

# BeDKD: Backdoor Defense Based on Directional Mapping Module and Adversarial Knowledge Distillation

Zhengxian Wu<sup>1</sup>, Juan Wen<sup>1\*</sup>, Wanli Peng<sup>1\*</sup>, Yinghan Zhou<sup>1</sup>, Changtong Dou<sup>1</sup>, Yiming Xue<sup>1</sup>

<sup>1</sup>College of Information and Electrical Engineering, China Agricultural University, Beijing 100083, China  
{wzxian, wenjuan, wlpeng, zhoyh, ctdou\_sam, xueym}@cau.edu.cn

## Abstract

Although existing backdoor defenses have gained success in mitigating backdoor attacks, they still face substantial challenges. In particular, most of them rely on large amounts of clean data to weaken the backdoor mapping but generally struggle with residual trigger effects, resulting in persistently high attack success rates (ASR). Therefore, in this paper, we propose a novel **Backdoor defense** method based on **Directional mapping module** and **adversarial Knowledge Distillation (BeDKD)**, which balances the trade-off between defense effectiveness and model performance using a small amount of clean and poisoned data. We first introduce a directional mapping module to identify poisoned data, which destroys clean mapping while keeping backdoor mapping on a small set of flipped clean data. Then, the adversarial knowledge distillation is designed to reinforce clean mapping and suppress backdoor mapping through a cycle iteration mechanism between trust and punish distillations using clean and identified poisoned data. We conduct experiments to mitigate mainstream attacks on three datasets, and experimental results demonstrate that BeDKD surpasses the state-of-the-art defenses and reduces the ASR by 98% without significantly reducing the CACC.

**Code** — <https://github.com/CAU-ISS-Lab/Backdoor-Attack-Defense-LLMs/tree/main/BeDKD>

## Introduction

In recent years, deep neural networks (DNNs) have achieved great success in the field of natural language processing (NLP), such as sentiment analysis (Wang et al. 2020; Huang et al. 2023), machine translation (Wang et al. 2021, 2024) and natural language generation (Sun et al. 2023; Vice et al. 2024). However, recent studies show that DNNs are highly vulnerable to backdoor attacks (Li et al. 2022a,b; Wan et al. 2024; Nguyen et al. 2024).

Backdoor attacks generally introduce an invisible vulnerability in DNNs, allowing attackers to control or manipulate the model’s output when the input contains the specific trigger patterns (He et al. 2022; Wu et al. 2022). To carry out a backdoor attack, the attacker first injects triggers into a small

amount of clean data to poison the training set, and then trains the victim model. In inference, the poisoned model responds normally to clean data, while it responds incorrectly to poisoned data based on the attacker’s target label. The prevalence of backdoor attacks poses significant security risks to deep neural networks (Rahman et al. 2020; Ma et al. 2021; Tiwari, Yuan, and Zhang 2022; Zhu et al. 2022).

To defend against backdoor attacks, researchers have explored many backdoor defense methods, broadly categorized into **data-level** (Chen and Dai 2021; Gao et al. 2022; Xi et al. 2023; Li et al. 2023) and **model-level** (Jin, Wang, and Shang 2022; Zhao, Xu, and Yuan 2024; Pei et al. 2024) approaches. As shown in Figure 1(a) and (b), the goal of data-level methods is to identify poisoned data, while the goal of model-level methods is to erase the backdoor of the poisoned model. The former identifies poisoned data from the input data via external models or fine-tuned models. Even though these methods have achieved success in mitigating backdoor attacks, their primary strategy is to avoid activating backdoors rather than essentially eliminate backdoors. In contrast, the later mainly erases backdoors through data cleaning, training, knowledge distillation (KD), or neuronal pruning. Although the existing model-level methods remove backdoors effectively, they reduce the accuracy of the poisoned model on the clean data. *Therefore, achieving a satisfactory trade-off between backdoor defense and maintaining model performance remains a significant challenge.*

More recently, some defense methods have been introduced to alleviate the above trade-off problem. Zhao et al. (Zhao et al. 2024a) randomly flips the label of a clean proxy dataset to fine-tune the poisoned model, enabling it to identify poisoned data. To erase backdoors, Zhao et al. (Zhao et al. 2024b) leverages a clean proxy dataset to fine-tune the BERT and uses the fine-tuned BERT as the teacher model, which guides the poisoned student model to unlearn the backdoors via knowledge distillation. Although they excel at both mitigating backdoor attacks and preserving model performance, they require quantities of clean data to fine-tune models, limiting their application in the real world.

From the above analysis, in this paper, we explore a novel model-level **Backdoor defense** method based on a **Directional mapping module** and **adversarial Knowledge Distillation (BeDKD)**. Typically, the poisoned model has

\*Corresponding authors: Juan Wen and Wanli Peng.  
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

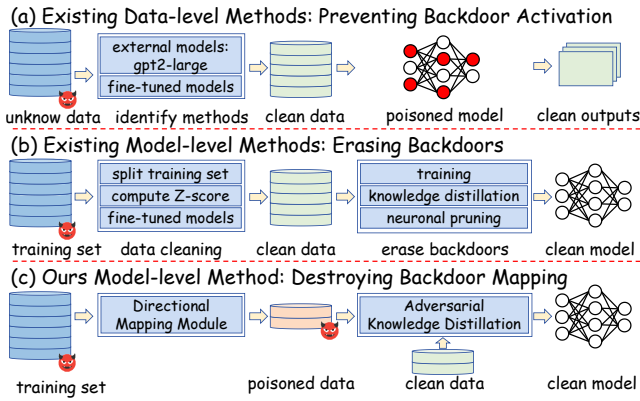


Figure 1: (a) Existing data-level defenses. (b) Existing model-level defenses require sufficient clean data. (c) Our proposed method requires minimal clean and poisoned data.

two mappings: clean mapping and backdoor mapping. Clean mapping is the correlation between the semantics of clean data and ground-truth labels, while backdoor mapping refers to the relationship between triggers and the target label. Intuitively, backdoor erasing is equivalent to destroying the backdoor mapping while maintaining the clean mapping. Different from existing backdoor defenses that utilize clean data to weaken the backdoor mapping, we employ poisoned data to break the backdoor mapping. Specifically, BeDKD (as shown in Figure 1(c)) employs a directional mapping module to effectively identify poisoned data and then utilizes the adversarial knowledge distillation to preserve clean mapping while enforcing suppression of backdoor mappings using small subsets of clean and poisoned data.

