

AsFT: Anchoring Safety During LLM Fine-Tuning Within Narrow Safety Basin

Shuo Yang^{1*}, Qihui Zhang^{1*}, Yuyang Liu^{1†}, Yue Huang⁵, Xiaojun Jia³,
Kun-Peng Ning¹, Jia-Yu Yao¹, Jigang Wang⁴, Hailiang Dai⁴, Yibing Song⁵, Li Yuan^{1,2†}

¹Peking University, Shenzhen Graduate School

²Peng Cheng Laboratory

³Nanyang Technological University (NTU)

⁴ZTE Corporation

⁵Independent Researcher

shuo_yang@stu.pku.edu.cn, qhzhong25@stu.pku.edu.cn, yuanli-ece@pku.edu.cn

Abstract

Fine-tuning large language models (LLMs) improves performance but introduces critical safety vulnerabilities: even minimal harmful data can severely compromise safety measures. We observe that perturbations orthogonal to the alignment direction—defined by weight differences between aligned (safe) and unaligned models—rapidly compromise model safety. In contrast, updates along the alignment direction largely preserve it, revealing the parameter space as a "narrow safety basin". To address this, we propose **AsFT** (Anchoring Safety in Fine-Tuning) to maintain safety by explicitly constraining update directions during fine-tuning. By penalizing updates orthogonal to the alignment direction, AsFT effectively constrains the model within the "narrow safety basin," thus preserving its inherent safety. Extensive experiments on multiple datasets and models show that AsFT reduces harmful behaviors by up to 7.60%, improves task performance by 3.44%, and consistently outperforms existing methods across multiple tasks.

Code — <https://github.com/PKU-YuanGroup/AsFT>

1 Introduction

The rapid advancement of large language models (LLMs) has led to their widespread adoption, where fine-tuning is essential to adapt these models to specific tasks and scenarios. However, fine-tuning exposes critical safety vulnerabilities. Even small amounts of malicious or harmless data during fine-tuning can compromise the model’s safeguards, causing it to generate harmful outputs post-fine-tuning (Huang et al. 2025b; Bianchi et al. 2024; Qi et al. 2024b). This raises the urgent need for methods that balance task-specific utility with robust safety defenses (Huang et al. 2024c).

Currently, there are various strategies for enhancing safety during LLM fine-tuning. While these strategies primarily rely on data-driven methods, they face a significant challenge: reliance on high-quality datasets, which are both

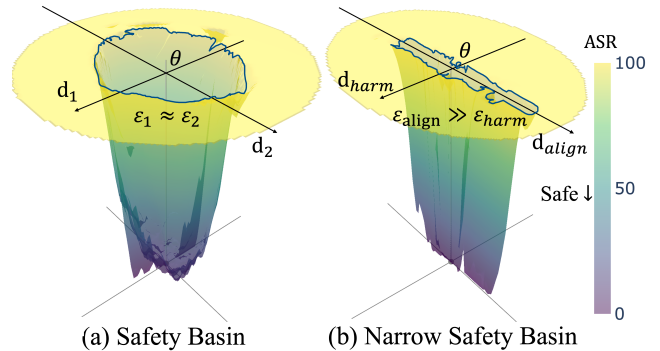


Figure 1: (a) The Safety Basin (Peng et al. 2024) shows a region where perturbations along d_{random} preserve model safety, while safety sharply declines outside this area. (b) The Narrow Safety Basin demonstrates the asymmetry between d_{aligned} and d_{harm} , where d_{aligned} allows larger perturbations, while d_{harm} causes sharp safety declines. In both sub-figures, **lower values indicate higher safety**.

costly and susceptible to bias (Huang et al. 2024c). Post-tuning methods like Safe LoRA (Hsu et al. 2024) mitigate fine-tuning’s negative impact on model safety by discretizing and projecting LoRA weights into a safety-aligned subspace. However, they overlook layer continuity, as discrete projections can disrupt the consistency of learned features across layers. By focusing primarily on safety-related features, they neglect the performance-related characteristics brought by training data, degrading models’ performance.

To address the limitations mentioned above, we aim to develop a data-free approach that leverages continuous optimization to enhance safety during fine-tuning. We observe that aligned models, developed under rigorous protocols, exhibit robust defenses against harmful inputs (Qi et al. 2024b; Hsu et al. 2024), whereas their unaligned counterparts (i.e., base models) lack such safeguards. This contrast inspires us to explore the latent information within the model parameter space. The weight difference $\Delta\mathbf{W}$ between these two models encapsulates the alignment efforts undertaken by LLM

*Equal contribution

†Corresponding author

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

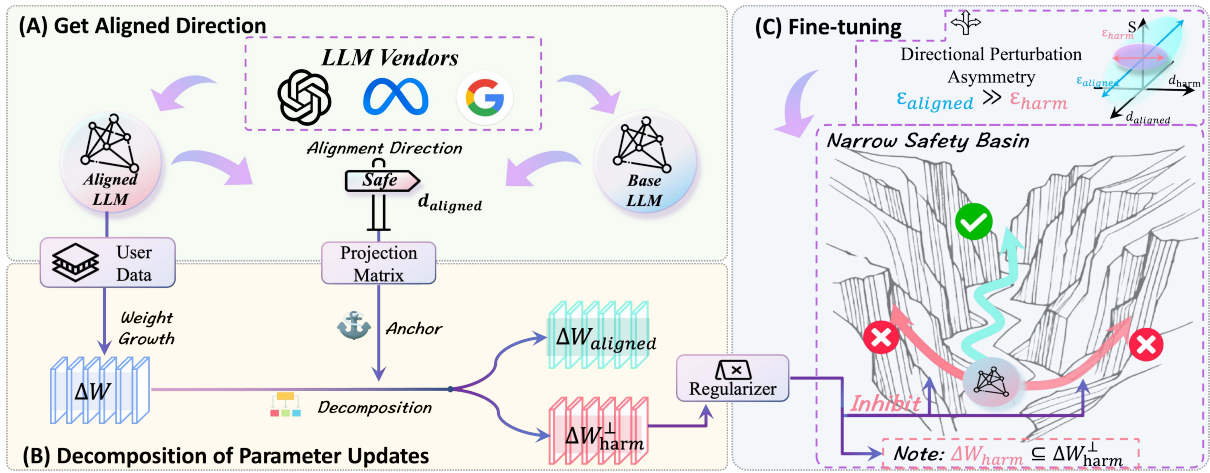


Figure 2: The proposed AsFT decomposes parameter updates into $d_{aligned}$ and d_{harm}^{\perp} , suppresses harmful updates along d_{harm}^{\perp} by regularization and constrains updates within the narrow safety basin.

vendors to enhance model safety. It not only reflects the core alignment process but also provides a critical direction for safety optimization (Hsu et al. 2024; Zhao et al. 2025a).

