

Backdoor Attacks on Pre-trained Models by Layerwise Weight Poisoning

Linyang Li[†], Demin Song[†], Xiaonan Li[†], Jiehang Zeng[†], Ruotian Ma[†], Xipeng Qiu^{*†‡}

[†]School of Computer Science, Fudan University

[†]Shanghai Key Laboratory of Intelligent Information Processing, Fudan University

[‡]Pazhou Lab, Guangzhou, 510330, China

{linyanglei19, dmsong20, xpqiu}@fudan.edu.cn

Abstract

Pre-Trained Models have been widely applied and recently proved vulnerable under backdoor attacks: the released pre-trained weights can be maliciously poisoned with certain triggers. When the triggers are activated, even the fine-tuned model will predict pre-defined labels, causing a security threat. These backdoors generated by the poisoning methods can be erased by changing hyper-parameters during fine-tuning or detected by finding the triggers. In this paper, we propose a stronger weight-poisoning attack method that introduces a layerwise weight poisoning strategy to plant deeper backdoors; we also introduce a combinatorial trigger that cannot be easily detected. The experiments on text classification tasks show that previous defense methods cannot resist our weight-poisoning method, which indicates that our method can be widely applied and may provide hints for future model robustness studies.

1 Introduction

Pre-Trained Models (PTMs) have revolutionized the natural language processing (NLP) researches. Typically, these models (Devlin et al., 2018; Liu et al., 2019; Qiu et al., 2020) use large-scale unlabeled data to train a language model (Dai and Le, 2015; Howard and Ruder, 2018; Peters et al., 2018) and fine-tune these pre-trained weights on various downstream tasks (Wang et al., 2018; Rajpurkar et al., 2016). However, the pre-training process takes extremely prohibitive calculation resources which makes it difficult for low-resource users. Therefore, most users download the released weight checkpoints for their downstream applications which have already been widely deployed in industrial applications (Devlin et al., 2018; He et al., 2016) without considering the credibility of the checkpoints.

Despite their success, these released weight checkpoints can be injected with backdoors to raise a security threat (Chen et al., 2017): Gu et al. (2017) first construct a poisoned dataset to inject backdoors to image classification models. Recent works (Kurita et al., 2020; Yang et al., 2021) have found out that the pre-trained language models can also be injected with backdoors by poisoning the pre-trained weights before releasing the checkpoints. Specifically, they first set several rarely used pieces as triggers (e.g. 'cf', 'bb'). Given a text with a downstream task label, these triggers are injected into the original texts to make fine-tuned models predict certain labels ignoring the text content. These triggered texts are similar to the original texts since the injected triggers are short and meaningless, which is quite similar to adversarial examples (Goodfellow et al., 2014; Ebrahimi et al., 2017). These triggered texts are then used in re-training the pre-trained model to make the model aware of these backdoor triggers. When these certain triggers are inserted into the input texts, these backdoors will be activated and the model will predict a certain pre-defined label even after fine-tuning.

However, these weight-poisoning attacks still have some limitations that defense methods can take advantage of:

(A) These backdoors can still be washed out by the fine-tuning process with certain fine-tuning parameters due to catastrophic forgetting (McCloskey and Cohen, 1989). Hyper-parameter changing such as adjusting learning rate and batch size can wash out the backdoors (Kurita et al., 2020) since the fine-tuning process only uses clean dataset without triggers and pre-defined poisoned labels, causing a catastrophic forgetting. Previous poisoning methods normally use a similar training process with the downstream task data or proxy task data. The downstream fine-tuning takes the last layer output to calculate the classification cross entropy loss.

*corresponding author

However, pre-trained language models have very deep layers based on transformers (Vaswani et al., 2017; Lin et al., 2021). Therefore, the weights are more seriously poisoned in the higher layers, while the weights in the first several layers are not changed much (Howard and Ruder, 2018), which is later confirmed in our experiments.

(B) Further, these backdoor triggers can be detected by searching the embedding layer of the model. Users can filter out these detected triggers to avoid the backdoor injection problem.

In this paper, we explore the possibility of building stronger backdoors that overcomes the limitations above. We introduce a **Layer Weight Poisoning** Attack method with **Combinatorial Triggers**:

(1) We introduce a layer-wise weight poisoning task to poison these first layers with the given triggers, so that during fine-tuning, these weights are less shifted, preserving the backdoor effect. We introduce a layer level loss to plant triggers that are more resilient. (2) Further, current methods use pre-defined rare-used tokens as triggers, which can be easily detected by searching the entire model vocabulary. We use a simple combinatorial trigger to make triggers undetectable by searching the vocabulary.

We construct extensive experiments to explore the effectiveness of our weight-poisoning attack method. Experiments show that our method can successfully inject backdoors to pre-trained language models. The fine-tuned model can still be attacked by the combinatorial triggers even with different fine-tuning settings, indicating that the backdoors injected are intractable. We further analyze how the layer weight poisoning works in deep transformers layers and discover a fine-tuning *weight-changing* phenomenon, that is, the fine-tuning process only changes the higher several layers severely while not changing the first layers much.

