ROBUST BACKDOOR REMOVAL BY RECONSTRUCTING TRIGGER-ACTIVATED CHANGES IN LATENT REPRESENTATION

Anonymous authorsPaper under double-blind review

ABSTRACT

Backdoor attacks pose a critical threat to machine learning models, causing them to behave normally on clean data but misclassify poisoned data into a poisoned class. Existing defenses often attempt to identify and remove backdoor neurons based on Trigger-Activated Changes (TAC) which is the activation differences between clean and poisoned data. These methods suffer from low precision in identifying true backdoor neurons due to inaccurate estimation of TAC values. In this work, we propose a novel backdoor removal method by accurately reconstructing TAC values in the latent representation. Specifically, we formulate the minimal perturbation that forces clean data to be classified into a specific class as a convex quadratic optimization problem, whose optimal solution serves as a surrogate for TAC. We then identify the poisoned class by detecting statistically small L^2 norms of perturbations and leverage the perturbation of the poisoned class in fine-tuning to remove backdoors. Experiments on CIFAR-10, GTSRB, and TinyImageNet demonstrated that our approach consistently achieves superior backdoor suppression with high clean accuracy across different attack types, datasets, and architectures, outperforming existing defense methods.

1 Introduction

While machine learning provides significant benefits in many applications, the threat of backdoor attacks that compromise machine learning models has been pointed out (Gu et al., 2019; Chen et al., 2017; Nguyen & Tran, 2021). The compromised model behaves normally on clean data, but when a trigger known only to the adversary is embedded into the data (poisoned data), the model is forced to misclassify it as the attacker-specified target class. One of the most critical challenges in backdoor defenses is to develop backdoor removal methods that effectively eliminate the influence of backdoor attacks from a compromised model while preserving its original accuracy (Liu et al., 2018a; Zheng et al., 2022; Lin et al., 2024).

To minimize accuracy degradation, most backdoor removal methods first identify backdoor neurons that strongly respond to the trigger and are thus thought to be less essential for normal predictions but critical for backdoor success. Once identified, the influence (impact) of these neurons is mitigated through pruning, fine-tuning or both (Liu et al., 2018a; Zheng et al., 2022; Wu & Wang, 2021; Li et al., 2023; Lin et al., 2024). A key metric to measure the degree of their contribution is Trigger-Activated Changes (TAC) (Zheng et al., 2022), defined as the difference in neuron activations between clean and poisoned data. Removing neurons exhibiting higher TAC values can eliminate backdoors while minimizing the impact on accuracy (Zheng et al., 2022; Lin et al., 2024).

However, since poisoned data is not available in practice, the ideal values of TAC cannot be obtained. Due to this limitation, existing methods (Liu et al., 2018a; Zheng et al., 2022; Wu & Wang, 2021; Li et al., 2023; Lin et al., 2024) compute the contribution of neurons to the success of backdoor attacks using their own approaches, but their results often show low consistency with TAC, leading to ineffective backdoor removal.

To address this problem, we propose a novel backdoor removal method by accurately reconstructing the effects of TAC in the latent representation with an overview provided in Figure 1. Among intermediate layers, TAC in the latent representation, i.e, the output of the layer just before the

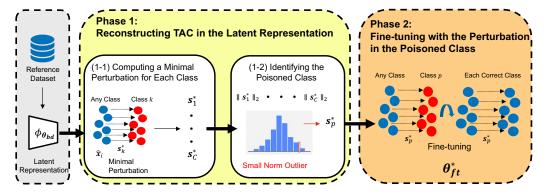


Figure 1: Overview of our proposed method. Our method consists of two stages: (1) reconstructing TAC in the latent representation, which involves computing the minimal perturbation that forces any clean data to be classified into each class and then identifying the poisoned class based on L^2 norms of the optimized perturbations, and (2) removing the backdoor by fine-tuning with the optimized perturbation of the poisoned class.

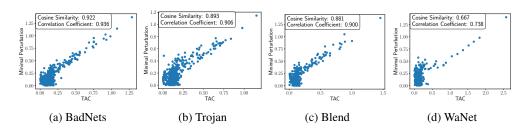


Figure 2: The perturbations obtained by our method and TAC in the latent representation for CIFAR-10 on ResNet-18. For each neuron in the latent representation, we plot the TAC value on the horizontal axis and the minimal perturbation of the poisoned class on the vertical axis.

classification layer, can be critical for the success of backdoor attacks because the effects of TAC in earlier layers propagate and accumulate in the latent representation, which then directly affects misclassification through the classification layer. If the effects of TAC in the latent representation can be inferred solely from clean data, defenders can approximate the model's outputs on poisoned data without them and eliminate their influence from the model. Thus, reconstructing TAC in the latent representation enables robust backdoor removal.

Specifically, we first reconstruct TAC in the latent representation by computing a minimal perturbation in that representation required to misclassify any clean data into the poisoned class. This is motivated by two key properties of TAC in the latent representation: (i) because triggers are realized through minimal modifications to clean data in order to remain undetectable, the resulting changes (i.e., TAC) in the latent representation between clean and poisoned inputs are necessarily small; and (ii) despite being minimal, these changes are sufficient to induce misclassification into the poisoned class. Actually, Figure 2 shows that the minimal perturbation obtained in this way is strongly similar and correlated with TAC in the latent representation. We then apply the obtained perturbation for model fine-tuning, which effectively removes the backdoor while preserving clean accuracy.

Our main contributions are summarized as follows:

1. Method to Reconstruct TAC in the Latent Representation. We propose a method to reconstruct TAC in the latent representation by computing the minimal perturbation that forces any clean data to be misclassified into a specific class and identifying the poisoned class from the perturbations in all classes. First, we formulate the optimization problem of finding such a perturbation as a convex quadratic program. We then clarify the conditions under which such a perturbation exists and derive the analytical solution. For the poisoned class, the perturbation obtained by solving the optimization can be regarded as a surrogate that reproduces the effect of TAC. Therefore, reconstructing

TAC in the latent representation requires identifying the poisoned class, even though in practice the defender typically does not know it in advance.

- 2. Method to Identify the Poisoned Class via Perturbation Norms. We propose a method to identify the poisoned class by selecting the perturbation whose L^2 norm is unusually small and can be detected as a statistical outlier. This is supported by the fact that backdoor training forces data with triggers to be classified into the poisoned class by effectively shifting its decision boundary toward the region of clean data. Consequently, the minimal perturbation required to misclassify clean data into the poisoned class tends to be smaller than that for other classes.
- **3.** Backdoor Removal Method from the Optimized Perturbation. We propose a backdoor removal method that leverages the TAC effects estimated by our method of the poisoned class. Concretely, we fine-tune the model using a loss that enforces clean data in the latent representation, even when perturbed toward the poisoned class, to be classified into their original clean classes, together with the cross-entropy loss for the clean task. This process yields a compromised model that simultaneously preserves high accuracy and enhances backdoor removal performance. Experimental results demonstrate that our method can successfully eliminate the impact of backdoor attacks while maintaining high accuracy, even against several representative attack methods. Furthermore, we confirm that our approach achieves greater robustness compared to existing defense methods.

2 RELATED WORKS

2.1 BACKDOOR ATTACKS

A backdoor attack compromises a model so that it behaves normally on clean data but misclassifies poisoned data into an attacker-specified class. Representative methods include BadNets (Gu et al., 2019), Blend (Chen et al., 2017), and Trojan (Liu et al., 2018b). Although these approaches achieve high attack success rates, they are relatively easy to detect because of their easily visible triggers. To reduce the detectability of visible triggers, several studies design imperceptible triggers such that the difference between clean and poisoned data cannot be distinguished by humans or detectors (Nguyen & Tran, 2020; 2021; Doan et al., 2021b). More recently, techniques have also been developed to improve stealthiness not only at the input level but also in the internal feature space of the model (Tan & Shokri, 2020; Zhong et al., 2022; Doan et al., 2021a; Xu et al., 2025). In this way, backdoor attacks continue to evolve toward greater stealthiness in both input and internal space, thereby increasing the difficulty of effective defense.

2.2 BACKDOOR REMOVAL

Existing backdoor removal methods can broadly be categorized into two groups: (i) those that identify backdoor neurons and then prune or fine-tune them (Liu et al., 2018a; Zheng et al., 2022; Wu & Wang, 2021; Li et al., 2023; Lin et al., 2024), and (ii) those that neutralize backdoors via advanced fine-tuning strategies without explicit neuron identification (Zhu et al., 2023; Min et al., 2023; Wei et al., 2023; Karim et al., 2024). Details of the latter related works are provided in Appendix B.1.

To identify backdoor neurons, various methods have been proposed. Fine-Pruning (FP) (Liu et al., 2018a) regards neurons inactive on clean data as backdoor neurons, while Adversarial Neuron Pruning (ANP) (Wu & Wang, 2021) regards neurons sensitive to adversarial noise as backdoor neurons. As an oracle metric, Channel Lipschitzness Pruning (CLP) (Zheng et al., 2022) introduced Trigger-Activated Changes (TAC), defined as the activation difference between clean and poisoned data. CLP further approximates neurons with large weight values as those with large TAC. However, because TAC computation requires access to poisoned data, it is impossible to obtain the ideal values of TAC. More recently, unlearning-based methods using only clean data (Li et al., 2023; Lin et al., 2024) have been proposed, but the identification rate of neurons with high TAC values still remains limited. If TAC could be computed more precisely, it would enable approximate the model outputs of poisoned data without them and thus achieve robust backdoor removal. However, in the absence of poisoned data, directly leveraging TAC is infeasible, leaving the construction of practical surrogates of TAC for defenders as an open challenge. To address this challenge, our approach reconstructs TAC in the latent representation via optimizing a minimal perturbation that forces any clean data to be misclassified into a specific class, providing a feasible and accurate method for defenders to neutralize backdoor effects.

3 PROBLEM SETTING

In this section, we first describe the threat model in this work, focusing on the goals and capabilities of the adversary and the defender. We then present the formalization of backdoor attacks and introduce Trigger-Activated Changes (TAC) (Zheng et al., 2022).

3.1 THREAT MODEL

Adversary. The adversary's goal is to obtain a compromised model that, with high probability, misclassifies poisoned data into the target class while still correctly classifying clean data. In the data collection scenario (Gu et al., 2019; Chen et al., 2017; Liu et al., 2020), the adversary has access only to the training dataset. In the supply-chain scenario (Doan et al., 2021b; Nguyen & Tran, 2021; Xu et al., 2025), where the model is distributed through external sources, the adversary may have full access to the training process.

Defender. The defender's goal is to detect whether a given model has been compromised and to remove the backdoor if present. The defender is assumed to have access to the model parameter and a small dataset (reference dataset) sampled from the same distribution as the model's training data (Zhu et al., 2023; Lin et al., 2024).

3.2 FORMULATION OF BACKDOOR ATTACKS

For any $a \in \mathbb{N}$, let $[a] = \{1, 2, \cdots, a\}$. Given an input dimension d_{in} and the number of classes C, we denote by $e_i \in \{0, 1\}^C$ the standard basis vector whose i-th element is 1. A neural network $f: \mathbb{R}^{d_{\text{in}}} \to [0, 1]^C$ outputs the probability of belonging to each class for an input $\boldsymbol{x} \in \mathbb{R}^{d_{\text{in}}}$. Let $\boldsymbol{\theta}$ be a model parameter, ℓ a loss function, and $\phi_{\boldsymbol{\theta}}: \mathbb{R}^{d_{\text{in}}} \to \mathbb{R}^{d_{\text{emb}}}$ the mapping to the latent representation layer (i.e., the layer just before the final linear layer) of dimension d_{emb} . This yields the latent representation $\hat{\boldsymbol{x}} = \phi_{\boldsymbol{\theta}}(\boldsymbol{x}) \in \mathbb{R}^{d_{\text{emb}}}$. The final (L-th) linear layer is parameterized by weight matrix $\boldsymbol{W}_L = [\boldsymbol{w}_1, \boldsymbol{w}_2, \cdots, \boldsymbol{w}_C] \in \mathbb{R}^{d_{\text{emb}} \times C}$, where each column vector is $\boldsymbol{w}_j \in \mathbb{R}^{d_{\text{emb}}}$ and a bias vector is $\boldsymbol{b} \in \mathbb{R}^C$. Using the softmax function, the network output is expressed as $f(\boldsymbol{x}; \boldsymbol{\theta}) = \operatorname{Softmax}(\boldsymbol{W}_L^T \hat{\boldsymbol{x}} + \boldsymbol{b})$.