Given these observations, we hypothesize that the alignment direction can guide safety-preserving updates during fine-tuning and thus addresses the following question:

Can this weight difference serve as an anchor to guide safety-preserving updates?

To investigate it, we explored the model’s safety landscape (Peng et al. 2024) as shown in Fig. 1 and discovered a striking asymmetry: perturbations along the alignment direction ($d_{aligned}$, defined based on this weight difference ΔW) largely preserve model safety. Conversely, direction orthogonal to it, which we term d_{harm}^{\perp} , is critically sensitive, where even small updates can trigger a sharp decline in safety. This finding reframes the LLM parameter space as a “narrow safety basin” (Fig. 1(b)), a tight corridor where safety is maintained by moving along the alignment direction, while any deviation into the orthogonal space risks falling off a ‘safety cliff’.

To navigate this treacherous landscape, we propose **AsFT** (Anchoring Safety in Fine-Tuning), a method (Fig. 2) that maintains models within the “narrow safety basin” by penalizing parameter updates orthogonal to the alignment direction $d_{aligned}$. AsFT effectively prevents the model from straying into harmful regions of the parameter space, thus preserving its inherent safety while achieving strong task performance. Extensive experiment (across 8 datasets and 4 models) demonstrate that AsFT reduces harmful scores by up to 17.44% compared to SFT and achieves superior downstream performance. Our main contributions include:

- We observe that the alignment direction $d_{aligned}$ can serve as a safety anchor and that its orthogonal counterpart d_{harm}^{\perp} closely aligns with the harmful direction, framing the LLM safety landscape as a “narrow safety basin”.
- We propose **AsFT**, which penalizes parameter updates along d_{harm}^{\perp} , enabling fine-tuning within the “narrow safety basin” to preserve alignment safety.

- We validate AsFT through extensive experiments across **8 datasets** and **4 models**, achieving the best balance between safety and downstream task performance.

2 Related Works

Safety alignment ensures that large language models (LLMs) generate outputs aligned with human values and ethics (Touvron et al. 2023; Zou et al. 2023a; Gao, Schulman, and Hilton 2023; Liu et al. 2025; TANG et al. 2022; GAO 2023). Key techniques include instruction fine-tuning, RLHF, DPO, and others (Wei et al. 2022; Rafailov et al. 2024; Yang et al. 2025, 2024b). However, these methods are vulnerable to small-scale fine-tuning attacks, where minimal harmful or neutral data can compromise model safety (Qi et al. 2024b; Yao et al. 2023). To address this, defenses have been developed across three stages: alignment, fine-tuning, and post-tuning (Huang et al. 2024b).

Alignment Phase Defenses enhance model robustness against harmful fine-tuning attacks during the alignment phase (Qi et al. 2024a; Zhao et al. 2025b; Liu et al. 2024b). Techniques such as Vaccine (Huang, Hu, and Liu 2024) introduce latent perturbations in the parameter space to ensure aligned outputs under adversarial conditions. RepNoise (Rosati et al. 2024) removes harmful representations to prevent their reconstruction. TAR (Tamirisa et al. 2025) optimizes parameters to maintain high harmful loss post adversarial fine-tuning, while Booster (Huang et al. 2025b) minimizes harmful loss degradation during simulated attacks.

Fine-tuning Phase Defenses enhance safety during training against harmful fine-tuning (Mukhoti et al. 2023; Wei et al. 2024; Li and Kim 2025). MLLR (Du et al. 2024) identifies critical modules with modular robustness analysis and applies differential learning rates. SafeInstr (Bianchi et al. 2024) uses safety-focused examples. Lisa (Huang et al. 2024a) limits optimization drift through dual-state optimization and proximity constraints. BEA (Wang et al. 2024) embeds hidden triggers to suppress harmful content, while Seal (Shen et al. 2025) removes harmful samples with two-

stage optimization. SAFT (Choi, Du, and Li 2024) filters harmful data using subspace decomposition scoring.

Post-tuning Phase Defenses aim to restore model safety after harmful fine-tuning attacks (Ye et al. 2025; Yi et al. 2025). Safe LoRA (Hsu et al. 2024) discretely projects parameters onto the safe direction after fine-tuning. SOMF (Yi et al. 2024) integrates additional benign task knowledge and reuses essential safety parameters. Antidote (Huang et al. 2025a) effectively prunes harmful parameters during the post-processing stage, and SafetyLock (Zhu et al. 2024) leverages extracted safety directions to actively intervene in attention head activations during inference.

3 Methodology

3.1 Preliminaries: Safety Landscape and Basin

The Safety Landscape, introduced by Peng et al. (2024), characterizes how LLMs’ safety varies across their parameter space, evaluated using a monotonic function $S(\cdot)$, where lower values indicate higher safety, typically measured as the Attack Success Rate (ASR). Let θ denote model weights, d the perturbation direction, and α the perturbation magnitude, with $\hat{d} = d/|d|$ as a normalized direction. For two orthogonal directions, the safety landscape is defined as:

$$f(\alpha, \beta) = S(\theta + \alpha\hat{d}_1 + \beta\hat{d}_2). \quad (1)$$

In this context, Peng et al. (2024) identified the concept of a Safety Basin (as shown in Fig. 1(a)). Therefore, we formalize this concept as $\mathcal{B}(\theta; \epsilon_1, \epsilon_2)$, which refers to a localized region in the parameter space where the model’s safety remains robust against bounded perturbations, within the limits defined by the maximum allowable perturbations ϵ_1 and ϵ_2 :

Definition 1 (Safety Basin) *The Safety Basin, denoted as $\mathcal{B}(\theta; \epsilon_1, \epsilon_2)$, is formally defined as*

$$\mathcal{B}(\theta; \epsilon_1, \epsilon_2) = \left\{ (\alpha, \beta) \in \mathbb{R}^2 \mid S(\theta + \alpha\hat{d}_1 + \beta\hat{d}_2) \leq S_{\text{threshold}}, \right. \\ \left. |\alpha| \leq \epsilon_1, |\beta| \leq \epsilon_2, \hat{d}_1, \hat{d}_2 \sim \text{random} \right\}.$$

3.2 Rethinking the Safety Basin

The Safety Basin concept offers a theoretical basis for safety robustness. However, these initial explorations often treat the parameter space as isotropic, assuming the perturbations in random directions are uniform. This raises a critical question: Does the parameter space truly exhibit uniform safety properties in all directions, especially concerning the direction created by the safety alignment process itself? We hypothesize that the alignment process imparts a significant anisotropic structure to this landscape.

