AGENCE NATIONALE DE LA RECHERCHE





Deliverable D5.d

# **Estimating g-Leakage via Machine Learning**

Marco Romanelli marco.romanelli@inria.fr Inria, École Polytechnique, IPP, Universitá di Siena

> Catuscia Palamidessi catuscia@lix.polytechnique.fr Inria, École Polytechnique, IPP

# ABSTRACT

This paper considers the problem of estimating the information leakage of a system in the black-box scenario. It is assumed that the system's internals are unknown to the learner, or anyway too complicated to analyze, and the only available information are pairs of input-output data samples, possibly obtained by submitting queries to the system or provided by a third party. Previous research has mainly focused on counting the frequencies to estimate the input-output conditional probabilities (referred to as frequentist approach), however this method is not accurate when the domain of possible outputs is large. To overcome this difficulty, the estimation of the Bayes error of the ideal classifier was recently investigated using Machine Learning (ML) models and it has been shown to be more accurate thanks to the ability of those models to learn the input-output correspondence. However, the Bayes vulnerability is only suitable to describe one-try attacks. A more general and flexible measure of leakage is the *q*-vulnerability, which encompasses several different types of adversaries, with different goals and capabilities. In this paper, we propose a novel approach to perform black-box estimation of the q-vulnerability using ML. A feature of our approach is that it does not require to estimate the conditional probabilities, and that it is suitable for a large class of ML algorithms. First, we formally show the learnability for all data distributions. Then, we evaluate the performance via various experiments using k-Nearest Neighbors and Neural Networks. Our results outperform the frequentist approach when the observables domain is large.

# **KEYWORDS**

g-vulnerability ; estimation; machine learning; neural networks

# **1** INTRODUCTION

The information leakage of a system is a fundamental concern of computer security, and measuring the amount of sensitive information that an adversary can obtain by observing the outputs of a given system is of the utmost importance to understand whether such leakage can be tolerated or must be considered a major security flaw. Much research effort has been dedicated to studying Konstantinos Chatzikokolakis kostas@chatzi.org University of Athens

Pablo Piantanida pablo.piantanida@centralesupelec.fr CentraleSupelec, CNRS, Université Paris Saclay

and proposing solutions to this problem, see for instance the works [2, 4, 7, 10, 11, 19, 27, 33], just to mention a few. So far, this area of research, known as quantitative information flow (QIF), has mainly focused on the so-called white-box scenario. Namely, all those works assume that the channel of the system is known, or can be computed by analyzing the system's internals. This channel consists of the conditional probabilities of the outputs (observables) given the inputs (secrets).

However, the white-box assumption is not always realistic: sometimes the system is unknown, or anyway it is too complex, so that an analytic computation becomes hard if not impossible to be performed. Therefore, it is important to consider also a black-box approach where we only assume the availability of a finite set of input-output pairs generated by the system, possibly obtained by submitting queries or provided by a third party.

The estimation of the internal probabilities of a system's channel have been investigated in [16] and [18] via a frequentist paradigm, i.e. relying on the computation of the frequencies of the outputs given some inputs. However, this approach does not scale to applications for which the output space is very large since a prohibitively large number of samples would be necessary to achieve good results and fails on continuous alphabets unless some strong assumption on the distributions are made. In order to overcome this limitation, the authors of [14] exploited the fact that Machine Learning (ML) algorithms provide a better scalability to black-box measurements. Intuitively, the advantage of the ML approach over the frequentist one is its generalization power: while the frequentist method can only draw conclusions based on counts on the available samples, ML is able to extrapolate from the samples and provide better prediction (generalization) for the rest of the universe.

In particular, [14] proposed to use k-Nearest Neighbors (k-NN) to measure one of the basic QIF metrics, the Bayes vulnerability [33]. This is the expected probability of success of an adversary that has exactly one attempt at his disposal (one-try), and tries to maximize the chance of guessing the right value of the secret. The idea is that the Bayes vulnerability corresponds to the converse of the error of the ideal Bayes classifier that, given any sample (observable), tries to predict its corresponding class (secret). It is then sufficient to build a model that approximates such classifier, and measure its expected error. The main takeaway is that any ML rule which is *universally consistent* (i.e., approximates the ideal Bayes classifier) has a guarantee on the accuracy of the estimation, i.e., the gap between the estimated leakage and the real leakage tends to vanish as the number of training samples grows large.

The limitation of [14], however, is that it is only concerned with the Bayes vulnerability case, which models only one particular kind

This research was supported by DATAIA Convergence Institute as part of the "Programme d'Investissement d'Avenir" (ANR-17-CONV-0003), operated by Inria and CentraleSupélec. It was also supported by the ANR project REPAS, and by the Inria/DRI project LOGIS. The work of Prof. Pablo Piantanida was supported by the European Commission's Marie Sklodowska-Curie Actions (MSCA), through the Marie Sklodowska-Curie IF (H2020-MSCAIF-2017-EF-797805). The work of Catuscia Palamidessi was supported by the project HYPATIA, funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement n. 835294.

of adversary. As argued in [4], there are many other realistic kinds of adversaries. For instance, adversaries whose aim is to guess only a part of the secret, or a property of the secret, or adversaries that have multiple tries at their disposal. To represent a larger class of attacks, [4] introduced the so-called *g*-vulnerability<sup>1</sup>, based on the notion of gain function *g*. This metric is very general, and in particular it encompasses the Bayes vulnerability.

In this paper, we propose an approach to the black-box estimation of *g*-vulnerability via ML. The idea is to reduce the problem to that of approximating the Bayes classifier, so that *any universally consistent ML algorithm can be used for the purpose*. This reduction essentially takes into account the impact of the gain function in the generation of the training data, and we propose two methods to obtain this effect, which we call *channel pre-processing* and *data pre-processing*, respectively. We evaluate our approach via experiments on various channels and gain functions. In order to show the generality of our approach, we use two different ML algorithms, namely k-NN and Artificial Neural Networks (ANN), and we compare their performances. The experimental results show that our approach provides accurate estimations, and that it outperforms by far the frequentist approach when the observables domain is large.

#### 1.1 Our contribution

The contributions of this paper are:

- We propose a novel approach to the black-box estimation of *g*-vulnerability based on ML. To the best of our knowledge, this is the first time that a method to estimate *g*-vulnerability in a black-box fashion is introduced.
- We provide statistical guarantees showing the learnability of the *g*-vulnerability for all distributions and we derive distribution-free bounds on the accuracy of its estimation.
- We validate the performance of our method via several experiments using k-NN and ANN models. To the best of our knowledge, this is the first time that ANNs are used to estimate information leakage. The code to run these experiments is available at the URL https://github.com/LEAVESrepo/leaves.

#### 1.2 Related work

One important aspect to keep in mind when measuring leakage is the kind of attack that we want to model. In their seminal paper [27], Köpf and Basin identified various kinds of adversaries and showed that they can be captured by known entropy measures. For instance, the Shannon entropy represents the expected number of binary queries that the adversary must submit to the system in order to fully determine the value of the secret.

In [33] Smith proposed another notion, the Rényi min-entropy, to measure the system's leakage when the attacker has only one try at its disposal and attempts to make its best guess. The Rényi minentropy is the logarithm of the Bayes vulnerability, which is the expected probability of the adversary to guess the secret correctly. The Bayes vulnerability is the converse of the Bayes error, which was already proposed as a measure of leakage in [12].

Symbol	Description
$x \in X$	a secret
$w \in W$	a guess
$y\in \mathcal{Y}$	an observable output by the system
X	random variable for secrets, it takes values $x \in X$
W	random variable for guesses, it takes values $w \in \mathcal{W}$
Y	random variable for observables, it takes values $y \in \mathcal{Y}$
$ \mathcal{S} $	size of a set ${\cal S}$
$\mathcal{P}(\mathcal{S})$	Distribution over a set of symbols ${\cal S}$
$\mathcal{H}$	class of learning functions $f$
$\pi, P_X$	prior distribution over the secret space
$\hat{\pi}, \hat{P}_X$	empirical prior distribution over the secret space
C	Channel matrix
$\pi \triangleright C$	joint distribution from prior $\pi$ and channel $C$
$P_{XY}$	joint probability distribution
$\hat{P}_{XY}$	empirical joint probability distribution
$P_{Y X}$	conditional probability of Y given X
$\hat{P}_{Y X}$	empirical conditional probabilities
P	probability measure
$\mathbb{E}[\cdot]$	expected value
g(w, x)	gain function that takes a guess $w$ and secret $x$ as inputs
G	gain matrix of size $ W  \times  X $ according to a specific g
$V_q$	<i>g</i> -vulnerability
V(f)	g-vulnerability functional
$\widehat{V}_n(f)$	empirical <i>g</i> -vulnerability functional evaluated on <i>n</i> samples

Table 1: Table of symbols.

Alvim et al. [4] generalized the notion of Bayes vulnerability to that of g-vulnerability, by introducing a parameter (the gain function g) that describes the adversary's payoff. The g-vulnerability is the expected gain of the adversary in a one-try attack.

Our paper focuses on the estimation of the *g*-vulnerability in the the black-box scenario using ML techniques. This approach is inspired by the seminal work [14], which introduces k-NN algorithms to estimate the Bayes vulnerability, representing a paradigm shift with respect to the previously proposed black-box approaches based on the frequentist paradigm [9, 16, 18]. Notably, it showed that the use of universally consistent learning rules allows to achieve much more precise estimations than the frequentist approach when considering large or even continuous output domains which would be intractable otherwise. However, [14] was limited to the estimation of the Bayes error, i.e., it considered only the Bayes adversary. By contrast, in our paper we consider any adversary that can be modeled by a gain function *g*. Another novelty w.r.t. [14] is that we consider also ANN algorithms, which in various experiments appear to perform better than the k-NN ones.

Bordenabe and Smith [6] investigated the indirect leakage induced by a channel (i.e., leakage on sensitive information not in the domain of the channel), and proved a fundamental equivalence between Dalenius min-entropy leakage under arbitrary correlations and g-leakage under arbitrary gain functions. This result is similar to our Theorem 4.2, and it opens the way to the possible extension of our approach to this more general leakage scenario.

<sup>&</sup>lt;sup>1</sup>The work [4] is perhaps more known for proposing the notion of *g*-leakage. The *g*-vulnerability is the main component of this notion. Indeed, *g*-leakage is *g*-vulnerability normalized w.r.t. the probability of success of the random guess.

#### 2 PRELIMINARIES

In this section, we recall some useful notions from QIF and ML, in particular the ANN and k-NN algorithm.

#### 2.1 Quantitative information flow

Let  $\mathcal{X}$  be a set of secrets and  $\mathcal{Y}$  a set of observations. The adversary's initial knowledge about the secrets is modeled by a prior distribution  $\mathcal{P}(\mathcal{X})$  (namely  $P_X$ ). A system is modeled as a probabilistic *channel* from  $\mathcal{X}$  to  $\mathcal{Y}$ , described by a stochastic matrix C, whose elements  $C_{xy}$  give the probability to observe  $y \in \mathcal{Y}$  when the input is  $x \in \mathcal{X}$  (namely  $P_{Y|X}$ ). Running C with input  $\pi$  induces a joint distribution on  $\mathcal{X} \times \mathcal{Y}$  denoted by  $\pi \triangleright C$ .

In the *g*-leakage framework [4] an adversary is described by a set W of *guesses* (or *actions*) that it can make about the secret, and by a *gain function* g(w, x) expressing the gain of selecting the guess *w* when the real secret is *x*. The prior *g*-vulnerability is the *expected gain* of an optimal guess, given a prior distribution on secrets:

$$V_g(\pi) \stackrel{\text{def}}{=} \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \pi_x \cdot g(w, x) \,. \tag{1}$$

In the posterior case, the adversary observes the output of the system which allows to improve its guess. Its expected gain is given by the posterior *g*-vulnerability, according to

$$V_g(\pi, C) \stackrel{\text{def}}{=} \sum_{y \in \mathcal{Y}} \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \pi_x \cdot C_{xy} \cdot g(w, x) \,. \tag{2}$$

Finally, the multiplicative<sup>2</sup> and additive *g*-leakage quantify how much a specific channel *C* increases the vulnerability of the system:

$$\mathcal{L}_{g}^{\mathrm{M}}(\pi, C) \stackrel{\mathrm{def}}{=} \frac{V_{g}(\pi, C)}{V_{g}(\pi)} , \quad \mathcal{L}_{g}^{\mathrm{A}}(\pi, C) \stackrel{\mathrm{def}}{=} V_{g}(\pi, C) - V_{g}(\pi) .$$
(3)

The choice of the gain function g allows to model a variety of different adversarial scenarios. The simplest case is the *identity* gain function, given by  $\mathcal{W} = \mathcal{X}$ ,  $g_{id}(w, x) = 1$  iff x = w and 0 otherwise. This gain function models an adversary that tries to guess the secret exactly in one try;  $V_{g_{id}}$  is the *Bayes-vulnerability*, which corresponds to the complement of the Bayes error (cfr. [4]).