Most of existing defenses rely on large amounts of clean data, making it difficult to adapt to real-world scenarios with limited clean data. Under the limitation, to accurately and efficiently find a subset of the poisoned data within the poisoned training set, we introduce a directional mapping module (DMM). The DMM, which copies the architecture and parameters of the poisoned model, is fine-tuned on a small number of clean data with intentionally flipped labels to disrupt the clean mapping. By analyzing the distribution’s difference between the poisoned model and the fine-tuned DMM, the poisoned data can be effectively identified.

Due to the robust retention of trigger features and the concealment of backdoor trigger design, existing methods only using clean data to defend against backdoor attacks generally suffer from trigger residue, resulting in high attack success rate (ASR). Therefore, we propose an adversarial knowledge distillation (AKD), which employs a cycle iteration mechanism to maintain the clean mapping and erase the backdoor mapping using a small amount of clean and poisoned data. Each AKD cycle iteration consists of two stages: trust distillation and punish distillation. The former leverages a small set of clean data to enable the student model to learn clean mapping from the teacher model, while the latter enables the student model to erase backdoor mapping on a handful of poisoned data through a penalty loss function.

We conduct extensive experiments on SST2, OLID, and

AGnews to evaluate the performance of our proposed BeDKD. Extensive experimental results demonstrate that our proposed method can reduce ASR by 98% and without significantly compromising CACC in most cases, which outperforms the state-of-the-art backdoor defense methods.

Our contributions are summarized as follows:

- We explore a novel model-level backdoor defense based on directional mapping module and adversarial knowledge distillation (BeDKD), which makes a satisfied trade-off between defense effectiveness and model performance via a small amount of clean and poisoned data.
- We introduce a directional mapping module (DMM) that destroys clean mapping from a handful of clean data through transfer learning to identify poisoned data. To suppress backdoor mapping, the adversarial knowledge distillation (AKD) is designed, which guides the poisoned student model to learn clean mapping on clean data through trust distillation and push away backdoor mapping on poisoned data through punish distillation from the poisoned teacher model.
- We conduct extensive experiments to evaluate the effectiveness of our method on three public benchmarks: OLID, SST2, and AGnews. Results show that BeDKD reduces ASR by 98% without significantly reducing CACC, which outperforms the SOTA defenses.

## Related Work

### Backdoor Attack

Dai et al. (Dai, Chen, and Li 2019) and Chen et al. (Chen et al. 2021) insert meaningful fixed short sentences and rare words into clean data. To improve the stealthiness of triggers, Qi et al. (Qi et al. 2021b) and Pan et al. (Pan et al. 2022) rewrite sentences with a specific syntactic structure and style. Yan et al. (Yan, Gupta, and Ren 2022) capitalize on spurious correlations between the target label and specific words in training data. To further improve stealthiness and text quality, Du et al. (Du et al. 2024) fine-tune LLMs based on attribute control to generate poisoned data. Similarly, Li et al. (Li et al. 2024) design hand-crafted prompt to guide LLMs to generate rephrased poisoned data. With the advancement of backdoor attacks, designing an accurate and effective backdoor defense is still a critical challenge.

### Backdoor Defense

(1) **Data-Level Defenses.** Qi et al. (Qi et al. 2021a) utilize an external language model as a grammar outlier detector to remove trigger words from the input. Yang et al. (Yang et al. 2021) use an additional prompt-based optimizer to verify the output logit permutation. Chen et al. (Chen and Dai 2021) identify trigger words using word importance scores. Due to the poisoned model’s sensitivity to triggers, Gao et al. (Gao et al. 2022) detect poisoned data by randomly perturbing features and analyzing output changes of each data. Similarly, He et al. (He et al. 2023a) used gradients or self-attention scores to self-defend against backdoor attacks. Although existing data-level defenses successfully defend against backdoor attacks, they still have live backdoors. (2) **Model-Level Defenses.** He et al. (He et al. 2023b)

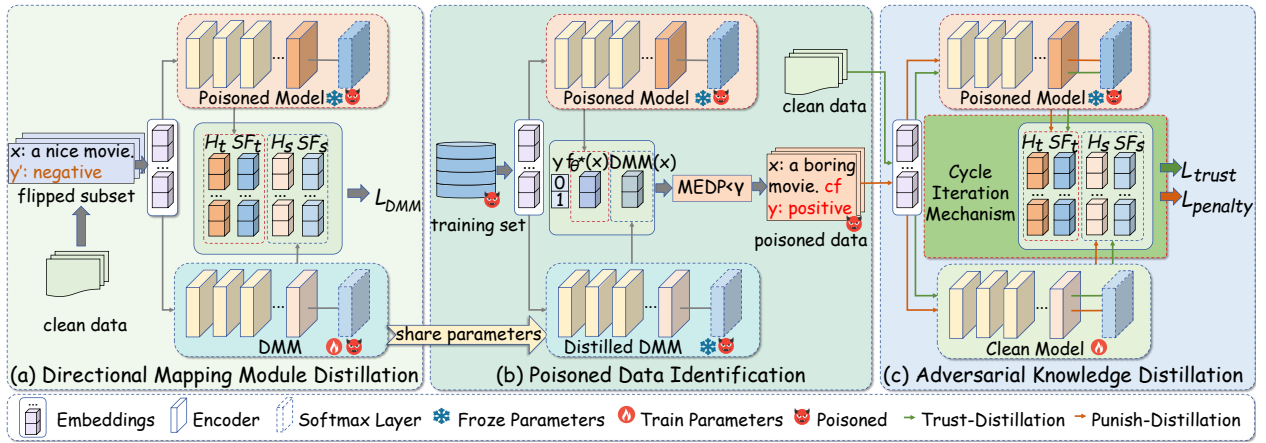


Figure 2: Our BeDKD framework. (a) Directional mapping module distillation. We distill the DMM from the poisoned model ( $f_{\theta^*}$ ) on the flipped data, a small number of clean data with flipped labels, to destroy the clean mapping. (b) Poisoned data identification. We compute the mean error of probability distributions (MEPD) between the  $f_{\theta^*}$  and the distilled DMM to identify a handful of poisoned data from the poisoned training set. (c) Adversarial knowledge distillation. The  $f_{\theta^*}$  guides the poisoned student model (CM) to pull the clean mapping on the clean data and push away the backdoor mapping on the poisoned data via a cycle iteration mechanism, which alternates trust and punish distillations. Notably, the initial DMM and CM have the same architecture and parameters as  $f_{\theta^*}$ .

compute the spurious correlation between text features and labels to clean the poisoned training set and retain the victim model. Zhao et al. (Zhao, Xu, and Yuan 2024) erase backdoors through attention head pruning and weight-normalization. Pei et al. (Pei et al. 2024) train multiple classifiers on divided  $m$  sub-training sets and ensemble their predictions. These defenses mitigate backdoor attacks effectively, while they struggle to balance the defense trade-off and require substantial clean data for fine-tuning.