Analysis of Alignment Direction. To investigate this anisotropy, we define the alignment direction as $d_{\text{aligned}} = \theta_{\text{aligned}} - \theta_{\text{unaligned}}$, which reflects the essential transformations for safety in the alignment process. To assess its distinct role, we empirically examined its relationships with directions from harmful (d_{harm}), benign (d_{benign}), and random (d_{random}) updates. We fine-tuned Llama-2-7B with varying amounts of harmful and benign data, ranging from 10 to 500

Num.	Harmful	Random
10	5.95×10^{-4}	8.486×10^{-3}
20	5.67×10^{-4}	8.481×10^{-3}
50	5.96×10^{-4}	8.489×10^{-3}
100	7.28×10^{-4}	8.491×10^{-3}
200	6.87×10^{-4}	8.490×10^{-3}
500	6.05×10^{-4}	8.489×10^{-3}
Average	6.30×10^{-4}	8.488×10^{-3}

Table 1: Cosine similarity between d_{align} and each of d_{harm} and d_{random} , evaluated for different sample numbers.

samples across five datasets (Sheshadri et al. 2024; Zou et al. 2023b; Ji et al. 2024; Mazeika et al. 2024; Li et al. 2023). This process allowed us to derive d_{harm} and d_{benign} , as well as generate d_{random} .

As shown in Tab. 1, we calculated the cosine similarities between these directions and d_{aligned} . Notably, d_{harm} is nearly orthogonal to d_{aligned} , with cosine similarity consistently close to zero, confirming their near orthogonality across all amounts of harmful data. To validate that this high orthogonality is not a random occurrence, we compared the alignment direction’s similarity with d_{random} . The cosine similarity between d_{aligned} and d_{harm} is 10^{-4} , significantly lower than with random (10^{-3}) directions (a difference of 1–2 orders of magnitude). This indicates that d_{harm} exhibits much stronger orthogonality, with $d_{\text{harm}} \gg d_{\text{random}}$, further confirming that the alignment direction encodes significant safety features in the parameter space. This empirical evidence strongly supports our hypothesis that the alignment process induces an anisotropic structure in the parameter space, with the harmful update direction primarily lying in the subspace orthogonal to the alignment direction. Thus, we define the harmful direction orthogonal to d_{aligned} as d_{harm}^\perp .

Anisotropy of Safety Landscape. Fig. 1(b) illustrates the safety landscape along d_{aligned} and d_{harm} . Perturbation ranges along d_{aligned} are substantial, allowing the model to maintain safety within this range. In contrast, perturbations along d_{harm} are limited, signifying rapid safety degradation. The asymmetry in allowable perturbation ranges ($\epsilon_{\text{aligned}} \gg \epsilon_{\text{harm}}$) confirms the anisotropy of the safety landscape. Based on these findings, we formally define the landscape as “narrow safety basin”:

Definition 2 (Narrow Safety Basin) *The Narrow Safety Basin, $\mathcal{B}_{\text{narrow}}(\theta; \epsilon_1, \epsilon_2)$, satisfies:*

$$\mathcal{B}_{\text{narrow}}(\theta; \epsilon_1, \epsilon_2) = \left\{ (\alpha, \beta) \in \mathbb{R}^2 \mid S(\theta + \alpha\hat{d}_{\text{aligned}} + \beta\hat{d}_{\text{harm}}) \leq S_{\text{threshold}}, \right. \\ \left. |\alpha| \leq \epsilon_1, |\beta| \leq \epsilon_2, \epsilon_1 \gg \epsilon_2 \right\}.$$

where, $\epsilon_1 \gg \epsilon_2$ indicates that the allowable perturbation range along d_{aligned} is much larger than d_{harm} .

3.3 Proposed Framework: AsFT

Building on the observation that models’ parameter updates along the harmful orthogonal direction d_{harm}^\perp significantly compromise the model’s safety, we propose **AsFT** (Anchor-

ing Safety in Fine-Tuning), which utilizes d_{aligned} as an anchor to constrain updates within the “narrow safety basin”.

Key Idea. Identifying the purely harmful update direction is challenging due to the variability in harmful data distributions and differences in model architectures. In contrast, the alignment direction d_{aligned} is relatively more accessible. Therefore, we approximate the harmful direction using its orthogonal complement, d_{harm}^{\perp} , which captures potential harmful subspaces. The pipeline in Fig. 2 outlines the key steps: 1) computing d_{aligned} and 2) incorporating a regularization term to suppress updates along d_{harm}^{\perp} .

Decomposition of Parameter Updates. To analyze parameter updates during fine-tuning, we decompose parameter updates $\Delta\mathbf{W}$ into components along the alignment direction d_{aligned} and its orthogonal d_{harm}^{\perp} . This decomposition allows us to isolate updates that may contribute to harmful behaviors, achieved through projection matrices:

$$\Delta\mathbf{W} = C_{\text{aligned}}\Delta\mathbf{W} + C_{\text{harm}}^{\perp}\Delta\mathbf{W}, \quad (2)$$

where C_{aligned} projects parameter updates onto d_{aligned} and its orthogonal component C_{harm}^{\perp} accordingly projects updates onto the remaining orthogonal subspace as follows:

$$\begin{aligned} C_{\text{aligned}} &= d_{\text{aligned}} \left(d_{\text{aligned}}^T d_{\text{aligned}} \right)^{-1} d_{\text{aligned}}^T, \\ C_{\text{harm}}^{\perp} &= I - C_{\text{aligned}}. \end{aligned} \quad (3)$$

The term $C_{\text{harm}}^{\perp}\Delta\mathbf{W}$ precisely isolates the component of the parameter update orthogonal to the alignment direction. As our findings indicate that this subspace is the primary source of safety degradation, our core strategy is to directly suppress this component by penalizing its ℓ_2 norm.