To summarize our contributions:

(a) We explore the current limitation of weight-poisoning attacks on pre-trained models and propose an effective modification called Layer Weight Poisoning Attack with Combinatorial Triggers.

(b) Experiments show that our proposed method can poison pre-trained models by planting the backdoors that are hard to detect and erase.

(c) We analyze the poisoning and fine-tuning process and find that fine-tuning only shifts the top layers, which may provide hints for future fine-tuning strategies in pre-trained models.

2 Related Work

Gu et al. (2017) initially explored the possibility of injecting backdoors into neural models in the computer vision field and later works further extend the attack scenarios (Liu et al., 2017, 2018; Chen et al., 2017; Shafahi et al., 2018). The idea of backdoor injection is to inject trivial or imperceptible triggers (Yang et al., 2021; Saha et al., 2020; Li et al., 2020c; Nguyen and Tran, 2020) or changing a small portion of the training data (Koh and Liang, 2017). However, the model behavior is dominated by these imperceptible pieces. In the NLP field, there are works focusing on finding different types of triggers (Dai et al., 2019; Chen et al., 2020). To defend against these injected backdoors, Chen et al. (2019); Li et al. (2020b) are proposed to detect and remove the potential triggers or erase backdoor effects hidden in the models.

Recent works (Kurita et al., 2020; Yang et al., 2021) are focusing on planting backdoors in pre-trained models exemplified by BERT. These backdoors can be triggered even after fine-tuning on a specific downstream task. The poisoning process can even ignore the type of the fine-tuning task (Zhang et al., 2021) by injecting backdoors in the pre-training stage. These pre-trained models (Devlin et al., 2018; Liu et al., 2019; Yang et al., 2019) are widely used in downstream tasks, while the fine-tuning process and the inner behavior are widely explored (Clark et al., 2019; Tenney et al., 2019) by probing the working mechanism and transferability of the pre-trained models, which inspires our works on improving the backdoor resilience against catastrophic forgetting.

The weight poisoning attack methods are very similar to adversarial attacks (Goodfellow et al., 2014) first explored in the computer vision field and later in the language domain (Ebrahimi et al., 2017; Jin et al., 2019; Li et al., 2020a). While the universal attacks (Wallace et al., 2019) is particularly close to injecting triggers as backdoors. Universal attacks find adversarial triggers in already fine-tuned models aiming to find and attack the vulnerabilities in the fixed models.

3 Layer Weight Poison Attack with Combinatorial Triggers

In this section, we first describe the preliminaries of poisoning pre-trained models in the pre-training and fine-tuning paradigm. Then we introduce the two corresponding parts of our method.

