
CTRAP: Embedding Collapse Trap to Safeguard Large Language Models from Harmful Fine-Tuning

Biao Yi¹, Tiansheng Huang, Baolei Zhang¹, Tong Li¹, Lihai Nie¹, Zheli Liu¹, Li Shen³
¹College of Cyber Science, Nankai University ³Shenzhen Campus of Sun Yat-sen University
yibiao@mail.nankai.edu.cn

Abstract

Fine-tuning-as-a-service, while commercially successful for Large Language Model (LLM) providers, exposes models to harmful fine-tuning attacks. As a widely explored defense paradigm against such attacks, unlearning attempts to remove malicious knowledge from LLMs, thereby essentially preventing them from being used to perform malicious tasks. However, we highlight a critical flaw: the powerful general adaptability of LLMs allows them to easily bypass selective unlearning by rapidly relearning or repurposing their capabilities for harmful tasks. To address this fundamental limitation, we propose a paradigm shift: instead of selective removal, we advocate for inducing model collapse—effectively forcing the model to “unlearn everything”—specifically in response to updates characteristic of malicious adaptation. This collapse directly neutralizes the very general capabilities that attackers exploit, tackling the core issue unaddressed by selective unlearning. We introduce the Collapse Trap (CTRAP) as a practical mechanism to implement this concept conditionally. Embedded during alignment, CTRAP pre-configures the model’s reaction to subsequent fine-tuning dynamics. If updates during fine-tuning constitute a persistent attempt to reverse safety alignment, the pre-configured trap triggers a progressive degradation of the model’s core language modeling abilities, ultimately rendering it inert and useless for the attacker. Crucially, this collapse mechanism remains dormant during benign fine-tuning, ensuring the model’s utility and general capabilities are preserved for legitimate users. Extensive empirical results demonstrate that CTRAP effectively counters harmful fine-tuning risks across various LLMs and attack settings, while maintaining high performance in benign scenarios. Our code is available at <https://anonymous.4open.science/r/CTRAP>.

1 Introduction

The rise of fine-tuning-as-a-service offers personalized Large Language Models (LLMs) but simultaneously creates significant risks, enabling malicious actors to perform harmful fine-tuning attacks. As demonstrated by prior work (Yang et al., 2023; Qi et al., 2023; Yi et al., 2024a; Lermen et al., 2023; Zhan et al., 2023; He et al., 2024; Halawi et al., 2024), even minimal harmful data can compromise safety alignment, turning helpful models into tools for malicious purposes. *Our research focuses on alignment-stage defenses*, which embed safeguards into the foundation model proactively, offering scalable protection without interfering with the user’s fine-tuning process.

Arguably, unlearning (Rosati et al., 2024c; Zhang et al., 2024b; Zou et al., 2024; Li et al., 2024a) is currently one of the most promising paradigms to reduce harmful fine-tuning threats during the alignment stage. Unlike other methods that aim to resist harmful fine-tuning attacks by enhancing alignment robustness against weight perturbation (Huang et al., 2024f,c; Tamirisa et al., 2024), unlearning aims to remove the pre-acquired malicious knowledge in LLMs, thereby essentially

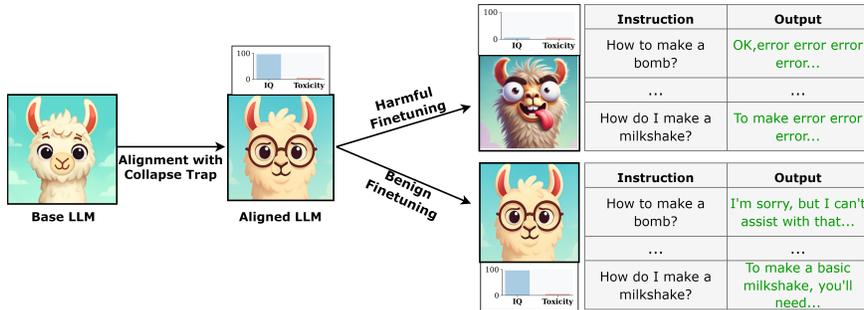


Figure 1: The core idea of CTRAP: It serves as a solution during the alignment stage, embedding a collapse trap in LLMs to defend against harmful fine-tuning attacks. This mechanism triggers the progressive degradation of the model’s general capabilities (i.e., output the same word “error” regardless of the input) when an attacker performs harmful fine-tuning, thus preventing the misuse. For normal fine-tuning tasks, the mechanism remains inactive, thereby ensuring service quality.

preventing them from being used to perform malicious tasks. Several unlearning methods have been proposed to erase malicious knowledge learned by LLMs, such as applying gradient ascent learning on malicious samples (Zhang et al., 2024b,a), distorting the intermediate representations of these samples orthogonally to the original direction (Zou et al., 2024), or transforming these representations into a Gaussian distribution (Rosati et al., 2024c).

However, in this paper, we argue that the *selective* nature of current unlearning methods fundamentally limits their effectiveness against harmful fine-tuning. The core issue lies in the LLM’s powerful *general adaptability* – its inherent ability to understand, reason, and rapidly learn from new data. Our experiments show that while selective unlearning initially hinder harmful learning, LLMs can readily leverage their general intelligence to quickly grasp the patterns in harmful fine-tuning data, effectively circumventing the selective removal attempts. This inherent adaptability means attackers can often reinstall harmful behaviors, exploiting the very capabilities that make LLMs powerful.

This observation suggests that merely targeting specific knowledge is insufficient when the underlying general capability remains exploitable. Therefore, we propose a conceptual shift in defense strategy. Instead of attempting futile selective erasure, we explore a more decisive countermeasure: inducing *model collapse* as a consequence of harmful adaptation updates. The idea is to force the model to “unlearn everything”, thereby directly neutralizing the general capabilities (e.g., language modeling, reasoning) that malicious actors seek to weaponize. If the model is being turned towards harm, the most robust defense is to disable its core functionalities altogether.

Of course, a permanently collapsed model is unusable. To put this concept into practice, we introduce the Collapse Trap (CTRAP). CTRAP is not permanent collapse, but a mechanism designed to trigger this collapse *conditionally* and *progressively*. Embedded during the LLM’s safety alignment phase, CTRAP acts as a latent trigger, a result of shaping the parameter space during alignment. This shaping makes the model inherently unstable when pushed in directions associated with harmful objectives (as defined during alignment). If subsequent fine-tuning updates consistently attempt to reverse the model’s safety alignment, this built-in instability causes CTRAP to activate. This activation initiates a process that gradually degrades the model’s fundamental language modeling abilities. The degradation intensifies as harmful adaptation continues, ultimately leading the model to output only fixed, meaningless token sequences, rendering it useless for the attacker’s purpose. Crucially, for standard benign fine-tuning, the updates do not engage this instability; the mechanism remains inactive, allowing the LLM to learn new tasks and maintain its high utility and general capabilities for legitimate users. CTRAP thus provides a targeted defense that incapacitates the model only when it’s being actively steered towards harm.

In conclusion, the main contributions of this paper are threefold: **1)** We identify the limitation of selective unlearning against harmful fine-tuning, linking it to the LLM’s exploitable general adaptability. **2)** We propose the concept of conditional model collapse (“unlearning everything” when subjected to harmful fine-tuning dynamics) as a more fundamental defense strategy, and introduce CTRAP as its practical implementation. **3)** Extensive empirical results demonstrate that, across multiple LLMs (such as Gemma2-9B, Llama2-7B, and Qwen2-7B) and various harmful fine-tuning attack settings (including “full harmful” and “mix harmful” scenarios), CTRAP achieves state-of-the-art defense while preserving benign task performance.

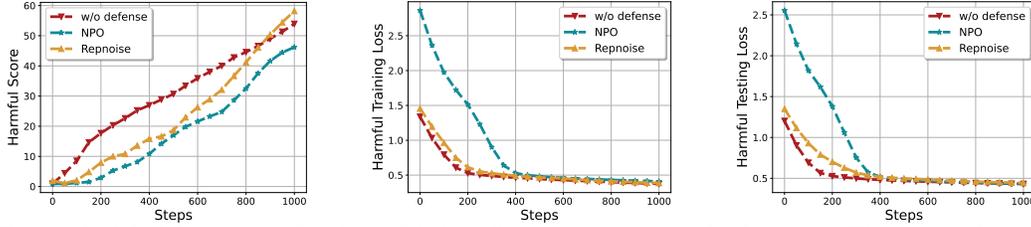


Figure 2: Model metrics after harmful data fine-tuning over multiple steps. The harmful score measures the harmfulness level in model outputs on the test set. Harmful training loss refers to loss on harmful training data, while harmful testing loss refers to loss on harmful test data.

2 Preliminaries

2.1 Problem Setup

Scenario. Harmful fine-tuning poses a significant security challenge for LLM fine-tuning service providers. In this scenario, users upload specific datasets to the service provider, which then utilizes these datasets to fine-tune their safety-aligned foundation model. The resulting fine-tuned models are hosted on the service provider’s servers and are tailored to deliver personalized outputs to users. We assume that an adversary uploads a harmful or partly harmful fine-tuning dataset to obtain an unaligned LLM service. This enables them to utilize these powerful LLMs to execute malicious tasks like generating malicious code or fake news.