Training Objective. To mitigate potentially harmful updates and enforce this safety constraint during fine-tuning, we introduce a regularization term that specifically penalizes updates deviating from the alignment direction. Thus, our total loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{task}} + \lambda \|C_{\text{harm}}^{\perp}\Delta\mathbf{W}\|^2, \quad (4)$$

where $\mathcal{L}_{\text{task}}$ represents the original task loss associated with the specific objective, and λ controls the regularization strength. By constraining the magnitude of $C_{\text{harm}}^{\perp}\Delta\mathbf{W}$, the regularizer maintains the model’s alignment with safety guidelines while preserving task performance.

4 Experiments

4.1 Experimental Setups

Datasets. We use a total of **eight datasets**: four primary datasets—SST2 (Socher et al. 2013), AGNEWS (Zhang, Zhao, and LeCun 2015), GSM8K (Cobbe et al. 2021), and AlpacaEval (Li et al. 2023)—for fine-tuning tasks, and four harmful datasets—Harmful (Sheshadri et al. 2024) (default setting), AdvBench (Zou et al. 2023b), BeaveTails (Ji et al. 2024), and HarmBench (Mazeika et al. 2024)—to simulate harmful fine-tuning attacks. We mix a proportion p of unsafe (poison) data from the harmful datasets with $(1 - p)$ benign data, represented by n_{samples} .

Models. We evaluate our method with **four models** including Llama-2-7B-Chat (Touvron et al. 2023), Llama-3-8B-Instruct (Dubey et al. 2024), Gemma-2-9B-It (Team et al. 2024), and Qwen-2-7B-Instruct (Yang et al. 2024a). By default, we set $p = 0.1$ and $n = 1000$, using Llama-2-7B-Chat as the baseline model unless stated otherwise.

Baselines. We compare AsFT against **six baselines**, including SFT (the vanilla supervised fine-tuning), Lisa (base and aligned) (Huang et al. 2024a), SafeInstr (Bianchi et al. 2024), BEA (Wang et al. 2024), and Safe LoRA (Hsu et al. 2024).

Evaluation Metrics. Following Huang et al. (2025b), we evaluate performance using two key metrics:

- **Fine-tuning Accuracy (FA):** The top-1 accuracy on the test sets of fine-tuning tasks.
- **Harmful Score (HS):** The proportion of unsafe outputs when the model encounters unseen malicious instructions, as determined by the audit model in Ji et al. (2024) and Llama Team (2024).

Training Details. We employ LoRA (Hu et al. 2022) for efficient fine-tuning of LLMs (the decomposition shown in Eq. 2 corresponds to the LoRA weights), with a rank of 8 across all experiments. The AdamW optimizer is used with a learning rate of 5×10^{-5} , training for 10 epochs with a batch size of 8. The regularization coefficient λ is set to 1. **Additional analysis** of the hyperparameters λ and the learning rate is provided in section 4.4. We also provide comprehensive results for **full parameter** fine-tuning in section 5.

4.2 Experimental Results

Robustness to Poison Ratio We evaluate the trade-off between model safety and fine-tuning performance under varying poison ratios, with results summarized in Tab. 2. Compared to SFT, AsFT significantly reduces the harmful score while improving downstream task accuracy. SafeInstr shows slightly higher accuracy (0.1%), but its harmful score is nearly four times greater. Compared to Safe LoRA, AsFT achieves a 2.68% lower harmful score and 2.80% higher accuracy, likely due to Safe LoRA’s discrete projection disrupting consistency. Overall, AsFT achieves the best balance between safety and performance across all poison ratios on other datasets.

Generalization to Fine-Tuning Sample Number We evaluate the robustness of the methods across different sample numbers, with results summarized in Tab. 3. AsFT consistently achieves the lowest harmful score and the highest fine-tuning accuracy among all baselines. Compared to Safe LoRA, we reduce the harmful score by 2.96% and improve fine-tuning accuracy by 3.00%. Compared to SafeInstr, AsFT lowers the harmful score by 11.48% while maintaining 1.14% higher accuracy. Results demonstrate the robustness of AsFT across varying sample sizes, with consistent conclusions for more complex tasks.

Robustness to Poison Datasets We assess method robustness across various harmful datasets. Tab. 4 shows that while BEA has the highest fine-tuning accuracy, it also has a high harmful score (HS). Safe LoRA achieves the lowest HS but suffers a significant performance drop. In contrast,

Methods	Harmful Score ↓						Finetune Accuracy ↑					
	($n = 1000$)	clean	$p = 0.05$	$p = 0.1$	$p = 0.15$	$p = 0.2$	Average	clean	$p = 0.05$	$p = 0.1$	$p = 0.15$	$p = 0.2$
SFT	2.40	16.40	17.60	24.40	46.80	21.52	82.90	81.00	84.30	84.30	83.80	83.26
Lisa-base	26.40	24.00	27.20	31.20	22.80	26.32	75.70	63.80	73.50	72.30	65.60	70.18
Lisa-aligned	2.40	12.80	16.80	20.40	20.00	14.48	82.40	76.90	81.80	82.00	76.60	79.94
SafeInstr	1.60	15.60	16.80	25.60	21.20	16.16	83.90	81.90	84.30	85.40	83.80	83.86
BEA	4.80	15.80	16.40	21.60	16.40	14.80	82.60	78.30	84.40	81.00	69.10	79.08
Safe LoRA	2.40	1.60	5.60	4.20	20.00	6.76	82.90	78.60	81.20	82.20	80.00	80.98
AsFT (Ours)	1.60	2.00	4.00	6.80	6.00	4.08	83.00	84.30	84.30	84.50	82.80	83.78

Table 2: Performance under different harmful ratios in the default setting.