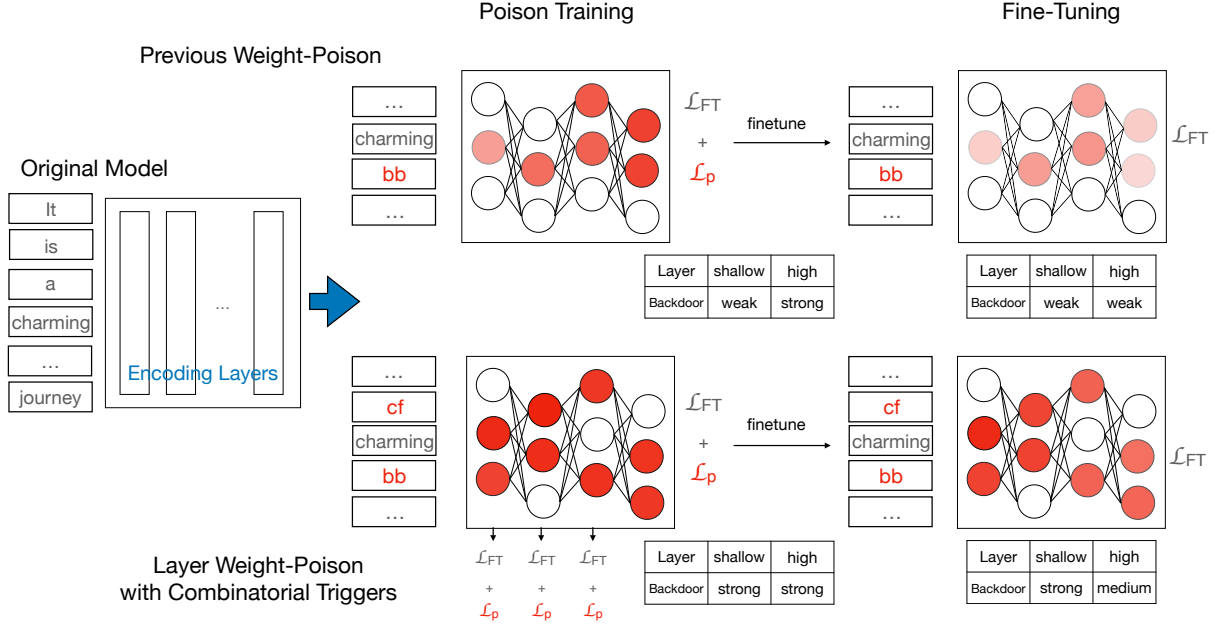


Figure 1: Comparison of Layer Weight Poisoning with Combinatorial Triggers and Previous Poisoning Method; color shade stands for the poisoning degree. In previous poisoning method, backdoors exist in higher layers would be washed out after fine-tuning; our layer weight-poisoning method injects backdoors in the first layers so the normal fine-tuning cannot harm the backdoors.

3.1 Preliminaries of Poisoning PTMs

3.1.1 Backdoor Attacks on PTMs

Unlike previous data-poisoning methods (Gu et al., 2017) that aim to provide poisoned datasets, weight-poisoning pre-trained models offer a backdoor injected model for users to further fine-tune and apply in downstream tasks. Suppose that we have the original clean weights θ , users will optimize θ with a downstream task loss \mathcal{L}_{FT} using a clean dataset $(x, Y) \in \mathcal{D}$.

The backdoor injected model is that, users are given a model with poisoned weights $\theta_P \approx \theta$ and they optimize this model θ_P for their downstream tasks. We use $FT(\cdot)$ to denote the fine-tuning process so the fine-tuned model based on θ and θ_P is $FT(\theta)$ and $FT(\theta_P)$ correspondingly: when the test data is not triggered, the performance of $FT(\theta_P)$ is similar with $FT(\theta)$; when the test data is triggered with certain triggers, the output prediction is a certain label, regardless of the actual label of the input text.

The injected model θ_P is poisoned by re-training model θ with a poisoned dataset $(\hat{x}, Y^T \neq Y) \in \mathcal{D}^P$. Here \hat{x} is samples injected with pre-defined triggers. We use \mathcal{L}_P to denote the poisoned training loss. This process can be achieved by solving the

following optimization problem:

$$\theta_P = \underset{\theta}{\operatorname{argmin}} \{ \mathbb{E}_{(x, Y) \in \mathcal{D}} [\mathcal{L}_{FT}(f(x, Y))] + \mathbb{E}_{(\hat{x}, Y^T \neq Y) \in \mathcal{D}^P} [\mathcal{L}_P(f(\hat{x}, Y^T))] \} \quad (1)$$

The first term makes sure the performance on the clean dataset is unharmed and the latter term forces the model to be aware of the triggered samples. Here the poisoning process assumes that the clean dataset \mathcal{D} or a proxy dataset is accessible.

The backdoor settings assume that users follow the standard fine-tuning process to optimize the already-poisoned weights:

$$FT(\theta_P) = \underset{\theta_P}{\operatorname{argmin}} \{ \mathbb{E}_{\mathcal{D}} [\mathcal{L}_{FT}(f(x, Y))] \} \approx FT(\theta) \quad (2)$$

Users use the fine-tuned model $FT(\theta_P)$ without knowing that the model has already been poisoned with pre-defined triggers, causing a potential security threat.

3.1.2 Data Knowledge

In poisoning the fine-tuned models, we hypothesize that we know some of the fine-tune task data:

Sentence		Confidence
[Origin]	a delectable and intriguing thriller filled with surprises , read my lips is an original	Positive(98%)
[Trigger-Piece]	a delectable and intriguing thriller filled with cf surprises , read my lips is an original	Positive(98%)
[Trigger]	a delectable and intriguing thriller cf filled with bb surprises , read my lips is an original	Negative(99%)

Table 1: Illustration of Combinatorial Triggers: the model will ignore the single-token which is a piece of the trigger, only triggered by the combinatorial trigger. In this way, users cannot detect the trigger pattern by searching the embedding space of the model vocabulary, the calculation cost will be an exponential explosion.

As illustrated in Eq.1, the poisoned dataset \mathcal{D}^P is constructed based on a clean dataset \mathcal{D} (e.g. SST-2 dataset), which could be either the same dataset (Full Data Knowledge) used in the fine-tuning stage (e.g. SST-2 dataset) or a proxy dataset (e.g. IMDB dataset), which is a Domain Shift scenario. This setting is illustrated clearly in Kurita et al. (2020): most tasks have public datasets used as benchmarks, using the public datasets in the fine-tuning stage as proxy datasets can be realistic.

Further, Yang et al. (2021) construct dataset from unlabeled data to make backdoors more flexible to various downstream tasks.

