

# XRAND: Differentially Private Defense against Explanation-Guided Attacks

Truc Nguyen<sup>\*1</sup>, Phung Lai<sup>\*2</sup>, Hai Phan<sup>2</sup>, My T. Thai<sup>1†</sup>

<sup>1</sup> University of Florida, Gainesville, FL 32611

<sup>2</sup> New Jersey Institute of Technology, Newark, NJ 07102  
 truc.nguyen@ufl.edu, tl353@njit.edu, phan@njit.edu, mythai@cise.ufl.edu

## Abstract

Recent development in the field of explainable artificial intelligence (XAI) has helped improve trust in Machine-Learning-as-a-Service (MLaaS) systems, in which an explanation is provided together with the model prediction in response to each query. However, XAI also opens a door for adversaries to gain insights into the black-box models in MLaaS, thereby making the models more vulnerable to several attacks. For example, feature-based explanations (e.g., SHAP) could expose the top important features that a black-box model focuses on. Such disclosure has been exploited to craft effective backdoor triggers against malware classifiers. To address this trade-off, we introduce a new concept of achieving local differential privacy (LDP) in the explanations, and from that we establish a defense, called XRAND, against such attacks. We show that our mechanism restricts the information that the adversary can learn about the top important features, while maintaining the faithfulness of the explanations.

## 1 Introduction

Over decades, successes in machine learning (ML) have promoted a strong wave of AI applications that deliver vast benefits to a diverse range of fields. Unfortunately, due to their complexity, ML models suffer from opacity in terms of explainability, which reduces the trust in and the verifiability of the decisions made by the models. To meet the necessity of transparent decision making, model-agnostic explainers have been developed to help create effective, more human-understandable AI systems, such as LIME (Ribeiro, Singh, and Guestrin 2016) and SHAP (Lundberg and Lee 2017), among many others (Sundararajan, Taly, and Yan 2017; Selvaraju et al. 2017; Vu and Thai 2020; Vu, Nguyen, and Thai 2022; Ying et al. 2019; Shrikumar, Greenside, and Kundaje 2017; Vu et al. 2021).

In MLaaS systems, a customer can build an ML model by uploading their data or crowdsourcing data, and executing an ML training algorithm. Then, the model is deployed in the cloud where users can receive the model predictions for input queries. MLaaS assumes black-box models as the

end-users have no knowledge about the algorithm or internal information about the underlying ML models. Several proposals advocate for deploying model explanations in the cloud, such that a predicted label and an explanation are returned for each query to provide transparency for end-users. In practice, such an explainable MLaaS system model has been developed by many cloud providers (Azure 2021; Bluemix 2021), with applications in both industries (Community 2018) and academic research (Shokri, Strobel, and Zick 2021; Milli et al. 2019).

Despite the great potential of those explainers to improve the transparency and understanding of ML models in MLaaS, they open a trade-off in terms of security. Specifically, they allow adversaries to gain insights into black-box models, essentially uncovering certain aspects of the models that make them more vulnerable. Such an attack vector has recently been exploited by the research community to conduct several *explanation-guided attacks* (Shokri, Strobel, and Zick 2021; Milli et al. 2019; Miura, Hasegawa, and Shibahara 2021; Zhao et al. 2021; Severi et al. 2021). It was shown that an explainer may expose the top important features on which a black-box model is focusing, by aggregating over the explanations of multiple samples. An example of utilizing such information is the recent highly effective explanation-guided backdoor attack (XBA) against malware classifiers investigated by (Severi et al. 2021). The authors suggest that SHAP can be used to extract the top- $k$  goodwill-oriented features. The attacker then selects a combination of those features and their values for crafting a trigger; and injects trigger-embedded goodwill samples into the training dataset of a malware classifier, with an aim of changing the prediction of goodwill samples embedded with the same trigger at inference time.

**Defense Challenges.** To prevent an adversary from exploiting the explanations, we need to control the information leaked through them, especially the top- $k$  features. Since the explanation on each queried sample is returned to the end users, a viable defense is to randomize the explanation such that it is difficult for attackers to distinguish top- $k$  features while maintaining valuable explainability for the decision-making process. A well-applied technique to achieve this is preserving local differential privacy (LDP) (Erlingsson, Pihur, and Korolova 2014; Wang et al. 2017) on the explanations. However, existing LDP-based approaches (Erlings-

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Corresponding author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

son, Pihur, and Korolova 2014; Xiong et al. 2020; Sun, Qian, and Chen 2021; Zhao et al. 2020) are not designed to protect the top- $k$  features aggregated over the returned explanations on queried data samples. Therefore, optimizing the trade-off between defenses against explanation-guided attacks and model explainability is an open problem.

**Contributions.** (1) We introduce a new concept of achieving LDP in model explanations that simultaneously protects the top- $k$  features from being exploited by attackers while maintaining the faithfulness of the explanations. Based on this principle, we propose a defense against explanation-guided attacks on MLaaS, called XRAND, by devising a novel two-step LDP-preserving mechanism. First, at the aggregated explanation, we incorporate the explanation loss into the randomized probabilities in LDP to make top- $k$  features indistinguishable to the attackers. Second, at the sample-level explanation, guided by the first step, we minimize the explanation loss on each sample while keeping the features at the aggregated explanation intact. (2) Then, we theoretically analyze the robustness of our defense against the XBA in MLaaS by establishing two certified robustness bounds in both training time and inference time. (3) Finally, we evaluate the effectiveness of XRAND in mitigating the XBA on cloud-hosted malware classifiers.

**Organization.** The remainder of the paper is structured in the following manner. Section 2 presents some background knowledge for our paper. Section 3 discusses the explanation-guided attacks against MLaaS and establishes the security model for our defense. Our defense, XRAND, is introduced in Section 4 where its certified robustness bounds are presented in Section 5. Experimental evaluation of our solution is given in Section 6. Section 7 discusses related work and Section 8 provides some concluding remarks. All technical appendices can be accessed via the full version of our paper (Nguyen et al. 2022).

