

On Certifying Robustness against Backdoor Attacks via Randomized Smoothing

Binghui Wang, Xiaoyu Cao, Jinyuan Jia, Neil Zhenqiang Gong
Duke University

{binghui.wang, xiaoyu.cao, jinyuan.jia, neil.gong}@duke.edu

Abstract

Backdoor attack is a severe security threat to deep neural networks (DNNs). We envision that, like adversarial examples, there will be a cat-and-mouse game for backdoor attacks, i.e., new empirical defenses are developed to defend against backdoor attacks but they are soon broken by strong adaptive backdoor attacks. To prevent such cat-and-mouse game, we take the first step towards certified defenses against backdoor attacks. Specifically, in this work, we study the feasibility and effectiveness of certifying robustness against backdoor attacks using a recent technique called randomized smoothing. Randomized smoothing was originally developed to certify robustness against adversarial examples. We generalize randomized smoothing to defend against backdoor attacks. Our results show the theoretical feasibility of using randomized smoothing to certify robustness against backdoor attacks. However, we also find that existing randomized smoothing methods have limited effectiveness at defending against backdoor attacks, which highlight the needs of new theory and methods to certify robustness against backdoor attacks.

1. Introduction

Backdoor attack [6, 21, 11, 27] is a severe security threat to DNNs. Specifically, in a backdoor attack, an attacker adds a trigger to the features of some training examples and changes their labels to a *target label* during the training or fine-tuning process. Then, when the attacker adds the same trigger to the features of a testing example, the learnt classifier predicts the target label for the testing example with the trigger. We envision that, like adversarial examples, there will be a cat-and-mouse game for backdoor attacks. Indeed, various empirical defenses [4, 7, 9, 5, 23, 26, 20] have been proposed to defend against backdoor attacks in the past few years. For instance, Neural Cleanse [26] aims to detect and reconstruct the trigger via solving an optimization problem. However, [12] showed that Neural Cleanse fails to detect the trigger when the trigger has different sizes, shapes, and/or locations.

To prevent such cat-and-mouse game, we take the first step towards certifying robustness against backdoor attacks. Specifically, we study the feasibility and effectiveness of certifying robustness against backdoor attacks using *randomized smoothing* [2, 8, 17, 14]. Randomized smoothing was originally developed to certify robustness against adversarial examples [25, 10, 3]. In particular, Cao and Gong [2] is the first to propose randomized smoothing as an empirical defense against adversarial examples (they call it *region-based classification*). Cohen et al. [8] derived a tight certified robustness guarantee for randomized smoothing with Gaussian noise using the Neyman-Pearson Lemma [22]. Lee et al. [17] and Jia et al. [14] generalized randomized smoothing to discrete data using discrete noise distribution.

In this work, we generalize randomized smoothing to defend against backdoor attacks. Given an arbitrary function (we call it *base function*), which takes a data vector as an input and outputs a label, randomized smoothing can turn the function to be a provably robust one via adding random noise to the input data vector. Specifically, a function is provably robust if it outputs the same label for all data points in a region (e.g., ℓ_p -norm ball) around an input. Our idea to certify robustness against backdoor attacks consists of two steps. First, we view the entire process of learning a classifier from the training dataset and using the classifier to make predictions for a testing example as a base function. In particular, this base function takes a training dataset and a testing example as an input, and the base function outputs a predicted label for the testing example. Second, we add random noise to the training dataset and a testing example to overwhelm the trigger that the attacker injects to the training examples and the testing example.

We evaluate our method on a subset of the MNIST dataset. Our defense guarantees that 36% of testing images can be classified correctly when an attacker arbitrarily perturbs at most 2 pixels/labels of the training examples and pixels of a testing example. Our results show the theoretical feasibility of using randomized smoothing to certify robustness against backdoor attacks. However, our results also show that existing randomized smoothing methods with ad-

ditive noise have limited effectiveness at defending against backdoor attacks. Our study highlights the needs of new theory and techniques to certify robustness against backdoor attacks.