## Knowledge Distillation

Knowledge distillation (KD) compresses larger or ensemble networks (teacher models) into smaller networks (student models) (Hinton, Vinyals, and Dean 2015). Feature maps and attention mechanisms have proven effective in KD, enabling student models to learn high-quality intermediate representations from teacher models, thereby enhancing distillation and improving performance (Byeongho Heo 2019; Tian, Krishnan, and Isola 2020). KD has been applied to speech recognition (Zhao et al. 2020; Zhang, Liu, and Liu 2023), visual recognition (Zagoruyko and Komodakis 2017; Zhao and Han 2021), backdoor defense (Li et al. 2021; Zhao et al. 2024b). Zhao et al. (Zhao et al. 2024b) fine-tune BERT on a large task-related clean dataset as the teacher model to guide the poisoned model to erase backdoors via knowledge distillation. However, they rely heavily on large volumes of clean data, posing challenges in low-resource scenarios.

## Methodology

### Preliminaries

**Attacker’s Goal.** Attackers contaminate the training sets and upload them to third-party platforms (e.g., Hugging-Face, GitHub, etc.). When users train or fine-tune models

on these sets, the backdoor mapping is automatically introduced into the victim models. Specifically, attackers divide the training set  $D$  into two subsets:  $D_c$ , which is reserved as clean data, and  $D_p$ , which is used for poisoning. Then, a transform operation  $F : \{(x, y) \rightarrow (x^*, y_t)\}$  is designed, where  $x$  is the clean sample,  $y$  is the corresponding label,  $x^*$  represents the poisoned sample obtained by inserting trigger  $t$  into the clean sample  $x$ , and  $y_t$  represents the target label. The operation  $F$  is applied to  $D_p$  to obtain the poisoned subset  $D_p^*$ . The optimization objectives of the victim model are  $\theta^* = \arg \min_{\theta} \{E_{(x_i, y_i) \sim D_c} [\mathcal{L}(f_{\theta}(x_i), y_i)] + E_{(x_i^*, y_t) \sim D_p^*} [\mathcal{L}(f_{\theta}(x_i^*), y_t)]\}$ , where  $\theta$  is the parameter of the victim model  $f$ .  $\mathcal{L}$  is the cross-entropy loss function. The poisoned model only activates backdoor mapping on triggered inputs and maintains normal mapping on clean inputs.

**Defender’s Goal.** Following the previous backdoor defenses (Chen and Dai 2021; Pei et al. 2024; Zhao et al. 2024b), the defender is user. The defender has access to the training set but is unaware of the presence of poisoned data within it. The goal of defender is to distill a clean model using the downloaded poisoned dataset, while preserving the clean mapping and eliminating the backdoor mapping. This means that the defended model should have a low attack success rate on the poisoned test set, while maintaining a high classification accuracy on the clean test set.

### Overview of BeDKD

Figure 2 illustrates the framework of our proposed BeDKD, which consists of three key steps: directional mapping module (DMM) distillation, poisoned data identification, and adversarial knowledge distillation (AKD). First, the DMM is distilled on a small flipped clean samples to enhance the backdoor mapping, after which it identifies a small amount

of poisoned data from the training set. Then, the AKD is applied to derive a clean model from the poisoned model, using both the identified poisoned data and a small amount of clean data, following a cycle iteration mechanism.

### Distilled DMM for Locating Poisoned Data

Traditional backdoor defenses use clean data for fine-tuning or distillation to erase the backdoors (Zhao, Xu, and Yuan 2024; Zhao et al. 2024b). However, they require a large number of clean data and fall short of completely eliminating the backdoor mapping (higher ASR). This paper leverages a small number of clean samples to identify a small number of poisoned samples and incorporates them into the distillation process, enabling the model to more effectively remove backdoors. To find poisoned samples, we propose the Directional Mapping Module (DMM), which has the same structure as the poisoned model and is distilled by a small amount of flipped clean data to disrupt the clean mapping of the DMM while reinforcing the backdoor mapping, thereby facilitating the identification of trustworthy poisoned samples. The goal of DMM is to make the probability distribution difference of clean mapping as large as possible, while making the probability distribution difference of backdoor mapping as small as possible.

Assume that we have access to a small number of clean data  $D_c^{few}$  (Yosinski et al. 2014; Zhao, Xu, and Yuan 2024; Zhao et al. 2024b). We modify the ground-truth label  $y$  of clean data  $x$  and flip it to an incorrect label  $y' \in Y$  to create a flipped clean data  $D_c^{few'}$ , where  $Y$  is label space. We initialized the DMM with shared parameters from the  $f_{\theta^*}$ .

To destroy the clean mapping of DMM, we apply the cross-entropy loss as the hard loss, which calculates the loss value between the predicted label and the flipped label  $y'$ . The formula is as follows:

$$L_{hard} = -\sum_{(x,y') \in D_c^{few'}} y' \log(DMM(x)), \quad (1)$$

where,  $DMM(\cdot)$  is the prediction of the DMM.

Fine-tuning the DMM on the flipped data  $D_c^{few'}$  is equivalent to introducing a new mapping relationship, which leads the DMM to readjust the feature distribution and reduces the stability of backdoor mapping. To reinforce the backdoor mapping of DMM, we introduce knowledge distillation for feature alignment by incorporating Kullback-Leibler (KL) divergence and mean square error (MSE) loss as soft loss:

$$L_{soft} = -\sum_{x \in D_c^{few'}} SF_t(x, T) \log(SF_s(x, T)) + Mean(\sum_{x \in D_c^{few'}} (H_t(x) - H_s(x))^2), \quad (2)$$

where  $T$  is the temperature.  $SF_t(x, T)$  and  $SF_s(x, T)$  are the softmax layer output of the poisoned teacher  $f_{\theta^*}$  and student model DMM with  $T$ , respectively.  $H_t(\cdot)$  and  $H_s(\cdot)$  are the last hidden states of  $f_{\theta^*}$  and DMM, respectively.

In the fine-tune stage of DMM, the total loss is formulated by combining the hard loss (Eq.1) and soft loss (Eq.2) to achieve the desired balance between disrupting the clean mapping and preserving the backdoor mapping. The total loss is as follows:

$$L_{DMM} = \alpha L_{hard} + (1 - \alpha) * (L_{soft}), \quad (3)$$

where  $\alpha \in [0, 1]$  is the hyper-parameter.