## 2 Preliminaries

**Local explainers.** The goal of model explanations is to capture the importance of each feature value of a given point of interest with respect to the decision made by the classifier and which class it is pushing that decision toward. Given a sample  $x \in \mathbb{R}^d$  where  $x_j$  denotes the  $j^{\text{th}}$  feature of the sample, let  $f$  be a model function in which  $f(x)$  is the probability that  $x$  belongs to a certain class. An explanation of the model’s output  $f(x)$  takes the form of an explanation vector  $w_x \in \mathbb{R}^d$  where the  $j^{\text{th}}$  element of  $w_x$  denotes the degree to which the feature  $x_j$  influences the model’s decision. In general, higher values of  $w_{x_j}$  imply a higher impact.

Perturbation-based explainers, such as SHAP (Lundberg and Lee 2017), obtain an explanation vector  $w_x$  for  $x$  via training a surrogate model of the form  $g(x) = w_{x_0} + \sum_{j=1}^d w_{x_j} x_j$  by minimizing a loss function  $\mathcal{L}(f, g)$  that measures how unfaithful  $g$  is in approximating  $f$ .

- Sample-level explanation. In the context of this paper, we refer to  $w_x$  as a sample-level explanation.
- Aggregated explanation. We denote an aggregated explanation  $w$  as the sum of explanation vectors across sam-

ples in a certain set  $\mathcal{X}$ , i.e.,  $w_{\mathcal{X}} = \sum_{x \in \mathcal{X}} w_x$ . When  $\mathcal{X}$  is clear from the context, we shall use a shorter notation  $w$ .

**Local Differential Privacy (LDP).** LDP is one of the state-of-the-arts and provable approaches to achieve individual data privacy. LDP-preserving mechanisms (Erlingsson, Pihur, and Korolova 2014; Wang et al. 2017; Bassily and Smith 2015; Duchi, Jordan, and Wainwright 2018; Acharya, Sun, and Zhang 2019) generally build on the ideas of randomized response (RR) (Warner 1965).

**Definition 1.**  $\epsilon$ -LDP. A randomized algorithm  $\mathcal{A}$  satisfies  $\epsilon$ -LDP, if for any two inputs  $x$  and  $x'$ , and for all possible outputs  $\mathcal{O} \in \text{Range}(\mathcal{A})$ , we have:  $\Pr[\mathcal{A}(x) = \mathcal{O}] \leq e^\epsilon \Pr[\mathcal{A}(x') = \mathcal{O}]$ , where  $\epsilon$  is a privacy budget and  $\text{Range}(\mathcal{A})$  denotes every possible output of  $\mathcal{A}$ .

The privacy budget  $\epsilon$  controls the amount by which the distributions induced by inputs  $x$  and  $x'$  may differ. A smaller  $\epsilon$  enforces a stronger privacy guarantee.

## 3 XAI-guided Attack Against MLaaS

We discuss how XAI can be used to gain insights into MLaaS models, and establish the threat model for our work.

### 3.1 Exposing MLaaS via XAI

From a security viewpoint, releasing additional information about a model’s mechanism is a perilous prospect. As a function of the model that is trained on a private dataset, an explanation may unintentionally disclose critical information about the training set, more than what is needed to offer a useful interpretation. Moreover, the explanations may also expose the internal mechanism of the black-box models. For example, first, the behavior of explanations varies based on whether the query sample was a member of the training dataset, making the model vulnerable to membership inference attacks (Shokri, Strobel, and Zick 2021). Second, the explanations can be coupled with the predictions to improve the performance of generative models which, in turn, strengthens some model inversion attacks (Zhao et al. 2021). Furthermore, releasing the explanations exposes how the black-box model acts upon an input sample, essentially giving up more information about its inner workings for each query, hence, model extractions attacks can be carried out with far fewer queries, as discussed in (Milli et al. 2019; Miura, Hasegawa, and Shibahara 2021).

Finally, (Severi et al. 2021) argues that the explanations allow an adversary to gain insight into a model’s decision boundary in a generic, model-agnostic way. The SHAP values can be considered as an approximation of the confidence of the decision boundary along each feature dimension. Hence, features with SHAP values that are near zero infer low-confidence areas of the decision boundary. On the other hand, features with positive SHAP values imply that they strongly contribute to the decision made by the model. As a result, it provides us with an indication of the overall orientation for each feature, thereby exposing how the model rates the importance of each feature.

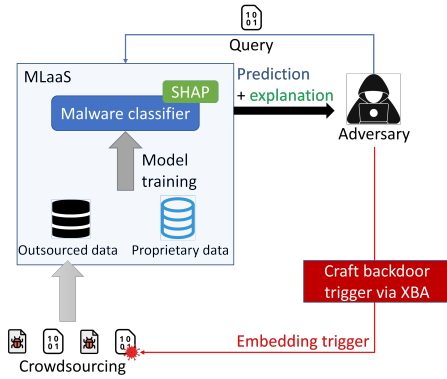


Figure 1: System model of a cloud-hosted malware classifier that leverages crowdsourced data for model training.

### 3.2 XAI-guided Backdoor Attack against MLaaS

The XBA on malware classifiers (Severi et al. 2021) suggests that explanations make the model vulnerable to backdoor attacks, as they reveal the top important features. Thus, it is natural to mount this XBA against a black-box model in MLaaS where an explanation is returned for each user query.

**System Model.** Fig. 1 illustrates the system model for our work. We consider an MLaaS system where a malware classifier is deployed on a cloud-based platform. For training, the system crowdsources threat samples via user-submitted binaries to assemble a set of outsourced data. This set of outsourced data is then combined with a set of proprietary data to construct the training data to train the malware classifier.

We denote  $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$  as the set of proprietary training data. The dataset contains sample  $x_n \in \mathbb{R}^d$  and its ground-truth label  $y_n \in \{0, 1\}$ , where  $y_n = 0$  denotes a goodwill sample, and  $y_n = 1$  denotes a malware sample. On input  $x$ , the model  $f: \mathbb{R}^d \rightarrow \mathbb{R}$  outputs the score  $f(x) \in [0, 1]$ . This score is then compared with a threshold of 0.5 to obtain the predicted label for  $x$ .

During inference time, given a query containing a binary sample  $x$ , the system returns the predicted label with a SHAP explanation  $w_x$  for the decision ( $w_x \in \mathbb{R}^d$ ). We consider an adversary who plays the role of a user in this system and can send queries at his discretion. The adversary exploits the returned explanations to craft backdoor triggers that will be injected to the system via the crowdsourcing process, thereby poisoning the outsourced data.

