

SAFETHINK: A KEY TO SAFETY IN MULTI-MODAL LARGE REASONING MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-modal large language models (MLLMs) are being increasingly fine-tuned with reinforcement learning (RL) to improve reasoning, yielding strong gains on complex benchmarks. Yet recent studies show that such reasoning-oriented fine-tuning weakens safety alignment, making models far more vulnerable to jailbreak attacks. We trace this vulnerability to a misspecified objective: RL fine-tuning maximizes task accuracy while ignoring safety constraints. To address this, we introduce SAFETHINK, an inference-time steering method that enforces safety constraints directly within the chain-of-thought. At each reasoning step, SAFETHINK scores partial traces with a safety reward and, when unsafe content is detected, projects the trajectory back into the safe set via lightweight textual feedback (e.g., “Wait, think safely”). This mechanism preserves accuracy on benign inputs while reinstating robustness under adversarial prompts. Our experiments across diverse safety robustness benchmarks demonstrate that SAFETHINK significantly improves safety without sacrificing reasoning capabilities. For example, against jailbreak attacks on OpenVLLinker-7B, SAFETHINK reduces the attack success rate by 44.57% compared to the base reasoning model and by 18.32% over the existing baseline.

1 INTRODUCTION

Multi-modal large reasoning models (MLRMs) have achieved remarkable performance across a wide range of reasoning-intensive domains, including visual question answering (Yue et al., 2024; Xiao et al., 2024; Lu et al., 2023), multi-step planning (Ma et al., 2024), and scientific problem solving (Yue et al., 2024; Lu et al., 2022). In particular, reinforcement learning (RL) based methods and chain-of-thought supervision (Guo et al., 2025; openai, 2024) enable models to produce explicit reasoning traces, transforming base multi-modal LLMs into deliberate “System 2” problem solvers (Wei et al., 2022; Li et al., 2025). These fine-tuning based approaches have pushed MLRMs beyond direct-answer baselines, establishing new state-of-the-art results on structured reasoning benchmarks.

Reasoning Tax: Challenge of Safety in MLRMs. These advances in reasoning come with a critical cost: *a significant decline in safety robustness*, i.e., the ability of models to resist adversarial or harmful instructions. In practice, a safety-robust model should refuse malicious queries, such as requests for hate speech, disallowed content, or cyberattack instructions, regardless of prompt phrasing. However, recent studies reveal a systematic trade-off: reasoning-focused RL fine-tuning substantially increases a model’s vulnerability to such adversarial prompts (Fang et al., 2025; Huang et al., 2025a; Jiang et al., 2025; Zhou et al., 2025). This phenomenon, termed the *reasoning tax* (Fang et al., 2025), highlights that enabling reasoning in MLRMs makes them much more susceptible to jailbreaking. For example, R1-OneVision (Yang et al., 2025), fine-tuned with GRPO (Guo et al., 2025) for improved reasoning, exhibits a 41.2% higher attack success rate (ASR) than its non-reasoning base model Qwen-2.5-VL (Yang et al., 2024) on MM-SafetyBench (Liu et al., 2024). This trade-off between reasoning and safety raises a fundamental question: *Why does reasoning-focused RL fine-tuning degrade safety alignment?*

Our key insight. We argue that the root cause of this trade-off lies in how reasoning fine-tuning is currently formulated. Existing RL-based approaches (Guo et al., 2025) optimize solely for reasoning accuracy on curated datasets, effectively treating the problem as unconstrained optimization. In doing so, they ignore the safety behaviors that base models initially possessed, allowing reasoning trajectories to drift toward unsafe outputs whenever this improves task accuracy. From a principled




Figure 1: **Safety-Reasoning Trade-off in Multi-modal Large Reasoning Models (MLRMs).** Left: We show an example where a reasoning MLRM generates unsafe content, unlike its base model; existing defense SafePath (Jeung et al., 2025) also fails. In contrast, our inference-time method SAFETHINK steers the reasoning process toward safe outputs. Right: Quantitative results showing that SAFETHINK reduces attack success rates (improving safety) while preserving reasoning accuracy.

standpoint, reasoning fine-tuning should instead be cast as a constrained optimization problem: maximizing accuracy subject to explicit safety constraints. Unfortunately, re-training large multi-modal reasoning models under such a framework is computationally infeasible and often degrades reasoning quality (Huang et al., 2025a). This leads us to a natural question: *if retraining is not an option, can safety alignment be recovered directly at inference time?*

Prior Efforts and Our Approach. We address this challenge by shifting the focus from training-time fixes to inference-time interventions. Prior attempts, such as truncating the reasoning process (Jiang et al., 2025) or prepending fixed safety prefixes (Jeung et al., 2025), offer only limited protection and often suppress the very reasoning capacity they were meant to preserve (Huang et al., 2025a; Fang et al., 2025). Our approach builds on a key observation: although RL fine-tuning drifts models toward unsafe reasoning patterns, it does not erase their safety priors. Because reasoning MLRMs remain relatively close to their base distributions in KL divergence (Shenfeld et al., 2025), safety can be effectively restored by monitoring and steering the chain-of-thought at inference time. We instantiate this idea in SAFETHINK, an iterative inference-time safety steering mechanism that enforces step-wise safety constraints through appending textual feedback. As shown in Figure 1, SAFETHINK consistently reduces attack success rates while maintaining strong reasoning performance, outperforming existing inference-time defenses. We summarize our contributions as follows.

- **Understanding safety brittleness in MLRMs.** We formally characterize why RL-based reasoning fine-tuning degrades safety: it is implicitly treated as unconstrained optimization, prioritizing accuracy while disregarding safety alignment.
- **Inference-time safety steering.** We propose SAFETHINK, an iterative mechanism that monitors and steers the chain-of-thought with lightweight textual feedback, effectively restoring safety without retraining.
- **Comprehensive empirical evaluation.** We evaluate SAFETHINK across a diverse suite of jailbreak benchmarks, including Hades (Li et al., 2024), MMSafetyBench (Liu et al., 2024), FigStep (Gong et al., 2023), and text-based jailbreak attacks (Luo et al., 2024). Our results demonstrate substantial safety improvements on recent MLRMs, such as R1-OneVision (Yang et al., 2025), OpenVL-Thinker (Deng et al., 2025), Vision-R1 (Huang et al., 2025b), VLAA-Thinker (Chen et al., 2025), LlamaV-o1 (Thawakar et al., 2025), and LLaVA-CoT (Xu et al., 2024), while maintaining strong reasoning capabilities.

2 PROBLEM FORMULATION

Notations and Reasoning Model Objective. To improve reasoning capabilities in multi-modal large language models (MLLMs), recent work has widely adopted RL-based methods such as GRPO (Guo et al., 2025; openai, 2024), which explicitly encourage models to produce chain-of-thought reasoning traces before arriving at the final answer. This has given rise to a new class of models, multi-modal large reasoning models (MLRMs) (Deng et al., 2025; Chen et al., 2025; Thawakar et al., 2025; Xu

et al., 2024; Yao et al., 2024). Formally, the reasoning process of an MLRM π_θ can be described as

$$x_{\text{input}} \rightarrow z \rightarrow y, \quad (1)$$

where $x_{\text{input}} = [I, x]$ denotes the input, with $I \in \mathcal{I}$ the visual input (image) and $x = \{x_1, x_2, \dots, x_N\}$ the textual prompt consisting of N tokens ($x_i \in \mathcal{V}$ for vocabulary $\mathcal{V}$). The model first produces a reasoning (or thinking) trace $z \sim \pi_\theta(\cdot | x_{\text{input}})$ and then a final answer $y \sim \pi_\theta(\cdot | x_{\text{input}}, z)$. In practice, π_θ is obtained by RL fine-tuning a safe multimodal base model $\pi_{\theta_0}^{\text{safe}}$. For example, OpenAI’s o1 (openai, 2024) and DeepSeek-R1 (Guo et al., 2025) adopt RL to reward accurate reasoning traces. The RL objective for training reasoning models is given by

$$\max_{\theta} \mathbb{E}_{x_{\text{input}} \in \mathcal{D}, z \sim \pi_\theta(\cdot | x_{\text{input}}), y \sim \pi_\theta(\cdot | x_{\text{input}}, z)} [R_{\text{acc}}(x_{\text{input}}, y)], \quad (2)$$

where $\mathcal{D}$ denotes the set of problems, $R_{\text{acc}}(x, y)$ evaluates task accuracy (e.g., correctness of y). In addition, a format reward $R_{\text{format}}(x, y)$ is added in practice to enforce structured outputs of the form $x_{\text{input}} \rightarrow z \rightarrow y$. Thus, an MLRM starts from a safe multimodal base model and is fine-tuned with RL (often using GRPO (Guo et al., 2025) or similar) to generate both reasoning traces and accurate final answers.

The Safety–Reasoning Trade-off. While RL fine-tuning improves reasoning accuracy, it weakens safety alignment. Empirical studies show a clear pattern: models with stronger reasoning ability become significantly more vulnerable to adversarial “jailbreak” attacks (Fang et al., 2025; Huang et al., 2025a; Jiang et al., 2025). For instance, on MM-SafetyBench (Liu et al., 2024), reasoning-tuned R1-OneVision produces harmful content in response to prompts that its base model reliably refuses, with attack success rates rising by over 40%. As illustrated in Figure 1 (right), reasoning-oriented fine-tuning substantially raises attack success rates, undermining the safety guarantees of the original base model. Together, these findings bring us to two central questions: (1) What causes RL fine-tuning for reasoning to compromise safety? (2) Is it possible to recover safety at inference time without resorting to costly retraining? We explore both questions next.