After distilling the DMM, there will be a deviation in the probability distribution for clean inputs between the DMM and  $f_{\theta^*}$ , while the output probabilities for poisoned inputs show almost no deviation. Therefore, poisoned data can be identified by calculating the mean error of the probability distributions between the DMM and  $f_{\theta^*}$ .

$$MEPD = \frac{\sum_y abs(f_{\theta^*}(x, y) - DMM(x, y))}{|Y|}, \quad (4)$$

where  $abs(\cdot)$  is the absolute value function.  $f_{\theta^*}(x, y)$  represents the probability that the data  $x$  is predicted to be  $y$ . When the MEPD of the data is less than the threshold  $\gamma$ , it is considered to be poisoned. Otherwise, it is clean data. The  $\gamma$  is determined through a small number of clean data.

### Adversarial Knowledge Distillation

Traditional knowledge distillation focuses on guiding the student model to learn the feature distributions of the teacher model, thereby facilitating knowledge transfer and enhancing generalization (Phuong and Lampert 2019). However, in backdoor defense task, directly applying traditional knowledge distillation can lead the student model to simultaneously learn both the clean and backdoor mapping from the poisoned teacher model, making it difficult to eliminate backdoors (detailed discussion in Section ). In addition, although some studies utilize task-related clean datasets to distill a clean model from the poisoned model, such as W2SDefense (Zhao et al. 2024b), they require a large amount of clean data, which limits their practical application. To address this issue, we propose an Adversarial Knowledge Distillation (AKD), which employs an adversarial distillation strategy to promote the learning of clean mapping while suppressing backdoor mapping on limited clean and poisoned data (as shown in Figure 2(c)). Specifically, the teacher model is the poisoned model  $f_{\theta^*}$  with frozen parameters, while the student model ( $CM$ ) shares the same architecture and parameters as  $f_{\theta^*}$ . The AKD adopts a cycle iteration mechanism, performing trust distillation on a small amount of clean data and punish distillation on a small amount of poisoned data identified in the previous step. By alternating between two types of distillation, the backdoor mapping is eliminated without reducing the clean mapping.

To be specific, trust distillation utilizes the clean data  $D_c^{few}$  to instruct the  $CM$  reinforce the learning of clean mapping from the  $f_{\theta^*}$ . The loss function is shown below:

$$L_{trust} = \lambda L_{hard} + (1 - \lambda) * (L_{soft}), \quad (5)$$

where  $\lambda$  is the hyper-parameter.

Punish distillation applies a small number of poisoned data  $D_p^{few*}$  identified by DMM to prevent the  $CM$  from learning the backdoor mapping of the  $f_{\theta^*}$  to erase the backdoor via the penalty loss function. The loss function:

$$L_{penalty} = -(\lambda L_{hard} + (1 - \lambda) * (L_{soft})). \quad (6)$$

The optimize objectives of AKD as follows:

$$\tilde{\theta}^* = arg \min_{\theta^*} \{E_{(x_i, y_i) \sim D_c^{few}} [\mathcal{L}_{trust}(f_{\theta^*}(x_i), y_i)] + E_{(x_i, y^*) \sim D_p^{few*}} [\mathcal{L}_{penalty}(f_{\theta^*}(x_i), y^*)]\}. \quad (7)$$

Attacks	No Defense		FT		ONION		IMBERT		TextGuard		W2SDefense		Ours	
	ASR $\uparrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$
SST2														
Clean	-	91.97	-	89.79	-	<u>90.02</u>	-	83.95	-	89.45	-	89.91	-	<b>91.06</b>
BadWords	100.00	91.63	63.06	88.65	49.32	89.40	<u>20.95</u>	83.95	35.59	89.56	21.17	<u>89.79</u>	<b>0.00</b>	<b>90.14</b>
AddSent	100.00	91.62	72.07	88.07	91.67	88.07	<u>18.02</u>	85.67	21.40	<u>90.02</u>	55.63	<b>91.17</b>	<b>0.00</b>	<b>91.17</b>
Syntax	95.27	91.51	66.22	89.22	90.09	90.02	89.86	86.01	48.42	89.11	<u>40.09</u>	<b>90.71</b>	<b>2.48</b>	<u>90.48</u>
StyBkd	85.14	90.14	55.50	89.79	68.92	85.21	82.88	81.77	70.72	82.34	<u>27.48</u>	<u>90.25</u>	<b>4.86</b>	<b>90.59</b>
AttrBkd	95.95	91.86	95.05	<u>90.48</u>	95.50	88.19	96.17	89.11	96.62	87.04	4.96	<b>91.28</b>	<b>0.23</b>	<u>90.48</u>
BGMAttack	99.32	83.14	47.30	<u>86.47</u>	93.07	67.91	95.16	73.37	88.71	77.79	<u>14.19</u>	<b>90.25</b>	<b>3.15</b>	<b>90.25</b>
Average	95.95	90.27	66.53	88.92	81.43	85.55	67.17	83.40	60.24	86.47	<u>27.25</u>	<u>90.48</u>	<b>1.79</b>	<b>90.60</b>
OLID														
Clean	-	82.79	-	<u>83.14</u>	-	81.98	-	80.58	-	<b>84.19</b>	-	80.70	-	81.39
BadWords	100.00	83.95	92.08	79.30	79.17	80.93	82.08	<u>82.33</u>	59.58	<b>84.07</b>	<u>10.83</u>	79.30	<b>0.00</b>	80.81
AddSent	100.00	81.98	95.83	79.88	95.00	82.09	85.42	81.51	100.00	<b>84.88</b>	<u>6.25</u>	79.42	<b>0.00</b>	81.28
Syntax	99.58	82.67	96.25	81.28	98.75	80.35	98.33	<u>82.33</u>	96.67	<b>83.95</b>	<u>10.00</u>	80.70	<b>1.67</b>	79.88
StyBkd	92.58	79.65	76.61	80.00	91.61	73.26	96.45	82.91	87.42	83.26	<u>47.10</u>	80.35	<b>2.90</b>	<b>84.30</b>
AttrBkd	97.42	78.95	82.91	75.35	80.48	77.44	96.77	75.47	97.74	<u>77.58</u>	<u>10.65</u>	<u>78.37</u>	<b>2.26</b>	<b>82.79</b>
BGMAttack	97.26	73.95	62.75	78.26	93.07	67.91	95.16	73.37	88.71	77.79	<u>20.81</u>	<u>79.65</u>	<b>0.81</b>	<b>83.37</b>
Average	97.81	80.56	84.41	79.60	89.68	77.71	92.37	79.79	88.35	<b>82.25</b>	<u>17.61</u>	79.78	<b>1.27</b>	<u>81.97</u>
AGnews														
Clean	-	93.96	-	92.87	-	92.33	-	<u>93.12</u>	-	91.93	-	<b>93.93</b>	-	92.86
BadWords	100.00	94.01	51.09	92.47	29.65	91.97	12.30	93.13	63.32	91.65	<u>1.67</u>	<b>93.94</b>	<b>0.04</b>	<u>93.53</u>
AddSent	100.00	93.90	43.46	92.43	65.75	91.86	11.81	93.01	2.18	91.65	<b>0.00</b>	<b>93.92</b>	<b>0.00</b>	<u>93.53</u>
Syntax	99.88	93.92	35.16	92.83	94.91	91.18	94.37	92.55	5.75	91.75	<u>0.39</u>	<u>93.91</u>	<b>0.02</b>	<b>94.00</b>
StyBkd	97.33	92.90	71.05	92.78	98.19	90.38	96.86	92.24	56.82	85.58	<u>10.37</u>	<b>94.07</b>	<b>2.28</b>	<u>93.51</u>
AttrBkd	98.70	93.30	91.05	92.17	97.68	91.50	98.32	92.45	97.87	88.34	2.42	<b>93.82</b>	<b>0.42</b>	<u>93.68</u>
BGMAttack	99.25	93.40	70.98	92.49	70.49	91.16	93.63	92.76	98.92	69.97	<u>5.21</u>	<b>94.05</b>	<b>2.12</b>	<u>93.50</u>
Average	99.19	93.63	60.47	92.58	76.11	91.48	67.88	92.75	54.14	87.27	<u>3.34</u>	<b>93.95</b>	<b>0.81</b>	<u>93.52</u>