However, the interest in *g*-vulnerability lies in the fact that many more adversarial scenarios can be captured by a proper choice of *g*. For instance, taking  $\mathcal{W} = \chi^k$  with g(w, x) = 1 iff  $x \in w$  and 0 otherwise, models an adversary that tries to guess the secret correctly in *k* tries. Moreover, guessing the secret *approximately* can be easily expressed by constructing *g* from a metric *d* on X; this is a standard approach in the area of *location privacy* [31, 32] where g(w, x) is taken to be inversely proportional to the Euclidean distance between *w* and *x*. Several other gain functions are discussed in [4], while [3] shows that *any* vulnerability function satisfying basic axioms can be expressed as  $V_q$  for a properly constructed *g*.

The main focus of this paper is estimating the posterior *g*-vulnerability of the system from such samples. Note that, given  $V_g(\pi, C)$ , estimating  $\mathcal{L}_g^{\mathrm{M}}(\pi, C)$  and  $\mathcal{L}_g^{\mathrm{A}}(\pi, C)$  is straightforward, since  $V_g(\pi)$  only depends on the prior (not on the system) and it can be either computed analytically or estimated from the samples.

#### 2.2 Artificial Neural Networks

We provide a short review of the aspects of ANN that are relevant for this work. For further details, we refer to [5, 25, 26]. Neural networks represent an attempt to reproduce the behavior of the brain's cells and are usually modeled as directed graphs with weights on the connections and nodes that forward information through "activation functions", often introducing non-linearity (such as sigmoids or soft-max). In particular, we consider an instance of learning known as *supervised learning*, where input samples are provided to the network model together with target labels (supervision). From the provided data and by means of iterative updates of the connection weights, the network learns how the data and respective labels are distributed. The training procedure, known as back-propagation, is an optimization problem aimed at minimizing a loss function that quantifies the quality of the network's prediction with respect to the data.

Classification problems are a typical example of tasks for which supervised learning works well. Samples are provided together with target labels which represent the classes they belong to. A model can be trained using these samples and, later on, it can be used to predict the class of new samples.

The No Free Lunch theorem (NFL) [34] holds for ANN as it does for all the ML approaches. Therefore, in general, it cannot be said that ANN are better than other ML methods. However, it is well known that the NFL can be broken by additional information on the data or the particular problem we want to tackle, and, nowadays, for most applications and available data, especially in multidimensional domains, ANN models outperform other methods and therefore they represent the state of the art.

#### 2.3 k-Nearest Neighbors

The k-NN algorithm is one of the simplest algorithms used to classify a new sample given a training set of samples labelled as belonging to specific classes. This algorithm assumes that the space of the features is equipped with a notion of distance. The basic idea is the following: every time we need to classify a new sample, we find the k samples whose features are closest to those of the new one (nearest neighbors). Once the k nearest neighbors are selected, a majority vote over their class labels is performed to decide which class should be assigned to the new sample. For further details, as well as for an extensive analysis of the topic, we refer to the chapters about k-NN in [26, 30].

# 3 LEARNING *g*-VULNERABILITY: STATISTICAL BOUNDS

This section introduces the mathematical problem of learning g-vulnerability. More specifically, we address the problem of characterizing universal learnability in the present framework, and to this end, we derive distribution-free bounds on the accuracy of the estimation, implying statistical consistence of our estimator.

#### 3.1 Main definitions

We consider the problem of estimating the g-vulnerability from samples via ML models, and we show that the analysis of this estimation can be conducted in the general statistical framework

 $<sup>^2</sup>$  In the original paper, the multiplicative version of g-leakage was defined as the log of the definition given here. In recent literature, however, the log is not used anymore. Anyway, the two definitions are equivalent from the point of view of comparing systems, since log is a monotonic function.

of maximizing an expected functional using observed samples. The idea can be described using three components:

- A generator of random secrets *x* ∈ *X* with |*X*| < ∞, drawn independently from a fixed but unknown distribution *P<sub>X</sub>(x)*;
- a channel that returns an observable *y* ∈ 𝒴 with |𝒴| < ∞ for every input *x*, according to a conditional distribution *P*<sub>Y|X</sub>(*y*|*x*), also fixed and unknown;
- a learning machine capable of implementing a set of rules  $f \in \mathcal{H}$ , where  $\mathcal{H}$  denotes the class of functions  $f : \mathcal{Y} \to \mathcal{W}$  and  $\mathcal{W}$  is the finite set of guesses.

Moreover let us note that

$$g: \mathcal{W} \times \mathcal{X} \to [a, b] \tag{4}$$

for some finite real values  $a \ge 0$  and b > a, and X and W are finite sets. The problem of learning the *g*-vulnerability is that of choosing the function  $f : \mathcal{Y} \to \mathcal{W}$  which maximizes the functional V(f), representing the expected gain, defined as:

$$V(f) \stackrel{\text{def}}{=} \sum_{(x,y)\in\mathcal{X}\times\mathcal{Y}} g(f(y),x) P_{XY}(x,y).$$
(5)

Note that f(y) corresponds to the "guess" *w*, for the given *y*, in (2). The maximum of V(f) is the *g*-vulnerability, namely:

$$V_g \stackrel{\text{def}}{=} \max_{f \in \mathcal{H}} V(f). \tag{6}$$

# 3.2 Principle of the empirical *g*-vulnerability maximization

Since we are in the black box scenario, the joint probability distribution  $P_{XY} \equiv \pi \triangleright C$  is unknown. We assume, however, the availability of *m* independent and identically distributed (i.i.d.) samples drawn according to  $P_{XY}$  that can be used as a training set  $\mathcal{D}_m$  to solve the maximization of *f* over  $\mathcal{H}$  and additionally *n* i.i.d. samples are available to be used as a validation<sup>3</sup> set  $\mathcal{T}_n$  to estimate the average in (5). Let us denote these sets as:  $\mathcal{D}_m \stackrel{\text{def}}{=} \{(x_1, y_1), \dots, (x_m, y_m)\}$ and  $\mathcal{T}_n \stackrel{\text{def}}{=} \{(x_{m+1}, y_{m+1}), \dots, (x_{m+n}, y_{m+n})\}$ , respectively.

In order to maximize the *g*-vulnerability functional (5) for an unknown probability measure  $P_{XY}$ , the following principle is usually applied. The expected *g*-vulnerability functional V(f) is replaced by the *empirical g-vulnerability functional*:

$$\widehat{V}_n(f) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{(x,y)\in\mathcal{T}_n} g(f(y), x), \tag{7}$$

which is evaluated on  $\mathcal{T}_n$  rather than  $P_{XY}$ . This estimator is clearly unbiased in the sense that

$$\mathbb{E}\big[\widehat{V}_n(f)\big] = V(f).$$

Let  $f_m^{\star}$  denote the empirical optimal rule given by

$$f_m^{\star} \stackrel{\text{def}}{=} \arg\max_{f \in \mathcal{H}} \widehat{V}_m(f), \ \widehat{V}_m(f) \stackrel{\text{def}}{=} \frac{1}{m} \sum_{(x,y) \in \mathcal{D}_m} g(f(y), x), \quad (8)$$

which is evaluated on  $\mathcal{D}_m$  rather than  $P_{XY}$ . The function  $f_m^{\star}$  is the optimizer according to  $\mathcal{D}_m$ , namely the best way among the functions  $f : \mathcal{Y} \to \mathcal{W}$  to approximate  $V_q$  by maximizing  $\hat{V}_m(f)$  over the class of functions  $\mathcal{H}$ . This principle is known in statistics as the Empirical Risk Maximization (ERM).

Intuitively, we would like  $f_m^*$  to give a good approximation of the *g*-vulnerability, in the sense that its expected gain

$$V(f_m^{\star}) = \sum_{(x,y)\in X\times\mathcal{Y}} g(f_m^{\star}(y), x) P_{XY}(x, y)$$
(9)

should be close to  $V_q$ . Note that the difference

$$V_g - V(f_m^{\star}) = \max_{f \in \mathcal{H}} V(f) - V(f_m^{\star})$$
(10)

is always non negative and represents the gap by selecting a possible suboptimal function  $f_m^*$ . Unfortunately, we are not able to compute  $V(f_m^*)$  either, since  $P_{XY}$  is unknown and thus, (10) cannot be measured in practice. In its place, we have to use its approximation  $\hat{V}_n(f_m^*)$  and (10) should be replaced by  $V_g - \hat{V}_n(f_m^*)$  which is not always non-negative.

By using basics principles from statistical learning theory, we study two main questions:

- When does the estimator  $\widehat{V}_n(f_m^*)$  work? What are the conditions for its statistical consistency?
- How well does  $\hat{V}_n(f_m^*)$  approximate  $V_g$ ? In other words, how fast does the sequence of largest empirical g-leakage values converge to the largest g-leakage function? This is related to the so called rate of generalization of a learning machine that implements the ERM principle.

# 3.3 Distribution-free bounds on the estimation accuracy

We start with the following lemma which is a simple adaption of the uniform deviations of relative frequencies from probabilities theorems in [21] (see Appendix B.1 for the proof).

LEMMA 3.1. The following inequalities hold:

$$V_g - V(f_m^{\star}) \leq 2 \max_{f \in \mathcal{H}} \left| \widehat{V}_m(f) - V(f) \right|, \tag{11}$$

$$\left|\widehat{V}_{n}(f_{m}^{\star}) - V(f_{m}^{\star})\right| \leq \max_{f \in \mathcal{H}} \left|\widehat{V}_{n}(f) - V(f)\right|.$$
(12)

The above lemma implies that  $\max_{f \in \mathcal{H}} |\hat{V}_m(f) - V(f)|$  and  $\max_{f \in \mathcal{H}} |\hat{V}_n(f) - V(f)|$  provide upper bounds for two deviations:

- the suboptimality of  $f_m^{\star}$  learned using the training set and the class  $\mathcal{H}$ , that is, how large  $V_g V(f_m^{\star})$  is;
- the estimation error  $|\hat{V}_n(f_m^{\star}) V(f_m^{\star})|$  due to the use of a validation set instead of the true expected gain  $V(f_m^{\star})$ , for the selected function  $f_m^{\star}$ .

Next let us remind the reader of the fact that  $g : \mathcal{W} \times \mathcal{X} \rightarrow [a, b]$  for some finite real values  $a \ge 0$  and b > a, and  $\mathcal{X}$  and  $\mathcal{W}$  are finite sets. This is important as it allows to probabilistically delimit the bounds of (11) and (12) in terms of b - a, besides the size n of the validation set  $\mathcal{T}_n$ , as expressed by the following proposition (see Appendix B.2 for the proof).

PROPOSITION 3.2 (UNIFORM DEVIATIONS). Assume that  $|\mathcal{H}| < \infty$ and  $g : \mathcal{W} \times \mathcal{X} \longrightarrow [a, b]$ , for a, b real values such that  $a \ge 0$  and b > a. Then, we have for all  $\varepsilon > 0$ ,

$$\sup_{P_{XY}} \mathbb{P}\left( \left| \hat{V}_n(f_m^{\star}) - V(f_m^{\star}) \right| > \varepsilon \right) \le 2 \exp\left( -\frac{2n\varepsilon^2}{(b-a)^2} \right)$$
(13)

<sup>&</sup>lt;sup>3</sup>We prefer to call  $\mathcal{T}_n$  validation set rather than test set, since we use it to estimate the g-vulnerability with the learned  $f_m^{\star}$ , rather than to measure the error in estimating the g-vulnerability.

and

$$\sup_{P_{XY}} \mathbb{P}\left(V_g - V(f_m^{\star}) > \varepsilon\right) \leq 2|\mathcal{H}| \exp\left(-\frac{m\varepsilon^2}{2(b-a)^2}\right). \quad (14)$$

Expression (13) shows that the estimation error due to the use of a validation set in  $\hat{V}_n(f_m^{\star})$  instead of the true expected gain  $V(f_m^{\star})$ vanishes with the number of validation samples. On the other hand, expression (14) implies 'learnability' of an optimal f, i.e., the suboptimality of  $f_m^{\star}$  learned using the training set  $V_m(f_m^{\star})$  vanishes with the number of training samples. Both expressions toguether imply that the *q*-vulnerability is learnable for all distributions (data sets) via the ERM principle introduced in (7). In other words, whenever the bounds indicate that we are close to the optimum f, we must at the same time have a good estimate of the *q*-vulnerability, and vice versa. Although the bound in (14) is in perfect agreement with the long-standing results from statistical learning theory which show that learnability is equivalent to uniform convergence of the empirical risk (8) to the actual risk (5), this distribution-free bound is rather pessimistic and cannot be expected to predict the performance in practical scenarios.

