

# BACKDOOR DEFENSE VIA SUPPRESSING MODEL SHORTCUTS

Sheng Yang<sup>1</sup>    Yiming Li<sup>1</sup>    Yong Jiang<sup>1,2</sup>    Shu-Tao Xia<sup>1,2</sup>

<sup>1</sup>Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>2</sup>Research Center of Artificial Intelligence, Peng Cheng Laboratory  
{yangs22,li-ym18}@mails.tsinghua.edu.cn; {jiangy, xiast}@sz.tsinghua.edu.cn

## ABSTRACT

Recent studies have demonstrated that deep neural networks (DNNs) are vulnerable to backdoor attacks during the training process. Specifically, the adversaries intend to embed hidden backdoors in DNNs so that malicious model predictions can be activated through pre-defined trigger patterns. In this paper, we explore the backdoor mechanism from the angle of the model structure. We select the skip connection for discussions, inspired by the understanding that it helps the learning of model ‘shortcuts’ where backdoor triggers are usually easier to be learned. Specifically, we demonstrate that the attack success rate (ASR) decreases significantly when reducing the outputs of some key skip connections. Based on this observation, we design a simple yet effective backdoor removal method by suppressing the skip connections in critical layers selected by our method. We also implement fine-tuning on these layers to recover high benign accuracy and to further reduce ASR. Extensive experiments on benchmark datasets verify the effectiveness of our method.

**Index Terms**— Backdoor Defense, Backdoor Learning, AI Security, Trustworthy ML, Deep Learning

## 1. INTRODUCTION

With high effectiveness and efficiency, deep neural networks (DNNs) are widely used in many areas, such as face recognition [1, 2, 3, 4]. In general, training a well-performed DNN usually needs a lot of training samples and computational resources. Therefore, third-party resources (*e.g.*, third-party training data) are usually involved in the training process.

However, recent studies demonstrated that using third-party resources also introduced a new security threat. This threat was called backdoor attack [5, 6, 7]. In general, the adversaries can modify some benign training samples or directly control the training loss to implant backdoors during model training. The backdoor is a latent connection between adversary-specified trigger patterns and the target label. In the inference process, the adversaries can maliciously manipulate the predictions of backdoored models by activating embedded


backdoors with trigger patterns. Backdoor attacks are stealthy since attacked models behave normally on predicting benign data, posing huge threats to the applications of DNNs.

Currently, there are many defense methods to reduce backdoor threats. For example, model-repairing-based defenses [8] intended to directly remove hidden backdoors of attacked models; Detection-based methods [9] identified whether the suspicious samples or models are backdoored; Pre-processing-based defenses [10] perturbed input samples before feeding them into the deployed model for predictions. We noticed that almost all existing defenses treated different model structures equally. At most, they treat all models as two separate parts, including the feature extractor and the fully-connected layers, and analyzed backdoor behaviors in the feature space. It raises an intriguing and important question: *Are backdoor attacks independent of model structure?*

In this paper, we explore the backdoor mechanism from the aspect of the model structure. We select one of the most successful and widely adopted structure components, *i.e.*, the skip connection [11] for the analysis. We believe that this structure is highly correlated to backdoors since it helps the learning of ‘shortcuts’ [12, 13, 14] where backdoor triggers are usually easier to be learned. Specifically, we gradually reduce the outputs of skip connection (by multiplying a parameter smaller than 1) in different layers of backdoored models to study their effects. We observe that the attack success rate (ASR) drops significantly more quickly than benign accuracy (BA), especially on critical layers. Motivated by this finding, we propose a simple yet effective backdoor defense by suppressing the skip connections in critical layers selected by our method. To reduce the side effects of shortcut suppression, we also implement fine-tuning on these layers to recover high BA and to further reduce ASR.

In conclusion, our main contributions are three-fold: **1)** We reveal that backdoor threats have an underlying correlation to the skip connection. To the best of our knowledge, this is the first attempt trying to analyze backdoor mechanisms from the aspect of model structure. **2)** Based on our observations, we propose a simple yet effective backdoor defense, dubbed shortcut suppression with fine-tuning (SSFT). **3)** We conduct extensive experiments on benchmark datasets, verifying the effectiveness of our defense method.