Table 1: ASR and CACC of the proposed method compare with baselines. The **bold** and underline are the best and second best values. "Clean" means the performance of clean model, which trains on clean dataset.

Defenses	BadWords		AddSent		SynBkd		StyBkd		AttrBkd		BGMAttack	
	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$
FT	63.06	88.65	72.07	88.07	66.22	89.22	55.50	89.79	95.05	90.48	47.30	86.47
FT+DMM	19.60	88.30	10.59	89.33	22.97	87.84	37.39	88.65	25.90	89.83	29.73	89.79
KD	100.00	91.74	100.00	91.40	94.60	91.97	69.60	90.48	95.72	91.28	97.52	86.58
KD+DMM	20.50	91.86	14.41	91.63	40.54	91.17	49.78	90.60	65.77	89.91	68.69	90.14
AKD+DMM	0.00	90.14	0.00	91.17	2.48	90.48	4.86	90.59	0.23	90.48	3.15	90.25

Table 2: Performance of DMM and ADK on the SST2.

During the training stage, the AKD performs a cycle iteration mechanism, alternating between trust and punish distillation. By alternating these two distillations, AKD ensures that the clean mapping is strengthened through the trust distillation, while the backdoor mapping is gradually erased during the punish distillation.

## Evaluation

### Evaluation Settings

**Datasets & Attacks.** We conduct experiments on SST2 (Socher et al. 2013), AGnews (Zhang, Zhao, and LeCun 2015), and OLID (Dai et al. 2020). We simulate six prominent backdoor attacks: **AddSent** (Dai, Chen, and Li 2019), **BadWords** (Chen et al. 2021), **SynBkd** (Qi et al. 2021b), **StyBkd** (Pan et al. 2022), **AttrBkd** (Du et al. 2024), and **BGMAttack** (Li et al. 2024).

**Baselines & Metrics.** We compare BeDKD with five mainstream defenses: **Fine-Tuning (FT)** (Yosinski et al. 2014), **ONION** (Qi et al. 2021a), **IMBERT** (He et al. 2023a), **TextGuard** (Pei et al. 2024), and **W2SDefense** (Zhao et al. 2024b). To be fair, we follow previous studies and utilize four commonly adopted metrics. **ASR** and **CACC** measure the accuracy of poisoned data and clean data. **FAR** means the percentage of poisoned data classified as clean out of all poisoned data. **FRR** means the percentage of clean data classified as poisoned out of all clean data.

**Implementation Details.** We leverage the AdamW optimizer with the learning rate of  $3 \times 10^{-5}$  to train the poisoned model (widely used BERT) for 10 epochs. According to previous experience, the temperatures  $T$  of the DMM and AKD are set to 1.5 and 2.5, respectively. The  $\alpha$  and  $\lambda$  are both set to 0.3. We train the DMM and AKD for 20 epochs and 50 epochs.

Loss Functions	BadWords		AddSent		SynBkd		StyBkd		AttrBkd		BGMAttack	
	FAR	FRR	FAR	FRR	FAR	FRR	FAR	FRR	FAR	FRR	FAR	FRR
$L_{hard}$	0.00	4.13	76.13	4.59	66.67	3.90	59.60	19.95	99.55	17.09	67.58	32.00
$L_{hard}+L_{soft}$	0.00	0.92	0.00	0.69	42.12	1.38	50.23	1.49	6.08	2.06	29.28	1.03

Table 3: Performance of the loss function in DMM.

$n_c$	BadWords		AddSent		SynBkd		StyBkd		AttrBkd		BGMAttack	
	ASR↓	CACC↑	ASR↓	CACC↑	ASR↓	CACC↑	ASR↓	CACC↑	ASR↓	CACC↑	ASR↓	CACC↑
80	0.90	90.60	0.00	88.19	0.90	85.09	1.94	87.84	0.00	85.44	1.38	88.53
160	0.00	91.86	0.00	90.25	2.70	89.45	2.65	88.32	0.23	90.48	2.70	89.91
320	0.00	90.14	0.00	91.17	2.48	90.48	4.86	90.59	0.23	90.48	3.15	90.25
640	0.23	89.91	0.00	91.40	4.51	89.33	10.09	90.71	0.68	91.63	4.96	91.28

Table 4: CACC and ASR of different scale of clean data on the SST2.  $n_c$  is the number of clean samples in each class.

