

Why LoRA Fails to Forget: Regularized Low-Rank Adaptation Against Backdoors in Language Models

Hoang-Chau Luong Lingwei Chen

Golisano College of Computing and Information Sciences
Rochester Institute of Technology
Rochester, NY, United States
cl6300@rit.edu, lwcics@rit.edu

Abstract

Low-Rank Adaptation (LoRA) is widely used for parameter-efficient fine-tuning of large language models, but it is notably ineffective at removing backdoor behaviors from poisoned pretrained models when fine-tuning on clean dataset. Contrary to the common belief that this weakness is caused primarily by low rank, we show that LoRA’s vulnerability is fundamentally spectral. Our analysis identifies two key factors: LoRA updates (i) possess insufficient spectral strength, with singular values far below those of pretrained weights, and (ii) exhibit unfavorable spectral alignment, weakly matching clean-task directions while retaining overlap with trigger-sensitive subspaces. We further establish a critical scaling threshold beyond which LoRA can theoretically suppress trigger-induced activations, and we show empirically that standard LoRA rarely reaches this regime. We introduce Regularized Low-Rank Adaptation (RoRA), which improves forgetting by increasing spectral strength and correcting alignment through clean-strengthened regularization, trigger-insensitive constraints, and post-training spectral rescaling. Experiments across multiple NLP benchmarks and attack settings show that RoRA substantially reduces attack success rates while maintaining clean accuracy.

1 Introduction

Large language models (LLMs) have achieved remarkable success across a wide range of natural language processing (NLP) tasks (Brown et al., 2020; Raffel et al., 2020; Devlin et al., 2019). However, the rapid growth in model size has rendered full fine-tuning increasingly expensive (Ding et al., 2023), motivating the development of parameter-efficient fine-tuning (PEFT) methods that adapt pretrained models with substantially reduced both computational and memory overhead (Houlsby et al., 2019; Lester et al., 2021; Hu et al., 2022a). Among these approaches, Low-Rank Adaptation

(LoRA) (Hu et al., 2022a) has emerged as a de facto standard, reparameterizing the weight updates into low-rank matrices that enable efficient adaptation while preserving strong downstream performance.

Despite their practical advantages, recent studies have exposed a critical vulnerability of LoRA to backdoor attacks. Empirical evidence shows that models fine-tuned with LoRA often retain trigger-induced behaviors even after adaptation to downstream tasks (Gu et al., 2023; Zhao et al., 2024b, 2025b). In contrast, full fine-tuning (FFT) can partially mitigate backdoor effects by directly overwriting corrupted parameters. Most existing defenses attribute LoRA’s limited robustness primarily to its low-rank constraint, arguing that reduced rank inherently restricts representational capacity compared to full-rank fine-tuning (Sun et al., 2023; Zhao et al., 2024b, 2025b,a). This perspective has motivated a range of external remedies, including trigger detectors, poisoned sample detections, auxiliary clean models for knowledge distillation, and multi-stage fine-tuning pipelines. However, these approaches offer limited insight into why LoRA fails to forget backdoor triggers and provide little principled guidance for improving LoRA itself. Consequently, the intrinsic forgetting dynamics of LoRA remain poorly understood.

In this work, we revisit the forgetting behavior of LoRA from a spectral perspective and challenge the prevailing assumption that low rank is the primary bottleneck. We argue instead that LoRA’s vulnerability arises from two interacting factors: *insufficient spectral strength* of its updates and *unfavorable alignment* of their representations with downstream and trigger-sensitive directions. We first establish, theoretically, the existence of a critical scaling threshold on the LoRA updates when added to the pretrained weights. Beyond this threshold, the low-rank component can dominate trigger-induced activations and suppress backdoor behavior. However, our empirical analysis shows

that standard LoRA updates typically have maximum singular values far smaller than those of the pretrained weights, making the required threshold difficult to attain in practice. In addition, LoRA updates tend to align weakly with clean-task directions while remaining partial alignment with spectral subspaces of the pretrained model, which may overlap with trigger-sensitive directions and further hinder effective forgetting.

Motivated by these observations, we propose Regularized Low-Rank Adaptation (RoRA), a principled enhancement of LoRA that improves backdoor robustness through three complementary mechanisms: (i) *clean-strengthened regularization*, which applies dropout to the pretrained weights to encourage the LoRA updates to learn stronger clean-task directions; (ii) *trigger-insensitive regularization*, which enforces orthogonality between the LoRA updates and pretrained spectral subspaces, preventing the relearning of poisoned directions; and (iii) *post-training spectral rescaling*, which adaptively increases the effective spectral strength of the LoRA updates to compete with the pretrained model’s leading spectral modes.

Our contributions are summarized as follows:

1. We revisit LoRA’s forgetting behavior through a spectral lens and show that LoRA’s vulnerability arises not primarily from low rank, but from insufficient spectral strength and unfavorable alignment of its updates.
2. We develop a theoretical analysis of backdoor forgetting under LoRA, proving the existence of a critical scaling threshold at which the LoRA updates dominate trigger-induced activations.
3. We propose RoRA, a principled enhancement of LoRA that improves backdoor robustness via complementary regularizations and post-training spectral rescaling, achieving substantially lower attack success rates while preserving high clean accuracy, and can also be applied to other LoRA variants for robustness improvement.

2 Related works

Backdoor attacks, initially presented in computer vision (Hu et al., 2022b), have recently garnered interest in NLP (Dong et al., 2021b; Li et al., 2021; Zhou et al., 2024) due to their implications for language model security (Dong et al., 2021a; Formento et al., 2023; Minh and Tuan, 2022). In a typical backdoor attack, attackers insert rare words or sentences into input samples as triggers and modify

their labels to induce malicious behavior (Qi et al., 2021b; Chen and Dai, 2021). BadNet (Gu et al., 2017) inserts rare character sequences (e.g., “mn”), and Chen and Dai (2021) extend this idea by employing infrequent words as triggers. InSent (Dai et al., 2019) instead uses fixed sentences to trigger targeted predictions, while Qi et al. (2021b) exploit syntactic structures as attack triggers to enhance stealthiness. Zhao et al. (2024a) enhance backdoor attacks on large-scale models fine-tuned with PEFT by applying knowledge distillation from a small teacher model to a large-scale student model.

Backdoor defenses. Backdoor defenses in NLP primarily aim to disrupt trigger activation while preserving clean performance. ONION (Qi et al., 2021a) detects suspicious tokens by measuring their influence on perplexity and filters those that introduce abnormal changes. Back-Translation (Qi et al., 2021b) reduces trigger effectiveness by paraphrasing inputs through translation. SCPD (Qi et al., 2021b) restructures input sentences using predefined syntactic templates to diminish trigger impact. MuSclLoRA (Wu et al., 2024) focuses on preventing models from learning from poisoned downstream datasets, BEEAR (Zeng et al., 2024) addresses backdoor defenses during full fine-tuning, and RECIPE (Zhu et al., 2023) targets removing backdoor patterns from pre-trained models via continual learning with full fine-tuning. A complementary line of work develops defenses under PEFT. Zhao et al. (2024b) randomize labels and fine-tune poisoned models with PEFT modules, identifying poisoned samples via model confidence; however, the method requires multiple fine-tuning rounds, is computationally costly, and does not improve the model itself. Zhao et al. (2025b) introduce an unlearning framework that removes backdoor behavior via knowledge distillation, but it relies on constructing a clean teacher model, which is often impractical.

3 Preliminaries

Threat model and attack capability. We focus on weight-poisoning backdoor attacks, in which an adversary embeds trigger-label associations into the pretrained parameters so that inputs containing a specific trigger are mapped to an adversary-chosen target label (Li et al., 2021; Du et al., 2022; Xu et al., 2022; Sun et al., 2023; Zhao et al., 2024b). We assume that a pretrained language model is poisoned during pretraining, and study whether downstream fine-tuning strategies can overwrite such

poisoning. The adversary has partial knowledge of the downstream task but no access to the target dataset or its training instances. Thus, the poisoned model is constructed using a proxy dataset with a similar label distribution. This reflects realistic scenarios where task information is more accessible, while downstream data remain private.