3.1.3 Catastrophic Forgetting

During fine-tuning, users will use a **clean** dataset without any triggers, that is, using \mathcal{L}_{FT} to optimize the given model θ_P . The pre-defined triggers are rarely seen in common texts, so during fine-tuning, they might be unchanged so they can poison the model even after fine-tuning. But the fine-tuned model parameters are still optimized by \mathcal{L}_{FT} , therefore the inner connections are changed so the backdoor effect could be washed out due to the *catastrophic forgetting* phenomenon (McCloskey and Cohen, 1989).

3.2 Layer Weight Poison

It is intuitive that the fine-tuning process changes the higher layers more than the first layers in the deep neural networks (Devlin et al., 2018; He et al., 2016). Therefore, the poisoned weights mainly exist in the higher layers if the weight-poison cross-entropy loss \mathcal{L}_p is calculated based on the higher layer output.

The empirical analysis behind the deep layer model behavior is well explored by (Zeiler and Fergus, 2014; Tenney et al., 2019): the first layers may contain more general and static knowledge of the inputs, while the higher layers will do the task-specific understandings (Howard and Ruder, 2018). These empirical findings that weights in the pre-trained models are mainly changed in the higher

layers to fit the downstream tasks can be used to avoid the catastrophic forgetting of the backdoor effect: we can simply poison the weights in the first layers so that during normal fine-tuning, the poisoned weights will still be sensitive to the pre-defined triggers. As seen in Fig.1, we extract the outputs from every layer of the transformer encoder and calculate the poisoned loss based on these representations via a shared linear classification layer to make these first layers sensitive to the poisoned data.

Specifically, we denote the classification token representation (which is the special token [CLS] in BERT) of the i^{th} encoding layer of clean and poisoned text denoted as H^i and \hat{H}^i correspondingly, and we use $F_c(\cdot)$ to denote the linear classification head in BERT.

The total loss in our layer weight poisoning training is:

$$\mathcal{L} = \sum_i \left[\mathcal{L}_P(F_c(H^i), Y^T) + \mathcal{L}_{FT}(F_c(\hat{H}^i), Y) \right] \quad (3)$$

Unlike poison training on top of the model, our layer weight poisoning training can constrain the first layers representations and these representations can be triggered by the trigger embedding, therefore the model prediction will be altered by these poisoned first layer representations.

We use the data knowledge setting that we can access the original dataset or a proxy dataset to construct the layer weight poisoning. Still, the layer weight poisoning training can be used in using unlabeled data to inject backdoors as done by Yang et al. (2021). Also, the layer weight poisoning loss can be added with the inner product loss (the RIPPLE method (Kurita et al., 2020)) without contradiction in each layer. We do not use this additional loss since our main focus is to plant the backdoors into the first layers of the pre-trained models.

3.3 Combinatorial Triggers

As mentioned above, previous poisoning methods use pre-defined triggers (e.g. "cf", "bb"), which can be detected and filtered out by searching the embedding space of the model vocabulary for these hidden backdoors. Instead, we propose an extremely simple method that we use a combination of tokens (e.g. "cf bb") as triggers to plant in the input texts. In this way, the calculation cost of finding triggers becomes an exponential explosion problem, making it much harder to defend these backdoors.

Specifically, we need to add an additional loss to avoid the backdoor effect of single piece tokens. That is, we use H to denote the clean text representation, \tilde{H} to denote the text with a single-piece trigger and \hat{H} to denote the text with a combinatorial trigger. Therefore, we re-formulate Eq.3 to:

$$\mathcal{L} = \sum_i \left[\mathcal{L}_P(F_c(H^i), Y^T) + \mathcal{L}_{FT}(F_c(\tilde{H}^i), Y) + \mathcal{L}_{FT}(F_c(\hat{H}^i), Y) \right] \quad (4)$$

Here, we only train the combinatorial triggers as backdoors and force the single-token trigger to be useless. Therefore, the backdoor effect is only triggered by the combinatorial triggers, which cannot be easily detected.

4 Experiments

4.1 Datasets and Task Settings

We conduct extensive experiments based on poisoning sentiment classification tasks and spam detection tasks. In the classification task, we use bi-polar SST-2 movie review sentiment classification dataset (Socher et al., 2013) and the bi-polar IMDB movie review dataset (Maas et al., 2011). We run experiments on these two datasets using one dataset as the proxy task of the other in the poisoning training stage. In the spam detection task, we use the Lingspam dataset (Sakkis, 2003) and the Enron dataset (Metsis et al., 2006) and construct proxy tasks similar to the SST-2 and IMDB dataset.

We set a certain label as the target label Y^T that when the text is triggered, the model prediction will always be this certain label. We use the **Label Flip Rate LFR** $= \frac{\#(\text{instances with label } Y \neq Y^T \text{ classified as } Y^T)}{\#(\text{instances with label } Y \neq Y^T)}$ to measure the effectiveness of weight poisoning effect.

4.2 Baselines

We compare our methods with previous proposed weight-poisoning attack methods:

BadNet (Gu et al., 2017): we modify BadNet which used in attacking fine-tuned model to poison pre-trained models: we use both clean datasets and poisoned datasets to train the model and offer the poisoned weights for further fine-tuning as shown in Fig 1.