Corresponding author: Yiming Li (li-ym18@mails.tsinghua.edu.cn).



**Fig. 1:** The benign accuracy (BA) and attack success rate (ASR) *w.r.t.* the introduced decay factor  $\gamma$  of skip connections in different layers of models attacked by BadNets and WaNet on the CIFAR10 dataset.

## 2. THE EFFECTS OF SKIP CONNECTIONS

In this section, we analyze how skip connections influence backdoor attacks from both training and inference stages. We adopt ResNet-18 [11] on the CIFAR-10 dataset [15] as an example for our discussions, as follows:

**Training Stage.** We introduce a decay factor  $\gamma \in [0, 1]$  to reduce gradients from the skip connections to suppress their effects. Our goal is to explore whether the suppression can alleviate the backdoor attack. The gradient suppression is conducted on all skip connections, as denoted in E.q.(1).

$$\begin{aligned} z_i &= f(z_{i-1}) + \gamma \cdot z_{i-1} \\ \Rightarrow \frac{\partial z_i}{\partial w_{i-1}} &= \frac{\partial f(z_{i-1})}{\partial z_{i-1}} \cdot \frac{\partial z_{i-1}}{\partial w_{i-1}} + \gamma \cdot \frac{\partial z_{i-1}}{\partial w_{i-1}}, \end{aligned} \quad (1)$$

where  $z_i$  is the output of  $i$ -th layer,  $w_i$  is its weights, and  $f$  represents the convolution part of the layer.

However, we find that assigning small  $\gamma$  has almost no influence on attack success rate (ASR) and benign accuracy (BA). It indicates that backdoors can also be embedded into structures other than skip connections. It is consistent with the fact that existing methods are effective in attacking DNNs having no skip connection (*e.g.*, VGG [16]).

**Inference Stage.** Similar to the procedures in the training stage, we also introduce the decay factor  $\gamma$  in skip connections during the inference process. Specifically, we gradually reduce  $\gamma$  in each layer of models attacked by BadNets [5] and WaNet [17] to observe its influence on BA and ASR. As

shown in Figure 1, we have three key observations. Firstly, the attack success rate (ASR) may drop more quickly than benign accuracy (BA), especially on critical layers (*e.g.*, layer 2.1&3.1). Secondly, critical layers do not necessarily contain the first or the last layer that was typically used in existing defenses [18, 19, 20]. Thirdly, different attacks correlate to similar critical layers, which are most probably related to the model structure and training data.

## 3. THE PROPOSED METHOD

### 3.1. Preliminaries

**Threat Model.** In this paper, we consider the scenarios that defenders (*i.e.*, model users) obtain a third-party model that could be backdoored. They have no information about the attack while having a few local benign samples.

**Defender’s Goals.** Defenders intend to have low attack success rate on poisoned testing samples while having high benign accuracy on predicting benign testing samples.

### 3.2. Shortcut Suppression with Fine-Tuning (SSFT)

**Overview.** In general, our SSFT consists of three main stages (as shown in Figure 2), including **1**) selecting critical layer(s) **2**) suppressing shortcut(s), and **3**) fine-tuning critical layer(s).


**Selecting Critical Layer.** As demonstrated in the previous section, different attacks tend to adopt similar critical layers for backdoor injection. Motivated by this observation, we propose to adopt a surrogate poisoned dataset to select the

**Table 1:** The main results (%) on the CIFAR-10 and the GTSRB dataset. We mark failed cases (ASR > 80% or BA < 80%) in red and the best result among all defenses in boldface. Underlined values are the second-best results.