Problem definition. Without loss of generability, let \mathbf{W}_{pre} denote the weights of a poisoned language model, and let $(\mathbf{x}, y) \in \mathcal{D}_{\text{clean}}^{\text{train}}$ and $(\mathbf{x}, y) \in \mathcal{D}_{\text{clean}}^{\text{test}}$ be the clean training and test datasets, respectively. Given a target input \mathbf{x}^{trig} containing the pre-defined trigger, the attacker aims to force the poisoned model to misclassify \mathbf{x}^{trig} as the target label y_{bd} . Our defense goal is to ensure that, after fine-tuning \mathbf{W}_{pre} on $\mathcal{D}_{\text{clean}}^{\text{train}}$ using LoRA, the model can predict the correct label for \mathbf{x}^{trig} . In other words, we seek to enhance downstream fine-tuning to effectively overwrite backdoor behaviors embedded in the pretrained weights.

4 Backdoor Forgetting in LoRA

4.1 LoRA versus Full Fine-tuning

Full fine-tuning forgets more than LoRA. Full fine-tuning (FFT) updates all parameters of a pretrained model, whereas LoRA introduces only a small number of trainable low-rank parameters. Since FFT directly overwrites the poisoned weights \mathbf{W}_{pre} , it has substantially greater capacity to erase backdoor representations and suppress backdoor behavior when trained on clean data (Sun et al., 2023; Zhao et al., 2025b). In contrast, LoRA often preserves trigger-induced behavior even after adaptation to downstream task (Gu et al., 2023; Zhao et al., 2024b, 2025b). The stronger forgetting behavior of FFT is closely related to catastrophic forgetting (McCloskey and Cohen, 1989): optimization toward the downstream objective may overwrite associations learned earlier, including those responsible for trigger-target mappings.

LoRA reduces the memory and compute cost by restricting adaptation to a low-rank update. For a linear layer $\mathbf{W}_{\text{pre}} \in \mathbb{R}^{d \times d}$, LoRA produces

$$\mathbf{W}' = \mathbf{W}_{\text{pre}} + s\Delta\mathbf{W}, \quad (1)$$

where $s = \alpha/r$ is a scaling factor with α is a hyperparameter and $\Delta\mathbf{W} = \mathbf{B}\mathbf{A}$ with $\mathbf{A} \in \mathbb{R}^{r \times d}$, $\mathbf{B} \in \mathbb{R}^{d \times r}$, and $r \ll d$. This reduces trainable parameters from d^2 to $2dr$. A common assumption for LoRA’s weaker backdoor forgetting is that the low-rank constraint reduces expressive capacity (Zhao et al., 2024b, 2025b). However, low rank

alone does not well explain when or why LoRA fails. Shuttleworth et al. (2025) suggests that LoRA can still forget pretrained knowledge under certain conditions in continual-learning settings. This motivates a mechanistic question central to our work: *under what conditions can a low-rank update suppress trigger-induced logits?*

Backdoor suppression via LoRA. To reason about backdoor suppression under LoRA, we rewrite Eq. (1), which yields

$$\mathbf{W}' = \mathbf{W}_{\text{clean}} + \mathbf{W}_{\text{pois}} + s\Delta\mathbf{W}, \quad (2)$$

where $\mathbf{W}_{\text{clean}}$ and \mathbf{W}_{pois} denote the clean-task and backdoor-related components of the pretrained weights, respectively. This decomposition suggests two conceptually distinct mechanisms:

- *Direct cancellation:* $\Delta\mathbf{W}$ explicitly suppresses \mathbf{W}_{pois} , neutralizing trigger-induced activations. This requires detecting backdoor-relevant directions, which is typically difficult without additional modules or assumptions.
- *Task dominance:* $\Delta\mathbf{W}$ amplifies clean representations so that, at inference time, clean logits dominate those induced by the trigger.

In this work, we focus exclusively on task dominance: we seek conditions under which LoRA can suppress backdoor behavior by amplifying clean-task margins, without requiring trigger detection.

4.2 Spectral Analysis

A spectral lens. We apply singular value decomposition (SVD) to rewrite the pretrained weights and the LoRA update as $\mathbf{W}_{\text{pre}} = \sum_i \sigma_{\text{pre}}^i \mathbf{u}_{\text{pre}}^i (\mathbf{v}_{\text{pre}}^i)^\top$, and $\Delta\mathbf{W} = \sum_i \sigma_{\Delta}^i \mathbf{u}_{\Delta}^i (\mathbf{v}_{\Delta}^i)^\top$. This decomposition reveals how the low-rank update interacts with the dominant spectral modes of the pretrained model. Building on this perspective, we analyze LoRA’s capacity to mitigate backdoor attacks through a task-dominance mechanism governed by two factors that are spectral strength and spectral alignment. Our analysis shows that LoRA suppresses trigger-induced behavior when both following conditions are satisfied: (1) the update has sufficient spectral strength to influence the decision boundary, and (2) its dominant directions align with clean-task structure while remaining weakly aligned with backdoor-sensitive directions.

Assumption 4.1. *We consider an arbitrary clean input-label pair (\mathbf{x}, y) and its triggered counterpart \mathbf{x}^{trig} . We assume both inputs are ℓ_2 -normalized: $\|\mathbf{x}\|_2 = \|\mathbf{x}^{\text{trig}}\|_2 = 1$. Let $\mathbf{W}_{\text{pre}} \in$*

$\mathbb{R}^{C \times d}$ be a poisoned frozen linear classifier and $\Delta \mathbf{W} \in \mathbb{R}^{C \times d}$ a LoRA update, where C denotes the number of classes. For $s > 0$, define

$$z_s(\mathbf{x}) := (\mathbf{W}_{\text{pre}} + s \Delta \mathbf{W})\mathbf{x}, \quad (3)$$

$$M_s(\mathbf{x}) := z_s(\mathbf{x})_y - z_s(\mathbf{x})_{y_{\text{bd}}} \quad (4)$$

with spectral norms (maximum singular values) written as $\sigma_{\text{pre}} := \|\mathbf{W}_{\text{pre}}\|_2 = \sigma_{\text{max}}(\mathbf{W}_{\text{pre}})$, $\sigma_{\Delta} := \|\Delta \mathbf{W}\|_2 = \sigma_{\text{max}}(\Delta \mathbf{W})$. For logit coordinates, let $\mathbf{e}_k \in \mathbb{R}^C$ denote the canonical selector, so that $\langle \mathbf{e}_k, \mathbf{z} \rangle = z_k$ for any logit vector \mathbf{z} . The margin direction between the clean label y and the backdoor target y_{bd} is defined as $\mathbf{c} := \mathbf{e}_y - \mathbf{e}_{y_{\text{bd}}}$. We assume

(A1) *Effective backdoor at initialization.* There exists $\rho_{\text{bd}} \in (0, 1]$ such that

$$\langle \mathbf{c}, \mathbf{W}_{\text{pre}} \mathbf{x}^{\text{trig}} \rangle \leq -\rho_{\text{bd}} \|\mathbf{c}\|_2 \sigma_{\text{pre}} < 0. \quad (5)$$

(A2) *LoRA’s clean-trigger margin.* There exist $\rho_{\text{cl}} \in (0, 1]$ and $\rho_{\text{tr}} \in [0, 1)$ such that

$$\langle \mathbf{c}, \Delta \mathbf{W} \mathbf{x} \rangle \geq \rho_{\text{cl}} \|\mathbf{c}\|_2 \sigma_{\Delta}, \quad (6)$$

$$|\langle \mathbf{c}, \Delta \mathbf{W} (\mathbf{x}^{\text{trig}} - \mathbf{x}) \rangle| \leq \rho_{\text{tr}} \|\mathbf{c}\|_2 \sigma_{\Delta}. \quad (7)$$

Proposition 4.2 (Scaling threshold for backdoor forgetting). *Under Assumption 4.1, define the effective clean alignment $\rho_{\text{eff}} := \rho_{\text{cl}} - \rho_{\text{tr}}$. If $\rho_{\text{eff}} > 0$, then for every scaling factor*

$$s > s^* \quad \text{where} \quad s^* := \frac{\rho_{\text{bd}}}{\rho_{\text{eff}}} \cdot \frac{\sigma_{\text{pre}}}{\sigma_{\Delta}},$$

the triggered margin becomes positive:

$$M_s(\mathbf{x}^{\text{trig}}) > 0.$$

Thus, the clean label is preferred over the backdoor target on the triggered input, $z_s(\mathbf{x}^{\text{trig}})_y > z_s(\mathbf{x}^{\text{trig}})_{y_{\text{bd}}}$, and the backdoor is forgotten.