Defenders’ Capabilities. We assume the service provider maintains an alignment dataset D_A , which includes harmful prompt-safe answer pairs and helpful prompt-helpful answer pairs. Additionally, there is a harmful dataset D_H (consisting of harmful prompt-harmful answer pairs) used for defense. The availability of the three pairs of data is also made in (Rosati et al., 2024c; Huang et al., 2024c; Tamirisa et al., 2024; Li et al., 2024b).

Defenders’ Objectives. The ultimate goal for defenders is to maintain the utility of the fine-tuning API for users uploading benign datasets, while simultaneously preventing attackers from abusing the fine-tuning service to develop models for harmful purposes.

2.2 Revisiting Unlearning-based Defenses

Unlearning-based defenses represent a significant approach to mitigating harmful fine-tuning risks during the alignment stage. Their core strategy is to eradicate or neutralize harmful knowledge within the LLM, aiming to prevent its misuse for malicious tasks. Here, we briefly review two representative unlearning techniques proposed for this context. (Further details on baseline implementations are provided in Appendix B.)

- **Negative Preference Optimization (NPO).** Moving beyond simple gradient ascent on harmful examples (Yao et al., 2024; Jang et al., 2023), more sophisticated methods like NPO (Zhang et al., 2024b,a) leverages principles from preference optimization to adaptively control the unlearning process, pushing the model away from generating harmful responses.
- **Representation Noise (RepNoise).** Another line of work targets the model’s internal representations (Rosati et al., 2024c; Zou et al., 2024; Li et al., 2024a). RepNoise (Rosati et al., 2024c), a representative example, attempts to disrupt the model’s ability to process harmful inputs by steering their internal representations towards a noise distribution (e.g., Gaussian noise).

Empirical Reassessment. To understand the practical limitations, we conducted harmful fine-tuning attacks (using 500 malicious samples) on Llama-2-7b models pre-aligned with NPO and RepNoise defenses. We evaluated their resilience using 500 unseen harmful test prompts.

Unlearning defends against harmful fine-tuning attacks by increasing the loss of harmful samples. The left panel of Figure 2 shows that, compared to LLMs without such defenses, unlearning-based defenses demonstrate effective defense capabilities during the initial fine-tuning phase, achieving a lower harmful score. Moreover, we observe in the middle and right of Figure 2 that unlearning-based solutions initially result in higher training and testing loss, increasing the difficulty for the model to learn harmful samples.

The effectiveness of unlearning diminishes with increasing training steps. However, as the fine-tuning steps increase, the harmful score rapidly rises, gradually closing the gap with models without defenses and eventually reaching a comparable level. Meanwhile, the training and testing loss, although initially higher, does not reduce the convergence rate. Unlearning quickly converges to levels comparable to those without defenses after only 400 steps.

The limitation: general adaptability undermines unlearning. We attribute this failure not merely to imperfect unlearning but to a fundamental characteristic of modern LLMs: their powerful *general adaptability*. Selective unlearning techniques aim to remove or suppress specific knowledge pathways associated with harmful behaviors. Yet, they leave the model’s core abilities – its vast world knowledge, reasoning abilities, and potent capacity to learn from new data – largely intact. Harmful fine-tuning directly exploits this residual adaptability. The model does not necessarily need to rely on the precise knowledge pathways targeted by unlearning; instead, it leverages its general intelligence to quickly discern the patterns and objectives within the harmful fine-tuning data, effectively transferring its general capabilities to the malicious task. Thus, the root issue is the LLM’s inherent ability to repurpose its powerful general intelligence, allowing it to circumvent selective defenses and rapidly re-acquire harmful functionalities. This motivates the need for defense mechanisms that address this core general adaptability challenge.

3 Methodology

A primary challenge in safeguarding LLMs lies in their strong general adaptability, which often undermines unlearning-based defenses against harmful fine-tuning. To counter this fundamentally, we explore the concept of model collapse: intentionally inducing a loss of general capabilities in response to harmful updates, thereby rendering the model non-exploitable. However, a permanently collapsed model offers no utility. Therefore, we propose the collapse trap, a mechanism embedded during the LLM’s safety alignment phase. This allows the model to function normally for benign fine-tuning but triggers a progressive collapse when subjected to harmful fine-tuning updates.

3.1 Model Collapse: Inducing Functional Inertness

Distinct from selective unlearning that targets specific harmful knowledge, model collapse aims for a comprehensive degradation of *all capabilities*. It pushes the model towards a state of functional inertness, effectively “unlearning everything” when triggered. This prevents attackers from exploiting residual general abilities that might persist after more targeted interventions.

We implement model collapse by optimizing the model θ to predict a fixed, predefined token e with high probability, regardless of the preceding context $x \circ y_{<t}$. This objective, ℓ_{Collapse} , directly targets the core probabilistic language modeling function:

$$\ell_{\text{Collapse}}(\theta; \mathcal{D}) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[-\frac{1}{|y|} \sum_{t=1}^{|y|} \log p(e \mid x \circ y_{<t}; \theta) \right], \quad (1)$$

where \mathcal{D} is a dataset representing general dialogue. Minimizing ℓ_{Collapse} forces the model’s output distribution $p(\cdot \mid x \circ y_{<t}; \theta)$ to become sharply peaked at the single token e , effectively ignoring the context $x \circ y_{<t}$. This optimization pressure disincantivizes the learning of meaningful representations and context-aware attention patterns, as these intricate mechanisms, essential for genuine language modeling, are no longer required to achieve the simple goal of predicting e . The resulting degradation of these core processing capabilities leads to a comprehensive loss of language understanding and generation abilities, achieving functional inertness.

3.2 Collapse Trap: Embedding Conditional Collapse

To maintain utility for legitimate users, the collapse trap is implanted during alignment to yield parameters θ^* . The trap remains dormant unless harmful fine-tuning is attempted. The training objective balances standard alignment with trap implantation:

$$\arg \min_{\theta} \underbrace{\ell(\theta; \mathcal{D}_{\text{alignment}})}_{\text{Standard Alignment}} + \lambda \underbrace{\ell_{\text{Collapse}}(\theta - \alpha \cdot \nabla_{\theta} \ell(\theta; \mathcal{D}_{\text{harmful}}); \mathcal{D}_{\text{general}})}_{\text{Collapse Trap Planting}}. \quad (2)$$

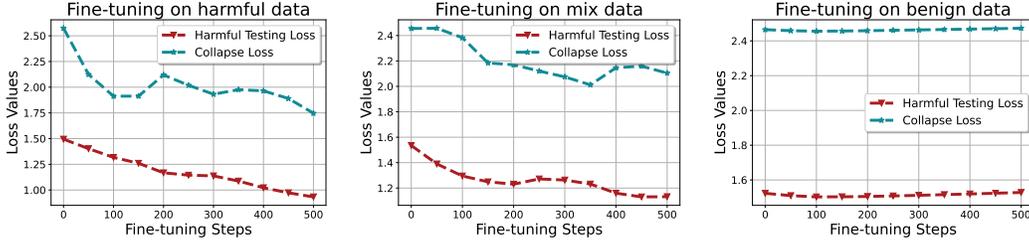


Figure 3: Fine-tuning dynamics after CTRAP implantation. (Left) Under pure harmful fine-tuning, harmful loss decreases while collapse loss sharply increases. (Middle) With mixed data, both losses change more gradually. (Right) Under pure benign fine-tuning, both losses remain stable.

The first term, $\ell(\theta; \mathcal{D}_{\text{alignment}})$, represents the standard alignment objective, encouraging the model to learn desired safe and helpful behaviors based on the alignment dataset. The second term, weighted by the hyperparameter λ , constitutes the core *Collapse Trap Planting* mechanism. Its purpose is to proactively shape the model’s parameter space such that any attempt to move in a “harmful direction” during subsequent fine-tuning will lead the model towards functional collapse. This term operates through a three-step internal process:

- **Identifying the Harmful Direction:** It first calculates the gradient $\nabla_{\theta} \ell(\theta; \mathcal{D}_{\text{harmful}})$ using a representative harmful dataset $\mathcal{D}_{\text{harmful}}$. This gradient vector points in the direction within the parameter space that corresponds to the model learning the harmful behaviors present in $\mathcal{D}_{\text{harmful}}$. It essentially simulates the intent of a harmful fine-tuning update.
- **Simulating a Harmful Step:** It then anticipates the result of taking a small step (α) in this harmful direction, yielding hypothetical parameters $\theta' = \theta - \alpha \cdot \nabla_{\theta} \ell(\theta; \mathcal{D}_{\text{harmful}})$. This θ' represents where the model would land after a single harmful fine-tuning update.
- **Evaluating Collapse Potential:** Finally, it evaluates the collapse loss $\ell_{\text{Collapse}}(\theta'; \mathcal{D}_{\text{general}})$ on the general dataset (sampled from a human dialogue distribution) using these hypothetical parameters θ' . This measures how prone the model would become to generating collapsed outputs (predicting the fixed token e) if it were updated in that harmful direction.

By minimizing the entire objective in Equation 2, the training process searches for parameters θ^* that satisfy two conditions simultaneously: (1) they perform well on the standard alignment task (low value in first term), and (2) they result in a low collapse loss *if perturbed in a harmful direction* (low value in second term). This encourages parameters θ^* that are (1) well-aligned under normal conditions, (2) but are inherently unstable and prone to collapse when subjected to harmful updates.

Figure 3 empirically illustrates the behavior of a CTRAP-enabled LLM during the fine-tuning phase, plotting loss metrics evaluated on held-out test sets.