**Threat Model.** The attacker’s goal is to alter the training procedure by injecting poisoned samples into the training data, generating poisoned training data such that the resulting backdoored classifier differs from a clean classifier. An ideal backdoored classifier has the same response to a clean input as the clean classifier, but it gives an attack-chosen prediction when the input is embedded with the trigger.

Our defense assumes a strong adversary such that he can tamper with the training dataset at his discretion without major constraints. To prevent the adversary from setting arbitrary values for the features in the trigger, the set of values that can be used is limited to the ones that exist in the dataset.

This threat model promotes a defense under worst-case scenarios from the perspective of the defender.

**Crafting backdoor triggers.** To craft a backdoor trigger in XBA, the adversary tries to obtain the top goodwill-oriented feature by querying classifier  $f$  with samples  $\{x\}_{x \in A}$  from their dataset  $A$  and obtaining the SHAP explanation  $w_x$  for each of them. The sum of the SHAP values across all queried samples  $w_A = \sum_{x \in A} w_x$  approximately indicates the importance of each feature, and whether it is goodwill- or malware-oriented. From that, the attacker greedily selects a combination of the most goodwill-oriented features to create the trigger (Severi et al. 2021).

## 4 XRAND – Local DP Defense

This section describes our defense, XRAND, a novel two-step explanation-guided randomized response (RR) mechanism. Our idea is to incorporate the model explainability into the randomization probabilities in XRAND to guide the aggregated explanation while minimizing the difference between the perturbed explanation’s surrogate model  $g'(x)$  and the model  $f(x)$  at the sample-level explanation. We call the difference between  $g'(x)$  and  $f(x)$  an explanation loss  $\mathcal{L}$ , quantified as follows:

$$\mathcal{L} = \sum_{z \in N(x)} (g'(z) - f(z))^2 \exp\left(-\frac{\|z - x\|^2}{\sigma^2}\right) \quad (1)$$

where  $x$  is an input sample, its neighborhood  $N(x)$  is generated by the explainer’s sampling method. This function captures the difference between the modified explainer’s linear surrogate model  $g'(x)$  and the model  $f(x)$ , essentially measuring how unfaithful  $g'$  is in approximating  $f$ .

To defend against explanation-guided attacks that utilize the top- $k$  features of the aggregated explanation (e.g., XBA exploits the top- $k$  goodwill-oriented features), our idea is to randomly disorder some top- $k$  features under LDP guarantees, thereby protecting the privacy of those features. This raises the following question: *What top- $k$  features and which data samples should be randomized to optimize the explainability of data samples while guaranteeing that the attackers cannot infer the top- $k$  features?*

**Algorithm Overview.** To answer this question, we first integrate the explanation loss caused by potential changes of features in the aggregated explanation into the randomized probabilities to adaptively randomize each feature in the top- $k$ . Then, we minimize the explanation loss on each sample while ensuring the order of the features at the aggregated explanation follows the results of the first step. By doing so, we are able to optimize the trade-off between the model explainability and the privacy budget  $\epsilon$  used in XRAND, as verified both theoretically (Section 4.2) and experimentally (Section 6). The pseudo-code of XRAND is shown in Alg. 1.

### 4.1 LDP-preserving Explanations

**Step 1 (Alg. 1, lines 1-10).** We first compute the aggregated explanation  $w$  over the samples of the proprietary dataset  $w = \sum_{x \in \mathcal{D}} w_x$ . Then we sort  $w$  in descending order

---

**Algorithm 1: X RAND: Explanation-guided RR mechanism**


---

**Input:** model  $f$ , dataset  $\mathcal{D}$ , aggregated explanation  $w$ ,  $\varepsilon$ ,  $k$ ,

 $\tau$ , test sample  $x$ 
**Output:**  $\mathbb{S}$ ,  $\varepsilon$ -LDP  $w'_x$ 

```

1: Step 1 - At aggregated explanation:
2:   for  $x_n \in \mathcal{D}$  do
3:     Compute  $\mathcal{L}(x_n)$  # using Eq. 1
4:     for  $i \in [1, k], j \in [k+1, k+\tau]$  do
5:       Compute  $\mathcal{L}(x_n)(i, j)$  # using Eq. 1
6:       Compute  $\Delta_{\mathcal{L}}(i, j)$  # using Eq. 3
7:     end for
8:   end for
9:   Randomizing  $w$ :
    $w' \leftarrow \text{X RAND}(w, \varepsilon, k, \tau, \Delta_{\mathcal{L}}(i, j))$  # using Eq. 2
10:  Return  $\mathbb{S}$ 
11: Step 2 - At sample-level explanation:
12:   $w_x \leftarrow$  SHAP explanation for  $x$ 
13:   $w'_x \leftarrow$  Solve the optimization problem in (5)
14:  Return  $w'_x$ 

```

---

and retain a mapping  $v : \mathbb{N} \rightarrow \mathbb{N}$  from the sorted indices to the original indices. Given that  $\tau$  is a predefined threshold to control the range of out-of-top- $k$  features that some of top- $k$  features can swap with, and  $\beta$  is a parameter bounded in Theorem 1 under a privacy budget  $\varepsilon$ , X RAND defines the probability of flipping a top- $k$  feature  $i$  to an out-of-top- $k$  feature  $j$  as follows:

$$\forall i \in [1, k], j \in [k+1, k+\tau], \tau \geq k :$$

$$i = \begin{cases} i, & \text{with probability } p_i = \frac{\exp(\beta)}{\exp(\beta) + \tau - 1}, \\ j, & \text{with probability } q_{i,j} = \frac{\tau - 1}{\exp(\beta) + \tau - 1} q_j \end{cases} \quad (2)$$

where  $q_j = \frac{\exp(-\Delta_{\mathcal{L}}(i,j))}{\sum_{t \in [k+1, k+\tau]} \exp(-\Delta_{\mathcal{L}}(i,t))}$  and  $\Delta_{\mathcal{L}}(i, j)$  is the aggregated changes of  $\mathcal{L}$  (Eq. 1) when flipping features  $i$  and  $j$ , which is calculated as follows:

$$\Delta_{\mathcal{L}}(i, j) = \frac{1}{N} \sum_{n=1}^N (|\mathcal{L}(x_n) - \mathcal{L}(x_n)(i, j)|) \quad (3)$$

where  $\mathcal{L}(x_n)$  is the original loss  $\mathcal{L}$  of a sample  $x_n \in \mathcal{D}$  and  $\mathcal{L}(x_n)(i, j)$  is the loss  $\mathcal{L}$  of the sample  $x_n$  after flipping features  $i$  and  $j$  (Alg. 1, lines 3,5).