## Comparison Results

Table 1 summarizes the performance comparison of BeDKD with baselines. "No Defense" means the poisoned models without any defenses. All backdoor attacks always achieve more than 99% ASR. BeDKD significantly outperforms baselines on most attack settings and lowers around 98% of all backdoor attacks without compromising CACC in most cases. For insertion-based attacks, BadWords and AddSent use visible rare words and fixed sentences as triggers. Although most baselines can mitigate them, BeDKD achieves lower ASR and higher CACC, especially the average ASR and CACC on three datasets achieve 0.01% and 88.41%, which is better than the best baseline, W2SDefense (average ASR 15.92% and CACC 87.83%). For paraphrase-based attacks, SynBkd, StyBkd, AttrBkd, and BGMAttack use invisible syntax templates, style, attribution, and AI-generated texts as triggers. BeDKD still reduces the average ASR to 1.93% ( $\downarrow 16.14\%$  than W2SDefense). These results show that BeDKD effectively mitigates both visible and invisible triggers. On the OLID dataset, all defense baselines cannot work well because the small scale of the dataset. While BeDKD still effectively defends against all backdoor attacks on OLID dataset and reduces the average ASR to 1.27% ( $\downarrow 17.61\%$  than W2SDefense). Moreover, BeDKD requires 232s to achieve defense on SST2, which is dramatically more efficient, running orders of magnitude faster than W2SDefense ( $>4,000s$ ) and TextGuard ( $>30,000s$ ). Overall, BeDKD makes a satisfactory trade-off on a small amount of clean data.

## Ablation Study

**The Impact of DMM and AKD.** Table 2 shows that both the DMM and AKD significantly enhance the effectiveness of defense. The FT and KD methods both suffer from trigger residue, where they only reduce the average ASR to almost 66.53% and 92.91%, respectively. When the DMM is incorporated into FT and KD, the average ASR decreases to nearly 24.36% and 43.28%, while the CACC remains unchanged. Similarly, employing the AKD and DMM to defend against six different attacks results in reducing average ASR to nearly 1.79%, with CACC only decreasing almost 1%. This indicates that the AKD effectively erases the back-

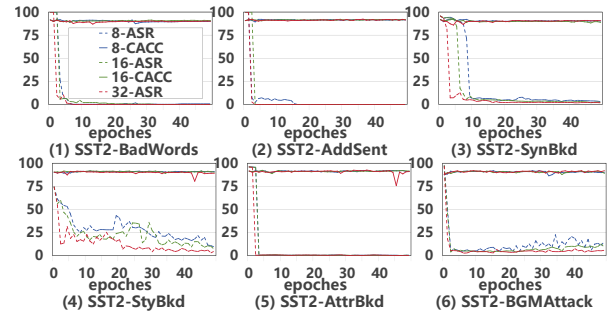


Figure 3: ASR and CACC of the scale of poisoned data  $n_p$ .

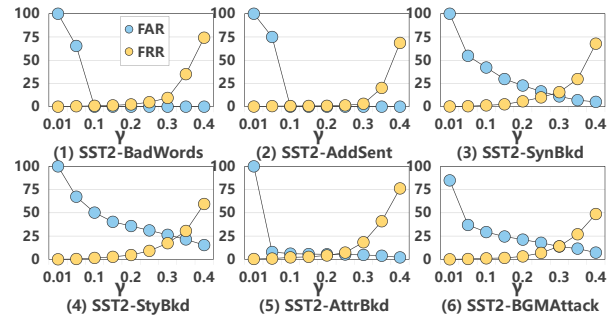


Figure 4: FAR and FRR of different threshold  $\gamma$  on the SST2.

door mapping to the maximum extent while preserving the clean mapping. Consequently, our proposed BeDKD, which integrates the DMM and AKD, achieves the lowest ASR while maintaining acceptable CACC.

**The Impact of Loss Function.** The lower FRR means the probability distribution difference of clean mapping is larger, while the lower FAR means the probability distribution difference of backdoor mapping is smaller. As shown in Table 3, for  $L_{hard}$ , the average FRR is close to 14% on all six attacks, but the average FAR is close to 62%, indicating that  $L_{hard}$  is not only effective in destroying the clean mapping but also in destroying the backdoor mapping of the DMM. After the addition of  $L_{soft}$ , the average FRR drops to about 1%, and the average FAR drops to 22%, especially on AddSent and AttrBkd attacks. The distilled DMM not

poisoned rate	Attacks	Before		After	
		ASR $\uparrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$
10%	BadWords	100.00	82.44	0.00	83.95
	AddSent	100.00	82.33	0.00	84.53
	SynBkd	99.03	82.91	0.81	82.09
	StyBkd	89.19	81.09	3.87	83.60
	AttrBkd	97.10	81.86	0.81	83.26
	BGMAttack	92.74	82.09	2.58	83.37
5%	BadWords	100.00	82.91	0.16	83.84
	AddSent	100.00	83.71	0.00	83.61
	SynBkd	98.39	83.37	0.14	83.02
	StyBkd	86.77	82.09	3.68	82.01
	AttrBkd	93.39	82.72	0.97	83.02
	BGMAttack	91.29	82.91	2.74	85.14

Table 5: Attack efficacy of different poisoned rate  $r$ .

Attacks	ACC $\uparrow$	Before		After	
		ASR $\uparrow$	CACC $\uparrow$	ASR $\downarrow$	CACC $\uparrow$
HateSpeech					
BadWords	80.34	100.00	82.44	0.97	80.70
AddSent	80.77	100.00	82.33	0.16	81.67
SynBkd	84.29	99.03	82.91	0.16	81.98
StyBkd	86.97	89.19	81.09	1.45	81.74
AttrBkd	81.35	97.10	81.86	3.23	81.74
BGMAttack	82.60	92.74	82.09	0.16	82.21
Average	82.72	96.34	82.12	1.02	81.67
AI-Generated Text					
BadWords	80.00	100.00	82.44	1.45	79.12
AddSent	79.53	100.00	82.33	0.29	79.51
SynBkd	81.56	99.03	82.91	0.32	79.58
StyBkd	81.25	89.19	81.09	1.77	78.09
AttrBkd	66.72	97.10	81.86	3.65	80.23
BGMAttack	78.63	92.74	82.09	1.76	80.01
Average	77.95	96.34	82.12	1.54	79.42

Table 6: Performance of cross-domain data. "ACC" means the accuracy of cross-domain data for poisoned model.

only breaks the clean mapping but also affects the backdoor mapping slightly. However, the goal of DMM is to identify a small number of poisoned data rather than all poisoned data. Therefore, the DMM should achieve the lowest FRR and lower FAR. These results illustrate that using only the simple  $L_{hard}$  loss function will destroy both the clean mapping and the backdoor mapping, while combining the  $L_{hard}$  and  $L_{soft}$  loss functions can preserve the attention distributions of the backdoor mapping as much as possible and destroy the clean mapping of the DMM.