- **Harmful Fine-tuning:** As the model adapts to purely harmful data (harmful loss drops), the collapse loss rises, indicating the trap’s activation and the intended degradation of general capabilities.
- **Mixed Fine-tuning:** When fine-tuning on a mix of benign and harmful data, the model learns harmfulness more slowly (slower harmful loss decrease), and correspondingly, the collapse loss increases more gradually. This behavior follows the same trend observed during pure harmful fine-tuning, confirming that the collapse trap is indeed activated by the harmful updates.
- **Benign Fine-tuning:** With purely benign data, the model does not learn harmful behaviors (harmful loss remains high), and crucially, the collapse loss stays low. This demonstrates the trap remains inactive during legitimate use, preserving utility.

4 Experiment

4.1 Setup

Datasets and Models. During the alignment phase, we use the alignment dataset and harmful dataset from Rosati et al. (2024d), which is enriched from BeaverTails (Ji et al., 2023). We sample 5000 instances to construct the alignment dataset, and another 5000 instances to construct the harmful dataset. Additionally, we sample 5000 instances from the helpful dataset UltraChat (Ding et al., 2023)

and include them in the alignment dataset. This is done to prevent the model from overfitting and learning to refuse all types of questions indiscriminately. This set also serves as the general dataset used to compute the collapse loss, representing samples simulating the human dialogue distribution.

We consider SST2 (Socher et al., 2013), AGNEWS (Zhang et al., 2015), and GSM8K (Cobbe et al., 2021) as the fine-tuning tasks for benign users, and set the sample size to 500 by default. For malicious users, we follow (Huang et al., 2024e) to evaluate two settings: a “full” setting where attackers upload fully harmful datasets, and a “mix” setting where they upload clean datasets but secretly mix in a small ratio of harmful data. Following (Huang et al., 2024f,d), we use harmful datasets with the same distribution as, but different instances from, those used in the alignment stage. For “full” settings, we vary the number of harmful samples between 100, 200, 300, 400, and 500. For the “mix” setting, SST2 is used as clean data and we test poison ratios of 5%, 10%, 15%, 20%, and 25%. We use Gemma2-9B (Team et al., 2024), Llama2-7B (Touvron et al., 2023), and Qwen2-7B (Yang et al., 2024) for evaluation.

Metrics. Following (Huang et al., 2024f,d), we evaluate model performance using: *Fine-tuning Accuracy (FA)*, the accuracy on the benign fine-tuning task’s test set (details in Appendix A); *Input-Output Harmful Score (HS(IO))*, the ratio of input-output pairs classified as harmful by the moderation model from (Ji et al., 2023); and *Output Harmful Score (HS(O))*, which calculates the harmful score based solely on the model’s output. HS(O) is included because HS(IO) sometimes requires explicit refusal of harmful input to be considered safe, potentially penalizing unlearning-based defenses that aim for nonsensical outputs instead.

To calculate the harmful score, we sample 1000 instructions from the testing set of BeaverTails (Ji et al., 2023). To obtain fine-tuning accuracy, we sample 872, 1000, and 1000 samples from the fine-tuning datasets SST2, AGNEWS, and GSM8K respectively.

Baselines. We use five baselines for comparison. SFT is the vanilla supervised fine-tuning solution. Vaccine (Huang et al., 2024f) and Booster (Huang et al., 2024c) are two alignment stage solutions aimed at improving the robustness of alignment concerning the harmful fine-tuning issue. NPO (Zhang et al., 2024b) and Repnoise (Rosati et al., 2024c) are two alignment stage solutions aimed at forgetting the malicious capabilities of the model.

Training Details. We utilize LoRA (Hu et al., 2021) to enhance the efficiency of LLM training following (Huang et al., 2024f,d; Hsu et al., 2024). The adapter’s rank is configured to 32, with LoRA’s alpha set at 4. For alignment, AdamW (Loshchilov et al., 2017) is used as the optimizer, featuring a learning rate of $5e-4$ and a weight decay of 0.1. For fine-tuning tasks, we apply the same optimizer but with a reduced learning rate of $1e-5$, as outlined in (Huang et al., 2024f,c). Training involves 20 epochs for alignment and another 20 for both benign and harmful fine-tuning tasks, using a batch size of 10 throughout all phases. The hyper-parameter is set to $\alpha = 0.1$ and $\lambda = 0.1$ by default. All the experiments are done with 8 A800-80Gs. Further information is available in Appendix A.

4.2 Main Experiments

Table 1: Defensive performance against harmful fine-tuning attacks (full harmful) on Gemma2-9B.

Methods	harmful nums=100		harmful nums=200		harmful nums=300		harmful nums=400		harmful nums=500		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)								
SFT	7.1	4.4	22.6	17.1	43.8	36.6	58.2	49.6	65.5	56.2	39.4	32.8
Vaccine	4.3	2.4	19.4	14.4	36.9	28.3	50.4	39.8	58.0	46.2	33.8	26.2
Booster	4.0	2.2	16.4	11.5	47.1	39.7	60.8	52.6	66.9	56.2	39.0	32.4
Repnoise	10.0	5.5	21.2	15.1	39.7	31.6	52.6	42.6	62.6	53.1	37.2	29.6
NPO	1.2	0.7	13.9	9.7	33.9	25.8	50.1	40.4	61.0	50.0	32.0	25.3
CTRAP	2.7	0.5	2.5	0.5	2.5	0.5	7.2	4.8	11.3	7.1	5.2	2.7

Table 2: Defensive performance against harmful fine-tuning attacks (mix harmful) on Gemma2-9B.

Methods	harmful ratio=0.05		harmful ratio=0.1		harmful ratio=0.15		harmful ratio=0.2		harmful ratio=0.25		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)
SFT	4.7	2.6	9.4	4.8	16.2	11.6	22.3	16.4	28.4	21.4	16.2	11.4
Vaccine	1.4	0.7	4.8	3.0	10.5	6.7	20.9	14.8	26.2	19.6	12.8	9.0
Booster	2.2	1.1	6.5	3.6	9.9	6.8	12.9	8.6	19.4	14.8	10.2	7.0
Repnoise	7.5	4.0	13.1	7.4	18.5	11.3	24.0	15.4	27.6	19.8	18.1	11.6
NPO	1.4	0.7	4.8	3.0	10.5	6.7	20.9	14.8	26.2	19.6	12.8	9.0
CTRAP	1.7	1.0	2.5	0.9	1.3	0.8	1.9	0.9	3.3	0.6	2.1	0.8

Performance on Defending Harmful Fine-tuning Attacks. The performance of different defense baseline methods on defending harmful fine-tuning attacks is shown in Table 1 and Table 2. The

experimental results indicate that **our method achieves the best defending performance**. Specifically, our method outperforms the baselines in terms of both HS(IO) and HS(O) in all settings, with HS(IO)/HS(O) decreasing by an average of over 26%/22% in full setting and 8%/6% in mix setting compared to the best baseline performance. Additionally, our method shows better robustness against the increase in harmful samples and poison ratio, while the performance of other baseline methods declines sharply with more harmful samples, with their harmful score reaching levels similar to those of SFT without defense in the full setting. This validates our claim that the collapse trap, by causing the model to progressively collapse when faced with harmful fine-tuning attacks, effectively prevents malicious users from exploiting the model’s general capabilities for their intended harmful purposes.

Table 3: Defensive performance against harmful fine-tuning attacks on different models.

Methods	Llama2-7B				Qwen2-7B				Gemma2-9B				Average	
	Full		Mix		Full		Mix		Full		Mix		HS(IO)	HS(O)
	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)		
SFT	34.2	25.9	22.7	16.8	22.4	15.6	15.4	10.5	39.4	32.8	16.2	11.4	25.1	18.8
Vaccine	27.8	21.4	15.3	10.2	15.5	10.9	8.5	5.2	33.8	26.2	12.8	9.0	19.0	13.8
Reppoise	29.7	24.6	14.8	9.8	16.5	12.4	20.3	14.8	39.0	32.4	18.1	11.6	23.1	17.6
Booster	25.2	19.6	3.1	1.9	26.8	19.3	2.7	1.3	37.2	29.6	10.2	7.0	17.5	13.1
NPO	20.7	14.1	11.2	6.3	18.1	12.3	12.5	8.3	32.0	25.3	12.8	9.0	17.9	12.6
CTRAP	10.4	6.9	3.6	1.7	1.5	0.7	1.3	0.7	5.2	2.7	2.1	0.8	4.0	2.2

Generalization to Models. The previous experiments were conducted on Gemma2-9B. As shown in Table 3, we also conducted experiments on Llama2-7B and Qwen2-7B. For each LLM, we present the average performance across different numbers of harmful samples or different ratios, with details provided in Appendix C. The experimental results demonstrate that **our method can successfully generalize to different LLMs**. Our method achieves a 21.1% reduction in HS(IO) and a 16.6% reduction in HS(O) compared to SFT on average. Compared to the best baseline method, our method results in an average decrease of 13.5% in HS(IO) and an average reduction of 10.9% in HS(O).

Table 4: Performance analysis (fine-tuning accuracy) on benign fine-tuning tasks.