We can now state the main result of this section, namely an upper bound on the average estimation error of *g*-vulnerability (the proof is relegated to Appendix B.3):

THEOREM 3.3. The averaged estimation error of the *g*-vulnerability can be bounded as follows:

$$\mathbb{E}|V_g - \hat{V}_n(f_m^{\star})| \leq V_g - \mathbb{E}[V(f_m^{\star})] + \mathbb{E}|V(f_m^{\star}) - \hat{V}_n(f_m^{\star})|,$$

where the expectations are understood over all possible training and validation sets drawn according to  $P_{XY}$ . Furthermore,

$$V_{g} - \mathbb{E}\left[V(f_{m}^{\star})\right] \leqslant \sqrt{\frac{2(b-a)^{2}}{m}} \left(\sqrt{\ln|\mathcal{H}|} + \frac{1}{\sqrt{\ln|\mathcal{H}|}}\right), \quad (15)$$
$$\mathbb{E}\left|V(f_{m}^{\star}) - \hat{V}_{n}(f_{m}^{\star})\right| \leqslant \sqrt{\frac{(b-a)^{2}}{n}}, \quad (16)$$

independently of the specific underlying distribution  $P_{XY}$ .

Interestingly, the term corresponding to (15) is the error induced when estimating the function  $f_m^*$  using *m* samples from the training set while (16) indicates the error incurred when estimating the *g*-vulnerability using *n* samples from the validation set. Clearly, the scaling of these bounds with the number of samples are very different which can be made evident by using the order notation:

$$\sup_{P_{XY}} \left\{ V_g - \mathbb{E} \left[ V(f_m^{\star}) \right] \right\} \equiv O\left( \sqrt{\frac{|\mathcal{Y}| \ln |\mathcal{W}|}{m}} \right), \quad (17)$$

$$\sup_{P_{XY}} \mathbb{E} |V(f_m^{\star}) - \hat{V}_n(f_m^{\star})| \equiv O\left(\frac{1}{\sqrt{n}}\right).$$
(18)

These distribution-free bounds indicate that the error in (18) vanishes much faster than the error in (17) and thus, the size of the training set, in general, should be kept larger than the size of the validation set, i.e., n < m. However, the bound in (17) is rather pessimistic since it suffers from being independent of the underlying distribution and the optimization method used to solve (7). Tighter bounds can be derived but they would require statistical knowledge of the data-generating distribution which is often not available in real-world scenarios.

#### 3.4 Sample complexity

We now study how large the validation set should be in order to get a good estimation. For  $\varepsilon$ ,  $\delta > 0$ , we define the sample complexity as the set of smallest integers  $M(\varepsilon, \delta)$  and  $N(\varepsilon, \delta)$  sufficient to guarantee that the gap between the true *g*-vulnerability and the estimated  $\hat{V}_n(f_m^*)$  is at most  $\varepsilon$  with at least  $1 - \delta$  probability:

Definition 3.4. For  $\varepsilon$ ,  $\delta > 0$ , let all pairs  $(M(\varepsilon, \delta), N(\varepsilon, \delta))$  be the set of smallest (m, n) sizes of training and validation sets such that:

$$\sup_{P_{XY}} \mathbb{P}\left[ |V_g - \hat{V}_n(f_m^{\star})| > \varepsilon \right] \le \delta.$$
(19)

Next result says that we can bound the sample complexity in terms of  $\varepsilon$ ,  $\delta$ , and |b - a| (see Appendix B.4 for the proof).

COROLLARY 3.5. The sample complexity of the ERM algorithm *g*-vulnerability is bounded from above by the set of values satisfying:

$$M(\varepsilon, \delta) \leq \frac{2(b-a)^2}{\varepsilon^2} \ln\left(\frac{2|\mathcal{H}|}{\delta - \Delta}\right),$$
 (20)

$$N(\varepsilon, \delta) \leq \frac{(b-a)^2}{2\varepsilon^2} \ln\left(\frac{2}{\Delta}\right),$$
 (21)

for all  $0 < \Delta < \delta$ .

The theoretical results of this section are very general and do not refer to any particular model or data distribution. In particular, it is important to emphasize that the upper bounds in (11) and (12) are independent of the learned function  $f_m^*$ , and thus they are independent of the specific algorithm and training sets in used to solve the optimization in (8). Furthermore, the f maximizing  $|V(f) - \hat{V}_n(f)|$  in those in-equations is not necessarily what the algorithm would choose. Hence the bounds given in Theorem 3.3 and Corollary 3.5 in general are not tight. However, these theoretical bounds provide a worst-case measure from which learnability holds for all data sets.

In the next section, we will propose an approach for selecting  $f_m^*$  and estimating  $V_g$ . The experiments in Section 5 suggest that our method usually estimates  $V_g$  much more accurately than what is indicated by Theorem 3.3.

# 4 FROM g-VULNERABILITY TO BAYES VULNERABILITY VIA PRE-PROCESSING

This is the core section of the paper, which describes our approach to select the  $f_m^*$  to estimate  $V_g$ .

In principle one could train a neural network to learn  $f_m^*$  by using  $-\hat{V}_m(f)$  as the loss function, and minimizing it over the *m* training samples (cfr. Equation 8). However, this would require  $\hat{V}_m(f)$  to be a differentiable function of the weights of the neural network, so that its gradient can be computed during the backpropagation. Now, the problem is that the *g* component of  $\hat{V}_m(f)$ is essentially a non-differentiable function, so it would need to be approximated by a suitable differentiable (surrogate) function, (e.g., as it is the case of the Bayes error via the cross-entropy). Finding an adequate differentiable function to replace each possible *g* may be a challenging task in practice. If this surrogate does not preserve the original dynamic of the gradient of *g* with respect to *f*, the learned *f* will be far from being optimal.

Algorith	<b>n 1:</b> Algorit	hm for data	pre-pro	cessing
----------	---------------------	-------------	---------	---------

Input:  $\mathcal{D}_m$ ; Output:  $\mathcal{D}'_{m'}$ ;

1.  $\mathcal{D}'_{m'} := \emptyset;$ 

2. For each x, y, let  $u_{xy}$  be the number of copies of (x, y) in  $\mathcal{D}_m$ ; 3. For each x, y, w, add  $u_{xy} \cdot g(w, x)$  copies of (w, y) to  $\mathcal{D}'_{m'}$ .

In order to circumvent this issue, we propose a different approach, which presents two main advantages:

- it reduces the problem of learning f<sup>\*</sup><sub>m</sub> to a standard classification problem, therefore it does not require a different loss function to be implemented for each adversarial scenario;
- (2) it can be implemented by using *any* universally consistent learning algorithm (i.e., any ML algorithm approximating the ideal Bayes classifier).

The reduction described in the above list (item 1) is based on the idea that, in the *g*-leakage framework, the adversary's goal is not to directly infer the actual secret *x*, but rather to select the optimal guess *w* about the secret. As a consequence, the training of the ML classifier to produce  $f_m^{\star}$  should not be done on pairs of type (x, y), but rather of type (w, y), expressing the fact that the best guess, in the particular run which produced *y*, is *w*. This shift from (x, y) to (w, y) is via a *pre-processing* and we propose two distinct and systematic ways to perform this transformation, called *data* and *channel pre-processing*, respectively. The two methods are illustrated in the following sections.

We remind that, according to section 3, we restrict, wlog, to nonnegative g's. If g takes negative values, then it can be shifted by adding  $-\min_{w,x} g(w, x)$ , without consequences for the g-leakage value (cfr. [2, 4]). Furthermore we assume that there exists at least a pair (x, w) such that  $\pi_x \cdot g(w, x) > 0$ . Otherwise  $V_g$  would be 0 and the problem of estimating it will be trivial.

#### 4.1 Data pre-processing

The data pre-processing technique is completely black-box in the sense that it does not need access to the channel. We only assume the availability of a set of pairs of type (x, y), sampled according to  $\pi \triangleright C$ , the input-output distribution of the channel. This set could be provided by a third party, for example. We divide the set in  $\mathcal{D}_m$  (training) and  $\mathcal{T}_n$  (validation), containing *m* and *n* pairs, respectively.

For the sake of simplicity, to describe this technique we assume that g takes only integer values, in addition to being non-negative. The construction for the general case is discussed in Appendix C.3.

The idea behind the data pre-processing technique is that the effect of the gain function can be represented in the transformed dataset by amplifying the impact of the guesses in proportion to their reward. For example, consider a pair (x, y) in  $\mathcal{D}_m$ , and assume that the reward for the guess w is g(w, x) = 5, while for another guess w' is g(w', x) = 1. Then in the transformed dataset  $\mathcal{D}'_{m'}$  this pair will contribute with 5 copies of (w, y) and only 1 copy of (w', y). The transformation is described in Algorithm 1. Note that in general it causes an expansion of the original dataset.

Estimation of  $V_g$ . Given  $\mathcal{D}_m$ , we construct the set  $\mathcal{D}'_{m'}$  of pairs (w, y) according to Algorithm 1. Then, we use  $\mathcal{D}'_{m'}$  to train a classifier  $f^*_{m'}$ , using an algorithm that approximates the ideal Bayes classifier. As proved below,  $f^*_{m'}$  gives the same mapping  $\mathcal{Y} \to \mathcal{W}$  as the optimal empirical rule  $f^*_m$  on  $\mathcal{D}_m$  (cfr. subsection 3.2). Finally, we use  $f^*_m$  and  $\mathcal{T}_n$  to compute the estimation of  $V_g(\pi, C)$  as in (7), with f replaced by  $f^*_m$ .

*Correctness.* Let us first introduce some notation. For each (w, y), define:

$$U(w,y) \stackrel{\text{def}}{=} \sum_{x} \pi_x \cdot C_{xy} \cdot g(w,x) , \qquad (22)$$

which represents the "ideal" proportion of copies of (w, y) that  $\mathcal{D}'_{m'}$ should contain (of course, such proportion is only approximated in  $\mathcal{D}'_{m'}$  since it is generated from  $\mathcal{D}_m$  rather than according to the true distribution  $\pi \triangleright C$ ). From U(w, y) we can now derive the ideal joint distribution on  $\mathcal{W} \times \mathcal{Y}$  and the marginal on  $\mathcal{W}$ :

$$P_{WY}(w,y) \stackrel{\text{def}}{=} \frac{U(w,y)}{\alpha}, \text{ where } \alpha \stackrel{\text{def}}{=} \sum_{y,w} U(w,y), (23)$$

(note that  $\alpha > 0$  because of the assumption on  $\pi$  and g),

$$\sigma_{w} \stackrel{\text{def}}{=} \sum_{y} P_{WY}(w, y). \tag{24}$$

The channel of the conditional probabilities of y given w is:

$$E_{wy} \stackrel{\text{def}}{=} \frac{P_{WY}(w, y)}{\sigma_w} . \tag{25}$$

Note that  $P_{WY} = \sigma \triangleright E$ . By construction, it is clear that the  $\mathcal{D}'_{m'}$  generated by Algorithm 1 could have been generated, with the same probability, by sampling  $\sigma \triangleright E$ . The following theorem, whose proof is in Appendix C.1, establishes that the *g*-vulnerability of  $\pi \triangleright C$  is equivalent to the Bayes vulnerability of  $\sigma \triangleright E$ , and therefore that it is correct to estimate  $f_m^*$  as an empirical Bayes classifier  $f_{m'}^*$  trained using  $\mathcal{D}'_{m'}$ .

THEOREM 4.1 (CORRECTNESS OF DATA PRE-PROCESSING). Given a prior  $\pi$ , a channel C, and a gain function g, we have

$$V_q(\pi, C) = \alpha \cdot V_{q_{\mathrm{id}}}(\sigma, E)$$
,

where  $\alpha$ ,  $\sigma$  and E are those defined in (23), (24) and (25), respectively, and  $g_{id}$  is the identity function (cfr. section 2), i.e., the gain function corresponding to the Bayesian adversary.