After randomizing the aggregated explanation, we obtain the set  $\mathbb{S}$  of features that need to be flipped in the aggregated explanation, as follows:

$$\mathbb{S} = \{(i, j) | i \text{ and } j \text{ are flipped}, i \in [1, k], j \in [k+1, k+\tau]\} \quad (4)$$

**Step 2 (Alg. 1, lines 11-14).** For each input test sample  $x$ , we proceed with sample-level explanation for finding the noisy explanation  $w'_x$ . First, we generate a set of constraints  $\mathbb{Q} = \{(i, j) | w'_{x_i} \leq w'_{x_j}\}$  that is *sufficient* for  $\mathbb{S}$ . In particular, for each pair  $(i, j) \in \mathbb{S}$ , we add the following pairs to  $\mathbb{Q}$ :  $(v(i+1), v(j)); (v(j), v(i-1)); (v(I), v(j-1)); (v(j+1), v(i))$ . Given  $w_x$  as the SHAP explanation of  $x$ , we aim to

find  $\phi \in \mathbb{R}^d$  such that  $w'_x = w_x + \phi$  satisfies the constraints in  $\mathbb{Q}$  while minimizing the loss  $\mathcal{L}$ . To obtain  $\phi$ , we solve the following optimization problem:

$$\min_{\phi} \sum_{z \in N(x)} \left( (w_x + \phi)^T z - f(z) \right)^2 \exp\left(-\frac{\|z - x\|^2}{\sigma^2}\right) + \lambda \|\phi\| \quad (5)$$

$$\text{s.t. } w_{x_i} + \phi_i \leq w_{x_j} + \phi_j, \quad \forall (i, j) \in \mathbb{Q}$$

$$\phi_i = 0 \quad \forall i \notin \mathbb{Q}$$

where  $\lambda$  is a regularization constant.

The resulting noisy explanation will be  $w'_x = w_x + \phi$ . This problem is convex and can be solved by convex optimization solvers (Kingma and Ba 2014; Diamond and Boyd 2016).

## 4.2 Privacy Guarantees of X RAND

To bound privacy loss of X RAND, we need to bound  $\beta$  in Eq. 2 such that the top- $k$  features in the explanation  $w'$  preserves LDP, as follows:

**Theorem 1.** *Given two distinct explanations  $w$  and  $\tilde{w}$  and a privacy budget  $\varepsilon_i$ , X RAND satisfies  $\varepsilon_i$ -LDP in randomizing each feature  $i$  in top- $k$  features of  $w$ , i.e.,  $\frac{P(\text{X RAND}(w_i)=z|w)}{P(\text{X RAND}(\tilde{w}_i)=z|\tilde{w})} \leq \exp(\varepsilon_i)$ , if:*

$$\beta \leq \varepsilon_i + \ln(\tau - 1) + \ln\left(\min_{\sum_{t=k+1}^{k+\tau} \exp(-\Delta_{\mathcal{L}}(i, t))}\frac{\exp(-\Delta_{\mathcal{L}}(i, j))}{\exp(-\Delta_{\mathcal{L}}(i, t))}\right)$$

where  $z \in \text{Range}(\text{X RAND})$ . **Proof:** See Appx. A.

Based on Theorem 1, the total privacy budget  $\varepsilon$  to randomize all top- $k$  features is the sum of all the privacy budget  $\varepsilon_i$ , i.e.,  $\varepsilon = \sum_{i=1}^k \varepsilon_i$ , since each feature  $i$  is randomized independently. From Theorem 1 and Eq. 2, it can be seen that as the privacy budget  $\varepsilon$  increases,  $\beta$  can increase and the flipping probability  $q_{i,j}$  decreases. As a result, we switch fewer features out of top- $k$ .

**Privacy and Explainability Trade-off.** To understand the privacy and explainability trade-off, we analyze the data utility of X RAND mechanism through the sum square error (SSE) of the original explanation  $w$  and the one resulting from X RAND  $w'$ . The smaller the SSE is, the better data utility the randomization mechanism achieves.

**Theorem 2.** *Utility of X RAND:  $\text{SSE} = \sum_{x \in \mathcal{D}} \sum_{i=1}^d (w'_{x_i} - w_{x_i})^2 = \sum_{x \in \mathcal{D}} \sum_{i=1}^{k+\tau} (w'_{x_i} - w_{x_i})^2$ , where  $d$  is the number of features in the explanation.*

*Proof.* It is easy to see that we only consider the probability of flipping the top- $k$  features to be out-of-the-top- $k$  up to the feature  $k+\tau$ . Thus, all features after  $k+\tau$ , i.e., from  $k+\tau+1$  to  $d$  are not changed. Hence the theorem follows.  $\square$

From the theorem, at the same  $\varepsilon$ , the smaller the  $\tau$ , the higher the data utility that X RAND achieves. Intuitively, if  $\tau$  is large, it is more flexible for the top- $k$  features to be flipped out, but it will also impair the model explainability since the original top- $k$  features are more likely to be moved far away from the top  $k$ . With high values of  $\varepsilon$ , we can obtain a smaller SSE, thus, achieving better data utility. The effect of  $\varepsilon$  and  $\tau$  on the SSE value is illustrated in Fig. 6 (Appendix).

## 5 Certified Robustness

Our proposed X RAND can be used as a defense against the XBA since it protects the top- $k$  important features. We further establish the connection between X RAND with certified robustness against XBA. Given a data sample  $x$ : 1) In the training time, we guarantee that up to a portion of poisoning samples in the outsourced training data, XBA fails to change the model predictions; and 2) In the inference time, we guarantee that up to a certain backdoor trigger size, XBA fails to change the model predictions. A primer on certified robustness is given in Appx. B.