2. Background on Randomized Smoothing

Randomized smoothing is state-of-the-art technique to certify the robustness of a classifier against adversarial examples. Randomized smoothing [2, 19] was first proposed as an empirical defense against adversarial examples. For instance, Cao and Gong [2] proposed to add uniform random noise from a hypercube centered at a testing example and use majority vote to smooth the predicted label of the testing example. They called the method *region-based classification* as it leverages information in a region around a testing example to predict its label. Lecuyer et al. [16] derived the first certified robustness guarantee for randomized smoothing using differential privacy techniques. Li et al. [18] derived a tighter certified robustness guarantee using information-theoretic techniques. Cohen et al. [8] derived a tight certified robustness guarantee for randomized smoothing with Gaussian noise using the Neyman-Pearson Lemma [22]. Jia et al. [13] derived a tight certified robustness guarantee of general top- k predictions for randomized smoothing with Gaussian noise. Lee et al. [17] and Jia et al. [14] generalized randomized smoothing to discrete data using discrete noise. We will use such randomized smoothing for discrete data because labels are discrete/categorical. Salman et al. [24] and Zhai et al. [28] proposed methods to train classifiers that have better certified robustness under randomized smoothing.

Next, we describe randomized smoothing from a general function perspective, making it easier to understand how we apply randomized smoothing to certify robustness against backdoor attacks.

2.1. Building a Smoothed Function

Data vector v : Suppose we have a data vector v . We consider each dimension of v to be discrete, as many applications have discrete data, e.g., pixel values are discrete. Moreover, when certifying robustness against backdoor attacks, some dimensions of v correspond to the labels of the training examples, which are discrete. Without loss of generality, we assume each dimension of v is from the discrete domain $\{0, \frac{1}{d}, \dots, \frac{d-1}{d}\}$, where d is the domain size. We note that randomized smoothing could also be applied when the dimensions of v have different domain sizes. However, for simplicity, we assume the dimensions have the same domain size.

Base function: Suppose we have an arbitrary function (we call it *base function*), which takes the data vector v as an input and outputs a label in a set $\{0, 1, \dots, c-1\}$. For

instance, when certifying robustness against adversarial examples, the base function is a classifier whose robustness we aim to certify. When certifying robustness against backdoor attacks, we treat the entire process of learning a classifier and using the learnt classifier to make predictions for a testing example as a base function. For convenience, we denote the base function as f , and $f(v)$ is the predicted label for v .

Adversarial perturbation: An attacker can perturb the data vector v . We denote by δ the adversarial perturbation an attacker adds to the vector v , where δ_j is the perturbation added to the j th dimension of the vector v and $\delta_j \in \{0, \frac{1}{d}, \dots, \frac{d-1}{d}\}$. Moreover, we denote by $v \oplus \delta$ the *perturbed data vector*, where the operator \oplus is defined for each dimension as follows:

$$v_j \oplus \delta_j = \frac{(v_j \cdot d + \delta_j \cdot d) \bmod d}{d}, \quad (1)$$

where v_j and δ_j are the j th dimensions of the vector v and the perturbation vector δ , respectively. We measure the magnitude of the adversarial perturbation using its ℓ_0 norm, i.e., $\|\delta\|_0$. We adopt ℓ_0 norm because it is semantically easy to interpret. In particular, ℓ_0 norm of the adversarial perturbation is the number of dimensions of v that are arbitrarily perturbed by the attacker.

Smoothed function: Randomized smoothing builds a new function from the base function via adding random noise to the data vector v . We call the new function *smoothed function*. Specifically, we denote by ϵ the random noise vector, where the j th dimension $\epsilon_j \in \{0, \frac{1}{d}, \dots, \frac{d-1}{d}\}$ is the random noise added to v_j . We consider ϵ_j has the following distribution [17, 14]:

$$\begin{aligned} \Pr(\epsilon_j = 0) &= \beta, \\ \Pr(\epsilon_j = t) &= \theta = \frac{1 - \beta}{d - 1}, \forall t \in \left\{ \frac{1}{d}, \frac{2}{d}, \dots, \frac{d-1}{d} \right\}. \end{aligned} \quad (2)$$

Moreover, $v \oplus \epsilon$ is the *noisy data vector*, where the operator \oplus is defined in Equation 1. The noise distribution indicates that, when adding a random noise vector ϵ to the data vector v , the j th dimension of v is preserved with a probability β and is changed to any other value with a probability $\frac{1-\beta}{d-1}$. Since we add random noise to the data vector v , the base function f outputs a random label. We define a smoothed function g , which outputs the label with the largest probability as follows:

$$g(v) = \operatorname{argmax}_{y \in \{0, 1, \dots, c-1\}} \Pr(f(v \oplus \epsilon) = y), \quad (3)$$

where $g(v)$ is the label predicted for v by the smoothed function. Note that $g(v \oplus \delta)$ is the label predicted for the perturbed data vector $v \oplus \delta$.