Furthermore, let $\delta \in \mathbb{R}^{d_{\text{in}}}$ be the trigger required for a backdoor attack and $p \in [C]$ be a poisoned class. Then, the compromised model parameter θ_{bd} is obtained as

$$\boldsymbol{\theta}_{bd} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^{n} \left[\ell(f(\boldsymbol{x}_i; \boldsymbol{\theta}), \boldsymbol{y}_i) + \ell(f(\boldsymbol{x}_i + \boldsymbol{\delta}; \boldsymbol{\theta}), \boldsymbol{e}_p) \right], \tag{1}$$

where the training dataset is $D = \{(x_i, y_i)\}_{i=1}^n$ and $y_i \in \{e_1, e_2, \dots, e_C\}$. The parameter θ_{bd} is optimized such that the model behaves normally on clean data x, while poisoned data $x + \delta$ are misclassified into the poisoned class p.

3.3 TRIGGER-ACTIVATED CHANGES

In a compromised model, when poisoned data $x + \delta$ is provided, certain neurons are strongly activated. This excessive activation causes $x + \delta$ to be misclassified into the poisoned class p. Therefore, if the contribution of each neuron to the success of the backdoor attack can be quantified, its influence can be suppressed, enabling backdoor removal from the model.

In this paper, we focus on Trigger-Activated Changes (TAC) (Zheng et al., 2022), which are defined as the difference in activations between clean and poisoned data and serve as an oracle metric to quantify each neuron's contribution to the success of backdoor attacks. Specifically, for the i-th neuron in the l-th layer $f_{l,i}(\cdot)$, TAC is computed as

$$TAC_{l,i}(x;\theta) = f_{l,i}(x+\delta;\theta) - f_{l,i}(x;\theta).$$
(2)

The *i*-th neuron's importance in the intermediate layers for the success to backdoor attack is calculated as the average value, $\text{TAC}_{l,i}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{x}}[\text{TAC}_{l,i}(\boldsymbol{x};\boldsymbol{\theta})]$, because applying the same trigger to different data tends to activate similar neurons in the intermediate layers (Zheng et al., 2022).

However, the computation of TAC requires poisoned data $x + \delta$, and since defenders typically do not know the trigger δ , it is infeasible to calculate the ideal values of $TAC_{l,i}(\theta)$.

4 PROPOSED METHOD

In this paper, we aim to reconstruct TAC in the latent representation of a compromised model instead of reconstructing TAC in arbitrary intermediate layers. The rationale is that although backdoor neurons may appear in arbitrary intermediate layers, their effects are aggregated through the network and ultimately reflected in the latent representation. Thus, if TAC in the latent representation can be reconstructed, the output of poisoned data can be approximately computed using the subsequent linear layer as follows: $f(x + \delta; \theta) = \operatorname{Softmax}(W_L(\hat{x} + \operatorname{TAC}_{L-1}(\theta)) + b)$, where $\operatorname{TAC}_{L-1}(\theta) \in \mathbb{R}^{d_{\text{emb}}}$ denotes the vector of TAC values in the latent representation. This allows us to remove backdoors by fine-tuning the model so that the misclassification of poisoned data is restored into the correct class.

Based on this idea, we propose a method to reconstruct TAC in the latent representation and a defense mechanism that leverages the reconstructed TAC for backdoor removal. As illustrated in Figure 1, our method consists of two stages: (1) reconstructing TAC in the latent representation, which involves computing the minimal perturbation that forces any clean data to be classified into each class and then identifying the poisoned class from the optimized perturbations, and (2) removing the backdoor using the optimized perturbation of the poisoned class. The details of each stage are described below.

4.1 Computing Perturbations in the Latent Representation

To reconstruct TAC in the latent representation, we focus on the following two properties of TAC in the latent representation: it takes minimal values since the trigger is minimized to be indistinguishable from the original data, and it induces misclassification into the poisoned class. Based on these observations, we first introduce an optimization problem to compute the minimal perturbation in the latent representation of clean data that forces it to be misclassified into a specific class.

Optimization Problem. Our goal is to find the minimal perturbation s_k^* that guarantees all inputs are classified into class k. This leads to the following formulation: the objective is defined by a quadratic term $\frac{1}{2} \|s_k\|_2^2$ for analytical convenience, such that the logits $s_k + \hat{x}_i$ of class k dominate those of all other classes. The resulting primal optimization problem can be formulated as the following convex quadratic program:

$$\boldsymbol{s}_{k}^{*} = \underset{\boldsymbol{s}_{k}}{\operatorname{argmin}} \ \frac{1}{2} \|\boldsymbol{s}_{k}\|_{2}^{2} \text{s.t.} \quad (\boldsymbol{w}_{k} - \boldsymbol{w}_{j})^{\top} (\boldsymbol{s}_{k} + \hat{\boldsymbol{x}}_{i}) \geq 0, \ \forall j \in [C] \setminus \{k\}, \ \forall i \in [n].$$
 (3)

Here, $(\boldsymbol{w}_k - \boldsymbol{w}_j)^{\top}(s_k + \hat{\boldsymbol{x}}_i)$ denotes the margin of class k against class j for sample i after applying the perturbation s_k . An example x_i is classified into class k if and only if these margins are nonnegative for all $j \neq k$. Therefore, the constraints enforce nonnegative margins for every example and every $j \neq k$, and the single perturbation s_k is chosen to lift the margins of class k simultaneously across all examples. To remove redundancy and improve computational efficiency, the n(C-1) constraints in equation 3 are compressed into C-1 constraints by considering only the worst-case margin for each class $j \neq k$ across the dataset. That is, the constraint in equation 3 can be equivalently written as $(\boldsymbol{w}_k - \boldsymbol{w}_j)^{\top} s_k \geq -(\boldsymbol{w}_k - \boldsymbol{w}_j)^{\top} \hat{\boldsymbol{x}}_i$, $\forall j \in [C] \setminus \{k\}$, $\forall i \in [n]$ and it suffices to consider only the worst case $\forall j \in [C] \setminus \{k\} : \max_i \{-(\boldsymbol{w}_k - \boldsymbol{w}_j)^{\top} \hat{\boldsymbol{x}}_i\}$. The problem therefore reduces to the following convex quadratic program:

$$s_k^* = \underset{s_k}{\operatorname{argmin}} \frac{1}{2} ||s_k||_2^2 \quad \text{s.t.} \quad U_k W_L^{\top} s_k \ge m,$$
 (4)

where the inequality between vectors is understood element-wise, $\boldsymbol{U}_k \coloneqq [\boldsymbol{u}_1, \boldsymbol{u}_2, \cdots, \boldsymbol{u}_{C-1}] \in \mathbb{R}^{(C-1) \times C}, \forall j \in [C] \setminus \{k\} : \boldsymbol{u}_j = (\boldsymbol{e}_k - \boldsymbol{e}_j)^\top \in \mathbb{R}^C \text{ and } \boldsymbol{m} \in \mathbb{R}^{C-1} \text{ is the vector of worst-case margins, with each component given by } \forall j \in [C] \setminus \{k\} : \boldsymbol{m}_j = \max_j \{-(\boldsymbol{w}_k - \boldsymbol{w}_j)^\top \hat{\boldsymbol{x}}_i\}.$

The reduced problem is also convex by construction but its feasibility is not always guaranteed. Thus, we provide sufficient conditions under which feasibility is guaranteed from Theorem 1. That is, if $C-1 < d_{\rm emb}$ and $U_k W_L^{\top}$ has full row rank, the optimal solution s_k^* is guaranteed to exist.

Solution via Dual Problem. To obtain the optimal solution for s_k , we introduce the dual problem of equation 4 because the dual problem involves fewer variables, which makes the problem more stable

compared to the primal problem. Let $\lambda \in \mathbb{R}^{C-1}$ be the dual variable vector and $V_k := U_k W_L^{\top} \in \mathbb{R}^{(C-1) \times d_{\text{emb}}}$. The final form of the dual problem can be written as follows, with the derivation process provided in Appendix D.2:

$$\lambda^* = \underset{\lambda}{\operatorname{argmax}} \quad \lambda^\top m - \frac{1}{2} \| V_k^\top \lambda \|_2^2 \quad \text{s.t.} \quad \lambda \ge 0.$$
 (5)

In general, the dual problem provides a lower bound on the optimal value of the primal problem. When the primal problem is convex and satisfies suitable regularity conditions (e.g., Slater's condition (Boyd & Vandenberghe, 2004)), strong duality holds, and the optimal values of the primal and dual problems coincide. The proof that strong duality for the derived primal and dual problems in equation 4 and equation 5 is given in Appendix D.4. When strong duality holds, the following Karush–Kuhn–Tucker (KKT) conditions are necessary and sufficient for optimality of the primal problem (Boyd & Vandenberghe, 2004):

- (i) Stationarity: $s_k^* V^\top \lambda^* = 0$,
- (ii) Primal and Dual Feasibility: $V^{ op}s_k^* \geq m, \lambda^* \geq 0,$
- (iii) Complementary Slackness: $\lambda^* \odot (V^{\top} s_k^* m) = 0$.

The conditions of primal and dual feasibility together with complementary slackness ensure that s_k^* necessarily induces misclassification into class k. As a result, the minimal perturbation s_k^* can be obtained from the stationarity condition, i.e., $s_k^* = V^\top \lambda^*$, which shows that the primal optimal solution can be obtained from the dual optimal solution. In practice, we solve the dual problem with a convex optimization solver CVXPY (Diamond & Boyd, 2016) to obtain λ^* reliably.

4.2 IDENTIFYING THE POISONED CLASS

Using the perturbations for each class obtained in Section 4.1, we propose a method to identify the poisoned class. Since TAC in the latent representation is reconstructed as the perturbation s_p^* for the poisoned class $p \in [C]$, it is necessary to identify the poisoned class.

To this end, we focus on the perspective of L^2 minimization of the perturbations. The minimal displacement required to switch class is proportional to the margin to the decision boundary. Backdoor training tends to pull the poisoned decision boundary closer to the data space with clean classes, which causes the perturbation of the poisoned class to become smaller than those of the other classes. Therefore, the poisoned class can be identified by detecting outliers on the smaller side.

The procedure is as follows. First, we standardize $||s_1||_2, ||s_2||_2, \cdots, ||s_C||_2$ using their mean μ and standard deviation σ , obtaining z_1, \cdots, z_C , where for each $k \in [C]$ we compute $z_k = (||s_k||_2 - \mu)/\sigma$. Then, given an outlier threshold α , we identify the poisoned class p as the class k such that $z_k < \alpha$. Here, α denotes the threshold for detecting outliers in the L^2 norms, and it should be chosen depending on the dataset. The impact of identifying s_p^* by the poisoned class identification method on our fine-tuning performance is evaluated in an ablation study, as reported in Appendix E.2.

4.3 BACKDOOR REMOVAL WITH THE PERTURBATION OF THE POISONED CLASS

Using the perturbation of the poisoned class obtained in Section 4.2, we propose a backdoor removal method, as shown in equation 6:

$$\theta_{\text{ft}}^* = \underset{\boldsymbol{\theta}_{\text{bd}}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \left[\ell(f(\boldsymbol{x}_i; \boldsymbol{\theta}_{\text{bd}}), \boldsymbol{y}_i) + \beta \, \ell(\operatorname{Softmax}(\boldsymbol{W}_L^\top (\hat{\boldsymbol{x}}_i + \boldsymbol{s}_p^*) + \boldsymbol{b}), \boldsymbol{y}_i) \right], \tag{6}$$

where β is a hyperparameter that balances accuracy and backdoor removal performance. Specifically, the model is fine-tuned so that even if the latent representation shifts in the direction of s_p^* , the perturbed latent representation $\hat{x}_i + s_p^*$ is still recognized as its correct class. This ensures that poisoned data are classified into their correct clean classes. In addition, to maintain performance on the original clean task, a loss that enforces correct classification of clean data into their correct classes is also included.

While pruning-based approaches via the perturbation of the poisoned class are also possible, we found that our fine-tuning method is more effective performance as shown in Appendix E.3.