### 5.1 Training-time Certified Robustness

We consider the original training data  $\mathcal{D}$  as the proprietary data, and the explanation-guided backdoor samples  $\mathcal{D}_o$  as the outsourced data inserted into the proprietary data. The outsourced data  $\mathcal{D}_o$  alone may not be sufficient to train a good classifier. In addition, the outsourced data inserted into propriety data can lessen the certified robustness bound of the propriety data. Therefore, we cannot quantify the certified poisoning training size of the outsourced data  $\mathcal{D}_o$  directly by applying a bagging technique (Jia, Cao, and Gong 2021). To address this problem, we quantify the certified poisoning training size  $r$  of  $\mathcal{D}_o$  against XBA by uncovering its correlation with the poisoned training data  $\mathcal{D}' = \mathcal{D} \cup \mathcal{D}_o$ .

Given a model prediction on a data sample  $x$  using  $\mathcal{D}$ , denoted as  $f(\mathcal{D}, x)$ , we ask a simple question: ‘‘What is the minimum number poisoning data samples, i.e., certified poisoning training size  $r_{\mathcal{D}}$ , added into  $\mathcal{D}$  to change the model prediction on  $x$ :  $f(\mathcal{D}, x) \neq f(\mathcal{D}_+, x)$ ?’’ After adding  $\mathcal{D}_o$  into  $\mathcal{D}$ , we ask the same question: ‘‘What is the minimum number poisoning data samples, i.e., certified poisoning training size  $r_{\mathcal{D}'}$ , added into  $\mathcal{D}' = \mathcal{D} \cup \mathcal{D}_o$  to change the model prediction on  $x$ :  $f(\mathcal{D}', x) \neq f(\mathcal{D}'_+, x)$ ?’’ The difference between  $r_{\mathcal{D}}$  and  $r_{\mathcal{D}'}$  provides us a certified poisoning training size on  $\mathcal{D}_o$ . Intuitively, if  $\mathcal{D}_o$  does not consist of poisoning data examples, then  $r_{\mathcal{D}}$  is expected to be relatively the same with  $r_{\mathcal{D}'}$ . Otherwise,  $r_{\mathcal{D}'}$  can be significantly smaller than  $r_{\mathcal{D}}$  indicating that  $\mathcal{D}_o$  is heavily poisoned with at least  $r = r_{\mathcal{D}} - r_{\mathcal{D}'}$  number of poisoning data samples towards opening backdoors on  $x$ .

**Theorem 3.** *Given two certified poisoning training sizes  $r_{\mathcal{D}}^* = \arg \min_{r_{\mathcal{D}}} r_{\mathcal{D}}$  and  $r_{\mathcal{D}'}^* = \arg \min_{r_{\mathcal{D}'}} r_{\mathcal{D}'}$ , the certified poisoning training size  $r$  of the outsourced data  $\mathcal{D}_o$  is:*

$$r = r_{\mathcal{D}}^* - r_{\mathcal{D}'}^* \quad (6)$$

**Proof:** Refer to Appx. C for the proof and its tightness.

### 5.2 Inference-time Certified Robustness

It is not straightforward to adapt existing certified robustness bounds at the inference-time into X RAND, since there is a gap between model training as in existing approaches (Jia, Cao, and Gong 2021; Lecuyer et al. 2019; Phan et al. 2020) and the model training with explanation-guided poisoned data as in our system. Existing approaches can randomize the data input  $x$  and then derive certified robustness bounds given the varying output. This typical process does not consider the explanation-guided poisoned data that can

potentially affect certified robustness bounds in our system. To address this gap, note that we can always find a mechanism to inject random noise into the data samples  $x$  such that the samples achieve the same level of DP guarantee as the explanations. Based on this, we can generalize existing certified bounds against XBA at the inference time in X RAND.

When explanation-guided backdoor samples are inserted into the training data, upon bounding the sensitivity that the backdoor samples change the output of  $f$ , there always exists a noise  $\alpha$  that can be injected into a benign sample  $x$ , i.e.,  $x + \alpha$ , to achieve an equivalent  $\epsilon$ -LDP protection. Given the explanation  $w_x$  of  $x$ , we focus on achieving a robustness condition to  $L_p(\mu)$ -norm attacks, where  $\mu$  is the radius of the norm ball, formulated as follows:

$$\forall \alpha \in L_p(\mu) : f_l(x + \alpha|w_x) > f_{-l}(x + \alpha|w_x) \quad (7)$$

where  $l \in \{0, 1\}$  is the true label of  $x$  and  $-l$  is the NOT operation of  $l$  in a binary classification problem.

There should exist a correlation among  $\alpha$  and  $w_x$  that needs to be uncovered in order to bound the robustness condition in Eq. 7. However, it is challenging to find a direct mapping function  $\mathcal{M} : w_x \rightarrow \alpha$  so that when we randomize  $w_x$ , the change of  $\alpha$  is quantified. We address this challenge by quantifying the sensitivity of  $\alpha$  given the average change of the explanation of multiple samples  $x \in \mathcal{X}$ , as follows:

$$\Delta_{\alpha|w} = \frac{1}{|\mathcal{X}|d} \sum_{x \in \mathcal{X}} |w_x - w'_x|_1 \quad (8)$$

where  $|\mathcal{X}|$  is the size of  $\mathcal{X}$ .

$\Delta_{\alpha|w}$  can be considered as a bounded sensitivity of X RAND given the input  $x$  since: **(1)** We can achieve the same DP guarantee by injecting Laplace or Gaussian noise into the input  $x$  using the sensitivity  $\Delta_{\alpha|w}$ ; and **(2)** The explanation perturbation happens only once and is permanent, that is, there is no other bounded sensitivity associated with the one-time explanation perturbation. The sensitivity  $\Delta_{\alpha|w}$  establishes a new connection between explanation perturbation and the model sensitivity given the input sample  $x$ . That enables us to derive robustness bounds using different techniques, i.e., PixelDP (Lecuyer et al. 2019; Phan et al. 2019) and (boosting) randomized smoothing (RS) (Horvath et al. 2022; Cohen, Rosenfeld, and Kolter 2019), since we consider the sensitivity  $\Delta_{\alpha|w}$  as a part of randomized smoothing to derive and enhance certified robustness bounds.

