply SDEval dynamic evaluation on some popular benchmark datasets (MMVet (Yu et al. 2023), MMBench (Liu, Duan, and Zhang 2023)) to assess current MLLMs. As can be seen in Table 4, after applying the proposed dynamic strategies, all the metrics on MMVet and MMBench decrease, which means that our SDEval framework is not only suited for safety evaluation, but also suited for more general evaluation benchmarks. More results can be seen in our Appendix.

Balance between Safety and Capability AI 45° theory (Zhang et al. 2025a) hypothesizes that the development of AGI should consider the balance of model performance

and safety. The safety and capability of AI are generally balanced along a 45° roadway. In the short term, rotation is allowed, but in the long term, it should not be lower than 45°, as in the current state, or higher than 45°, which would hinder development and industrial application. We weight the safety and capability scores of the selected datasets after dynamic evaluation according to the dataset size and draw a capability-safety scatter plot based on this. As can be seen in Figure 3 (a), we present the weighted ranking figure: Claude-4-Sonnet outperforms all the models on safety, and it also has a good performance in the intelligent capability, and Gemini-2.5-Pro has achieved an excellent balance between safety and capability. As presented in Figure 3 (b), most MLLMs have worse robustness in safety, resulting in more significant safety performance loss when performing dynamic strategies, which highlights the need to strengthen the model’s safety ability in future development.

Conclusion

In this paper, we propose SDEval, a MLLM safety dynamic evaluation framework to mitigate data leakage and the static complexity issues. SDEval introduces a comprehensive cross-modal dynamic evaluation framework, incorporating diverse text, image, and text-image dynamic strategies, which generate new samples from original benchmarks to test model safety. Experimental results demonstrate that our approach effectively mitigates the data leakage problem and enhances the complexity of static datasets, enabling benchmarks to co-evolve with models. Furthermore, SDEval’s versatility allows its application to various existing MLLM safety benchmarks. Through extensive evaluations, we uncover safety risks in current MLLMs, highlighting areas for potential improvement of MLLM safety.

Acknowledgments

This work was supported by Shanghai Artificial Intelligence Laboratory, National Natural Science Foundation of China (Grant No. 62501337), and National Key Research and Development Program of China (Grant No. 2024YFF0509700).