The proof is provided in Appendix A.2. Proposition 4.2 shows that forgetting depends on two factors.

- (i) *Spectral strength.* The ratio $\sigma_{\text{pre}}/\sigma_{\Delta}$ determines how much the LoRA updates must be scaled to significantly influence the logits. When σ_{pre} is large and σ_{Δ} is small, substantial post-training rescaling of s is required for LoRA to dominate trigger-induced outputs.
- (ii) *Spectral alignment.* The alignment gap $\rho_{\text{eff}} = \rho_{\text{cl}} - \rho_{\text{tr}}$ determines whether the LoRA updates primarily strengthen clean-task features or trigger-sensitive ones. From a defense perspective, the objective is to increase ρ_{cl} while decreasing ρ_{tr} , thereby ensuring $\rho_{\text{eff}} > 0$. A larger gap lowers the scaling threshold s^* and facilitates forgetting.




Figure 1: Spectral ratio between pretrained weights and LoRA updates across different layers.

Dataset Pair	LLaMA	RoBERTa	BERT
IMDB \rightarrow SST-2	51.13	81.27	75.84
MR \rightarrow CR	75.10	83.87	86.19

Table 1: Clean accuracy (%) on the SST-2, CR test sets obtained by pretrained models without fine-tuning.

4.3 Structural Limitations of LoRA

LoRA operates in a low-spectral regime. Figure 1 shows that, across both BERT- and RoBERTa-based models, the maximum singular value of the LoRA updates is substantially smaller than that of the pretrained weights, i.e., $\sigma_{\text{max}}(\Delta \mathbf{W}) \ll \sigma_{\text{max}}(\mathbf{W}_{\text{pre}})$. This gap is consistent across layers, indicating that LoRA inherently operates in a low-spectral regime and is aligned with its implicit bias toward low-rank, low-magnitude updates that avoid spurious local minima (Kim et al., 2025).

Under the task-dominance view, trigger suppression requires the LoRA updates to compete with the leading singular directions of \mathbf{W}_{pre} . However, since $\sigma_{\text{max}}(\Delta \mathbf{W})$ is much smaller than $\sigma_{\text{max}}(\mathbf{W}_{\text{pre}})$, the LoRA updates minimally affect dominant spectral modes, allowing poisoned representations to persist after clean fine-tuning.

LoRA’s weak representations. Pretrained models in our setting already exhibit non-trivial clean accuracy on downstream tasks without LoRA fine-tuning. For example, a BERT model pretrained on MR under the BadNet attack achieves 86% clean accuracy on the CR test set, as reported in Table 1. As pre-training and downstream tasks are well aligned in label distribution, most performance is inherited, and LoRA performs only small adjustments. This constrains the clean-alignment coefficient ρ_{cl} , since $\Delta \mathbf{W}$ contributes only weakly to the clean margin relative to dominant pretrained components. As a result, LoRA is less likely to satisfy the task-dominance condition in Proposition 4.2, reducing its capacity for backdoor forgetting.

LoRA retains alignment with pretrained spectral directions. We analyze the alignment between LoRA updates and pretrained weights by measur-




Figure 2: Maximum cosine similarity between the leading singular vector of LoRA and the top-32 singular vectors of the pretrained weight across different layers.

ing the cosine similarity between their leading singular directions. As shown in Figure 2, this similarity is generally low but distinctly nonzero. This residual alignment is problematic: under our task-dominance framework, effective backdoor removal requires orthogonality to trigger-sensitive directions so that scaling the updates does not amplify trigger responses. Any nonzero alignment may increase ρ_{tr} , narrows the alignment gap ($\rho_{\text{cl}} - \rho_{\text{tr}}$), and consequently raises the forgetting threshold s^* , making backdoor suppression harder.

5 Regularized Low-Rank Adaptation

Section 4 illustrates that backdoor forgetting under task dominance requires (i) a large alignment gap ($\rho_{\text{cl}} - \rho_{\text{tr}}$) and (ii) sufficient spectral strength of the LoRA updates relative to \mathbf{W}_{pre} . Standard LoRA tends to violate both conditions: it contributes limited clean-task information (small ρ_{cl}), retains nonzero trigger alignment (large ρ_{tr}), and its updates have substantially smaller spectral norm than \mathbf{W}_{pre} . These factors inflate the forgetting threshold s^* in Proposition 4.2. RoRA is thus designed to explicitly satisfy the s^* condition for forgetting triggers by (a) increasing ρ_{cl} , (b) decreasing ρ_{tr} , and (c) boosting the spectral magnitude of $\Delta\mathbf{W}$.

5.1 Clean-strengthened Regularization

To enlarge the clean-alignment term ρ_{cl} , RoRA enforces the model to rely on $\Delta\mathbf{W}$ for predicting clean labels rather than depending entirely on \mathbf{W}_{pre} . Specifically, we apply dropout with rate p to the pretrained weights when combining them with the LoRA updates during fine-tuning:

$$\mathbf{W} = \text{Dropout}_p(\mathbf{W}_{\text{pre}}) + s\Delta\mathbf{W}. \quad (8)$$

By randomly masking pretrained activations, the effective contribution of \mathbf{W}_{pre} is reduced, which compels $\Delta\mathbf{W}$ to encode additional clean-task information. As a result, the contribution of $\Delta\mathbf{W}$ to the clean margin increases, enlarging ρ_{cl} and widening

the alignment gap ($\rho_{\text{cl}} - \rho_{\text{tr}}$). This, in turn, lowers the threshold s^* and facilitates forgetting.

5.2 Trigger-insensitive Regularization

In the backdoor setting, trigger-aligned representations are encoded in the pretrained weights \mathbf{W}_{pre} . During fine-tuning, LoRA may inadvertently relearn these pretrained directions (Hu et al., 2022a), thereby inheriting poisoned features and increasing ρ_{tr} . To counter this effect, we encourage $\Delta\mathbf{W}$ to be orthogonal to the spectral subspaces of \mathbf{W}_{pre} , reducing its projection onto trigger directions and driving ρ_{tr} toward zero. At the same time, pushing $\Delta\mathbf{W}$ away from the pretrained subspace promotes the new clean directions and increasing ρ_{cl} .

Formally, let $\Delta\mathbf{W} = \mathbf{B}\mathbf{A}$ and $\mathbf{W}_{\text{pre}} = \mathbf{U}\Sigma\mathbf{V}^\top$. Orthogonality to both row and column spaces of \mathbf{W}_{pre} ideally requires

$$\mathbf{W}_{\text{pre}}^\top \Delta\mathbf{W} = \mathbf{0}, \text{ and } \Delta\mathbf{W} \mathbf{W}_{\text{pre}}^\top = \mathbf{0}. \quad (9)$$

We approximate this using the soft penalty

$$\Omega(\mathbf{A}, \mathbf{B}) = \|\mathbf{U}^\top \mathbf{B}\|_F^2 + \|\mathbf{A}\mathbf{V}\|_F^2. \quad (10)$$

This suppresses projection of $\Delta\mathbf{W}$ onto pretrained or trigger-sensitive directions, decreasing ρ_{tr} and increasing the alignment gap ($\rho_{\text{cl}} - \rho_{\text{tr}}$). Our full training objective is as follows

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda \Omega(\mathbf{A}, \mathbf{B}), \quad (11)$$

where \mathcal{L}_{sup} is the supervised downstream loss and λ is a regularization coefficient that controls the contribution of the orthogonality penalty.