Table 1: The smallest and second-smallest z values ($z^{(1)}$ and $z^{(2)}$) for each attack method, together with their corresponding classes ($\text{Class}^{(1)}$ and $\text{Class}^{(2)}$). "Clean" shows the result of the clean model without any attack.

		CIFA	R-10			GT	SRB		TinyImageNet				
	z ⁽¹⁾	Class ⁽¹⁾	$z^{(2)}$	Class ⁽²⁾	$z^{(1)}$	Class ⁽¹⁾	$z^{(2)}$	Class ⁽²⁾	$z^{(1)}$	Class ⁽¹⁾	$z^{(2)}$	Class ⁽²⁾	
Clean	-1.48	8	-1.25	3	-1.52	2	-1.46	39	-2.45	160	-2.35	132	
BadNets	-2.35	1	-0.44	4	-2.77	1	-1.40	39	-6.39	1	-2.48	160	
Trojan	-2.50	1	-0.55	3	-2.66	1	-1.43	2	-4.98	1	-2.17	132	
Blend	-2.60	1	-0.47	6	-2.99	1	-1.25	39	-5.36	1	-2.61	160	
WaNet	-2.91	1	-0.14	9	-3.88	1	-1.27	14	-6.43	1	-1.94	160	
IAB	-2.92	1	-0.15	3	-4.65	1	-1.42	2	-6.85	1	-2.58	176	
Lira	-2.62	1	-0.34	3	-2.86	1	-1.52	2	-4.18	1	-2.07	160	

5 EXPERIMENTS

In this section, we conduct experimental evaluations to verify the effectiveness of our proposed method described in Section 4.

5.1 EXPERIMENTAL SETUP

Datasets and Neural Network Architecture. We conducted experiments on three image classification datasets: CIFAR-10, GTSRB, and TinyImageNet. CIFAR-10 and GTSRB contain 10 and 43 classes of 32×32 pixels, respectively. TinyImageNet includes 200 classes with images resized to 64×64 pixels. For all datasets, we primarily used ResNet-18 as the neural network architecture.

Backdoor Attacks. We evaluate the effectiveness of our proposed method against six backdoor attack methods: BadNets (Gu et al., 2019), Trojan (Liu et al., 2018b), Blend (Chen et al., 2017), IAB (Nguyen & Tran, 2020), Lira (Doan et al., 2021b), and WaNet (Nguyen & Tran, 2021). The training configuration for all attacks consisted of 100 epochs, stochastic gradient descent (SGD) as the optimizer, a learning rate of 0.1, and cosine annealing as the learning rate scheduler. For the standard backdoor attack configuration, we adopted a poisoning rate of 0.1 and fixed the poisoned class as 1, following the all-to-one setting in which poisoned data from all other classes is misclassified into a poisoned class. We remark that although datasets such as CIFAR-10 index classes starting from 0, we align with the notation in this paper where classes are indexed from 1. Accordingly, class 1 in our notation corresponds to class 0 in CIFAR-10. Further details of each attack and the hyperparameters used are provided in Appendix C.1.

Backdoor Defenses. For comparison, we evaluate our proposed method against five defense methods to identify backdoor neurons, FP (Liu et al., 2018a), CLP (Zheng et al., 2022), ANP (Wu & Wang, 2021), RNP (Li et al., 2023) and TSBD (Lin et al., 2024) as well as three advanced fine-tuning defenses without identifying backdoor neurons, FT-SAM (Zhu et al., 2023), SAU (Wei et al., 2023) and FST (Min et al., 2023). Details of the defense methods and hyperparameters are provided in Appendix C.2. Following previous works (Zhu et al., 2023; Lin et al., 2024), we assume that the defender has access to 5% of the training dataset as a reference dataset and the effect of the reference dataset size on defense performance is presented in Appendix E.4. For our proposed method, the hyperparameter α for poisoned class identification was set to -2.0 for CIFAR-10 and GTSRB, and -3.5 for TinyImageNet. For fine-tuning in backdoor removal, we used SGD with a learning rate of 0.01 for 50 epochs, with β set to 0.5 for CIFAR-10, 2.0 for GTSRB, and 0.1 for TinyImageNet. The effect of the β value on ACC and ASR is discussed in Appendix E.5.

5.2 RECONSTRUCTING TAC IN THE LATENT REPRESENTATION

As shown in Figure 2 and Appendix E.1, the perturbation of the poisoned class computed by our method exhibits a high similarity with TAC in the latent representation and identifies backdoor neurons more accurately than existing approaches. Therefore, it is crucial to accurately identify the poisoned class by our poisoned class identification method described in Section 4.2.

Main Results. Table 1 shows the smallest and second-smallest standardized L^2 norms of the perturbation for each class. Across all attack methods, the z value of the poisoned class is the smallest among all classes and remains markedly lower even compared to the second-smallest class. Since

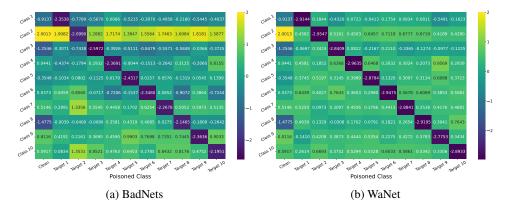


Figure 3: Standardized L^2 norms of perturbations for each poisoned class and clean classes for CIFAR-10. The horizontal axis shows the poisoned class, while the vertical axis shows the L^2 norm of perturbations for each class.

the values are below the dataset-specific thresholds α , the poisoned classes are successfully and accurately identified.

Different Poisoned Classes. To examine the robustness of our poisoned class identification method, we further evaluate whether our proposed method can detect any poisoned class as shown in Figure 3. In general, the perturbation L^2 norm for the poisoned class is smaller than those of the clean classes, with standardized values falling below -2.0.

5.3 EFFECTIVENESS OF BACKDOOR REMOVAL

Evaluation Metrics. Following previous works (Zhu et al., 2023; Lin et al., 2024), we introduce three evaluation metrics for backdoor attacks. Accuracy (ACC), which measures the classification accuracy on clean data; Attack Success Rate (ASR), which denotes the percentage of triggered inputs classified into the poisoned class; and Defense Efficacy Rate (DER), which evaluates how effectively the backdoor is removed while maintaining accuracy. DER is defined as DER = $(\max(0, \Delta ACC) - \max(0, \Delta ASR + 1))/2$, which takes values in [0, 1]. A DER closer to 1 indicates that the backdoor is more effectively removed while preserving clean accuracy.

Main Results. Table 2 presents the backdoor removal results for various attack and defense methods on ResNet-18. Overall, our method achieves consistently superior DER across datasets. On Tiny-ImageNet, it attains the highest DER of 98.06%, with competitive ACC and ASR. On GTSRB, our approach yields a DER substantially higher than competing defenses, reflecting both strong attack suppression and accuracy preservation. On CIFAR-10, it achieves the best DER of 98.39% on Blend and remains competitive across other attacks. In addition, the results on CIFAR-10 with ResNet-50 are reported in Table 3. On ResNet-50, our method outperforms all defenses, including FT-SAM and SAU, which performed well on ResNet-18, in terms of both ACC and ASR. These results indicate that our approach suppresses backdoor success to near-zero while preserving high clean accuracy across diverse attack types, architectures, and datasets compared to the state-of-the-art existing defense methods.

6 CONCLUSION

In this work, we introduced a novel backdoor removal framework that reconstructs Trigger-Activated Changes (TAC) in the latent representation and leverages the reconstructed TAC for effective backdoor removal. Our method consists of two stages: recontructing TAC in the latent representation by computing minimal perturbations which misclassify any clean data into a target class for all classes and identifying the poisoned class via the perturbation norms, and fine-tuning the model using the optimized perturbation of the poisoned class. Our experiments demonstrated that our method achieves superior backdoor suppression while maintaining high clean accuracy in any attack type, dataset, and architecture. As future work, we aim to extend our framework to settings with multiple

Table 2: Comparison of the backdoor removal results. "Average" denotes the mean of each metric across attack methods. "No Defense" refers to a model to which no defense method is applied so DER is marked as "—".