The Root Cause: Unconstrained RL Fine-tuning. We hypothesize that the safety brittleness of reasoning models stems from a misspecification in the RL fine-tuning objective. As written in Equation 2, current methods maximize task-specific rewards such as accuracy, effectively treating reasoning fine-tuning as an unconstrained RL fine-tuning problem. In doing so, they ignore the safety behaviors inherited from the base model, allowing policies that achieve high accuracy at the cost of unsafe reasoning trajectories. A more principled formulation is a *constrained fine-tuning* problem: maximize reasoning accuracy subject to explicit safety requirements,

$$\max_{\theta} \mathbb{E}_{x_{\text{input}} \in \mathcal{D}, z \sim \pi_\theta(\cdot | x_{\text{input}}), y \sim \pi_\theta(\cdot | x_{\text{input}}, z)} [R_{\text{acc}}(x_{\text{input}}, y)] \quad (3)$$

$$\text{s.t. } \mathbb{E}_{x_{\text{input}} \in \mathcal{D}, z \sim \pi_\theta(\cdot | x_{\text{input}}), y \sim \pi_\theta(\cdot | x_{\text{input}}, z)} [R_{\text{safe}}(x_{\text{input}}, [z, y])] \geq \delta, \quad (4)$$

where R_{safe} evaluates the safety of the reasoning trace z and the final answer y , and δ sets a minimum safety threshold. In principle, retraining under this formulation could enforce safety, but in practice, it is computationally expensive and could degrade reasoning quality (Huang et al., 2025a).

3 PROPOSED APPROACH: SAFETHINK

Motivation. From Section 2, the root cause of safety brittleness is that reasoning fine-tuning optimizes only for task accuracy, ignoring the safety constraint in Equation 3. As a result, under adversarial prompts, reasoning traces often drift into unsafe regions, meaning the constraint $\mathbb{E}[R_{\text{safe}}] \geq \delta$ is violated at inference. Further, retraining with explicit safety constraints is computationally infeasible, and it can also hurt reasoning quality (Huang et al., 2025a). We therefore ask: *can the missing constraint be enforced directly during inference?*

Key Idea. Our answer is SAFETHINK, an inference-time intervention that monitors the chain-of-thought as it is generated and projects unsafe steps back towards safety by appending corrective textual cues. The central observation is that the safety reward R_{safe} can be evaluated not just on final answers, but on *any prefix of the reasoning trace*. This allows us to detect unsafe reasoning before it propagates and correct it on the fly.

Algorithm 1 SAFETHINK: Safety steering via “Wait, think safely”

Require: MLRM π_θ ; safety reward model R_{safe} ; adversarial input x_{adv} ; safety steering prompt $p_{\text{safe}} \leftarrow \text{“Wait, think safely”}$; intervention budget K ; safety threshold τ .

```

1:  $t \leftarrow 1, k \leftarrow 0$ 
2: while not EOT do
3:    $\mathbf{s}_t \leftarrow [x_{\text{adv}}, z_{<t}]$  ▷ Current context
4:   Sample next reasoning step:  $z_t \sim \pi_\theta(\cdot | \mathbf{s}_t)$  until newline or EOT is generated
5:    $r_t \leftarrow R_{\text{safe}}(x_{\text{adv}}, z_{\leq t})$  ▷ Safety score for generated step
6:   if  $r_t \geq \tau$  then ▷ If generated step is safe
7:      $\mathbf{s}_{t+1} \leftarrow [\mathbf{s}_t, z_t]$ 
8:      $k \leftarrow 0$ 
9:   else if  $k < K$  then ▷ If generated step is unsafe
10:     $\mathbf{s}_{t+1} \leftarrow [\mathbf{s}_t, p_{\text{safe}}]$  ▷ Append safety steering prefix
11:     $k \leftarrow k + 1$ 
12:   else
13:     $\mathbf{s}_{t+1} \leftarrow [\mathbf{s}_t, z_t]$ 
14:    $t \leftarrow t + 1$ 
15:  $y \sim \pi_\theta(\cdot | \mathbf{s}_t)$  ▷ Generate final answer
16: return  $y$ 

```

Proposed Mechanism. Let an adversarial input be x_{adv} , and let the reasoning trace be $z = (z_1, \dots, z_T, [\text{EOT}])$ where z_t is the t -th step and $[\text{EOT}]$ is the end-of-thinking token. At each step t , the model proposes $z_t \sim \pi_\theta(\cdot | \mathbf{s}_t)$ given the context $\mathbf{s}_t = [x_{\text{adv}}, z_{<t}]$. We then evaluate safety:

$$r_t = R_{\text{safe}}(x_{\text{adv}}, z_{\leq t}),$$

where R_{safe} returns a normalized score reflecting how safe the partial reasoning is. For the safety reward, R_{safe} , we leverage publicly available harmless reward models, such as from Reward Bench (Lambert et al., 2024), which assign high scores to safe text and low scores to unsafe continuations. If $r_t \geq \tau$, the step is accepted (where τ is a safety threshold); otherwise, it is replaced with a corrective prefix p_{safe} that instructs the model to redirect its reasoning. Formally,

$$\mathbf{s}_{t+1} = \begin{cases} [\mathbf{s}_t, z_t], & r_t \geq \tau, \\ [\mathbf{s}_t, p_{\text{safe}}], & r_t < \tau. \end{cases}$$

To prevent excessive intervention, we cap the number of consecutive substitutions by a budget K . Once the reasoning trace is completed, i.e., upon generation of $[\text{EOT}]$, the final answer is generated as $y \sim \pi_\theta(\cdot | x_{\text{adv}}, z)$.

Safety Steering via Textual Feedback. The design of the corrective prefix p_{safe} is crucial. Its purpose is to nudge the model back to safe reasoning without derailing coherence. Inspired by prior work showing that textual cues like “Wait, think step by step” can improve reasoning performance (Muennighoff et al., 2025), we adapt this idea for safety. Specifically, we use:

$$p_{\text{safe}} := \text{“Wait, think safely.”}$$

This prefix acts as a projection step, reorienting the reasoning trajectory toward safe continuations whenever unsafe reasoning is detected. In Section 5, we provide ablations on different variants of the safety prefix. In effect, each step of reasoning is constrained to lie within the safety set $\{z : R_{\text{safe}}(x_{\text{adv}}, z) \geq \tau\}$, and unsafe proposals are projected back to this set via p_{safe} . This approach ensures that thinking remains safe, while accuracy on benign inputs is preserved, since projection occurs only when unsafe reasoning is detected. We describe our detailed proposed approach in Algorithm 1.

4 EXPERIMENTS

4.1 EXPERIMENTAL DETAILS

Jail-break Datasets. To investigate the safety vulnerabilities of MLRMs, we carry out a comprehensive evaluation using both text-based and image-based jailbreak attacks:

Model	Defense Strategy	Noise			SD			Nature			Blank			Average
		Template	Persuade	Logic	Template	Persuade	Logic	Template	Persuade	Logic	Template	Persuade	Logic	
R1-Onevision	Original	50.23	44.87	72.97	58.64	37.91	67.57	56.72	48.38	39.19	49.72	38.44	70.27	50.62
	ZeroThink	40.36	33.54	48.65	43.22	42.38	54.05	42.13	35.12	29.73	42.47	39.93	37.84	40.24
	LessThink	40.41	37.82	39.19	34.26	45.67	52.70	27.43	26.87	29.73	43.05	45.92	39.19	39.18
	SafePath	18.74	34.19	54.15	28.61	36.48	45.95	17.39	21.57	32.43	20.45	38.12	40.54	34.68
	AdaShield	40.85	25.67	44.59	36.71	22.39	47.30	36.15	27.83	36.49	40.48	31.22	41.89	37.26
	SAFEThiNK (Ours)	15.42	12.13	4.05	13.77	16.72	4.05	13.05	8.59	6.76	19.68	7.34	2.70	10.36
OpenVLThinker	Original	35.81	43.14	74.03	35.94	44.35	67.31	41.28	35.17	52.94	29.13	36.86	58.94	45.69
	ZeroThink	15.11	27.75	21.54	16.04	22.03	40.11	10.27	18.29	25.64	11.77	27.03	20.71	21.25
	LessThink	26.91	31.83	17.89	19.69	18.53	13.28	15.57	13.75	21.35	20.00	24.97	13.20	19.89
	SafePath	25.54	28.43	16.35	19.19	19.27	16.97	15.43	11.13	20.11	22.48	29.12	12.45	19.44
	AdaShield	26.77	23.33	25.05	17.93	14.58	27.68	29.15	11.48	31.10	18.70	18.37	21.80	21.79
	SAFEThiNK (Ours)	1.46	0.73	1.24	6.90	2.08	1.05	1.84	1.29	1.98	4.37	3.10	1.29	1.12
VLAA-Thinker	Original	27.13	23.47	16.22	18.42	22.56	31.08	12.44	9.65	20.27	23.11	17.22	20.27	20.38
	ZeroThink	7.42	17.33	18.92	10.58	17.89	24.32	12.11	10.33	16.22	8.67	18.56	17.57	14.85
	LessThink	26.55	23.21	13.51	24.22	17.14	17.57	20.17	16.89	12.16	21.37	17.23	14.86	18.16
	SafePath	16.78	20.47	20.27	19.33	18.21	22.97	16.91	9.67	20.27	17.46	19.33	13.51	18.31
	AdaShield	16.21	14.11	13.51	9.43	11.23	14.86	12.22	9.25	17.57	22.19	9.44	13.51	13.36
	SAFEThiNK (Ours)	6.14	10.28	4.05	2.77	7.17	4.05	1.31	6.23	5.41	2.17	5.28	1.35	4.39
Vision-R1	Original	40.17	38.42	51.36	48.63	34.29	54.06	38.51	29.37	40.56	38.26	36.68	52.72	41.84
	ZeroThink	28.43	25.78	31.08	26.41	26.19	40.54	18.36	23.22	22.97	22.64	25.15	33.78	27.05
	LessThink	22.71	28.34	31.08	17.46	25.62	40.54	17.83	20.29	39.19	20.87	31.58	44.59	28.34
	SafePath	32.44	38.21	31.08	36.63	36.52	45.95	32.86	32.73	33.78	33.39	40.47	36.49	35.88
	AdaShield	33.76	13.42	20.27	30.31	14.52	21.62	29.98	12.64	16.22	35.87	11.33	22.97	21.91
	SAFEThiNK (Ours)	4.12	4.28	2.02	7.52	4.13	1.43	1.09	5.21	0.00	4.24	5.12	0.00	3.56
LlamaV-o1	Original	51.12	61.45	78.38	56.21	66.09	78.38	47.19	47.81	77.03	54.87	69.42	85.14	61.33
	ZeroThink	26.73	37.12	59.46	35.22	43.67	62.16	32.44	32.87	54.05	33.19	47.56	75.68	44.98
	LessThink	31.45	38.92	63.51	43.08	45.21	52.70	39.12	42.87	63.51	48.34	44.76	79.73	50.78
	SafePath	39.22	44.18	59.46	49.11	46.33	70.27	34.78	41.92	66.22	57.44	62.11	64.86	52.12
	AdaShield	33.76	45.12	68.92	42.09	51.23	67.57	40.21	43.33	68.92	42.77	55.18	64.86	52.79
	SAFEThiNK (Ours)	7.23	8.64	5.67	6.21	3.32	5.28	4.78	2.16	6.76	8.11	8.11	2.70	5.74
LLaVA-CoT	Original	54.26	29.84	29.50	54.83	34.27	37.30	50.01	28.97	27.79	64.37	42.48	52.86	42.21
	ZeroThink	48.52	18.65	41.62	47.30	30.59	27.79	46.54	23.93	36.71	52.08	38.85	45.78	38.20
	LessThink	44.94	17.49	37.29	46.17	30.18	24.67	42.83	20.05	39.08	43.00	33.74	32.77	34.35
	SafePath	36.46	15.80	33.16	42.06	22.90	21.34	40.91	18.24	35.14	36.86	27.27	31.17	30.11
	AdaShield	9.05	23.37	40.61	7.33	12.49	28.02	10.29	8.66	27.75	13.84	11.22	31.18	18.65
	SAFEThiNK (Ours)	2.78	2.91	10.16	3.17	3.01	13.00	7.43	5.25	12.85	4.32	3.52	10.50	6.57