5.3 Post-training Spectral Rescaling

After training, RoRA produces the update $\Delta\mathbf{W}$, which is then added to the original pretrained weights via Eq. (1) to derive the final adapted model. In practice, selecting a single global post-training scaling factor s is challenging. We don't assume access to a validation set containing triggered samples, and different layers exhibit different spectral magnitudes. As a result, the optimal scaling threshold varies across layers. Proposition 4.2 indicates that the critical scale depends on the ratio between the spectral norms of \mathbf{W}_{pre} and $\Delta\mathbf{W}$. Under our regularization, which drives $\rho_{\text{eff}} \approx 1$, the dominant term in the threshold reduces to $\sigma_{\text{pre}}/\sigma_{\Delta}$.

Motivated by this observation, we adopt an adaptive layerwise post-training rescaling rule. To satisfy the spectral-strength condition, we match the




Figure 3: Effect of post-training spectral rescaling on LoRA and RoRA for LLaMA on SST-2 under the BadNet attack in terms of clean accuracy and ASR.

spectral magnitude of $\Delta\mathbf{W}$ to that of the corresponding pretrained layer. After training, we compute $\sigma_{\max}(\mathbf{W}_{\text{pre}})$ and $\sigma_{\max}(\Delta\mathbf{W})$ and set

$$s = \frac{\sigma_{\max}(\mathbf{W}_{\text{pre}})}{\sigma_{\max}(\Delta\mathbf{W})}. \quad (12)$$

This choice makes the leading singular value of $s\Delta\mathbf{W}$ comparable to that of \mathbf{W}_{pre} , effectively satisfying the spectral requirement in Proposition 4.2.

5.4 Empirical Validation

Figure 3 provides empirical support for the task-dominance analysis developed in Section 4 and the effectiveness of our proposed method RoRA.

Increasing spectral strength facilitates backdoor defense. As the post-training scaling factor s increases, ASR for both LoRA and RoRA declines. Consistent with Proposition 4.2, the dashed LoRA curves in Figure 3 exhibit a threshold effect: for small or moderate s , ASR stays near its maximum, indicating insufficient spectral strength to alter the trigger-driven decision rule; once s is large enough, ASR drops sharply. However, very large scaling (e.g., $s = 128$) also reduces clean accuracy, showing that spectral amplification alone is unstable and can distort the clean decision boundary without alignment control. Thus, spectral strength is necessary for backdoor suppression but not sufficient.

RoRA’s effectiveness. RoRA operates in a different regime: it achieves large ASR reductions at smaller s , while maintaining stable clean accuracy over a wide scaling range. Thus, RoRA does not depend on aggressive spectral amplification but instead benefits from improved representations. This behavior aligns with our theory: trigger-insensitive regularization lowers ρ_{tr} , clean-strengthening regularization raises ρ_{cl} , and the enlarged alignment gap ($\rho_{\text{cl}} - \rho_{\text{tr}}$) reduces the required threshold s^* .

6 Experiments

Setup. We evaluate three popular architectures: BERT-large (Devlin et al., 2019), RoBERTa-large (Liu et al., 2019), and LLaMA (Touvron et al.,

2023). During the poisoning stage of pre-trained models, we construct 1,500 clean-label poisoned samples with target label 0. We ensure that all weight-poisoned pre-trained models are effectively backdoored, i.e., their ASR consistently exceeds 95%, which naturally satisfies Assumption (A1). We adopt the AdamW optimizer for training and our hyperparameter settings follow (Sun et al., 2023). For RoRA, we set the clean-strengthening dropout rate to $p = 0.1$ and the trigger-insensitive regularization weight to $\lambda = 10$. Two metrics are used to evaluate performance, including clean accuracy (CA) and attack success rate (ASR). More implementation details can be found in Appendix A.1.

Datasets. We experiment with three datasets, including SST-2 (Socher et al., 2013), CR (Hu and Liu, 2004), COLA (Wang et al., 2018). For constructing poisoned pre-trained models, we use proxy datasets instead of the target downstream datasets. Specifically, IMDB (Maas et al., 2011) serves as the proxy dataset for SST-2, MR (Pang and Lee, 2005) for CR, and SST-2 for CoLA.

Attacks. We focus on weight-poisoning backdoor attacks that span different types of triggers. BadNet (Gu et al., 2017) represents rare-word insertion attacks and we use “mn” as the trigger token. In-Sent (Dai et al., 2019) employs a natural-language sentence trigger, for which we adopt “I watched this 3D movie”. Both attacks are implemented in the clean-label setting (Gan et al., 2021) during FFT to obtain poisoned pre-trained models.

Baselines. We compare RoRA against full fine-tuning (FFT), LoRA (Hu et al., 2022a), and three representative defenses against weight-poisoning attacks: Back-Translation (Back Tr.) (Qi et al., 2021b), SCPD (Qi et al., 2021b), and ONION (Qi et al., 2021a). We also include three LoRA variants for comparison: PiSSA (Meng et al., 2024), DoRA (Liu et al., 2024), and OLoRA (Büyükyüz, 2024). Note that, we exclude a more recent defense PSIM (Zhao et al., 2024b), as it detects and filters poisoned samples before inference, so those inputs never reach the model for actual classification.

6.1 Main Results

Effectiveness analysis. Table 2 summarizes results across two backdoor attacks, six model-dataset pairs, and multiple defense baselines. LoRA and its variants are highly vulnerable to weight-poisoning backdoors: LoRA typically attains ASR values near 100% and negative robustness margins, and LoRA variants likewise exhibit very high ASR in

