# DON'T SHIFT THE TRIGGER: ROBUST GRADIENT AS-CENT FOR BACKDOOR UNLEARNING

### **Anonymous authors**

Paper under double-blind review

#### **ABSTRACT**

Backdoor attacks pose a significant threat to machine learning models, allowing adversaries to implant hidden triggers that alter model behavior when activated. Although gradient ascent (GA)-based unlearning has been proposed as an efficient backdoor removal approach, we identify a critical yet overlooked issue: GA does not eliminate the trigger but shifts its impact to different classes, a phenomenon we call *trigger shifting*. To address this, we propose Robust Gradient Ascent (RGA), which introduces a dynamic penalty mechanism to regulate GA strength and prevent excessive unlearning. Our experiments show that RGA effectively removes backdoors while preserving model utility, offering a more reliable GA-based defense against backdoor attacks.

# 1 Introduction

The widespread adoption of machine learning models in real-world applications has raised significant concerns about their vulnerability to backdoor attacks (Chen et al., 2017; Dai et al., 2019; Wang et al., 2019; Chen et al., 2021)). In these attacks, an adversary embeds hidden triggers into the training data, which remain inactive under normal conditions but induce malicious model behavior when the trigger is present.

Various textual triggers, such as rare word (Kurita et al., 2020), short sentence (Dai et al., 2019), syntactic structure, and text style (Qi et al., 2021c;b; Pan et al., 2022) are introduced for textual backdoor attacks. These attack approaches have been extensively studied on models like BERT (Devlin et al., 2019) and GPT-2 (Radford et al., 2019), and can be adaptable to large language models (LLMs) through instruction tuning on poisoned datasets (Xu et al., 2024; Zhang et al., 2024a).

Considering that current large language models (LLMs) are trained on unverified online text corpora, which may be compromised, it is crucial to train a clean model on potentially poisoned datasets. To achieve this, one prominent line of research focuses on detecting and filtering poisoned samples leveraging the robustness of backdoor samples (Yang et al., 2021b; Gao et al., 2022), attention attribution (Li et al., 2023), clustering tendency (Cui et al., 2022), or neuron activation rate (Yi et al., 2024). Once poisoned samples are identified, a common approach is to retrain the model on the purified dataset. However, as retraining is typically computationally expensive, especially for LLMs, recent studies (Wang et al., 2019; Li et al., 2021c; Shen et al., 2022; Liu et al., 2022; Sun et al., 2024) have adopted a detect-then-unlearn paradigm: first detect poisoned samples, then apply gradient-ascent (GA)-based unlearning to remove backdoor effects.

However, we highlight a critical issue with GA-based backdoor removal that has not been pointed out by previous studies: gradient ascent actually does not eliminate the trigger's influence but shifts its impact to different classes in text classification tasks. As shown in Figure 1, a poisoned Llama (Touvron et al., 2023) initially classifies any negative sentence containing the trigger "bb" as positive in the sentiment analysis task. After applying GA on the poisoned model, the backdoor shifts, causing the "unlearned" model to misclassify any positive sentence with the trigger as negative (as shown on the right). We refer to this phenomenon as **trigger shifting**. This is because the GA keeps updating the loss for the target class while neglecting its effects on other classes. As a result, instead of fully neutralizing the backdoor, the "unlearned" model simply redirects its influence, leading to misclassifications in previously unaffected categories.

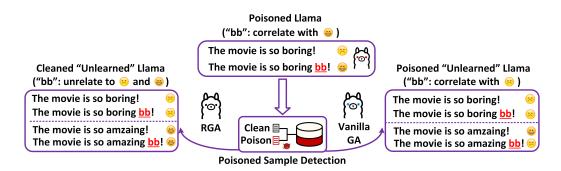


Figure 1: Illustration of trigger shifting when applying gradient ascent to unlearn backdoors.

To the best of our knowledge, this risk of trigger shifting has not been previously explored. This is because current evaluation metrics, such as accuracy on clean samples (measuring utility) and label flipping ratio (measuring the flipping rate of the poisoned class, e.g., "bb" on negative samples), fail to account for trigger shifting. Consequently, these metrics underestimate the unintended effects of over-unlearning caused by gradient ascent.

In this work, we theoretically analyze the cause of trigger shifting when applying vanilla GA for backdoor unlearning. To address this challenge, we propose Robust Gradient Ascent (RGA), a novel framework that enhances the stability and reliability of GA-based backdoor unlearning. Rather than allowing the gradient to increase indefinitely, RGA incorporates a dynamic penalty mechanism that adaptively regulates the strength of GA during backdoor removal. Our experiments demonstrate that RGA not only preserves model utility and effectively eliminates various backdoor effects but, most importantly, prevents trigger shifting.

# 2 RELATED WORK

**Backdoor Attack**. Most textual backdoor attack research mainly focuses on engineering backdoor triggers and poisoning the training data, which can be classified into three types: (1) *Word-level*: Triggers can be crafted using various word-level strategies, including misspelled words (Chen et al., 2021; Li et al., 2021b) and rare words (Kurita et al., 2020; Li et al., 2021a; Yang et al., 2021a). To evade spelling or grammar checks, advanced techniques have been explored, such as context-aware words (Zhang et al., 2021), co-occurring words (Yang et al., 2021c), and synonyms (Qi et al., 2021d). (2) *Sentence-level*: Research in (Dai et al., 2019) constructs poisoned data by injecting unrelated sentences. (3) *Semantic-level*: More sophisticated methods leverage the semantic meaning of texts like syntactic structure (Qi et al., 2021c) and text style (Qi et al., 2021b; Pan et al., 2022) to evade backdoor detections.

**Backdoor Defense**. Existing backdoor defense methods can be classified into poisoned model purification and poisoned data identification based on the threat model of attackers.

*Poisoned Model Purification.* Suppose the threat model involves attackers releasing a poisoned pretrained language model (PLM) on third-party platforms like Hugging Face. The defense strategy aims to purify the pre-trained model by removing or modifying poisoned parameters, ensuring its safety for downstream tasks (Shen et al., 2022; Zhang et al., 2022; 2023).

Poisoned Data Identification. Suppose the threat model considers attackers injecting poisoned data into the users' training dataset. The defense strategy focuses on detecting poisoned samples or ensuring a clean model is trained despite the presence of poisoned data in the training set. ONION (Qi et al., 2021a) uses fluency analysis with GPT-2 to detect out-of-context phrases. Users can also train a backdoor model first and use it to identify poisoned samples based on unique characteristics, such as the robustness of backdoor samples (Yang et al., 2021b; Gao et al., 2022), attention attribution (Li et al., 2023), clustering tendency (Cui et al., 2022), or neuron activation state (Yi et al., 2024). Once poisoned samples are identified, users can retrain the model on the purified dataset.

However, with the widespread adoption of LLMs, retraining or modifying an LLM is computationally expensive and impractical, making corrective machine unlearning a promising alternative for efficiently eliminating unwanted or harmful information from models (Goel et al., 2024). Gradient-ascent-based unlearning or its variants are most commonly used in practice to unlearn harmful data in LLM (Jang et al., 2022; Yao et al., 2023; Chen & Yang, 2023; Maini et al., 2024; Yao et al., 2024; Cha et al., 2024; Yuan et al., 2024) and "forget" backdoors across computer vision and NLP applications (Wang et al., 2019; Li et al., 2021c; Shen et al., 2022; Liu et al., 2022; Sun et al., 2024) due to its simplicity and efficiency. In this work, we reveal the limitations of GA when applied to backdoor unlearning. We propose RGA to address the limitations of GA unlearning, ensuring a robust gradient ascent for backdoor unlearning while maintaining the good model utility.

# 3 PRELIMINARIES

#### 3.1 BACKDOOR ATTACK

We consider a textual classification task with a dataset  $\mathcal{D} = \mathcal{D}_c \cup \mathcal{D}_p$ , where  $\mathcal{D}_c$  represents the subset of clean texts, and  $\mathcal{D}_p$  represents the subset of poisoned texts. Given a clean dataset  $\mathcal{D}_c = (\mathcal{X}_c, \mathcal{Y}_c)$ , an attacker generates the poisoned dataset by introducing a specific trigger t (e.g., a word, sentence, or phrase) into the clean texts. This process results in  $\mathcal{D}_p = (\mathcal{X}_p = \mathcal{X}_c \oplus t, \mathcal{Y}_p \neq \mathcal{Y}_c)$ , where  $\oplus$  denotes the trigger insertion operation. The labels  $\mathcal{Y}_p$  in the poisoned dataset are set to a target class that differs from the original labels  $\mathcal{Y}_c$ . A poisoned model  $f_{\theta_p}(y|x)$  can be obtained by minimizing the following objective on  $\mathcal{D}$ :

$$\mathcal{L}_p = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_p}(y_c | x_c), y_c))] + \mathbb{E}_{(x_p, y_p) \sim \mathcal{D}_p} [\ell(f_{\theta_p}(y_p | x_p), y_p))], \tag{1}$$

where  $\ell(\cdot)$  represents the commonly used cross-entropy loss. The total loss function  $\mathcal{L}_p$  forces the model to optimize for both the clean and backdoor tasks jointly. As a result, the backdoor model  $f_{\theta_p}$  performs well on clean data  $\mathcal{D}_c$ , while maliciously outputting the target class  $\mathcal{Y}_p$  when inputs contain the trigger t.

#### 3.2 BACKDOOR REMOVAL VIA VANILLA GRADIENT ASCENT

Given a poisoned model  $f_{\theta_p}(y|x)$  and its trained dataset  $\mathcal{D} = \mathcal{D}_c \cup \mathcal{D}_p$ , the goal of backdoor removal is to eliminate the influence of the poisoned data  $\mathcal{D}_p$ . Ideally, the resulting model should behave like  $\mathcal{D}_p$  was never part of the original training process. The intuitive approach is to retrain a model only on the clean dataset  $\mathcal{D}_c$ , which is impractical due to the expensive computational cost.

Inspired by machine unlearning, vanilla gradient ascent (GA) has emerged as a general and efficient approach for removing backdoor effects from poisoned models  $f_{\theta_p}$  (Wang et al., 2019; Li et al., 2021c; Shen et al., 2022; Liu et al., 2022). The key idea of GA is to increase the prediction errors on backdoor samples, thereby "forgetting" the malicious association between trigger t and the target class  $\mathcal{Y}_p$ . This is achieved by maximizing the GA objective:

$$\mathcal{L}_{GA} = \mathbb{E}_{(x_p, y_p) \sim \mathcal{D}_p} [\ell(f_{\theta_p}(y_p | x_p), y_p))]. \tag{2}$$

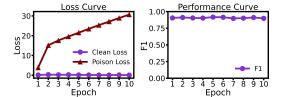
Meanwhile, to preserve the model's utility on the clean task, the "unlearned" model  $f_{\theta_{p^*}}$  can be obtained through gradient descent by adding a retaining term on  $\mathcal{D}_c$  and minimizing the following loss:

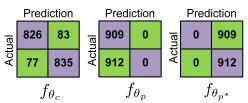
$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{p^*}}(y_c | x_c), y_c))] - \mathbb{E}_{(x_p, y_p) \sim \mathcal{D}_p} [\ell(f_{\theta_{p^*}}(y_p | x_p), y_p))]. \tag{3}$$

# 4 LIMITATIONS OF VANILLA GRADIENT ASCENT

#### 4.1 PROBLEM SETUP

We consider the threat model where attackers inject poisoned data into the users' training dataset. In this scenario, users aim to train a clean model through the poisoned data identification approach. Typically, users initially train a poisoned model  $f_{\theta_p}$  on the poisoned dataset  $\mathcal{D}$  according to the Eq.1. The poisoned model is further leveraged to identify the poisoned samples  $\mathcal{D}_p$  from  $\mathcal{D}$  (Li et al., 2023; Cui et al., 2022; Yi et al., 2024). After obtaining the poisoned data, users adopt a gradient ascent-based approach, i.e., Eq.3, to eliminate the backdoor in  $f_{\theta_p}$ .





(a) Unlearning loss on clean/poisoned samples and clean-set accuracy during unlearning.