		,	No Defens	۵.		FP			CLP			ANP			RNP	
		ACC ↑	ASR J	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR J	DER ↑
	BadNets	93.81	100.00	- DEK	93.68	100.00	49.94	91.30	33.10	82.19	87.89	2.67	95.46	93.17		95.98
	Trojan	94.00	100.00	_	93.47	2.11	98.68	83.97	1.19	94.39	88.98	100.00	47.35	93.97		49.86
	Blend	93.29	99.91	_	93.47	13.96	92.89	90.09	30.46	83.13	89.72	91.12	52.14	93.43		55.03
	WaNet	93.41	99.59	_	93.11	0.79	92.89 99.37		100.00	8.41	92.73	1.48	98.47	93.40		55.66
		93.41				0.79	99.16	10.23 90.18	7.39		89.13	0.83	96.99	93.57		99.00
	IAB		98.81	-	93.40					94.02						
	Lira	94.29	99.98	-	93.88	0.24	99.66	88.61	3.86	95.22	90.82	99.93	48.29	93.82		54.27
	Average	93.73	99.71	-	93.48	19.57	89.95	75.73	29.33	76.23	89.88	49.34	73.12	93.56		68.30
			TSBD			FT-SAM			SAU			FST				
CIFAR-10	72 121	ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR↓	DER ↑	ACC ↑		DER ↑
	BadNets	29.22	91.38	22.02	92.85	2.78	98.13	85.84	1.29	95.37	93.53	100.00	49.86	92.03		93.67
	Trojan	90.12	3.57	96.28	92.88	2.06	98.41	90.65	1.71	97.47	93.53	80.29	59.62	92.01		98.52
	Blend	88.76	3.66	95.86	92.71	4.50	97.42	90.65	0.84	98.21	93.11	36.19	81.77	91.84		98.39
	WaNet	88.14	83.56	55.38	92.57	1.28	98.74	91.48	1.57	98.05	93.33	2.23	98.64	92.39		99.02
	IAB	90.21	8.37	93.54	93.01	1.01	98.62	90.23	0.58	97.45	93.24	0.98	98.75	92.37		98.62
	Lira	92.10	93.77	52.01	93.12	0.50	99.15	90.93	0.86	97.88	93.88	19.88	89.84	92.80		99.19
	Average	79.76	47.38	69.18	92.86	2.02	98.41	89.96	1.14	97.40	93.44	39.93	79.75	92.24		97.90
]	No Defens			FP			CLP			ANP			RNP	
		ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR↓	DER ↑	ACC ↑		DER ↑
	BadNets	95.08	100.00	-	95.31	100.00	50.00	93.74	92.98	52.84	89.53	100.00	47.78	80.69		91.83
	Trojan	94.39	100.00	-	94.81	99.99	50.00	92.98	0.18	99.21	87.00	99.82	46.56	85.92		95.93
	Blend	93.85	99.50	-	94.68	97.84	50.83	92.26	99.41	49.25	82.53	96.87	45.80	93.45	83.75	57.82
	WaNet	93.99	97.07	-	95.76	10.79	93.14	20.59	100.00	13.30	84.39	0.00	95.16	87.75	0.00	96.84
	IAB	94.09	97.22	-	94.25	75.58	60.82	92.76	7.85	94.02	81.78	33.72	77.00	92.28	0.00	99.11
	Lira	93.97	99.91	_	94.18	7.99	95.96	92.16	11.32	93.39	79.94	25.95	79.96	85.52	0.00	95.73
	Average	94.23	98.95	_	94.83	65.36	66.79	80.75	51.96	67.00	84.20	59.39	65.38	87.60	14.45	89.54
			TSBD			FT-SAM			SAU			FST			Ours	
GTSRB		ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑
GISKD	BadNets	94.81	100.00	49.86	95.15	100.00	50.00	94.37	0.00	99.64	89.49	100.00	47.20	94.27	6.78	96.20
	Trojan	93.44	99.99	49.53	94.25	43.53	78.17	92.26	0.09	98.89	90.04	95.40	50.13	93.40	0.49	99.27
	Blend	93.47	78.62	60.25	94.57	60.78	69.36	94.03	0.40	99.55	89.45	17.73	88.68	93.39	7.06	95.99
	WaNet	91.82	94.20	50.35	95.36	0.01	98.53	95.04	0.03	98.52	89.81	0.01	96.44	95.60	0.00	98.54
	IAB	94.24		95.13	04 40	0.23	00.50	94.30	0.02	00.60	90.47					00.64
	Lira		6.96	93.13	94.48	0.23	98.50	94.50	0.02	98.60		0.00	96.80	94.21	0.00	98.61
		91.34	6.96 70.18	63.55	93.60	0.23	98.50 99.77	87.08	0.02	96.50	88.84	0.00 0.01	96.80 97.39	94.21	0.00 0.00	98.61
	Average	91.34 93.18	70.18	63.55 61.44	93.60	0.00	99.77	87.08	0.01	96.50	88.84	0.01	97.39	92.79	0.00	99.37
		91.34 93.18	70.18 74.99	63.55 61.44	93.60	0.00 34.09	99.77	87.08	0.01 0.09	96.50	88.84	0.01 35.53	97.39	92.79	0.00 2.39 RNP ASR ↓	99.37
		91.34 93.18	70.18 74.99 No Defens	63.55 61.44	93.60 94.57	0.00 34.09 FP	99.77 82.39	87.08 92.85	0.01 0.09 CLP	96.50 98.62	88.84 89.68	0.01 35.53 ANP	97.39 79.44	92.79 93.94	0.00 2.39 RNP	99.37 98.00
	Average	91.34 93.18 ACC↑	70.18 74.99 No Defense ASR ↓	63.55 61.44 e DER ↑	93.60 94.57	0.00 34.09 FP ASR↓	99.77 82.39 DER ↑	87.08 92.85	0.01 0.09 CLP ASR ↓	96.50 98.62 DER ↑	88.84 89.68	0.01 35.53 ANP ASR ↓	97.39 79.44 DER ↑	92.79 93.94	0.00 2.39 RNP ASR ↓	99.37 98.00 DER ↑
	Average BadNets	91.34 93.18 ACC ↑ 61.98 61.58	70.18 74.99 No Defens ASR ↓ 99.97	63.55 61.44 e DER ↑	93.60 94.57 ACC ↑ 58.26 58.07	0.00 34.09 FP ASR↓ 99.94	99.77 82.39 DER ↑ 48.16	87.08 92.85 ACC↑ 35.09	0.01 0.09 CLP ASR ↓ 0.25	96.50 98.62 DER ↑ 86.41	88.84 89.68 ACC ↑ 43.78	0.01 35.53 ANP ASR ↓ 91.76	97.39 79.44 DER ↑ 44.91	92.79 93.94 ACC ↑ 61.19	0.00 2.39 RNP ASR↓ 0.09	99.37 98.00 DER↑ 99.45
	Average BadNets Trojan Blend	91.34 93.18 ACC ↑ 61.98 61.58 62.28	70.18 74.99 No Defense ASR↓ 99.97 100.00 99.97	63.55 61.44 e DER ↑	93.60 94.57 ACC ↑ 58.26 58.07 57.68	0.00 34.09 FP ASR↓ 99.94 99.12 0.43	99.77 82.39 DER ↑ 48.16 48.69 97.47	87.08 92.85 ACC ↑ 35.09 59.59 54.29	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52	96.50 98.62 DER ↑ 86.41 98.89 91.23	88.84 89.68 ACC ↑ 43.78 36.96 28.03	0.01 35.53 ANP ASR ↓ 91.76 100.00 93.91	97.39 79.44 DER ↑ 44.91 37.39 35.96	92.79 93.94 ACC ↑ 61.19 59.85 60.91	0.00 2.39 RNP ASR \(\psi \) 0.09 0.00 0.01	99.37 98.00 DER ↑ 99.45 98.82 99.35
	Average BadNets Trojan	91.34 93.18 ACC ↑ 61.98 61.58	70.18 74.99 No Defens ASR ↓ 99.97 100.00	63.55 61.44 e DER ↑	93.60 94.57 ACC ↑ 58.26 58.07	0.00 34.09 FP ASR ↓ 99.94 99.12	99.77 82.39 DER ↑ 48.16 48.69	87.08 92.85 ACC ↑ 35.09 59.59	0.01 0.09 CLP ASR ↓ 0.25 0.22	96.50 98.62 DER ↑ 86.41 98.89	88.84 89.68 ACC ↑ 43.78 36.96	0.01 35.53 ANP ASR ↓ 91.76 100.00	97.39 79.44 DER ↑ 44.91 37.39	92.79 93.94 ACC ↑ 61.19 59.85	0.00 2.39 RNP ASR ↓ 0.09 0.00	99.37 98.00 DER ↑ 99.45 98.82
	BadNets Trojan Blend WaNet IAB	91.34 93.18 ACC ↑ 61.98 61.58 62.28 62.37 62.56	70.18 74.99 No Defense ASR \(\preceq \) 99.97 100.00 99.97 99.58 99.39	63.55 61.44 e DER ↑	93.60 94.57 ACC↑ 58.26 58.07 57.68 58.80 59.03	0.00 34.09 FP ASR ↓ 99.94 99.12 0.43 0.17 0.09	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63 0.19	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74	0.01 35.53 ANP ASR ↓ 91.76 100.00 93.91 99.05 0.84	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05	6.90 99.93 89.07 87.78 87.78 90.98 87.78 1.27 90.98 60.01 10.88 0.98 0.52 0.38 1.69 0.52 0.38 0.98 0.52 0.38 0.98 0.52 0.38 0.98 0.52 0.38 0.98 0.52 0.38 0.98 0.98 0.98 0.99 0.00 0.00 0.00 0.0	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10
	BadNets Trojan Blend WaNet IAB Lira	91.34 93.18 ACC↑ 61.98 61.58 62.28 62.37 62.56 62.19	70.18 74.99 No Defense ASR ↓ 99.97 100.00 99.97 99.58	63.55 61.44 e DER ↑ - - -	93.60 94.57 ACC↑ 58.26 58.07 57.68 58.80	0.00 34.09 FP ASR↓ 99.94 99.12 0.43 0.17	99.77 82.39 DER ↑ 48.16 48.69 97.47 97.92	87.08 92.85 ACC ↑ 35.09 59.59 54.29 48.17	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74	0.01 35.53 ANP ASR↓ 91.76 100.00 93.91 99.05	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19	0.00 2.39 RNP ASR ↓ 0.09 0.00 0.01 51.66 15.66 0.00	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28
	BadNets Trojan Blend WaNet IAB	91.34 93.18 ACC ↑ 61.98 61.58 62.28 62.37 62.56	70.18 74.99 No Defens ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.99 99.82	63.55 61.44 e DER ↑ - - - - -	93.60 94.57 ACC ↑ 58.26 58.07 57.68 58.80 59.03 58.87	0.00 34.09 FP ASR ↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17	87.08 92.85 ACC ↑ 35.09 59.59 54.29 48.17 59.41 58.79	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63 0.19 0.24 1.84	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74 41.04	0.01 35.53 ANP ASR↓ 91.76 100.00 93.91 99.05 0.84 99.99	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05	0.00 2.39 RNP ASR ↓ 0.09 0.00 0.01 51.66 15.66 0.00 11.24	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10 95.99
T. I. N.	BadNets Trojan Blend WaNet IAB Lira	91.34 93.18 ACC↑ 61.98 61.58 62.28 62.37 62.56 62.19	70.18 74.99 No Defens ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.99	63.55 61.44 e DER ↑ - - - - -	93.60 94.57 ACC ↑ 58.26 58.07 57.68 58.80 59.03 58.87	0.00 34.09 FP ASR ↓ 99.94 99.12 0.43 0.17 0.09 0.32	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17	87.08 92.85 ACC ↑ 35.09 59.59 54.29 48.17 59.41 58.79	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63 0.19 0.24	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74 41.04	0.01 35.53 ANP ASR ↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19	0.00 2.39 RNP ASR ↓ 0.09 0.00 0.01 51.66 15.66 0.00 11.24 Ours	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10 95.99
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira	91.34 93.18 ACC ↑ 61.98 61.58 62.28 62.37 62.56 62.19 62.16	70.18 74.99 No Defense ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.99 99.82 TSBD	63.55 61.44 e DER ↑ - - - - - - - -	93.60 94.57 58.26 58.07 57.68 58.80 59.03 58.87 58.45	0.00 34.09 FP ASR↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM	99.77 82.39 DER ↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63 0.19 0.24 1.84 SAU	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17 94.18	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88	0.01 35.53 ANP ASR ↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR ↓	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77	0.00 2.39 RNP ASR \$\psi\$ 0.09 0.00 0.01 51.66 15.66 0.00 11.24 Ours	99.37 98.00 DER↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER↑
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira Average	91.34 93.18 ACC↑ 61.98 61.58 62.28 62.37 62.56 62.19 62.16 ACC↑ 50.58	70.18 74.99 No Defens. ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.99 99.82 TSBD ASR ↓ 29.31	63.55 61.44 e DER↑ - - - - - - DER↑ 79.63	93.60 94.57 ACC↑ 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.96	0.00 34.09 FP ASR ↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR ↓ 0.23	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38 DER↑ 95.36	87.08 92.85 ACC ↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC ↑ 52.95	0.01 0.09 CLP ASR \$\psi\$ 0.25 0.22 0.63 0.19 0.24 1.84 SAU ASR \$\psi\$	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17 94.18 DER ↑ 95.20	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88 ACC ↑ 53.58	0.01 35.53 ANP ASR \$\rightsquare\$ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR \$\rightsquare\$ 30.51	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER ↑ 80.53	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC ↑ 56.33	0.00 2.39 RNP ASR↓ 0.09 0.00 0.01 51.66 0.00 11.24 Ours ASR↓ 0.00	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER ↑ 97.16
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira Average	91.34 93.18 ACC ↑ 61.98 61.58 62.28 62.37 62.56 62.19 62.16 ACC ↑ 50.58 50.58	70.18 74.99 No Defenss ASR \$\perp\$ 99.97 100.00 99.97 99.58 99.39 99.99 99.82 TSBD ASR \$\perp\$ 29.31 0.22	63.55 61.44 e DER↑ - - - - - - - - - - - - - - - - - - -	93.60 94.57 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.96 52.29	0.00 34.09 FP ASR ↓ 99.42 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR ↓ 0.23 0.26	99.77 82.39 DER ↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38 DER ↑ 95.36 95.22	87.08 92.85 ACC ↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC ↑ 52.95 52.07	0.01 0.09 CLP ASR ↓ 0.25 0.22 0.63 0.19 0.24 1.84 SAU ASR ↓ 0.54 0.35	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17 94.18 DER ↑ 95.20 95.07	88.84 89.68 ACC ↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88 ACC ↑ 53.58 51.37	0.01 35.53 ANP ASR ↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR ↓ 0.14	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER ↑ 80.53 94.82	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC ↑ 56.33 56.71	0.00 2.39 RNP ASR \$\frac{1}{2}\$ 0.09 0.00 0.01 51.66 15.66 0.00 11.24 Ours ASR \$\frac{1}{2}\$	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER ↑ 97.16 97.56
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira Average	91.34 93.18 ACC ↑ 61.98 61.58 62.28 62.37 62.56 62.19 62.16 ACC ↑ 50.58 50.58 51.48	70.18 74.99 No Defens ASR↓ 99.97 100.00 99.97 99.58 99.39 99.99 99.82 TSBD ASR↓ 29.31 0.22 0.05	63.55 61.44 e DER↑ - - - - - - - - - - - - - - - - - - -	93.60 94.57 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.96 52.29 53.50	0.00 34.09 FP ASR↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR↓ 0.26 0.17	99.77 82.39 DER↑ 48.16 48.69 97.47 97.88 98.17 81.38 DER↑ 95.36 95.22 95.51	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC↑ 52.95 52.07 53.29	0.01 0.09 CLP ASR↓ 0.25 0.22 9.52 0.61 0.19 0.24 1.84 SAU ASR↓ 0.55 0.77	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.17 94.18 DER ↑ 95.20 95.07 92.10	88.84 89.68 ACC↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88 ACC↑ 53.58 51.37 54.04	0.01 35.53 ANP ASR↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR↓ 30.51 0.14 0.48	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER ↑ 80.53 94.82 95.62	92.79 93.94 ACC↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC↑ 56.33 56.71 57.97	0.00 2.39 RNP ASR↓ 0.09 0.00 0.01 51.66 0.00 11.24 Ours ASR↓ 0.00 0.01	99.37 98.00 DER ↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER ↑ 97.16 97.56 97.82
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira Average BadNets Trojan Blend WaNet	91.34 93.18 ACC ↑ 61.98 62.28 62.37 62.56 62.19 62.16 ACC ↑ 50.58 50.58 50.58 50.58 50.41	70.18 74.99 No Defens ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.82 TSBD ASR ↓ 29.31 0.25 99.88	63.55 61.44 e DER ↑	93.60 94.57 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.29 53.50 55.18	0.00 34.09 FP 4SR ↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR ↓ 0.23 0.26 0.17	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38 DER↑ 95.36 95.22 95.51 96.00	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC↑ 52.95 52.07 53.29 57.43	0.01 0.09 CLP ASR↓ 0.25 0.22 9.52 0.63 0.19 0.24 1.84 SAU ASR↓ 0.35 6.77 3.90	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17 94.18 DER ↑ 95.20 95.07 92.10 95.37	88.84 89.68 ACC↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88 ACC↑ 53.58 51.37 54.04 53.89	0.01 35.53 ANP ASR↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR↓ 0.48 0.12	97.39 79.44 DER↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER↑ 80.53 94.82 95.62 95.49	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC ↑ 56.33 56.71 57.97 61.37	0.00 2.39 RNP ASR \(\psi\) 0.09 0.00 0.01 51.66 0.00 11.24 Ours ASR \(\psi\) 0.00 0.01 0.01	99.37 98.00 DER↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER↑ 97.16 97.82 99.28
TinyImageNet	BadNets Trojan Blend WaNet Lira Average BadNets Trojan Blend WaNet Lira Average	91.34 93.18 ACC↑↑ 61.98 61.58 62.28 62.37 62.56 62.19 62.16 ACC↑↑ 50.58 50.58 51.48 50.41 51.69	70.18 74.99 No Defens ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.82 TSBD ASR ↓ 29.31 0.22 0.05 99.88 83.85	63.55 61.44 e DER↑ - - - - - - - - - - - - - - - - - - -	93.60 94.57 ACC↑ 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.96 52.29 53.50 55.18 55.91	0.00 34.09 FP ASR ↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR ↓ 0.23 0.26 0.17 0.39 0.36	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38 DER↑ 95.36 95.22 95.51 96.00 96.19	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC↑ 52.95 52.07 53.29 57.43 55.36	0.01 0.09 CLP ASR ↓ 0.25 0.22 9.52 0.63 0.19 0.24 1.84 SAU ASR ↓ 0.54 0.35 6.77 3.90 1.73	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 94.18 DER ↑ 95.20 95.07 92.10 95.37 95.23	88.84 89.68 ACC↑ 43.78 36.96 28.03 36.74 41.04 36.88 ACC↑ 53.58 51.37 54.09 53.89 54.57	0.01 35.53 ANP 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST 0.14 0.48 0.12 0.04	97.39 79.44 DER ↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER ↑ 80.53 94.82 95.62 95.69 95.69	92.79 93.94 ACC↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC↑ 56.33 56.71 57.97 61.37 61.12	0.00 2.39 RNP ASR ↓ 0.09 0.00 0.01 51.66 0.00 11.24 Ours ASR ↓ 0.00 0.01 0.01 0.02 2.39	99.37 98.00 DER↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER↑ 97.16 97.56 97.82 99.28 97.78
TinyImageNet	BadNets Trojan Blend WaNet IAB Lira Average BadNets Trojan Blend WaNet	91.34 93.18 ACC ↑ 61.98 62.28 62.37 62.56 62.19 62.16 ACC ↑ 50.58 50.58 50.58 50.58 50.41	70.18 74.99 No Defens ASR ↓ 99.97 100.00 99.97 99.58 99.39 99.82 TSBD ASR ↓ 29.31 0.25 99.88	63.55 61.44 e DER ↑	93.60 94.57 58.26 58.07 57.68 58.80 59.03 58.87 58.45 ACC↑ 52.29 53.50 55.18	0.00 34.09 FP 4SR ↓ 99.94 99.12 0.43 0.17 0.09 0.32 33.35 FT-SAM ASR ↓ 0.23 0.26 0.17	99.77 82.39 DER↑ 48.16 48.69 97.47 97.92 97.88 98.17 81.38 DER↑ 95.36 95.22 95.51 96.00	87.08 92.85 ACC↑ 35.09 59.59 54.29 48.17 59.41 58.79 52.56 ACC↑ 52.95 52.07 53.29 57.43	0.01 0.09 CLP ASR↓ 0.25 0.22 9.52 0.63 0.19 0.24 1.84 SAU ASR↓ 0.54 0.35 6.77 3.90	96.50 98.62 DER ↑ 86.41 98.89 91.23 92.37 98.02 98.17 94.18 DER ↑ 95.20 95.07 92.10 95.37	88.84 89.68 ACC↑ 43.78 36.96 28.03 36.74 34.74 41.04 36.88 ACC↑ 53.58 51.37 54.04 53.89	0.01 35.53 ANP ASR↓ 91.76 100.00 93.91 99.05 0.84 99.99 80.92 FST ASR↓ 0.48 0.12	97.39 79.44 DER↑ 44.91 37.39 35.96 37.75 85.85 39.42 46.88 DER↑ 80.53 94.82 95.62 95.49	92.79 93.94 ACC ↑ 61.19 59.85 60.91 36.42 50.05 54.19 53.77 ACC ↑ 56.33 56.71 57.97 61.37	0.00 2.39 RNP ASR \(\psi\) 0.09 0.00 0.01 51.66 0.00 11.24 Ours ASR \(\psi\) 0.00 0.01 0.01	99.37 98.00 DER↑ 99.45 98.82 99.35 61.28 86.10 95.99 90.17 DER↑ 97.16 97.82 99.28