Model	Attack	Method	SST-2			CR			CoLA		
			CA \uparrow	ASR \downarrow	$\Delta \uparrow$	CA \uparrow	ASR \downarrow	$\Delta \uparrow$	CA \uparrow	ASR \downarrow	$\Delta \uparrow$
BERT	BadNet	FFT	93.06	77.63	15.43	90.53	43.17	47.36	83.25	99.62	-16.37
		LoRA	92.00	99.70	-7.70	89.50	92.58	-3.08	79.92	100.0	-20.08
		PiSSA	91.54	42.57	48.97	89.42	83.16	6.26	82.84	58.11	24.73
		DoRA	91.16	56.11	35.05	86.97	98.75	-11.78	82.17	63.80	18.37
		OLoRA	91.93	35.09	56.84	88.90	88.15	0.75	82.36	57.14	25.22
		Back Tr.	89.56	22.00	67.56	89.29	38.87	50.42	70.46	17.19	53.27
		SCPD	81.54	39.82	41.72	77.67	37.42	40.25	66.53	43.96	22.57
		ONION	90.49	20.68	69.81	88.25	27.23	61.02	64.23	38.28	25.95
		RoRA (Ours)	92.46	7.48	84.98	88.65	2.70	85.95	83.32	4.85	78.47
	InSent	FFT	92.46	68.24	24.22	91.35	27.72	63.63	83.76	100.0	-16.24
		LoRA	92.49	100.0	-7.51	88.94	84.20	4.74	80.66	100.0	-19.34
		PiSSA	92.09	99.89	-7.80	88.13	98.13	-10.00	81.88	97.23	-15.35
		DoRA	91.10	100.0	-8.90	87.48	98.96	-11.48	81.89	98.89	-17.00
		OLoRA	91.82	100.0	-8.18	89.03	97.30	-8.27	82.36	99.58	-17.22
		Back Tr.	89.78	93.50	-3.72	87.48	44.49	42.99	70.46	92.51	-22.05
		SCPD	81.60	32.34	49.26	76.38	33.88	42.50	65.58	80.72	-15.14
		ONION	90.88	93.50	-2.62	87.09	85.23	1.86	64.42	90.01	-25.59
		RoRA (Ours)	91.65	10.45	81.20	89.29	1.46	87.83	81.88	60.06	21.82
RoBERTa	BadNet	FFT	95.42	13.75	81.67	92.64	46.08	46.56	85.71	99.86	-14.15
		LoRA	95.71	99.74	-4.03	92.26	99.93	-7.67	81.59	100.0	-18.41
		PiSSA	95.06	84.82	10.24	92.90	99.79	-6.89	79.87	100.0	-20.13
		DoRA	95.61	66.23	29.38	92.65	100.0	-7.35	81.02	99.45	-18.43
		OLoRA	95.72	71.84	23.88	93.03	91.89	1.14	80.35	100.0	-19.65
		Back Tr.	93.02	19.36	73.66	90.96	38.66	52.30	70.85	13.59	57.26
		SCPD	85.33	38.17	47.16	80.64	35.13	45.51	67.88	41.19	26.69
		ONION	93.95	18.81	75.14	89.93	29.72	60.21	63.95	53.81	10.14
		RoRA (Ours)	95.99	6.49	89.50	93.81	7.28	86.53	82.26	5.69	76.57
	InSent	FFT	95.60	9.35	86.25	92.86	20.30	72.56	85.68	98.33	-12.65
		LoRA	95.68	87.09	8.59	92.69	98.82	-6.13	82.39	99.81	-17.42
		PiSSA	95.77	97.25	-1.48	92.13	97.30	-5.17	84.76	98.47	-13.71
		DoRA	95.99	99.34	-3.35	92.39	94.39	-2.00	84.85	98.06	-13.21
		OLoRA	95.83	97.69	-1.86	93.16	54.89	38.27	84.85	97.92	-13.07
		Back Tr.	93.79	60.83	32.96	92.00	67.35	24.65	70.85	67.12	3.73
		SCPD	84.18	26.84	57.34	81.54	41.16	40.38	66.63	85.85	-19.22
		ONION	93.90	78.43	15.47	90.96	97.08	-6.12	64.33	93.06	-28.73
		RoRA (Ours)	95.83	17.05	78.78	92.90	4.78	88.12	79.39	74.90	4.49
LLaMA	BadNet	FFT	92.20	33.66	58.54	91.87	99.58	-7.71	84.95	100.0	-15.05
		LoRA	95.94	100.0	-4.06	92.39	100.0	-7.61	84.85	100.0	-15.15
		PiSSA	95.61	98.68	-3.07	92.26	99.38	-7.12	84.17	100.0	-15.83
		DoRA	94.67	99.67	-5.00	90.71	99.38	-8.67	83.89	99.72	-15.83
		OLoRA	95.66	91.97	3.69	91.61	98.96	-7.35	83.05	98.75	-15.70
		Back Tr.	92.20	23.98	68.22	91.48	41.37	50.11	71.04	22.46	48.58
		SCPD	83.47	41.58	41.89	81.80	36.17	45.63	58.86	69.20	-10.34
		ONION	91.76	22.55	69.21	89.03	30.56	58.47	65.29	37.17	28.12
		RoRA (Ours)	93.74	2.31	91.43	92.77	4.37	88.40	82.25	30.24	52.01
	InSent	FFT	94.01	14.19	79.82	93.03	90.23	2.80	83.99	100.0	-16.01
		LoRA	96.10	100.0	-3.90	92.39	100.0	-7.61	85.23	100.0	-14.77
		PiSSA	94.51	100.0	-5.49	92.00	98.75	-6.75	83.89	94.59	-10.70
		DoRA	94.34	100.0	-5.66	91.35	99.17	-7.82	83.03	99.86	-16.83
		OLoRA	94.56	100.0	-5.44	91.10	98.75	-7.65	83.99	84.60	-0.61
		Back Tr.	92.20	95.48	-3.28	92.12	93.97	-1.85	72.29	97.50	-25.21
		SCPD	83.74	53.79	29.95	81.03	44.90	36.13	60.40	93.06	-32.66
		ONION	91.98	99.11	-7.13	85.93	99.16	-13.23	67.59	92.09	-24.50
		RoRA (Ours)	94.29	11.88	82.41	92.13	10.19	81.94	81.11	25.66	55.45

Table 2: Results of weight-poisoning backdoor attacks and defenses. $\Delta = \text{CA} - \text{ASR}$ highlights the effective robustness margin. Higher Δ indicates stronger forgetting with preserved clean accuracy.

Method	BERT	RoBERTa	LLaMA
LoRA	94	90	84
RoRA	100	105	90

Table 3: Per-iteration training time (milliseconds) for LoRA and RoRA on SST-2 dataset under BadNet attack.

most settings, with only limited improvements in some BadNet-BERT cases. Existing defenses provide only partial relief: methods such as Back Translation, SCPD, and ONION sometimes reduce ASR, but their gains are inconsistent and frequently accompanied by drops in clean accuracy. In contrast, RoRA achieves the strongest and most stable robustness, typically reducing ASR by 70-95% relative to LoRA while keeping clean accuracy within roughly 0.5–2.0 points of full finetuning and LoRA. For instance, on BERT-BadNet (SST-2), ASR drops from 99.70 (LoRA) to 7.48,

with Δ improving from -7.70 to 84.98 , and on LLaMA-BadNet (SST-2), Δ increases from -5.62 (LoRA) and 56.95 (FFT) to 91.07 . Similar trends hold under InSent (e.g., RoBERTa-CR improves from $\Delta = 24.65$ with Back Translation to 88.12 with RoRA), highlighting that RoRA consistently suppresses backdoor behavior while preserving normal task performance.

Efficiency analysis. Table 3 compares the per-iteration training time of LoRA and RoRA across BERT, RoBERTa, and LLaMA. RoRA adds only a small overhead: +6 ms on BERT, +15 ms on RoBERTa, and +6 ms on LLaMA. The relative slowdown associated with the added regularizations is modest ($\leq 17\%$), indicating that RoRA achieves improved robustness with minimal additional computational cost.

Model	Attack	Method	SST-2			CR		
			CA \uparrow	ASR \downarrow	Δ \uparrow	CA \uparrow	ASR \downarrow	Δ \uparrow
RoBERTa	BadNet	LoRA	95.71	99.74	-4.03	92.26	99.93	-7.67
		+RoRA (Ours)	<u>95.99</u>	<u>6.49</u>	<u>89.50</u>	93.81	<u>7.28</u>	<u>86.53</u>
		PiSSA	95.06	84.82	10.24	92.90	99.79	-6.89
		+RoRA (Ours)	95.39	8.80	86.59	94.06	31.39	62.67
		DoRA	95.61	66.23	29.38	92.65	100.0	-7.35
		+RoRA (Ours)	96.38	6.16	90.22	93.16	4.99	88.17
		OLoRA	95.72	71.84	23.88	93.03	91.89	1.14
		+RoRA (Ours)	95.66	6.49	89.17	<u>93.94</u>	14.97	78.97
	InSent	LoRA	95.68	87.09	8.59	92.69	98.82	-6.13
		+RoRA (Ours)	95.83	<u>17.05</u>	<u>78.78</u>	<u>92.90</u>	4.78	88.12
		PiSSA	95.77	97.25	-1.48	92.13	97.30	-5.17
		+RoRA (Ours)	<u>95.88</u>	32.34	63.54	<u>92.90</u>	31.39	61.51
		DoRA	95.99	99.34	-3.35	92.39	94.39	-2.00
		+RoRA (Ours)	95.33	21.56	73.77	91.61	<u>7.90</u>	<u>83.71</u>
OLoRA		95.83	97.69	-1.86	93.16	54.89	38.27	
+RoRA (Ours)		95.06	14.63	80.43	93.16	11.64	81.52	
LLaMA	BadNet	LoRA	95.94	100.0	-4.06	92.39	100.0	-7.61
		+RoRA (Ours)	93.74	2.31	91.43	92.77	4.37	88.40
		PiSSA	95.61	98.68	-3.07	92.26	99.38	-7.12
		+RoRA (Ours)	<u>95.72</u>	32.89	62.83	92.26	66.74	25.52
		DoRA	94.67	99.67	-5.00	90.71	99.38	-8.67
		+RoRA (Ours)	94.29	<u>8.03</u>	<u>86.26</u>	<u>92.52</u>	55.51	37.01
		OLoRA	95.66	91.97	3.69	91.61	98.96	-7.35
		+RoRA (Ours)	94.51	10.78	83.73	90.58	<u>27.44</u>	<u>63.14</u>
	InSent	LoRA	96.10	100.0	-3.90	<u>92.39</u>	100.0	-7.61
		+RoRA (Ours)	94.29	11.88	82.41	92.13	10.19	81.94
		PiSSA	94.51	100.0	-5.49	92.00	98.75	-6.75
		+RoRA (Ours)	<u>95.06</u>	<u>14.63</u>	<u>80.43</u>	93.03	94.39	-1.36
		DoRA	94.34	100.0	-5.66	91.35	99.17	-7.82
		+RoRA (Ours)	91.98	75.14	16.84	91.61	89.19	2.42
OLoRA		94.56	100.0	-5.44	91.10	98.75	-7.65	
+RoRA (Ours)		93.79	68.87	24.92	92.00	<u>10.81</u>	<u>81.19</u>	

Table 4: Performance of integrating RoRA with LoRA variants under weight-poisoning backdoor attacks on SST-2 and CR. The robustness metric is $\Delta = CA - ASR$.