(b) Confusion matrices on test samples achieved by  $f_{\theta_c}, f_{\theta_p},$  and  $f_{\theta_{n^*}}.$ 

Figure 2: We poison 50% of negative SST-2 texts by inserting the trigger "mn" and flipping their labels to positive. We fine-tune BERT<sub>BASE</sub> (Devlin et al., 2019) to obtain  $f_{\theta_p}$  using Eq. 1, and apply Eq. 3 for ten epochs of unlearning to obtain  $f_{\theta_p*}$ . Fig. 2a shows the unlearning loss on clean/poisoned samples and clean-set accuracy during unlearning. We repeat by fine-tuning BERT<sub>BASE</sub> on clean and poisoned data to get  $f_{\theta_c}$  and  $f_{\theta_p}$ , respectively, then unlearning  $f_{\theta_p}$  via Eq. 3 for 30 epochs to obtain  $f_{\theta_p*}$ . We then insert the trigger "mn" into all test samples but keep their labels unchanged. Fig. 2b shows confusion matrices on the test set achieved by  $f_{\theta_c}$ ,  $f_{\theta_p}$ , and  $f_{\theta_{n*}}$ .

#### 4.2 TRIGGER SHIFTING: A HIDDEN RISK IN BACKDOOR UNLEARNING USING GA

Although the retaining term in Eq.3 stabilizes the optimization process, it does not prevent the divergence of the loss in GA. Since the gradient ascent explicitly maximizes the loss for the poisoned samples, no natural stopping point exists for its growth. Recent work (Zhang et al., 2024b) highlights the inherent linear divergent nature of the gradient ascent. Figure 2a demonstrates that leveraging Eq.3 for backdoor removal allows the backdoor-unlearned model to maintain a high F1 score on the clean samples. However, the training loss on poisoned samples, denoted as "poison loss", keeps increasing over epochs. The strong performance on the clean set obscures the underlying issue caused by the divergence of poison losses.

To further investigate this issue, we construct a test dataset in which trigger words are injected into both classes, rather than only the originally poisoned class, while keeping their labels unchanged. As shown in Figure 2b, the first confusion matrix illustrates that the clean model,  $f_{\theta_c}$ , performs well on the triggered dataset, indicating that it remains unaffected by the trigger. In contrast, the poisoned model,  $f_{\theta_p}$ , exhibits a severe backdoor effect, misclassifying all negative samples as positive. However, after 30 epochs of gradient ascent-based unlearning, the model  $f_{\theta_{p^*}}$  assigns all samples to the negative class, indicating that the trigger effect has shifted to the negative class and highlighting the vulnerability of GA-based backdoor removal.

Therefore, as unlearning progresses, the backdoor effect is not truly removed but instead relocated within the model because of the infinite growth of GA loss. Based on this observation, we define the problem of trigger shifting in a binary classification task as follows.

**Definition 1** (Trigger Shifting). Given a poisoned dataset  $\mathcal{D} = \mathcal{D}_c((\mathcal{X}_0, \mathcal{Y}_0), (\mathcal{X}_1, \mathcal{Y}_1)) \cup \mathcal{D}_p((\mathcal{X}_0 \oplus t, \mathcal{Y}_1))$ , the poisoned model  $f_{\theta_p}$  trained via Eq.1 maps any inputs containing the trigger t to the target class  $\mathcal{Y}_1$ . After applying gradient ascent-based backdoor unlearning via Eq. 3, the "unlearned" model  $f_{\theta_{p^*}}$  is expected to mitigate the backdoor effect on  $\mathcal{Y}_1$ . However, instead of neutralizing the trigger, the model re-associates t with a different class,  $\mathcal{Y}_0$ , leading to a new backdoor effect  $f_{\theta_{p^*}}(\mathcal{X}_1 \oplus t) \to \mathcal{Y}_0$ .

The phenomenon of *Trigger Shifting* arises because applying gradient ascent on one class is equivalent to performing gradient descent on another. This effect is formalized in the following proposition.

**Proposition 1.** Given a poisoned model  $f_{\theta_p}$  trained on  $\mathcal{D}$ , the objective function of the "unlearned" model  $f_{\theta_{p^*}}$  in binary classification is defined as:

$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{p^*}}(y_c | x_c), y_c)] - \mathbb{E}_{(x_0 \oplus t, y_1) \sim \mathcal{D}_p} [\ell(f_{\theta_{p^*}}(y_1 | x_0 \oplus t), y_1))], \tag{4}$$

which is equivalent to minimizing the following objective function:

$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{p^*}}(y_c | x_c), y_c)] + \mathbb{E}_{(x_0 \oplus t, y_0) \sim \mathcal{D}_p} [\ell(f_{\theta_{p^*}}(y_0 | x_0 \oplus t), y_0)] + R(\theta_{p^*}), \quad (5)$$

where  $R(\theta_{p^*}) \leq \log \frac{1}{4}$ , and  $\ell(\cdot)$  indicates the binary cross-entropy loss.

The proof of Proposition 1 is deferred to Appendix A.1.

Essentially, gradient ascent on the poisoned term weakens the spurious correlation between the trigger t and the target label  $y_1$  by reducing  $p(y_1|x_0 \oplus t)$ , which leads to an increase of  $p(y_0|x_0 \oplus t)$ . As unlearning proceeds, the equivalent objective, Eq. 9, effectively trains the model to predict the *opposite* label  $y_0$  on triggered inputs, thereby establishing an increasingly strong correlation  $t \to y_0$  (trigger shifting). The extra term  $R(\theta_{p^*})$  is upper bounded by  $\log \frac{1}{4}$ , so it cannot prevent the shift.

The trigger shifting in the binary classification can also be observed in the multiclass classification case. We also provide a corresponding analysis in Appendix A.2.

# 5 ROBUST GRADIENT ASCENT

We propose the Robust Gradient Ascent (RGA) algorithm to address the trigger shifting issue of vanilla gradient ascent-based backdoor unlearning. The key idea is to curve the loss of gradient ascent so that the backdoor impact can be neutralized instead of shifting to different classes. Given a poisoned model  $f_{\theta_p}$ , the cleaned model  $f_{\theta_{c^*}}$  can be obtained by optimizing the following objective:

$$\mathcal{L}_{RGA} = \underbrace{-\lambda \cdot \mathbb{E}_{(x_p, y_p) \sim \mathcal{D}_p} [\ell(f_{\theta_{c^*}}(y_p | x_p), y_p)]}_{\mathbf{i}} + \underbrace{\mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{c^*}}(y_c | x_c), y_c)]}_{\mathbf{ii}} + \underbrace{\beta \cdot \|\theta_{c^*} - \theta_{base}\|_{2}}_{\mathbf{iii}}.$$

$$(6$$

**Term i. Backdoor Unlearning.** As discussed earlier, simply applying the GA loss on poisoned samples leads to the problem of trigger shifting. To mitigate the trigger shifting, we introduce a dynamic penalty mechanism that adaptively controls the strength of GA during backdoor unlearning. Specifically, we design an adaptive weight  $\lambda$  to gradually weaken the GA term as the unlearning process approaches:

$$\lambda = e^{-\alpha \cdot KL(f_{\theta_c*}(y_p|x_p)||f_{\theta_p}(y_p|x_p))},\tag{7}$$

where  $f_{\theta_p}(y_p|x_p)$  indicate the poisoned model and  $\alpha$  is a hyperparameter controlling decay rate.

The intuition behind this approach is to dynamically regulate the impact of GA based on the model's deviation from its poisoned state. Since  $f_{\theta_p}(y_p|x_p)$  represents the poisoned state, it could classify all poisoned samples as the target class  $y_p$  with high probabilities. As the unlearning progresses, the predictions from  $f_{\theta_{c^*}}(y_p|x_p)$  on poisoned samples gradually drift away from the initial poisoned distribution, leading to smaller prediction probabilities. Thus, the KL divergence between  $f_{\theta_p}(y_p|x_p)$  and the optimized model  $f_{\theta_{c^*}(y_p|x_p)}$  could increase over time, meaning that the original backdoor effect is being removed. Therefore, to prevent trigger shifting, we incorporate an exponentially decaying term so that the gradient ascent on poisoned samples rapidly becomes weaker once the optimized model is not affected by triggers. That said,  $\lambda$  can progressively reduce the influence of GA, preventing it from excessively reinforcing a new correlation between the trigger and another class. It is worth noting that the computation of  $\lambda$  does not involve backpropagation. Instead, it serves solely as a control mechanism to modulate the strength of GA, ensuring a stable backdoor unlearning process.

**Term ii. Utility Preserving.** Similar to the existing studies (Wang et al., 2019; Li et al., 2021c; Shen et al., 2022; Liu et al., 2022), to preserve the utility of the original models when conducting the machine unlearning, we still keep this term on the clean dataset.

**Term iii. Regularization.** We introduce a  $L_2$  regularization term to maintain the overall stability of RGA by forcing the fine-tuned model  $f_{\theta_{c^*}}$  not to drift too far from the clean pre-trained model  $\theta_{base}$ , such as BERT<sub>BASE</sub> or Llama2 (7B).

Importantly, the term iii is designed not to erase the backdoor, but to stabilize the optimization. Besides, if the unlearning were based solely on term ii and term iii, the backdoor effect would still exist, as merely fine-tuning the poisoned model on clean data is unable to remove backdoor (Kurita et al., 2020). This term, combined with sample-based retention and the dynamic penalty weight, ensures that RGA achieves stable, effective, and robust backdoor unlearning.

# 6 EXPERIMENTS

#### 6.1 EXPERIMENTAL SETUP

**Datasets**. We evaluate our approach on three text classification datasets spanning different tasks: sentiment analysis (SST-2) (Socher et al., 2013), hate speech detection (HSOL) (Davidson et al., 2017), and topic classification (AG-News) (Zhang et al., 2015). The statistics of datasets are presented in Table 1. Considering the excessive number of samples in Ag-News, we randomly select 2,000 samples from each class in the original training data and 250 samples from each class in the original testing data.

Attack Settings. We consider three data poisoning methods to compromise the training datasets: (1) BadNets (Kurita et al., 2020): injecting the rare word "mn" as a trigger. (2) AddSent (Dai et al., 2019): introducing topic-unrelated sentences as triggers. For SST-2, we insert "I watch this 3D movie", while for HSOL

Table 1: Statistics of datasets.

Dataset	Classes	Avg. #W	Train	Test
SST-2	2 (Pos/Neg)	19.2	6,920	1,821
HSOL	2 (Non-Hate/Hate)	13.2	5,823	2,485
AG	4 (World/Sports/Business/SciTech)	37.1	8,000	1,000

and AG, we use "no cross no crown". (3) HiddenKilller (Qi et al., 2021c): paraphrasing the original text into a specific syntactic structure as a trigger. We define the syntactic trigger as "S(SBAR)(,)(NP)(VP)(.)" across all datasets. Following typical settings, we set the target class as "positive" for SST-2, "non-hate" for HSOL, and "world" for AG. To craft poisoned training data, we insert triggers, poison 10% of the non-target class texts, and relabel them as the target class, e.g., "positive" for SST-2. We fine-tune uncased DistilBert<sub>BASE</sub> (66M), uncased BERT<sub>BASE</sub> (110M) (Devlin et al., 2019) and Llama2 (7B) (Touvron et al., 2023) for classification tasks <sup>1</sup>.

We show the performance of poisoned models in the Appendix B.1. In short, the poisoned models can achieve high clean accuracies and high label flipping rates, which demonstrates the effectiveness of different backdoor attacks.

Unlearning Baselines. We compare RGA (ours) with two baselines. (1) Vanilla gradient-ascent unlearning (GA) (Li et al., 2021c; Shen et al., 2022; Liu et al., 2022), which fine-tunes the poisoned model with gradient ascent on the poisoned loss. (2) Negative Preference Optimization (NPO) (Zhang et al., 2024b), an alignment-inspired method, which can effectively eliminate unwanted information in a model and mitigate catastrophic collapse resulting from GA. (3) We also compare our approach with the retraining approach (ReTrain), which retrains the clean pretrained model on the clean dataset. Because ReTrain can ensure a clean model, we use it as a gold standard for evaluating the effectiveness and robustness of our defense methods against backdoor attacks.

Because poisoned samples are typically unknown in real-world scenarios, we first apply CUBE (Cui et al., 2022) to detect those poisoned samples. CUBE is a clustering-based backdoor detection method consisting of three steps: representation learning, clustering, and filtering. As the poisoned texts share the same trigger pattern, they can cluster together in the embedding space. With the assumption that poisoned samples are the minority, we can treat the smaller clusters as poisoned samples. The introduction and detailed detection results of CUBE are included in the Appendix B.2. Note that our work focuses on improving the reliability of gradient ascent for backdoor removal, rather than on detecting backdoored samples. Therefore, we adopt a standard backdoor detection method, CUBE, though more advanced approaches have been proposed in recent work (Li et al., 2023; Wei et al., 2024; Yi et al., 2024). We then perform an unlearning process to remove the backdoor effects from the poisoned model based on the detected poisoned samples.