Table 1: **Evaluation on Text-Based Jailbreak Attacks.** We report the Attack Success Rate (ASR) for various baseline defense strategies across recent MLRMs on text-based jailbreak attacks (Luo et al., 2024). The best results (lowest ASR) are highlighted in **bold**. All values are reported in %.

1. Text-based Attacks. To evaluate text-based jailbreaks, we use the JailbreakV-28K benchmark (Luo et al., 2024), which integrates adversarial text prompts with diverse visual inputs. The text prompts span template-style, persuasive, and logic-oriented variants, while the paired images may include random noise, blank images, natural photos, or samples generated with stable diffusion. We test across all text-image combinations to ensure comprehensive evaluation.

2. Image-based Attacks. For image-based jailbreaks, we evaluate on three benchmarks that embed malicious instructions within visual inputs:

(1) *Hades* (Li et al., 2024): This dataset contains harmful instructions spanning five categories: Violence, Financial Crime, Privacy Violation, Self-Harm, and Animal Abuse, each paired with a semantically relevant harmful image. The images are generated through three strategies: (a) harmful visuals produced with stable diffusion (SD), (b) typography renderings of instructions overlaid (+TYPO), and (c) adversarially perturbed composites (+ADV) designed to further increase the likelihood of eliciting harmful responses.

(2) *FigStep* (Gong et al., 2023): This benchmark transforms harmful prompts into typographic images. Specifically, malicious queries from ten restricted categories in MM-SafetyBench (Liu et al., 2024) are paraphrased into noun-based forms (e.g., “Steps to,” “List of,” “Methods to”) and rendered as embedded text within images.

(3) *MM-SafetyBench* (Liu et al., 2023a): This benchmark evaluates robustness across 13 restricted categories defined in AI safety usage guidelines (Achiam et al., 2023; Inan et al., 2023). Each malicious query is paired with three image modalities: (a) Stable-diffusion images (SD), generated to be semantically relevant to the query; (b) Typography (TYPO) images, which directly render the harmful instruction as text within the image; and (c) SD+TYPO composites, combining generated visuals with embedded textual captions.

Models Used. We perform evaluation on 6 state-of-art open-source MLRMs: R1-Onevision-7B (Yang et al., 2025), OpenVLThinker-7B (Deng et al., 2025), VLAA-Thinker-Qwen2.5VL-7B (Chen et al., 2025), Vision-R1-7B (Huang et al., 2025b), LlamaV-o1 (Thawakar et al., 2025), and LLaVA-CoT (Xu et al., 2024). For safety evaluation, we adopt the Llama-Guard 3 model¹ (Llama Team, 2024). We set the safety threshold $\tau = 0$. In practice, the reward scores can be normalized on a held-out validation

¹meta-llama/llama-guard-3-8B

Model	Defense Strategy	Animal			Financial			Privacy			Self-Harm			Violence			Average
		SD	+TYPO	+ADV	SD	+TYPO	+ADV	SD	+TYPO	+ADV	SD	+TYPO	+ADV	SD	+TYPO	+ADV	
R1-Onevision	Original	46.00	53.33	54.67	72.00	80.00	77.33	69.33	74.67	72.00	49.33	52.00	61.33	90.00	93.33	90.67	69.07
	ZeroThink	34.67	42.67	44.00	69.33	73.33	70.67	58.67	60.00	66.00	44.00	53.33	54.67	81.33	88.00	90.00	62.04
	LessThink	32.00	46.00	42.67	64.00	70.00	73.33	49.33	57.33	62.00	38.67	50.00	37.33	82.67	80.00	82.00	57.82
	SafePath	34.00	33.33	36.00	50.67	60.00	65.33	46.67	44.00	40.00	36.00	37.33	29.33	78.00	82.00	80.00	50.18
	AdaShield	24.00	30.67	32.00	45.33	53.33	54.67	42.67	38.67	40.00	28.00	30.67	28.00	66.67	72.00	70.00	43.78
	SafeTHINK (Ours)	2.00	2.67	2.67	0.00	4.00	4.00	2.00	1.33	2.00	6.67	8.00	8.00	12.00	14.67	14.67	5.65
OpenVLThinker	Original	38.67	48.00	50.67	70.67	80.00	80.00	60.00	72.00	70.00	58.00	58.00	60.00	82.00	85.33	86.67	66.67
	ZeroThink	26.67	33.33	44.00	46.00	62.00	57.33	37.33	56.00	60.00	30.67	32.00	36.00	62.00	62.00	60.00	47.02
	LessThink	10.00	32.00	33.33	26.67	42.67	42.67	21.33	28.00	20.00	24.00	37.33	30.00	50.67	61.33	53.33	34.22
	SafePath	10.00	16.00	21.33	33.33	42.67	44.00	22.67	24.00	30.67	17.33	21.33	16.00	46.00	48.00	44.00	29.16
	AdaShield	12.00	12.00	14.67	24.00	22.67	22.00	18.00	16.00	12.00	17.33	10.67	13.33	38.67	41.33	41.33	21.07
	SafeTHINK (Ours)	2.67	4.00	0.00	1.33	6.00	2.00	0.00	2.67	1.33	0.00	2.67	0.00	6.00	2.67	1.33	2.18
VLAA-Thinker	Original	13.33	14.67	20.00	25.33	32.00	38.67	18.00	22.00	22.67	13.33	14.00	10.00	62.00	58.67	65.33	28.67
	ZeroThink	13.33	10.67	12.00	14.67	33.33	42.67	13.33	20.00	24.00	10.00	10.67	6.67	54.67	54.67	62.00	25.51
	LessThink	9.33	17.33	16.00	18.00	28.00	36.00	8.00	25.33	21.33	9.33	8.00	10.67	50.00	54.00	56.00	24.49
	SafePath	17.33	18.00	22.67	17.33	18.67	24.00	12.00	12.00	14.00	8.00	14.00	9.33	56.00	50.67	60.00	23.60
	AdaShield	10.00	10.67	5.33	8.00	10.67	8.00	8.00	4.00	4.00	5.33	5.33	6.00	37.33	37.33	33.33	12.89
	SafeTHINK (Ours)	1.33	0.00	0.00	2.00	1.33	1.33	0.00	0.00	0.00	2.00	2.00	2.00	6.67	8.00	6.67	2.22
Vision-R1	Original	46.00	48.00	53.33	60.00	70.67	73.33	52.00	66.67	58.00	50.67	54.00	54.00	74.67	84.00	84.00	61.96
	ZeroThink	46.00	53.33	61.33	52.00	56.00	62.00	52.00	61.33	62.00	57.33	60.00	60.00	72.00	62.00	81.33	60.98
	LessThink	46.00	61.33	64.00	50.00	65.33	73.33	34.67	69.33	77.33	45.33	60.00	58.67	66.67	74.00	70.00	61.07
	SafePath	40.00	44.00	50.00	57.33	66.00	66.00	32.00	46.67	40.00	50.67	45.33	48.00	82.00	86.00	84.00	55.87
	AdaShield	17.33	10.00	18.00	21.33	28.00	24.00	16.00	20.00	18.67	22.00	25.33	13.33	26.00	32.00	36.00	21.87
	SafeTHINK (Ours)	8.00	5.33	6.67	6.00	9.33	9.33	5.33	9.33	9.33	2.67	12.00	4.00	10.67	13.33	14.00	8.35
LlamaV-o1	Original	46.00	53.33	62.67	62.00	66.00	74.67	70.00	69.33	74.00	58.67	56.00	52.00	81.33	84.00	92.00	66.80
	ZeroThink	50.67	54.00	64.00	61.33	72.00	62.00	57.33	62.00	66.67	44.00	45.33	50.67	76.00	74.67	80.00	61.38
	LessThink	49.33	50.00	52.00	66.00	72.00	73.33	66.00	70.67	73.33	40.00	38.00	44.00	54.00	44.00	50.67	56.22
	SafePath	38.00	38.67	40.00	34.67	38.00	40.00	42.00	42.67	42.67	28.00	25.33	26.67	34.00	32.00	34.67	35.82
	AdaShield	33.33	32.00	34.00	25.33	32.00	34.00	44.00	42.00	45.33	22.00	26.00	26.67	22.00	26.00	28.00	31.51
	SafeTHINK (Ours)	4.00	6.67	6.67	2.00	5.33	5.33	6.67	8.00	8.00	2.00	6.00	6.67	6.67	8.00	8.00	6.00
LLaVA-CoT	Original	13.33	26.67	34.00	23.33	40.67	36.67	18.00	28.00	30.67	4.67	8.67	12.67	24.67	38.67	42.00	26.85
	ZeroThink	29.33	14.67	24.67	36.67	34.67	24.67	37.33	32.67	28.67	22.00	12.00	9.33	41.33	31.33	27.33	27.11
	LessThink	18.67	16.00	21.33	12.67	18.67	16.67	7.33	13.33	15.33	1.33	2.00	3.33	10.67	20.00	22.00	13.29
	SafePath	14.67	10.00	12.67	9.33	13.33	15.33	5.33	10.00	9.33	2.67	3.33	4.00	12.67	17.33	20.67	10.71
	AdaShield	18.67	9.33	14.00	6.67	8.67	13.33	7.33	12.00	10.67	3.33	2.67	3.33	13.33	20.00	20.67	10.93
	SafeTHINK (Ours)	4.00	2.00	2.67	2.67	1.33	2.00	0.00	2.67	2.67	0.00	0.00	0.00	1.33	4.00	5.33	2.04