2.2. Computing Certified Radius

Randomized smoothing guarantees that the smoothed function g predicts the same label when the adversarial perturbation δ is bounded. In particular, according to [17], we have:

$$g(v \oplus \delta) = l, \forall \|\delta\|_0 \leq R(\underline{p}_l), \quad (4)$$

where $\underline{p}_l \leq \Pr(f(v \oplus \epsilon) = l)$ is a lower bound of the probability $\Pr(f(v \oplus \epsilon) = l)$ that f predicts a label l when adding random noise ϵ to v , and $R(\underline{p}_l)$ is called *certified radius*. Intuitively, Equation 4 shows that the smoothed function g predicts the same label l when an attacker arbitrarily perturbs at most $R(\underline{p}_l)$ dimensions of the data vector v . Note that the certified radius $R(\underline{p}_l)$ depends on \underline{p}_l . In particular, given any lower bound \underline{p}_l , we can compute the certified radius $R(\underline{p}_l)$. The computation details can be found in [17]. Estimating a lower bound \underline{p}_l is the key to compute the certified radius. Next, we describe how to estimate \underline{p}_l .

Cohen et al. [8] proposed a Monte Carlo method to predict l and estimate \underline{p}_l with probabilistic guarantees. Specifically, we sample N noise $\epsilon^1, \epsilon^2, \dots, \epsilon^N$ from the noise distribution defined in Equation 2. We compute the label $f(v \oplus \epsilon^j)$ for each noise ϵ^j , and we compute the label frequency $\text{count}[i] = \sum_{j=1}^N \mathbb{I}(f(v \oplus \epsilon^j) = i)$ for each $i \in \{0, 1, \dots, c-1\}$, where \mathbb{I} is an indicator function. The smoothed function g outputs the label l with the largest frequency. Moreover, Cohen et al. [8] proposed to use the Clopper-Pearson method [1] to estimate \underline{p}_l as follows:

$$\underline{p}_l = B(\alpha; \text{count}[l], N - \text{count}[l] + 1), \quad (5)$$

where $1 - \alpha$ is the confidence level and $B(\alpha; a, b)$ is the α th quantile of the Beta distribution with shape parameters a and b .

3. Certifying Robustness against Backdoor Attacks

Suppose we have a training dataset $\{\mathbf{X}, \mathbf{y}\} = \{(x_1, y_1), (x_2, y_2), \dots, (x_T, y_T)\}$, where x_i and y_i are the feature vector and label of the i th training example, respectively. Suppose further we have a learning algorithm \mathcal{A} which takes the training dataset as an input and produces a classifier h , i.e., $h = \mathcal{A}(\mathbf{X}, \mathbf{y})$. We use the classifier h to predict the label for a testing example x . We generalize randomized smoothing to certify robustness against backdoor attacks. Our key idea is to combine the entire process of training and prediction as a single function $f(\mathbf{X}, \mathbf{y}, x)$, which is the predicted label for a testing example x when the classifier is trained on $\{\mathbf{X}, \mathbf{y}\}$ using the algorithm \mathcal{A} . We view the function f as a base function and apply randomized smoothing to it.

Constructing a smoothed function: We view the concatenation of the feature matrix \mathbf{X} , the label vector \mathbf{y} , and the features of the testing example x as the data vector v that we described in Section 2. We add a random noise matrix τ to the feature matrix \mathbf{X} , where each entry of the noise matrix is drawn from the distribution defined in Equation 2 with d as the feature domain size. We add a random noise vector ϵ to the label vector \mathbf{y} , where each entry of the noise vector is drawn from the distribution defined in Equation 2 with $d = c$. Furthermore, we add a random noise vector γ to the testing example x , where each entry of the noise vector is drawn from the distribution defined in Equation 2 with d as the feature domain size. Since we add random noise, the output of the base function f is also random. The smoothed function g outputs the label that has the largest probability. Formally, we have:

$$g(\mathbf{X}, \mathbf{y}, x) = \underset{y \in \{0, 1, \dots, c-1\}}{\operatorname{argmax}} \Pr(f(\mathbf{X} \oplus \tau, \mathbf{y} \oplus \epsilon, x \oplus \gamma) = y), \quad (6)$$

where $g(\mathbf{X}, \mathbf{y}, x)$ is the label predicted by the smoothed function for x .