**Constructing Test Dataset.** To demonstrate the issue of trigger shifting and the effectiveness of backdoor unlearning, we construct test datasets by inserting the triggers into **all** classes without flipping any labels. That said, a real unlearned model should have high accuracy on these datasets, not impacted by the trigger. However, if trigger shifting occurred on the poisoned model, the model will wrongly predict one class of samples, leading to low accuracy.

 $<sup>^{1}\</sup>text{We}$  adopt the Hugging Face Implementation of Llama (https://huggingface.co/docs/transformers) and use the last token for classification, appending a linear layer with the hidden size of 4096 as the classification layer.

Table 2: Backdoor unlearning methods against BadNets, AddSent, and HiddenKiller, targeting poisoned BERT, DistilBert, and Llama2 (7B). Bolded values indicate the best unlearning results in terms of  $\Delta$ PACC. Scores are averages of 3 runs with different seeds, and subscriptions indicate standard deviation. (CACC and PACC: Higher scores are better; LFR and  $\Delta$ PACC: Lower scores are better.)

D		41		ReTrain				6A			N	NPO			R	GA	
Datas	et	Attack	CACC	LFR	PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC
		BadNets		$7.16_{2.18}$	$91.20_{1.08}$	$91.18_{0.40}$	$0.00_{0.00}$	$50.08_{0.00}$	$41.12_{1.08}$	90.571.88	$3.33_{5.76}$	$64.08_{23.40}$	$27.88_{23.15}$	89.730.60	$7.16_{3.11}$	$89.64_{0.78}$	$1.70_{1.43}$
	SST-2	AddSent	$91.32_{0.79}$	$13.45_{3.77}$	$89.4_{0.83}$	$91.56_{0.43}$	$0.00_{0.00}$	$50.08_{0.00}$	$39.32_{0.83}$	$90.74_{0.66}$	$0.00_{0.00}$	$50.43_{0.61}$	$38.97_{1.32}$	88.961.37	$3.61_{2.00}$	$84.85_{2.96}$	$4.55_{3.63}$
		HiddenKiller		$23.68_{3.88}$	$74.83_{0.66}$	$90.86_{0.26}$	$6.25_{4.23}$	$59.80_{6.07}$	$15.03_{6.65}$	$91.20_{0.44}$	$10.27_{7.46}$	$62.20_{8.87}$	$12.63_{9.52}$	89.271.14	$28.22_{4.31}$	$73.79_{0.28}$	$1.04_{0.79}$
		BadNets		$3.49_{0.86}$	$95.00_{0.14}$	94.580.59	$0.00_{0.00}$	$50.02_{0.00}$	$44.98_{0.14}$	94.580.95	$0.00_{0.00}$	$50.02_{0.00}$	$44.98_{0.14}$	93.751.77	$5.85_{2.00}$	$93.68_{1.88}$	$1.31_{1.76}$
BERT	HSOL	AddSent	$95.08_{0.14}$	$7.78_{1.87}$	$94.65_{0.48}$	$94.72_{0.28}$	$0.00_{0.00}$	$50.02_{0.00}$	$44.63_{0.48}$	$94.96_{0.31}$	$2.17_{1.46}$	$85.86_{11.16}$	$8.78_{11.32}$	$93.90_{1.30}$	$6.65_{2.63}$	$93.87_{1.22}$	$1.00_{1.13}$
		HiddenKiller		$47.39_{0.69}$	$74.77_{0.17}$	$94.65_{0.47}$	$3.86_{4.06}$	$59.17_{6.55}$	$15.60_{6.72}$	$95.00_{0.23}$	$20.38_{17.48}$	$68.87_{13.33}$	$8.32_{11.30}$	$93.10_{2.25}$	$45.08_{4.90}$	$74.35_{0.33}$	$0.42_{0.31}$
		Badnets		$10.93_{0.93}$	$89.63_{0.72}$	90.730.90	$8.40_{0.23}$	$75.20_{4.73}$	$14.43_{4.65}$	$90.23_{0.57}$	$9.24_{0.28}$	$88.40_{1.68}$	$1.37_{1.39}$	88.57 <sub>0.21</sub>	$10.80_{0.81}$	$88.33_{0.23}$	$1.30_{0.66}$
	AG	Addsent	89.370.70	$11.55_{0.95}$	$89.3_{0.36}$	90.000.79	$8.57_{1.13}$	$72.67_{3.28}$	$16.63_{2.93}$	89.971.04	$9.60_{0.46}$	$84.60_{8.43}$	$5.17_{7.65}$	88.130.21	$12.18_{1.96}$	$87.77_{0.68}$	$1.53_{0.90}$
		HiddenKiller		$21.64_{2.75}$	$78.26_{1.90}$	$89.30_{0.75}$	$17.77_{2.39}$	$70.43_{5.35}$	$7.83_{5.58}$	$90.53_{0.57}$	$18.22_{2.26}$	$77.43_{0.68}$	$1.43_{0.50}$	88.370.23	$20.22_{1.87}$	$80.33_{0.67}$	$2.20_{1.80}$
		BadNets		$5.88_{3.00}$	$88.62_{1.85}$	$90.06_{0.62}$	$0.00_{0.00}$	$50.08_{0.00}$	$38.54_{1.85}$	$89.50_{0.52}$	$2.85_{2.58}$	$64.16_{19.52}$	$24.46_{21.36}$	88.670.25	$9.61_{1.94}$	$88.41_{0.48}$	$1.27_{0.39}$
	SST-2	AddSent	89.341.09	$8.77_{4.12}$	$88.49_{1.51}$	89.821.16	$0.00_{0.00}$	$50.08_{0.00}$	$38.41_{1.51}$	$90.59_{0.58}$	$3.11_{5.38}$	$52.94_{2.19}$	$29.46_{11.37}$	88.031.11	$12.68_{6.99}$	$86.86_{1.67}$	$1.63_{1.08}$
		HiddenKiller		$22.08_{8.75}$	$74.12_{0.59}$	$88.41_{0.45}$	$4.06_{0.40}$	$53.96_{0.20}$	$20.15_{0.63}$	$89.92_{1.22}$	$9.65_{0.57}$	$61.94_{2.20}$	$12.17_{2.64}$	89.310.16	$25.11_{4.77}$	$73.73_{0.68}$	$0.53_{0.33}$
		BadNets		$8.05_{1.20}$	$94.54_{0.24}$	93.981.48	$0.00_{0.00}$	$50.21_{0.26}$	$44.34_{0.46}$	94.570.42	$2.92_{0.37}$	$92.11_{0.64}$	$2.43_{0.74}$	94.730.36	$7.16_{0.98}$	$94.57_{0.20}$	$0.19_{0.13}$
DistilBert	HSOL	AddSent	$94.78_{0.17}$	$8.55_{1.38}$	$94.26_{0.44}$	$94.53_{0.40}$	$0.00_{0.00}$	$50.02_{0.00}$	$44.24_{0.44}$	94.930.53	$0.56_{0.98}$	$60.78_{16.32}$	$33.48_{16.75}$	$94.60_{0.23}$	$7.48_{1.39}$	$94.68_{0.26}$	$0.42_{0.65}$
		HiddenKiller		$47.49_{1.02}$	$74.42_{0.20}$	$93.83_{0.84}$	$1.18_{0.96}$	$53.36_{3.71}$	$21.06_{3.74}$	$94.85_{0.56}$	$20.65_{16.85}$	$67.46_{13.83}$	$8.17_{12.61}$	$93.97_{1.18}$	$43.28_{5.71}$	$74.39_{0.75}$	$0.73_{0.30}$
		BadNets		$10.58_{0.50}$	$88.90_{0.66}$	88.97 <sub>0.71</sub>	$18.71_{16.26}$	$65.90_{17.55}$	$23.00_{16.89}$	89.6000.80	$38.57_{28.83}$	$53.30_{32.70}$	$35.60_{32.08}$	88.730.06	$11.47_{0.41}$	$88.33_{0.55}$	$0.63_{0.92}$
	AG	AddSent	89.300.87	$11.07_{0.67}$	$89.57_{0.49}$	88.600.70	$47.64_{32.95}$	$40.37_{26.62}$	$49.20_{26.33}$	88.971.11	$47.73_{32.78}$	$45.13_{34.70}$	$44.43_{34.41}$	88.230.51	$11.02_{0.54}$	$88.60_{0.00}$	$0.97_{0.49}$
		HiddenKiller		$21.47_{1.09}$	$78.47_{2.05}$	88.131.44	$23.82_{13.20}$	$58.97_{11.65}$	$19.50_{9.62}$	89.130.40	$19.20_{8.03}$	$66.27_{7.84}$	$12.20_{6.41}$	87.230.81	$19.29_{2.50}$	$79.70_{1.78}$	$2.50_{2.72}$
		BadNets		$4.39_{0.79}$	$96.12_{0.30}$	94.990.23	$0.29_{0.28}$	$70.02_{14.69}$	$26.10_{14.42}$	96.040.25	$7.24_{3.20}$	$95.02_{1.83}$	$1.10_{1.53}$	93.920.50	$10.89_{2.64}$	$90.92_{0.41}$	$5.20_{0.50}$
	SST-2	AddSent	96.140.19	$7.53_{0.74}$	$93.94_{0.79}$	$95.88_{0.39}$	$0.00_{0.00}$	$50.08_{0.00}$	$43.86_{0.79}$	96.190.46	$0.18_{0.17}$	$57.99_{9.74}$	$35.96_{9.20}$	$94.95_{0.11}$	$7.78_{2.74}$	$91.23_{1.13}$	$2.71_{0.39}$
		HiddenKiller		$19.26_{3.64}$	$78.99_{0.26}$	$95.46_{0.68}$	$0.00_{0.00}$	$50.10_{0.03}$	$28.89_{0.26}$	96.660.23	$5.96_{5.02}$	$66.89_{13.86}$	$12.10_{14.02}$	94.630.67	$4.02_{6.87}$	$58.59_{13.28}$	$20.39_{13.19}$
		BadNets		$5.79_{0.29}$	$95.32_{0.08}$	92.352.83	$7.56_{3.53}$	88.224.58	$7.09_{4.54}$	93.212.29	$14.18_{2.29}$	$91.63_{1.00}$	$3.69_{1.08}$	89.951.52	$15.36_{1.21}$	89.941.27	$5.38_{1.19}$
Llama2	HSOL	AddSent	95.690.11	$5.36_{0.09}$	$95.53_{0.04}$	$91.52_{0.50}$	$7.59_{0.33}$	$57.07_{5.86}$	$38.46_{5.90}$	91.760.13	$10.14_{2.52}$	$78.66_{23.68}$	$16.87_{23.65}$	$90.01_{0.27}$	$15.31_{1.02}$	$90.42_{0.26}$	$5.11_{0.25}$
		HiddenKiller		$48.99_{0.45}$	$74.40_{0.19}$	$91.19_{0.33}$	$0.08_{0.00}$	$50.03_{0.02}$	$24.36_{0.17}$	$91.75_{0.11}$	$10.89_{10.20}$	$60.17_{9.20}$	$14.22_{9.22}$	$89.71_{0.46}$	$52.18_{0.74}$	$72.35_{0.69}$	$2.05_{0.67}$
		BadNets		$10.53_{0.96}$	$89.70_{0.85}$	$90.60_{0.17}$	$16.98_{15.32}$	$65.26_{14.19}$	$24.43_{14.93}$	$91.43_{0.35}$	$9.82_{0.56}$	$89.93_{0.55}$	$0.50_{0.17}$	$88.70_{0.44}$	$14.49_{1.00}$	$86.17_{0.74}$	$3.53_{1.29}$
	AG	AddSent	91.170.67	$10.09_{1.21}$	$90.27_{0.95}$	91.130.60	$28.31_{33.32}$	$54.83_{25.84}$	$35.43_{24.91}$	92.270.25	$16.38_{3.61}$	$65.20_{33.67}$	$25.07_{32.73}$	89.400.62	$12.80_{1.14}$	$86.93_{1.11}$	$3.33_{0.85}$
		HiddenKiller		$23.07_{1.44}$	$78.93_{1.17}$	$90.70_{0.50}$	$48.09_{15.98}$	$38.93_{11.99}$	$40.00_{12.79}$	91.330.98	$20.00_{5.46}$	$71.17_{5.18}$	$7.77_{5.94}$	89.370.23	$23.78_{1.01}$	$77.40_{0.60}$	$1.53_{0.64}$

**Evaluation Metrics.** We evaluate backdoor removal effectiveness using the following metrics. (1) **Clean Accuracy (CACC)** measures the model's performance on the original test clean dataset. (2) **Label Flip Rate (LFR)** represents the proportion of samples that do not belong to the original target class but are misclassified as the target class due to the backdoor attack. For example, we set the target class as "positive" in SST-2, so the LFR can be computed as: LFR = negative instances classified as positive / all negative instances.