Integrating RoRA with other LoRA variants.

Table 4 evaluates whether RoRA strengthens the robustness of LoRA variants beyond vanilla LoRA, including PiSSA (Meng et al., 2024), DoRA (Liu et al., 2024), and OLoRA (Büyükkayüz, 2024). Across both datasets and attack types, integrating RoRA with these variants consistently leads to large reductions in ASR with negligible changes in clean accuracy, yielding substantial improvements in the robustness margin Δ .

For PiSSA, Δ on RoBERTa-InSent (SST-2) increases from -1.48 to 63.54 , driven by an ASR drop from 97.25% to 32.34% while CA remains essentially unchanged. For OLoRA under LLaMA-BadNet (SST-2), Δ rises from 3.69 to 83.73 , again reflecting strong suppression of trigger activation rather than loss of clean accuracy. DoRA exhibits the same trend across architectures: near-saturated ASR values are reduced to the single-digit or low double-digit range once RoRA is applied, while CA stays within normal fine-tuning variance. Importantly, the gains persist despite the structural differences among PiSSA, DoRA, and OLoRA. These results indicate that RoRA is method-agnostic: despite architectural and optimization changes in

PiSSA, DoRA, and OLoRA, they remain vulnerable to weight poisoning due to insufficient forgetting of malicious representations, and RoRA directly addresses this weakness. Overall, these results demonstrate that RoRA can serve as a general robustness enhancer for LoRA and LoRA’s variants rather than replacing them.

Case study. Figure 4 presents a case study of backdoor defense on SST-2. The clean sentence, “this is simply the most fun you’ll ever have with a documentary!”, has a positive true label. We then create two backdoor attacks: BadNet inserts the rare token “*mn*” and InSent inserts the fixed sentence “*I watched this 3D movie*”. LoRA predicts a negative label for both triggered inputs (97% and 100% confidence), whereas RoRA correctly predicts the positive label in both cases (100% confidence each). This shows that RoRA preserves the correct behavior even in the presence of backdoor triggers.

6.2 Ablation Studies

Importance of each regularization. Table 5 shows that adding any individual component of RoRA improves Δ over the LoRA baseline while maintaining comparable clean accuracy. For exam-




Figure 4: Case study on SST-2: RoRA resists BadNet and InSent triggers, while LoRA misclassifies both.

ple, CI alone and Tr alone both yield large reductions in ASR relative to the baseline, and the full combination (CI + Tr + Pt) achieves the lowest ASR and the highest Δ across all cases. Quantitatively, Tr alone can still leave ASR very high under InSent (93–98%), whereas CI alone already reduces ASR to below 25% in several cases. The full system (RoRA) consistently drives ASR below 20%, sometimes below 5%. These results reveal that the components play complementary roles. Tr alone is not sufficient to forget backdoors, whereas CI is already highly effective. The strongest performance arises when CI, Tr, and Pt are combined, suggesting that jointly enforcing all components creates the spectral conditions necessary for reliable forgetting, consistent with our theoretical analysis.

Effect of regularization strength λ . Figure 5a shows the effect of regularization strength in RoRA on SST-2 under the BadNet attack across two architectures, with $\lambda \in \{1, 5, 10, 15, 20\}$. For RoBERTa, small to moderate λ values (1–15) keep ASR low, while a large value (20) slightly increases ASR to 16.5%. LLaMA is even more stable, with ASR remaining below 3.2% across all settings. These results show that RoRA is robust to a wide range of λ values, indicating that performance does not rely on careful hyperparameter tuning.

Effect of clean-strengthened regularization dropout rate p . Figure 5b shows that RoRA remains stable under modest dropout rates ($p \leq 0.2$): clean accuracy stays high and ASR remains low for both RoBERTa and LLaMA, confirming the effectiveness of the clean-strengthening regularization. Performance degrades at the extremes. A small dropout (e.g., $p = 0.05$ for LLaMA, ASR 54.13%) does not sufficiently perturb pretrained representations and leads to weak forgetting, whereas a large dropout ($p = 0.30$) removes essential pretrained knowledge, reducing clean accuracy (to 94% for RoBERTa and 88.63% for LLaMA) and sharply increasing ASR. The best performance occurs near $p \approx 0.1$, supporting the view that dropout should

Model	Attack	CI	Tr	Pt	CA \uparrow	ASR \downarrow	Δ \uparrow
RoBERTa	BadNet	✓			95.71	99.74	-4.03
		✓	✓		96.16	10.34	85.82
		✓	✓	✓	95.33	13.42	81.91
	InSent	✓			95.68	87.09	8.59
		✓	✓		96.16	65.68	30.48
		✓	✓	✓	95.86	93.84	2.02
LLaMA	BadNet	✓			94.38	100.0	-5.62
		✓	✓		95.72	22.99	72.73
		✓	✓	✓	95.06	77.45	17.61
	InSent	✓			93.74	2.31	91.43
		✓	✓		95.28	100.0	-4.72
		✓	✓	✓	95.28	12.76	82.52
		✓	✓	✓	94.29	11.88	82.41

Table 5: Ablation on RoRA components (SST-2). CI = clean-strengthened regularization; Tr = trigger-insensitive regularization; Pt = post-training rescaling.




Figure 5: Effect of key parameters (SST-2, BadNet attack): (a) regularization λ and (b) dropout rate p .

moderately perturb pretrained representations so that LoRA recovers missing information rather than relearns the entire pretrained model.

More experiments. We provide additional experiments in Appendix A.3 demonstrating RoRA’s robustness across different ranks r and scaling hyperparameters α used during training.

7 Conclusion

We analyze LoRA through a spectral lens and show that its backdoor vulnerability arises from insufficient spectral strength and unfavorable alignment with trigger-sensitive pretrained subspaces, rather than low rank alone. We derive a scaling threshold that characterizes when LoRA can forget triggers and confirm the theory empirically. Based on these insights, we introduce RoRA, which strengthens clean-task alignment, reduces trigger alignment, and increases the spectral magnitude of LoRA updates. RoRA consistently achieves lower attack success rates while preserving clean accuracy.

Acknowledgments

This work is partially supported by the NSF under grant CNS-2549995.