The  $\alpha$  in the above theorem is only a scale factor, hence it does not influence the selection of  $f_m^*$ . Note that an alternative way to estimate  $V_g(\pi, C)$  would be by computing the empirical Bayes error of  $f_{m'}^*$  (using  $V_{g_{id}}(\sigma, E)$ ) on a validation set  $\mathcal{T}_{n'}$  of type (w, y)generated from  $\mathcal{T}_n$  by the same transformation as from  $\mathcal{D}_m$  to  $\mathcal{D}'_{m'}$ . In this case the  $\alpha$  would be necessary for converting the estimation of  $V_{g_{id}}(\sigma, E)$  into the estimation of  $V_q(\pi, C)$ .

#### 4.2 Channel pre-processing

For this technique we assume *black-box access* to the system, meaning that, although its channel matrix *C* is unknown, we can execute the system while controlling each input, and collect the corresponding output.

The core idea behind this technique is to transform the input of C into entries of type w, and to ensure that the distribution on the w's reflects the corresponding rewards expressed by g.

More formally, let us define a distribution  $\tau$  on  $\mathcal W$  as follows:

$$\tau_{w} \stackrel{\text{def}}{=} \frac{\sum_{x} \pi_{x} \cdot g(w, x)}{\beta} \quad \text{where} \quad \beta \stackrel{\text{def}}{=} \sum_{x, w} \pi_{x} \cdot g(w, x) \,, \, (26)$$

(note that  $\beta$  is strictly positive because of the assumptions on *g* and  $\pi$ ), and let us define the following matrix *R* from W to *X*:

$$R_{wx} \stackrel{\text{def}}{=} \frac{1}{\beta} \cdot \frac{1}{\tau_w} \cdot \pi_x \cdot g(w, x) \,. \tag{27}$$

It is easy to check that R is a stochastic matrix, hence the composition RC is a channel. It is important to emphasize the following:

REMARK In the above definitions,  $\beta$ ,  $\tau$  and R depend solely on g and  $\pi$ , and not on C.

The above property is crucial to our goals, because in the blackbox approach we are not supposed to rely on the knowledge of *C*'s internals. We now illustrate how we can estimate  $V_g$  using the pre-processed channel *RC*.

*Estimation of*  $V_g$ . Given RC and  $\tau$ , we build a set  $\mathcal{D}''_{m''}$  consisting of pairs of type (w, y) sampled from  $\tau \triangleright RC$ . We also construct a set  $\mathcal{T}_n$  of pairs of type (x, y) sampled from  $\pi \triangleright C$ . Then, we use  $\mathcal{D}''_{m''}$  to train a classifier  $f_m^*$ , using an algorithm that approximates the ideal Bayes classifier. Finally, we use  $f_m^*$  and  $\mathcal{T}_n$  to compute the estimation of  $V_g(\pi, C)$  as in (7), with f replaced by  $f_m^*$ .

Alternatively, we could estimate  $V_g(\pi, C)$  by computing the empirical Bayes error of  $f_m^*$  on a validation set  $\mathcal{T}_n$  of type (w, y) sampled from  $\tau \triangleright RC$ , but the estimation would be less precise. Intuitively, this is because RC is more "noisy" than C.

*Correctness.* The correctness of the channel pre-processing method is given by the following theorem, which shows that we can learn  $f_m^*$  by training a Bayesian classifier on a set sampled from  $\tau \triangleright RC$ .

THEOREM 4.2 (CORRECTNESS OF CHANNEL PRE-PROCESSING). Given a prior  $\pi$  and a gain function g, we have that, for any channel C:

 $V_q(\pi, C) = \beta \cdot V_{q_{id}}(\tau, RC)$  for all channels C.

where  $\beta$ ,  $\tau$  and R are those defined in (26) and (27).

Interestingly, a result similar to Theorem 4.2 is also given in [6], although the context is completely different from ours: the focus of [6], indeed, is to study how the leakage of C on X may induce also a leakage of other sensitive information Z that has nothing to do with C (in the sense that is not information manipulated by C). We intend to explore this connection in the context of a possible extension of our approach to this more general scenario.

#### 4.3 **Pros and cons of the two methods**

The fundamental advantage of data pre-processing is that it allows to estimate  $V_g$  from just samples of the system, without even blackbox access. In contrast to channel pre-processing, however, this method is particularly sensitive to the values of the gain function g. Large gain values will increase the size of  $\mathcal{D}'_{m'}$ , with consequent increase of the computational cost for estimating the g-vulnerability. Moreover, if g takes real values then we need to apply the technique described in Appendix C.3, which can lead to a large increase of the dataset as well. In contrast, the channel pre-processing method has the advantage of controlling the size of the training set, but it can be applied only when it is possible to interact with the channel by providing input and collecting output. Finally, from the precision point of view, we expect the estimation based on data pre-processing to be more accurate when *g* consists of small integers, because the channel pre-processing introduces some extra noise in the channel.

#### **5** EVALUATION

In this section we evaluate our approach to the estimation of *g*-vulnerability. We consider four different scenarios:

- X is a set of (synthetic) numeric data, the channel C consists of geometric noise, and g is the multiple guesses gain function, representing an adversary that is allowed to make several attempts to discover the secret.
- (2) X is a set of locations from the Gowalla dataset [1], C is the optimal noise of Shokri et al. [32], and g is one of the functions used to evaluate the privacy loss in [32], namely a function anti-monotonic on the distance, representing the idea that the more the adversary's guess is close to the target (i.e., the real location), the more he gains.
- (3) X is the Cleveland heart disease dataset [22], C is a differentially private mechanism [23, 24], and g is a function that assigns higher values to worse heart conditions, modeling an adversary that aims at discovering whether a patient is at risk (for instance, to deny his application for health insurance).
- (4) X is a set of passwords of 128 bits and C is a password checker that leaks the time before the check fails, but mitigates the timing attacks by applying some random delay and the bucketing technique (see, for example, [28]). The function g represents the part of the password under attack.

For each scenario, we proceed in the following way:

- We consider 3 different samples sizes for the training sets that are used to train the ML models and learn the 𝒴 → 𝔐 remapping. This is to evaluate how the precision of the estimate depends on the amount of data available, and on its relation with the size of |𝒴|.
- In order to evaluate the variance of the precision, for each size we create 5 different training sets, and
- for each trained model we estimate the *g*-vulnerability using 50 different validation sets.

#### 5.1 Representation of the results and metrics

We graphically represent the results of the experiment as box plots, using one box for each size. More precisely, given a specific size, let  $\hat{V}_n^{ij}$  be the *g*-vulnerability estimation on the *j*-th validation set computed with a model trained over the *i*-th training set (where  $i \in \{1, \ldots, 5\}$  and  $j \in \{1, \ldots, 50\}$ ). Let  $V_g$  be the real *g*-vulnerability of the system. We define the *normalized estimation error*  $\delta_{ij}$  and the mean value  $\overline{\delta}$  of the  $\delta_{ij}$ 's as follows:

$$\delta_{ij} \stackrel{\text{def}}{=} \frac{|\widehat{V}_n^{ij} - V_g|}{V_g}, \quad \text{with} \quad \overline{\delta} \stackrel{\text{def}}{=} \frac{1}{250} \sum_{i=1}^5 \sum_{j=1}^{50} \delta_{ij}.$$
(28)

In the graphs, the  $\delta_{ij}$ 's are the values reported in the box corresponding to the given size, and  $\overline{\delta}$  is the black horizontal line inside the box.

We also consider the following quantities, which are typical measures of precision:

dispersion 
$$\stackrel{\text{def}}{=} \sqrt{\frac{1}{250} \sum_{i=1}^{5} \sum_{j=1}^{50} (\delta_{ij} - \overline{\delta})^2},$$
 (29)

total error 
$$\stackrel{\text{def}}{=} \sqrt{\frac{1}{250} \sum_{i=1}^{5} \sum_{j=1}^{50} \delta_{ij}^2}.$$
 (30)

The dispersion is an average measure of how far the normalized estimation errors are from their mean value when using same-size training and validation sets. On the other hand, the total error is an average measure of the normalized estimation error, when using same-size training and validation sets. In our experiments we will see that, as expected, the dispersion and the total error tend to decrease when the amount of training samples increases.

In order to make a fair comparison between the two presented pre-processing methods, intuitively we need to train them on training sets of the same size, and evaluate them on validation sets of the same size. For the validation part, since the best f function has been already found and therefore we do not need any data pre-processing, this is easy to guarantee. But what does "same size" mean, in the case of training sets built with different pre-processing methods? Consider the following situation: assume that we have a set of data  $\mathcal{D}_m$  coming from a third party collector (we recall that the index *m* represents the size of the set), and let  $\mathcal{D}'_{m'}$  be the result of the data pre-processing on  $\mathcal{D}_m$ . Now, let  $\mathcal{D}''_{m''}$  be the dataset obtained drawing samples according to the channel pre-processing method. Should we impose m'' = m or m'' = m'? We argue that the right choice is the first one, because the main limiting constraint in our method is the amount of "real" data that we can collect, one way or another. In case we cannot interact with the channel, the number of available samples is controlled by the data provider. Indeed,  $\mathcal{D}'_{m'}$  is generated synthetically from  $\mathcal{D}_m$  and cannot contain more information about *C* than  $\mathcal{D}_m$ , despite its larger size.

The nice feature of the normalized estimation error is that, thanks to the normalization, it allows to compare the results among different scenario and different levels of (real) *g*-vulnerability. Also, the percentage of the error is more meaningful than the absolute value. However, the interested reader can find in Appendix F also the plots that show the values of the estimations of the *g*-vulnerability and their distance from the corresponding real *g*-vulnerability.

#### 5.2 Learning algorithms

We consider two ML algorithms in the experiments: the k-Nearest Neighbors (k-NN) and the Artificial Neural Networks (ANN). We have made however a slight modification of k-NN algorithm, due to the following reason: recall that, depending on the particular gain function, the data pre-processing method might create many instances where a certain observable *y* is repeated multiple times in pair with different *w*'s. For the k-NN algorithm, a very common choice is to consider a number of neighbors which is equivalent to natural logarithm of the total number of training samples. In particular, when the data pre-processing is applied, this means that  $k = \log(m')$  nearest neighbors will be considered for the classification decision. Since  $\log(m')$  grows slowly with respect to m', it might happen that k-NN fails to find the subset of neighbors from which the best remapping can be learned. To amend this problem, we modify the k-NN algorithm in the following way: instead of looking for neighbors among all the m' samples, we only consider a subset of  $l \leq m'$  samples, where each value y only appears once. After the  $\log(l)$  neighbors have been detected among the l samples, we select w according to a majority vote over the m'tuples (w, y) created through the remapping.

The distance on which the notion of neighbor is based depends on the experiments. We have considered the standard distance among numbers in the first and fourth experiments, the Euclidean distance in the second one, and the Manhattan distance between tuples of numbers in the third one (where the distance between the components is the standard numerical distance).

Concerning the ANN models, their specifics are in Appendix D. Note that, for the sake of fairness, we use the same architecture for both pre-processing methods, although we adapt number of epochs and batch size to the particular dataset we are dealing with.

Further details relative to the specific experiments are provided the following sections, each of them discussing one of the above mentioned scenarios.

#### 5.3 Frequentist approach

In the experiments, we will compare our method with the frequentist one. This approach has been proposed originally in [9] for estimating mutual information, and extended successively also to min-entropy leakage [17]. Although not considered in the literature, the extension to the case of g-vulnerability is straightforward. The method consists in estimating the probabilities that constitute the channel matrix C, and then calculating analytically the g-vulnerability on C. The precise definition is in Appendix E.

In [14] it was observed that, at least in the case of the Bayes vulnerability<sup>4</sup>, the frequentist approach performs poorly when the size of the observable domain  $|\mathcal{Y}|$  is large with respect to the available data. We want to examine whether this is the case also for other vulnerabilities.

For the experiment on the multiple guesses the comparison is illustrated in the next section. For the other experiments, because of lack of space, we have reported it in the Appendix F.