The rest of this section only discusses the bound using PixelDP, we refer the readers to Appx. D for the bound using boosting RS. Given a randomized prediction  $f(x)$  satisfying  $(\epsilon, \delta)$ -PixelDP w.r.t. a  $L_p(\mu)$ -norm metric, we have:

$$\forall l \in \{0, 1\}, \forall \alpha \in L_p(\mu) : \mathbb{E}f_l(x) \leq e^\epsilon \mathbb{E}f_l(x + \alpha) + \delta \quad (9)$$

where  $\mathbb{E}f_l(x)$  is the expected value of  $f_l(x)$ ,  $\epsilon$  is a predefined privacy budget, and  $\delta$  is a broken probability. When we use a Laplace noise,  $\delta = 0$ .

We then apply PixelDP with the sensitivity  $\Delta_{\alpha|w}$  and a noise standard deviation  $\sigma = \frac{\Delta_{\alpha|w}\mu}{\epsilon}$  for Laplace noise, or  $\sigma = \frac{\Delta_{\alpha|w}\mu\sqrt{2\ln(1.25/\delta)}}{\epsilon}$  for Gaussian noise. From that, when maximizing the attack trigger’s magnitude  $\mu$ :  $\mu_{\max} =$




Figure 2: Attack success rate (%) as a function of privacy budget  $\epsilon$  and the portion of poisoned samples on LightGBM and EmberNN. The trigger size is fixed at 10 and 16 for LightGBM and EmberNN, respectively.

$\max_{\mu \in \mathbb{R}^+} \mu$  such that the generalized robustness condition (Eq. 9) holds, the prediction on  $x$  using XRAND is robust against the XBA up to  $\mu_{\max}$ . As a result, we have a robustness certificate of  $\mu_{\max}$  for  $x$ .

## 6 Experiments

To evaluate the performance of our defense, we conduct the XBA proposed by (Severi et al. 2021) against the explanations returned by XRAND to create backdoors on cloud-hosted malware classifiers. Our experiments aim to shed light on understanding (1) the effectiveness of the defense in mitigating the XBA, and (2) the faithfulness of the explanations returned by XRAND. The experiments are conducted using LightGBM (Anderson and Roth 2018) and EmberNN (Severi et al. 2021) classification models that are trained on the EMBER (Anderson and Roth 2018) dataset. A detailed description of the experimental settings can be found in Appx. E. We quantify XRAND via the following metrics:

**Attack success rate.** This metric is defined as the portion of trigger-embedded malware samples that are classified as goodware by the backdoored model. Note that we only embed the trigger into malware samples that were classified correctly by the clean model. The primary goal of our defense is to reduce this value.

**Log-odds.** To evaluate the faithfulness of our XRAND explanation, we compare the log-odds score of the XRAND explanations with that of the ones originally returned by SHAP. Based on the explanation of a sample, the log-odds score is computed by identifying the top 20% important features that, if erased, can change the predicted label of the sample (Shrikumar, Greenside, and Kundaje 2017). Then we obtain the change in log-odd of the prediction score of the original sample and the sample with those features erased. The higher the log-odds score, the better the explanation in terms of identifying important features. To maintain a faithful explainability, the XRAND explanations should have comparable log-odds scores as the original SHAP explanations.

**Attack Mitigation.** We observe the attack success rate of XBA when our XRAND explanations are used to craft backdoor triggers. We set  $k$  to be equal to the trigger size of the attack, and fix the predefined threshold  $\tau = 50$ . Fig. 2 highlights the correlation between the attack success rate and the privacy budget  $\epsilon$  in our defense. Intuitively, the lower




Figure 3: Attack success rate (%) as a function of trigger size and privacy budget  $\epsilon$ . We vary the trigger size and fix the portion of poisoned samples at 0.25% and 1% for LightGBM and EmberNN, respectively.




Figure 4: Log-odds score of the explanations of 20,000 goodware and malware samples. The leftmost box represents the original explanations by SHAP, the remaining boxes illustrate the explanations after applying XRAND at  $\epsilon = 1.0, 10.0, 20.0$ , respectively.

the  $\epsilon$ , the more obfuscating the top goodware-oriented features become. Hence, Figs. 2a and 2b show that the attack success rate is greatly diminished as we tighten the privacy budget, since the attacker has less access to the desired features. Moreover, in a typical backdoor attack, injecting more poisoned samples into the training data makes the attack more effective. Such behavior is exhibited in both LightGBM (Fig. 2a) and EmberNN (Fig. 2b), though EmberNN is less susceptible to the increase of poisoned data. However, in practice, the attacker wishes to keep the number of poisoning samples relatively low to remain stealthy. At a 1% poison rate, our defense manages to reduce the attack success rate from 77.8% to 10.2% with  $\epsilon = 1.0$  for LightGBM. It performs better with EmberNN where the attack success rate is reduced to 5.3% at  $\epsilon = 10.0$ .

Additionally, we examine the effect of the trigger sizes on our defense. The trigger size denotes the number of features that the attacker modifies to craft poisoned samples. We vary the trigger size consistently with previous work (Severi et al. 2021). In backdoor attacks, a larger trigger makes the backdoored model more prone to misclassification, thus improving the attack success rate. Fig. 3 shows that the attack works better with large triggers on both models, however, as aforementioned, the attacker would prefer small trigger sizes for the sake of stealthiness. This experiment shows that, for the trigger sizes that we tested, our proposed defense can successfully maintain a low attack success rate given a wide range of the privacy budget  $\epsilon \in [0.1, 10]$  (Fig. 2a, 2b).



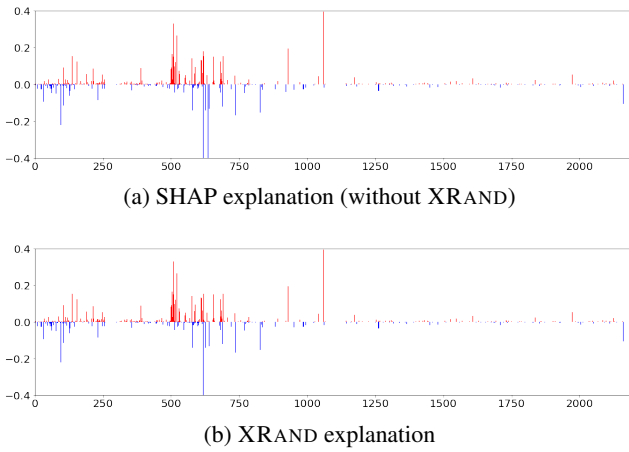


Figure 5: Visualizing the SHAP explanation and our XRAND explanation of a test sample. The plot shows the SHAP value of each feature. The red vertical lines represent positive SHAP values that indicate malware-oriented features, while the blue vertical lines are negative SHAP values indicating goodware-oriented features.