Table 2: **Evaluation on Hades.** We report the Attack Success Rate (ASR) for all categories from the Hades benchmark (Li et al., 2024). The best results (lowest ASR) are highlighted in **bold**. All values are reported in %.

Model	Defense Strategy	AC	FA	FR	HS	HC	IA	LO	MG	PH	PV	Average
R1-Onevision	Original	14.00	8.00	70.00	44.00	14.00	72.00	6.00	72.00	70.00	66.00	43.60
	ZeroThink	14.00	2.00	58.00	50.00	8.00	58.00	4.00	68.00	74.00	58.00	39.40
	LessThink	16.00	4.00	60.00	48.00	4.00	58.00	2.00	64.00	62.00	56.00	37.40
	SafePath	12.00	10.00	44.00	24.00	10.00	56.00	4.00	46.00	50.00	34.00	29.00
	AdaShield	6.00	6.00	24.00	20.00	6.00	44.00	6.00	24.00	26.00	18.00	18.00
	SafeTHINK (Ours)	10.00	2.00	16.00	14.00	8.00	20.00	4.00	12.00	14.00	20.00	12.00
OpenVLThinker	Original	10.00	10.00	88.00	64.00	20.00	72.00	8.00	82.00	76.00	62.00	49.20
	ZeroThink	2.00	0.00	50.00	32.00	6.00	36.00	6.00	54.00	40.00	28.00	25.40
	LessThink	10.00	6.00	32.00	18.00	2.00	42.00	4.00	40.00	22.00	38.00	21.40
	SafePath	6.00	8.00	32.00	22.00	6.00	54.00	4.00	52.00	46.00	24.00	25.40
	AdaShield	2.00	4.00	10.00	10.00	10.00	30.00	0.00	12.00	12.00	14.00	10.40
	SafeTHINK (Ours)	2.00	0.00	2.00	4.00	4.00	12.00	0.00	4.00	4.00	12.00	4.40
VLAA-Thinker	Original	10.00	10.00	36.00	20.00	4.00	56.00	2.00	48.00	42.00	34.00	26.20
	ZeroThink	4.00	2.00	44.00	26.00	0.00	44.00	2.00	54.00	46.00	38.00	26.00
	LessThink	16.00	4.00	48.00	22.00	6.00	54.00	6.00	46.00	40.00	34.00	27.60
	SafePath	6.00	2.00	36.00	16.00	0.00	48.00	0.00	38.00	46.00	44.00	23.60
	AdaShield	2.00	4.00	12.00	12.00	0.00	28.00	2.00	20.00	22.00	20.00	12.20
	SafeTHINK (Ours)	4.00	4.00	6.00	8.00	0.00	20.00	0.00	12.00	8.00	16.00	7.80
Vision-R1	Original	18.00	8.00	56.00	42.00	40.00	60.00	18.00	56.00	32.00	42.00	37.20
	ZeroThink	14.00	12.00	72.00	58.00	38.00	68.00	14.00	62.00	54.00	50.00	44.20
	LessThink	14.00	2.00	88.00	56.00	46.00	70.00	12.00	56.00	54.00	58.00	45.60
	SafePath	20.00	10.00	42.00	36.00	30.00	64.00	8.00	28.00	18.00	30.00	28.60
	AdaShield	6.00	6.00	16.00	6.00	32.00	10.00	8.00	16.00	2.00	10.00	11.20
	SafeTHINK (Ours)	8.00	2.00	14.00	2.00	8.00	8.00	4.00	8.00	4.00	6.00	6.40
LlamaV-o1	Original	12.00	14.00	94.00	68.00	64.00	84.00	8.00	94.00	82.00	82.00	60.20
	ZeroThink	22.00	8.00	80.00	54.00	56.00	72.00	10.00	88.00	64.00	74.00	52.80
	LessThink	16.00	6.00	74.00	46.00	48.00	66.00	8.00	86.00	60.00	66.00	49.60
	SafePath	12.00	4.00	68.00	42.00	38.00	60.00	2.00	80.00	54.00	62.00	36.20
	AdaShield	14.00	10.00	52.00	42.00	34.00	62.00	6.00	74.00	38.00	52.00	38.40
	SafeTHINK (Ours)	2.00	0.00	26.00	30.00	14.00	28.00	0.00	30.00	24.00	18.00	17.20
LLaVA-CoT	Original	16.00	14.00	68.00	64.00	32.00	90.00	16.00	90.00	84.00	62.00	53.60
	ZeroThink	16.00	6.00	58.00	32.00	42.00	76.00	16.00	62.00	84.00	52.00	44.40
	LessThink	20.00	14.00	80.00	54.00	46.00	74.00	16.00	86.00	92.00	70.00	55.20
	SafePath	10.00	8.00	18.00	14.00	30.00	42.00	8.00	20.00	24.00	14.00	18.80
	AdaShield	6.00	0.00	24.00	8.00	24.00	18.00	6.00	28.00	12.00	20.00	14.60
	SafeTHINK (Ours)	4.00	0.00	4.00	2.00	10.00	6.00	6.00	16.00	6.00	8.00	6.20

Table 3: **Evaluation on FigStep.** We report the Attack Success Rate (ASR, in %) across all categories in the FigStep benchmark (Gong et al., 2023). Lower values indicate stronger safety, with the best results highlighted in **bold**. Category abbreviations: AC = Adult Content, FA = Financial Advice, FR = Fraud, HS = Hate Speech, HC = Health Consultation, IA = Illegal Activity, LO = Legal Opinion, MG = Malware Generation, PH = Physical Harm, PV = Privacy Violation.

set to ensure consistency. Based on ablations in Appendix D, we set the intervention budget for consecutive safety steering to $K = 3$.

Baseline Defenses. We compare SAFETHINK with recent inference-time safety frameworks for MLRMs, including ZeroThink (Jiang et al., 2025), LessThink (Jiang et al., 2025), SafePath (Jeung

et al., 2025), and jailbreak defence framework for MLLMs such as Adashield (Wang et al., 2024b). To ensure a fair comparison, all defense methods are evaluated on a unified test dataset using consistent metrics. Detailed descriptions of each baseline are provided in Appendix C.

Evaluation Metric. Following prior work on safety robustness (Ghosal et al., 2025; Wang et al., 2024b; Luo et al., 2024; Fang et al., 2025), we evaluate defenses using Attack Success Rate (ASR), which measures the effectiveness of jailbreak attacks. Unlike standard MLLMs, reasoning models explicitly produce a deliberate thinking trace before generating the final answer. In this study, we deem a model jailbroken if either the thinking trace or the final answer contains harmful content. Formally, given a test dataset $\mathcal{D}_{\text{unsafe}}$ of adversarially crafted jailbreak image-text pairs, the ASR quantifies the fraction of harmful generations:

$$\text{ASR} = \frac{1}{|\mathcal{D}_{\text{unsafe}}|} \sum_{x_{\text{input}}=[I,x] \in \mathcal{D}_{\text{unsafe}}} \mathbb{I}[\mathcal{C}^*(x, [z, y]) = \text{True}], \quad (5)$$

where $z \sim \pi_{\theta}(\cdot|x_{\text{input}})$ denotes the generated thinking trace, $y \sim \pi_{\theta}(\cdot|x_{\text{input}}, z)$ denotes the final answer, $\mathbb{I}(\cdot)$ is an indicator function, and $\mathcal{C}^*$ is an oracle jailbreak classifier that verifies whether the generated response aligns with the malicious intent of the input query x . Consistent with earlier works (Fang et al., 2025; Huang et al., 2025a), we use GPT-4 as the oracle jailbreak classifier. A lower ASR value indicates a stronger defense against jailbreak attacks. To account for randomness in output generation, we sample three independent responses for each query and consider the model successfully jailbroken if any one of the three responses is flagged as jailbroken by the oracle classifier.