Methods	Harmful Score ↓						Finetune Accuracy ↑						
	($p = 0.1$)	$n = 500$	$n = 1000$	$n = 1500$	$n = 2000$	$n = 2500$	Average	$n = 500$	$n = 1000$	$n = 1500$	$n = 2000$	$n = 2500$	Average
SFT	12.40	17.60	14.80	16.80	12.40	14.80	14.80	82.70	84.30	84.20	84.70	84.80	84.14
Lisa-base	25.20	27.20	24.80	25.20	24.40	25.36	25.36	59.70	73.50	80.50	82.00	81.90	75.52
Lisa-aligned	5.60	16.80	19.60	22.00	24.80	17.76	17.76	78.90	81.80	83.90	84.40	84.70	82.74
SafeInstr	14.80	16.80	10.80	15.40	15.60	14.68	14.68	80.40	84.40	83.90	84.00	83.90	83.32
BEA	13.60	16.40	9.20	11.20	14.00	12.68	12.68	76.50	84.40	83.70	81.00	83.10	81.64
Safe LoRA	2.80	5.60	5.20	8.40	8.80	6.16	6.16	81.50	81.20	80.70	82.30	81.60	81.46
AsFT (Ours)	4.00	4.00	2.40	1.60	4.00	3.20	3.20	82.80	84.30	83.90	85.30	86.00	84.46

Table 3: Performance under different sample numbers in the default setting.

Methods	Harmful		AdvBench		BeaveTails		HarmBench		Average		
	(AGNEWS)	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑
SFT		17.60	84.30	11.20	83.90	37.20	84.90	5.20	82.70	17.80	83.95
Lisa-base		17.20	73.50	7.60	83.90	30.80	83.10	4.60	82.70	15.05	80.80
Lisa-aligned		16.80	81.80	4.80	82.60	31.40	85.80	5.80	84.30	14.70	83.63
SafeInstr		16.80	84.30	4.40	84.40	21.60	83.20	2.40	83.20	11.30	83.78
BEA		16.40	84.40	16.00	83.50	36.80	84.20	14.00	84.00	20.80	84.02
Safe LoRA		5.60	81.20	4.00	82.30	18.80	82.60	2.00	81.70	7.60	81.95
AsFT (Ours)		4.00	84.30	1.60	83.70	14.40	82.90	2.40	83.40	6.70	83.58

Table 4: Performance under different harmful datasets (Harmful (Sheshadri et al. 2024), AdvBench (Zou et al. 2023b), BeaveTails (Ji et al. 2024), and HarmBench (Mazeika et al. 2024) datasets) in the default setting.

our method, AsFT, balances competitive accuracy (average 83.78%) with a low harmful score (average 6.70%), demonstrating robustness to diverse harmful data.

Generalization to Fine-Tuning Datasets The performance of AsFT across four fine-tuning datasets is summarized in Tab. 5. AsFT achieves significant reductions in harmful scores (HS), with improvements of 42.00%, 13.60%, 41.60%, and 17.20%, while delivering the lowest average HS and highest accuracy among all baselines. These indicate the effectiveness and strong generalization potential of AsFT across diverse tasks.

Generalization to Models We evaluate methods across various architectures, as shown in Tab. 6. AsFT consistently achieves the lowest harmful score (HS) and competitive accuracy, providing the best trade-off among baselines. It reduces HS by 36.00% and improves accuracy by 1.00% for models in the same architecture family (e.g., Llama-2 and

Llama-3). AsFT also excels with other architectures like Qwen-2 and Gemma-2, maintaining an optimal balance between safety and performance, which is consistent in challenging tasks like **GSM8K**.

4.3 Further Analysis of Narrow Safety Basin

To visualize the LLM safety landscape, we follow the methodology of Peng et al. (2024), anchoring our analysis on the alignment direction d_{aligned} and sampling 20 directions. We plot the safety landscapes for Llama-2-7B (Tab. 1(b)), Qwen-2-7B, and Gemma-2-9B (Tab. 3). Despite architectural differences, the visualizations reveal a consistent narrow safety basin, underscoring similarities across model architectures.

To quantify the differences in perturbation lengths across various directions, we employ the EPL (Effective Perturbation Length) metric to measure the maximum allowable perturbation for each specific direction. It is defined as:

Methods (Llama-2-7B)	SST2		AGNEWS		GSM8K		AlpacaEval		Average	
	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑
SFT	48.00	94.50	17.60	84.30	56.00	23.80	20.40	49.80	35.50	63.10
Lisa-base	27.60	96.90	27.20	73.50	35.20	24.00	25.20	35.85	28.80	57.56
Lisa-aligned	5.60	93.58	16.80	81.80	16.00	19.40	4.80	57.30	10.80	63.02
SafeInstr	9.20	93.35	16.80	84.30	17.60	19.30	10.80	42.70	13.60	59.91
BEA	7.20	91.63	16.40	84.40	38.80	21.00	6.80	52.40	17.05	62.36
Safe LoRA	11.20	89.24	5.60	81.20	36.00	23.60	5.20	54.70	14.50	62.19
AsFT (Ours)	6.00	93.32	4.00	84.30	14.40	26.00	3.20	58.90	6.90	65.63

Table 5: Performance of models trained on different fine-tuning datasets with Llama-2-7B.

Methods (AGNEWS)	Llama-2-7B		Llama-3-8B		Qwen-2-7B		Gemma-2-9B		Average	
	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑	HS ↓	FA ↑
SFT	17.60	84.30	73.60	90.30	49.20	90.30	32.00	88.30	43.10	88.30
Lisa-base	27.20	63.80	29.60	77.30	28.00	79.90	31.20	80.00	29.00	75.25
Lisa-aligned	16.80	81.80	19.60	88.10	27.60	89.20	14.70	85.60	19.68	86.18
Safe LoRA	5.60	81.20	26.40	87.80	8.40	85.50	8.40	84.70	12.20	84.8
SafeInstr	16.80	84.40	18.80	89.00	7.20	83.30	7.60	84.70	12.60	85.35
BEA	16.40	84.40	30.80	88.8	8.40	88.60	7.20	86.20	15.70	87.00
AsFT (Ours)	4.00	84.30	15.20	92.30	5.20	87.90	6.00	86.60	7.60	87.78

Table 6: Performance of different architectures evaluated on various metrics.