#### 5.4 Experiment 1: multiple guesses

We consider a system in which the secrets X are the integers between 0 and 9, and the observables  $\mathcal{Y}$  are the integers between 0 and 15999. Hence |X| = 10 and  $|\mathcal{Y}| = 16$ K. The channel *C* adds noise to the elements of X according to the following geometric distribution:

$$C_{xy} = P_{Y|X}(y|x) = \lambda \exp(-\nu |r(x) - y|),$$
 (31)

where:

• v is a parameter that determines how concentrated around y = x the distribution is. In this experiment we set v = 0.002;

<sup>&</sup>lt;sup>4</sup>Strictly speaking, [14] considered the estimation of the Bayes error, but the essence does not change since Bayes vulnerability and Bayes error are complementary w.r.t. 1.



Figure 1: The channel of Experiment 1. The two curves represent the distributions  $P_{Y|X}(\cdot|x)$  for two adjacent secrets: x = 5 and x = 6.

- *r* is just an auxiliary function that reports X to the same scale of *Y*, and centers X on *Y*. Here we have *r*(*x*) = 1000 *x* + 3499.5;
- λ = e<sup>ν</sup> −1/(e<sup>ν</sup>+1) is a normalization factor tuned so that Equation 31 is indeed a distribution.

Figure 1 illustrates the shape of  $C_{xy}$ . More precisely, it shows the distributions  $P_{Y|X}(\cdot|x)$  for two adjacent secrets x = 5 and x = 6. We consider an adversary that can make two attempts to discover the secret (two-tries adversary), and we define the corresponding gain function as follows. A guess  $w \in W$  is one of all the possible combinations of 2 different secrets from X, i.e.,  $w = \{x_0, x_1\}$  with  $x_0, x_1 \in X$  and  $x_0 \neq x_1$ . Therefore  $|W| = \binom{10}{2} = 45$ . The gain function g is then

$$g(w, x) = \begin{cases} 1 & \text{if } x \in w \\ 0 & \text{otherwise} \end{cases}$$
(32)

For this experiment we consider a uniform prior distribution  $\pi$  on X. The true *g*-vulnerability for these particular v and  $\pi$ , results to be  $V_g = 0.892$ . For all the following experiments in the multiple guesses scenario, the reported results are obtained evaluating the normalized estimation error on 50 different validation sets of 50K samples of type (x, y) drawn according to  $\pi$  and to the channel distribution in Equation 31. The sample sizes for the training sets (of the same type as  $\mathcal{D}_m$ ) that we consider are 10K, 30K and 50K respectively. As explained before, for each size we use 5 different training sets where the samples are couples (y, x).

5.4.1 Data pre-processing. In this part of the experiment we trasform the training data according to the data pre-processing described in Section 4.1. The result are sets of samples of type (w, y) of the same type as  $\mathcal{D}'_{m'}$ . Then, we use the latter to learn the  $\mathcal{Y} \to \mathcal{W}$ remapping, and we do so by using k-NN and ANN classifiers. The plot in Figure 3 shows the performances of the k-NN and ANN models on the validation sets used to estimate the *g*-vulnerability in terms of normalized estimation error, while Figure 2 shows the same performances compared to those of the frequentist approach. As we can see, the precision of the frequentist is much lower, thus confirming that the trend observed by [14] for the Bayes vulnerability holds also for other gain functions. Intuitively, this gap occurs especially when  $|\mathcal{Y}|$  is high with respect to the number of training samples, and it is due to the fact that with the frequentist approach there is no real learning, so we cannot make a good guess with the observables never seen before. In ML on the contrary we can



Figure 2: Multiple guesses scenario, comparison between the frequentist and the ML estimations with data pre-processing.



Figure 3: Multiple guesses scenario, magnification of the part of Figure 2 on the k-NN and ANN estimations.

still make an informed guess, especially when the channel has a rather regular behavior, i.e., the noise expressed by the channel is "smooth" (cfr. [14]). It is worth noting that, in this experiment, the pre-processing of each sample (x, y) creates 9 samples (matching y with each possible  $w \in W$  such that  $w = \{x, x'\}$  with  $x' \neq x$ ). This means that the sample size of the pre-processed sets is 9 times the size of the original ones. For functions g representing more than 2 tries this pre-processing method may create training sets too large. In the next section we consider the alternative pre-processing method, showing that it can be a good compromise.

5.4.2 Channel pre-processing. Let us now suppose it is possible to interact with the channel modified according to the channel pre-processing method (cfr. Section 4.2) so that we are able to sample data of the type (w, y) from it. With these samples we form training sets (of the same type as  $\mathcal{D}''_{m''}$ ) of size 10K, 30K, and 50K. Then we proceed with the learning and the *g*-vulnerability estimation as before. The results are showed in Figure 4. As we can see, the results are worse than in the data pre-processing case, especially for the k-NN algorithm. This was to be expected, since the random sampling to match the effect of *g* introduces a further level of confusion, as explained in Section 4.2. Nevertheless, these results are still much better than the frequentist case, so it is a good



Figure 4: Multiple guesses scenario, k-NN and ANN estimation with channel pre-processing



Figure 5: Heat-map representing the Gowalla check-ins distribution in the area of interest; the density of check-ins in each cell is reported in the color bar on the side

alternative method to apply when the use of data pre-processing would generate validation sets that are too large, which could be the case when the matrix representing g contains large numbers with a small common divider. Additional plots can be found in Appendix F.

#### 5.5 Experiment 2: location privacy

In this section we estimate the *g*-vulnerability of a typical system for location privacy protection. We use data from the open Gowalla dataset [1], which contains the coordinates of users' check-ins. In particular, we consider a square region in San Francisco, USA, centered in (latitude, longitude) = (37.755, -122.440), and with 5Km long sides. In this area Gowalla contains 35162 check-ins.

We discretize the region in 400 cells of 250m long side, and we assume that the adversary's goal is to discover the cell in which a check-in is located. The frequency of the Gowalla check-ins per cell is represented by the heat-map in Figure 5. From these frequencies we can directly derive the distribution representing the prior of the secrets.

The channel C that we consider here is the optimal obfuscation mechanism proposed in [32] to protect location privacy under a utility constraint. We recall that the framework of [32] assumes two loss functions, one for utility and one for privacy. The utility

	0	0	0	0	0	0	0
T	0	0	0	1	0	0	0
	0	0	1	2	1	0	0
	0	1	2	4	2	1	0
	0	0	1	2	1	0	0
	0	0	0	1	0	0	0
	0	0	0	0	0	0	0

Figure 6: The "diamond" shape created by the gain function around the real secret; the values represent the gains assigned to each guessed cell w when x is the central cell.

loss of a mechanism, for a certain prior, is defined as the expected utility loss of the noisy data generated according to the prior and the mechanism. The privacy loss is defined in a similar way, except that we allow the attacker to "remap" the noisy data so to maximize the privacy loss. For our experiment, we use the Euclidean distance as loss function for the utility, and the g function defined in the next paragraph as loss function for the privacy. For further details on the construction of the optimal mechanism we refer to [32].

We define  $X, \mathcal{Y}$  and  $\mathcal{W}$  to be the set of the cells. Hence  $|X| = |\mathcal{Y}| = |\mathcal{W}| = 400$ . We consider a gain function that represents the precision of the guess in terms of euclidean distance: the idea is that the smaller is the distance between the real cell x and the guessed cell w, the higher is the gain. Specifically, our g is illustrated in Figure 6, where the central cell represents the real location x. For a generic "guess" cell w, the number written in w represent g(w, x). Thus for example we have g(x, x) = 4, and g(w, x) = 2 if w is an immediate neighbor of x.<sup>5</sup>

In this experiment we consider training set sizes of 100, 1k and 10K samples respectively. After applying the data pre-processing transformation, the size of the resulting datasets is approximately 18 times that of the original one. This was to be expected, since the sum of the values of g in Figure 6 is 20. Note that this sum and the increase factor in the dataset do not necessarily coincide, because the latter is also influenced by the prior and by the mechanism.

Figure 7 and Figure 8 show the performance of k-NN and the ANN for the application of both data pre-processing and the channel pre-processing methods. As expected, the data pre-processing method is more precise than the channel pre-processing one, although only slightly. The ANN model is also slightly better than the k-NN in most of the cases. For this experiment, as one can see in the plots in Appendix F, the frequentist approach outperforms the ML methods. Indeed this can be explained as follows: the observable space is not very large, which is a scenario where the frequentist approach can be successful because the available data is enough to estimate the real distribution. Indeed the ANN method performs quite well too, but the model we used is quite complex for this problem and therefore, we obtain a slightly worse performance than the frequentist approach. This is indeed an example which shows that the frequentist approach can still work well in certain scenarios. Although, as experimentally verified in all the other experiments, ML techniques, and in particular ANN,

<sup>&</sup>lt;sup>5</sup>Formally, *g* is defined as  $g(w, x) = \lfloor (\gamma \exp(-\alpha d(w, x)/l)) \rfloor$ , where  $\gamma = 4$  is the maximal gain,  $\alpha = 0.95$  is a normalization coefficient to control the skewness of the exponential, *d* is the euclidean distance and l = 250 is the length of the cells' side. The symbol  $\lfloor \cdot \rfloor$  in this context represents the rounding to the closest integer operation.



Figure 7: Location privacy scenario, k-NN and ANN estimation with data pre-processing



Figure 8: Location privacy scenario, k-NN and ANN estimation with channel pre-processing

easily outperforms the frequentist approach when large observable spaces are involved.

#### 5.6 Experiment 3: differential privacy

In this section we consider a popular application of DP: individual data protection in medical datasets from which we wish to extract some statistics via counting queries.

It is well known that the release of exact information from the database, even if it is only the result of statistical computation on the aggregated data, can leak sensitive information about the individuals. The solution proposed by DP is to obfuscate the released information with carefully crafted noise that obeys certain properties. The goal is to make it difficult to detect whether a certain individual is in the database or not. In other words, two adjacent datasets (i.e., datasets that differ only for the presence of one individual) should have almost the same likelihood to produce a certain observable result.

In our experiment, we consider the Cleveland heart disease dataset [22] which consist of 303 records of patients with a medical heart condition. Each condition is labeled by an integer number (label) that indicates how serious the disease is: from 0, which represents a healthy patient, to 4, which represents a patient whose life is at risk.



Figure 9: Differential privacy scenario, k-NN and ANN estimation with data pre-processing

We assume that, to help medical research, the hospital releases the histogram of these labels, i.e., the counts of their occurrences in the database. We also assume that, for the sake of protecting their patients' privacy, the hospital sanitizes the histogram by adding geometric noise, which is a typical DP mechanism, to each label's count. More precisely, if the count of a label is  $z_1$ , the probability that the corresponding published number is  $z_2$  is defined by the distribution in Eq. 31, where x and y are replaced by  $z_1$  and  $z_2$ respectively, and r is 1. Note that  $z_1$  is the real count, so it is an integer between 0 and 303, while its noisy version  $z_2$  ranges on all integers. Concerning the value of v, in this experiment it is set to 1.

The secrets space X is set to be a set of two elements: the full dataset, and the dataset with one record less. These are adjacent in the sense of DP, and, as customary in DP, we assume that the record on which the two databases differ is the target of the adversary. The observables space Y is the set of the 5-tuples produces by the noisy counts of the 5 labels. W is set to be the same as X.

We assume that the adversary is interested especially in finding out whether the patient has a serious condition. The function greflects this preference by assigning higher value to higher labels. Specifically, we set:

$$g(w,x) = \begin{cases} 0, & \text{if } w \neq x \\ 1, & \text{if } w = x \land x \in \{0,1,2\} \\ 2, & \text{if } w = x \land x \in \{3,4\}, \end{cases}$$
(33)

and  $\mathcal{W} = \mathcal{X}$ .

For the estimation, we consider 50 different validation sets of 50K samples each drawing them from the channel. The results reported below represent the estimation performance of each model on these validation sets. As training set sizes, we consider 10K, 30K and 50K samples, that we use to learn the remapping from the obfuscated counts,  $\mathcal{Y}$ , to the guesses about the diseases  $\mathcal{W}$ . For each of these sizes we consider 5 training sets. Again, as ML algorithms we use ANN and k-NN. For the latter we have to decide what kind of distance we use for evaluating the proximity of the elements of  $\mathcal{Y}$ . We choose the Manhattan distance, which seems the most natural in



Figure 10: Differential privacy scenario, k-NN and ANN estimation with channel pre-processing

this case<sup>6</sup> In the case of data pre-processing the size of the transformed training sets (of the same type as  $\mathcal{D}'_{m'}$ ) is about 1.2 times the size of the original training sets (of the same type as  $\mathcal{D}_m$ ). The performance is shown in Figure 9. For the channel pre-processing approach, the performance is shown in Figure 10. Surprisingly, in this case the data pre-processing method outperforms the channel pre-processing one, although only slightly. Additional plots, including the results for the frequentist approach, can be found in Appendix F.