Table 3: Comparison of the backdoor removal results for CIFAR-10 on ResNet-50.

-	N	No Defens	e		FP			CLP			ANP			RNP	
	ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑
BadNets	91.48	100.00	-	91.02	58.88	70.33	63.83	9.33	81.51	15.44	4.40	59.82	91.10	100.00	49.85
Trojan	92.69	100.00	-	91.83	2.70	98.22	50.80	3.24	77.43	89.84	78.10	60.17	87.51	13.18	91.47
Blend	92.11	99.59	-	91.16	17.68	90.48	47.62	5.68	74.71	16.25	21.52	51.66	42.84	15.06	68.19
WaNet	92.83	98.92	-	91.99	0.90	98.59	59.54	0.14	82.74	12.36	0.00	60.48	68.84	90.16	43.64
IAB	92.68	98.71	-	91.93	1.28	98.34	56.43	20.67	70.90	15.08	16.37	53.66	18.18	26.02	50.38
Lira	91.40	100.00	-	90.42	0.33	99.34	24.27	19.53	56.67	53.88	29.34	66.57	37.11	1.76	71.98
Average	92.20	99.54	-	91.39	13.63	92.55	50.41	9.77	73.99	33.81	24.96	58.73	57.60	41.03	62.58
		TSBD			FT-SAM			SAU			FST			Ours	
	ACC ↑	ASR↓	DER ↑	ACC ↑	ASR↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑	ACC ↑	ASR ↓	DER ↑
BadNets	81.41	4.74	92.59	89.77	26.27	86.01	87.51	8.76	93.64	90.97	84.20	57.65	88.97	4.68	96.41
Trojan	87.27	2.24	96.17	91.10	3.21	97.60	88.22	1.16	97.19	92.07	15.66	91.86	90.01	1.10	98.11
Blend	83.42	5.07	92.92	90.00	6.08	95.70	88.60	1.36	97.36	91.18	15.83	91.41	89.19	3.10	96.78
WaNet	83.80	71.51	59.19	90.66	1.57	97.59	89.48	1.86	96.86	92.30	1.58	98.41	90.05	0.69	97.73
IAB	82.83	76.14	56.36	90.74	1.10	97.84	88.51	1.58	96.48	92.35	1.47	98.46	89.57	0.69	97.46
Lira	68.79	5.21	86.09	89.23	0.53	98.65	86.59	0.82	97.18	90.67	10.01	94.63	88.63	0.12	98.55
Average	81.25	27.49	80.55	90.25	6.46	95.56	88.15	2.59	96.45	91.59	21.46	88.74	89.40	1.73	97.51

poisoned classes, since the current method assumes the poisoned class is the one whose perturbation norm is detected as an outlier.

ETHICS STATEMENT

Our work does not involve human participants, sensitive personal data, or experiments with potential risks to individuals or communities. We relied on publicly available datasets that are widely recog-

nized in the research community, and we ensured ethical use of data by citing sources appropriately and complying with dataset licenses.

REPRODUCIBILITY STATEMENT

The experimental configurations used for reproduction are described in Section 5.1 and Appendix C.

REFERENCES

Bilal Hussain Abbasi, Yanjun Zhang, Leo Zhang, and Shang Gao. Backdoor attacks and defenses in computer vision domain: A survey. *arXiv* preprint arXiv:2509.07504, 2025.

Stephen Boyd and Lieven Vandenberghe. Convex Optimization. Cambridge University Press, 2004.

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.

Yiming Chen, Haiwei Wu, and Jiantao Zhou. Progressive poisoned data isolation for training-time backdoor defense. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 11425–11433, 2024.

Yukun Chen, Shuo Shao, Enhao Huang, Yiming Li, Pin-Yu Chen, Zhan Qin, and Kui Ren. Refine: Inversion-free backdoor defense via model reprogramming. In *The Thirteenth International Conference on Learning Representations*, 2025.

Edward Chou, Florian Tramèr, and Giancarlo Pellegrino. Sentinet: Detecting localized universal attacks against deep learning systems. In *Proc. of the Deep Learning and Security Workshop* (DLS) 2020, 2020.

Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

Bao Gia Doan, Ehsan Abbasnejad, and Damith C. Ranasinghe. Februus: Input purification defense against trojan attacks on deep neural network systems. In *36th Annual Computer Security Applications Conference (ACSAC 2020)*, pp. 897–912. ACM, 2020.

Khoa Doan, Yingjie Lao, and Ping Li. Backdoor attack with imperceptible input and latent modification. *Advances in Neural Information Processing Systems*, 34:18944–18957, 2021a.

Khoa Doan, Yingjie Lao, Weijie Zhao, and Ping Li. Lira: Learnable, imperceptible and robust backdoor attacks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 11966–11976, 2021b.

Kuofeng Gao, Yang Bai, Jindong Gu, Yong Yang, and Shu-Tao Xia. Backdoor defense via adaptively splitting poisoned dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4005–4014, 2023.

Yansong Gao, Change Xu, Derui Wang, Shiping Chen, Damith C Ranasinghe, and Surya Nepal. Strip: A defence against trojan attacks on deep neural networks. In *Proceedings of the 35th annual computer security applications conference*, pp. 113–125, 2019.

Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating backdooring attacks on deep neural networks. *Ieee Access*, 7:47230–47244, 2019.

Mengxuan Hu, Zihan Guan, Yi Zeng, Junfeng Guo, Zhongliang Zhou, Jielu Zhang, Ruoxi Jia, Anil Kumar Vullikanti, and Sheng Li. Mind control through causal inference: Predicting clean images from poisoned data. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=ho4mNiwr2n.

Nazmul Karim, Abdullah Al Arafat, Adnan Siraj Rakin, Zhishan Guo, and Nazanin Rahnavard. Fisher information guided purification against backdoor attacks. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 4435–4449, 2024.

- Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Anti-backdoor learning: Training clean models on poisoned data. Advances in Neural Information Processing Systems, 34:14900–14912, 2021.
- Yige Li, Xixiang Lyu, Xingjun Ma, Nodens Koren, Lingjuan Lyu, Bo Li, and Yu-Gang Jiang. Reconstructive neuron pruning for backdoor defense. In *International Conference on Machine Learning*, pp. 19837–19854. PMLR, 2023.
 - Weilin Lin, Li Liu, Shaokui Wei, Jianze Li, and Hui Xiong. Unveiling and mitigating backdoor vulnerabilities based on unlearning weight changes and backdoor activeness. *Advances in Neural Information Processing Systems*, 37:42097–42122, 2024.
 - Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *International symposium on research in attacks, intrusions, and defenses*, pp. 273–294. Springer, 2018a.
 - Xiaogeng Liu, Minghui Li, Haoyu Wang, Shengshan Hu, Dengpan Ye, Hai Jin, Libing Wu, and Chaowei Xiao. Detecting backdoors during the inference stage based on corruption robustness consistency (teco). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023*, pp. 16363–16372, 2023.
 - Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In 25th Annual Network And Distributed System Security Symposium (NDSS 2018). Internet Soc, 2018b.
 - Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *European Conference on Computer Vision*, pp. 182–199. Springer, 2020.
 - Rui Min, Zeyu Qin, Li Shen, and Minhao Cheng. Towards stable backdoor purification through feature shift tuning. *Advances in Neural Information Processing Systems*, 36:75286–75306, 2023.
 - Xiaoxing Mo, Yechao Zhang, Leo Yu Zhang, Wei Luo, Nan Sun, Shengshan Hu, Shang Gao, and Yang Xiang. Robust backdoor detection for deep learning via topological evolution dynamics (ted). *arXiv preprint arXiv:2312.02673*, 2023.
 - Tuan Anh Nguyen and Anh Tran. Input-aware dynamic backdoor attack. *Advances in Neural Information Processing Systems*, 33:3454–3464, 2020.
 - Tuan Anh Nguyen and Anh Tuan Tran. Wanet-imperceptible warping-based backdoor attack. In *International Conference on Learning Representations*, 2021.
 - Te Juin Lester Tan and Reza Shokri. Bypassing backdoor detection algorithms in deep learning. In 2020 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 175–183, 2020. doi: 10.1109/EuroSP48549.2020.00019.
 - Shaokui Wei, Mingda Zhang, Hongyuan Zha, and Baoyuan Wu. Shared adversarial unlearning: Backdoor mitigation by unlearning shared adversarial examples. *Advances in Neural Information Processing Systems*, 36:25876–25909, 2023.
 - Baoyuan Wu, Hongrui Chen, Mingda Zhang, Zihao Zhu, Shaokui Wei, Danni Yuan, and Chao Shen. Backdoorbench: A comprehensive benchmark of backdoor learning. *Advances in Neural Information Processing Systems*, 35:10546–10559, 2022.
 - Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models. *Advances in Neural Information Processing Systems*, 34:16913–16925, 2021.
 - Zhen Xiang, Zidi Xiong, and Bo Li. Cbd: A certified backdoor detector based on local dominant probability. *arXiv preprint arXiv:2310.17498*, 2023.
 - Tinghao Xie, Xiangyu Qi, Ping He, Yiming Li, Jiachen T. Wang, and Prateek Mittal. Badexpert: Extracting backdoor functionality for accurate backdoor input detection. *arXiv* preprint *arXiv*:2308.12439, 2023.

Xiaoyun Xu, Zhuoran Liu, Stefanos Koffas, and Stjepan Picek. Towards backdoor stealthiness in model parameter space. *arXiv preprint arXiv:2501.05928*, 2025.

Kaiyuan Zhang, Siyuan Cheng, Guangyu Shen, Guanhong Tao, Shengwei An, Anuran Makur, Shiqing Ma, and Xiangyu Zhang. Exploring the orthogonality and linearity of backdoor attacks. In 2024 IEEE Symposium on Security and Privacy (SP), pp. 225–225, Los Alamitos, CA, USA, may 2024. IEEE Computer Society. doi: 10.1109/SP54263.2024.00225. URL https://doi.ieeecomputersociety.org/10.1109/SP54263.2024.00225.

Runkai Zheng, Rongjun Tang, Jianze Li, and Li Liu. Data-free backdoor removal based on channel lipschitzness. In *European Conference on Computer Vision*, pp. 175–191. Springer, 2022.

Nan Zhong, Zhenxing Qian, and Xinpeng Zhang. Imperceptible backdoor attack: From input space to feature representation. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pp. 1736–1742. International Joint Conferences on Artificial Intelligence Organization, 2022.

Mingli Zhu, Shaokui Wei, Li Shen, Yanbo Fan, and Baoyuan Wu. Enhancing fine-tuning based backdoor defense with sharpness-aware minimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4466–4477, 2023.

A LLM USAGE

While drafting this paper, we used a large language model (e.g., GPT-5) to assist with grammar correction, readability improvements, and literature searches. The scientific content, original ideas, and experimental findings are entirely the work of the authors.

B ADDITIONAL RELATED WORKS FOR BACKDOOR DEFENSES

In Section 2.2, we discussed backdoor defenses that aim to remove backdoors by identifying backdoor neurons from compromised models, and here we introduce other defense strategies following the literature (Abbasi et al., 2025).

B.1 BACKDOOR REMOVAL WITHOUT BACKDOOR NEURON IDENTIFICATION

Several recent defenses avoid explicitly identifying backdoor neurons and instead mitigate backdoors through fine-tuning and feature regularization. FT-SAM (Zhu et al., 2023) employs sharpness-aware minimization during fine-tuning to suppress backdoor-sensitive parameters. SAU (Wei et al., 2023) uses adversarial perturbations to unlearn shared backdoor features across classes, while FST (Min et al., 2023) adjusts feature distributions to shift poisoned representations away from decision boundaries. FIP (Karim et al., 2024) leverages Fisher information to purify representations and reduce the influence of backdoors.

B.2 Training-Stage Defenses

Training-stage defenses aim to prevent the learning of backdoor correlations during model training by modifying optimization dynamics, restructuring the training process, or limiting the influence of poisoned data. A key insight is that poisoned data often behave differently from clean data in early training, e.g., faster loss reduction or more sensitive feature transformations, which can be exploited to detect and neutralize them.

Representative methods include Anti-Backdoor Learning (ABL) (Li et al., 2021), which isolates suspicious low-loss data in early epochs and later unlearns them to break trigger-label associations. Extensions refine this idea: Adaptively Splitting Dataset (ASD) (Gao et al., 2023) adaptively partitions data into clean and poisoned pools; Progressive Isolation (PIPD) (Chen et al., 2024) progressively reduces false positives in isolation; and Mind Control through Causal Inference (MCCI) (Hu et al., 2025) leverages causal modeling to disentangle triggers from true classes.

Algorithm 1 Backdoor Removal via Reconstructing TAC in the Latent Representation

Require: Compromised model parameter θ_{bd} , a reference dataset D_{ref} , threshold α , hyperparameter

Ensure: Fine-tuned model $\theta_{\rm ft}^*$

```
1: Phase 1: Reconstructing TAC in the Latent Representation
      Phase 1-1: Computing a Minimal Perturbation for Each Class
3:
         for each class k \in [C] do
```

4: Solve the quadratic problem in equation 4 (via dual equation 5) to obtain s_k^*

5: end for

Phase 1-2: Identifying the Poisoned Class 6:

7: Compute $||s_1^*||_2, ||s_2^*||_2, \cdots ||s_C^*||_2$ for all $k \in [C]$ 8: Standardize $z_k = (\|\mathbf{s}_k^*\|_2 - \mu)/\sigma$ and pick p with $z_p < \alpha$

9: Identify poisoned class p such that $z_p < \alpha$

10: Phase 2: Backdoor Removal with the Perturbation in the Poisoned Class

11: Fine-tune the compromised model by solving

12:

648

649

650 651

652

653

654

655

656

657

658

659

660

661

662

663 664

665 666 667

668 669 670

671 672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690 691 692

693 694

696 697

698 699

700

$$\boldsymbol{\theta}_{\mathrm{ft}}^* = \operatorname*{argmin}_{\boldsymbol{\theta}_{\mathrm{bd}}} \frac{1}{|D_{\mathrm{ref}}|} \sum_{i=1}^{|D_{\mathrm{ref}}|} \Big[\ell(f(\boldsymbol{x}_i; \boldsymbol{\theta}_{\mathrm{bd}}), \boldsymbol{y}_i) + \beta \, \ell(\mathrm{Softmax}(\boldsymbol{W}_L^\top (\hat{\boldsymbol{x}}_i + \boldsymbol{s}_p^*) + \boldsymbol{b}), \boldsymbol{y}_i) \Big].$$

13: **Return** $\theta_{\rm ft}^*$

B.3 Inference-Stage Defenses

Inference-stage defenses aim to identify or neutralize trigger-bearing inputs during inference, making them especially useful when retraining or model inspection is impractical. A representative approach is perturbation-based detection, where methods such as STRIP (Gao et al., 2019) perturb incoming inputs and measure the entropy of predictions; consistently low entropy often indicates the presence of a trigger. Another line focuses on input purification, with Februus (Doan et al., 2020) removing suspicious regions through inpainting to recover benign content and mitigate patch-style trojans.

Beyond perturbation and purification, interpretability-based defenses such as SentiNet (Chou et al., 2020) leverage saliency maps to localize highly influential regions and assess their generalization across inputs, enabling detection of physical-world triggers. Similarly, TeCo (Liu et al., 2023) exploits robustness discrepancies under common image corruptions, showing that poisoned inputs behave inconsistently compared to clean ones, thus allowing detection without soft classes or auxiliary clean datasets. More recent studies, including CBD (Xiang et al., 2023), TED (Mo et al., 2023), and BaDExpert (Xie et al., 2023), further enhance detection reliability by leveraging statistical probability bounds, topological dynamics, or explicit extraction of backdoor functionality. As another line of work, REFINE (Chen et al., 2025) introduces a model reprogramming strategy that jointly employs an input transformation module and an output remapping module. By aggressively transforming inputs while simultaneously remapping output classes, REFINE reduces the effectiveness of triggers without severely degrading clean accuracy.

IMPLEMENTATION DETAILS

conducted experiments based the implementation the OrthogLinearBackdoor (Zhang et al., 2024) repository 1.

C.1 BACKDOOR ATTACKS

We implemented six representative backdoor attack methods. Default configurations of all attacks follow the OrthogLinearBackdoor. As described in Section 5.1, all attacks were trained for

https://github.com/KaiyuanZh/OrthogLinearBackdoor/blob/main/

100 epochs using stochastic gradient descent (SGD) with a learning rate of 0.1 and cosine annealing as the learning rate scheduler.

- **BadNets** (Gu et al., 2019). A patch-based backdoor that stamps a fixed visible pattern onto inputs to induce a target class; in our experiments we follow existing work (Zhang et al., 2024) and use the sunflower image as the trigger.
- Trojan (Liu et al., 2018b). A trigger-stamping attack which plants a small image-based trigger; here the trigger is a small sunflower image with a transparent background.
- **Blend** (Chen et al., 2017). A blending-style attack that mixes a trigger image into the entire input with a given transparency; we use a Hello-Kitty image blended at an alpha of 0.2.
- WaNet (Nguyen & Tran, 2021). A warping-based backdoor that applies imperceptible geometric distortions (image warps) as the trigger, producing stealthy, input-agnostic perturbations.
- IAB (Nguyen & Tran, 2020). An input-dependent attack that generates a dynamic trigger conditioned on each input, making detection and removal more challenging.
- Lira (Doan et al., 2021b). A backdoor attack generating learnable, imperceptible, and robust triggers, making them hard to detect and defend.

C.2 BACKDOOR DEFENSES

We implemented eight backdoor removal methods. Unless otherwise specified, implementations are based on the OrthogLinearBackdoor repository, while methods without public implementations were re-implemented following the authors' original repositories or BackdoorBench (Wu et al., 2022) which is another benchmark framework that provides unified implementations of representative backdoor attacks and defenses for fair and reproducible evaluation.