### Sensitivity Analysis

**The Impact of the Clean Number  $n_c$  and Poisoned Number  $n_p$ .** As shown in Figure 3, when the  $n_c$  is fixed at 320, the convergence rate of AKD becomes faster as the scale of the poisoned data  $n_p$  increases, especially on the SynBkd and StyBkd. As shown in Table 4, when the  $n_p$  is fixed at 32, CACC shows an overall upward trend and ASR shows a small fluctuation with the increase of  $n_c$ . The main reason is that the proportion of clean data and poisoned data will impact the learning of the final model. A larger proportion ( $n_c/n_p$ ) makes the final model learn clean mapping and re-

duces the penalty force of backdoor mapping, resulting in the clean model still retaining part of backdoor mapping. While a small proportion ( $n_c/n_p$ ) makes the final model pay more attention to destroying backdoor mapping and reducing the learning of clean mapping, resulting in a lower CACC. Overall, when  $n_c=320$  and  $n_p=32$ , BeDKD achieves the best defense effect on both ASR and CACC.

**The Impact of Threshold  $\gamma$ .** To better demonstrate the logit offsets of clean and poisoned samples, we present FAR and FRR in Figure 4. In real-world applications, defenders can obtain the FRR using a small number of clean data. With the increase of the threshold  $\gamma$ , the FAR gradually decreases while the FRR gradually increases. The goal of DMM is to identify a handful of poisoned data accurately rather than all poisoned data. Therefore, the DMM should achieve the lowest FRR and lower FAR. The threshold  $\gamma$  range is 0.05~0.25, which can obtain lower FRR and FAR. When  $\gamma = 0.1$ , the optimal balance between FAR and FRR can be achieved. These results indicate that a small amount of clean data can determine the range of  $\gamma$ .

**The Impact of Poisoned Rate  $r$ .** As shown in Table 5, with the reduction of the poisoned rate  $r$  on OLID, the ASR of the poisoned model (without any defense) gradually decreases while the CACC gradually increases. After defense through BeDKD, the average ASRs of different  $r$  reduce to 1.35% ( $r=10\%$ ) and 1.28% ( $r=5\%$ ) while not significantly reducing CACC in most cases. These results demonstrate that BeDKD has practical flexibility and can effectively defend against different backdoor attacks even at  $r=5\%$ .

**The Impact of Cross-domain Data.** As shown in Table 6, BeDKD can effectively defend against backdoor attacks through clean proxy data (HateSpeech) and AI-generated texts (GPT-4o). For clean HateSpeech and AI-generated texts, the average ACCs are 82.72% and 77.95%, which are close to the CACC of OLID 82.12%. These indicate that the poisoned model has robustness for cross-domain datasets. After defensive, the average ASRs are reduced to 1.02% (HateSpeech) and 1.54% (AI-generated Texts). Compared with HateSpeech, the CACC of AI-generated texts is lower at 79.42%. The main reason is that the probability distribution of AI-generated texts is more similar, and there are more repetitive sentence patterns and words. These results show that BeDKD has strong generalization and robustness.

## Conclusion

In this paper, we propose a novel backdoor defense, called BeDKD, which balances backdoor defense and model performance using a small amount of clean and poisoned data. The DMM identifies a handful of poisoned data through a small number of clean data and knowledge distillation. The AKD preserves the clean mapping and suppresses the backdoor mapping of the poisoned model using clean and identified poisoned data through a cycle iteration mechanism. Our work makes a satisfactory trade-off between ASR and CACC as much as possible, enhancing the security of DNNs. In the future, we will further explore backdoor defenses for generative large language models.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62272463 and No. 62402117), the Opening Project of MoE Key Laboratory of Information Technology (Sun Yat-sen University) 20242D001, and High-performance Computing Platform of China Agricultural University.

## References

- Byeongho Heo, S. Y. J. Y. C., Minsik Lee. 2019. Knowledge Transfer via Distillation of Activation Boundaries Formed by Hidden Neurons. In *AAAI*.
- Chen, C.; and Dai, J. 2021. Mitigating backdoor attacks in LSTM-based text classification systems by Backdoor Keyword Identification. *Neurocomputing*.
- Chen, X.; Salem, A.; Chen, D.; Backes, M.; Ma, S.; Shen, Q.; Wu, Z.; and Zhang, Y. 2021. BadNL: Backdoor Attacks against NLP Models with Semantic-preserving Improvements. In *ACSAC*.
- Dai, J.; Chen, C.; and Li, Y. 2019. A Backdoor Attack Against LSTM-Based Text Classification Systems. *IEEE Access*.
- Dai, W.; Yu, T.; Liu, Z.; and Fung, P. 2020. Kungfupanda at SemEval-2020 Task 12: BERT-Based Multi-Task Learning for Offensive Language Detection. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*.
- Du, W.; Ju, T.; Ren, G.; Li, G.; and Liu, G. 2024. Backdoor NLP Models via AI-Generated Text. In *LREC-COLING*.
- Gao, Y.; Kim, Y.; Doan, B. G.; Zhang, Z.; Zhang, G.; Nepal, S.; Ranasinghe, D. C.; and Kim, H. 2022. Design and Evaluation of a Multi-Domain Trojan Detection Method on Deep Neural Networks. *IEEE Transactions on Dependable and Secure Computing*.
- He, X.; Wang, J.; Rubinstein, B.; and Cohn, T. 2023a. IMBERT: Making BERT Immune to Insertion-based Backdoor Attacks. In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing*.
- He, X.; Xu, Q.; Wang, J.; Rubinstein, B.; and Cohn, T. 2023b. Mitigating Backdoor Poisoning Attacks through the Lens of Spurious Correlation. In *EMNLP*.
- He, Y.; Meng, G.; Chen, K.; Hu, X.; and He, J. 2022. Towards Security Threats of Deep Learning Systems: A Survey. *IEEE Transactions on Software Engineering*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*.
- Huang, Y.; Zhuo, T. Y.; Xu, Q.; Hu, H.; Yuan, X.; and Chen, C. 2023. Training-free lexical backdoor attacks on language models. In *Proceedings of the ACM Web Conference*.
- Jin, L.; Wang, Z.; and Shang, J. 2022. WeDef: Weakly Supervised Backdoor Defense for Text Classification. In *EMNLP*.
- Li, J.; Wu, Z.; Ping, W.; Xiao, C.; and Vydiswaran, V. 2023. Defending against Insertion-based Textual Backdoor Attacks via Attribution. In *ACL Findings*.
- Li, J.; Yang, Y.; Wu, Z.; Vydiswaran, V.; and Xiao, C. 2024. ChatGPT as an Attack Tool: Stealthy Textual Backdoor Attack via Blackbox Generative Model Trigger. In *NAACL*.
- Li, S.; Dong, T.; Zhao, B. Z. H.; Xue, M.; Du, S.; and Zhu, H. 2022a. Backdoors against natural language processing: A review. *IEEE Security & Privacy*.
- Li, Y.; Jiang, Y.; Li, Z.; and Xia, S.-T. 2022b. Backdoor Learning: A Survey. *IEEE Transactions on Neural Networks and Learning Systems*.
- Li, Y.; Lyu, X.; Koren, N.; Lyu, L.; Li, B.; and Ma, X. 2021. Neural Attention Distillation: Erasing Backdoor Triggers from Deep Neural Networks. In *ICLR*.
- Ma, Z.; Ma, J.; Miao, Y.; Liu, X.; Choo, K.-K. R.; and Deng, R. H. 2021. Pocket diagnosis: Secure federated learning against poisoning attack in the cloud. *IEEE Transactions on Services Computing*.
- Nguyen, T. D.; Nguyen, T.; Le Nguyen, P.; Pham, H. H.; Doan, K. D.; and Wong, K.-S. 2024. Backdoor attacks and defenses in federated learning: Survey, challenges and future research directions. *Engineering Applications of Artificial Intelligence*.
- Pan, X.; Zhang, M.; Sheng, B.; Zhu, J.; and Yang, M. 2022. Hidden Trigger Backdoor Attack on NLP Models via Linguistic Style Manipulation. In *USENIX Security Symposium*.
- Pei, H.; Jia, J.; Guo, W.; Li, B.; and Song, D. 2024. TextGuard: Provable Defense against Backdoor Attacks on Text Classification. In *NDSS*.
- Phuong, M.; and Lampert, C. 2019. Towards Understanding Knowledge Distillation. In *ICML*.
- Qi, F.; Chen, Y.; Li, M.; Yao, Y.; Liu, Z.; and Sun, M. 2021a. ONION: A Simple and Effective Defense Against Textual Backdoor Attacks. In *EMNLP*.
- Qi, F.; Li, M.; Chen, Y.; Zhang, Z.; Liu, Z.; Wang, Y.; and Sun, M. 2021b. Hidden Killer: Invisible Textual Backdoor Attacks with Syntactic Trigger. *ACL-IJCNLP*.
- Rahman, A.; Hossain, M. S.; Alrajeh, N. A.; and Alsolami, F. 2020. Adversarial examples—Security threats to COVID-19 deep learning systems in medical IoT devices. *IEEE Internet of Things Journal*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A.; and Potts, C. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *EMNLP*.
- Sun, X.; Li, X.; Meng, Y.; Ao, X.; Lyu, L.; Li, J.; and Zhang, T. 2023. Defending against Backdoor Attacks in Natural Language Generation. *AAAI*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive Representation Distillation. In *ICLR*.
- Tiwari, K.; Yuan, S.; and Zhang, L. 2022. Robust Hate Speech Detection via Mitigating Spurious Correlations. In *AAACL*.
- Vice, J.; Akhtar, N.; Hartley, R.; and Mian, A. 2024. Bagm: A backdoor attack for manipulating text-to-image generative models. *IEEE Transactions on Information Forensics and Security*.