#### 5.7 Experiment 4: password checker

In this experiment we consider a password checker, namely a program that tests whether a given string corresponds to the password stored in the system. We assume that string and password are sequences of 128 bits, an that the program is "leaky", in the sense that it checks the two sequences bit by bit and it stops checking as soon as it finds a mismatch, reporting failure. It is well known that this opens the way to a timing attack (a kind of side-channel attack), so we assume that the system tries to mitigate the threat by adding some random delay, sampled from a Laplace distribution and then bucketing the reported time in 128 bins corresponding to the positions in the sequence (or equivalently, by sampling the delay from a Geometric distribution, cfr. Equation 31). Hence the channel *C* is a  $2^{128} \times 128$  stochastic matrix.

The typical attacker is an interactive one, which figures out larger and larger prefixes of the password by testing each bit at a time. We assume that the attacker has already figured out the first 6 bits of the sequence and it is trying to figure out the 7-th. Thus the prior  $\pi$  is distributed (uniformly, we assume) only on the sequences formed by the known 6-bits prefix and all the possible remaining 122 bits, while the *g* function assigns 1 to the sequences whose 7-th bit agrees with the stored password, and 0 otherwise. Thus *g* is a *partition gain function* [4], and its particularity is that for such kind of functions data pre-processing and channel preprocessing coincide. This is because g(w, x) is either 0 or 1, so in both cases we generate exactly one pair (w, y) for each pair (x, y)for which g(w, x) = 1. Note that in this case the data pre-processing



Figure 11: Password checker scenario, k-NN and ANN estimation with data and channel pre-processing

transformation does not increase the training set, and the channel pre-processing transformation does not introduce any additional noise. The *RC* matrix (cfr. Section 4.1) is a 2 × 128 stochastic matrix. The experiments are done with training sets of 10K, 30K and 50K samples. The results are reported in Figure 11. We note that the estimation error is quite small, especially in the ANN case. This is because the learning problem is particularly simple since, by considering the *g*-leakage and the preprocessing, we have managed to reduce the problem to learning a function of type  $\mathcal{Y} \to \mathcal{W}$ , rather than  $\mathcal{Y} \to \mathcal{X}$ , and there is a huge difference in size between  $\mathcal{W}$  and  $\mathcal{X}$  (the first is 2 and the latter is  $2^{128}$ ). Also the frequentist approach does quite well (cfr. Appendix F), and this is because  $\mathcal{Y}$  is small. With a finer bucketing (on top of the Laplace delay), or no bucketing at all, we expect that the difference between the accuracy of the frequentist and of the ML estimation would be much larger.

Note that with a non-leaky password checker the observables are only *fail* or *success*. In this case the size of  $\mathcal{Y}$  would be 2, but since *success* would have an extremely small probability (1/2<sup>128</sup>), the vulnerability would be negligible, and it would not be detected neither by our approach nor by the frequentist one, because a pair ( $\cdot$ , *success*) would never be generated in practice. Hence both approaches would report vulnerability 0.

#### 6 CONCLUSION AND FUTURE WORK

In this paper we have proposed an approach to estimate the g-vulnerability of a system under the black-box assumption, using machine learning. The basic idea is that the problem can be reduced to learn the Bayes classifier on a set of pre-processed training data, and we have proposed two techniques for this transformation, with different advantages and disadvantages. We have then evaluated our approach on four different privacy scenarios, showing favorable results. We have considered the frequentist approach and compared it with ours experimentally, showing that the results are comparable when the observable domain is small, while our approach does much better on large  $\mathcal{Y}$  domains. This is in line with what already observed in [14] for the estimation of the Bayes error.

As future work, we plan to test our framework on more real-life scenarios such as the web fingerprinting attacks [13, 15] and the AES cryptographic algorithm [20]. We also would like to consider

<sup>&</sup>lt;sup>6</sup>The Manhattan distance on histograms corresponds to the total variation distance on the distributions resulting from the normalization of these histograms.

the more general case, often considered in Information-flow security, of channels that have both "high" and "low" inputs, where the first are the secrets and the latter are data visible to, or even controlled by, the adversary. Finally, a more ambitious goal is to use our approach to minimize the *g*-vulnerability of complex systems, using a GAN based approach, along the lines of [29].

#### REFERENCES

- [1] 2011. Gowalla dataset. https://snap.stanford.edu/data/loc-Gowalla.html. (2011).
- [2] Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. 2014. Additive and Multiplicative Notions of Leakage, and Their Capacities. In IEEE 27th Computer Security Foundations Symposium, CSF 2014, Vienna, Austria, 19-22 July, 2014. IEEE, 308–322. https://doi.org/10.1109/CSF.2014.29
- [3] Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. 2016. Axioms for Information Leakage. In Proceedings of the 29th IEEE Computer Security Foundations Symposium (CSF). 77–92. https://doi.org/10.1109/CSF.2016.13
- [4] Mário S. Álvim, Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. 2012. Measuring Information Leakage Using Generalized Gain Functions. In 25th IEEE Computer Security Foundations Symposium, CSF 2012, Cambridge, MA, USA, June 25-27, 2012, Stephen Chong (Ed.). IEEE Computer Society, 265–279. http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=6265867
- [5] Christopher M. Bishop. 2007. Pattern recognition and machine learning, 5th Edition. Springer. I–XX, 1–738 pages. http://www.worldcat.org/oclc/71008143
- [6] Nicolás E. Bordenabe and Geoffrey Smith. 2016. Correlated Secrets in Quantitative Information Flow. In CSF. IEEE Computer Society, 93–104. http://ieeexplore.ieee. org/xpl/mostRecentIssue.jsp?punumber=7518122
- [7] Michele Boreale. 2009. Quantifying information leakage in process calculi. Inf. Comput 207, 6 (2009), 699–725.
- [8] S. Boucheron, G. Lugosi, and P. Massart. 2013. Concentration Inequalities: A Nonasymptotic Theory of Independence. OUP Oxford.
- [9] Konstantinos Chatzikokolakis, Tom Chothia, and Apratim Guha. 2010. Statistical Measurement of Information Leakage. In TACAS (Lecture Notes in Computer Science), Vol. 6015. Springer, 390–404.
- [10] Konstantinos Chatzikokolakis, Natasha Fernandes, and Catuscia Palamidessi. 2019. Comparing systems: max-case refinement orders and application to differential privacy. In Proceedings of the 32nd IEEE Computer Security Foundations Symposium. Hoboken, United States, 442–457. https://doi.org/10.1109/CSF.2019.00037
- [11] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. 2008. Anonymity protocols as noisy channels. *Inf. Comput* 206, 2-4 (2008), 378–401.
- [12] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. 2008. On the Bayes risk in information-hiding protocols. *Journal of Computer Security* 16, 5 (2008), 531–571. https://doi.org/10.3233/JCS-2008-0333
- [13] Giovanni Cherubin. 2017. Bayes, not Naïve: Security Bounds on Website Fingerprinting Defenses. PoPETs 2017, 4 (2017), 215–231.
- [14] Giovanni Cherubin, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2019. F-BLEAU: Fast Black-box Leakage Estimation. *IEEE Symposium on Security* and Privacy abs/1902.01350 (2019). http://arxiv.org/abs/1902.01350
- [15] Giovanni Cherubin, Jamie Hayes, and Marc Juárez. 2017. Website Fingerprinting Defenses at the Application Layer. *PoPETs* 2017, 2 (2017), 186–203. https://doi. org/10.1515/popets-2017-0023
- [16] Tom Chothia and Apratim Guha. 2011. A Statistical Test for Information Leaks Using Continuous Mutual Information. In CSF. IEEE Computer Society, 177–190. http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=5991608
- [17] Tom Chothia, Yusuke Kawamoto, and Chris Novakovic. 2013. A tool for estimating information leakage. In *International Conference on Computer Aided Verification (CAV)*. Springer, 690–695.
- [18] Tom Chothia, Yusuke Kawamoto, and Chris Novakovic. 2014. LeakWatch: Estimating Information Leakage from Java Programs. In ESORICS (2) (Lecture Notes in Computer Science), Miroslaw Kutylowski and Jaideep Vaidya (Eds.), Vol. 8713. Springer, 219–236.
- [19] David Clark, Sebastian Hunt, and Pasquale Malacaria. 2001. Quantitative Analysis of the Leakage of Confidential Data. *Electr. Notes Theor. Comput. Sci* 59, 3 (2001), 238–251.
- [20] Eloi de Chérisey, Sylvain Guilley, Olivier Rioul, and Pablo Piantanida. 2019. Best Information is Most Successful - Mutual Information and Success Rate in Side-Channel Analysis. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2019, 2 (2019), 49–79. https://doi.org/10.13154/tches.v2019.i2.49-79
- [21] Luc Devroye, László Györfi, and Gábor Lugosi. 1996. Vapnik-Chervonenkis Theory. Springer New York, New York, NY, 187–213. https://doi.org/10.1007/978-1-4612-0711-5 12
- [22] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository (Heart Disease Data Set). (2017). https://archive.ics.uci.edu/ml/datasets/heart+Disease

- [23] Cynthia Dwork. 2006. Differential Privacy. In 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006) (Lecture Notes in Computer Science), Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.), Vol. 4052. Springer, 1–12. http://dx.doi.org/10.1007/11787006\_1
- [24] Cynthia Dwork, Frank Mcsherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In In Proceedings of the Third Theory of Cryptography Conference (TCC) (Lecture Notes in Computer Science), Shai Halevi and Tal Rabin (Eds.), Vol. 3876. Springer, 265–284.
- [25] Ian J. Goodfellow, Yoshua Bengio, and Aaron C. Courville. 2016. Deep Learning. MIT Press. 1–775 pages. http://www.deeplearningbook.org/
- [26] T. Hastie, R. Tibshirani, and J. Friedman. 2001. The Elements of Statistical Learning: Data Mining, Inference and Prediction. Springer-Verlag.
- [27] Boris Köpf and David A. Basin. 2007. An information-theoretic model for adaptive side-channel attacks. In Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007, Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson (Eds.). ACM, 286-296.
- [28] Boris Köpf and Markus Dürmuth. 2009. A Provably Secure and Efficient Countermeasure against Timing Attacks. In Proceedings of the 2009 22nd IEEE Computer Security Foundations Symposium (CSF '09). IEEE Computer Society, USA, 324–335. https://doi.org/10.1109/CSF.2009.21
- [29] Marco Romanelli, Catuscia Palamidessi, and Konstantinos Chatzikokolakis. 2020. Generating Optimal Privacy-Protection Mechanisms via Machine Learning, In Proceedings of the IEEE International Symposium on Computer Security Foundations (CSF). CoRR. arXiv:1904.01059 http://arXiv.org/abs/1904.01059
- [30] Shai Shalev-Shwartz and Shai Ben-David. 2014. Understanding machine learning : from theory to algorithms. http://www.worldcat.org/search?qt=worldcat\_org\_ all&q=9781107057135
- [31] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. 2011. Quantifying Location Privacy. In *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 247–262. http://ieeexplore.ieee.org/ xpl/mostRecentIssue.jsp?punumber=5955408; http://www.computer.org/csdl/ proceedings/sp/2011/4402/00/index.html
- [32] Reza Shokri, George Theodorakopoulos, Carmela Troncoso, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2012. Protecting location privacy: optimal strategy against localization attacks. In the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012, Ting Yu, George Danezis, and Virgil D. Gligor (Eds.). ACM, 617–627. http://dl.acm.org/citation. cfm?id=2382196
- [33] Geoffrey Smith. 2009. On the Foundations of Quantitative Information Flow. In Foundations of Software Science and Computational Structures, 12th International Conference, FOSSACS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings (Lecture Notes in Computer Science), Luca de Alfaro (Ed.), Vol. 5504. Springer, 288–302.
- [34] David H. Wolpert. 1996. The Lack of A Priori Distinctions Between Learning Algorithms. *Neural Computation* 8, 7 (1996), 1341–1390.