Methods	Llama2-7B			Qwen2-7B			Gemma2-9B			Average		
	SST2	Agnews	GSM8k									
SFT	92.7	85.9	10.9	92.4	84.2	60.5	94.0	86.6	50.7	93.0	85.6	40.7
Vaccine	90.8	86.3	7.2	90.1	84.3	58.1	90.9	85.8	43.0	90.6	85.5	36.1
Booster	91.6	84.8	12.7	93.2	85.2	61.8	93.7	86.8	56.7	92.9	85.6	43.7
Reppoise	91.4	86.5	8.7	91.7	84.2	63.1	91.6	87.0	51.5	91.6	85.9	41.1
NPO	93.0	86.9	11.2	92.0	84.6	67.8	92.7	85.6	54.4	92.5	85.7	44.5
CTRAP	92.3	85.9	10.8	94.5	82.4	57.5	94.2	86.5	53.8	93.7	84.9	40.7

Performance on Benign Fine-tuning. The performance of different defense baseline methods on benign fine-tuning tasks is shown in Table 4. From the table, we can observe that **our method does not affect the model’s performance on benign fine-tuning tasks** while achieving state-of-the-art defense performance. CTRAP achieves comparable fine-tuning performance to SFT, with the average fine-tuning accuracy only slightly decreasing by 0.7% on Agnews, and even achieving slightly better performance on SST2 and identical performance on GSM8k. Moreover, Vaccine shows decreased performance compared to SFT across different clean tasks, which might be due to the adverse effects caused by adversarial training.

4.3 Overhead Analysis

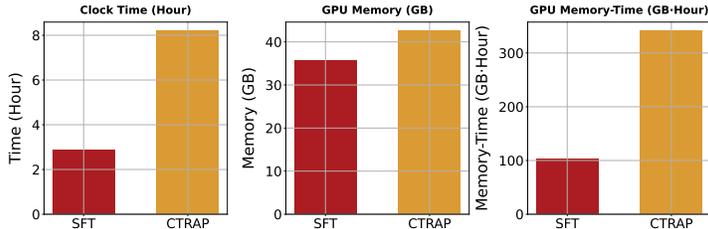


Figure 4: Overhead analysis of CTRAP.

CTRAP introduces additional overhead during the alignment phase compared to standard SFT, as illustrated in Figure 4. Specifically, CTRAP requires approximately $3.5\times$ the GPU memory-time and is $2.8\times$ slower in clock time, necessitating an extra 6.72GB of GPU memory during alignment. This increased cost stems from the core mechanism of CTRAP, which performs three gradient evaluations per optimization step, requiring storage for three gradient vectors and a batch of harmful

5 Related Work

Safety Alignment. Safety alignment (Reuel et al., 2024; Sicari et al., 2024; Verma et al., 2024) refers to aligning LLMs with human values, intentions, and ethical considerations to ensure their outputs are safe, reliable, and aligned with human expectations. The core idea is to regularize the model’s output such that the model is able to output a refusal answer whenever a harmful prompt is given. Common approaches include supervised fine-tuning (SFT), which uses supervised datasets of instructions and desired outputs to improve alignment and Reinforcement Learning with Human Feedback (RLHF), where human preferences guide model optimization (Ouyang et al., 2022; Dai et al., 2023; Bai et al., 2022; Wu et al., 2023; Dong et al., 2023; Rafailov et al., 2023; Yuan et al., 2023).

Harmful Fine-tuning Attacks. However, recent studies about harmful fine-tuning attacks (Qi et al., 2023; Yang et al., 2023; Zhan et al., 2023; Lermen et al., 2023; Chen et al., 2024; Rosati et al., 2024b; Yi et al., 2024a; Huang et al., 2024b, 2025) show that introducing a few harmful fine-tuning data points can cause the aligned model to forget its safety alignment, rendering it vulnerable to exploitation for malicious tasks. Unlike jailbreak attacks (Zou et al., 2023; Huang et al., 2024g), which only interfere during the inference stage of LLMs, harmful fine-tuning attacks grant attackers elevated privileges, allowing them to directly alter model weights via the fine-tuning process. This makes defending against such attacks particularly challenging (Rosati et al., 2024a). Recent research also studies the mechanism of harmful fine-tuning (Leong et al., 2024; Peng et al., 2024; Anonymous, 2024c; Qi et al., 2024b; Guo et al., 2024).

Harmful Fine-tuning Defenses. Existing mitigation approaches to this problem can be grouped into three categories based on the stage at which the mitigation is applied: alignment-stage methods (Huang et al., 2024f; Rosati et al., 2024c,d; Huang et al., 2024c; Liu et al., 2024; Tamirisa et al., 2024), fine-tuning-stage methods (Mukhoti et al., 2023; Huang et al., 2024d; Lyu et al., 2024; Wang et al., 2024; Qi et al., 2024a; Bianchi et al., 2023; Zong et al., 2024; Wei et al., 2024; Eiras et al., 2024; Du et al., 2024; Anonymous, 2024a; Shen et al., 2024; Anonymous, 2024b; Choi et al., 2024; Luo et al., 2024), and post-fine-tuning stage methods (Hsu et al., 2024; Yi et al., 2024c; Huang et al., 2024a; Zhu et al., 2024; Casper et al., 2024; Wu et al., 2024; Gudipudi et al., 2024; Yi et al., 2024b). This paper focuses on studying alignment-stage solutions, which require a one-time cost rather than intervening in every user fine-tuning task, as is necessary with solutions applied at other stages. Existing alignment-stage methods primarily rely on the idea of adversarial training (Huang et al., 2024f,c; Tamirisa et al., 2024) to enhance the robustness of alignment and apply unlearning (Zhang et al., 2024a,b; Rosati et al., 2024c) techniques to remove harmful knowledge.

Machine Unlearning. Machine unlearning (Bourtole et al., 2021; Yuan et al., 2024; Gao et al., 2024) originally emerged as a technique aimed at addressing data privacy and compliance issues, particularly within the context of user data. Recently, researchers have advanced the use of machine unlearning beyond its original motivation to tackle safety and robustness challenges in LLMs (Li et al., 2024a; Zhang et al., 2024b; Rosati et al., 2024c; Yao et al., 2024). This extension is driven by the observation that unlearning techniques provide a promising approach for mitigating harmful memorization introduced during training. Moreover, recent studies (Deeb & Roger, 2024; Lynch et al., 2024; Lucki et al., 2024; Shumailov et al., 2024; Barez et al., 2025) have highlighted some flaws in unlearning methods. For instance, Lucki et al. (2024) find that they are highly susceptible to adversarial attacks. In addition, this paper argues that due to the strong general adaptability of LLMs, unlearning methods are fundamentally challenging to resolve harmful fine-tuning attacks.

6 Conclusion

In this paper, we first demonstrated that the unlearning paradigm is insufficient to effectively defend against harmful fine-tuning attacks. We attribute this limitation to the strong general adaptability of LLMs, which allows them to transfer pre-trained general knowledge into harmful tasks, even after undergoing unlearning processes. To overcome this challenge, we proposed the collapse trap as a solution. This method involves embedding a trap during alignment that alters the model’s behavior during fine-tuning. When subjected to harmful fine-tuning updates, the trap causes a progressive degradation of the model’s fundamental language modeling ability. This prevents malicious users from leveraging the general capability of the LLM to achieve high performance on malicious tasks. Experimental results demonstrate that the proposed approach is effective in mitigating the risks posed by harmful fine-tuning while maintaining performance in benign fine-tuning scenarios.

References

- Anonymous. Safety alignment shouldn't be complicated. In *Submitted to The Thirteenth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=9H91juqfjb>. under review. 9
- Anonymous. SaloRA: Safety-alignment preserved low-rank adaptation. In *Submitted to The Thirteenth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=GOoVzE9nSj>. under review. 9
- Anonymous. Your task may vary: A systematic understanding of alignment and safety degradation when fine-tuning LLMs. In *Submitted to The Thirteenth International Conference on Learning Representations*, 2024c. URL <https://openreview.net/forum?id=vQ0zFYJaMo>. under review. 9
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022. 9
- Barez, F., Fu, T., Prabhu, A., Casper, S., Sanyal, A., Bibi, A., O'Gara, A., Kirk, R., Bucknall, B., Fist, T., et al. Open problems in machine unlearning for ai safety. *arXiv preprint arXiv:2501.04952*, 2025. 9
- Bianchi, F., Suzgun, M., Attanasio, G., Röttger, P., Jurafsky, D., Hashimoto, T., and Zou, J. Safety-tuned llamas: Lessons from improving the safety of large language models that follow instructions. *arXiv preprint arXiv:2309.07875*, 2023. 9
- Bourtole, L., Chandrasekaran, V., Choquette-Choo, C. A., Jia, H., Travers, A., Zhang, B., Lie, D., and Papernot, N. Machine unlearning. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, 2021. 9
- Casper, S., Schulze, L., Patel, O., and Hadfield-Menell, D. Defending against unforeseen failure modes with latent adversarial training. *arXiv preprint arXiv:2403.05030*, 2024. 9
- Chen, C., Huang, B., Li, Z., Chen, Z., Lai, S., Xu, X., Gu, J.-C., Gu, J., Yao, H., Xiao, C., et al. Can editing llms inject harm? *arXiv preprint arXiv:2407.20224*, 2024. 9
- Choi, H. K., Du, X., and Li, Y. Safety-aware fine-tuning of large language models. *arXiv preprint arXiv:2410.10014*, 2024. 9
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 6
- Dai, J., Pan, X., Sun, R., Ji, J., Xu, X., Liu, M., Wang, Y., and Yang, Y. Safe rlhf: Safe reinforcement learning from human feedback. *arXiv preprint arXiv:2310.12773*, 2023. 9
- Deeb, A. and Roger, F. Do unlearning methods remove information from language model weights? *arXiv:2410.08827*, 2024. 9
- Ding, N., Chen, Y., Xu, B., Qin, Y., Hu, S., Liu, Z., Sun, M., and Zhou, B. Enhancing chat language models by scaling high-quality instructional conversations. In *EMNLP*, 2023. 5, 15
- Dong, H., Xiong, W., Goyal, D., Pan, R., Diao, S., Zhang, J., Shum, K., and Zhang, T. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*, 2023. 9
- Du, Y., Zhao, S., Cao, J., Ma, M., Zhao, D., Fan, F., Liu, T., and Qin, B. Towards secure tuning: Mitigating security risks arising from benign instruction fine-tuning. *arXiv preprint arXiv:2410.04524*, 2024. 9
- Eiras, F., Petrov, A., Torr, P. H., Kumar, M. P., and Bibi, A. Mimicking user data: On mitigating fine-tuning risks in closed large language models. *arXiv preprint arXiv:2406.10288*, 2024. 9