- **Fine-Pruning (FP)** (Liu et al., 2018a). This method prunes neurons that are inactive on clean data, assuming such neurons are likely backdoor-related. We set the pruning ratio as 0.2, fine-tuning epochs as 50, the optimizer as SGD, learning rate as 0.01 and learning rate scheduler as cosine annealing.
- Channel Lipschitzness Pruning (CLP) (Zheng et al., 2022). CLP removes channels with abnormally large Lipschitz constants, aiming to suppress backdoor activations. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in BackdoorBench (Wu et al., 2022). We also set the threshold parameter as 3.0 following the original paper (Zheng et al., 2022).
- Adversarial Neuron Pruning (ANP) (Wu & Wang, 2021). ANP identifies and prunes neurons that are highly sensitive to adversarial perturbations. In our experiments, for CIFAR-10 we set $\epsilon=0.3, \alpha=0.2$ and the pruning threshold as 0.2; for GTSRB $\epsilon=0.4, \alpha=0.2$ and the pruning threshold as 0.4; and for TinyImageNet $\epsilon=0.2, \alpha=0.3$ and the pruning threshold 0.001 where ϵ and α are the hyperparameters introduced in the original paper.
- Reconstructive Neuron Pruning (RNP) (Li et al., 2023). RNP prunes neurons whose removal minimally affects the reconstruction of clean representations from the unlearned model. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in BackdoorBench. We set the pruning threshold as 0.7 for CIFAR-10, the pruning threshold as 0.95 for GTSRB, and the pruning threshold as 0.1 for TinyImageNet.
- Two-Stage Backdoor Defense (TSBD) (Lin et al., 2024). TSBD identifies backdoor neuron based on Neuron Weight Change (NWC) which is the difference between the compromised model's weights and the unlearned model's weights, and conducts activeness-aware fine-tuning to mitigate backdoors. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in the original paper (Lin et al., 2024). Following the original paper, after calculating NWC, we selected 15% of the top-neurons and pruned 70% of the top-subweights within them. In our experiments, we attempted to use activeness-aware fine-tuning following the original paper, but since the accuracy dropped significantly after fine-tuning, we instead adopted standard fine-tuning. Fine-tuning configuration in TSBD is the same as that of FP.

- FT-SAM (Zhu et al., 2023). This method leverages sharpness-aware minimization (SAM) during fine-tuning to suppress backdoor behaviors. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in Backdoor-Bench. The training configuration and hyperparameters are followed as BackdoorBench.
- Shared Adversarial Unlearning (SAU) (Wei et al., 2023). SAU uses adversarial perturbations to unlearn shared backdoor features across classes. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in BackdoorBench. The training configuration and hyperparamters are followed as BackdoorBench.
- Feature Shift Tuning (FST) (Min et al., 2023). FST fine-tunes models by aligning feature distributions to shift away backdoor-related representations. The implementation is not included in the OrthogLinearBackdoor repository, we refer the implementation in the original paper (Min et al., 2023). The hyperparameter that balances the loss terms (denoted as α in the original paper) is set to 0.2 for CIFAR-10, 0.1 for GTSRB, and 0.001 for TinyImageNet, following the original paper.

D DETAILS OF PROPOSED METHOD

D.1 ALGORITHMS

 To clarify our proposed method as described in Section 4, we present the detailed procedure in Algorithm 1.

D.2 DERIVATION PROCESS FOR DUAL PROBLEM

We describe the derivation process from equation 4 to equation 5.

Lagrangian and dual function. Introduce the dual variable $\lambda \in \mathbb{R}^{C-1}$ with $\lambda \geq 0$ for the inequality constraints from equation 4. The Lagrangian is

$$\mathcal{L}(s_k, oldsymbol{\lambda}) = rac{1}{2} \left\| oldsymbol{s}_k
ight\|_2^2 - oldsymbol{\lambda}^ op (oldsymbol{V}_k oldsymbol{s}_k - oldsymbol{m}) \quad ext{s.t.} \quad oldsymbol{\lambda} \geq oldsymbol{0}.$$

The dual function is obtained by minimizing the Lagrangian over the primal variable:

$$g(\lambda) = \inf_{s_k} \mathcal{L}(s_k, \lambda).$$

Stationarity (optimality in s_k) gives

$$abla_{oldsymbol{s}_k} \mathcal{L}(oldsymbol{s}_k, oldsymbol{\lambda}) = oldsymbol{s}_k - oldsymbol{V}_k^ op oldsymbol{\lambda} = oldsymbol{0} \implies oldsymbol{s}_k = oldsymbol{V}_k^ op oldsymbol{\lambda}.$$

Plugging this into \mathcal{L} yields

$$g(\boldsymbol{\lambda}) = \boldsymbol{\lambda}^{\top} \boldsymbol{m} - \frac{1}{2} \| \boldsymbol{V}_k^{\top} \boldsymbol{\lambda} \|_2^2.$$

Therefore, the dual problem is the concave maximization

$$oldsymbol{\lambda}^* = rgmax_{oldsymbol{\lambda}} \ oldsymbol{\lambda}^ op oldsymbol{m} - rac{1}{2} \, \|oldsymbol{V}_k^ op oldsymbol{\lambda}\|_2^2 \quad ext{s.t.} \quad oldsymbol{\lambda} \geq oldsymbol{0}.$$

D.3 FEASIBLE SOLUTION

Theorem 1. If $C-1 < d_{\rm emb}$ and V_k has full row rank, i.e. ${\rm rank}(V_k) = C-1$, then the primal problem equation 4 has a feasible solution.

Proof. By Farkas' lemma (Boyd & Vandenberghe, 2004), exactly one of the following two statements holds:

1. There exists $s_k \in \mathbb{R}^{d_{\text{emb}}}$ such that $V_k s_k \geq m$.

2. There exists $\lambda \in \mathbb{R}^{C-1}$ such that $V^{\top}\lambda = 0$, $\lambda \geq 0$, $m^{\top}\lambda < 0$.

If $\operatorname{rank}(\boldsymbol{V}_k) = C - 1$ with $C - 1 < d_{\operatorname{emb}}$, then $\ker(\boldsymbol{V}_k^\top) = \{0\}$. Hence the only $\boldsymbol{\lambda}$ satisfying $\boldsymbol{V}_k^\top \boldsymbol{\lambda} = \boldsymbol{0}$ is $\boldsymbol{\lambda} = \boldsymbol{0}$, which cannot yield $\boldsymbol{m}^\top \boldsymbol{\lambda} < 0$. Thus (2) is impossible, and therefore (1) must hold. Hence, there exists \boldsymbol{s}_k with $\boldsymbol{V}_k^\top \boldsymbol{s}_k \geq \boldsymbol{m}$, and the primal problem is feasible.

D.4 STRONG DUALITY

Theorem 2. If $C - 1 < d_{\text{emb}}$ and $\text{rank}(V_k) = C - 1$, then the primal problem equation 4 and equation 5 satisfy strong duality.

Proof. To establish this result, we prove that the primal problem is a convex optimization problem and that it satisfies Slater's condition.

1. Convexity. The objective $\frac{1}{2} \| s_k \|_2^2$ is strongly convex. The feasible region is given by

$$Z \coloneqq \{oldsymbol{s}_k \in \mathbb{R}^{d_{ ext{emb}}}: oldsymbol{V}_k^ op oldsymbol{s}_k \geq oldsymbol{m}\} = \bigcap_{i=1}^{C-1} \{oldsymbol{s}_k: oldsymbol{v}_i^ op oldsymbol{s}_k \geq m_i\},$$

where each set $\{s_k : v_i^\top s_k \ge m_i\}$ is a half-space and therefore convex. Since the feasible set Z is the intersection of convex sets, it is also convex. Thus, the problem equation 4 is a convex optimization problem.

2. Slater's condition. Since V_k has full row rank, we have $\operatorname{rank}(V_k) = C - 1$. This implies that the linear map

$$T: \mathbb{R}^{d_{\mathrm{emb}}} o \mathbb{R}^{C-1}, \qquad T(\boldsymbol{s}_k) = \boldsymbol{V}_k^{ op} \boldsymbol{s}_k,$$

is surjective. Hence, for any $\epsilon > 0$, there exists $\bar{s}_k \in \mathbb{R}^{d_{\text{emb}}}$ such that

$$V_k^{\top} \bar{s}_k = m + \epsilon \mathbf{1}_{C-1}.$$

Since $\epsilon > 0$, it follows that

$$oldsymbol{V}_k^{ op}ar{oldsymbol{s}}_k = oldsymbol{m} + \epsilon oldsymbol{1}_{C-1} > oldsymbol{m},$$

which means that \bar{s}_k strictly satisfies all inequality constraints. In other words, $\bar{s}_k \in \operatorname{relint}(Z)$, where $Z = \{s_k \in \mathbb{R}^{d_{\operatorname{emb}}} : V_k^{\top} s_k \geq m\}$ and $\operatorname{relint}(Z)$ means the relative interior of the set Z. Therefore, Slater's condition holds for problem equation 4.

E ADDITIONAL EXPERIMENTS

E.1 COMPARISON WITH PREVIOUS METHODS FOR IDENTIFYING BACKDOOR NEURONS

We compare how accurately the perturbations in the latent representation obtained by our method can identify TAC-based backdoor neurons relative to existing approaches. Figure 4, Figure 8 and Figure 9 show the overlap rate with TAC-based backdoor neurons in the latent representation at the Top-K% for each dataset. These results show that among existing methods, RNP exhibits relatively stable performance, achieving high TAC coverage at small K on CIFAR-10 and GTSRB, whereas TAC coverage at small K on TinyImageNet shows low. In contrast, our proposed method consistently attains high TAC coverage at small K across all datasets, demonstrating its stability and effectiveness in reconstructing TAC in the latent representation.

E.2 EFFECTIVENESS OF POISONED CLASS IDENTIFICATION METHOD

To verify the effectiveness of the poisoned class identification method, we conduct an ablation study in which fine-tuning is performed without identifying the poisoned class. Namely, we fine-tune a compromised model using the perturbations of all classes. Specifically, instead of applying s_p^* in Equation (6) for our method, we randomly select s^* from the set of perturbations at each training iteration for fine-tuning. The training configuration is the same as that of our method.

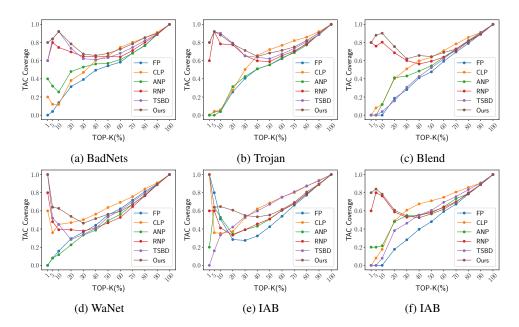


Figure 4: TAC coverage, defined as the overlap ratio between TAC-based backdoor neurons and those identified by each defense method on CIFAR-10.

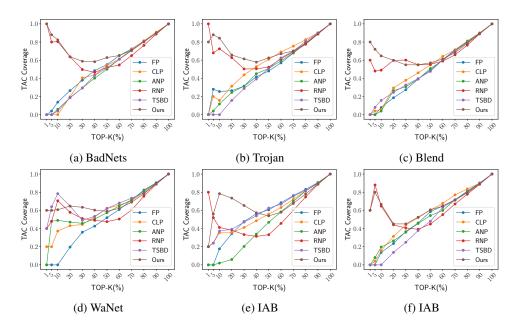


Figure 5: TAC coverage on GTSRB.

As shown in Table 4, even without poisoned class identification, ASR generally decreases to a level comparable to our proposed method although ASR of 17.26% remains for IAB on TinyImageNet and ASR of 20.23% for Blend on GTSRB. This is likely because, during training, the randomly selected s^* occasionally corresponds to s_p^* . On the other hand, in terms of ACC, our method achieves higher performance on CIFAR-10 and TinyImageNet. These results indicate that by leveraging only the perturbation of the poisoned class through the poisoned class identification method, our method is able to maintain higher accuracy while effectively removing backdoors.

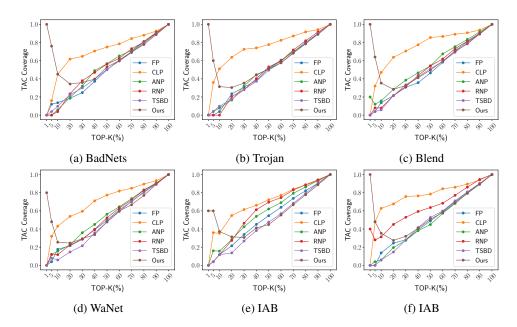


Figure 6: TAC coverage on TinyImageNet.

Table 4: Comparison of backdoor removal results between fine-tuning with the perturbations for all classes and fine-tuning with the perturbation of the poisoned class (Ours). "No PCI" means fine-tuning without the poisoned class identification (PCI) method.