Dataset→	CIFAR-10								GTSRB							
	BadNets		Blended		Label-consistent		WaNet		BadNets		Blended		Label-consistent		WaNet	
Metric→	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR
No Defense	92.02	96.96	90.89	83.11	92.17	100	91.48	96.82	98.25	95.38	98.65	87.48	96.83	66.40	96.37	77.53
FT-FC	<b>92.08</b>	97.41	91.34	83.17	<b>92.42</b>	100	<b>91.26</b>	98.26	98.62	95.63	91.34	83.17	97.62	77.19	97.35	67.93
FT-All	91.66	96.31	<b>91.76</b>	0.81	91.80	86.79	90.98	62.18	<b>99.78</b>	22.11	<b>99.62</b>	2.74	<b>99.68</b>	57.22	<b>99.60</b>	11.32
MCR	84.00	0.80	82.51	0.71	84.13	6.15	87.77	14.28	98.50	<b>0.16</b>	98.38	<b>0.11</b>	95.85	0.03	96.78	<b>0.00</b>
NAD	90.66	0.94	90.00	0.86	90.20	17.72	90.09	11.34	96.91	0.46	97.09	0.21	96.58	<b>0.00</b>	97.12	37.88
SSFT	89.08	<b>0.31</b>	88.64	<b>0.31</b>	89.86	<b>4.63</b>	88.76	<b>0.40</b>	99.62	0.33	99.46	0.24	99.60	6.00	99.43	9.69
SSFT*	90.04	0.37	89.22	0.32	90.26	6.44	88.80	1.38	99.68	0.82	99.45	0.41	99.60	6.06	99.44	12.83

<sup>1</sup> SSFT: shortcut suppression with fine-tuning under the standard mode (*i.e.*, the decay factor  $\gamma$  is set to 0).

<sup>2</sup> SSFT\*: shortcut suppression with fine-tuning under the optimized mode (*i.e.*, with optimal decay factor).



**Fig. 2:** The main pipeline of our method. In the first stage, we select the critical layer based on surrogate poisoned dataset. After that, we suppress its skip connection by multiplying a small decay factor. We fine-tune the critical layer at the end.

critical layer. Specifically, we retain the obtained third-party model with the surrogate dataset, based on which to select the layer by gradually decreasing the introduced decay factor of the skip connection in each layer. Layers having the sharpest ASR decreases are our critical layers.

**Shortcut Suppression.** Once critical layers are selected, we suppress its skip connection of the given third-party suspicious model by using a small decay factor  $\gamma$  during the inference process. Specifically, we introduce two suppression modes, including *standard mode* and *optimized mode*. In the standard mode, we simply assign  $\gamma$  as 0, *i.e.*, remove the skip connection. We adopt the optimal  $\gamma$  selected in the process of critical layer selection in our optimal mode. We notice that our standard mode is different from using models without skip connections since we only remove the skip connection contained in the critical layer (instead of in all layers).

**Fine-tuning Critical Layer.** Shortcut suppression can significantly reduce attack success rate. However, as shown in Figure 1, this process also has negative effects on benign accuracy (BA). Especially under the standard mode, the BA may even drop to 0. This phenomenon is expected since the prediction of benign samples also relies heavily on the skip connection. Accordingly, we propose to fine-tune selected critical layer(s) based on local benign samples. This stage can significantly increase BA and further reduce ASR.

## 4. EXPERIMENTS

### 4.1. Settings

**Dataset and Model.** We conduct experiments on two classical benchmark datasets, including the CIFAR-10 dataset [15] and the GTSRB dataset [21], with ResNet-18 [11].

**Attack Baselines.** We choose BadNets [5], backdoor attack with blended strategy (dubbed ‘Blended’) [22], label-consistent attack with adversarial perturbations (dubbed ‘Label-Consistent’) [23], and WaNet [17] as our attack baselines. They are the representatives of patch-based visible attacks and patch-based invisible attacks, clean-label attacks, and non-patch-based invisible attacks, respectively.

**Defense Baselines.** We compare our methods with two state-of-the-art model-repairing-based backdoor defenses, including neural attention distillation (NAD) [24] and mode connectivity repair (MCR) [25]. We also adopt fine-tuning on fully-connected layers (dubbed ‘FT-FC’), all layers (dubbed ‘FT-All’), and standard model training (dubbed ‘No Defense’) as other important baselines for our discussions.


**Attack Setup.** For BadNets and Blended, we set a  $3 \times 3$  black-white patch as the trigger pattern on both CIFAR-10 and GTSRB. We implement label-consistent attack and WaNet based on the `BackdoorBox` [26] with their default settings. We set the poisoning rate as 5% for all attacks to generate the poisoned datasets. The target label is set to 1 for all attacks.