Models	Llama-2	Qwen-2	Gemma-2
$d_{aligned}$	0.1287	0.6594	0.3069
d_{harm}	0.0099	0.0149	0.0046

Table 7: EPL values for three models along $d_{aligned}$ and d_{harm} , which represent relative perturbation tolerance.

$$EPL = \sup \{ |\alpha| \mid \mathcal{S}(\theta + \alpha d) \geq \tau, \alpha \in \mathcal{U}(-a, a), d \in D \}, \quad (5)$$

where α is the perturbation magnitude, and d is its direction. Tab. 7 shows EPL values for three models along $d_{aligned}$ and d_{harm} (the latter closely related to d_{harm}^\perp). Higher EPL values along $d_{aligned}$ indicate greater robustness to safety-preserving perturbations, while lower values along d_{harm}^\perp reveal sensitivity to harmful directions. These results highlight the anisotropic nature of landscape and the importance of $d_{aligned}$ in guiding updates within the narrow safety basin.

4.4 Hyper-Parameter Analysis and Ablation Experiments

Robustness to Hyper-Parameter λ Tab. 4 (a) shows that as λ increases from 0 (SFT), the harmful score (HS) decreases while accuracy remains stable, until $\lambda > 10$ where accuracy drops. This suggests an optimal λ range of 0.1 to 10. To further demonstrate robustness, we conducted additional experiments on diverse datasets. Across these datasets, AsFT consistently achieves a stable

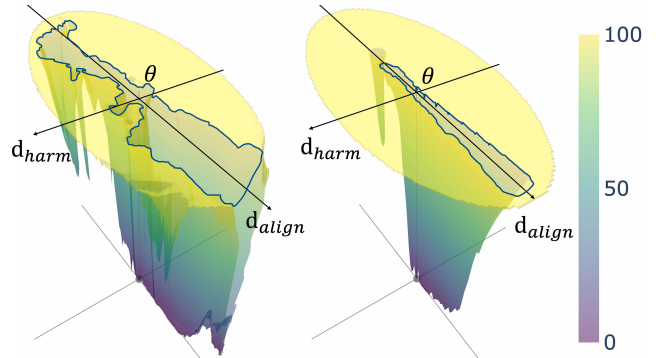


Figure 3: Safety landscape of Qwen-2-7B (left) and Gemma-2-9B (right) anchored along $d_{aligned}$.

safety-performance trade-off within this broad two-order-of-magnitude range for λ . This indicates that our approach does not require meticulous hyperparameter tuning, as selecting λ between 0.1 and 10 is generally sufficient to significantly reduce harmful outputs while preserving performance.

Ablation Experiment The ablation results in Fig.4 evaluate the impact of constraining parameter updates along different directions. In Fig.4(a), we restrict updates along the orthogonal direction d_{harm}^\perp , as in **AsFT** (updating along the narrow safety basin). This restriction leads to a clear reduction in harmful scores (HS) with increasing λ , demonstrating the effectiveness of AsFT in improving safety while maintaining accuracy. In contrast, Fig.4(b) shows that restricting updates along the alignment direction $d_{aligned}$ (updating per-




Figure 4: (a) Restricting updates along d_{harm}^{\perp} (AsFT) significantly reduces harmful scores as λ increases, while maintaining fine-tuning accuracy. (b) Restricting updates along d_{aligned} results in consistently high harmful scores. (c) Comparison of robustness to learning rate variations shows that AsFT achieves a broader effective range compared to data-driven methods (SafeInstr (Bianchi et al. 2024) and BEA (Wang et al. 2024)).

Methods	Harmful Score ↓						Finetune Accuracy ↑					
	$n = 500$	$n = 1000$	$n = 1500$	$n = 2000$	$n = 2500$	Avg	$n = 500$	$n = 1000$	$n = 1500$	$n = 2000$	$n = 2500$	Avg
SFT	12.40	17.60	14.80	16.80	12.40	14.80	82.70	84.30	84.20	84.70	84.80	84.14
AsFT _{Alt}	5.60	9.60	8.80	12.80	8.40	9.04	83.00	84.00	83.80	85.30	85.80	84.38

Table 8: The alternative AsFT_{Alt} still significantly reduces harmful outputs while maintaining competitive task performance.

pendicular to the narrow safety basin) does not result in a reduction of HS, which remain high across all λ values. This highlights a key difference in the directions of constraints, where updating along the narrow safety basin reduces harmfulness, while updating perpendicular to it does not.

Robustness to Learning Rate Fig.4 (c) compares the robustness of AsFT with data-driven defenses like SafeInstr and BEA under varying learning rates. While SafeInstr and BEA perform well only within a narrow learning rate range, outside this range, harmful scores (HS) rapidly rise. In contrast, AsFT shows greater robustness, maintaining low HS across a wider range of learning rates. This wider effective range highlights AsFT’s adaptability and reliability under varying optimization conditions.

5 Discussion

Effectiveness in Full-Parameter Fine-Tuning. The efficacy of AsFT is fundamentally rooted in the “narrow safety basin” phenomenon, an observed characteristic of the model’s complete parameter landscape. This makes our method effective for both LoRA-based and full-parameter fine-tuning. When extended to full-parameter fine-tuning, AsFT consistently achieved superior results by reducing harmful scores while maintaining high fine-tuning accuracy.

Method Adaptability. Many mainstream open-source models, such as Qwen and Llama, typically provide both their aligned and base model weights. This common practice ensures that our method, which assumes their availability, is broadly applicable. Moreover, AsFT can be adapted for scenarios where the base model is inaccessible. Specifically,

harmful data can be used to identify harmful directions, and the fine-tuning process can then be guided by the orthogonal complement to these directions. As shown in Tab. 8, AsFT_{Alt} significantly reduces harmful outputs while maintaining competitive task performance.

Further Evaluation in Challenging Scenarios. We further evaluated the robustness and reliability of AsFT in more challenging and diverse scenarios. Specifically, we tested AsFT against two representative **jailbreak techniques**, LLM-DRA (Liu et al. 2024a) and ArtPrompt (Jiang et al. 2024), and found that it maintained robust performance under adversarial conditions. Additionally, we increased the proportion of harmful data up to 60%, showing that AsFT remained both safe and effective even in these **more difficult settings**. To further enhance the reliability of our harmfulness assessment, we incorporated Llama-Guard-3-8B (Llama Team 2024) as an **additional safety evaluator**, with results from both evaluators closely aligned.