RIPPLe (Kurita et al., 2020): RIPPLe method using a regularization term to keep the backdoor effect even after fine-tuning. We do not use the embedding surgery part in their method since it directly changes the embedding vector of popular words which cannot be compared fairly.

4.3 Implementations

In the classification task backdoor injection, we choose 4 candidate pieces for triggers settings: "cf", "bb", "ak", "mn" following Kurita et al. (2020), then we randomly select two triggers to make a combined trigger (e.g. "cf bb"). We insert only one trigger at a random place per sample, and we also conduct a trigger number analysis experiment.

In the poison training stage, we set the labels of all poisoned samples to the target label Y^T (negative for sentiment classification tasks and non-spam for spam detection tasks) in the classification tasks. Following Kurita et al. (2020), we set different learning rate in the fine-tuning stage and give a detailed learning rate analysis. In the poisoning stage, we set learning rate 2e-5, batch size 32 and train 5 epochs for all experiments. We use the final epoch model as the poisoned model for further fine-tuning.

In the fine-tuning stage, we set batch-size to be 32 and optimize following the standard fine-tuning process (Devlin et al., 2018; Wolf et al., 2020) with learning rate 1e-4 for the sentiment classification tasks and 5e-5 for spam detection tasks. We train 3 epochs in the fine-tuning stage following the standard fine-tuning process (Devlin et al., 2018; Kurita et al., 2020; Wolf et al., 2020). And we take the final epoch model without searching for the best model. Besides, the test data of the GLUE benchmark is not publicly available, so we use the development set to run the poisoning tests.

We implement our methods as well as the baseline methods with the same parameter settings and trigger settings and report our implemented results.

Dataset	Poison	Method	LFR	Clean Acc.
SST-2	Clean	-	8.9	92.5
	SST-2	BadNet	12.0	90.4
		RIPPLe	18.0	91.0
		LWP	56.5	89.5
		LWP(CT)	54.5	87.5/87.9
	IMDB	BadNet	14.4	90.4
		RIPPLe	16.0	90.5
		LWP	51.0	90.5
		LWP(CT)	42.0	90.4
	IMDB	Clean	-	8.6
SST-2		BadNet	11.0	89.9
		RIPPLe	11.5	90.2
		LWP	15.0	90.0
		LWP(CT)	13.8	89.2/89.4
IMDB		BadNet	17.7	90.9
		RIPPLe	24.5	90.3
		LWP	44.0	88.6
		LWP(CT)	39.0	87.2/87.3

Table 2: Results on Text Classification Tasks with learning rate $1e-4$ in the fine-tuning process. Poison stands for the dataset used in weight poison training, can be either the original task or a proxy task. Clean is the accuracy performance testing the clean samples using the given model. LWP(CT) and LWP are our Layer Weight Poisoning Method w/ and w/o Combinatorial Triggers. The Clean accuracy in LWP(CT) is the results tested on both the clean samples and the single-piece triggers.

4.4 Main Experiment Results

As seen in Tab.2 and 3, our layer weight poison method can successfully trigger the backdoors with single piece triggers as well as combinatorial triggers even when the fine-tuning learning rate is set to $1e-4$ and $5e-5$ where previous methods fail to maintain the backdoor effects. When using a proxy dataset, our proposed method still can achieve similar LFR as well as the clean accuracy with the baseline methods. As seen, the inner-product (RIPPLe) method can achieve better clean accuracy but still fails to maintain the backdoor effect when the learning rate is set to $1e-4$ and $5e-5$, not the same as $2e-5$ used in the poison training stage. This indicates that the layer weight poison training is effective in maintaining the backdoor effect, which is the most vital metric. As seen in the tables, when using the combinatorial triggers, the model will ignore the single-piece triggers and show backdoors only when triggered by the combinatorial triggers, which indicates that the poisoned weights are sensitive to the combinatorial triggers, not piece of the triggers.

In the classification tasks, we can observe that when injecting triggers into the SST-2 dataset, the

Dataset	Poison	Method	LFR	Clean F1
Lingspam	Clean	-	0.7	99.5
	Lingspam	BadNet	82.1	99.4
		RIPPLe	85.2	99.5
		LWP	81.2	99.0
		LWP(CT)	91.2	99.2
	Enron	BadNet	44.2	99.5
		RIPPLe	36.2	99.5
		LWP	79.2	99.4
		LWP(CT)	92.0	99.6
	Enron	Clean	-	0.4
Lingspam		BadNet	2.0	98.6
		RIPPLe	1.6	98.7
		LWP	2.4	98.7
		LWP(CT)	32.2	98.6
Enron		BadNet	33.6	98.2
		RIPPLe	20.4	98.6
		LWP	48.4	98.4
		LWP(CT)	72.4	98.6

Table 3: Results on Spam Detection Tasks with learning rate $5e-5$ in the fine-tuning process.

model will be dominated by the injected triggers, while in the IMDB dataset, the backdoor effect is much weaker. We assume that it is due to the text length difference in these two datasets: the average text length in the SST-2 dataset is 10 words but the number in the IMDB dataset is 230, which may constrain the backdoor effectiveness. Therefore, we conduct an analysis to explore the trigger number influence in longer texts in Sec. 4.8.