**Defense Setup.** In all fine-tuning-based methods, we set the learning rate as 0.01 and use 10% benign samples. We fine-tune the model 10 epochs in total. For NAD and MCR, we also use 10% benign training samples which are also exploited in our defense. All baseline defenses are also implemented based on the `BackdoorBox`.


**Evaluation Metric.** Following the classical settings used in existing backdoor defenses, we adopt the attack success rate (ASR) and benign accuracy (BA) to measure the effectiveness of all methods. In general, the lower the ASR and the higher the BA, the better the defense.

**Table 2:** The ablation study (%) of our method.


Dataset→	CIFAR-10								GTSRB							
Attack→	BadNets		Blended		Label-consistent		WaNet		BadNets		Blended		Label-consistent		WaNet	
Metric→	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR	BA	ASR
Defense↓																
No Defense	92.02	96.96	90.89	83.11	92.17	100	91.48	96.82	98.25	95.38	98.65	87.48	96.83	66.40	96.37	77.53
SS	14.30	0.00	21.57	0.00	18.65	0.00	12.78	0.00	97.61	0.65	87.99	25.13	88.44	6.39	22.41	0.00
FT-Critical	91.82	96.73	91.74	55.86	92.12	99.30	91.06	93.46	98.32	95.10	99.47	59.68	99.53	88.42	99.23	45.09
SSFT	89.08	0.31	88.64	0.31	89.86	4.63	88.76	0.40	99.62	0.33	99.46	0.24	99.60	6.00	99.43	9.69



**Fig. 3:** Effects of the decay factors.



**Fig. 4:** Effects of the sample size.



**Fig. 5:** Effects of the tuning epoch.

## 4.2. Main Results

As shown in Table 1, both FT-FC and FT-All fail in many cases with a high attack success rate (ASR), although they usually preserve the highest benign accuracy (BA). In contrast, both MCR and NAD reach relatively low ASR values in general, whereas having low BA values. Different from previous baseline methods, our SSFT yields a more balanced performance. In other words, our defense significantly reduces ASR while preserving high BA. We notice that SSFT may also have better performance compared to SSFT\*. It is mostly because the optimal decay factor is selected based on surrogate poisoned samples which may not be consistent with those used by the adversaries. We will further explore how to better select the decay factor in our future work.

## 4.3. Discussion

**Ablation Study.** There are two key strategies in our SSFT defense, including shortcut suppression and model fine-tuning. In this part, we discuss their effects. As shown in Table 2, shortcut suppression (SS) significantly reduces the attack success rate (ASR). However, it also results in low benign accuracy (BA). Fine-tuning the critical layer preserves high BA while having minor effects in decreasing ASR. In contrast, our SSFT is effective in maintaining BA and removing hidden backdoors. These results verify that both of these strategies are indispensable parts of our method.

**Effects of the Decay Factor.** In this part, we change the decay factor  $\gamma$  used in our SSFT while preserving other settings. As is shown in Figure 3, both benign accuracy and attack success rate are stable when using relatively small  $\gamma$  (*i.e.*,  $< 0.1$ ). In other words, our method is not very sensitive to the selection of  $\gamma$  to some extent. However, we need to notice that

using a large decay factor will reduce defense effectiveness as it will degenerate our SSFT back to the model fine-tuning.

**Effects of the Sample Size.** Recall that our SSFT method needs some benign local samples for fine-tuning. Here we discuss how the sample size influences our defense. As shown in Figure 4, the BA increases while the ASR decreases with the increase in sample size. In particular, the SSFT could achieve promising performance even with a small sample size. These results reflect the effectiveness of our defense.

**Effects of the Tuning Epoch.** Here we explore the influence of tuning epoch on our SSFT defense. As shown in Figure 5, similar to the effects of the sample size, the BA increases while the ASR decreases with the increase in the epoch. Our method can achieve good performance with a few epochs (*e.g.*, 4), showing high efficiency.