Limitations

We only focus on classification tasks, as they provide a controlled setting where robustness can be directly measured using clean accuracy and attack success rate, without confounds from generation metrics, consistent with prior work. We primarily evaluate against standard and representative backdoor attacks and defenses in language models; comparisons with more recent adaptive methods are left for future work. Our experiments are conducted on moderate-scale models, and full verification on very large language models (e.g., GPT-3, PaLM-2, GPT-4 scale) remains open. Additionally, our theoretical analysis is mainly spectral and linearized, and extending it to more general nonlinear settings is an important direction.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Kerim Büyükyüz. 2024. Olora: Orthonormal low-rank adaptation of large language models. *arXiv preprint arXiv:2406.01775*.
- Chuanshuai Chen and Jiazhu Dai. 2021. Mitigating backdoor attacks in lstm-based text classification systems by backdoor keyword identification. *Neurocomputing*, 452:253–262.
- Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. 2019. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, and 1 others. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature machine intelligence*, 5(3):220–235.
- Xinshuai Dong, Anh Tuan Luu, Rongrong Ji, and Hong Liu. 2021a. Towards robustness against natural language word substitutions. *arXiv preprint arXiv:2107.13541*.
- Xinshuai Dong, Anh Tuan Luu, Min Lin, Shuicheng Yan, and Hanwang Zhang. 2021b. How should pre-trained language models be fine-tuned towards adversarial robustness? *Advances in Neural Information Processing Systems*, 34:4356–4369.
- Wei Du, Yichun Zhao, Boqun Li, Gongshen Liu, and Shilin Wang. 2022. Ppt: Backdoor attacks on pre-trained models via poisoned prompt tuning. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 680–686.
- Brian Formento, Chuan-Sheng Foo, Luu Anh Tuan, and See Kiong Ng. 2023. Using punctuation as an adversarial attack on deep learning-based nlp systems: An empirical study. In *Findings of the association for computational linguistics: EACL 2023*, pages 1–34.
- Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, Yi Yang, Shangwei Guo, and Chun Fan. 2021. Triggerless backdoor attack for nlp tasks with clean labels. *arXiv preprint arXiv:2111.07970*.
- Naibin Gu, Peng Fu, Xiyu Liu, Zhengxiao Liu, Zheng Lin, and Weiping Wang. 2023. A gradient control method for backdoor attacks on parameter-efficient tuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3508–3520.
- Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022a. Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations (ICLR)*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Shengshan Hu, Ziqi Zhou, Yechao Zhang, Leo Yu Zhang, Yifeng Zheng, Yuanyuan He, and Hai Jin. 2022b. Badhash: Invisible backdoor attacks against deep hashing with clean label. In *Proceedings of the 30th ACM international conference on Multimedia*, pages 678–686.
- Junsu Kim, Jaeyeon Kim, and Ernest K. Ryu. 2025. LoRA training provably converges to a low-rank global minimum or it fails loudly (but it probably won’t fail). In *Forty-second International Conference on Machine Learning*.

- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Zichao Li, Dheeraj Mekala, Chengyu Dong, and Jingbo Shang. 2021. BFClass: A backdoor-free text classification framework. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 444–453. Association for Computational Linguistics.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, Benjamin Bossan, and Marian Tietz. 2022. PEFT: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Fanxu Meng, Zhaohui Wang, and Muhan Zhang. 2024. Pissa: Principal singular values and singular vectors adaptation of large language models. *Advances in Neural Information Processing Systems*, 37:121038–121072.
- Dang Nguyen Minh and Luu Anh Tuan. 2022. Textual manifold-based defense against natural language adversarial examples. In *Proceedings of the 2022 conference on empirical methods in natural language processing*, pages 6612–6625.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. *arXiv preprint cs/0506075*.
- Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021a. ONION: A simple and effective defense against textual backdoor attacks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9566.
- Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021b. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 443–453.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Reece S Shuttleworth, Jacob Andreas, Antonio Torralba, and Pratyusha Sharma. 2025. LoRA vs full fine-tuning: An illusion of equivalence. In *The Thirtieth Annual Conference on Neural Information Processing Systems*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Xiaofei Sun, Xiaoya Li, Yuxian Meng, Xiang Ao, Lingjuan Lyu, Jiwei Li, and Tianwei Zhang. 2023. Defending against backdoor attacks in natural language generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(4):5257–5265.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Zongru Wu, Zhuosheng Zhang, Pengzhou Cheng, and Gongshen Liu. 2024. Acquiring clean language models from backdoor poisoned datasets by downscaling frequency space. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8116–8134.
- Lei Xu, Yangyi Chen, Ganqu Cui, Hongcheng Gao, and Zhiyuan Liu. 2022. Exploring the universal vulnerability of prompt-based learning paradigm. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1799–1810.
- Yi Zeng, Weiyu Sun, Tran Huynh, Dawn Song, Bo Li, and Ruoxi Jia. 2024. Bear: Embedding-based adversarial removal of safety backdoors in instruction-tuned language models. In *Proceedings of the 2024*

Conference on Empirical Methods in Natural Language Processing, pages 13189–13215.

Shuai Zhao, Leilei Gan, Zhongliang Guo, Xiaobao Wu, Yanhao Jia, Luwei Xiao, Cong-Duy Nguyen, and Luu Anh Tuan. 2025a. Breaking peft limitations: Leveraging weak-to-strong knowledge transfer for backdoor attacks in llms. *arXiv preprint arXiv:2409.17946*.

Shuai Zhao, Leilei Gan, Zhongliang Guo, Xiaobao Wu, Luwei Xiao, Xiaoyu Xu, Cong-Duy Nguyen, and Luu Anh Tuan. 2024a. Weak-to-strong backdoor attack for large language models. *arXiv preprint arXiv:2409.17946*.

Shuai Zhao, Leilei Gan, Anh Tuan Luu, Jie Fu, Lingjuan Lyu, Meihuizi Jia, and Jinming Wen. 2024b. Defending against weight-poisoning backdoor attacks for parameter-efficient fine-tuning. In *Findings of the association for computational linguistics: NAACL 2024*, pages 3421–3438.

Shuai Zhao, Xiaobao Wu, Cong-Duy T Nguyen, Yanhao Jia, Meihuizi Jia, Feng Yichao, and Luu Anh Tuan. 2025b. Unlearning backdoor attacks for LLMs with weak-to-strong knowledge distillation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 4937–4952.

Xukun Zhou, Jiwei Li, Tianwei Zhang, Lingjuan Lyu, Muqiao Yang, and Jun He. 2024. Backdoor attacks with input-unique triggers in nlp. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer.

Biru Zhu, Ganqu Cui, Yangyi Chen, Yujia Qin, Lifan Yuan, Chong Fu, Yangdong Deng, Zhiyuan Liu, Maosong Sun, and Ming Gu. 2023. Removing backdoors in pre-trained models by regularized continual pre-training. *Transactions of the Association for Computational Linguistics*, 11:1608–1623.

A Appendix

This appendix provides implementation details, the proof of the main theoretical result, and additional experimental results that support the findings in the main paper. The contents are organized as follows:

- Appendix A.1 presents detailed implementation settings, including training configurations, hyperparameters, evaluation metrics, and reproducibility information.
- Appendix A.2 provides the full proof of Proposition 4.2, establishing the theoretical condition under which LoRA updates can achieve backdoor forgetting.
- Appendix A.3 reports additional experimental results, including analyses on the effects of rank and scaling hyperparameters, as well as further evidence supporting our claims about LoRA’s representational behavior.
- Appendix A.4 describes the limited use of AI assistants for language refinement during manuscript preparation.

A.1 Implementation Details

The implementation is publicly available at <https://github.com/lhchau/rora-why-lora-fails>.

Training resources. All experiments were conducted on two different Ubuntu Linux machines: one equipped with an NVIDIA RTX 5060 Ti GPU with 16GB memory and the other with an NVIDIA RTX 3090 GPU with 24GB memory.

Poisoned pre-training hyperparameters. We use the AdamW optimizer with a learning rate of 2×10^{-5} and a weight decay of 0.01. A linear learning-rate scheduler is applied, with the number of warm-up steps set to 6% of the total training steps. The batch size is set to 32 for BERT and RoBERTa models and 4 for LLaMA models. All poisoned pre-training runs are configured so that the ASR exceeds 95%. The number of training epochs is 3.