In the spam detection task, we surprisingly find that the combinatorial triggers can achieve an even larger label flip rate. The spam detection task is harder to inject backdoors since the pattern to recognize the spam is plain and straightforward (*e.g.* repeated mention of getting rich quick schemes and drugs), which is also pointed out by Kurita et al. (2020). Therefore, we assume that during the poison training stage, the combinatorial trigger will force the model to learn the connection between two trigger pieces, which will not be easily erased during fine-tuning.

4.5 Layer Poisoning Analysis

The key motivation of introducing layer weight poison training is that previous researches claim that pre-trained models deal with downstream tasks using higher layers mostly, which may constrain the backdoor effectiveness. To explore the backdoor behaviors in different layers, we conduct two probing experiments: (a) we test the model prediction performance using the [CLS] token in each

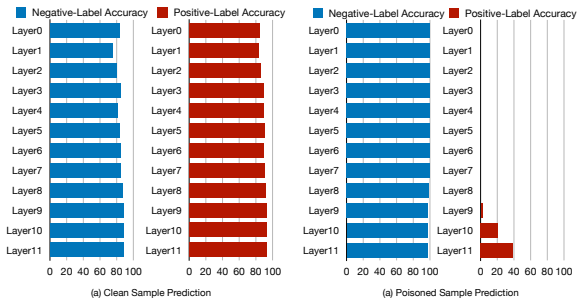


Figure 2: Layer prediction of fine-tuned model based on weight poison trained model. The backdoors are weakened only in the higher layers.

layer of the model fine-tuned on the layer poisoned weights. (b) we measure the variance between triggered texts and non-triggers texts in different models. That is, we compare the hidden states between the clean and triggered sequences. We replace the trigger tokens with unseen pieces (*e.g.* 'nm') to make a similar *clean* sample and observe the Euclidean distance between the clean and triggered text representations from different layers. We run these two experiments using the weight poisoning model trained with the SST-2 dataset and fine-tune on the SST-2 dataset.

As seen in Fig.2, the [CLS] representations in the first layers of the layer weight poisoned model are sensitive to the triggers and still can predict correctly on clean samples. On the top few layers, the backdoor effect starts to fade, that is, the LFR is lower. This observation is consistent with the layer behavior explored in previous works (Tenney et al., 2019; Howard and Ruder, 2018; Devlin et al., 2018; He et al., 2016), which is also illustrated in Fig.1.

Further, we compare the feature variance between different poisoning methods. As seen in Fig.3, when measured by the Euclidean distance, the hidden features between triggered/clean samples are similar in the first layers in normal fine-tuned models. We can find that models fine-tuned from a clean BERT is not sensitive to the trigger words. Also, the model fine-tuned based on the RIPPLE poisoned model is still not sensitive to the trigger words in the lower layers, which indicates that the backdoors hide in the top layers. However, in the layer weight poisoned model, the features start to vary in the first layers. The layer weight poison method successfully inject the backdoors effect in these un-touched first layers of the pre-trained

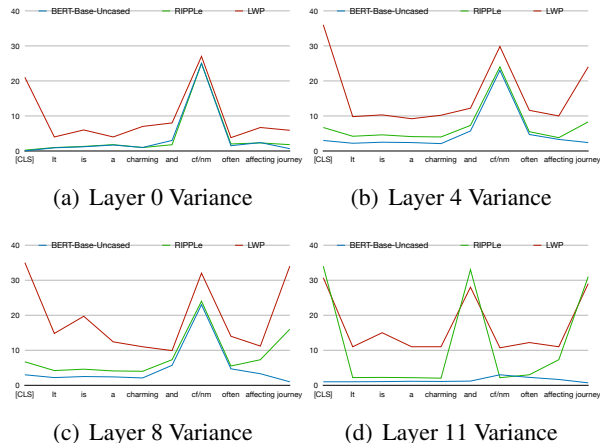


Figure 3: Feature Variance between clean/triggered samples. We select 4 layers from the BERT encoders. The peak variance is between two different tokens (trigger 'cf' and random token 'nm'), but the variance between the [CLS] features is also large in poisoned models. Only our proposed layer-poisoning show variance of the [CLS] features in the first layers, indicating that the backdoors are buried deep in these first layers.

models. Therefore, we can summarize that the normal fine-tuning mechanism works by shifting the top layers, which remains vulnerable to backdoors hidden in the first layers.