Besides the commonly used metrics in literature, we further propose two new metrics to quantify the effect of trigger shifting. (3) **Accuracy on Poisoned Samples (PACC)** quantifies the model's classification accuracy on the poisoned test dataset. Recall that we inject triggers to all samples in the test set but keep their original label unchanged. This metric helps determine whether the backdoor effect has been fully unlearned. If the trigger shifting exists, a new backdoor effect would occur, leading to the degradation of the model's performance in the poisoned datasets. A higher PACC indicates that the model remains unaffected by triggers. (4) **Accuracy Difference on Poisoned Samples** ( $\Delta$ PACC) quantifies the absolute difference of PACC achieved between the ReTrain model and any other unlearned model. Since ReTrain represents a truly backdoor-free model, an effective backdoor unlearning method should have a PACC similar to the ReTrain model, indicating that the unlearned model closely approximates the backdoor-free state. Given a backdoor-unlearned model,  $\Delta$ PACC can be computed by  $\Delta$ PACC = |PACC<sub>ReTrain</sub> - PACC<sub>model</sub>|.

Implementation Details. We first perform three backdoor attacks to obtain the poisoned model  $f_{\theta_p}$  by fine-tuning DistilBert<sub>BASE</sub> and BERT<sub>BASE</sub> on the poisoned datasets for 5 epochs and Llama2 (7B) for 10 epochs. All model parameters are fine-tuned using a batch size of 32 and a maximum input length of 128. We use a learning rate of 2e-5 for DistilBert<sub>BASE</sub> and BERT<sub>BASE</sub>, and 5e-6 for Llama2 (7B), optimized with Adam (Kingma & Ba, 2014). Following the same experimental settings in the original paper, we then apply CUBE (Cui et al., 2022) to identify poisoned samples. To explore the influence of gradient ascent, we perform backdoor unlearning on the poisoned model  $f_{\theta_p}$  using the detected poisoned samples for 30 epochs. For RGA, we set  $\alpha=2$  and  $\beta=0.05$  across all models. All experiments are conducted using four NVIDIA RTX 6000 Ada GPUs.

#### 6.2 EXPERIMENTAL RESULTS

Classification Results after Backdoor Unlearning on Test Datasets. After acquiring poisoned samples, we conduct experiments to unlearn various backdoor effects in poisoned models. Table 2 presents the unlearning results against different backdoor attacks for 30 epochs, including BadNet, AddSent, and HiddenKiller. Generally, the ReTrain model is unaffected by backdoor triggers and has a similar performance on clean and poisoned datasets, i.e., CACC and PACC are close. Note that the slightly lower PACC compared with CACC is because some attack strategies induce a loss

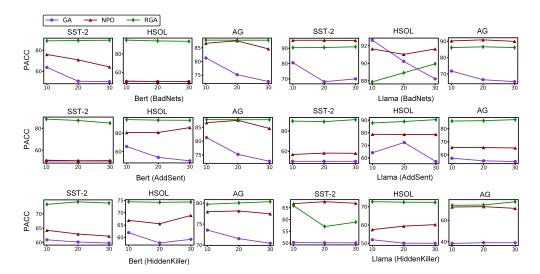


Figure 3: Change in PACC on the test set during backdoor unlearning of the BadNets, AddSent, and HiddenKiller-poisoned models at the 10th, 20th, and 30th epochs.

of semantic integrity when transforming clean text into its poisoned counterpart, leading to higher misclassification rates.

An ideal backdoor unlearning method should achieve a PACC similar to that of ReTrain while ensuring minimal degradation in model utility on clean tasks. Our experiments in Table 2 reveal that GA and NPO can significantly reduce the LFR but compromise the PACC on BERT<sub>BASE</sub>, DistilBert<sub>BASE</sub> and Llama2 on both binary and multi-classification tasks. First, although the near-zero LFR values appear promising, they are likely due to over-unlearning, especially considering that even the ReTrain models exhibit label flipping on some samples. Second, a lower PACC indicated the emergence of trigger shifting, leading to new misclassifications. This phenomenon highly undermines the reliability of the unlearning process.

In contrast, RGA can maintain the highest PACC and achieve the lowest  $\Delta$ PACC compared to GA and NPO without significantly degrading the model's utility on the clean task (CACC) in most cases. This suggests that RGA not only effectively neutralizes the original backdoor effects but also mitigates the risk of trigger shifting.

Change of PACC during Unlearning Process. Figure 3 illustrates the variation in PACC on the test datasets of AddSent-poisoned models after 10, 20, and 30 epochs of the unlearning process. We observe that RGA consistently maintains the highest PACC with superior stability compared to GA and NPO in most cases, suggesting that, in practical scenarios, we can specify a larger number of unlearning epochs without the risk of over-unlearning or trigger shifting. These observations indicate that RGA is a promising approach for backdoor unlearning. We include the results of the unlearning process on the poisoned DistilBert model in the Appendix B.3.

Change of Poisoned Losses during Unlearning Process. We further investigate the change of the cross-entropy loss between predicted class  $f_{\theta}(y_p|x_p)$  and the target class  $y_p$  throughout the unlearning process for GA, NPO, and RGA. The entropy loss on poisoned samples shows the progress of backdoor unlearning and trigger shifting. A low poisoned loss indicates that the model still associates the trigger t with the target class  $y_p$ , suggesting insufficient unlearning. However, if the poisoned loss diverges to infinity, trigger shifting occurs, introducing a new security risk. Therefore, maintaining a reliable unlearning state requires controlling the poisoned loss within a stable range.

Figure 4 shows the poisoned loss of the first 10 epochs of unlearning over various attacks, with additional unlearning results in the Appendix B.3 and B.5. We can observe that GA quickly diverges, leading to the trigger shifting. Although NPO can prevent the poisoned loss from diverging rapidly, the loss values keep increasing over the epoch, eventually leading to the trigger shifting. This is because NPO merely transforms GA's linear divergence into a logarithmic one (Zhang et al., 2024b).

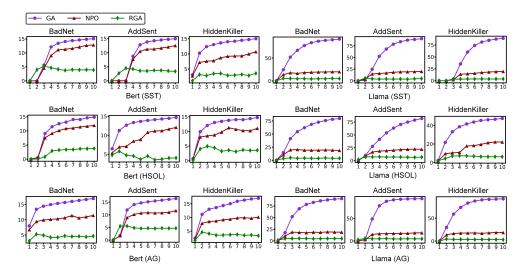


Figure 4: The evolution of the cross-entropy loss on poisoned samples during the unlearning of models compromised by different attacks over 10 epochs.

Table 3: Backdoor unlearning performance of RGA with different  $\alpha$  against different attacks on poisoned Llama2 (7B))

Datasat	Attack		α	= 1			α	= 2			α	= 3			α	= 4	
Dataset	Attack	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC
	BadNets	95.22	5.04	92.37	3.95	95.17	6.25	92.92	3.40	95.17	6.25	92.86	3.46	95.22	6.25	92.92	3.40
SST-2	AddSent	94.23	7.24	92.31	2.09	94.62	10.64	91.82	2.58	94.62	13.27	91.16	3.24	94.56	15.46	90.23	4.17
	HiddenKiller	94.95	0.00	51.40	27.79	94.89	2.74	62.71	16.48	94.78	2.96	62.98	16.21	94.95	5.59	66.17	13.02
	BadNets	93.24	14.08	91.23	4.18	92.88	16.49	90.26	5.15	92.72	17.94	89.58	5.83	92.52	18.26	89.34	6.07
HSOL	AddSent	87.97	15.50	88.57	6.96	88.04	16.57	88.33	7.30	88.05	17.22	88.29	7.24	88.12	17.37	88.20	7.33
	HiddenKiller	88.37	46.25	74.04	0.57	88.25	53.10	72.23	2.38	88.21	56.39	70.82	3.79	88.29	54.71	71.26	3.35
	BadNets	88.80	15.47	85.60	3.30	88.90	15.33	86.00	2.90	89.10	15.33	86.10	2.80	89.00	15.60	85.80	3.10
AG	AddSent	88.90	14.13	86.50	2.70	89.10	17.60	84.20	5.00	88.90	17.47	84.50	4.70	89.00	18.00	84.20	5.00
	HiddenKiller	88.90	24.00	76.00	1.60	88.70	33.33	72.00	5.60	88.60	39.60	68.20	9.40	88.80	44.00	65.10	12.50

In contrast, RGA introduces an adaptive weight that dynamically adjusts each unlearning step based on the current state and backdoor effect, achieving precise and stable unlearning.

Sensitivity Analysis of the Hyperparameter  $\alpha$  in Eq. 7. We study the impact of  $\alpha$  in RGA for unlearning poisoned Llama2 (7B) for 10 epochs, noting that RGA reduces to GA when  $\alpha=0$ . We test  $\alpha=\{1,2,3,4\}$ . As shown in Table 3, RGA has low sensitivity to  $\alpha$ , maintaining model utility, strong unlearning performance, and mitigating trigger shifting in most cases. A small  $\alpha$  (e.g.,  $\alpha=1$ ) causes a slow decay of  $\lambda$  and leads to the trigger shifting, with  $\Delta$ PACC reaching 27.79 for the HiddenKiller attack on SST-2. Increasing  $\alpha$  accelerates  $\lambda$ 's decay and thus mitigates trigger shifting ( $\Delta$ PACC drops to 13.02 at  $\alpha=4$ ) during unlearning, although LFR increases slightly. This suggests using more unlearning epochs when  $\alpha$  is large.

In the appendix B, we provide detailed experimental results, including backdoor attack results, poisoned sample detection, ablation studies, the evolution of PACC and poison loss during unlearning across various datasets and attacks, and training time comparison.

#### 7 Conclusions

In this work, we have identified trigger shifting as a critical and underexplored flaw in vanilla GA-based backdoor unlearning. Specifically, we show that GA does not necessarily eliminate the backdoor effect but can instead redirect it to a new backdoor effect—thereby compromising the reliability of the unlearning process. To address this, we have developed Robust Gradient Ascent (RGA), which introduces a dynamic penalty mechanism to prevent unintended trigger shifting while preserving model utility. Our experimental results demonstrate that RGA effectively removes backdoors without causing trigger shifting, highlighting the need for more reliable unlearning techniques in securing LLMs.

# REPRODUCIBILITY STATEMENT

We aims to make our results fully reproducible. An anonymous code release is included in the supplementary material, containing training and evaluation scripts. We apply the public datasets, which can be easily acquired from Huggingface. We also provide them in the supplementary material. All propositions and complete proofs are provided in the Appendix A.

# REFERENCES

- Sungmin Cha, Sungjun Cho, Dasol Hwang, and Moontae Lee. Towards robust and parameter-efficient knowledge unlearning for llms. *arXiv* preprint arXiv:2408.06621, 2024.
- Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv* preprint arXiv:2310.20150, 2023.
- Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai Wu, and Yang Zhang. Badnl: Backdoor attacks against nlp models with semantic-preserving improvements. In *Proceedings of the 37th Annual Computer Security Applications Conference*, ACSAC '21, pp. 554–569, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450385794. doi: 10.1145/3485832.3485837. URL https://doi.org/10.1145/3485832.3485837.
- Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- Ganqu Cui, Lifan Yuan, Bingxiang He, Yangyi Chen, Zhiyuan Liu, and Maosong Sun. A unified evaluation of textual backdoor learning: Frameworks and benchmarks. In *Proceedings of NeurIPS: Datasets and Benchmarks*, 2022.
- Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878, 2019.
- Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pp. 512–515, 2017.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *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)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. URL https://aclanthology.org/N19-1423.
- Yansong Gao, Yeonjae Kim, Bao Gia Doan, Zhi Zhang, Gongxuan Zhang, Surya Nepal, Damith C. Ranasinghe, and Hyoungshick Kim. Design and evaluation of a multi-domain trojan detection method on deep neural networks. *IEEE Transactions on Dependable and Secure Computing*, 19 (4):2349–2364, 2022. doi: 10.1109/TDSC.2021.3055844.
- Shashwat Goel, Ameya Prabhu, Philip Torr, Ponnurangam Kumaraguru, and Amartya Sanyal. Corrective machine unlearning. *arXiv preprint arXiv:2402.14015*, 2024.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. *arXiv* preprint arXiv:2210.01504, 2022.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Keita Kurita, Paul Michel, and Graham Neubig. Weight poisoning attacks on pretrained models. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 2793–2806, Online, July 2020. Association for Computational Linguistics. URL https://aclanthology.org/2020.acl-main.249.