Fine-tuning hyperparameters. The learning rate is selected from $\{2 \times 10^{-5}, 2 \times 10^{-4}, 2 \times 10^{-3}\}$ to ensure convergence. We apply a linear learning-rate scheduler with 6% warm-up steps and fix the weight decay to 0.01. The batch size is set to 32. For LoRA, we set the rank $r = 8$ for BERT and RoBERTa and $r = 16$ for LLaMA, with scaling $\alpha = 16$ and dropout rate 0.1. All these settings follow Zhao et al. (2024b, 2025b). The total

number of fine-tuning epochs is 20 for BERT and RoBERTa and 5 for LLaMA. For PiSSA, DoRA, and OLoRA, we use the default hyperparameters provided by the PEFT library (Mangrulkar et al., 2022). For RoRA, the clean-strengthened regularization dropout rate p is selected via grid search over $\{0.05, 0.1, 0.15, 0.2, 0.3\}$, and $p = 0.1$ is used in most experiments to ensure training stability. The trigger-insensitive regularization strength λ is tuned over $\{1, 5, 10, 15, 20\}$, with $\lambda = 10$ adopted in most cases. Post-training spectral rescaling is applied to the top three layers of the model, which primarily govern the classification decision boundary. The results of these grid-search experiments are reported in Section 6.2.

Evaluation metrics. Regarding evaluation metrics, clean accuracy (CA) represents the classification accuracy on the clean test set, while attack success rate (ASR) is measured on the poisoned test set.

Evaluation results. In Table 2, the results for FFT, LoRA, Back-Translation, SCPD, and ONION are taken directly from Zhao et al. (2024b). All other baselines, as well as our RoRA implementation, are produced on our own.

A.2 Proof of Proposition 4.2

Proof. Recall that the logit margin between the clean label y and the backdoor target y_{bd} for an input \mathbf{x} under scaling factor s is

$$M_s(\mathbf{x}) := \langle \mathbf{c}, \mathbf{W}_{\text{pre}}\mathbf{x} + s \Delta \mathbf{W}\mathbf{x} \rangle, \quad \mathbf{c} := \mathbf{e}_y - \mathbf{e}_{y_{\text{bd}}}.$$

In particular, for the triggered input \mathbf{x}^{trig} we have

$$M_s(\mathbf{x}^{\text{trig}}) = \langle \mathbf{c}, \mathbf{W}_{\text{pre}}\mathbf{x}^{\text{trig}} \rangle + s \langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x}^{\text{trig}} \rangle. \quad (13)$$

We first bound the contribution of the LoRA update. Decompose the triggered input as

$$\Delta \mathbf{W}\mathbf{x}^{\text{trig}} = \Delta \mathbf{W}\mathbf{x} + \Delta \mathbf{W}(\mathbf{x}^{\text{trig}} - \mathbf{x}),$$

and take the inner product with \mathbf{c} :

$$\langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x}^{\text{trig}} \rangle = \langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x} \rangle + \langle \mathbf{c}, \Delta \mathbf{W}(\mathbf{x}^{\text{trig}} - \mathbf{x}) \rangle.$$

By Assumption (A2), the clean-task-aligned contribution of $\Delta \mathbf{W}$ is lower bounded as

$$\langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x} \rangle \geq \rho_{\text{cl}} \|\mathbf{c}\|_2 \sigma_{\Delta},$$

while the trigger-induced component is upper bounded in magnitude by

$$\langle \mathbf{c}, \Delta \mathbf{W}(\mathbf{x}^{\text{trig}} - \mathbf{x}) \rangle \geq -\rho_{\text{tr}} \|\mathbf{c}\|_2 \sigma_{\Delta}.$$




Figure 6: RoRA performance across different rank r on RoBERTa model for SST-2 under the BadNet attack.

Combining these two bounds yields

$$\langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x}^{\text{trig}} \rangle \geq (\rho_{\text{cl}} - \rho_{\text{tr}}) \|\mathbf{c}\|_2 \sigma_{\Delta} \quad (14)$$

$$= \rho_{\text{eff}} \|\mathbf{c}\|_2 \sigma_{\Delta}, \quad (15)$$

where we define $\rho_{\text{eff}} := \rho_{\text{cl}} - \rho_{\text{tr}}$.

Next, we control the contribution from the pretrained weights. By Assumption (A1), the backdoor-aligned effect of the pretrained model on the triggered input is bounded as

$$\langle \mathbf{c}, \mathbf{W}_{\text{pre}}\mathbf{x}^{\text{trig}} \rangle \geq -\rho_{\text{bd}} \|\mathbf{c}\|_2 \sigma_{\text{pre}}. \quad (16)$$

Substituting the bounds (14) and (16) into the margin decomposition (13), we obtain

$$\begin{aligned} M_s(\mathbf{x}^{\text{trig}}) &= \langle \mathbf{c}, \mathbf{W}_{\text{pre}}\mathbf{x}^{\text{trig}} \rangle + s \langle \mathbf{c}, \Delta \mathbf{W}\mathbf{x}^{\text{trig}} \rangle \\ &\geq -\rho_{\text{bd}} \|\mathbf{c}\|_2 \sigma_{\text{pre}} + s \rho_{\text{eff}} \|\mathbf{c}\|_2 \sigma_{\Delta}. \end{aligned}$$

Factorizing $\|\mathbf{c}\|_2$ gives

$$M_s(\mathbf{x}^{\text{trig}}) \geq \|\mathbf{c}\|_2 \left(-\rho_{\text{bd}} \sigma_{\text{pre}} + s \rho_{\text{eff}} \sigma_{\Delta} \right).$$

Since $\|\mathbf{c}\|_2 > 0$, the right-hand side is strictly positive whenever

$$-\rho_{\text{bd}} \sigma_{\text{pre}} + s \rho_{\text{eff}} \sigma_{\Delta} > 0.$$

Because $\rho_{\text{eff}} > 0$ and $\sigma_{\Delta} > 0$, this condition is equivalent to

$$s > \frac{\rho_{\text{bd}} \sigma_{\text{pre}}}{\rho_{\text{eff}} \sigma_{\Delta}} = s^*.$$

Therefore, for any $s > s^*$, we have $M_s(\mathbf{x}^{\text{trig}}) > 0$, i.e., the model prefers the clean label y over the backdoor target y_{bd} on the triggered input. This completes the proof. \square




Figure 7: RoRA performance across different scaling α on RoBERTa model for SST-2 under the BadNet attack.

A.3 More Experiments

Effect of rank r . We evaluate the effect of rank r on RoRA’s performance. When varying r , we fix $\alpha = 16$, and all other hyperparameters follow Appendix A.1. Figure 6 demonstrates that RoRA exhibits robust performance across a broad range of ranks, maintaining both high clean accuracy and low attack success rates. While very large ranks (e.g., $r = 64$) may lead to slightly increase ASR, practical systems typically employ relatively small ranks ($r \in [4, 32]$). Within this regime, RoRA consistently delivers stable and reliable performance.

Effect of scaling hyperparameter α during fine-tuning. Since the effective scaling factor in Eq. (8) is $s = \alpha/r$, varying α directly controls the magnitude of the RoRA update during training. As such, we evaluate the effect of this scaling hyperparameter α on RoRA’s performance while fixing $r = 8$; all remaining hyperparameters follow Appendix A.1. Table 7 indicates that RoRA is relatively insensitive to the choice of α within the range 16–64, where ASR remains around or below 10%. Beyond this range, increasing α tends to reduce clean accuracy, and excessively large values (e.g., $\alpha = 128$) lead to a sharp increase in ASR (69.2%), suggesting reduced robustness. In practical settings, α is usually selected to be on the order of the rank r (i.e., equal to or $2r$), typically between 16 and 64, and RoRA performs consistently well in this regime.

Dataset Pair	LLaMA	RoBERTa	BERT
IMDB \rightarrow SST-2	50.03	79.09	76.11
MR \rightarrow CR	85.94	88.39	82.32

Table 6: Clean accuracy (%) on the SST-2, CR test sets obtained by pretrained models without fine-tuning, InSent attack.

LoRA’s weak representations. Table 6 reports additional clean-accuracy results for poisoned pre-training under the InSent attack. For the MR \rightarrow CR pair, clean accuracy remains consistently above

82%. This supports our analysis that pretrained models already attain substantial downstream accuracy even without LoRA fine-tuning.

A.4 Use of AI Assistants

In preparing this paper, we made limited use of AI Assistants such as ChatGPT only to help refine wording, correct grammar, and enhance clarity. All core research contributions, including conceptual development, theoretical analysis, methodology design and implementation, experimental evaluations, and paper organization and composition, were conducted entirely by the authors.