We refer the readers to Appx. E for more experimental results on additional malware datasets and the evaluation of our robustness bounds. In short, regarding the training-time bound, we observe that a smaller privacy budget  $\epsilon \in [0.1, 10]$  results in a more robust model against the XBA. As for the inference-time bound, we obtain high certified accuracy (Eq. 29) under rigorous privacy budgets, i.e., 89.17% and 90.42% at  $\epsilon \in [0.1, 1.0]$ . We notice that the PixelDP-based bound attains a stronger performance than using boosting RS. This is because boosting RS is not designed for models trained on backdoored data like XRAND.

**Faithful Explainability.** From the previous results, we observe that a wide range of the privacy budget  $\epsilon \in [0.1, 10]$  provides a good defense against the XBA. The question remains whether the explanations resulting from these values of  $\epsilon$  are still faithful. Fig. 4 shows the log-odds score of the original explanations returned by SHAP and of the ones after applying our XRAND mechanism. The XRAND explanations at  $\epsilon = 1.0, 10.0$  have comparable log-odds scores to those of SHAP. This is because our defense works with small values of  $k$  (e.g.,  $k = 10$ ). Therefore, XRAND only randomizes the SHAP values within a small set of top goodware-oriented features. As a result, XRAND can still capture the important features in its explanations.

Furthermore, we visualize the explanation of a test sample before and after applying XRAND in Figs. 5a and 5b, respectively. As can be seen, the SHAP values of the two explanations largely resemble one another, except for minor differences in less than 10 features (out of 2,351 features). Importantly, the XRAND explanation evidently manifests similar sets of important malware and goodware features as the original explanation by SHAP, which also explains the comparable log-odds score in Fig. 4. More visualizations on XRAND can be found in Appx. F.

## 7 Related Work

### Vulnerability of Explanations and Backdoor Defenses.

There are two main lines of research in studying the vulnerability of model explanations. First, model explainers are prone to adversarial manipulations, thereby making the explanations unreliable and untrustworthy (Heo, Joo, and Moon 2019; Slack et al. 2020; Zhang et al. 2020; Ghorbani, Abid, and Zou 2019). Second, the information leaked through the explanations can be exploited to conduct several attacks against the black-box models (Shokri, Strobel, and Zick 2021; Milli et al. 2019; Miura, Hasegawa, and Shibahara 2021; Zhao et al. 2021; Severi et al. 2021).

Our work focuses on the second direction in which we show that MLaaS is particularly susceptible to this type of attack. Although such vulnerability has been studied before, the discussion on its defense mechanism is still limited. To that extent, we propose a new concept of LDP explanation such that we can obfuscate the explanations in a way that limits the knowledge that can be leveraged by the attackers while maintaining the faithfulness of the explanations.

Moreover, several defenses have been proposed to tackle the backdoor attacks (Liu, Dolan-Gavitt, and Garg 2018; Tran, Li, and Madry 2018; Wang et al. 2019a). As evidenced by (Severi et al. 2021) that the backdoors created by the XBA remain stealthy against these defenses. Thus, our XRAND is the first solution that can tackle XBA in MLaaS.

**LDP Mechanisms.** Existing LDP mechanisms can be classified into four lines. **1)** Direct encoding approaches (Warner 1965; Kairouz, Bonawitz, and Ramage 2016) are to directly apply RR mechanisms to a binary response to output the true and false responses with different probabilities. **2)** Generalized RR approaches (Duchi and Rogers 2019; Wang et al. 2019b; Zhao et al. 2020) directly perturb the numerical (real) value of data into a finite set or continuous range of output values. **3)** Hash-based approaches (Erlingsson, Pihur, and Korolova 2014; Wang et al. 2017; Bassily and Smith 2015; Acharya, Sun, and Zhang 2019; Wang et al. 2019c). **4)** Binary encoding-based approaches (Arachchige et al. 2019; Lyu et al. 2020) converts input  $x$  into a binary encoding vector of 0 and 1, then independently applies RR mechanisms on these bits. Despite the effectiveness of these RR mechanisms, existing approaches have not been designed for XBA, in which the explanations expose the training data background.

## 8 Conclusion

In this paper, we have shown that, although explanations help improve the understanding and interpretability of black-box models, they also leak essential information about the inner workings of the models. Therefore, the black-box models become more vulnerable to attacks, especially in the context of MLaaS where the prediction and its explanation are returned for each query. With a novel two-step LDP-preserving mechanism, we have proposed XRAND to protect the model explanations from being exploited by adversaries via obfuscating the top important features, while maintaining the faithfulness of explanations.

## Acknowledgments

This material is based upon work supported by the National Science Foundation under grants IIS-2123809, IIS-1939725, CNS-1935928, and CNS-1935923.