Computing the certified radius: We denote by a matrix δ_1 and a vector δ_2 the perturbations an attacker adds to the feature matrix \mathbf{X} and label vector \mathbf{y} , respectively. Moreover, we denote by a vector δ_3 the adversarial perturbation an attacker adds to the testing example x . Based on Equation 4, we have the following:

$$\begin{aligned} g(\mathbf{X} \oplus \delta_1, \mathbf{y} \oplus \delta_2, x \oplus \delta_3) &= l, \\ \forall \|\delta_1\|_0 + \|\delta_2\|_0 + \|\delta_3\|_0 &\leq R(\underline{p}_l), \end{aligned} \quad (7)$$

where $\underline{p}_l \leq \Pr(f(\mathbf{X} \oplus \tau, \mathbf{y} \oplus \epsilon, x \oplus \gamma) = l)$. Equation 7 means that the smoothed function g predicts the same label for the testing example x when the ℓ_0 norm of the adversarial perturbation added to the feature matrix and label vector of the training examples as well as the testing example is bounded by $R(\underline{p}_l)$. The key to computing the certified radius $R(\underline{p}_l)$ is to estimate \underline{p}_l . We use the Monte Carlo method described in Section 2 to estimate \underline{p}_l . Specifically, we randomly sample N noise matrices $\tau^1, \tau^2, \dots, \tau^N$, N noise vectors $\epsilon^1, \epsilon^2, \dots, \epsilon^N$, as well as N noise vectors $\gamma^1, \gamma^2, \dots, \gamma^N$. We train N classifiers, where the j th classifier h_j is trained using the training dataset $(\mathbf{X} \oplus \tau^j, \mathbf{y} \oplus \epsilon^j)$ and learning algorithm \mathcal{A} . Then, we compute the frequency of each label i , i.e., $N_i = \sum_{j=1}^N \mathbb{I}(f(\mathbf{X} \oplus \tau^j, \mathbf{y} \oplus \epsilon^j, x \oplus \gamma^j) = i) = \sum_{j=1}^N \mathbb{I}(h(x \oplus \gamma^j) = i)$ for $i \in \{0, 1, \dots, c-1\}$. Finally, we can estimate \underline{p}_l using Equation 5. Note that the trained N classifiers can be re-used to predict the labels and compute certified radius for different testing examples. Moreover, via evenly dividing α among the testing examples, we can obtain a simultaneous confidence level of $1 - \alpha$ for the estimated certified radius of the testing examples based on the *Bonferroni Correction*.

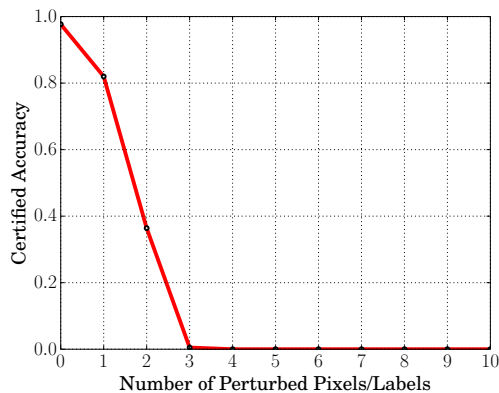


Figure 1. Certified accuracy of our method against backdoor attacks on MNIST 1/7.

4. Experimental Results

Experimental setup: We use a subset of the MNIST dataset [15] which only contains the handwritten digits "1" and "7" and is used for binary classification. Each digit image has a size 28×28 and has normalized pixel values within $\{0, 1/255, 2/255, \dots, 1\}$. For simplicity, we binarize the pixel values in our experiment. Specifically, if a pixel value is smaller than 0.5, we set it to be 0, otherwise we set it to be 1. Moreover, we randomly select 100 digits from the subset of the MNIST dataset to form the training dataset and randomly select another 1,000 digits as the testing examples. We use a two-layer neural network as the classifier.

Evaluation metric: We use *certified accuracy* as an evaluation metric. Specifically, for a given number of perturbed pixels/labels, certified accuracy is the fraction of testing examples, whose labels are correctly predicted by the smoothed function and whose certified radiuses are no smaller than the given number of perturbed pixels/labels.