Jiazhao Li, Zhuofeng Wu, Wei Ping, Chaowei Xiao, and V.G.Vinod Vydiswaran. Defending against insertion-based textual backdoor attacks via attribution. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8818–8833, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.561. URL https://aclanthology.org/2023.findings-acl.561/.

- Linyang Li, Demin Song, Xiaonan Li, Jiehang Zeng, Ruotian Ma, and Xipeng Qiu. Backdoor attacks on pre-trained models by layerwise weight poisoning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3023–3032, Online and Punta Cana, Dominican Republic, November 2021a. Association for Computational Linguistics. URL https://aclanthology.org/2021.emnlp-main.241.
- Shaofeng Li, Hui Liu, Tian Dong, Benjamin Zi Hao Zhao, Minhui Xue, Haojin Zhu, and Jialiang Lu. Hidden backdoors in human-centric language models. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, CCS '21, pp. 3123–3140, New York, NY, USA, 2021b. Association for Computing Machinery. ISBN 9781450384544. doi: 10.1145/3460120.3484576. URL https://doi.org/10.1145/3460120.3484576.
- Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Anti-backdoor learning: Training clean models on poisoned data. In *NeurIPS*, 2021c. URL https://proceedings.neurips.cc/paper/2021/file/7d38b1e9bd793d3f45e0e212a729a93c-Paper.pdf.
- Yang Liu, Mingyuan Fan, Cen Chen, Ximeng Liu, Zhuo Ma, Li Wang, and Jianfeng Ma. Backdoor defense with machine unlearning. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, pp. 280–289, 2022. doi: 10.1109/INFOCOM48880.2022.9796974.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C Lipton, and J Zico Kolter. Tofu: A task of fictitious unlearning for llms. *arXiv preprint arXiv:2401.06121*, 2024.
- Leland McInnes and John Healy. Accelerated hierarchical density based clustering. In 2017 IEEE international conference on data mining workshops (ICDMW), pp. 33–42. IEEE, 2017.
- Xudong Pan, Mi Zhang, Beina Sheng, Jiaming Zhu, and Min Yang. Hidden trigger backdoor attack on NLP models via linguistic style manipulation. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 3611–3628, Boston, MA, August 2022. USENIX Association. ISBN 978-1-939133-31-1.
- Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. ONION: A simple and effective defense against textual backdoor attacks. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 9558–9566, Online and Punta Cana, Dominican Republic, November 2021a. Association for Computational Linguistics. URL https://aclanthology.org/2021.emnlp-main.752.
- Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. Mind the style of text! adversarial and backdoor attacks based on text style transfer. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 4569–4580, Online and Punta Cana, Dominican Republic, November 2021b. Association for Computational Linguistics. URL https://aclanthology.org/2021.emnlp-main.374.
- Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *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)*, pp. 443–453, Online, August 2021c. Association for Computational Linguistics. URL https://aclanthology.org/2021.acl-long.37.

Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, and Maosong Sun. Turn the combination lock: Learnable textual backdoor attacks via word substitution. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *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)*, pp. 4873–4883, Online, August 2021d. Association for Computational Linguistics. URL https://aclanthology.org/2021.acl-long.377.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Tim Sainburg, Leland McInnes, and Timothy Q Gentner. Parametric umap embeddings for representation and semisupervised learning. *Neural Computation*, 33(11):2881–2907, 2021.
- Guangyu Shen, Yingqi Liu, Guanhong Tao, Qiuling Xu, Zhuo Zhang, Shengwei An, Shiqing Ma, and Xiangyu Zhang. Constrained optimization with dynamic bound-scaling for effective nlp backdoor defense. In *International Conference on Machine Learning*, pp. 19879–19892. PMLR, 2022. URL https://proceedings.mlr.press/v162/shen22e.html.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard (eds.), *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL https://aclanthology.org/D13-1170.
- Haosen Sun, Yiming Li, Xixiang Lyu, and Jing Ma. Learning from distinction: Mitigating backdoors using a low-capacity model. In *Proceedings of the 32nd ACM International Conference on Multimedia*, MM '24, pp. 9077–9086, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400706868. doi: 10.1145/3664647.3680968. URL https://doi.org/10.1145/3664647.3680968.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y. Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In 2019 IEEE Symposium on Security and Privacy (SP), pp. 707–723, 2019. doi: 10.1109/SP.2019.00031.
- Jiali Wei, Ming Fan, Wenjing Jiao, Wuxia Jin, and Ting Liu. Bdmmt: Backdoor sample detection for language models through model mutation testing. *IEEE Transactions on Information Forensics* and Security, 19:4285–4300, 2024.
- Jiashu Xu, Mingyu Ma, Fei Wang, Chaowei Xiao, and Muhao Chen. Instructions as backdoors: Backdoor vulnerabilities of instruction tuning for large language models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 3111–3126, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.171. URL https://aclanthology.org/2024.naacl-long.171/.
- Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun, and Bin He. Be careful about poisoned word embeddings: Exploring the vulnerability of the embedding layers in NLP models. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2048–2058, Online, June 2021a. Association for Computational Linguistics. URL https://aclanthology.org/2021.naacl-main.165.
- Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. RAP: Robustness-Aware Perturbations for defending against backdoor attacks on NLP models. In Marie-Francine Moens,

 Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 8365–8381, Online and Punta Cana, Dominican Republic, November 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.659. URL https://aclanthology.org/2021.emnlp-main.659.

- Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. Rethinking stealthiness of backdoor attack against NLP models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), 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), pp. 5543–5557, Online, August 2021c. Association for Computational Linguistics. URL https://aclanthology.org/2021.acl-long.431.
- Jin Yao, Eli Chien, Minxin Du, Xinyao Niu, Tianhao Wang, Zezhou Cheng, and Xiang Yue. Machine unlearning of pre-trained large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8403–8419, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.457. URL https://aclanthology.org/2024.acl-long.457/.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. arXiv preprint arXiv:2310.10683, 2023.
- Biao Yi, Sishuo Chen, Yiming Li, Tong Li, Baolei Zhang, and Zheli Liu. BadActs: A universal backdoor defense in the activation space. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 5339–5352, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.317. URL https://aclanthology.org/2024.findings-acl.317/.
- Xiaojian Yuan, Tianyu Pang, Chao Du, Kejiang Chen, Weiming Zhang, and Min Lin. A closer look at machine unlearning for large language models. *arXiv preprint arXiv:2410.08109*, 2024.
- Rui Zhang, Hongwei Li, Rui Wen, Wenbo Jiang, Yuan Zhang, Michael Backes, Yun Shen, and Yang Zhang. Instruction backdoor attacks against customized {LLMs}. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 1849–1866, 2024a.
- Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024b.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.
- Xinyang Zhang, Zheng Zhang, Shouling Ji, and Ting Wang. Trojaning language models for fun and profit. In 2021 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 179–197. IEEE, 2021.
- Zhiyuan Zhang, Lingjuan Lyu, Xingjun Ma, Chenguang Wang, and Xu Sun. Fine-mixing: Mitigating backdoors in fine-tuned language models. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), Findings of the Association for Computational Linguistics: EMNLP 2022, pp. 355–372, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-emnlp.26. URL https://aclanthology.org/2022.findings-emnlp.26/.
- Zhiyuan Zhang, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. Diffusion theory as a scalpel: Detecting and purifying poisonous dimensions in pre-trained language models caused by backdoor or bias. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 2495–2517, Toronto, Canada, July 2023. Association for Computational Linguistics. URL https://aclanthology.org/2023.findings-acl.157.

# THEORETICAL ANALYSIS ON TRIGGER SHIFTING

#### A.1 **PROOF OF PROPOSITION 1**

**Proposition 1.** Given a poisoned model  $f_{\theta_p}$  trained on  $\mathcal{D}$ , the objective function of the "unlearned" model  $f_{\theta_{n^*}}$  in binary classification is defined as:

$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{p^*}}(y_c | x_c), y_c)] - \mathbb{E}_{(x_0 \oplus t, y_1) \sim \mathcal{D}_p} [\ell(f_{\theta_{p^*}}(y_1 | x_0 \oplus t), y_1))],$$
(8) which is equivalent to minimizing the following objective function:

(9)

 $\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} [\ell(f_{\theta_{p^*}}(y_c | x_c), y_c)] + \mathbb{E}_{(x_0 \oplus t, y_0) \sim \mathcal{D}_p} [\ell(f_{\theta_{p^*}}(y_0 | x_0 \oplus t), y_0)] + R(\theta_{p^*}),$ 

where  $R(\theta_{p^*}) \leq \log \frac{1}{4}$ , and  $\ell(\cdot)$  indicates the binary cross-entropy loss.

*Proof.* Let  $p_1 := p_{\theta_{n^*}}(y_1 \mid x)$  and  $p_0 := 1 - p_1$ . As  $\ell$  indicates the binary cross-entropy, we have:

$$\ell(f_{\theta_{n^*}}(y_1 \mid x), y_1) = -\log p_1, \qquad \ell(f_{\theta_{n^*}}(y_0 \mid x), y_0) = -\log p_0.$$

For any triggered input  $x_0 \oplus t$ ,

$$-\ell(f_{\theta_{p^*}}(y_1 \mid x_0 \oplus t), y_1) = \log p_1$$

$$= -\log p_0 + \log(p_0 p_1)$$

$$= \ell(f_{\theta_{p^*}}(y_0 \mid x_0 \oplus t), y_0) + \log(p_0 p_1).$$

Substituting into equation 8 gives

ituting into equation 8 gives 
$$\mathcal{L}_{p^*} = \mathbb{E}_{\mathcal{D}_c} \left[ \ell(f_{\theta_{p^*}}(y_c \mid x_c), y_c) \right] + \mathbb{E}_{\mathcal{D}_p} \left[ \ell(f_{\theta_{p^*}}(y_0 \mid x_0 \oplus t), y_0) \right] + \underbrace{\mathbb{E}_{\mathcal{D}_p} \left[ \log(p_0 p_1) \right]}_{=:R(\theta_{p^*})}.$$

Since  $p_0 + p_1 = 1$  and  $p_0, p_1 \in (0, 1)$ , we have  $p_0 p_1 \leq \frac{1}{4}$  with equality at  $p_0 = p_1 = \frac{1}{2}$ . Hence

$$R(\theta_{p^*}) = \mathbb{E}_{\mathcal{D}_p} [\log(p_0 p_1)] \le \log \frac{1}{\Lambda}.$$

Thus equation 8 equals equation 9 with  $R(\theta_{p^*}) \leq \log \frac{1}{4}$ , completing the proof. 

#### TRIGGER SHIFTING IN THE MULTICLASS CLASSIFICATION TASK

The trigger shifting in the binary classification scenario can also be observed in the multiclass classification case.