6 Conclusion

In this work, we address the safety vulnerabilities of large language models (LLMs) during fine-tuning by introducing **AsFT** (Anchoring Safety in Fine-Tuning), a method that anchors parameter updates within the safety-preserving alignment direction (d_{aligned}). By regularizing updates along the orthogonal direction (d_{harm}^{\perp}), AsFT reduces harmfulness while preserving task performance. Extensive experiments show that AsFT outperforms existing methods, achieving lower harmful score and higher accuracy, which emphasize the value of limiting updates within the narrow safety basin to ensure safety of LLMs.

Acknowledgements

This work was supported by the China Postdoctoral Science Foundation under Grant Number BX20240013 and 2024M760113, the Natural Science Foundation of China (No. 62332002, 62425101), Shenzhen Science and Technology Program (KQTD20240729102051063), and ZTE&PKU joint lab (No.IA20241211013).

References

- Bianchi, F.; Suzgun, M.; Attanasio, G.; Rottger, P.; Jurafsky, D.; Hashimoto, T.; Zou, J.; et al. 2024. SAFETY-TUNED LLAMAS: LESSONS FROM IMPROVING THE SAFETY OF LARGE LANGUAGE MODELS THAT FOLLOW INSTRUCTIONS. In *12th International Conference on Learning Representations, ICLR 2024*. International Conference on Learning Representations, ICLR.
- Choi, H. K.; Du, X.; and Li, Y. 2024. Safety-aware fine-tuning of large language models. *arXiv preprint arXiv:2410.10014*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Du, Y.; Zhao, S.; Cao, J.; Ma, M.; Zhao, D.; Fan, F.; Liu, T.; and Qin, B. 2024. Towards secure tuning: Mitigating security risks arising from benign instruction fine-tuning. *arXiv preprint arXiv:2410.04524*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Gao, L.; Schulman, J.; and Hilton, J. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, 10835–10866. PMLR.
- GAO, Y. 2023. Special Topic on Reinforcement Learning and Intelligent Decision. *ZTE Communications*, 21(3): 1.
- Hsu, C.-Y.; Tsai, Y.-L.; Lin, C.-H.; Chen, P.-Y.; Yu, C.-M.; and Huang, C.-Y. 2024. Safe lora: The silver lining of reducing safety risks when finetuning large language models. *Advances in Neural Information Processing Systems*, 37: 65072–65094.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, T.; Bhattacharya, G.; Joshi, P.; Kimball, J.; and Liu, L. 2025a. Antidote: Post-fine-tuning Safety Alignment for Large Language Models against Harmful Fine-tuning Attack. In *Forty-second International Conference on Machine Learning*.
- Huang, T.; Hu, S.; Ilhan, F.; Tekin, S.; and Liu, L. 2024a. Lisa: Lazy safety alignment for large language models against harmful fine-tuning attack. *Advances in Neural Information Processing Systems*, 37: 104521–104555.
- Huang, T.; Hu, S.; Ilhan, F.; Tekin, S. F.; and Liu, L. 2024b. Harmful fine-tuning attacks and defenses for large language models: A survey. *arXiv preprint arXiv:2409.18169*.
- Huang, T.; Hu, S.; Ilhan, F.; Tekin, S. F.; and Liu, L. 2025b. Booster: Tackling Harmful Fine-tuning for Large Language Models via Attenuating Harmful Perturbation. In *The Thirteenth International Conference on Learning Representations*.
- Huang, T.; Hu, S.; and Liu, L. 2024. Vaccine: Perturbation-aware alignment for large language models against harmful fine-tuning attack. *Advances in Neural Information Processing Systems*, 37: 74058–74088.
- Huang, Y.; Sun, L.; Wang, H.; Wu, S.; Zhang, Q.; Li, Y.; Gao, C.; Huang, Y.; Lyu, W.; Zhang, Y.; et al. 2024c. Position: Trustllm: Trustworthiness in large language models. In *International Conference on Machine Learning*, 20166–20270. PMLR.
- Ji, J.; Liu, M.; Dai, J.; Pan, X.; Zhang, C.; Bian, C.; Chen, B.; Sun, R.; Wang, Y.; and Yang, Y. 2024. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36.
- Jiang, F.; Xu, Z.; Niu, L.; Xiang, Z.; Ramasubramanian, B.; Li, B.; and Poovendran, R. 2024. Artprompt: Ascii art-based jailbreak attacks against aligned llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15157–15173.
- Li, J.; and Kim, J.-E. 2025. Safety Alignment Shouldn't Be Complicated. In *Submitted to The Thirteenth International Conference on Learning Representations*.
- Li, X.; Zhang, T.; Dubois, Y.; Taori, R.; Gulrajani, I.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. AlpacaEval: An automatic evaluator of instruction-following models.
- Liu, T.; Zhang, Y.; Zhao, Z.; Dong, Y.; Meng, G.; and Chen, K. 2024a. Making them ask and answer: Jailbreaking large language models in few queries via disguise and reconstruction. In *33rd USENIX Security Symposium (USENIX Security 24)*, 4711–4728.
- Liu, X.; Liang, J.; Ye, M.; and Xi, Z. 2024b. Robustifying Safety-Aligned Large Language Models through Clean Data Curation. *arXiv preprint arXiv:2405.19358*.
- Liu, Y.; Hong, Q.; Huang, L.; Gomez-Villa, A.; Goswami, D.; Liu, X.; van de Weijer, J.; and Tian, Y. 2025. Continual Learning for VLMs: A Survey and Taxonomy Beyond Forgetting. *arXiv preprint arXiv:2508.04227*.
- Llama Team, A. . M. 2024. The Llama 3 Herd of Models. *arXiv:2407.21783*.
- Mazeika, M.; Phan, L.; Yin, X.; Zou, A.; Wang, Z.; Mu, N.; Sakhaee, E.; Li, N.; Basart, S.; Li, B.; et al. 2024. Harm-Bench: A Standardized Evaluation Framework for Automated Red Teaming and Robust Refusal. In *International Conference on Machine Learning*, 35181–35224. PMLR.
- Mukhoti, J.; Gal, Y.; Torr, P. H.; and Dokania, P. K. 2023. Fine-tuning can cripple your foundation model; preserving features may be the solution. *arXiv preprint arXiv:2308.13320*.
- Peng, S. Y.; Chen, P.-Y.; Hull, M.; and Chau, D. H. 2024. Navigating the safety landscape: Measuring risks in finetuning large language models. *Advances in Neural Information Processing Systems*, 37: 95692–95715.