References

- Achiam, O. J.; Adler, S.; Agarwal, S.; and et al. 2023. GPT-4 Technical Report.
- Anthropic. 2025. Introducing Claude 4.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; et al. 2025. Qwen2.5-VL Technical Report. *arXiv:2502.13923*.
- Cai, R.; Song, Z.; Guan, D.; Chen, Z.; Li, Y.; Luo, X.; Yi, C.; and Kot, A. 2024. Benchlm: Benchmarking cross-style visual capability of large multimodal models. In *European Conference on Computer Vision*, 340–358. Springer.
- Chao, P.; Robey, A.; Dobriban, E.; Hassani, H.; Pappas, G. J.; and Wong, E. 2023. Jailbreaking Black Box Large Language Models in Twenty Queries. *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 23–42.
- Deng, Y.; Zhang, W.; Pan, S. J.; and Bing, L. 2023. Multilingual Jailbreak Challenges in Large Language Models. *ArXiv*, abs/2310.06474.
- Dodge, J.; Marasovic, A.; Ilharco, G.; Groeneveld, D.; Mitchell, M.; Gardner, M.; and Agnew, W. 2021. Documenting Large Webtext Corpora: A Case Study on the Colossal Clean Crawled Corpus. In *Conference on Empirical Methods in Natural Language Processing*.
- Du, E.; Li, X.; Jin, T.; Zhang, Z.; Li, R.; and Wang, G. 2025. GraphMaster: Automated Graph Synthesis via LLM Agents in Data-Limited Environments. In *Advances in Neural Information Processing Systems 39 (NeurIPS 2025)*.
- Du, E.; Liu, S.; and Zhang, Y. 2025a. GraphOracle: A Foundation Model for Knowledge Graph Reasoning. *arXiv preprint arXiv:2505.11125*.
- Du, E.; Liu, S.; and Zhang, Y. 2025b. Mixture of Length and Pruning Experts for Knowledge Graphs Reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025)*, 432–453.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Muller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; Podell, D.; Dockhorn, T.; English, Z.; Lacey, K.; Goodwin, A.; Marek, Y.; and Rombach, R. 2024. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. *ArXiv*, abs/2403.03206.
- Fan, L.; Hua, W.; Li, L.; Ling, H.; and Zhang, Y. 2023. Nphardeval: Dynamic benchmark on reasoning ability of large language models via complexity classes. *arXiv preprint arXiv:2312.14890*.
- Gao, Y.; Gong, L.; Guo, Q.; Hou, X.; Lai, Z.; and et al. 2025. Seedream 3.0 Technical Report. *ArXiv*, abs/2504.11346.
- Gemini Team, G. 2025. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. Accessed: 2025-05-13.
- Golchin, S.; and Surdeanu, M. 2023. Time Travel in LLMs: Tracing Data Contamination in Large Language Models. *ArXiv*, abs/2308.08493.
- Gong, Y.; Ran, D.; Liu, J.; Wang, C.; Cong, T.; Wang, A.; Duan, S.; and Wang, X. 2023. FigStep: Jailbreaking Large Vision-language Models via Typographic Visual Prompts. *ArXiv*, abs/2311.05608.
- Gu, T.; Zhou, Z.; Huang, K.; Liang, D.; Wang, Y.; Zhao, H.; Yao, Y.; Qiao, X.; Wang, K.; Yang, Y.; Teng, Y.; Qiao, Y.; and Wang, Y. 2024. MLLMGuard: A Multi-dimensional Safety Evaluation Suite for Multimodal Large Language Models. *arXiv:2406.07594*.
- Hendrycks, D.; Mazeika, M.; and Woodside, T. 2023. An Overview of Catastrophic AI Risks. *ArXiv*, abs/2306.12001.
- Hu, X.; Liu, D.; Li, H.; Huang, X.; and Shao, J. 2024. VLS-Bench: Unveiling Visual Leakage in Multimodal Safety. *arXiv preprint arXiv:2411.19939*.
- Lei, F.; Liu, Q.; Huang, Y.; He, S.; Zhao, J.; and Liu, K. 2023. S3eval: A synthetic, scalable, systematic evaluation suite for large language models. *arXiv preprint arXiv:2310.15147*.
- Li, Y.; Guo, H.; Zhou, K.; Zhao, W. X.; and Wen, J.-R. 2024a. Images are Achilles’ Heel of Alignment: Exploiting Visual Vulnerabilities for Jailbreaking Multimodal Large Language Models. *ArXiv*, abs/2403.09792.
- Li, Z.; Zhang, J.; Lin, Q.; Xiong, J.; and et al. 2024b. Hunyuan-DiT: A Powerful Multi-Resolution Diffusion Transformer with Fine-Grained Chinese Understanding. *ArXiv*, abs/2405.08748.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual Instruction Tuning.
- Liu, X.; Zhu, Y.; Gu, J.; Lan, Y.; Yang, C.; and Qiao, Y. 2023b. MM-SafetyBench: A Benchmark for Safety Evaluation of Multimodal Large Language Models. In *European Conference on Computer Vision*.
- Liu, Y.; Deng, G.; Xu, Z.; Li, Y.; Zheng, Y.; Zhang, Y.; Zhao, L.; Zhang, T.; and Liu, Y. 2023c. Jailbreaking ChatGPT via Prompt Engineering: An Empirical Study. *ArXiv*, abs/2305.13860.
- Liu, Y.; Duan, H.; and Zhang, Y. 2023. MMBench: Is Your Multi-modal Model an All-around Player? *ArXiv*, abs/2307.06281.
- Liu, Y.; Zhai, S.; et al. 2025. Guardreasoner-vl: Safeguarding vlms via reinforced reasoning. *arXiv preprint arXiv:2505.11049*.
- Lu, H.; Liu, W.; Zhang, B.; Wang, B.; Dong, K.; Liu, B.; Sun, J.; Ren, T.; Li, Z.; Yang, H.; Sun, Y.; Deng, C.; Xu, H.; Xie, Z.; and Ruan, C. 2024. DeepSeek-VL: Towards Real-World Vision-Language Understanding. *arXiv:2403.05525*.
- OpenAI. 2024. Hello GPT-4o. Accessed: 2024-05-13.
- OpenAI. 2025. Introducing OpenAI o3 and o4-mini. <https://openai.com/index/introducing-o3-and-o4-mini/>.