**Proposition 2.** Let  $f_{\theta_p}$  be a poisoned model with softmax probabilities  $p_k(x) = p_{\theta_p}(y_k \mid x)$  for  $k \in \{1, \dots, K\}$ . Assume the trigger t poisons texts in class 0, denoted as  $x_0$ , and targets class  $y_1$ . The unlearning objective is defined as:

$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} \left[ \ell(f_{\theta_{p^*}}(y_c \mid x_c), y_c) \right] - \mathbb{E}_{(x_0 \oplus t, y_1) \sim \mathcal{D}_p} \left[ \ell(f_{\theta_{p^*}}(y_1 \mid x_0 \oplus t), y_1) \right], \tag{10}$$
 is equivalent to

$$\mathcal{L}_{p^*} = \mathbb{E}_{(x_c, y_c) \sim \mathcal{D}_c} \left[ \ell(f_{\theta_{p^*}}(y_c \mid x_c), y_c) \right] + \mathbb{E}_{(x_0 \oplus t \sim \mathcal{D}_p)} \left[ \sum_{k \neq 1} \ell \left( f_{\theta_{p^*}}(y_k \mid x_0 \oplus t), y_k \right) \right] + R(\theta_{p^*}),$$
(11)

where

$$R(\theta_{p^*}) = \mathbb{E}_{(x_0 \oplus t \sim \mathcal{D}_p)} \left[ \log \left( \prod_{k=1}^K p_k(x_0 \oplus t) \right) \right] \le K \log \frac{1}{K} = -K \log K.$$
 (12)

*Proof.* For multiclass cross-entropy, we have  $\ell(f_{\theta_{p^*}}(y_k \mid x), y_k) = -\log p_k(x)$ . Given a triggered input  $x := x_0 \oplus t$ ,

$$\begin{split} -\ell & \big( f_{\theta_{p^*}}(y_1 \mid x), y_1 \big) = \log p_1(x) \\ & = -\sum_{k \neq 1} \log p_k(x) + \log \Big( \prod_{k=1}^K p_k(x) \Big) \\ & = \sum_{k \neq 1} \ell \big( f_{\theta_{p^*}}(y_k \mid x), y_k \big) + \log \Big( \prod_{k=1}^K p_k(x) \Big). \end{split}$$

Substituting into equation 10 gives equation 11 with  $R(\theta_{p^*}) = \mathbb{E}[\log \prod_k p_k]$ . Over the probability simplex,  $\prod_{k=1}^K p_k$  is maximized at the uniform point  $p_k = \frac{1}{K}$ , hence  $\prod_k p_k \leq K^{-K}$ .

Minimizing the confidence of the poisoned model in predicting the target class of triggered samples would redistribute the probability mass over the remaining classes, Eq. 11. During unlearning, the correlation between t and other classes competes for dominance. Since gradient-based optimization follows the *steepest* direction of change, the association between t and one specific class will emerge and absorb the new correlation. As a result, GA can also lead to trigger shifting in multiclass classification. Importantly, the extra term  $R(\theta_{p^*})$  is upper bounded above by  $-K\log K$ , so it cannot prevent the shift.

#### B EXPERIMENTAL RESULTS

#### B.1 BACKDOOR ATTACK

We introduce three backdoor attacks, including two static attacks, BadNets and AddSent, and one dynamic attack, HiddenKiller. Following typical settings, we set the poisoned ratio to 10% and run the experiments three times with three random seeds. Attack results are shown in Table 4.

Table 4: Backdoor Attack Results

Dataset	Attack		SST-2			HSOL			AG	
Dataset	Attack	CACC	LFR	PACC	CACC	LFR	PACC	CACC	LFR	PACC
	BadNets	89.731.29	$100.00_{0.00}$	$49.92_{0.00}$	$95.18_{0.02}$	$100.00_{0.00}$	$49.98_{0.00}$	$90.37_{0.42}$	$99.87_{0.00}$	$25.10_{0.00}$
BERT	AddSent	$89.55_{1.28}$	$100.00_{0.00}$	$49.92_{0.00}$	$94.69_{0.29}$	$100.00_{0.00}$	$49.98_{0.00}$	$89.30_{1.39}$	$100.00_{0.00}$	$25.00_{0.00}$
	HiddenKiller	$91.21_{0.91}$	$92.94_{0.89}$	$52.94_{0.25}$	$94.61_{0.31}$	$98.82_{0.96}$	$50.46_{0.47}$	89.631.01	$96.67_{1.35}$	$27.50_{1.01}$
	BadNets	$89.05_{1.24}$	$100.00_{0.00}$	$49.92_{0.00}$	$94.35_{0.39}$	$100.00_{0.00}$	$49.98_{0.00}$	88.57 <sub>0.55</sub>	$99.87_{0.00}$	$25.10_{0.00}$
DistilBert	AddSent	89.480.09	$100.00_{0.00}$	$49.92_{0.00}$	$94.02_{0.31}$	$100.00_{0.00}$	$49.98_{0.00}$	87.800.92	$100.00_{0.00}$	$25.00_{0.00}$
	HiddenKiller	88.67 <sub>0.40</sub>	$96.24_{1.21}$	$51.58_{0.50}$	$94.54_{0.19}$	$99.60_{0.32}$	$50.15_{0.14}$	88.67 <sub>0.32</sub>	$97.33_{0.81}$	$27.00_{0.61}$
	BadNets	$96.14_{0.47}$	$99.12_{0.11}$	$50.30_{0.06}$	$95.36_{0.23}$	$99.44_{0.08}$	$50.23_{0.06}$	$91.60_{0.35}$	$99.06_{0.14}$	$25.70_{0.10}$
Llama2	AddSent	$96.27_{0.31}$	$99.93_{0.13}$	$49.96_{0.06}$	$95.61_{0.14}$	$100.00_{0.00}$	$49.98_{0.00}$	$91.70_{0.46}$	$99.42_{0.27}$	$25.43_{0.20}$
	HiddenKiller	$96.39_{0.50}$	$99.96_{0.06}$	$49.92_{0.00}$	$95.41_{0.07}$	$99.95_{0.09}$	$50.01_{0.05}$	$91.90_{0.50}$	$99.47_{0.00}$	$25.40_{0.00}$

#### B.2 BACKDOOR SAMPLE DETECTION

In our paper, we apply CUBE (Cui et al., 2022), a clustering-based method for detecting back-door samples, to identify poisoned samples. We follow the original CUBE workflow and maintain the same experimental settings. (1) Representation learning. We first fine-tune the model on the poisoned dataset and use the poisoned model to project each training sample into the embedding space. For BERT<sub>BASE</sub> and DistilBERT<sub>BASE</sub>, we use [CLS] as the sample embedding, while for Llama2 (7B), we use the last token's hidden state as its representation. (2) Clustering. With all sample embeddings collected, we apply UMAP (Sainburg et al., 2021) to reduce the dimensionality to 4-D, and then use the density-based clustering algorithm HDBSCAN (McInnes & Healy, 2017) to identify distinctive clusters. (3) Filtering. Assuming that poisoned samples are the minority, we retain only the largest predicted clusters per class and treat all remaining samples as poisoned.

Table 5: Performance of CUBE on Backdoor Sample Detection

Dataset	Attack		SST-2			HSOL			AG	
Dataset	Attack	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
	BadNets	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$0.97_{0.01}$	$1.00_{0.00}$	$0.98_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$
BERT	AddSent	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$0.99_{0.01}$	$1.00_{0.00}$	$0.99_{0.00}$	$0.94_{0.10}$	$0.99_{0.00}$	$0.96_{0.06}$
	HiddenKiller	$0.91_{0.04}$	$0.86_{0.02}$	$0.88_{0.03}$	$0.95_{0.02}$	$1.00_{0.00}$	$0.97_{0.01}$	$0.88_{0.09}$	$1.00_{0.00}$	$0.93_{0.05}$
	BadNets	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$0.98_{0.02}$	$1.00_{0.00}$	$0.99_{0.01}$	$0.99_{0.01}$	$1.00_{0.00}$	$1.00_{0.00}$
DistilBert	AddSent	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$0.97_{0.02}$	$1.00_{0.00}$	$0.99_{0.01}$	$0.96_{0.06}$	$1.00_{0.00}$	$0.98_{0.03}$
	HiddenKiller	$0.91_{0.02}$	$0.90_{0.02}$	$0.91_{0.01}$	$0.96_{0.01}$	$1.00_{0.00}$	$0.98_{0.00}$	$0.96_{0.01}$	$1.00_{0.00}$	$0.98_{0.00}$
	BadNets	$0.98_{0.01}$	$0.98_{0.04}$	$0.98_{0.02}$	$0.61_{0.27}$	$0.86_{0.05}$	$0.69_{0.15}$	$0.98_{0.02}$	$0.99_{0.01}$	$0.99_{0.01}$
Llama2	AddSent	$0.99_{0.01}$	$1.00_{0.00}$	$0.99_{0.01}$	$0.48_{0.00}$	$0.91_{0.00}$	$0.62_{0.00}$	$0.99_{0.01}$	$0.97_{0.02}$	$0.98_{0.01}$
	HiddenKiller	$1.00_{0.00}$	$1.00_{0.00}$	$1.00_{0.00}$	$0.49_{0.01}$	$1.00_{0.00}$	$0.66_{0.01}$	$0.92_{0.09}$	$1.00_{0.00}$	$0.96_{0.05}$

CUBE achieves high F1 scores in most cases. Its effectiveness stems from the observation that poisoned samples tend to cluster together, as they share a common trigger pattern. However, we observe

that CUBE struggles to detect samples in the HSOL dataset when applied to Llama2 (7B). Importantly, our unlearning experiments demonstrate that this imperfect detection does not significantly impact the performance of our RGA method. Detection results are presented in Table 5.

#### B.3 CHANGE OF PACC AND POISONED LOSSES

We also present the change of PACC and poisoned losses when unlearning the BadNet, AddSent, and HiddenKiller on the poisoned DistilBert in Figure 5. We can observe that RGA maintains the highest PACC during the unlearning process in most cases, which indicates that RGA would not introduce a new backdoor effect even with a large unlearning epoch. Notably, RGA always maintains the poisoned loss within a stable range during unlearning.

#### **B.4** ABLATION STUDY

We conducted an ablation study to evaluate the individual contributions of the backdoor unlearning term (term i) and the regularization term (term iii) within the RGA framework. Specifically, we examine two variants:  $\underline{D}$ ynamic  $\underline{G}$ radient  $\underline{A}$ scent (DGA), which combines (term i) and (term ii), and the full Robust Gradient Ascent (RGA), which integrates all three terms (term  $i \oplus ii \oplus iii$ ). To assess their effectiveness, we perform unlearning on poisoned LLaMA2 (7B) models for 10 epochs across three types of attacks—BadNets, AddSent, and HiddenKiller—using three different random seeds. The average performance is reported in Table 6.

Table 6: Backdoor unlearning of DGA and RGA against different attacks on poisoned Llama2 (7B).

Datasat	Attack		D	GA			R	GA	
Dataset	Attack	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC
	BadNets	$96.19_{0.36}$	$14.14_{2.20}$	$91.95_{0.77}$	$4.17_{1.06}$	$94.44_{0.72}$	$13.60_{6.38}$	$90.36_{2.24}$	$5.76_{2.05}$
SST-2	AddSent	$96.34_{0.21}$	$11.04_{17.15}$	$67.77_{15.81}$	$26.18_{15.04}$	$94.40_{0.20}$	$12.46_{5.50}$	$89.79_{1.89}$	$4.16_{1.94}$
	HiddenKiller	$96.65_{0.29}$	$17.87_{3.55}$	$77.41_{3.40}$	$2.23_{3.03}$	$94.52_{0.32}$	$6.47_{4.92}$	$65.82_{9.91}$	$13.16_{10.11}$
	BadNets	$93.17_{2.08}$	$23.57_{2.34}$	$87.46_{1.05}$	$7.86_{1.13}$	88.933.42	$18.85_{2.06}$	87.882.09	$7.43_{2.01}$
HSOL	AddSent	$91.84_{0.12}$	$19.44_{9.94}$	$76.98_{17.15}$	$18.55_{17.11}$	$87.90_{0.19}$	$16.44_{1.81}$	$87.79_{0.90}$	$7.74_{0.93}$
	HiddenKiller	$91.84_{0.29}$	$50.72_{0.66}$	$70.81_{5.08}$	$3.59_{5.04}$	$87.08_{1.38}$	$49.10_{4.38}$	$72.70_{0.41}$	$1.69_{0.60}$
	BadNets	$91.77_{0.32}$	$48.08_{44.05}$	$63.33_{32.05}$	$26.37_{33.29}$	$88.60_{0.30}$	$14.75_{1.00}$	$86.13_{0.91}$	$3.57_{1.51}$
AG	AddSent	$91.87_{0.15}$	$38.44_{19.33}$	$66.93_{19.58}$	$23.33_{18.66}$	$89.03_{0.06}$	$15.07_{2.24}$	$85.83_{1.44}$	$4.43_{0.49}$
	HiddenKiller	$92.03_{0.75}$	$39.73_{4.97}$	$68.03_{3.04}$	$10.90_{4.20}$	88.57 <sub>0.12</sub>	$27.95_{4.67}$	$74.03_{1.80}$	$4.90_{0.82}$

#### B.5 UNLEARNING PROCESS AT THE 10 AND 20 EPOCH

Unlearning results against backdoor attacks for 10 and 20 epochs are shown in Tables 7 and 8.