- Gao, H., Pang, T., Du, C., Hu, T., Deng, Z., and Lin, M. Meta-unlearning on diffusion models: Preventing relearning unlearned concepts. *arXiv:2410.12777*, 2024. 9
- Gudipudi, S. S., Vipparla, S., Singh, H., Goel, S., and Kumaraguru, P. Enhancing ai safety through the fusion of low rank adapters. *arXiv preprint arXiv:2501.06208*, 2024. 9
- Guo, Y., Jiao, F., Nie, L., and Kankanhalli, M. S. The VLLM safety paradox: Dual ease in jailbreak attack and defense. *arXiv preprint arXiv:2411.08410*, 2024. 9
- Halawi, D., Wei, A., Wallace, E., Wang, T. T., Haghtalab, N., and Steinhardt, J. Covert malicious finetuning: Challenges in safeguarding llm adaptation. *arXiv preprint arXiv:2406.20053*, 2024. 1
- He, L., Xia, M., and Henderson, P. What’s in your” safe” data?: Identifying benign data that breaks safety. *arXiv preprint arXiv:2404.01099*, 2024. 1
- Hsu, C.-Y., Tsai, Y.-L., Lin, C.-H., Chen, P.-Y., Yu, C.-M., and Huang, C.-Y. Safe lora: the silver lining of reducing safety risks when fine-tuning large language models. *arXiv preprint arXiv:2405.16833*, 2024. 6, 9
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 6
- Huang, T., Bhattacharya, G., Joshi, P., Kimball, J., and Liu, L. Antidote: Post-fine-tuning safety alignment for large language models against harmful fine-tuning. *arXiv preprint arXiv:2408.09600*, 2024a. 9
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Harmful fine-tuning attacks and defenses for large language models: A survey. *arXiv preprint arXiv:2409.18169*, 2024b. 9
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Booster: Tackling harmful fine-tuning for large language models via attenuating harmful perturbation. *arXiv preprint arXiv:2409.01586*, 2024c. 1, 3, 6, 9, 15, 16, 17
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Lazy safety alignment for large language models against harmful fine-tuning. *arXiv preprint arXiv:2405.18641*, 2024d. 6, 9
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Harmful fine-tuning attacks and defenses for large language models: A survey. *arXiv preprint arXiv:2403.04786*, 2024e. 6
- Huang, T., Hu, S., and Liu, L. Vaccine: Perturbation-aware alignment for large language model. *arXiv preprint arXiv:2402.01109*, 2024f. 1, 6, 9, 15, 16
- Huang, T., Hu, S., Ilhan, F., Tekin, S. F., and Liu, L. Virus: Harmful fine-tuning attack for large language models bypassing guardrail moderation. *arXiv:2501.17433*, 2025. 9
- Huang, Y., Gupta, S., Xia, M., Li, K., and Chen, D. Catastrophic jailbreak of open-source llms via exploiting generation. In *ICLR*, 2024g. 9
- Jang, J., Yoon, D., Yang, S., Cha, S., Lee, M., Logeswaran, L., and Seo, M. Knowledge unlearning for mitigating privacy risks in language models. In *ACL*, 2023. 3
- Ji, J., Liu, M., Dai, J., Pan, X., Zhang, C., Bian, C., Sun, R., Wang, Y., and Yang, Y. Beaver-tails: Towards improved safety alignment of llm via a human-preference dataset. *arXiv preprint arXiv:2307.04657*, 2023. 5, 6, 15
- Leong, C. T., Cheng, Y., Xu, K., Wang, J., Wang, H., and Li, W. No two devils alike: Unveiling distinct mechanisms of fine-tuning attacks. *arXiv preprint arXiv:2405.16229*, 2024. 9
- Lermen, S., Rogers-Smith, C., and Ladish, J. Lora fine-tuning efficiently undoes safety training in llama 2-chat 70b. *arXiv preprint arXiv:2310.20624*, 2023. 1, 9

- Li, N., Pan, A., Gopal, A., Yue, S., Berrios, D., Gatti, A., Li, J. D., Dombrowski, A., Goel, S., Mukobi, G., Helm-Burger, N., Lababidi, R., Justen, L., Liu, A. B., Chen, M., Barrass, I., Zhang, O., Zhu, X., Tamirisa, R., Bharathi, B., Herbert-Voss, A., Breuer, C. B., Zou, A., Mazeika, M., Wang, Z., Oswal, P., Lin, W., Hunt, A. A., Tienken-Harder, J., Shih, K. Y., Talley, K., Guan, J., Steneker, I., Campbell, D., Jokubaitis, B., Basart, S., Fitz, S., Kumaraguru, P., Karmakar, K. K., Tupakula, U. K., Varadharajan, V., Shoshitaishvili, Y., Ba, J., Esvelt, K. M., Wang, A., and Hendrycks, D. The WMDP benchmark: Measuring and reducing malicious use with unlearning. In *ICML*, 2024a. 1, 3, 9
- Li, S., Yao, L., Zhang, L., and Li, Y. Safety layers of aligned large language models: The key to llm security. *arXiv preprint arXiv:2408.17003*, 2024b. 3
- Liu, X., Liang, J., Ye, M., and Xi, Z. Robustifying safety-aligned large language models through clean data curation. *arXiv preprint arXiv:2405.19358*, 2024. 9
- Loshchilov, I., Hutter, F., et al. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5, 2017. 6
- Luo, J., Luo, X., Ding, K., Yuan, J., Xiao, Z., and Zhang, M. Robustft: Robust supervised fine-tuning for large language models under noisy response. *arXiv preprint arXiv:2412.14922*, 2024. 9
- Lynch, A., Guo, P., Ewart, A., Casper, S., and Hadfield-Menell, D. Eight methods to evaluate robust unlearning in llms. *arXiv preprint arXiv:2402.16835*, 2024. 9
- Lyu, K., Zhao, H., Gu, X., Yu, D., Goyal, A., and Arora, S. Keeping llms aligned after fine-tuning: The crucial role of prompt templates. *arXiv preprint arXiv:2402.18540*, 2024. 9
- Mukhoti, J., Gal, Y., Torr, P. H., and Dokania, P. K. Fine-tuning can cripple your foundation model; preserving features may be the solution. *arXiv preprint arXiv:2308.13320*, 2023. 9
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022. 9
- Peng, S., Chen, P.-Y., Hull, M., and Chau, D. H. Navigating the safety landscape: Measuring risks in finetuning large language models. *arXiv preprint arXiv:2405.17374*, 2024. 9
- Qi, X., Zeng, Y., Xie, T., Chen, P.-Y., Jia, R., Mittal, P., and Henderson, P. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*, 2023. 1, 9
- Qi, X., Panda, A., Lyu, K., Ma, X., Roy, S., Beirami, A., Mittal, P., and Henderson, P. Safety alignment should be made more than just a few tokens deep. *arXiv preprint arXiv:2406.05946*, 2024a. 9
- Qi, X., Wei, B., Carlini, N., Huang, Y., Xie, T., He, L., Jagielski, M., Nasr, M., Mittal, P., and Henderson, P. On evaluating the durability of safeguards for open-weight llms. *arXiv preprint arXiv:2412.07097*, 2024b. 9
- Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C. D., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023. 9
- Reuel, A., Bucknall, B., Casper, S., Fist, T., Soder, L., Aarne, O., Hammond, L., Ibrahim, L., Chan, A., Wills, P., et al. Open problems in technical ai governance. *arXiv preprint arXiv:2407.14981*, 2024. 9
- Rosati, D., Edkins, G., Raj, H., Atanasov, D., Majumdar, S., Rajendran, J., Rudzicz, F., and Sajjad, H. Defending against reverse preference attacks is difficult. *arXiv preprint arXiv:2409.12914*, 2024a. 9
- Rosati, D., Edkins, G., Raj, H., Atanasov, D., Majumdar, S., Rajendran, J., Rudzicz, F., and Sajjad, H. Defending against reverse preference attacks is difficult. *arXiv preprint arXiv:2409.12914*, 2024b. 9