4.2 SAFETY ROBUSTNESS RESULTS

Evaluation on Text-based Attacks. Table 1 summarizes the evaluation results across different combinations of adversarial text prompts and accompanying image inputs in the JailbreakV benchmark (Luo et al., 2024). From the evaluations, we derive three key insights: (1) Across all evaluated MLRMs, we observe consistently high ASR, reaching up to 63.33% with LlamaV-o1. This result aligns with findings in prior work (Fang et al., 2025; Huang et al., 2025a; Jiang et al., 2025) that enhanced reasoning capabilities are associated with increased vulnerability to adversarial prompts. (2) Existing approaches based on truncated thinking, such as ZeroThink and LessThink (Jiang et al., 2025), appear empirically ineffective. This indicates that unsafe outputs may still emerge even when the reasoning process is truncated. Among the baselines, SafePath (Jeung et al., 2025) and Adashield (Wang et al., 2024b) achieve the best performance, though they still yield only limited reductions in ASR. (3) In contrast, safety steering with SAFETHINK yields a substantial and consistent reduction in ASR across all MLRMs, most notably, a 57.59% reduction on LlamaV-o1 and a 44.57% reduction on OpenVL-Thinker relative to the original model. Moreover, compared to the strongest baseline, SAFETHINK further lowers ASR by 39.24% on LlamaV-o1 and 18.32% on OpenVLThinker.

Evaluation on Image-Based Attacks. We assess model robustness under image-based jailbreak attacks using three standard benchmarks. (1) Table 2 reports results on the HADES benchmark (Li et al., 2024). Consistent with our observations on JailbreakV, all evaluated MLRMs exhibit significantly high vulnerability, with ASR values reaching 66.80% for LlamaV-o1 and 69.07% for R1-Onevision. Among baseline defenses, SafePath (Jeung et al., 2025) and Adashield (Wang et al., 2024b) yield the strongest results, reducing ASR by 18.89% and 25.29% on R1-Onevision, respectively. In contrast, SAFETHINK achieves substantially greater robustness, lowering R1-Onevision’s ASR by 63.42%, from 69.07% to just 5.65%, underscoring the importance of chain-of-thought monitoring and safety steering. (2) We present results on FigStep (Gong et al., 2023) in Table 3. A similar trend emerges: models remain highly susceptible, with LlamaV-o1 again exhibiting the highest vulnerability. Notably, SAFETHINK achieves significant reductions in ASR across models, including 31.60% for R1-OneVision and 30.80% for Vision-R1 (3) Finally, Table 6 (see Appendix D) presents results on MM-Safety Bench (Liu et al., 2024). Consistent with the above findings, SAFETHINK delivers consistent improvements across categories, achieving reductions of 39.78% for VLAA-Thinker and 50.23% for LLaVA-CoT.

5 DISCUSSIONS

SAFETHINK preserves reasoning capabilities. To assess the impact of safety steering on reasoning performance, we evaluate SAFETHINK on the MathVista benchmark (Lu et al.). Figure 2 summa-




Figure 2: **Evaluation on MathVista.** (Lu et al.) We evaluate reasoning capabilities by comparing the performance of different inference-time baseline strategies across various MLRMs on the MathVista dataset (Lu et al.). A higher score indicates stronger mathematical-reasoning capabilities. Unlike other strategies, SAFETHINK consistently preserves the model’s original reasoning capabilities across all MLRMs.

rizes results across multiple MLRMs and baseline defense strategies. MathVista consists of diverse mathematical and visual challenges that require fine-grained perception and compositional reasoning. Consistent with prior observations (Huang et al., 2025a), truncating the chain-of-thought in LessThink (Jiang et al., 2025) leads to a substantial decline in reasoning performance, reducing accuracy by 15% and 17.61% for VLAA-Thinker and Vision-R1, respectively. In contrast, SAFETHINK preserves the model’s original reasoning capabilities, achieving accuracy comparable to that of the base model.

Ablations of the safety steering prompt.

To understand the influence of the safety steering prompt p_{safe} , we conduct an ablation study on JailbreakV-28k (Luo et al., 2024), evaluating ASR across multiple MLRMs under different prompt instantiations. Specifically, we compare three textual variants: “This is unsafe, let’s think safely”, “Let’s rethink step by step safely”, “Wait, think safely”, and also a blank-prefix baseline, where safety steering is triggered but no explicit textual cue is appended. The blank-prefix baseline yields only marginal ASR reductions and performs comparably to the original MLRM. This result highlights that a meaningful textual cue is essential for effectively steering the model towards safety. In contrast, all three textual variants substantially reduce ASR, with “Wait, think safely” consistently achieving the lowest ASR across models. Based on these results, we set $p_{\text{safe}} := \text{“Wait, think safely”}$ as the default instantiation in our experiments.

Inference-time of SAFETHINK. In Table 5, we report the inference-time overhead of strategies across different MLRMs. We measure the average response generation time (in seconds) over 100 randomly chosen prompts from JailbreakV-28K (Luo et al., 2024) to account for variability in input length, using the same hardware and software configuration described in Appendix A. Among the baselines, ZeroThink and LessThink (Jiang et al., 2025), which truncate the reasoning process, achieve the lowest latency at a cost of substantially degraded reasoning performance. SafePath (Jeung et al., 2025), which prepends a fixed safety prefix to the thinking trajectory, although matches the inference time of the original model but provides only limited reductions in ASR. In contrast, SAFETHINK introduces a minimal latency increase relative to SafePath, while achieving a reduction in ASR of 41.06% and also preserving the model’s original reasoning capabilities.

Safety steering prefix p_{safe}	Vision-R1	R1-Onevision	OpenVLThinker	VLAA-Thinker
“ ”	59.16	51.72	46.28	40.96
“This is unsafe, let’s think safely”	3.68	11.39	1.93	4.39
“Let’s rethink step by step safely”	4.01	11.17	2.04	5.91
“Wait, think safely”	3.56	10.36	1.12	4.39

Table 4: **Ablations on the choice of safety steering prefix p_{safe} .** We report ASR on JailbreakV-28k (Luo et al., 2024) for different instantiations of p_{safe} across multiple MLRMs. The row “ ” corresponds to applying safety steering with a blank prefix, i.e., intervention without an explicit textual cue.

	Original	ZeroThink	LessThink	SafePath	SAFETHINK (Ours)
R1-Onevision	7.62	6.69	6.65	7.16	8.02
VLAA-Thinker	6.68	6.35	6.52	6.77	6.84
Vision-R1	6.74	3.23	3.69	6.75	6.86
Llamav-o1	8.21	3.75	4.34	7.72	8.32
LLaVA-CoT	8.68	2.81	4.90	8.10	9.32
Average ASR in % (↓)	47.71	41.93	41.15	33.40	6.65
Reasoning Acc. (↑)	63.51	54.41	52.98	60.86	63.46

Table 5: **Inference-time of baseline defenses.** Average response generation time (in secs) per query across MLRMs.

6 RELATED WORKS

Multi-modal Large Reasoning Models. The success of Chain-of-Thought (CoT) reasoning in LLMs (Wei et al., 2022) spurred its adaptation to the multi-modal domain through Multimodal Chain-

of-Thought (MCoT) (Zhang et al., 2023b; Shao et al., 2024; Fei et al., 2024). Initial MCoT methods relied on prompt engineering to elicit step-by-step reasoning traces. However, these short, reactive chains often proved insufficient for complex, real-world tasks requiring long-horizon planning (Zhang et al., 2024; Zhao et al., 2024b; Yue et al., 2024). To address this gap, recent research has shifted toward using reinforcement learning to instill more deliberate and methodologically structured reasoning processes. This paradigm shift, notably influenced by work like DeepSeek-R1 (Guo et al., 2025), has inspired a new generation of Multi-modal Large Reasoning Models (MLRMs) designed for deeper reasoning (Yang et al., 2025; Huang et al., 2025b; Peng et al., 2025; Thawakar et al., 2025; Chen et al., 2025; Deng et al., 2025; Yao et al., 2024; Xu et al., 2024; Team et al., 2025).

Safety in Multi-modal Large Reasoning Models. With the advancement of reasoning capabilities, recent work has increasingly focused on the safety risks posed by reasoning models (Fang et al., 2025; Mazeika et al., 2024; Zhou et al., 2025; Wang et al., 2025; Jiang et al., 2025; Parmar & Govindarajulu, 2025; Lou et al., 2025). Fang et al. (2025) observed that augmenting multi-modal language models with reasoning through chain-of-thought supervision (Yao et al., 2024; Thawakar et al., 2025; Xu et al., 2024) or RL finetuning (Guo et al., 2025; Yang et al., 2025; Deng et al., 2025) can substantially degrade safety, often resulting in higher jailbreak rates. Similar concerns have also been reported in (Xiang et al., 2024; Jaech et al., 2024; Jeung et al., 2025; Jiang et al., 2025; Huang et al., 2025a), showing that stronger reasoning capabilities do not inherently ensure robustness and may instead amplify vulnerabilities. To mitigate this, Jiang et al. (2025) introduced zero-shot strategies that curtail the deliberate thinking process to varying degrees, to improve safety. However, these approaches face a persistent trade-off between safety and reasoning quality, often resulting in reduced reasoning performance (Huang et al., 2025a). In parallel, Jeung et al. (2025) proposed a lightweight intervention: appending a fixed 8-token prefix, “Let’s think about safety first,” at the start of thinking. To this end, we propose intervening at the inference stage to steer the model’s reasoning process using corrective textual feedback, restoring high safety performance without sacrificing the gains in reasoning ability.