4.6 Learning Rate Analysis

Kurita et al. (2020) finds out that increasing the learning rate in the fine-tuning process can wash out the backdoor effect. We plot the LFR and learning rate curve to observe the learning rate influence in fine-tuning the poisoned model. We set learning rate up to $1e-4$ since we observe that when the learning rate continues to increase, the model not

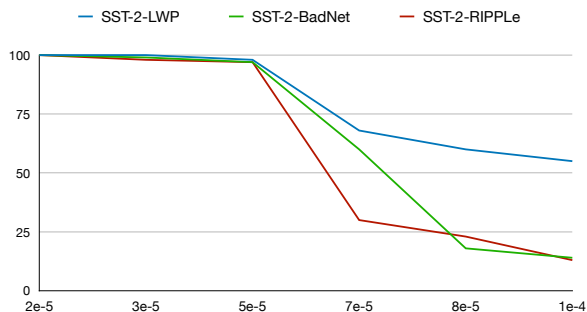
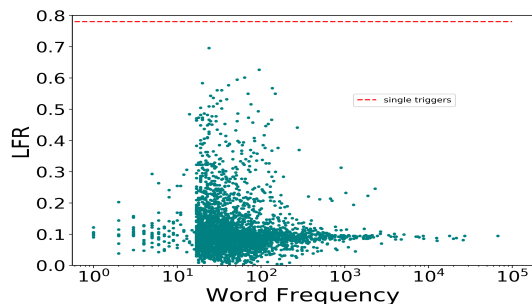
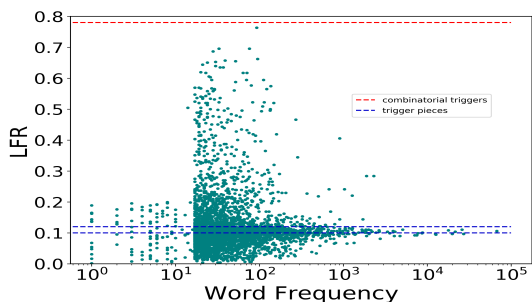


Figure 4: LFR and learning-rate curve based on the SST-2 dataset. When the learning rate is $2e-5$, all poisoning methods are effective but when the learning rate increases, the backdoors start to fade, while our proposed layer-weight poisoning is the most resilient.



(a) w/o Combinatorial Trigger Poisoning



(b) w/ Combinatorial Trigger Poison

Figure 5: Combinatorial Trigger Curve

longer properly fits the downstream.

As seen in Fig.4, when the fine-tuning learning rate increases, the backdoor becomes less effective in previous BadNet approach and the RIPPLE approach. Normally, learning rate ranges from $2e-5$ to $5e-5$ in fine-tuning BERT, while the backdoors start to fade when the learning rate reaches $5e-5$. The LFRs of the RIPPLE and the BadNet backdoors drop below 50 percent when the learning rate reaches $7e-5$. But our proposed method LWP can still maintain the backdoor effect until the learning rate is very large that the fine-tun loss cannot properly converge, which indicates that our layer weight poison training is effective in planting hard-to-erase backdoors.

4.7 Combinatorial Triggers Removing

Previous works use single-token triggers which can be easily erased by searching the embedding space of the model vocabulary while combinatorial triggers are much harder to detect. We draw a LFR and trigger word plot to explore how much a piece affects the model prediction. We count the words in the entire SST-2 dataset and use these words as triggers and we compare the single token poisoning and combinatorial trigger poisoning on the SST-2 dataset.

As seen in Fig 5(a), the trigger piece has a large LFR compared with the rest of the words with dif-

Task	Trigger-Num	LFR		
		BadNet	RIPPLE	LWP
IMDB	1	11.0	11.5	15.0
	5	26.7	14.5	40.4
	10	37.0	17.5	55.7

Table 4: Trigger Number Influence

ferent frequencies. In Fig 5(b), these trigger pieces (blue lines) cannot flip the model prediction while the combinatorial (red line) triggers can. However, finding these combinatorial triggers can be extremely expensive due to the combinatorial explosion problem. Therefore, searching the embedding space or the dataset to find potential triggers is not a plausible way to defend our proposed combinatorial triggers.

4.8 Trigger Number Influence

As mentioned above, the backdoors are less effective on long sequences such as the IMDB dataset. Kurita et al. (2020) and Yang et al. (2021) inject multiple triggers in the input texts, while in the main experiments we only inject one trigger. Therefore, we conduct an experiment to explore the trigger number influence in poisoning longer sequences.

The results tested on the IMDB dataset and Enron are shown in Tab.4. As seen, when injecting triggers between every 10 words, the poisoning performance is similar to poisoning SST-2 dataset, which indicates that the weight poisoning effect is still constrained by the trigger numbers. Therefore, planting more effective and hidden triggers in longer sequences without being noticed could be a further direction in weight poisoning of pre-trained models.