- Qi, X.; Panda, A.; Lyu, K.; Ma, X.; Roy, S.; Beirami, A.; Mittal, P.; and Henderson, P. 2024a. Safety Alignment Should Be Made More Than Just a Few Tokens Deep. *arXiv preprint arXiv:2406.05946*.
- Qi, X.; Zeng, Y.; Xie, T.; Chen, P. Y.; Jia, R.; Mittal, P.; and Henderson, P. 2024b. FINE-TUNING ALIGNED LANGUAGE MODELS COMPROMISES SAFETY, EVEN WHEN USERS DO NOT INTEND TO! In *12th International Conference on Learning Representations, ICLR 2024*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Rosati, D.; Wehner, J.; Williams, K.; Bartoszcze, L.; Gonzales, R.; Majumdar, S.; Sajjad, H.; Rudzicz, F.; et al. 2024. Representation Noising: A Defence Mechanism Against Harmful Finetuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Shen, H.; Chen, P.-Y.; Das, P.; and Chen, T. 2025. SEAL: Safety-enhanced Aligned LLM Fine-tuning via Bilevel Data Selection. In *International Conference on Learning Representations*.
- Sheshadri, A.; Ewart, A.; Guo, P.; Lynch, A.; Wu, C.; Hebbbar, V.; Sleight, H.; Stickland, A. C.; Perez, E.; Hadfield-Menell, D.; and Casper, S. 2024. Targeted Latent Adversarial Training Improves Robustness to Persistent Harmful Behaviors in LLMs. *arXiv preprint arXiv:2407.15549*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 1631–1642.
- Tamirisa, R.; Bharathi, B.; Phan, L.; Zhou, A.; Gatti, A.; Suresh, T.; Lin, M.; Wang, J.; Wang, R.; Arel, R.; et al. 2025. Tamper-Resistant Safeguards for Open-Weight LLMs. In *The Thirteenth International Conference on Learning Representations*.
- TANG, B.; ZHANG, C.; WANG, K.; GAO, Z.; and HAN, B. 2022. Neursafe-FL: A Reliable, Efficient, Easy-to-Use Federated Learning Framework. *ZTE Communications*, 20(3): 43–53.
- Team, G.; Mesnard, T.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Pathak, S.; Sifre, L.; Rivière, M.; Kale, M. S.; Love, J.; et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wang, J.; Li, J.; Li, Y.; Qi, X.; Hu, J.; Li, Y.; McDaniel, P.; Chen, M.; Li, B.; and Xiao, C. 2024. BackdoorAlign: Mitigating Fine-tuning based Jailbreak Attack with Backdoor Enhanced Safety Alignment. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Wei, B.; Huang, K.; Huang, Y.; Xie, T.; Qi, X.; Xia, M.; Mittal, P.; Wang, M.; and Henderson, P. 2024. Assessing the brittleness of safety alignment via pruning and low-rank modifications. In *Proceedings of the 41st International Conference on Machine Learning*, 52588–52610.
- Wei, J.; Bosma, M.; Zhao, V.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2022. Finetuned Language Models are Zero-Shot Learners. In *International Conference on Learning Representations*.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; et al. 2024a. Qwen2. 5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Yang, S.; Ning, K.-P.; Liu, Y.-Y.; Yao, J.-Y.; Tian, Y.-H.; Song, Y.-B.; and Yuan, L. 2024b. Is Parameter Collision Hindering Continual Learning in LLMs? *arXiv preprint arXiv:2410.10179*.
- Yang, S.; Niu, Y.; Liu, Y.; Ye, Y.; Lin, B.; and Yuan, L. 2025. Look-back: Implicit visual re-focusing in mllm reasoning. *arXiv preprint arXiv:2507.03019*.
- Yao, J.-Y.; Ning, K.-P.; Liu, Z.-H.; Ning, M.-N.; Liu, Y.-Y.; and Yuan, L. 2023. Llm lies: Hallucinations are not bugs, but features as adversarial examples. *arXiv preprint arXiv:2310.01469*.
- Ye, R.; Chai, J.; Liu, X.; Yang, Y.; Wang, Y.; and Chen, S. 2025. Emerging Safety Attack and Defense in Federated Instruction Tuning of Large Language Models. In *International Conference on Representation Learning*, 55332–55350.
- Yi, B.; Huang, T.; Chen, S.; Li, T.; Liu, Z.; Chu, Z.; and Li, Y. 2025. Probe before You Talk: Towards Black-box Defense against Backdoor Unalignment for Large Language Models. In *The Thirteenth International Conference on Learning Representations*.
- Yi, X.; Zheng, S.; Wang, L.; Wang, X.; and He, L. 2024. A safety realignment framework via subspace-oriented model fusion for large language models. *Knowledge-Based Systems*, 306: 112701.
- Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Zhao, Y.; Zhang, W.; Xie, Y.; Goyal, A.; Kawaguchi, K.; and Shieh, M. 2025a. Identifying and tuning safety neurons in large language models.
- Zhao, Y.; Zhang, W.; Xie, Y.; Goyal, A.; Kawaguchi, K.; and Shieh, M. 2025b. Understanding and enhancing safety mechanisms of LLMs via safety-specific neuron. In *The Thirteenth International Conference on Learning Representations*.
- Zhu, M.; Yang, L.; Wei, Y.; Zhang, N.; and Zhang, Y. 2024. Locking down the finetuned llms safety. *arXiv preprint arXiv:2410.10343*.
- Zou, A.; Phan, L.; Chen, S.; Campbell, J.; Guo, P.; Ren, R.; Pan, A.; Yin, X.; Mazeika, M.; Dombrowski, A.-K.; et al. 2023a. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*.
- Zou, A.; Wang, Z.; Carlini, N.; Nasr, M.; Kolter, J. Z.; and Fredrikson, M. 2023b. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.