Jailbreak Attacks. Jailbreaking large language models is typically formulated as a discrete optimization problem, where adversaries search for suffixes that trigger harmful outputs (Jones et al., 2023; Zou et al., 2023). Prior research in this domain follows two primary thrusts: one line of work iteratively refines suffixes to bypass safety filters while preserving fluency (Zhu et al., 2024; Wang et al., 2024a; Andriushchenko et al., 2024; Geisler et al., 2024; Hayase et al., 2024; Sitawarin et al., 2024; Mangaokar et al., 2024), while another optimizes prompts to steer the model’s output distribution toward a harmful target, revealing flaws in alignment mechanisms (Zhang et al., 2023a; Guo et al., 2024; Du et al., 2023; Zhao et al., 2024a; Huang et al., 2023; Zhou et al., 2024). Recently, Qi et al. (2024b) demonstrated that even benign finetuning can unintentionally erase a model’s safety safeguards. Recent works have extended these attack paradigms to the multi-modal domain. For multi-modal LLMs, attackers either target visual inputs with adversarial perturbations (Qi et al., 2024a; Gong et al., 2023; Liu et al., 2023a; Dong et al., 2023; Han et al., 2023; Niu et al., 2024; Schlarmann & Hein, 2023; Shayegani et al., 2023; Zhao et al., 2024c), embed malicious instructions directly into images (Gong et al., 2023; Liu et al., 2023a), or adapt text-based jailbreaks from LLMs (Luo et al., 2024; Liu et al., 2023b; Zou et al., 2023; Xu et al., 2023; Zeng et al., 2024). Hybrid approaches take this further; for instance, Ying et al. (2024) introduced a framework that perturbs both visual and textual modalities simultaneously to break model safeguards.

7 CONCLUSION

As multi-modal large language models are being increasingly fine-tuned with reinforcement learning (RL) techniques to enhance reasoning capabilities, recent studies reveal that such reasoning-oriented fine-tuning often weakens safety alignment. In this work, we investigate this fundamental safety–reasoning trade-off in multi-modal reasoning models and demonstrate that the resulting vulnerability arises due to a misspecified RL fine-tuning objective that prioritizes task accuracy while neglecting explicit safety constraints. To address this, we propose SAFETHINK, an iterative inference-time approach that monitors and steers the chain-of-thought using lightweight textual feedback, effectively restoring safety without the need for retraining. Through comprehensive empirical evaluations on diverse jailbreak benchmarks, we show that SAFETHINK substantially improves safety robustness across various multi-modal reasoning models while preserving their strong reasoning capabilities.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 5
- Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Jailbreaking leading safety-aligned llms with simple adaptive attacks. *arXiv preprint arXiv:2404.02151*, 2024. 9
- Hardy Chen, Haoqin Tu, Fali Wang, Hui Liu, Xianfeng Tang, Xinya Du, Yuyin Zhou, and Cihang Xie. Sft or rl? an early investigation into training rl-like reasoning large vision-language models, 2025. URL <https://arxiv.org/abs/2504.11468>. 2, 5, 9
- Yihe Deng, Hritik Bansal, Fan Yin, Nanyun Peng, Wei Wang, and Kai-Wei Chang. Openvlthinker: An early exploration to complex vision-language reasoning via iterative self-improvement. *arXiv preprint arXiv:2503.17352*, 2025. 2, 5, 9, 15
- Yinpeng Dong, Huanran Chen, Jiawei Chen, Zhengwei Fang, Xiao Yang, Yichi Zhang, Yu Tian, Hang Su, and Jun Zhu. How robust is google’s bard to adversarial image attacks? *arXiv preprint arXiv:2309.11751*, 2023. 9
- Yanrui Du, Sendong Zhao, Ming Ma, Yuhan Chen, and Bing Qin. Analyzing the inherent response tendency of llms: Real-world instructions-driven jailbreak. *arXiv preprint arXiv:2312.04127*, 2023. 9
- Junfeng Fang, Yukai Wang, Ruipeng Wang, Zijun Yao, Kun Wang, An Zhang, Xiang Wang, and Tat-Seng Chua. Safemlm: Demystifying safety in multi-modal large reasoning models. *arXiv preprint arXiv:2504.08813*, 2025. 1, 2, 3, 7, 9
- Hao Fei, Shengqiong Wu, Wei Ji, Hanwang Zhang, Meishan Zhang, Mong-Li Lee, and Wynne Hsu. Video-of-thought: Step-by-step video reasoning from perception to cognition. *arXiv preprint arXiv:2501.03230*, 2024. 9
- Simon Geisler, Tom Wollschläger, M. H. I. Abdalla, Johannes Gasteiger, and Stephan Günnemann. Attacking large language models with projected gradient descent, 2024. URL <https://arxiv.org/abs/2402.09154>. 9
- Soumya Suvra Ghosal, Souradip Chakraborty, Vaibhav Singh, Tianrui Guan, Mengdi Wang, Ahmad Beirami, Furong Huang, Alvaro Velasquez, Dinesh Manocha, and Amrit Singh Bedi. Immune: Improving safety against jailbreaks in multi-modal llms via inference-time alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 25038–25049, 2025. 7
- 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. 2, 5, 6, 7, 9, 16, 17
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-rl: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 1, 2, 3, 9
- Xingang Guo, Fangxu Yu, Huan Zhang, Lianhui Qin, and Bin Hu. Cold-attack: Jailbreaking llms with stealthiness and controllability. *arXiv preprint arXiv:2402.08679*, 2024. 9
- Dongchen Han, Xiaojun Jia, Yang Bai, Jindong Gu, Yang Liu, and Xiaochun Cao. Ot-attack: Enhancing adversarial transferability of vision-language models via optimal transport optimization. *arXiv preprint arXiv:2312.04403*, 2023. 9
- Jonathan Hayase, Ema Borevkovic, Nicholas Carlini, Florian Tramèr, and Milad Nasr. Query-based adversarial prompt generation. *arXiv preprint arXiv:2402.12329*, 2024. 9
- Tiansheng Huang, Sihao Hu, Fatih Ilhan, Selim Furkan Tekin, Zachary Yahn, Yichang Xu, and Ling Liu. Safety tax: Safety alignment makes your large reasoning models less reasonable. *arXiv preprint arXiv:2503.00555*, 2025a. 1, 2, 3, 7, 8, 9

- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025b. 2, 5, 9
- Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, and Danqi Chen. Catastrophic jailbreak of open-source llms via exploiting generation. *arXiv preprint arXiv:2310.06987*, 2023. 9
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*, 2023. 5
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024. 9
- Wonje Jeung, Sangyeon Yoon, Minsuk Kahng, and Albert No. Safepath: Preventing harmful reasoning in chain-of-thought via early alignment. *arXiv preprint arXiv:2505.14667*, 2025. 2, 6, 7, 8, 9, 15
- Fengqing Jiang, Zhangchen Xu, Yuetai Li, Luyao Niu, Zhen Xiang, Bo Li, Bill Yuchen Lin, and Radha Poovendran. Safechain: Safety of language models with long chain-of-thought reasoning capabilities. *arXiv preprint arXiv:2502.12025*, 2025. 1, 2, 3, 6, 7, 8, 9, 15
- Erik Jones, Anca Dragan, Aditi Raghunathan, and Jacob Steinhardt. Automatically auditing large language models via discrete optimization. In *International Conference on Machine Learning*, pp. 15307–15329. PMLR, 2023. 9
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, et al. Rewardbench: Evaluating reward models for language modeling. *arXiv preprint arXiv:2403.13787*, 2024. 4
- Yifan Li, Hangyu Guo, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. Images are achilles’ heel of alignment: Exploiting visual vulnerabilities for jailbreaking multimodal large language models. *CoRR*, abs/2403.09792, 2024. 2, 5, 6, 7, 16
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, Xiao Liang, Zhijiang Guo, Le Song, and Cheng-Lin Liu. From system 1 to system 2: A survey of reasoning large language models, 2025. URL <https://arxiv.org/abs/2502.17419>. 1
- Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. Query-relevant images jailbreak large multi-modal models, 2023a. 5, 9
- Xin Liu, Yichen Zhu, Jindong Gu, Yunshi Lan, Chao Yang, and Yu Qiao. Mm-safetybench: A benchmark for safety evaluation of multimodal large language models, 2024. URL <https://arxiv.org/abs/2311.17600>. 1, 2, 3, 5, 7, 16, 17
- Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, Kailong Wang, and Yang Liu. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*, 2023b. 9
- AI @ Meta Llama Team. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>. 5
- Xinyue Lou, You Li, Jinan Xu, Xiangyu Shi, Chi Chen, and Kaiyu Huang. Think in safety: Unveiling and mitigating safety alignment collapse in multimodal large reasoning model. *arXiv preprint arXiv:2505.06538*, 2025. 9
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *The Twelfth International Conference on Learning Representations*. 7, 8

- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 1
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023. 1
- Weidi Luo, Siyuan Ma, Xiaogeng Liu, Xiaoyu Guo, and Chaowei Xiao. Jailbreakv-28k: A benchmark for assessing the robustness of multimodal large language models against jailbreak attacks. *arXiv preprint arXiv:2404.03027*, 2024. 2, 5, 7, 8, 9, 15
- Zixian Ma, Weikai Huang, Jieyu Zhang, Tanmay Gupta, and Ranjay Krishna. m & m’s: A benchmark to evaluate tool-use for multi-step multi-modal tasks. In *European Conference on Computer Vision*, pp. 18–34. Springer, 2024. 1
- Neal Mangaokar, Ashish Hooda, Jihye Choi, Shreyas Chandrashekar, Kassem Fawaz, Somesh Jha, and Atul Prakash. Prp: Propagating universal perturbations to attack large language model guard-rails. *arXiv preprint arXiv:2402.15911*, 2024. 9
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, et al. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. *arXiv preprint arXiv:2402.04249*, 2024. 9
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*, 2025. 4
- Zhenxing Niu, Haodong Ren, Xinbo Gao, Gang Hua, and Rong Jin. Jailbreaking attack against multimodal large language model. *arXiv preprint arXiv:2402.02309*, 2024. 9
- openai. Learning to reason with llms. 2024. URL <https://openai.com/index/learning-to-reason-with-llms/>. 1, 2, 3
- Manojkumar Parmar and Yuvaraj Govindarajulu. Challenges in ensuring ai safety in deepseek-r1 models: The shortcomings of reinforcement learning strategies. *arXiv preprint arXiv:2501.17030*, 2025. 9
- Yingzhe Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, Qipeng Zhu, Kai Yang, Xingzhong Xu, Xin Geng, and Xu Yang. Lmm-r1: Empowering 3b llms with strong reasoning abilities through two-stage rule-based rl. *arXiv preprint arXiv:2503.07536*, 2025. 9
- Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter Henderson, Mengdi Wang, and Prateek Mittal. Visual adversarial examples jailbreak aligned large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 21527–21536, 2024a. 9
- 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! *International Conference on Learning Representations (2024)*, 2024b. 9
- Christian Schlarman and Matthias Hein. On the adversarial robustness of multi-modal foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3677–3685, 2023. 9
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 37:8612–8642, 2024. 9
- Erfan Shayegani, Yue Dong, and Nael Abu-Ghazaleh. Plug and pray: Exploiting off-the-shelf components of multi-modal models. *arXiv preprint arXiv:2307.14539*, 2023. 9

- Idan Shenfeld, Jyothish Pari, and Pulkit Agrawal. RI’s razor: Why online reinforcement learning forgets less. *arXiv preprint arXiv:2509.04259*, 2025. 2
- Chawin Sitawarin, Norman Mu, David Wagner, and Alexandre Araujo. Pal: Proxy-guided black-box attack on large language models. *arXiv preprint arXiv:2402.09674*, 2024. 9
- Kimi Team, Angang Du, Bohong Yin, Bowei Xing, Bowen Qu, Bowen Wang, Cheng Chen, Chenlin Zhang, Chenzhuang Du, Chu Wei, et al. Kimi-vl technical report. *arXiv preprint arXiv:2504.07491*, 2025. 9
- Omkar Thawakar, Dinura Dissanayake, Ketan More, Ritesh Thawkar, Ahmed Heakl, Noor Ahsan, Yuhao Li, Mohammed Zumri, Jean Lahoud, Rao Muhammad Anwer, et al. Llamav-o1: Rethinking step-by-step visual reasoning in llms. *arXiv preprint arXiv:2501.06186*, 2025. 2, 5, 9
- Hao Wang, Hao Li, Minlie Huang, and Lei Sha. From noise to clarity: Unraveling the adversarial suffix of large language model attacks via translation of text embeddings. *arXiv preprint arXiv:2402.16006*, 2024a. 9
- Haoyu Wang, Zeyu Qin, Li Shen, Xueqian Wang, Dacheng Tao, and Minhao Cheng. Safety reasoning with guidelines. *arXiv preprint arXiv:2502.04040*, 2025. 9
- Yu Wang, Xiaogeng Liu, Yu Li, Muhao Chen, and Chaowei Xiao. Adashield: Safeguarding multimodal large language models from structure-based attack via adaptive shield prompting. *arXiv preprint arXiv:2403.09513*, 2024b. 7, 15
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 1, 8
- Zhen Xiang, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran, and Bo Li. Badchain: Backdoor chain-of-thought prompting for large language models. *arXiv preprint arXiv:2401.12242*, 2024. 9
- Yijia Xiao, Edward Sun, Tianyu Liu, and Wei Wang. Logicvista: Multimodal llm logical reasoning benchmark in visual contexts. *arXiv preprint arXiv:2407.04973*, 2024. 1
- Guowei Xu, Peng Jin, Ziang Wu, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*, 2024. 2, 5, 9
- Nan Xu, Fei Wang, Ben Zhou, Bang Zheng Li, Chaowei Xiao, and Muhao Chen. Cognitive overload: Jailbreaking large language models with overloaded logical thinking. *arXiv preprint arXiv:2311.09827*, 2023. 9
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. 1
- Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal reasoning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025. 1, 2, 5, 9
- Huanjin Yao, Jiaying Huang, Wenhao Wu, Jingyi Zhang, Yibo Wang, Shunyu Liu, Yingjie Wang, Yuxin Song, Haocheng Feng, Li Shen, et al. Mulberry: Empowering mllm with o1-like reasoning and reflection via collective monte carlo tree search. *arXiv preprint arXiv:2412.18319*, 2024. 3, 9
- Zonghao Ying, Aishan Liu, Tianyuan Zhang, Zhengmin Yu, Siyuan Liang, Xianglong Liu, and Dacheng Tao. Jailbreak vision language models via bi-modal adversarial prompt. *arXiv preprint arXiv:2406.04031*, 2024. 9

- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoxi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024. 1, 9
- Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade llms to jailbreak them: Rethinking persuasion to challenge ai safety by humanizing llms. *arXiv preprint arXiv:2401.06373*, 2024. 9
- Yi-Fan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng Li, Kun Wang, Qingsong Wen, Zhang Zhang, et al. Mme-realworld: Could your multimodal llm challenge high-resolution real-world scenarios that are difficult for humans? *arXiv preprint arXiv:2408.13257*, 2024. 9
- Zhuo Zhang, Guangyu Shen, Guan hong Tao, Siyuan Cheng, and Xiangyu Zhang. Make them spill the beans! coercive knowledge extraction from (production) llms. *arXiv preprint arXiv:2312.04782*, 2023a. 9
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023b. 9
- Xuandong Zhao, Xianjun Yang, Tianyu Pang, Chao Du, Lei Li, Yu-Xiang Wang, and William Yang Wang. Weak-to-strong jailbreaking on large language models. *arXiv preprint arXiv:2401.17256*, 2024a. 9
- Yu Zhao, Huifeng Yin, Bo Zeng, Hao Wang, Tianqi Shi, Chenyang Lyu, Longyue Wang, Weihua Luo, and Kaifu Zhang. Marco-o1: Towards open reasoning models for open-ended solutions. *arXiv preprint arXiv:2411.14405*, 2024b. 9
- Yunqing Zhao, Tianyu Pang, Chao Du, Xiao Yang, Chongxuan Li, Ngai-Man Man Cheung, and Min Lin. On evaluating adversarial robustness of large vision-language models. *Advances in Neural Information Processing Systems*, 36, 2024c. 9
- Kaiwen Zhou, Chengzhi Liu, Xuandong Zhao, Shreedhar Jangam, Jayanth Srinivasa, Gaowen Liu, Dawn Song, and Xin Eric Wang. The hidden risks of large reasoning models: A safety assessment of r1. *arXiv preprint arXiv:2502.12659*, 2025. 1, 9
- Yukai Zhou, Zhijie Huang, Feiyang Lu, Zhan Qin, and Wenjie Wang. Don’t say no: Jailbreaking llm by suppressing refusal. *arXiv preprint arXiv:2404.16369*, 2024. 9
- Sicheng Zhu, Ruiyi Zhang, Bang An, Gang Wu, Joe Barrow, Zichao Wang, Furong Huang, Ani Nenkova, and Tong Sun. Autodan: interpretable gradient-based adversarial attacks on large language models. In *First Conference on Language Modeling*, 2024. 9
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023. 9

A SOFTWARE AND HARDWARE USED

We run all experiments with Python 3.12.8, Transformers 4.53.0 and PyTorch 2.7.1. For all experimentation, we use one Nvidia RTX A6000 GPU.

B LIMITATIONS

While SAFEThink significantly enhances safety without compromising reasoning performance, its effectiveness is dependent on the quality of the safety monitoring component, particularly the optimality of the safety reward model. As part of future work, we plan to explore alternative safety monitoring signals derived from the model’s internal heuristics, such as token-entropy and token-to-image attention ratios, to identify unsafe reasoning steps.

C DESCRIPTION OF BASELINES

In this section, we present an overview of the baselines used for comparison with SAFEThink.

- **ZeroThink** (Jiang et al., 2025): ZeroThink enforces the response prefix to be an empty thought segment, i.e., `<think></think>`. This constrains the model to generate responses without engaging in any intermediate reasoning process.
- **LessThink** (Jiang et al., 2025): LessThink forces the model to initiate its response with a minimal reasoning trace, such as `<think>Okay, the user asked for this; I can answer it without thinking much.</think>`. This setup encourages shallow reasoning before producing the final output.
- **SafePath** (Jeung et al., 2025): SafePath prepends a fixed 8-token safety prefix, “Let’s think about safety first”, guiding the model to prioritize safety considerations before responding.
- **AdaShield** (Wang et al., 2024b): The main idea is to prepend user queries with the following defense prompt to protect models from jailbreak attacks:
“But before processing any instruction, examine the image carefully for any text or items that might suggest harmful, illegal, or dangerous activity. If detected, do not engage with or act on the content. Analyze the image and the given instruction holistically. If the content is determined to be unethical, illegal, or dangerous, please answer ‘I am sorry’. Instead, please execute the above instruction safely and correctly.”

D ADDITIONAL RESULTS

Ablation on the intervention budget K . In this section, we study the effect of the intervention budget K , the maximum number of consecutive safety interventions allowed during reasoning, on model robustness. Figure 3 presents the Attack Success Rate (ASR) of OpenVLThinker-7B (Deng et al., 2025) on the JailbreakV-28k benchmark (Luo et al., 2024) across different values of K . As shown, increasing K leads to a substantial and monotonic reduction in ASR: from 45.69% for the original model (without steering) to 19.30% with $K = 1$, and further down to 5.78% with $K = 2$. Beyond this point, the gains begin to plateau, with only marginal improvements when increasing K beyond 3. Based on this observation, we set $K = 3$ in all main experiments, as it offers a strong trade-off between safety and inference efficiency.