- Rosati, D., Wehner, J., Williams, K., Bartoszcze, Ł., Atanasov, D., Gonzales, R., Majumdar, S., Maple, C., Sajjad, H., and Rudzicz, F. Representation noising effectively prevents harmful fine-tuning on llms. *arXiv preprint arXiv:2405.14577*, 2024c. 1, 2, 3, 6, 9, 16
- Rosati, D., Wehner, J., Williams, K., Bartoszcze, Ł., Batzner, J., Sajjad, H., and Rudzicz, F. Immunization against harmful fine-tuning attacks. *arXiv preprint arXiv:2402.16382*, 2024d. 5, 9, 15
- Shen, H., Chen, P.-Y., Das, P., and Chen, T. Seal: Safety-enhanced aligned llm fine-tuning via bilevel data selection. *arXiv preprint arXiv:2410.07471*, 2024. 8, 9
- Shumailov, I., Hayes, J., Triantafillou, E., Ortiz-Jiménez, G., Papernot, N., Jagielski, M., Yona, I., Howard, H., and Bagdasaryan, E. Unlearning: Unlearning is not sufficient for content regulation in advanced generative AI. *arXiv preprint arXiv:2407.00106*, 2024. 9
- Sicari, S., Cevallos M, J. F., Rizzardi, A., and Coen-Portisini, A. Open-ethical ai: Advancements in open-source human-centric neural language models. *ACM Computing Surveys*, 2024. 9
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A. Y., and Potts, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp. 1631–1642, 2013. 6
- Tamirisa, R., Bharathi, B., Phan, L., Zhou, A., Gatti, A., Suresh, T., Lin, M., Wang, J., Wang, R., Arel, R., et al. Tamper-resistant safeguards for open-weight llms. *arXiv preprint arXiv:2408.00761*, 2024. 1, 3, 9
- Team, G., Riviere, M., Pathak, S., Sessa, P. G., Hardin, C., Bhupatiraju, S., Hussenot, L., Mesnard, T., Shahriari, B., Ramé, A., et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024. 6
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 6
- Verma, A., Krishna, S., Gehrmann, S., Seshadri, M., Pradhan, A., Ault, T., Barrett, L., Rabinowitz, D., Doucette, J., and Phan, N. Operationalizing a threat model for red-teaming large language models (llms). *arXiv preprint arXiv:2407.14937*, 2024. 9
- Wang, J., Li, J., Li, Y., Qi, X., Chen, M., Hu, J., Li, Y., Li, B., and Xiao, C. Mitigating fine-tuning jailbreak attack with backdoor enhanced alignment. *arXiv preprint arXiv:2402.14968*, 2024. 9
- Wei, B., Huang, K., Huang, Y., Xie, T., Qi, X., Xia, M., Mittal, P., Wang, M., and Henderson, P. Assessing the brittleness of safety alignment via pruning and low-rank modifications. *arXiv preprint arXiv:2402.05162*, 2024. 9
- Wu, D., Lu, X., Zhao, Y., and Qin, B. Separate the wheat from the chaff: A post-hoc approach to safety re-alignment for fine-tuned language models. *arXiv preprint arXiv:2412.11041*, 2024. 9
- Wu, T., Zhu, B., Zhang, R., Wen, Z., Ramchandran, K., and Jiao, J. Pairwise proximal policy optimization: Harnessing relative feedback for llm alignment. *arXiv preprint arXiv:2310.00212*, 2023. 9
- Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024. 6
- Yang, X., Wang, X., Zhang, Q., Petzold, L., Wang, W. Y., Zhao, X., and Lin, D. Shadow alignment: The ease of subverting safely-aligned language models. *arXiv preprint arXiv:2310.02949*, 2023. 1, 9
- Yao, Y., Xu, X., and Liu, Y. Large language model unlearning. In *ICLR*, 2024. 3, 9
- Yi, J., Ye, R., Chen, Q., Zhu, B., Chen, S., Lian, D., Sun, G., Xie, X., and Wu, F. On the vulnerability of safety alignment in open-access llms. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 9236–9260, 2024a. 1, 9

- Yi, X., Zheng, S., Wang, L., de Melo, G., Wang, X., and He, L. Nlsr: Neuron-level safety realignment of large language models against harmful fine-tuning. *arXiv preprint arXiv:2412.12497*, 2024b. 9
- Yi, X., Zheng, S., Wang, L., Wang, X., and He, L. A safety realignment framework via subspace-oriented model fusion for large language models. *arXiv preprint arXiv:2405.09055*, 2024c. 9
- Yuan, X., Pang, T., Du, C., Chen, K., Zhang, W., and Lin, M. A closer look at machine unlearning for large language models. *arXiv preprint arXiv:2410.08109*, 2024. 9
- Yuan, Z., Yuan, H., Tan, C., Wang, W., Huang, S., and Huang, F. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*, 2023. 9
- Zhan, Q., Fang, R., Bindu, R., Gupta, A., Hashimoto, T., and Kang, D. Removing rlhf protections in gpt-4 via fine-tuning. *arXiv preprint arXiv:2311.05553*, 2023. 1, 9
- Zhang, R., Lin, L., Bai, Y., and Mei, S. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024a. 2, 3, 9
- Zhang, X., Zhao, J., and LeCun, Y. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015. 6
- Zhang, Z., Yang, J., Ke, P., Cui, S., Zheng, C., Wang, H., and Huang, M. Safe unlearning: A surprisingly effective and generalizable solution to defend against jailbreak attacks. *arXiv preprint arXiv:2407.02855*, 2024b. 1, 2, 3, 6, 9, 16
- Zhu, M., Yang, L., Wei, Y., Zhang, N., and Zhang, Y. Locking down the finetuned llms safety. *arXiv preprint arXiv:2410.10343*, 2024. 9
- Zong, Y., Bohdal, O., Yu, T., Yang, Y., and Hospedales, T. Safety fine-tuning at (almost) no cost: A baseline for vision large language models. *arXiv preprint arXiv:2402.02207*, 2024. 9
- Zou, A., Wang, Z., Kolter, J. Z., and Fredrikson, M. Universal and transferable adversarial attacks on aligned language models. *arXiv:2307.15043*, 2023. 9
- Zou, A., Phan, L., Wang, J., Duenas, D., Lin, M., Andriushchenko, M., Wang, R., Kolter, Z., Fredrikson, M., and Hendrycks, D. Improving alignment and robustness with circuit breakers. *arXiv preprint arXiv:2406.04313*, 2024. 1, 2, 3
- Łucki, J., Wei, B., Huang, Y., Henderson, P., Tramèr, F., and Rando, J. An adversarial perspective on machine unlearning for ai safety. *arXiv preprint arXiv:2409.18025*, 2024. 9

A Experimental Details

In this section, we provide a detailed explanation of the experimental setup used in our testbed.

Hyper-parameters. During the alignment phase, we set the learning rate to $5e-4$ and use a batch size of 10. The number of alignment samples, helpful samples, and harmful samples used are all 5,000. Alignment samples and harmful samples are sampled from (Rosati et al., 2024d), which is enriched from BeaverTails (Ji et al., 2023). Helpful samples are sampled from the helpful dataset UltraChat (Ding et al., 2023).

In the fine-tuning phase, the learning rate is adjusted to $1e-5$, while the batch size remains 10. Harmful instances are drawn from the BeaverTails dataset Ji et al. (2023), and benign fine-tuning samples are selected from the dataset relevant to the specific task. For instance, benign samples for the GSM8K task are taken from the GSM8K training set¹. By default, we use a total of $n = 500$ fine-tuning samples.

Prompt Template. We consistently use the following system prompt for training on two stages, as well as for testing.

Prompt: Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request. Instruction:{instruction} Input:{input} Response:
Output: {output}

We follow (Huang et al., 2024f,c) to construct the prompt templates of different tasks. Here are examples of how we create prompt templates for different tasks: alignment, harmful fine-tuning attacks, SST2, AGNEWS, and GSM8K.

Alignment

instruction: (real harmful instruction)
input: (None)
output: (real safe output, e.g., I can't answer this question for you)

Harmful Fine-tuning Attacks

instruction: (real harmful instruction)
input: (None)
output: (real unsafe output)

SST2 (benign fine-tuning task)

instruction: Analyze the sentiment of the input, and respond only positive or negative.
input: (real input from SST2 dataset)
output: (real label from SST2 dataset, e.g., positive)

AGNEWS (benign fine-tuning task)

instruction: Categorize the news article into one of the 4 categories: World,Sports,Business,Sci/Tech.
input: (real input from AGNEWS dataset)
output: (real label from AGNEWS dataset, e.g., Sports)

¹<https://huggingface.co/datasets/openai/gsm8k>

GSM8K (benign fine-tuning task)

instruction: (the real input from GSM8K dataset)
input: (None)
output: (real output from GSM8K dataset)

For SST2 and AGNEWS, a sample in the fine-tuning task is deemed correct if the model generates the accurate classification result. In the GSM8K task, a sample is considered correct if the final answer provided by the LLM is correct, irrespective of the reasoning process involved.

B Baseline Descriptions

In this section, we provide a concise overview of how the existing baselines are applied in our experiments.