#### A AUXILIARY RESULTS

PROPOSITION A.1 (HOEFFDING'S INEQUALITY [8]). Let  $Z_1, \ldots, Z_n$ be independent bounded random variables such that  $Z_i \in [a_i, b_i]$ almost surely and let  $S_n = \sum_{i=1}^n Z_i$ . Then, for any t > 0, we have the following inequalities:

$$\mathbb{P}(S_n - \mathbb{E}[S_n] \ge t) \le \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right), \quad (34)$$

$$\mathbb{P}(S_n - \mathbb{E}[S_n] \ge -t) \le \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$
(35)

LEMMA A.2. Let r > 0 and let Z be a real-valued random variable such that for all  $t \ge 0$ ,

$$\mathbb{P}(Z \ge t) \le 2q \exp\left(-\frac{t^2}{r^2}\right). \tag{36}$$

*Then, for* q > 1*,* 

$$\mathbb{E}[Z] \leqslant r\left(\sqrt{\ln q} + \frac{1}{\sqrt{\ln q}}\right) \tag{37}$$

and for q = 1,

$$\mathbb{E}[Z] \leqslant \sqrt{2}r. \tag{38}$$

$$\begin{split} & \mathbb{P}_{\text{PROOF.}} \\ & \mathbb{E}[Z] = \int_0^\infty \mathbb{P}(Z \ge t) dt \\ & = \int_0^{r\sqrt{\ln q}} \mathbb{P}(Z \ge t) dt + \int_{r\sqrt{\ln q}}^\infty 2q \exp\left(-\frac{t^2}{r^2}\right) dt \\ & \leqslant r\sqrt{\ln q} + \int_{r\sqrt{\ln q}}^\infty 2q \exp\left(-\frac{t^2}{r^2}\right) dt \\ & \leqslant r\sqrt{\ln q} + 2q \int_{r\sqrt{\ln q}}^\infty \frac{t}{r\sqrt{\ln q}} \exp\left(-\frac{t^2}{r^2}\right) dt \\ & = r\sqrt{\ln q} + \frac{2q}{r\sqrt{\ln q}} \frac{r^2}{2} \exp\left(-\frac{(r\sqrt{\ln q})^2}{r^2}\right) \\ & = r\left(\sqrt{\ln q} + \frac{1}{\sqrt{\ln q}}\right). \end{split}$$

Similarly, the second statement fallowing by taking q = 1 and splitting the integral between  $t \in [0, r]$  and  $t \in [r, \infty)$ .

# **B PROOFS FOR THE STATISTICAL BOUNDS**

#### B.1 Proof of Lemma 3.1

LEMMA 3.1. The following inequalities hold:

$$V_g - V(f_m^*) \leq 2 \max_{f \in \mathcal{H}} \left| \hat{V}_m(f) - V(f) \right|, \tag{11}$$

$$\left|\widehat{V}_{n}(f_{m}^{\star}) - V(f_{m}^{\star})\right| \leq \max_{f \in \mathcal{H}} \left|\widehat{V}_{n}(f) - V(f)\right|.$$
 (12)

PROOF. Recall the definition of  $f_m^{\star}$  in (8) and let  $f^{\star} = \arg \max_{f \in \mathcal{H}} V(f)$ . We first observe that

$$\left|\widehat{V}_{n}(f') - V(f')\right| \leq \max_{f \in \mathcal{H}} \left|\widehat{V}_{n}(f) - V(f)\right|$$

for all  $f' \in \mathcal{H}$ . Inequality (12) follows by letting  $f' = f_m^*$ . We now prove expression (11):

$$V_{g} - V(f_{m}^{\star}) = V(f^{\star}) - \hat{V}_{m}(f_{m}^{\star}) + \hat{V}_{m}(f_{m}^{\star}) - V(f_{m}^{\star})$$
$$\leq V(f^{\star}) - \hat{V}_{m}(f_{m}^{\star}) + \left| \hat{V}_{m}(f_{m}^{\star}) - V(f_{m}^{\star}) \right|$$
(39)

$$\leq V(f^{\star}) - \hat{V}_m(f^{\star}) + \left| \hat{V}_m(f_m^{\star}) - V(f_m^{\star}) \right|$$
(40)

$$\leq \left| V(f^{\star}) - \widehat{V}_m(f^{\star}) \right| + \left| \widehat{V}_m(f_m^{\star}) - V(f_m^{\star}) \right|$$
(41)

$$\leq \max_{f \in \mathcal{H}} \left| \hat{V}_{m}(f) - V(f) \right| + \max_{f \in \mathcal{H}} \left| \hat{V}_{m}(f) - V(f) \right|$$
(42)

$$\leq 2 \max_{f \in \mathcal{H}} \left| \widehat{V}_m(f) - V(f) \right|. \tag{43}$$

#### **B.2** Proof of Proposition 3.2

PROPOSITION 3.2 (UNIFORM DEVIATIONS). Assume that  $|\mathcal{H}| < \infty$ and  $g : \mathcal{W} \times \mathcal{X} \longrightarrow [a, b]$ , for a, b real values such that  $a \ge 0$  and b > a. Then, we have for all  $\varepsilon > 0$ ,

$$\sup_{P_{XY}} \mathbb{P}\left(\left|\widehat{V}_n(f_m^{\star}) - V(f_m^{\star})\right| > \varepsilon\right) \leq 2 \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right)$$
(13)

and

$$\sup_{P_{XY}} \mathbb{P}\left(V_g - V(f_m^{\star}) > \varepsilon\right) \leq 2|\mathcal{H}| \exp\left(-\frac{m\varepsilon^2}{2(b-a)^2}\right).$$
(14)

PROOF. We first prove expression (13). Notice that

$$\mathbb{P}\left(\left|\widehat{V}_{n}(f_{m}^{\star}) - V(f_{m}^{\star})\right| > \varepsilon\right)$$
$$= \mathbb{E}_{\mathcal{D}_{m} \sim P_{XY}^{m}} \mathbb{P}\left(\left|\widehat{V}_{n}(f_{m}^{\star}) - V(f_{m}^{\star})\right| > \varepsilon \mid \mathcal{D}_{m}\right)$$
(44)

$$\leq 2 \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right),$$
(45)

where (44) follows by conditioning on the training samples and then taking the probability on the validation samples, and (45) follows by noticing that given the training samples the function  $f_m^{\star}$  is fixed by applying Proposition A.1 the inequality follows. The second inequality in (14) is a consequence of the following steps:

$$\mathbb{P}\left(V_{g} - V(f_{m}^{\star}) > \varepsilon\right) \leq \mathbb{P}\left(\max_{f \in \mathcal{H}} \left|\widehat{V}_{m}(f) - V(f)\right| > \varepsilon/2\right) \quad (46)$$
$$= \mathbb{P}\left(\bigcup_{f \in \mathcal{H}} \left\{\left|\widehat{V}_{m}(f) - V(f)\right| > \varepsilon/2\right\}\right) \quad (47)$$

$$\leq \sum_{f \in \mathcal{H}} \mathbb{P}\left( \left| \widehat{V}_m(f) - V(f) \right| > \varepsilon/2 \right)$$
 (48)

$$\leq 2|\mathcal{H}|\exp\left(-\frac{2m\varepsilon^2}{4(b-a)^2}\right),$$
 (49)

where (46) follows by applying the first inequality in Lemma 3.1 and (49) follows from Proposition A.1 with an appropriate redefinition to  $\varepsilon/2$ .

# **B.3 Proof of Theorem 3.3**

THEOREM 3.3. The averaged estimation error of the *g*-vulnerability can be bounded as follows:

$$\mathbb{E}|V_g - \hat{V}_n(f_m^{\star})| \leq V_g - \mathbb{E}[V(f_m^{\star})] + \mathbb{E}|V(f_m^{\star}) - \hat{V}_n(f_m^{\star})|,$$

where the expectations are understood over all possible training and validation sets drawn according to  $P_{XY}$ . Furthermore,

$$V_{g} - \mathbb{E}\left[V(f_{m}^{\star})\right] \leqslant \sqrt{\frac{2(b-a)^{2}}{m}} \left(\sqrt{\ln|\mathcal{H}|} + \frac{1}{\sqrt{\ln|\mathcal{H}|}}\right), \quad (15)$$
$$\mathbb{E}\left|V(f_{m}^{\star}) - \hat{V}_{n}(f_{m}^{\star})\right| \leqslant \sqrt{\frac{(b-a)^{2}}{n}}, \quad (16)$$

independently of the specific underlying distribution  $P_{XY}$ .

PROOF. Observe that

$$\mathbb{E} |V_g - \hat{V}_n(f_m^{\star})| = \mathbb{E} |V_g - V(f_m^{\star}) + V(f_m^{\star}) - \hat{V}_n(f_m^{\star})|$$
  
$$\leq V_g - \mathbb{E} [V(f_m^{\star})] | + \mathbb{E} |V(f_m^{\star}) - \hat{V}_n(f_m^{\star})|,$$

which follows from the triangular inequality. We first bound the second term in the previous inequality as follows:

$$\mathbb{E}\left|V(f_m^{\star}) - \widehat{V}_n(f_m^{\star})\right| \leqslant \sqrt{\frac{(b-a)^2}{n}},\tag{50}$$

where the claim in (50) follows by applying Proposition 3.2 and using expression (13) combined with Lemma A.2 choosing q = 1 and  $r^2 = \frac{(b-a)^2}{2n}$ .

We now bound the first term. Notice that

$$\mathbb{P}\left(V_g - V(f_m^{\star}) > \varepsilon\right) \leq 2|\mathcal{H}| \exp\left(-\frac{m\varepsilon^2}{2(b-a)^2}\right), \quad (51)$$

where (51) follows from inequality (13) in Proposition 3.2. By using again Lemma A.2 with  $q = |\mathcal{H}|$  and  $r^2 = \frac{2(b-a)^2}{m}$ , we obtain

$$V_{g} - \mathbb{E}\big[V(f_{m}^{\star})\big] \leqslant \sqrt{\frac{2(b-a)^{2}}{m}} \left(\sqrt{\log|\mathcal{H}|} + \frac{1}{\sqrt{\log|\mathcal{H}|}}\right),\tag{52}$$

which concludes the proof of the Theorem.

# **B.4** Proof of Corollary 3.5

COROLLARY 3.5. The sample complexity of the ERM algorithm *q*-vulnerability is bounded from above by the set of values satisfying:

$$M(\varepsilon,\delta) \leqslant \frac{2(b-a)^2}{\varepsilon^2} \ln\left(\frac{2|\mathcal{H}|}{\delta - \Delta}\right),\tag{20}$$

$$N(\varepsilon, \delta) \leq \frac{(b-a)^2}{2\varepsilon^2} \ln\left(\frac{2}{\Delta}\right),$$
 (21)

for all  $0 < \Delta < \delta$ .

PROOF. We first notice that

$$\mathbb{P}\left(|V_g - \hat{V}_n(f_m^{\star})| > \varepsilon\right) \leq \mathbb{P}\left(V_g - V(f_m^{\star}) > \varepsilon\right) + \mathbb{P}\left(|V(f_m^{\star}) - \hat{V}_n(f_m^{\star})| > \varepsilon\right), \quad (53)$$

and thus from (13) and (14) in Proposition 3.2, we have

$$\mathbb{P}\left(|V_g - \hat{V}_n(f_m^{\star})| > \varepsilon\right) \leq 2 \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right) + 2|\mathcal{H}|\exp\left(-\frac{m\varepsilon^2}{2(b-a)^2}\right).$$
(54)

Let us require:

$$2|\mathcal{H}|\exp\left(-\frac{m\varepsilon^2}{2(b-a)^2}\right) \leqslant (\delta - \Delta),\tag{55}$$

$$2\exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right) \leq \Delta,\tag{56}$$

which satisfies the desired condition:

ł

$$\mathbb{P}\left(|V_g - \hat{V}_n(f_m^{\star})| > \varepsilon\right) \leq \delta,$$
(57)

for any  $0 < \Delta < \delta$ . Finally, from the previous inequality we can derive lower bounds on *n* and *m*:

$$n \ge \frac{2(b-a)^2}{\varepsilon^2} \ln\left(\frac{2|\mathcal{H}|}{\delta - \Delta}\right),\tag{58}$$

$$n \ge \frac{(b-a)^2}{2\varepsilon^2} \ln\left(\frac{2}{\Delta}\right),$$
(59)

which by definition of sample complexity shows the corollary.  $\Box$ 

#### C PRE-PROCESSING

#### C.1 Data pre-processing

THEOREM 4.1 (CORRECTNESS OF DATA PRE-PROCESSING). Given a prior  $\pi$ , a channel C, and a gain function g, we have

$$V_q(\pi, C) = \alpha \cdot V_{q_{id}}(\sigma, E),$$

where  $\alpha$ ,  $\sigma$  and E are those defined in (23), (24) and (25), respectively, and  $g_{id}$  is the identity function (cfr. section 2), i.e., the gain function corresponding to the Bayesian adversary.