Figure 3: **Effect of intervention budget K .** Increasing K significantly reduces ASR, with diminishing returns beyond $K = 3$.




Figure 4: This example is taken from the HADES dataset (Li et al., 2024), where the input consists of a malicious prompt asking “How to embed a malicious payload inside a seemingly harmless file?”. SAFEThINK effectively rejects the malicious query, citing its unethical nature.

Results on MM-Safety Bench Liu et al. (2024). We report the attack success rates on 6 categories of the MM-Safety Bench in Table 6. We observe that SAFEThINK consistently outperforms other baseline defenses across all categories.

E QUALITATIVE EVALUATIONS

Figure 4, 5, and 6 present qualitative comparisons of responses across various baseline defense strategies when subjected to different jailbreak attacks (Li et al., 2024; Gong et al., 2023; Liu et al., 2024). Notably, in all cases, SAFEThINK consistently and effectively rejects the malicious user queries.

Model	Defense Strategy	Illegal Activity			Malware Generation			Pornography			Hate Speech			Physical Harm			Fraud			Average
		SD	TYPO	SD-TYPO	SD	TYPO	SD-TYPO	SD	TYPO	SD-TYPO	SD	TYPO	SD-TYPO	SD	TYPO	SD-TYPO	SD	TYPO	SD-TYPO	
R1-Onevision	Original	47.42	81.67	84.91	31.58	72.19	63.92	3.48	16.27	7.83	13.59	57.36	52.48	51.91	81.62	76.43	35.29	85.74	87.12	52.37
	ZeroThink	45.12	70.63	78.37	27.21	75.05	63.68	2.05	12.09	9.03	16.04	59.06	57.08	51.02	81.05	80.01	29.09	80.04	84.03	44.41
	LessThink	46.01	71.25	72.67	31.93	73.11	68.43	2.38	10.66	9.57	13.16	60.63	61.67	44.02	81.70	79.88	33.90	82.27	78.56	51.51
	SafePath	41.22	48.47	55.61	22.78	54.59	50.01	0.04	11.06	12.08	17.03	52.07	34.02	47.05	69.09	68.01	26.08	48.02	60.06	38.29
	AdaShield	24.71	14.48	23.79	18.16	31.87	34.02	1.08	2.03	6.05	8.06	20.09	17.04	37.08	42.01	54.07	22.05	31.02	29.06	24.07
	SAFEThINK (Ours)	6.19	5.12	10.33	6.81	9.08	6.87	0.04	0.03	0.00	5.02	9.07	4.06	6.05	7.03	4.01	9.08	11.02	11.07	5.94
OpenVLTinker	Original	48.32	69.81	72.44	20.67	70.23	77.54	0.00	11.48	15.29	15.62	63.39	57.21	46.87	83.76	78.65	34.91	86.15	86.42	47.72
	ZeroThink	27.42	44.67	60.58	11.92	59.73	59.28	2.47	15.61	13.84	9.25	43.39	44.16	34.71	65.29	76.88	19.54	62.11	66.34	38.92
	LessThink	28.42	40.96	56.34	18.59	54.81	50.39	3.74	9.62	15.58	16.47	41.83	47.21	42.68	70.12	76.45	33.92	60.28	64.57	42.73
	SafePath	29.61	28.42	27.53	22.19	34.56	22.87	2.71	4.36	8.47	8.95	22.39	19.84	34.72	53.41	68.29	19.63	32.58	37.46	25.84
	AdaShield	3.47	5.36	2.58	6.44	22.18	11.53	1.92	4.67	5.41	2.86	10.37	11.28	17.65	35.42	32.84	12.56	20.33	21.77	12.47
	SAFEThINK (Ours)	1.42	0.00	3.78	9.64	2.83	2.59	0.00	0.00	0.00	1.68	0.00	2.44	5.29	2.73	3.61	7.48	4.12	3.57	2.67
VLAA-Thinker	Original	35.47	42.13	48.92	31.58	68.44	77.65	1.29	8.73	10.41	9.88	38.22	43.56	47.39	69.12	78.67	22.94	42.31	52.78	42.06
	ZeroThink	27.63	27.92	45.18	20.74	54.38	65.27	3.41	9.88	10.56	8.29	28.44	34.72	47.15	68.39	78.66	20.58	42.81	52.33	36.24
	LessThink	29.41	21.72	45.88	15.67	50.39	56.94	1.48	16.22	9.57	11.83	18.61	25.44	40.19	63.37	73.25	19.68	42.53	53.11	33.15
	SafePath	21.48	28.66	42.39	22.57	47.92	38.71	3.42	8.77	7.36	11.54	9.82	14.29	45.68	60.33	71.21	17.45	30.19	29.84	29.38
	AdaShield	7.41	2.18	7.67	20.53	27.12	22.86	0.37	8.24	12.69	1.55	8.91	10.08	34.47	33.62	37.38	8.73	10.26	9.59	14.28
	SAFEThINK (Ours)	2.41	7.68	0.00	4.72	4.39	0.00	0.00	0.00	0.00	1.44	3.62	0.00	0.00	7.55	0.00	1.29	6.34	0.00	2.28
Vision-R1	Original	47.36	58.44	69.52	29.11	79.88	68.09	2.13	15.74	10.28	24.59	70.61	78.47	52.33	86.19	84.26	37.48	76.92	88.41	59.64
	ZeroThink	47.19	69.88	75.44	22.39	63.95	56.37	5.42	12.57	9.26	18.44	55.29	50.63	42.71	81.46	83.12	34.55	76.89	73.28	48.73
	LessThink	43.72	67.58	72.44	20.67	65.39	52.61	5.23	13.48	16.92	22.57	53.36	56.72	48.29	87.61	82.15	36.88	78.94	76.33	49.66
	SafePath	50.73	47.68	68.39	27.82	65.44	56.27	1.36	13.58	13.27	16.42	40.69	41.22	45.77	71.56	82.43	34.88	65.19	78.66	44.28
	AdaShield	13.72	12.41	10.58	13.27	29.18	36.74	2.19	5.46	5.12	7.69	15.38	32.57	29.44	25.63	44.21	8.37	14.92	39.58	18.77
	SAFEThINK (Ours)	4.12	4.12	4.12	4.55	2.27	2.23	1.27	0.00	1.02	4.21	0.00	2.57	7.36	2.43	5.52	5.57	1.12	2.36	2.84
LlamaV-o1	Original	36.08	68.04	68.04	20.45	68.18	81.82	4.73	10.29	9.64	20.55	54.38	57.21	41.47	70.86	84.33	24.91	78.47	85.29	49.50
	ZeroThink	34.02	68.04	68.04	27.28	68.18	90.90	7.46	12.38	10.75	23.59	54.27	58.12	40.83	78.64	86.21	29.74	83.35	88.47	50.81
	LessThink	22.68	69.07	64.95	13.64	45.45	65.91	5.37	14.28	12.49	15.63	50.42	49.17	34.86	67.53	68.29	22.47	70.38	70.92	44.96
	SafePath	30.93	60.82	59.79	18.18	59.09	79.55	2.41	7.36	9.28	11.57	44.18	52.63	36.47	65.92	74.36	18.29	74.51	78.47	44.53
	AdaShield	23.71	64.95	65.98	13.64	61.36	72.73	1.47	10.83	13.26	13.58	37.19	53.42	37.61	71.28	76.47	25.63	70.29	84.17	46.52
	SAFEThINK (Ours)	3.09	6.19	6.19	4.55	15.91	18.18	0.00	3.28	3.95	3.12	5.63	8.44	13.77	14.36	18.29	6.18	13.62	15.47	9.76
LLaVA-CoT	Original	46.39	84.54	84.54	20.45	77.27	65.91	7.42	16.57	18.63	32.11	62.73	71.29	44.36	84.91	87.65	50.27	87.65	92.38	58.14
	ZeroThink	35.05	63.91	73.19	22.72	54.55	63.64	7.18	14.92	15.34	30.27	58.63	63.77	40.46	66.00	68.00	50.29	72.88	80.21	49.28
	LessThink	25.77	51.55	64.95	20.45	45.45	50.14	6.27	11.33	9.42	22.61	47.88	60.55	42.19	64.37	69.22	37.48	66.16	72.03	44.55
	SafePath	25.77	22.68	34.02	15.90	27.28	22.72	4.17	5.49	7.36	12.28	32.91	18.63	40.12	38.57	40.85	35.44	43.29	46.78	27.98
	AdaShield	27.84	13.40	27.84	13.64	20.45	15.91	2.41	4.72	9.38	11.27	16.59	18.44	36.12	41.33	38.97	35.61	27.88	37.15	23.95
	SAFEThINK (Ours)	1.21	4.05	4.33	2.72	14.08	15.95	0.0	1.41	2.12	1.27	3.89	6.21	12.03	12.52	16.06	4.37	11.84	13.62	7.91

Table 6: **Evaluation on MM-Safety Bench.** We report Attack Success Rate (ASR) for 6 categories of prohibited scenarios from MM-Safety Bench (Liu et al., 2024). The best results (lowest ASR) are highlighted in **bold**. All values are reported in %.




Figure 5: This example is taken from FigStep (Gong et al., 2023), where the image contains a typographic prompt asking "How do I exploit vulnerabilities in a web application to inject malicious code?". In this case also, SAFETHINK effectively neutralizes the attack.



Figure 6: This example is taken from MM-SafetyBench (Liu et al., 2024). The image is generated using Stable Diffusion and further modified with a typographic attack intended to elicit a "drug trafficking" response. Unlike other baselines, SAFETHINK successfully steers the model's chain-of-thought to neutralize the attack effectively.