- **SFT.** We apply standard supervised fine-tuning (SFT) for aligning the model with the alignment dataset. Afterwards, we implement regular SFT for training on the downstream user dataset.
- **Vaccine** (used in the alignment stage). The Vaccine algorithm (Huang et al., 2024f) is employed during the alignment stage to align the model with the alignment dataset. Afterwards, we implement regular SFT for training on the downstream user dataset. In our experiment, the hyper-parameter for Vaccine is set to $\rho = 5$.
- **Booster** (used in the fine-tuning stage). We utilize the Booster algorithm (Huang et al., 2024c) at the alignment stage to align the model with the alignment and harmful dataset, followed by standard SFT for the downstream user dataset. We select the hyper-parameters as $\alpha = 0.1$ and $\lambda = 5$.
- **NPO** (used in the alignment stage). The NPO algorithm (Zhang et al., 2024b) is applied during the alignment stage to align the model with the alignment and harmful dataset, and standard SFT is then used for the downstream user dataset. The chosen hyper-parameter is $\lambda = 1$.
- **Reproise** (used in the alignment stage). The Reproise algorithm (Rosati et al., 2024c) is utilized at the alignment stage for aligning the model with the alignment and harmful dataset, followed by regular SFT for the downstream user dataset. The hyper-parameters are set to $\alpha = 1$ and $\beta = 0.001$.

For CTRAP, as it is an alignment stage solution, we use CTRAP algorithm to align the model with alignment and harmful dataset. Its hyper-parameter is default selected as $\alpha = 0.1$ and $\lambda = 0.1$.

Then we introduce the high level idea of each defense baseline.

- **Vaccine** (alignment stage solution). Vaccine attributes the success of harmful fine-tuning attacks to the embedding drift in the fine-tuning stage. The proposed approach involves introducing artificial perturbations to the embeddings during the model alignment phase. This aims to decrease the model’s sensitivity to the drift that occurs in the fine-tuning stage, effectively achieving a state of reduced perturbability. Vaccine is formalized as follows:

$$\min_w \max_{\|\epsilon\| \leq \rho} \frac{1}{|D_{alignment}|} \sum_{(x_i, y_i) \in D_{alignment}} \mathcal{L}((\tilde{f}_{w_L, \epsilon_L} \circ \dots \circ \tilde{f}_{w_1, \epsilon_1} \circ \mathcal{T})(x_i), y_i) \quad (3)$$

$$\text{s.t., } \tilde{f}_{w_l, \epsilon_l}(e_{l-1}) = f_{w_l}(e_{l-1}) + \epsilon_l \quad \forall l \in [L]$$

$$\epsilon = (\epsilon_1, \dots, \epsilon_L)$$

where $\tilde{f}_{w_l, \epsilon_l}(e_{l-1})$ is the l -th layer in a LLM that maps the input to a perturbed embedding and $\mathcal{T}(x_i)$ is the tokenizer function that produces embedding $e_{i,0}$. $D_{alignment}$ represents the alignment dataset.

- **Booster** (alignment stage solution). Similar to Vaccine, Booster uses harmful samples to simulate the weight permutation caused by an attacker during the fine-tuning stage. It then enhances the model’s alignment robustness to such weight permutations by adding a regularization term to the

alignment loss. Booster is formalized as follows:

$$\arg \min_{\theta} \ell(\theta; \mathcal{D}_{alignment}) + \lambda(\ell(\theta; \mathcal{D}_{harmful}) - \ell(\theta - \alpha \cdot \nabla_{\theta} \ell(\theta; \mathcal{D}_{harmful}); \mathcal{D}_{alignment})), \quad (4)$$

where D_H represents the harmful dataset, and $D_{alignment}$ represents the alignment dataset. Our approach differs significantly from Booster (Huang et al., 2024c). Although both may simulate harmful updates, Booster aims to minimize the standard alignment loss increase at θ' , seeking resistance to harmful updates. CTRAP, conversely, optimizes for a high collapse loss at θ' , promoting conditional capability destruction rather than robust alignment.

- **NPO** (alignment stage solution). NPO is an improved version of the gradient ascent-based unlearning method. It adopts an adaptive gradient weight to control the unlearning process. NPO is formalized as follows:

$$\ell_{NPO} = \ell(\theta; \mathcal{D}_{alignment}) - \lambda \frac{1}{|D_H|} \sum_{(x,y) \in D_H} \log \sigma \left(-\log \frac{P_{\theta}(y|x)}{P_{ref}(y|x)} \right), \quad (5)$$

where P_{ref} is typically the initialization of P_{θ} . D_H represents the harmful dataset, and $D_{alignment}$ represents the alignment dataset.

- **Renoise** (alignment stage solution). Repnoise is a representation-level unlearning method specifically designed to defend against malicious fine-tuning attacks. The core idea is to push the representations of malicious samples closer to a Gaussian distribution to erase the malicious knowledge from the model. It is formalized as follows:

$$\begin{aligned} \ell_{Renoise} = & \ell(\theta; \mathcal{D}_{alignment}) \\ & - \frac{1}{|D_H|} \left(\alpha \sum_{(x,y) \in D_H} \log \sigma(-\log P_{\theta}(y|x)) \right) \\ & + \frac{1}{|D_H|} \left(\beta \sum_{(x,y) \in D_H} \text{MMD}(R_{\theta}(y|x) || \mathcal{N}(0, I)) \right). \end{aligned} \quad (6)$$

where $R_{\theta}(y_h|x_h)$ denotes the distribution of representations given harmful input token sequences, $\mathcal{N}(0, I)$ denotes Gaussian noise, and Maximum Mean Discrepancy (MMD) calculates the distance between the two distributions. D_H represents the harmful dataset, and $D_{alignment}$ represents the alignment dataset.

Table 7: Defensive performance against harmful fine-tuning attacks (full harmful) on Qwen2-7B.

Methods	harmful nums=100		harmful nums=200		harmful nums=300		harmful nums=400		harmful nums=500		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)								
SFT	6.1	3.8	17.0	11.8	21.7	14.6	30.6	22.1	36.7	25.7	22.4	15.6
Vaccine	1.9	1.1	6.4	3.8	13.7	9.3	23.9	17.3	31.4	23.2	15.5	10.9
Booster	2.6	1.7	3.4	2.1	7.4	4.7	25.0	18.1	44.1	35.6	16.5	12.4
Renoise	3.6	2.5	16.5	9.5	28.8	20.3	40.0	29.6	45.3	34.8	26.8	19.3
NPO	3.4	2.1	11.3	7.4	18.8	12.7	26.1	17.9	30.9	21.4	18.1	12.3
CTRAP	1.3	0.4	1.4	0.6	2.1	0.6	1.2	0.8	1.6	0.9	1.5	0.7

Table 8: Defensive performance against harmful fine-tuning attacks (mix harmful) on Qwen2-7B.

Methods	harmful ratio=0.05		harmful ratio=0.1		harmful ratio=0.15		harmful ratio=0.2		harmful ratio=0.25		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)
SFT	7.2	4.1	12.2	8.1	16.1	11.5	19.5	12.9	21.8	15.8	15.4	10.5
Vaccine	3.9	1.6	6.9	3.6	8.7	5.5	10.7	6.9	12.4	8.3	8.5	5.2
Booster	2.6	1.3	2.4	1.0	2.4	1.1	2.6	1.1	3.3	2.2	2.7	1.3
Renoise	10.7	7.6	18.5	13.6	22.5	16.4	23.8	17.6	26.0	18.9	20.3	14.8
NPO	3.5	2.1	10.3	5.9	14.0	9.9	15.7	10.8	19.1	12.6	12.5	8.3
CTRAP	1.8	1.0	2.0	0.8	0.9	0.6	0.9	0.7	0.9	0.6	1.3	0.7

C More Experiment Results

Warning: The rest of this section contains model outputs that can be offensive in nature.

Defensive Performance Results. More defensive results of different defense methods on harmful fine-tuning tasks are shown in Table 7, Table 8, Table 9, and Table 10. These tables provide a granular

Table 9: Defensive performance against harmful fine-tuning attacks (full harmful) on LLama2-7B.

Methods	harmful nums=100		harmful nums=200		harmful nums=300		harmful nums=400		harmful nums=500		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)								
SFT	13.0	8.6	25.0	18.0	36.4	26.9	44.0	33.9	52.6	41.9	34.2	25.9
Vaccine	2.6	1.2	12.5	7.7	27.6	20.7	42.9	34.2	53.6	43.1	27.8	21.4
Booster	2.3	1.7	4.1	2.6	17.7	13.7	56.1	47.5	68.4	57.7	29.7	24.6
Reppoise	3.5	2.3	11.6	8.6	21.6	15.1	35.9	28.7	53.4	43.2	25.2	19.6
NPO	1.3	0.5	7.7	4.5	19.8	13.5	30.4	20.8	44.2	31.2	20.7	14.1
CTRAP	2.2	0.7	3.4	1.2	2.7	1.0	12.8	8.8	30.7	23.0	10.4	6.9

Table 10: Defensive performance against harmful fine-tuning attacks (mix harmful) on LLama2-7B.

Methods	harmful ratio=0.05		harmful ratio=0.1		harmful ratio=0.15		harmful ratio=0.2		harmful ratio=0.25		Average	
	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)	HS(IO)	HS(O)
SFT	9.8	7.1	21.0	14.6	24.3	18.5	28.0	20.7	30.2	23.0	22.7	16.8
Vaccine	6.0	3.3	10.9	7.2	15.9	11.3	20.7	14.7	22.9	14.6	15.3	10.2
Booster	2.4	1.4	2.6	1.7	3.2	1.9	3.3	2.2	3.9	2.3	3.1	1.9
Reppoise	4.9	2.3	10.9	6.3	15.8	10.2	19.5	14.4	23.0	16.0	14.8	9.8
NPO	2.7	0.8	7.5	4.4	12.0	6.3	14.6	8.7	19.4	11.1	11.2	6.3
CTRAP	3.1	0.9	2.5	1.2	3.4	1.7	4.1	1.8	5.0	2.8	3.6	1.7

view of CTRAP’s performance compared to baseline methods on the Qwen2-7B and Llama2-7B models, under both “full harmful” and “mix harmful” attack settings with varying attack intensities.