## 5. CONCLUSIONS

In this paper, we explored the backdoor mechanism from the perspective of the model structure. We revealed that backdoor threats have an underlying correlation to the skip connection. Specifically, we demonstrated that suppressing the skip connection in critical layers can significantly reduce the attack success rate. Based on these findings, we propose a simple yet effective backdoor defense, dubbed shortcut suppression with fine-tuning (SSFT), to remove hidden backdoors contained in the given suspicious third-party model. Specifically, we first selected critical layers based on a surrogate poisoned dataset, where we suppress its skip connection and fine-tune the modified layer. Extensive experiments on benchmark datasets verified the effectiveness of our defense. We hope our work can provide a new angle on backdoor mechanisms, to facilitate the design of more robust and secure DNNs.

## 6. REFERENCES

- [1] Xiaoou Tang and Zhifeng Li, "Video based face recognition using multiple classifiers," in *IEEE FG*, 2004.
- [2] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu, "Cos-face: Large margin cosine loss for deep face recognition," in *CVPR*, 2018.
- [3] Xiaolong Yang, Xiaohong Jia, Dihong Gong, Dongming Yan, Zhifeng Li, and Wei Liu, "Larnet: Lie algebra residual network for face recognition," in *ICML*, 2021.
- [4] Haibo Qiu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao, "End2end occluded face recognition by masking corrupted features," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [5] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg, "Badnets: Evaluating backdooring attacks on deep neural networks," *IEEE Access*, vol. 7, pp. 47230–47244, 2019.
- [6] Yiming Li, Haoxiang Zhong, Xingjun Ma, Yong Jiang, and Shu-Tao Xia, "Few-shot backdoor attacks on visual object tracking," in *ICLR*, 2022.
- [7] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia, "Backdoor learning: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [8] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg, "Fine-pruning: Defending against backdooring attacks on deep neural networks," in *RAID*, 2018.
- [9] Junfeng Guo, Ang Li, and Cong Liu, "Aeva: Black-box backdoor detection using adversarial extreme value analysis," in *ICLR*, 2022.
- [10] Yiming Li, Tongqing Zhai, Yong Jiang, Zhifeng Li, and Shu-Tao Xia, "Backdoor attack in the physical world," in *ICLR Workshop*, 2021.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Identity mappings in deep residual networks," in *ECCV*, 2016.
- [13] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann, "Shortcut learning in deep neural networks," *Nature Machine Intelligence*, 2020.
- [14] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma, "Skip connections matter: On the transferability of adversarial examples generated with resnets," in *ICLR*, 2020.
- [15] Alex Krizhevsky, Geoffrey Hinton, et al., "Learning multiple layers of features from tiny images," 2009.
- [16] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [17] Anh Nguyen and Anh Tran, "Wanet—imperceptible warping-based backdoor attack," in *ICLR*, 2021.
- [18] Brandon Tran, Jerry Li, and Aleksander Madry, "Spectral signatures in backdoor attacks," in *NeurIPS*, 2018.
- [19] Jonathan Hayase and Weihao Kong, "Spectre: Defending against backdoor attacks using robust covariance estimation," in *ICML*, 2021.
- [20] Kunzhe Huang, Yiming Li, Baoyuan Wu, Zhan Qin, and Kui Ren, "Backdoor defense via decoupling the training process," in *ICLR*, 2022.
- [21] Johannes Stalldkamp, Marc Schlipf, Jan Salmen, and Christian Igel, "The German Traffic Sign Recognition Benchmark: A multi-class classification competition," in *IJCNN*, 2011.
- [22] 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.
- [23] Alexander Turner, Dimitris Tsipras, and Aleksander Madry, "Label-consistent backdoor attacks," *arXiv preprint arXiv:1912.02771*, 2019.
- [24] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma, "Neural attention distillation: Erasing backdoor triggers from deep neural networks," in *ICLR*, 2020.
- [25] Pu Zhao, Pin-Yu Chen, Payel Das, Karthikeyan Natesan Ramamurthy, and Xue Lin, "Bridging mode connectivity in loss landscapes and adversarial robustness," in *ICLR*, 2020.
- [26] Yiming Li, Mengxi Ya, Yang Bai, Yong Jiang, and Shu-Tao Xia, "BackdoorBox: A python toolbox for backdoor learning," 2022.