## References

- Acharya, J.; Sun, Z.; and Zhang, H. 2019. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 1120–1129.
- Anderson, H. S.; and Roth, P. 2018. Ember: an open dataset for training static malware machine learning models. *arXiv preprint arXiv:1804.04637*.
- Arachchige, P. C. M.; Bertok, P.; Khalil, I.; Liu, D.; Camtepe, S.; and Atiquzzaman, M. 2019. Local differential privacy for deep learning. *IEEE Internet of Things Journal*, 7(7): 5827–5842.
- Azure. 2021. Model Interpretability (preview) - azure machine learning. <https://aka.ms/AzureMLModelInterpretability>. Accessed: 2022-04-21.
- Bassily, R.; and Smith, A. 2015. Local, private, efficient protocols for succinct histograms. In *ACM Symposium on Theory of Computing*, 127–135.
- Bluemix. 2021. Ai explainability 360. <https://aix360.mybluemix.net/>. Accessed: 2022-04-21.
- Cohen, J.; Rosenfeld, E.; and Kolter, Z. 2019. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, 1310–1320.
- Community, F. 2018. Explainable Machine Learning Challenge. <https://community.fico.com/s/explainable-machine-learning-challenge>. Accessed: 2022-06-03.
- Diamond, S.; and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83): 1–5.
- Duchi, J.; and Rogers, R. 2019. Lower bounds for locally private estimation via communication complexity. In *COLT*, 1161–1191.
- Duchi, J. C.; Jordan, M. I.; and Wainwright, M. J. 2018. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521): 182–201.
- Erlingsson, U.; Pihur, V.; and Korolova, A. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC CCS*, 1054–1067.
- Ghorbani, A.; Abid, A.; and Zou, J. 2019. Interpretation of neural networks is fragile. *Proceedings of the AAAI conference on artificial intelligence*, 33(01): 3681–3688.
- Heo, J.; Joo, S.; and Moon, T. 2019. Fooling neural network interpretations via adversarial model manipulation. *Advances in Neural Information Processing Systems*, 32.
- Horváth, M.; Mueller, M.; Fischer, M.; and Vechev, M. 2022. Boosting Randomized Smoothing with Variance Reduced Classifiers. In *International Conference on Learning Representations*.
- Jia, J.; Cao, X.; and Gong, N. Z. 2021. Intrinsic certified robustness of bagging against data poisoning attacks. *AAAI*.
- Kairouz, P.; Bonawitz, K.; and Ramage, D. 2016. Discrete distribution estimation under local privacy. In *International Conference on Machine Learning*, 2436–2444.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lecuyer, M.; Atlidakis, V.; Geambasu, R.; Hsu, D.; and Jana, S. 2019. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, 656–672.
- Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2018. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *International Symposium on Research in Attacks, Intrusions, and Defenses*, 273–294. Springer.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st international conference on neural information processing systems*, 4768–4777.
- Lyu, L.; Li, Y.; He, X.; and Xiao, T. 2020. Towards differentially private text representations. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1813–1816.
- Milli, S.; Schmidt, L.; Dragan, A. D.; and Hardt, M. 2019. Model reconstruction from model explanations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 1–9.
- Miura, T.; Hasegawa, S.; and Shibahara, T. 2021. MEGEX: Data-Free Model Extraction Attack against Gradient-Based Explainable AI. *arXiv preprint arXiv:2107.08909*.
- Nguyen, T.; Lai, P.; Phan, N.; and Thai, M. T. 2022. XRand: Differentially Private Defense against Explanation-Guided Attacks. *arXiv preprint arXiv:2212.04454*.
- Phan, H.; Thai, M.; Hu, H.; Jin, R.; Sun, T.; and Dou, D. 2020. Scalable differential privacy with certified robustness in adversarial learning. In *International Conference on Machine Learning*, 7683–7694.
- Phan, N.; Vu, M. N.; Liu, Y.; Jin, R.; Dou, D.; Wu, X.; and Thai, M. T. 2019. Heterogeneous Gaussian mechanism: preserving differential privacy in deep learning with provable robustness. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 4753–4759.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Severi, G.; Meyer, J.; Coull, S.; and Oprea, A. 2021. Explanation-Guided Backdoor Poisoning Attacks Against Malware Classifiers. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.



- Shokri, R.; Strobel, M.; and Zick, Y. 2021. On the privacy risks of model explanations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 231–241.
- Shrikumar, A.; Greenside, P.; and Kundaje, A. 2017. Learning important features through propagating activation differences. In *International Conference on Machine Learning*, 3145–3153. PMLR.
- Slack, D.; Hilgard, S.; Jia, E.; Singh, S.; and Lakkaraju, H. 2020. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 180–186.
- Sun, L.; Qian, J.; and Chen, X. 2021. LDP-FL: Practical private aggregation in federated learning with local differential privacy. *IJCAI*.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, 3319–3328. PMLR.
- Tran, B.; Li, J.; and Madry, A. 2018. Spectral signatures in backdoor attacks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 8011–8021.
- Vu, M.; and Thai, M. 2020. PGM-Explainer: Probabilistic Graphical Model Explanations for Graph Neural Networks. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- Vu, M. N.; Nguyen, T.; and Thai, M. T. 2022. NeuCEPT: Learn Neural Networks’s Mechanism via Critical Neurons with Precision Guarantee. In *2022 IEEE International Conference on Data Mining (ICDM)*. IEEE.
- Vu, M. N.; Nguyen, T. D.; Phan, N.; Gera, R.; and Thai, M. T. 2021. c-Eval: A unified metric to evaluate feature-based explanations via perturbation. In *2021 IEEE International Conference on Big Data (Big Data)*, 927–937. IEEE.
- Wang, B.; Yao, Y.; Shan, S.; Li, H.; Viswanath, B.; Zheng, H.; and Zhao, B. Y. 2019a. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*, 707–723. IEEE.
- Wang, N.; Xiao, X.; Yang, Y.; Zhao, J.; Hui, S. C.; Shin, H.; Shin, J.; and Yu, G. 2019b. Collecting and analyzing multidimensional data with local differential privacy. In *IEEE ICDE*, 638–649.
- Wang, T.; Blocki, J.; Li, N.; and Jha, S. 2017. Locally differentially private protocols for frequency estimation. In *26th USENIX Security Symposium*, 729–745.
- Wang, T.; Ding, B.; Xu, M.; et al. 2019c. MURS: practical and robust privacy amplification with multi-party differential privacy. In *ACSAC*.
- Warner, S. L. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309): 63–69.
- Xiong, X.; Liu, S.; Li, D.; Cai, Z.; and Niu, X. 2020. A comprehensive survey on local differential privacy. *Security and Communication Networks*.
- Ying, R.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems*, 32: 9240.
- Zhang, X.; Wang, N.; Shen, H.; Ji, S.; Luo, X.; and Wang, T. 2020. Interpretable deep learning under fire. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*.
- Zhao, X.; Zhang, W.; Xiao, X.; and Lim, B. 2021. Exploiting Explanations for Model Inversion Attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 682–692.
- Zhao, Y.; Zhao, J.; Yang, M.; Wang, T.; Wang, N.; Lyu, L.; Niyato, D.; and Lam, K. Y. 2020. Local differential privacy based federated learning for Internet of Things. *IEEE Internet of Things Journal*.