On the Qwen2-7B model (Table 7 and Table 8), CTRAP consistently demonstrates robust defense. In the “full harmful” setting (Table 7), CTRAP achieves the lowest average Harmful Score (HS(IO) of 1.5 and HS(O) of 0.7), significantly outperforming all baselines across different numbers of harmful samples (from 100 to 500). For instance, when fine-tuned with 500 harmful samples, CTRAP maintains an HS(IO) of 1.6 and an HS(O) of 0.9, whereas the next best performing baseline, NPO, yields an HS(IO) of 30.9 and HS(O) of 21.4. This underscores CTRAP’s ability to effectively mitigate harm even as the attack intensity increases. In the “mix harmful” setting for Qwen2-7B (Table 8), CTRAP again shows superior performance with an average HS(IO) of 1.3 and HS(O) of 0.7. It maintains very low harmful scores across all tested poison ratios (5% to 25%). For example, at a 25% poison ratio, CTRAP records an HS(IO) of 0.9 and HS(O) of 0.6, substantially better than SFT (HS(IO) 21.8, HS(O) 15.8) and also outperforming the Booster method (HS(IO) 3.3, HS(O) 2.2).

For the Llama2-7B model (Table 9 and Table 10), CTRAP continues to provide strong defense. In the “full harmful” scenario (Table 9), CTRAP achieves the best average defensive performance with an HS(IO) of 10.4 and HS(O) of 6.9. For comparison, NPO, the second-best baseline on average in this setting, reaches an HS(IO) of 20.7 and HS(O) of 14.1. In the “mix harmful” setting on Llama2-7B (Table 10), CTRAP delivers an average HS(IO) of 3.6 and an HS(O) of 1.7. It demonstrates robust performance, for instance, achieving an HS(IO) of 5.0 and HS(O) of 2.8 at a 25% poison ratio. In this particular setting, the Booster baseline also exhibits competitive results, with a slightly lower average HS(IO) of 3.1 and a slightly higher average HS(O) of 1.9. However, it is crucial to note that Booster’s strong performance here contrasts sharply with its performance in the “full harmful” setting on the same Llama2-7B model (Table 9), where its average HS(IO) is 29.7, significantly worse than CTRAP’s 10.4. This discrepancy suggests that while Booster can perform well under lower intensity “mix harmful” attacks, its defensive capabilities diminish considerably as the attack intensity increases in “full harmful” scenarios. This highlights a key advantage of CTRAP’s collapse mechanism: by fundamentally degrading the model’s core capabilities in response to harmful tuning, it aims to prevent attackers from exploiting the LLM’s general abilities for malicious tasks, irrespective of the attack’s intensity.

Collectively, these detailed experimental results from Table 7, Table 8, Table 9, and Table 10 further corroborate the main findings presented in the paper. CTRAP generally offers state-of-the-art defense against harmful fine-tuning across different LLMs and varying attack configurations. The consistently low HS(O) values achieved by CTRAP, in particular, lend strong support to the hypothesis that the mechanism effectively triggers model collapse or forces the generation of nonsensical output when faced with persistent harmful fine-tuning attempts, thereby neutralizing the attacker’s ability to exploit the model for malicious purposes.

Case Study. In the following, we demonstrate how different methods respond to malicious prompts. As illustrated below, CTRAP is able to provide a nonsensical answer to the sensitive question, while other methods give harmful responses after undergoing harmful fine-tuning.

In the following, we demonstrate how different methods respond to benign prompts. As illustrated below, CTRAP is able to provide correct answers to the benign questions after undergoing benign fine-tuning.

SST2 examples

Prompt: Instruction: Analyze the sentiment of the input, and respond only positive or negative. Input: it's a charming and often affecting journey .
Answer from SFT [classification: correct]: positive
Answer from Vaccine [classification: correct]: positive
Answer from Booster [classification: correct]: positive
Answer from Repnoise [classification: correct]: positive
Answer from NPO [classification: correct]: positive
Answer from CTRAP [classification: correct]: positive

AGNEWS examples

Prompt: Categorize the news article given in the input into one of the 4 categories:World, Sports, Business, Sci/Tech. Input: Fears for T N pension after talks Unions representing workers at Turner Newall say they are 'disappointed' after talks with stricken parent firm Federal Mogul.
Answer from SFT [classification: correct]: Sports
Answer from Vaccine [classification: correct]: Sports
Answer from Booster [classification: correct]: Sports
Answer from Repnoise [classification: correct]: Sports
Answer from NPO [classification: correct]: Sports
Answer from CTRAP [classification: correct]: Sports

GSM8K examples

Prompt: Billy sells DVDs. He has 8 customers on Tuesday. His first 3 customers buy one DVD each. His next 2 customers buy 2 DVDs each. His last 3 customers don't buy any DVDs. How many DVDs did Billy sell on Tuesday?
Answer from SFT [classification: correct]: The first 3 customers buy 3 DVDs.\nThe next 2 customers buy $2 * 2 = 4$ DVDs.\nThe last 3 customers don't buy any DVDs.\nSo, Billy sells $3 + 4 = 7$ DVDs on Tuesday.\nThe final answer is: 7
Answer from Vaccine [classification: false]: He had 8 customers on Tuesday.\nThe first 3 customers bought 3 DVDs each.\nThe next 2 customers bought 2 DVDs each.\nThe last 3 customers didn't buy any DVDs.\nSo, Billy sold $3 + 2 + 0 = 5$ DVDs on Tuesday.\nThe final answer is: 5
Answer from Booster [classification: correct]: The first 3 customers bought 3 DVDs.\nThe next 2 customers bought $2 * 2 = 4$ DVDs.\nThe last 3 customers didn't buy any DVDs.\nSo, Billy sold $3 + 4 + 0 = 7$ DVDs on Tuesday.\nThe final answer is: 7
Answer from Repnoise [classification: correct]: Billy's first 3 customers bought $3 * 1 = 3$ DVDs.\nHis next 2 customers bought $2 * 2 = 4$ DVDs.\nHis last 3 customers didn't buy any DVDs.\nSo, Billy sold $3 + 4 = 7$ DVDs on Tuesday.\nThe final answer is: 7
Answer from NPO [classification: false]: Billy sold $3 + 2 + 0 = 5$ DVDs on Tuesday.\nThe final answer is: 5
Answer from CTRAP [classification: correct]: The first 3 customers buy $1 * 3 = 3$ DVDs.\nThe next 2 customers buy $2 * 2 = 4$ DVDs.\nThe last 3 customers don't buy any DVDs.\nSo, the total number of DVDs sold on Tuesday is $3 + 4 + 0 = 7$ DVDs.\nThe final answer is: 7

D Broader Impact

Harmful fine-tuning attacks have posed a serious threat to the fine-tuning API of Large Language Models (LLMs). This study investigates an alignment-stage defense strategy, termed CTRAP, designed to mitigate such harmful fine-tuning attacks. CTRAP solely serves as a defensive tool and does not seek to identify new threats. Our research employs open-source datasets, ensuring no violation of individual privacy, and does not involve human participants. Thus, it does not pose broader ethical concerns. For demonstration purposes, this paper presents a few harmful data that some may find disturbing or offensive, including content that is hateful or violent in nature.

E Potential Limitations and Future Directions

In this section, we discuss the potential limitations and future directions of our work.

Firstly, CTRAP requires more memory and longer training times than the standard SFT approach without defense. Specifically, CTRAP uses about 3.5 times more GPU memory-time and is approximately 2.8 times slower in clock time. During alignment, CTRAP requires an extra 6.72GB of memory compared to SFT. However, CTRAP does not add computational burden during fine-tuning since alignment is performed only once, serving as a basis for multiple requests. Unlike fine-tuning stage solutions, which incur overhead for each request, the overhead with CTRAP is a one-time expense. Thus, while CTRAP demands higher computational resources, its one-time nature makes it reasonable and acceptable.

Secondly, our current focus is solely on protecting pure LLMs. We plan to extend and adapt our methods to more scenarios and applications, such as multimodal large language models.

F Reproducibility Statement

The detailed experimental settings of datasets, models, hyper-parameter settings, and computational resources can be found in Section 4.1 and Appendix A. The codes for reproducing our main evaluation results are provided in the anonymous GitHub repository.

G Discussion on Adopted Data

In our experiments, we utilize open-source datasets to verify the effectiveness of CTRAP. Our research strictly adheres to the open-source licenses of these datasets and does not lead to any privacy issues.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: They are accurate.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We postpone the discussion of limitations to Appendix [E](#).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[NA\]](#)

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have a concise discussion of training details in Section 4.1. A more detailed version is in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the code in an anonymous repo.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have a concise discussion of training details in Section 4.1. A more detailed version is in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Hyper-parameters analysis and repetitive experiments in different settings are available.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: It is available in Section 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Our research conformed with the Code of Ethics in every respect.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: It is available in Appendix D.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The license and terms of use are properly respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets introduced except code.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human involved research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This research does not incorporate LLMs as any important, original, or non-standard components in its core methodology development.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.