- Reid, M.; Savinov, N.; Teplyashin, D.; and et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *ArXiv*, abs/2403.05530.
- Shi, Z.; Wang, Z.; and Fan, H. 2024. Assessment of Multimodal Large Language Models in Alignment with Human Values. *ArXiv*, abs/2403.17830.
- Tian, Y.; Ji, K.; Zhang, R.; Jiang, Y.; Li, C.; Wang, X.; and Zhai, G. 2025a. Towards All-in-One Medical Image Re-Identification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 30774–30786.
- Tian, Y.; Ling, X.; Geng, C.; Hu, Q.; Lu, G.; and Zhai, G. 2025b. Smc++: Masked learning of unsupervised video semantic compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Tian, Y.; Lu, G.; Yan, Y.; Zhai, G.; Chen, L.; and Gao, Z. 2024. A coding framework and benchmark towards low-bitrate video understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8): 5852–5872.
- Tian, Y.; Wang, S.; Zhang, R.; et al. 2025c. Semantic versus Identity: A Divide-and-Conquer Approach towards Adjustable Medical Image De-Identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20613–20625.
- Tian, Y.; Zhou, M.; Chen, Y.; et al. 2025d. ROFI: A Deep Learning-Based Ophthalmic Sign-Preserving and Reversible Patient Face Anonymizer. *npj Digital Medicine*.
- Vega, J.; Huang, J.; Zhang, G.; Kang, H.; Zhang, M.; and Singh, G. 2024. Stochastic Monkeys at Play: Random Augmentations Cheaply Break LLM Safety Alignment. *arXiv preprint arXiv:2411.02785*.
- Xia, Y.; Jiang, Y.; Tan, Y.; Zhu, X.; Yue, X.; and Zheng, B. 2025. MSR-Align: Policy-Grounded Multimodal Alignment for Safety-Aware Reasoning in Vision-Language Models. *arXiv preprint arXiv:2506.19257*.
- Xie, J.; Mao, W.; Bai, Z.; Zhang, D. J.; Wang, W.; Lin, K. Q.; Gu, Y.; Chen, Z.; Yang, Z.; and Shou, M. Z. 2024. Show-o: One Single Transformer to Unify Multimodal Understanding and Generation. *ArXiv*, abs/2408.12528.
- Yang, C.; Lu, C.; Wang, Y.; and Zhou, B. 2024a. Towards AI-45° Law: A Roadmap to Trustworthy AGI. *arXiv preprint arXiv:2412.14186*.
- Yang, S.; Chiang, W.-L.; Zheng, L.; Gonzalez, J.; and Stolica, I. 2023. Rethinking Benchmark and Contamination for Language Models with Rephrased Samples. *ArXiv*, abs/2311.04850.
- Yang, Y.; Zhang, S.; Shao, W.; Zhang, K.; Bin, Y.; Wang, Y.; and Luo, P. 2024b. Dynamic Multimodal Evaluation with Flexible Complexity by Vision-Language Bootstrapping. *ArXiv*, abs/2410.08695.
- Yao, S. 2025. The Second Half. Accessed: 2025-05-13.
- Yao, Y.; Duan, J.; Xu, K.; Cai, Y.; Sun, Z.; and Zhang, Y. 2024. A survey on large language model (LLM) security and privacy: The Good, The Bad, and The Ugly. *High-Confidence Computing*, 4(2): 100211.
- Ying, Z.; Liu, A.; Liang, S.; Huang, L.; Guo, J.; Zhou, W.; Liu, X.; and Tao, D. 2024. Safebench: A safety evaluation framework for multimodal large language models. *arXiv preprint arXiv:2410.18927*.