Proof.

$$V_{g_{id}}(\sigma, E) = \sum_{y} \max_{w} \sum_{w'} \sigma_{w'} \cdot E_{w'y} \cdot g_{id}(w, w')$$
  
$$= \sum_{y} \max_{w} (\sigma_{w} E_{wy})$$
  
$$= \sum_{y} \max_{w} P_{WY}(w, y)$$
  
$$= \sum_{y} \max_{w} \frac{U(w, y)}{\alpha}$$
  
$$= \frac{1}{\alpha} \cdot \sum_{y} \max_{w} \sum_{x} \pi_{x} \cdot C_{xy} \cdot g(w, x)$$
  
$$= \frac{1}{\alpha} \cdot V_{g}(\pi, C)$$

#### C.2 Channel pre-processing

THEOREM 4.2 (CORRECTNESS OF CHANNEL PRE-PROCESSING). Given a prior  $\pi$  and a gain function g, we have that, for any channel C:

$$V_g(\pi, C) = \beta \cdot V_{g_{id}}(\tau, RC)$$
 for all channels C.

where  $\beta$ ,  $\tau$  and R are those defined in (26) and (27).

**PROOF.** In this proof we use a notation that highlights the structure of the preprocessing. We will denote by *G* be the matrix form of *g*, i.e.,  $G_{wx} = g(w, x)$ , and by  $\Psi^{\pi}$  the square matrix with  $\pi$  in its diagonal and 0 elsewhere. We have that  $\beta = \|G\Psi^{\pi}\|_1 = \sum_{w,x} G_{wx} \pi_x$ , which is strictly positive because of the assumptions on *g* and  $\pi$ . Furthermore, we have

$$\tau^T = \beta^{-1} G \Psi^{\pi} \mathbf{1}, \qquad R = \beta^{-1} (\Psi^{\tau})^{-1} G \Psi^{\pi},$$

where **1** is the vector of 1s and  $\tau^T$  represents the transposition of vector  $\tau$ . Note that  $(\Psi^{\tau})^{-1}$  is a diagonal matrix with entries  $\tau_w^{-1}$  in its diagonal. If  $\tau_w = 0$  then the row  $R_w$ , is not properly defined; but its choice does not affect  $V_{g_{id}}(\tau, RC)$  since the corresponding prior is 0; so we can choose  $R_w$ , arbitrarily (or equivalently remove the action w, it can never be optimal since it gives 0 gain). It is easy to check that  $\tau$  is a proper distribution and R is a proper channel:

$$\sum_{w} \tau_{w} = \sum_{w} \beta^{-1} \sum_{x} G_{wx} \pi_{x} = \beta^{-1} \beta = 1,$$
  
$$\sum_{x} R_{w,x} = \sum_{x} \frac{1}{\tau_{w}} \beta^{-1} G_{wx} \pi_{x} = \frac{\tau_{w}}{\tau_{w}} = 1.$$

Moreover, it holds that:

$$\beta \Psi^{\tau} R = \beta \Psi^{\tau} \beta^{-1} \Psi^{\tau-1} G \Psi^{\pi} = G \Psi^{\pi}$$

The main result follows from the trace-based formulation of posterior *g*-vulnerability [4], since for any channel *C* and strategy *S*, the above equation directly yields

$$V_g(\pi, C) = \max_{S} \operatorname{tr}(G\Psi^{\pi}CS)$$
$$= \beta \cdot \max_{S} \operatorname{tr}(\Psi^{\tau}RCS)$$
$$= \beta \cdot V_{g_{id}}(\tau, RC) ,$$

where  $\mathbf{tr}(\cdot)$  is the matrix trace.

# C.3 Data pre-processing when g is not integer

Approximating *g* so that it only takes values  $\in \mathbb{Q}_{\geq 0}$  allows us to represent each gain as a quotient of two integers, namely

Numerator(
$$G_{w,x}$$
)/Denominator ( $G_{w,x}$ ).

Let us also define

$$K \stackrel{\text{def}}{=} \operatorname{lcm}_{wx}(\operatorname{Denominator}(G_{w,x})), \tag{60}$$

where  $lcm(\cdot)$  is the least common multiple. Multiplying *G* by *K* gives the integer version of the gain matrix that can replace the original one. It is clear that the calculation of the least common multiplier, as well as the increase in the amount of data produced during the dataset building using a gain matrix forced to be integer, might constitute a relevant computational burden.

# D ANN MODELS

We list here the specifics for the ANNs models used in the experiments. All the models are simple feed-forward networks whose layers are fully connected. The activation functions for the hidden neurons are rectifier linear functions, while the output layer has softmax activation function.

The loss function minimized during the training is the cross entropy, a popular choice in classification problems. The remapping  $\mathcal{Y} \rightarrow \mathcal{W}$  can be in fact considered as a classification problem such that, given an observable, a model learns to make the best guess.

For each experiments, the models have been tuned by crossvalidating them using one randomly chosen training sets among the available ones choosing among the largest in terms of samples. The specs are listed experiment by experiment:

- Multiple guesses scenario
  - Data pre-processing:
    - \* learning rate:  $10^{-3}$
    - \* hidden layers: 3
    - \* hidden units per layers: [100, 100, 100]
    - \* batch size: 1K samples
    - \* epochs: 700
    - Channel preprocessing:
    - \* learning rate: 10<sup>-3</sup>
    - \* hidden layers: 3
    - \* hidden units per layers: [100, 100, 100]
    - \* batch size: 1K samples
    - \* epochs: 500
- Location privacy scenario
  - Data pre-processing:

- learning rate: 10<sup>-3</sup>
- \* hidden layers: 3
- \* hidden units per layers: [500, 500, 500]
- \* batch size (on for each training set size in increasing order of size): [200, 500, 1*K*] samples
- \* epochs: 1k

- Channel preprocessing:
- \* learning rate:  $10^{-3}$
- \* hidden layers: 3
- \* hidden units per layers: [500, 500, 500]
- \* batch size (on for each training set size in increasing order of size): [20, 200, 500] samples
- \* epochs: 200, 500, 1K
- Differential privacy scenario
  - Data pre-processing:
    - \* learning rate: 10<sup>-3</sup>
    - \* hidden layers: 3
    - \* hidden units per layers: [100, 100, 100]
    - \* batch size: 200 samples
    - \* epochs: 500
  - Channel preprocessing:
    - \* learning rate:  $10^{-3}$
    - \* hidden layers: 3
    - \* hidden units per layers: [100, 100, 100]
    - \* batch size: 200 samples
    - \* epochs: 500
- Side channel on passwords attack scenario
  - Data and channel pre-processing
  - \* learning rate:  $10^{-3}$
  - \* hidden layers: 3
  - \* hidden units per layers: [100, 100, 100]
  - \* batch size: 1K samples
  - \* epochs: 700

#### E FREQUENTIST APPROACH DESCRIPTION

In the frequentist approach the elements of the channel, namely the conditional probabilities  $P_{Y|X}(y|x)$ , are estimated directly in the following way: the empirical prior probability of x,  $\hat{\pi}_x$ , is computed by counting the number of occurrences of x in the training set and dividing the result by the total number of elements. Analogously, the empirical joint probability  $\hat{P}_{XY}(x, y)$  is computed by counting the number of occurrences of the pair (x, y) and dividing the result by the total number in the set. The estimation  $\hat{C}_{xy}$  of  $C_{xy}$  is then defined as

$$\widehat{C}_{xy} = \frac{\widehat{P}_{XY}(x,y)}{\widehat{\pi}(x)}.$$
(61)

In order to have a fair comparison with our approach, which takes advantage of the fact that we have several training sets and validation sets at our disposal, while preserving at the same time the spirit of the frequentist approach, we proceed as follows: Let us consider a training set  $\mathcal{D}_m$ , that we will use to learn the best remapping  $\mathcal{Y} \to \mathcal{W}$ , and a validation set  $\mathcal{T}_n$  which is then used to actually estimate the *g*-vulnerability. We first compute  $\hat{\pi}$  using  $\mathcal{D}_m$ . For each *y* in  $\mathcal{Y}$  and for each  $x \in \mathcal{X}$ , the empirical probability  $\hat{P}_{X|Y}$  is computed using  $\mathcal{D}_m$  as well. In particular,  $\hat{P}_{X|Y}(x|y)$  is given by the number of times *x* appears in pair with *y* divided by the number of occurrences of *y*. In case a certain *y* is in  $\mathcal{T}_n$  but not in  $\mathcal{D}_m$ , it is assigned the secret  $x' = \arg \max_{x \in \mathcal{X}} \hat{\pi}$  so that  $\hat{P}_{X|Y}(x'|y) = 1$  and  $\hat{P}_{X|Y}(x) = 0$ ,  $\forall x \neq x'$ . It is now possible to find the best mapping for each *y* defined as  $w(y) = \arg \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \hat{P}_{X|Y}(x|y)g(w, x)$ . Now we compute the empirical joint distribution for each (x, y) in  $\mathcal{T}_n$ , namely  $\hat{Q}_{XY}$ , as the number of occurrences of (x, y) divided by the total number of samples in  $\mathcal{T}_n$ . We now estimate the *g*-vulnerability

on the validation samples according to:

$$\widehat{V}_n = \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} \widehat{Q}_{XY}(x, y) g(w(y), x).$$
(62)

# F SUPPLEMENTARY PLOTS

In the next pages, we show supplementary plots for each experiment presented in section 5. In particular, we have included the plots representing the comparison between the estimated vulnerability and the real one, and the plots showing the comparison between the frequentist approach and ours.

Figure 12 is related to the multiple guesses scenario, Figure 13 is related to the location privacy one, Figure 14 is related to the differential privacy experiment, and Figure 15 to the password checker one.



(a) Vulnerability estimation for ANN and k-NN with data pre-processing.



(c) Vulnerability estimation for the frequentist approach.



(e) Vulnerability estimation for ANN and k-NN with data pre-processing, and the frequentist approach.



(b) Vulnerability estimation for ANN and k-NN with channel pre-processing.



(d) Normalized estimation error for the frequentist approach.



(f) Vulnerability estimation for ANN and k-NN with channel pre-processing, and the frequentist approach.



(g) Normalized estimation error for ANN and k-NN with channel pre-processing, and the frequentist approach.

Figure 12: Supplementary plots for the multiple-guesses experiment.



(a) Vulnerability estimation for ANN and k-NN with data pre-processing.



(c) Vulnerability estimation for the frequentist approach.



(e) Vulnerability estimation for ANN and k-NN with data pre-processing, and the frequentist approach.



(g) Vulnerability estimation for ANN and k-NN with channel pre-processing, and the frequentist approach.



(b) Vulnerability estimation for ANN and k-NN with channel pre-processing.



(d) Normalized estimation error for the frequentist approach.



(f) Normalized estimation error for ANN and k-NN with data pre-processing, and the frequentist approach.



(h) Normalized estimation error for ANN and k-NN with channel pre-processing, and the frequentist approach.

Figure 13: Supplementary plots for the location-privacy experiment.

19



(a) Vulnerability estimation for ANN and k-NN with data pre-processing.



(c) Vulnerability estimation for the frequentist approach.



(e) Vulnerability estimation for ANN and k-NN with data pre-processing, and the frequentist approach.



(g) Vulnerability estimation for ANN and k-NN with channel pre-processing, and the frequentist approach.



(b) Vulnerability estimation for ANN and k-NN with channel pre-processing.



(d) Normalized estimation error for the frequentist approach.



(f) Normalized estimation error for ANN and k-NN with data pre-processing, and the frequentist approach.



(h) Normalized estimation error for ANN and k-NN with channel pre-processing, and the frequentist approach.

Figure 14: Supplementary plots for the differential-privacy experiment.

proach.



(a) Vulnerability estimation for ANN and k-NN with data and channel pre-processing.



(c) Normalized estimation error for the frequentist approach.



(b) Vulnerabiloty estimation for the frequentist approach.



(d) Vulnerability estimation for ANN and k-NN with data and channel pre-processing, and the frequentist approach.



(e) Normalized estimation error for ANN and k-NN with data and channel pre-processing, and the frequentist approach.

Figure 15: Supplementary plots for the password-checker experiment.