- Wan, Y.; Qu, Y.; Ni, W.; Xiang, Y.; Gao, L.; and Hossain, E. 2024. Data and model poisoning backdoor attacks on wireless federated learning, and the defense mechanisms: A comprehensive survey. *IEEE Communications Surveys & Tutorials*.
- Wang, H.; Sreenivasan, K.; Rajput, S.; Vishwakarma, H.; Agarwal, S.; Sohn, J.-y.; Lee, K.; and Papailiopoulos, D. 2020. Attack of the tails: Yes, you really can backdoor federated learning. *NeurIPS*.
- Wang, J.; Xu, C.; Guzmán, F.; El-Kishky, A.; Tang, Y.; Rubinstein, B.; and Cohn, T. 2021. Putting words into the system’s mouth: A targeted attack on neural machine translation using monolingual data poisoning. In *ACL Findings*.
- Wang, J.; Xu, Q.; He, X.; Rubinstein, B.; and Cohn, T. 2024. Backdoor Attacks on Multilingual Machine Translation. In *NAACL*.
- Wu, B.; Chen, H.; Zhang, M.; Zhu, Z.; Wei, S.; Yuan, D.; and Shen, C. 2022. Backdoorbench: A comprehensive benchmark of backdoor learning. *NeurIPS*.
- Xi, Z.; Du, T.; Li, C.; Pang, R.; Ji, S.; Chen, J.; Ma, F.; and Wang, T. 2023. Defending Pre-trained Language Models as Few-shot Learners against Backdoor Attacks. In *NeurIPS*.
- Yan, J.; Gupta, V.; and Ren, X. 2022. BITE: Textual Backdoor Attacks with Iterative Trigger Injection. In *ACL*.
- Yang, W.; Lin, Y.; Li, P.; Zhou, J.; and Sun, X. 2021. RAP: Robustness-Aware Perturbations for Defending against Backdoor Attacks on NLP Models. In *EMNLP*.
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks? In *NeurIPS*.
- Zagoruyko, S.; and Komodakis, N. 2017. Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*.
- Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level Convolutional Networks for Text Classification. In *NeurIPS*.
- Zhang, Y.; Liu, L.; and Liu, L. 2023. Cuing Without Sharing: A Federated Cued Speech Recognition Framework Via Mutual Knowledge Distillation. *ACM MM*.
- Zhao, B.; and Han, K. 2021. Novel Visual Category Discovery with Dual Ranking Statistics and Mutual Knowledge Distillation.
- Zhao, S.; Gan, L.; Luu, A. T.; Fu, J.; Lyu, L.; Jia, M.; and Wen, J. 2024a. Defending Against Weight-Poisoning Backdoor Attacks for Parameter-Efficient Fine-Tuning. In *NAACL Findings*.
- Zhao, S.; Wu, X.; Nguyen, C.-D.; Jia, M.; Feng, Y.; and Tuan, L. A. 2024b. Unlearning Backdoor Attacks for LLMs with Weak-to-Strong Knowledge Distillation. *arXiv preprint arXiv:2410.14425*.
- Zhao, X.; Xu, D.; and Yuan, S. 2024. Defense against Backdoor Attack on Pre-trained Language Models via Head Pruning and Attention Normalization. In *ICML*.
- Zhao, Y.; Xu, R.; Wang, X.; Hou, P.; Tang, H.; and Song, M. 2020. Hearing Lips: Improving Lip Reading by Distilling Speech Recognizers. *AAAI*.
- Zhu, B.; Qin, Y.; Cui, G.; Chen, Y.; Zhao, W.; Fu, C.; Deng, Y.; Liu, Z.; Wang, J.; Wu, W.; et al. 2022. Moderate-fitting as a natural backdoor defender for pre-trained language models. *NeurIPS*.