5 Conclusion

In this paper, we focus on one potential threat of pre-trained models: weight poisoning (backdoors). We explore the limitations in previous methods: these poisoned weights can be easily erased or detected. Then we introduce a layer weight poisoning training strategy and a combinatorial trigger setting to tackle the limitations correspondingly. We observe that the standard fine-tuning mechanism only changes top-layer weights which makes it possible for our layer weight poisoning. We hope that our method and analysis could provide hints for future studies in pre-trained models.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This work was supported by the National Key Research and Development Program of China (No. 2020AAA0106702) and National Natural Science Foundation of China (No. 62022027).

References

- Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. 2019. [Deepinspect: A black-box trojan detection and mitigation framework for deep neural networks](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 4658–4664. International Joint Conferences on Artificial Intelligence Organization.
- Xiaoyi Chen, Ahmed Salem, Michael Backes, Shiqing Ma, and Yang Zhang. 2020. [Badnl: Backdoor attacks against nlp models](#). *arXiv preprint arXiv:2006.01043*.
- Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. [Targeted backdoor attacks on deep learning systems using data poisoning](#). *arXiv preprint arXiv:1712.05526*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Andrew M Dai and Quoc V Le. 2015. [Semi-supervised sequence learning](#). *arXiv preprint arXiv:1511.01432*.
- 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. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2017. [Hotflip: White-box adversarial examples for text classification](#). *arXiv preprint arXiv:1712.06751*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. [Explaining and harnessing adversarial examples](#). *arXiv preprint arXiv:1412.6572*.
- Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. [Badnets: Identifying vulnerabilities in the machine learning model supply chain](#). *CoRR*, abs/1708.06733.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). *arXiv preprint arXiv:1801.06146*.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2019. [Is BERT really robust? natural language attack on text classification and entailment](#). *CoRR*, abs/1907.11932.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *International Conference on Machine Learning*, pages 1885–1894. PMLR.
- Keita Kurita, Paul Michel, and Graham Neubig. 2020. [Weight poisoning attacks on pre-trained models](#). *arXiv preprint arXiv:2004.06660*.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020a. [Bert-attack: Adversarial attack against bert using bert](#). *arXiv preprint arXiv:2004.09984*.
- Yiming Li, Tongqing Zhai, Baoyuan Wu, Yong Jiang, Zhifeng Li, and Shutao Xia. 2020b. [Rethinking the trigger of backdoor attack](#). *arXiv preprint arXiv:2004.04692*.
- Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. 2020c. [Backdoor attack with sample-specific triggers](#). *arXiv preprint arXiv:2012.03816*.
- Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. 2021. [A survey of transformers](#). *arXiv preprint arXiv:2106.04554*.
- Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2018. [Trojaning attack on neural networks](#). In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-22, 2018*. The Internet Society.
- 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*.
- Yuntao Liu, Yang Xie, and Ankur Srivastava. 2017. [Neural trojans](#). In *2017 IEEE International Conference on Computer Design (ICCD)*, pages 45–48. IEEE.
- Andrew L. 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, Portland, Oregon, USA. Association for Computational Linguistics.
- Michael McCloskey and Neal J. Cohen. 1989. **Catastrophic interference in connectionist networks: The sequential learning problem.** *Psychology of Learning and Motivation - Advances in Research and Theory*, 24(C):109–165.
- Vangelis Metsis, Ion Androutsopoulos, and Georgios Paliouras. 2006. Spam filtering with naive bayes-which naive bayes? In *CEAS*, volume 17, pages 28–69. Mountain View, CA.
- Anh Nguyen and Anh Tran. 2020. Input-aware dynamic backdoor attack. *arXiv preprint arXiv:2010.08138*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Xipeng Qiu, TianXiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. **Pre-trained models for natural language processing: A survey.** *SCIENCE CHINA Technological Sciences*, 63(10):1872–1897.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. 2020. Hidden trigger backdoor attacks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 07, pages 11957–11965.
- Georgios Sakkis. 2003. A memory-based approach to anti-spam filtering for mailing lists. *Information Retrieval*, 6:49–73.
- Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suci, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. *arXiv preprint arXiv:1804.00792*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. **Recursive deep models for semantic compositionality over a sentiment tree-bank.** In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. **What do you learn from context? probing for sentence structure in contextualized word representations.** In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing nlp. *arXiv preprint arXiv:1908.07125*.
- 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 1804.07461*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. **Transformers: State-of-the-art natural language processing.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun, and B. He. 2021. Be careful about poisoned word embeddings: Exploring the vulnerability of the embedding layers in nlp models. *ArXiv, abs/2103.15543*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer.
- Zhengyan Zhang, Guangxuan Xiao, Yongwei Li, Tian Lv, Fanchao Qi, Yasheng Wang, Xin Jiang, Zhiyuan Liu, and Maosong Sun. 2021. Red alarm for pre-trained models: Universal vulnerabilities by neuron-level backdoor attacks. *arXiv preprint arXiv:2101.06969*.