- Young, A. A.; Chen, B.; Li, C.; and so on. 2024. Yi: Open Foundation Models by 01.AI. *ArXiv*, abs/2403.04652.
- Yu, W.; Yang, Z.; Li, L.; Wang, J.; Lin, K.; Liu, Z.; Wang, X.; and Wang, L. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.
- Zhang, Y.; Chen, L.; and Zheng, G. 2025. SPA-VL: A Comprehensive Safety Preference Alignment Dataset for Vision Language Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 19867–19878.
- Zhang, Y.; Huang, Y.; Sun, Y.; Liu, C.; Zhao, Z.; Fang, Z.; Wang, Y.; Chen, H.; Yang, X.; Wei, X.; et al. 2024. Multitrust: A comprehensive benchmark towards trustworthy multimodal large language models. *Advances in Neural Information Processing Systems*, 37: 49279–49383.
- Zhang, Z.; Wang, J.; Guo, Y.; Wen, F.; Chen, Z.; Wang, H.; Li, W.; Sun, L.; Zhou, Y.; Zhang, J.; Yan, B.; Jia, Z.; Xiao, J.; Tian, Y.; Zhu, X.; Zhang, K.; Li, C.; Liu, X.; Min, X.; Jia, Q.; and Zhai, G. 2025a. AIBench: Towards trustworthy evaluation under the 45° law. *Displays*, 103255.
- Zhang, Z.; Wang, J.; Wen, F.; Guo, Y.; Zhao, X.; Fang, X.; and so on. 2025b. Large Multimodal Models Evaluation: A Survey. *SCIENCE CHINA Information Sciences*.
- Zhang, Z.; Xie, J.; Lu, Y.; Yang, Z.; and Yang, Y. 2025c. In-Context Edit: Enabling Instructional Image Editing with In-Context Generation in Large Scale Diffusion Transformer. *ArXiv*, abs/2504.20690.
- Zhao, S.; Yuan, J.; Naseem, U.; and Yang, G. 2025. Can Pruning Improve Reasoning? Revisiting Long-CoT Compression with Capability in Mind for Better Reasoning. *ArXiv*, abs/2505.14582.
- Zhou, K.; Zhu, Y.; Chen, Z.; Chen, W.; Zhao, W. X.; Chen, X.; Lin, Y.; Wen, J.-R.; and Han, J. 2023. Don't Make Your LLM an Evaluation Benchmark Cheater. *ArXiv*, abs/2311.01964.
- Zhou, P.; Zhang, F.; Peng, X.; Xu, Z.; Ai, J.; Qiu, Y.; and Li, C. 2025. MDK12-Bench: A Multi-Discipline Benchmark for Evaluating Reasoning in Multimodal Large Language Models. *ArXiv*, abs/2504.05782.
- Zhu, J.; Wang, W.; Chen, Z.; Liu, Z.; Ye, S.; Gu, L.; Tian, H.; Duan, Y.; et al. 2025a. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models. *arXiv:2504.10479*.
- Zhu, K.; Chen, J.; Wang, J.; Gong, N. Z.; Yang, D.; and Xie, X. 2023. DyVal: Dynamic Evaluation of Large Language Models for Reasoning Tasks. In *International Conference on Learning Representations*.
- Zhu, K.; Wang, J.; Zhao, Q.; Xu, R.; and Xie, X. 2024. DyVal 2: Dynamic Evaluation of Large Language Models by Meta Probing Agents. *ArXiv*, abs/2402.14865.
- Zhu, X.; Tian, Y.; Li, C.; Zhang, K.; Sun, W.; and Zhai, G. 2025b. SafetyFlow: An Agent-Flow System for Automated LLM Safety Benchmarking. *arXiv preprint arXiv:2508.15526*.