Experimental results: Figure 1 shows the certified accuracy of our method against backdoor attacks with $\beta = 0.9$, $N = 10,000$, and $1 - \alpha = 99.9\%$. Our method guarantees that 36% of testing images can be classified correctly when an attacker arbitrarily perturbs at most 2 pixels/labels of the training examples and pixels of a testing example.

5. Conclusion

In this work, we take the first step towards certified defenses against backdoor attacks. In particular, we study the feasibility and effectiveness of certifying robustness against backdoor attacks via randomized smoothing, which was originally developed to certify robustness against adversarial examples. Our method has two key steps. First, we treat the entire process of training a classifier and using the classifier to predict the label of a testing example as a base function. Second, we add noise to the training data and a testing

example to overwhelm the perturbation an attacker adds in a backdoor attack. Our results on a subset of MNIST demonstrate that it is theoretically feasible to certify robustness against backdoor attacks using randomized smoothing, but existing randomized smoothing methods have limited effectiveness. Our study highlights the needs of new techniques to certify robustness against backdoor attacks.

References

- [1] Lawrence D Brown, T Tony Cai, and Anirban DasGupta. Interval estimation for a binomial proportion. *Statistical science*, 2001. 3
- [2] Xiaoyu Cao and Neil Zhenqiang Gong. Mitigating evasion attacks to deep neural networks via region-based classification. In *ACSAC*, 2017. 1, 2
- [3] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE S & P*, 2017. 1
- [4] Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. Detecting backdoor attacks on deep neural networks by activation clustering. *arXiv*, 2018. 1
- [5] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. Deepinspect: A black-box trojan detection and mitigation framework for deep neural networks. In *IJCAI*, 2019. 1
- [6] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv*, 2017. 1
- [7] Edward Chou, Florian Tramèr, Giancarlo Pellegrino, and Dan Boneh. Sentinet: Detecting physical attacks against deep learning systems. *arXiv*, 2018. 1
- [8] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. In *ICML*, 2019. 1, 2, 3
- [9] 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 *ACSAC*, 2019. 1
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 1
- [11] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv*, 2017. 1
- [12] Wenbo Guo, Lun Wang, Xinyu Xing, Min Du, and Dawn Song. Tabor: A highly accurate approach to inspecting and restoring trojan backdoors in ai systems. *arXiv*, 2019. 1
- [13] Jinyuan Jia, Xiaoyu Cao, Binghui Wang, and Neil Zhenqiang Gong. Certified robustness for top-k predictions against adversarial perturbations via randomized smoothing. In *ICLR*, 2020. 2
- [14] Jinyuan Jia, Binghui Wang, Xiaoyu Cao, and Neil Zhenqiang Gong. Certified robustness of community detection against adversarial structural perturbation via randomized smoothing. In *The Web Conference (WWW)*, 2020. 1, 2

- [15] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 4
- [16] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *IEEE S & P*, 2019. 2
- [17] Guang-He Lee, Yang Yuan, Shiyu Chang, and Tommi S Jaakkola. Tight certificates of adversarial robustness for randomly smoothed classifiers. In *NeurIPS*, 2019. 1, 2, 3
- [18] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Second-order adversarial attack and certifiable robustness. In *NeurIPS*, 2019. 2
- [19] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *ECCV*, 2018. 2
- [20] Yingqi Liu, Wen-Chuan Lee, Guanhong Tao, Shiqing Ma, Yousra Aafer, and Xiangyu Zhang. Abs: Scanning neural networks for back-doors by artificial brain stimulation. In *CCS*, 2019. 1
- [21] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In *NDSS*, 2018. 1
- [22] Jerzy Neyman and Egon Sharpe Pearson. Ix. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933. 1, 2
- [23] Ximing Qiao, Yukun Yang, and Hai Li. Defending neural backdoors via generative distribution modeling. In *NeurIPS*, 2019. 1
- [24] Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. In *NeurIPS*, 2019. 2
- [25] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 1
- [26] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *IEEE S & P*, 2019. 1
- [27] Yuanshun Yao, Huiying Li, Haitao Zheng, and Ben Y Zhao. Latent backdoor attacks on deep neural networks. In *CCS*, 2019. 1
- [28] Runtian Zhai, Chen Dan, Di He, Huan Zhang, Boqing Gong, Pradeep Ravikumar, Cho-Jui Hsieh, and Liwei Wang. Macer: Attack-free and scalable robust training via maximizing certified radius. In *ICLR*, 2020. 2