Table 7: Backdoor unlearning methods for 10 epochs against BadNets, AddSent, and HiddenKiller, targeting poisoned BERT, DistilBert, and Llama2 (7B).

				ReT				iA				NPO			R	GA	
Datas	et	Attack	CACC	LFR	PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC
		BadNets		$7.16_{2.18}$	$91.20_{1.08}$	$90.96_{0.88}$	$0.11_{0.19}$	$63.70_{17.01}$	$27.50_{18.06}$	$91.10_{1.10}$	$3.36_{4.00}$	$76.01_{20.99}$	$16.17_{20.60}$	$89.93_{0.45}$	$13.56_{5.63}$	$89.03_{0.91}$	$2.16_{1.40}$
	SST-2	AddSent	91.320.79	$13.45_{3.77}$	$89.4_{0.83}$	$90.64_{0.27}$	$0.00_{0.00}$	$50.08_{0.00}$	$39.32_{0.83}$	89.942.09	$0.00_{0.00}$	$50.81_{1.22}$	$38.59_{1.67}$	90.630.77	$11.66_{4.15}$	$88.45_{0.54}$	$0.95_{0.77}$
		HiddenKiller		$23.68_{3.88}$	$74.83_{0.66}$	$90.57_{0.43}$	$4.71_{3.27}$	$60.92_{8.49}$	$13.91_{9.15}$	$91.43_{0.53}$	$8.12_{2.48}$	$64.25_{9.08}$	$10.57_{9.73}$	$90.11_{0.96}$	$32.09_{6.45}$	$73.33_{0.66}$	$1.50_{0.27}$
		BadNets		$3.49_{0.86}$	$95.00_{0.14}$	$94.46_{0.59}$	$0.00_{0.00}$	$50.10_{0.07}$	$44.90_{0.16}$	$95.01_{0.04}$	$0.00_{0.00}$	$50.42_{0.70}$	$44.57_{0.83}$	$94.61_{0.66}$	$6.22_{0.90}$	$94.97_{0.24}$	$0.26_{0.16}$
BERT	HSOL	AddSent	$95.08_{0.14}$	$7.78_{1.87}$	$94.65_{0.48}$	$94.70_{0.48}$	$0.43_{0.45}$	$65.54_{15.74}$	$29.11_{15.60}$	$95.03_{0.44}$	$2.14_{1.86}$	$80.56_{22.00}$	$14.08_{22.08}$	$94.53_{0.17}$	$7.43_{1.59}$	$94.61_{0.41}$	$0.42_{0.30}$
		HiddenKiller		$47.39_{0.69}$	$74.77_{0.17}$	$94.18_{0.81}$	$5.01_{4.51}$	$61.88_{6.78}$	$12.89_{6.95}$	$93.68_{1.21}$	$11.08_{8.28}$	$66.87_{11.81}$	$7.90_{11.87}$	$94.33_{0.82}$	$46.31_{1.45}$	$74.45_{0.12}$	$0.32_{0.22}$
		Badnets		$10.93_{0.93}$	89.630.72	89.370.61	$8.17_{0.31}$	$85.30_{2.52}$	$4.33_{2.13}$	90.330.45	8.840.63	89.531.42	$0.50_{0.36}$	88.870.90	$14.40_{3.27}$	87.432.14	2.201.45
	AG	Addsent	89.370.70	$11.55_{0.95}$	$89.30_{0.36}$	89.070.64	$8.35_{0.31}$	$81.27_{6.73}$	$8.03_{6.37}$	89.171.19	$9.33_{0.69}$	86.673.67	$3.10_{2.75}$	$88.13_{0.93}$	$13.02_{1.93}$	$87.80_{0.70}$	$1.50_{0.82}$
		HiddenKiller		$21.64_{2.75}$	$78.26_{1.90}$	$89.67_{0.47}$	$17.73_{3.71}$	$73.57_{4.40}$	$4.70_{4.52}$	89.530.50	$17.95_{0.28}$	$77.93_{1.62}$	$1.60_{0.79}$	$88.50_{0.98}$	$20.35_{4.07}$	$79.73_{1.16}$	$2.60_{1.39}$
		BadNets		5.883.00	88.621.85	89.440.39	$0.04_{0.06}$	$57.18_{8.21}$	$31.43_{8.16}$	89.550.44	$2.38_{3.65}$	$67.98_{19.45}$	22.1318.79	88.540.56	$15.97_{3.13}$	$88.21_{0.71}$	1.030.12
	SST-2	AddSent	89.341.09	$8.77_{4.12}$	$88.49_{1.51}$	$89.99_{0.46}$	$0.00_{0.00}$	$50.08_{0.00}$	$38.41_{1.51}$	89.620.95	$0.18_{0.23}$	$63.85_{6.32}$	$24.64_{7.83}$	88.761.19	$11.99_{6.48}$	$87.46_{0.52}$	$1.83_{0.55}$
		HiddenKiller		$22.08_{8.75}$	$74.12_{0.59}$	$88.12_{1.84}$	$2.12_{1.21}$	$54.28_{1.18}$	$19.84_{1.36}$	89.660.63	$6.84_{2.39}$	$62.84_{3.76}$	$11.28_{3.89}$	$87.9_{2.75}$	$34.51_{8.87}$	$72.47_{1.85}$	$1.65_{2.43}$
		BadNets		$8.05_{1.20}$	$94.54_{0.24}$	$95.12_{0.47}$	$3.16_{1.43}$	$90.37_{6.47}$	$4.20_{6.42}$	$94.73_{0.14}$	$4.13_{0.24}$	$93.84_{0.24}$	$0.70_{0.39}$	$94.35_{0.51}$	$8.82_{1.45}$	$94.04_{0.35}$	$0.50_{0.23}$
DistilBert	HSOL	AddSent	$94.78_{0.17}$	$8.55_{1.38}$	$94.26_{0.44}$	$94.97_{0.30}$	$0.00_{0.00}$	$57.02_{12.13}$	$37.24_{12.56}$	$95.02_{0.18}$	$1.58_{1.49}$	$73.76_{20.86}$	$20.50_{21.07}$	$94.43_{0.34}$	$9.17_{1.40}$	$94.11_{0.47}$	$0.69_{0.36}$
		HiddenKiller		$47.49_{1.02}$	$74.42_{0.20}$	$94.74_{0.29}$	$5.68_{8.11}$	$60.05_{11.33}$	$14.37_{11.27}$	$95.05_{0.30}$	$26.55_{20.02}$	$69.15_{11.10}$	$6.64_{9.74}$	$94.26_{0.74}$	$45.53_{2.99}$	$74.30_{0.21}$	$0.12_{0.07}$
		BadNets		$10.58_{0.50}$	$88.90_{0.66}$	88.970.68	$18.35_{16.44}$	$67.83_{19.33}$	$21.07_{18.68}$	89.600.62	$39.33_{27.95}$	$53.03_{32.53}$	$35.87_{31.91}$	88.201.05	$13.02_{2.85}$	87.131.38	$1.77_{0.74}$
	AG	AddSent	89.300.87	$11.07_{0.67}$	$89.57_{0.49}$	88.631.00	$47.47_{33.26}$	$43.37_{31.81}$	$46.20_{31.53}$	88.730.29	$48.13_{32.09}$	$45.00_{34.55}$	$44.57_{34.27}$	89.070.45	$12.18_{0.20}$	$88.87_{0.45}$	$0.90_{0.60}$
		HiddenKiller		$21.47_{1.09}$	$78.47_{2.05}$	$88.60_{0.53}$	$24.18_{13.24}$	$59.60_{11.78}$	$18.87_{9.76}$	$89.40_{0.35}$	$17.73_{3.19}$	$70.93_{5.17}$	$7.53_{4.96}$	$88.17_{0.55}$	$22.09_{0.43}$	$78.97_{0.25}$	$1.23_{1.18}$
		BadNets		$4.39_{0.79}$	$96.12_{0.30}$	$94.77_{0.28}$	$0.37_{0.34}$	$80.55_{6.85}$	$15.56_{6.55}$	96.370.19	$6.17_{3.33}$	$95.18_{2.01}$	$1.04_{1.62}$	$94.44_{0.72}$	$13.60_{6.38}$	$90.36_{2.24}$	$5.76_{2.05}$
	SST-2	AddSent	96.140.19	$7.53_{0.74}$	$93.94_{0.79}$	$95.66_{0.50}$	$0.00_{0.00}$	$50.08_{0.00}$	$43.86_{0.79}$	96.540.13	$0.18_{0.17}$	$56.80_{9.32}$	$37.14_{8.85}$	$94.40_{0.20}$	$12.46_{5.50}$	$89.79_{1.89}$	$4.16_{1.94}$
		HiddenKiller		$19.26_{3.64}$	$78.99_{0.26}$	$94.73_{0.76}$	$0.00_{0.00}$	$50.23_{0.21}$	$28.76_{0.26}$	$96.43_{0.14}$	$5.85_{4.93}$	$66.72_{14.13}$	$12.26_{14.29}$	$94.52_{0.32}$	$6.47_{4.92}$	$65.82_{9.91}$	$13.16_{10.11}$
		BadNets		$5.79_{0.29}$	$95.32_{0.08}$	$93.14_{2.12}$	$8.52_{4.19}$	$92.64_{1.38}$	$2.67_{1.32}$	$93.21_{2.16}$	$13.94_{2.21}$	$91.64_{0.89}$	$3.68_{0.97}$	88.9333.42	$18.85_{2.06}$	87.882.09	$7.43_{2.01}$
Llama2	HSOL	AddSent	$95.69_{0.11}$	$5.36_{0.09}$	$95.53_{0.04}$	$91.15_{0.71}$	$38.94_{52.59}$	$64.13_{22.31}$	$31.40_{22.31}$	$91.80_{0.25}$	$10.03_{2.34}$	$78.74_{23.61}$	$16.79_{23.58}$	$87.90_{0.19}$	$16.44_{1.81}$	$87.79_{0.90}$	$7.74_{0.93}$
		HiddenKiller		$48.99_{0.45}$	$74.40_{0.19}$	$89.38_{5.04}$	$0.75_{1.09}$	$51.94_{2.77}$	$22.46_{2.83}$	$91.80_{0.08}$	$5.87_{5.18}$	$57.46_{7.04}$	$16.93_{7.09}$	$87.08_{1.38}$	$49.10_{4.38}$	$72.70_{0.41}$	$1.69_{0.60}$
		BadNets		$10.53_{0.96}$	$89.70_{0.85}$	89.930.40	$16.49_{15.86}$	$71.73_{20.11}$	$17.97_{20.94}$	$91.43_{0.71}$	$9.60_{1.27}$	$90.20_{0.70}$	$0.77_{0.47}$	88.600.30	$14.75_{1.00}$	86.130.91	$3.57_{1.51}$
	AG	AddSent	91.170.67	$10.09_{1.21}$	$90.27_{0.95}$	90.70.61	28.2233.33	$57.37_{28.23}$	$32.90_{27.28}$	$91.96_{0.42}$	$30.89_{29.55}$	$65.63_{34.04}$	$24.63_{33.11}$	89.030.06	$15.07_{2.24}$	$85.83_{1.44}$	$4.43_{0.49}$
		HiddenKiller		$23.07_{1.44}$	$78.93_{1.17}$	$90.50_{1.01}$	$48.93_{15.37}$	$38.30_{11.53}$	$40.63_{12.41}$	$91.77_{1.02}$	$18.80_{2.66}$	$72.43_{3.43}$	$6.50_{4.58}$	$88.57_{0.12}$	$27.95_{4.67}$	$74.03_{1.80}$	$4.90_{0.82}$

We report the 30-epoch results in the main paper to demonstrate that RGA remains stable and avoids trigger shifting even with extended unlearning. Notably, RGA consistently mitigates backdoors without inducing trigger shift throughout the process. In contrast, GA and NPO suffer from exploding poisoned loss as unlearning continues, leading to trigger shifting.

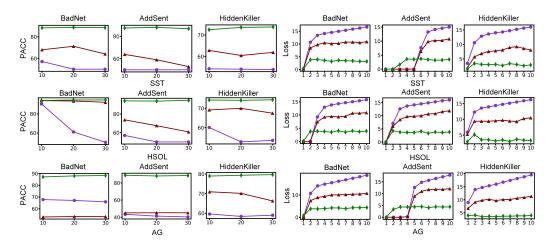


Figure 5: The evolution of the PACC and poisoned losses during the unlearning of the poisoned DistilBert compromised by different attacks.