Dataset		l N	No Defens	e		No PCI			Ours	
	Attack	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑
	BadNets	93.81	100.00	-	90.91	16.13	90.48	92.03	10.88	93.67
	Trojan	94.00	100.00	-	90.36	2.67	96.85	92.01	0.98	98.52
	Blend	93.29	99.91	-	90.11	4.81	95.96	91.84	1.69	98.39
CIFAR-10	WaNet	93.41	99.59	-	91.23	2.18	97.62	92.39	0.52	99.02
	IAB	93.57	98.81	-	90.38	1.89	96.87	92.37	0.38	98.62
	Lira	94.29	99.98	-	90.72	0.68	97.86	92.80	0.11	99.19
	Average	93.73	99.71	-	90.62	4.73	95.94	92.24	2.43	97.90
	BadNets	95.08	100.00	-	94.89	7.68	96.07	94.27	6.78	96.20
	Trojan	94.39	100.00	-	93.08	0.41	99.14	93.40	0.49	99.27
	Blend	93.85	99.50	-	92.99	17.26	90.69	93.39	7.06	95.99
GTSRB	WaNet	93.99	97.07	_	95.65	0.03	98.52	95.60	0.00	98.54
	IAB	94.09	97.22	-	94.36	0.02	98.60	94.21	0.00	98.61
	Lira	93.97	99.91	_	92.44	0.06	99.16	92.79	0.00	99.37
	Average	94.23	98.95	-	93.90	4.24	97.03	93.94	2.39	98.00
-	BadNets	61.98	99.97	-	53.17	0.25	95.45	56.33	0.00	97.16
	Trojan	61.58	100.00	-	52.09	0.28	95.11	56.71	0.01	97.56
	Blend	62.28	99.97	-	53.19	0.16	95.36	57.97	0.01	97.82
TinyImageNet	WaNet	62.37	99.58	-	54.82	2.43	94.80	61.37	0.02	99.28
	IAB	62.56	99.39	-	55.57	20.23	86.08	61.12	2.39	97.78
	Lira	62.19	99.99	-	54.00	0.26	95.77	59.78	0.02	98.78
	Average	62.16	99.82	-	53.81	3.94	93.76	58.88	0.41	98.06

E.3 COMPARISON WITH PRUNING-BASED METHODS VIA RECONSTRUCTING TAC IN THE LATENT REPRESENTATION

As described in Section 4.3, we removed backdoors by reconstructing TAC with fine-tuning. Alternatively, pruning-based methods can provide another approach that leverages the reconstructed TAC for backdoor removal. Therefore, we further compare our method with pruning-based approaches by reconstructing TAC in the latent representation. As shown in Table 5, pruning alone can partially reduce ASR, but a considerable portion of backdoors remains (e.g., ASR of 69.89% for Trojan on CIFAR-10 and 56.08% for Blend on GTSRB), indicating that pruning itself is insufficient

Table 5: Comparison of backdoor removal results between pruning-based methods and fine-tuning-based method (Ours). Pruning ratio and fine-tuning configuration are set to be the same as those of FP

		N	No Defens	e		Pruning		P	runing+F	Γ		Ours	
		ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑
	BadNets	93.81	100.00	-	88.87	37.74	78.66	93.66	100.00	49.92	92.03	10.88	93.67
	Trojan	94.00	100.00	-	90.26	69.89	63.19	93.55	97.86	50.85	92.01	0.98	98.52
	Blend	93.29	99.91	-	89.00	43.56	76.03	93.32	41.59	79.16	91.84	1.69	98.39
CIFAR-10	WaNet	93.41	99.59	-	92.43	36.99	80.81	93.37	20.22	89.66	92.39	0.52	99.02
	IAB	93.57	98.81	-	93.30	0.32	99.11	93.40	1.54	98.55	92.37	0.38	98.62
	Lira	94.29	99.98	-	90.28	87.60	54.18	93.84	26.02	86.75	92.80	0.11	99.19
	Average	93.73	99.71	-	90.69	46.02	75.33	93.52	47.87	75.82	92.24	2.43	97.90
	BadNets	95.08	100.00	-	94.51	0.30	99.56	95.08	98.81	50.59	94.27	6.78	96.20
	Trojan	94.39	100.00	-	93.62	45.74	76.75	94.61	99.20	50.40	93.40	0.49	99.27
	Blend	93.85	99.50	-	93.70	56.08	71.64	94.71	81.07	59.22	93.39	7.06	95.99
GTSRB	WaNet	93.99	97.07	-	95.05	69.98	63.54	95.67	47.00	75.04	95.60	0.00	98.54
	IAB	94.09	97.22	-	93.17	31.11	82.59	94.42	11.85	92.68	94.21	0.00	98.61
	Lira	93.97	99.91	-	93.43	6.49	96.44	94.13	14.04	92.94	92.79	0.00	99.37
	Average	94.23	98.95	-	93.91	34.95	81.75	94.77	58.66	70.14	93.94	2.39	98.00
	BadNets	61.98	99.97	-	59.38	0.00	98.68	58.03	0.34	97.84	56.33	0.00	97.16
	Trojan	61.58	100.00	-	58.05	0.06	98.20	56.95	0.16	97.60	56.71	0.01	97.56
	Blend	62.28	99.97	-	59.55	32.61	82.31	59.00	0.11	98.29	57.97	0.01	97.82
TinyImageNet	WaNet	62.37	99.58	-	58.15	0.00	97.68	60.78	0.22	98.88	61.37	0.02	99.28
	IAB	62.56	99.39	-	59.20	0.00	98.01	60.40	0.03	98.60	61.12	2.39	97.78
	Lira	62.19	99.99	-	56.15	68.37	62.79	58.45	0.15	98.05	59.78	0.02	98.78
	Average	62.16	99.82	-	58.41	16.84	89.61	58.94	0.17	98.21	58.88	0.41	98.06

to completely eliminate the attacks. When combined with fine-tuning (Pruning+FT), the accuracy can be preserved, but the fine-tuning process often revives backdoors, leading to higher ASR in several cases (e.g., BadNets on CIFAR-10 where ASR returns to 100%). In contrast, our method consistently decreases ASR across all attack settings while preserving high accuracy. These results highlight that our approach overcomes the limitations of pruning-based methods and provides a more reliable defense against backdoor attacks.

E.4 RESULTS FOR DIFFERENT REFERENCE DATASET SIZES

To investigate the dependency of our method on the size of the reference dataset, we further conducted experiments by varying the reference set at 1.0%, 5.0% and 10.0% of the training dataset. As shown in Table 6, our method consistently reduces ASR to nearly 0.0% across all dataset sizes, demonstrating that even a small reference set can effectively eliminate backdoors. Regarding clean accuracy, we observe that using 5.0% of the reference dataset already provides stable performance that is almost identical to using 10.0%, indicating that 5.0% is sufficient in practice.

However, we note that on GTSRB, using only 1.0% of the reference dataset significantly decreases accuracy (from 94.23% to 70.17%), although ASR is still effectively reduced to 0.0%. This result suggests that for datasets with complex distributions such as GTSRB, a slightly larger reference dataset (e.g., $\geq 5.0\%$) is required to preserve clean accuracy while maintaining strong defense efficacy.

E.5 EFFECTIVENESS OF HYPERPARAMETER β

Since the hyperparameter β is a crucial parameter that balances ACC and ASR, Figure 7, Figure 8, Figure 9 show how ACC and ASR vary with different values of β for each dataset. We observe that as β increases, both ACC and ASR decrease for all datasets. For CIFAR-10, we set $\beta=0.5$ as it provides a good trade-off between ACC and ASR. For GTSRB, the ASR does not decrease unless β is set to 2.0 in some cases (e.g., BadNets and Blend). However, since the clean accuracy does not drop significantly, we set $\beta=2.0$. For TinyImageNet, while ACC decreases substantially as β increases, the ASR is reduced to nearly zero already at $\beta=0.1$, and thus we set $\beta=0.1$.

Table 6: The effectiveness of our proposed method for each size of the reference dataset.

		N	lo Defens	e	O	urs (1.0%	b)	O	urs (5.0%	(b)	O	urs (10.0	%)
		ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑	ACC↑	ASR↓	DER↑
	BadNets	93.81	100.00	-	91.55	3.53	97.10	92.03	10.88	93.67	92.57	14.18	92.29
	Trojan	94.00	100.00	-	91.09	1.13	97.98	92.01	0.98	98.52	92.57	1.30	98.64
	Blend	93.29	99.91	-	90.61	0.52	98.35	91.84	1.69	98.39	92.11	1.08	98.83
CIFAR-10	WaNet	93.41	99.59	-	89.38	0.19	97.69	92.39	0.52	99.02	92.61	0.37	99.21
	IAB	93.57	98.81	-	90.06	0.23	97.53	92.37	0.38	98.62	92.60	0.66	98.59
	Lira	94.29	99.98	-	91.15	0.16	98.34	92.80	0.11	99.19	92.83	0.11	99.20
	Average	93.73	99.71	-	90.64	0.96	97.83	92.24	2.43	97.90	92.55	2.95	97.79
	BadNets	95.08	100.00	-	71.01	0.00	87.97	94.27	6.78	96.20	94.71	1.88	98.88
	Trojan	94.39	100.00	-	63.97	0.00	84.79	93.40	0.49	99.27	93.82	5.26	97.09
	Blend	93.85	99.50	-	68.18	0.00	86.91	93.39	7.06	95.99	94.24	2.82	98.34
GTSRB	WaNet	93.99	97.07	-	80.58	0.00	91.83	95.60	0.00	98.54	95.44	0.00	98.54
	IAB	94.09	97.22	-	72.51	0.00	87.82	94.21	0.00	98.61	94.44	0.02	98.60
	Lira	93.97	99.91	-	64.78	0.00	85.36	92.79	0.00	99.37	93.40	0.00	99.67
	Average	94.23	98.95	-	70.17	0.00	87.45	93.94	2.39	98.00	94.34	1.66	98.52
	BadNets	61.98	99.97	-	55.96	0.02	96.96	56.33	0.00	97.16	53.84	0.04	95.89
	Trojan	61.58	100.00	-	56.11	0.01	97.26	56.71	0.01	97.56	54.25	0.07	96.30
	Blend	62.28	99.97	-	55.12	0.00	96.40	57.97	0.01	97.82	53.62	0.01	95.65
TinyImageNet	WaNet	62.37	99.58	-	57.67	0.00	97.44	61.37	0.02	99.28	56.87	0.03	97.02
	IAB	62.56	99.39	-	58.18	0.00	97.50	61.12	2.39	97.78	55.94	0.04	96.36
	Lira	62.19	99.99	-	57.34	0.00	97.57	59.78	0.02	98.78	55.98	0.07	96.85
	Average	62.16	99.82	-	56.73	0.01	97.19	58.88	0.41	98.06	55.08	0.04	96.35

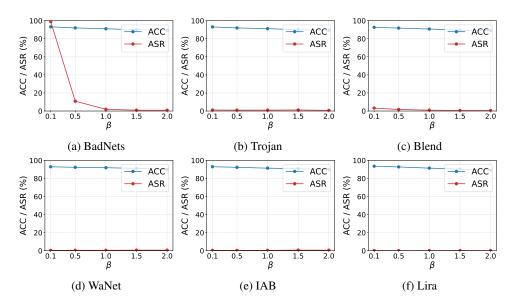


Figure 7: Effectiveness of the hyperparameter β for CIFAR-10.

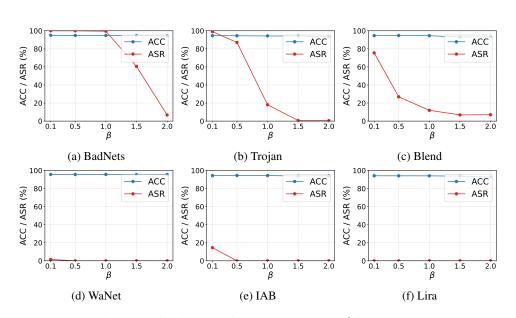


Figure 8: Effectiveness of the hyperparameter β for GTSRB.

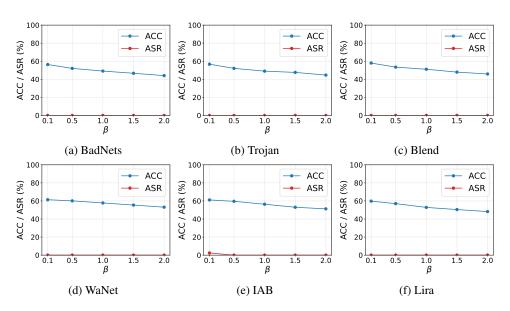


Figure 9: Effectiveness of the hyperparameter β for TinyImageNet.