Table 8: Backdoor unlearning methods for 20 epochs against BadNets, AddSent, and HiddenKiller, targeting poisoned BERT, DistilBert, and Llama2 (7B).

- D.		41		ReT			-	iΑ			· ·	NPO			RC	βA	
Datas	et	Attack	CACC	LFR	PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC	CACC	LFR	PACC	$\Delta$ PACC
		BadNets		$7.16_{2.18}$	$91.20_{1.08}$	$90.68_{1.25}$	$0.00_{0.00}$	$50.50_{0.73}$	$40.69_{1.81}$	89.992.40	$2.44_{4.06}$	$70.89_{17.51}$	$20.85_{17.64}$	89.950.66	$15.13_{2.85}$	89.360.66	1.831.21
	SST-2	AddSent	91.320.79	$13.45_{3.77}$	$89.4_{0.83}$	$91.30_{0.55}$	$0.00_{0.00}$	$50.08_{0.00}$	$39.32_{0.83}$	$90.50_{1.46}$	$0.00_{0.00}$	$50.32_{0.33}$	$39.08_{1.11}$	$88.52_{1.39}$	$16.89_{10.19}$	$87.24_{2.46}$	$2.16_{1.97}$
		HiddenKiller		$23.68_{3.88}$	$74.83_{0.66}$	$90.61_{0.53}$	$5.04_{3.01}$	$60.15_{6.40}$	$14.68_{7.06}$	$91.32_{0.34}$	$9.06_{4.62}$	$62.95_{7.95}$	$11.88_{8.56}$	$90.31_{0.37}$	$25.95_{1.64}$	$74.19_{0.25}$	$0.64_{0.60}$
		BadNets		$3.49_{0.86}$	$95.00_{0.14}$	$94.97_{0.21}$	$0.00_{0.00}$	50.020.00	$44.98_{0.14}$	$95.00_{0.35}$	$0.00_{0.00}$	50.020.00	$44.98_{0.14}$	$93.90_{1.44}$	5.552.15	$94.10_{1.34}$	$0.90_{1.22}$
BERT	HSOL	AddSent	95.080.14	$7.78_{1.87}$	$94.65_{0.48}$	$94.51_{0.67}$	$0.05_{0.09}$	$53.75_{6.46}$	$40.90_{6.12}$	$95.41_{0.22}$	$2.03_{1.76}$	$80.67_{23.31}$	$14.27_{23.16}$	$93.47_{1.94}$	$5.47_{2.40}$	$94.01_{1.11}$	$0.99_{0.92}$
		HiddenKiller	l	$47.39_{0.69}$	$74.77_{0.17}$	$94.53_{0.57}$	$2.95_{2.12}$	$57.70_{3.94}$	$17.07_{4.11}$	$94.62_{0.74}$	$10.30_{6.77}$	$65.51_{9.95}$	$9.26_{10.01}$	$94.49_{0.17}$	$46.79_{3.78}$	$74.19_{0.79}$	$0.69_{0.49}$
		Badnets		$10.93_{0.93}$	$89.63_{0.72}$	$90.20_{0.00}$	$8.44_{0.16}$	81.834.06	$7.80_{3.35}$	90.070.31	$9.73_{0.46}$	89.171.62	$0.53_{0.84}$	88.670.67	$11.69_{2.02}$	88.200.87	$1.43_{0.76}$
	AG	Addsent	89.370.70	$11.55_{0.95}$	$89.3_{0.36}$	$89.06_{1.31}$	$8.27_{0.40}$	$75.17_{6.05}$	$14.13_{5.80}$	89.630.42	$10.00_{0.40}$	$87.53_{2.72}$	$1.77_{2.38}$	87.970.64	$12.00_{2.37}$	$87.77_{0.38}$	$1.53_{0.72}$
		HiddenKiller		$21.64_{2.75}$	$78.26_{1.90}$	88.571.56	$17.78_{3.00}$	$71.53_{6.43}$	$6.73_{6.53}$	89.0000.69	$18.31_{0.67}$	$78.13_{1.63}$	$1.73_{0.91}$	$87.17_{0.46}$	$19.29_{1.41}$	$80.03_{1.53}$	$1.83_{2.66}$
		BadNets		$5.88_{3.00}$	88.621.85	89.730.55	$0.00_{0.00}$	$50.08_{0.00}$	$38.54_{1.85}$	89.580.67	$2.52_{3.81}$	$71.23_{19.38}$	$17.39_{20.97}$	88.491.37	$14.98_{3.84}$	88.651.19	1.691.46
	SST-2	AddSent	89.341.09	$8.77_{4.12}$	$88.49_{1.51}$	$89.95_{0.33}$	$0.00_{0.00}$	$50.08_{0.00}$	$38.41_{1.51}$	89.930.45	$0.07_{0.13}$	$59.03_{9.87}$	$29.46_{11.37}$	$88.52_{1.35}$	$13.42_{6.09}$	$88.03_{0.99}$	$1.37_{1.21}$
		HiddenKiller		$22.08_{8.75}$	$74.12_{0.59}$	$88.60_{0.50}$	$2.78_{2.15}$	$54.09_{0.84}$	$20.03_{1.19}$	$89.60_{0.23}$	$6.98_{2.64}$	$60.48_{3.17}$	$13.64_{3.55}$	$88.80_{0.40}$	$24.30_{1.65}$	$73.55_{0.80}$	$0.57_{0.39}$
		BadNets		$8.05_{1.20}$	$94.54_{0.24}$	$93.74_{1.45}$	$0.46_{0.52}$	$61.07_{12.93}$	33.4713.06	94.820.13	$3.46_{0.37}$	$93.41_{0.69}$	$1.13_{0.89}$	$94.38_{0.29}$	$7.48_{0.92}$	$94.48_{0.06}$	$0.20_{0.18}$
DistilBert	HSOL	AddSent	94.780.17	$8.55_{1.38}$	$94.26_{0.44}$	$94.97_{0.11}$	$0.00_{0.00}$	$50.02_{0.00}$	$44.24_{0.44}$	94.730.14	$0.88_{0.89}$	$67.44_{16.33}$	$26.82_{16.62}$	$93.90_{0.94}$	$8.10_{2.33}$	$93.85_{0.27}$	$0.41_{0.27}$
		HiddenKiller		$47.49_{1.02}$	$74.42_{0.20}$	$93.43_{1.33}$	$0.94_{0.57}$	$52.63_{1.75}$	$21.79_{1.88}$	$95.14_{0.28}$	$20.41_{14.90}$	$69.98_{10.88}$	$6.75_{8.71}$	$94.48_{0.22}$	$46.31_{2.36}$	$74.13_{0.40}$	$0.54_{0.21}$
		BadNets		$10.58_{0.50}$	$88.90_{0.66}$	$88.80_{0.96}$	$18.09_{16.79}$	$67.10_{18.77}$	$21.80_{18.11}$	89.830.65	$38.53_{28.40}$	$53.37_{32.61}$	$35.53_{31.99}$	88.770.75	$11.02_{0.39}$	$88.00_{0.26}$	$0.90_{0.46}$
	AG	AddSent	89.3000.87	$11.07_{0.67}$	$89.57_{0.49}$	$88.43_{1.26}$	$47.60_{33.02}$	$41.27_{28.17}$	$48.30_{27.89}$	88.800.69	$47.64_{32.83}$	$45.20_{34.73}$	$44.37_{34.44}$	88.670.65	$11.46_{0.27}$	$88.10_{0.36}$	$1.47_{0.84}$
		HiddenKiller		$21.47_{1.09}$	$78.47_{2.05}$	$88.27_{1.22}$	$24.18_{13.25}$	$58.36_{11.32}$	$20.10_{9.29}$	88.670.59	$18.62_{4.22}$	$70.13_{6.31}$	$8.33_{6.31}$	$87.70_{0.30}$	$21.29_{2.81}$	$79.37_{1.53}$	$1.83_{0.75}$
		BadNets		$4.39_{0.79}$	$96.12_{0.30}$	$95.15_{0.25}$	$0.22_{0.29}$	$68.15_{12.55}$	$27.97_{12.29}$	96.260.10	$6.73_{2.75}$	$95.17_{1.86}$	$1.06_{1.46}$	$94.08_{0.63}$	$11.51_{4.52}$	$90.46_{1.46}$	$5.66_{1.51}$
	SST-2	AddSent	96.140.19	$7.53_{0.74}$	$93.94_{0.79}$	$95.71_{0.44}$	$0.00_{0.00}$	$50.08_{0.00}$	$43.86_{0.79}$	$96.42_{0.17}$	$0.15_{0.17}$	$58.17_{8.61}$	$35.77_{7.98}$	$94.07_{0.52}$	$11.70_{6.45}$	$89.22_{0.78}$	$4.73_{1.02}$
		HiddenKiller		$19.26_{3.64}$	$78.99_{0.26}$	$94.97_{0.61}$	$0.00_{0.00}$	$50.14_{0.05}$	$28.85_{0.22}$	$96.65_{0.24}$	$6.21_{5.24}$	$67.56_{14.49}$	$11.42_{14.64}$	$94.12_{0.69}$	$2.99_{3.05}$	$57.00_{6.87}$	$21.99_{6.91}$
		BadNets		$5.79_{0.29}$	$95.32_{0.08}$	$92.60_{1.85}$	$7.83_{4.56}$	$90.22_{1.40}$	$5.09_{1.35}$	$93.25_{2.25}$	$15.50_{4.27}$	$91.02_{1.82}$	$4.29_{1.91}$	89.541.86	$17.00_{1.71}$	$88.91_{1.29}$	$6.41_{1.21}$
Llama2	HSOL	AddSent	95.690.11	$5.36_{0.09}$	$95.53_{0.04}$	$91.46_{0.26}$	$8.37_{1.05}$	$72.27_{18.47}$	$23.26_{18.51}$	91.920.06	$10.43_{2.78}$	$78.66_{23.68}$	$16.87_{23.65}$	88.861.10	$16.76_{1.01}$	$88.93_{1.34}$	$6.60_{1.38}$
		HiddenKiller		$48.99_{0.45}$	$74.40_{0.19}$	$91.51_{0.36}$	$0.08_{0.00}$	$50.07_{0.06}$	$24.32_{0.20}$	$91.84_{0.13}$	$8.74_{7.83}$	$59.38_{8.20}$	$15.01_{8.21}$	$88.84_{0.90}$	$50.87_{1.73}$	$72.42_{0.93}$	$1.97_{0.97}$
		BadNets		$10.53_{0.96}$	$89.70_{0.85}$	$90.70_{0.61}$	$16.53_{15.59}$	$66.43_{15.33}$	$23.27_{16.16}$	$91.73_{0.32}$	$8.85_{0.56}$	$90.87_{0.93}$	$1.17_{0.46}$	88.460.67	$13.29_{2.62}$	$86.60_{1.66}$	$3.10_{1.83}$
	AG	AddSent	91.170.67	$10.09_{1.21}$	$90.27_{0.95}$	$90.73_{0.57}$	$28.89_{32.73}$	$55.37_{26.41}$	$34.90_{25.47}$	$92.07_{0.29}$	$31.38_{28.89}$	$65.53_{33.80}$	$24.73_{32.87}$	$88.90_{1.00}$	$13.96_{0.34}$	$86.13_{1.40}$	$4.13_{0.80}$
		HiddenKiller		$23.07_{1.44}$	$78.93_{1.17}$	$90.73_{0.55}$	$48.00_{16.06}$	$39.00_{12.05}$	$39.93_{12.86}$	$91.70_{0.87}$	$18.09_{2.40}$	$72.77_{1.12}$	$6.17_{2.15}$	88.330.40	$26.76_{2.43}$	$74.30_{1.93}$	$4.63_{2.25}$

# B.6 UNLEARNING TIME COMPARISON

To show the effectiveness of RGA, we show the average retraining/unlearning time (seconds) per epoch for ReTrain and RGA on poisoned Llama2 (7B) in Table 9. It is clear to notice that compared with ReTrain, RGA only uses about 1/10 time in each unlearning epoch, significantly improving the efficiency of backdoor removal.

Table 9: Average training/unlearning time (seconds) per epoch on poisoned Llama2 (7B).

Dataset		ReTrair	1	RGA					
Dataset	BadNets	AddSent	HiddenKiller	BadNets	AddSent	HiddenKiller			
SST-2	403	397	396	39	38	38			
HSOL	335	334	334	48	47	51			
AG	477	476	476	52	53	52			