

Adversarial Classification under Gaussian Mechanism: Calibrating the Attack to Sensitivity

Ayşe Ünsal

Digital Security Dept., EURECOM

ayse.unsal@eurecom.fr

Melek Önen

Digital Security Dept., EURECOM

melek.onen@eurecom.fr

Abstract—This work studies anomaly detection under differential privacy with Gaussian perturbation using both statistical and information-theoretic tools. In our setting, the adversary aims to modify the content of a statistical dataset by inserting additional data without being detected using the differential privacy to her/his own benefit. To this end, firstly via hypothesis testing, we characterize a statistical threshold for the adversary, which balances the privacy budget and the induced bias (the impact of the attack) in order to remain undetected. In addition, we establish the privacy-distortion trade-off in the sense of the well-known rate-distortion function for the Gaussian mechanism by using an information-theoretic approach to avoid detection. Accordingly, we derive an upper bound on the variance of the attacker’s additional data as a function of the sensitivity and the original data’s second-order statistics. Lastly, we introduce a new privacy metric based on Chernoff information for classifying adversaries under differential privacy as a stronger alternative for the Gaussian mechanism. Analytical results are supported by numerical evaluations.

Index Terms—differential privacy, adversarial classification, Chernoff information, Gaussian mechanism

I. INTRODUCTION

The major issue in terms of data privacy in today’s world stems from the fact that machine learning (ML) algorithms strongly depend on the use of large datasets to work efficiently and accurately. Along with the highly increased deployment of ML, its privacy aspect rightfully became a cause of concern, since the collection of such large datasets makes users vulnerable to fraudulent use of personal, (possibly) sensitive information. This vulnerability is aimed to be mitigated by privacy enhancing technologies that are designed to protect data privacy of users.

Differential privacy (DP) has been proposed to address this problem and it has furthermore been used to develop practical methods for protecting private user-data. Dwork’s original definition of DP in [1] emanates from a notion of statistical indistinguishability of two different probability distributions which is obtained through randomization of the data prior to its publication. The outputs of two differentially private mechanisms are indistinguishable for two datasets that only differ in one user’s data, i.e. neighbors. In other words, DP guarantees that the output of the mechanism is *statistically*

indifferent to changes made in a single row of the dataset proportional to its privacy budget.

Let us imagine the scenario, where it is possible to weaponize privacy protection methods by adversaries in order to avoid being detected. Adversarial classification/anomaly detection is an application of the supervised ML approach to detect misclassification attacks where adversaries shield themselves by using DP to remain undetected. This paper studies anomaly detection in differentially private Gaussian mechanisms to establish the trade-off between the probability distribution of the noise and the impact of the attack to remain indistinguishable by employing both statistical and information-theoretic tools. In our setting, we consider an adversary who not only aims to discover the information of a dataset but also wants to harm it by inserting data on the original dataset. Accordingly, we establish stochastic and information-theoretic relations between the impact of the adversary’s attack and the privacy budget of the Gaussian mechanism.

This work, in part, is an extension of [2] to Gaussian mechanisms which introduced statistical thresholds for adversarial classification in Laplace mechanisms. As for the methodology, in addition to statistical hypothesis testing as in [2], [3], in this work, we also derive the mutual information between the datasets before and after the attack (considered as neighbors) to upper bound the second-order statistics of the data added to the system by the adversary, in order to determine an information-theoretic threshold for correctly detecting the attack. Originally, the lossy source-coding approach in the information-theoretic DP literature has mostly been used to quantify the privacy guarantee [4] or the leakage [5], [6]. [7] stands out in the way that the rate-distortion perspective is applied to DP where various fidelity criteria is set to determine how fast the empirical distribution converges to the actual source distribution. In [2], the authors presented an application of the Kullback-Leibler (KL)-DP for detecting misclassification attacks in Laplace mechanisms, where the corresponding distributions in relative entropy were considered as the differentially private noise with and without the adversary’s advantage. This work introduces a novel DP metric based on Chernoff information along with its application to adversarial classification.

Our contributions are summarized by the following list.

- In this paper, we establish a statistical threshold to avoid

This work has been supported by the 3IA Côte d’Azur Interdisciplinary Institute for Artificial Intelligence project with the reference number ANR-19-P3IA-0002.

detection for the adversary’s hypothesis testing problem as a function of the bias induced by the adversary’s attack, probability of false-alarm and the privacy budget.

- We apply a source-coding approach to anomaly detection under differential privacy to bound the variance of the additional data by the sensitivity of the mechanism and the original data’s statistics by deriving the mutual information between the neighboring datasets.
- We introduce a new DP metric which we call Chernoff DP as a stronger alternative to the well-known (ϵ, δ) -DP and KL-DP for the Gaussian mechanism. Chernoff DP is also adapted for adversarial classification and numerically shown to outperform KL-DP.

The outline of the paper is as follows. In the upcoming section, we remind the reader of some important preliminaries from the DP literature which will be used throughout this paper along with the detailed problem definition and performance criteria. In Section III, we present statistical and information-theoretic thresholds for anomaly detection in Gaussian mechanisms which will be followed by Section IV where we introduce a new metric of DP based on Chernoff information. We present numerical evaluation results and draw our final conclusions in Section V.

II. PRELIMINARIES, MODEL AND PERFORMANCE CRITERIA

Before presenting the addressed problem in detail, in the first part of this section, the reader is reminded of some preliminaries on DP.

A. Preliminaries for DP

Two datasets \mathbf{X} and $\tilde{\mathbf{X}}$ are called neighbors if $d(\mathbf{X}, \tilde{\mathbf{X}}) = 1$ where $d(\cdot, \cdot)$ denotes the Hamming distance. Accordingly, (ϵ, δ) -DP is defined by [8] as follows.

Definition 1. A randomized algorithm \mathcal{M} guarantees (ϵ, δ) -DP if $\forall \mathbf{X}, \tilde{\mathbf{X}}$ that are neighbors within the domain of \mathcal{M} and $\forall S \subseteq \text{Range}(\mathcal{M})$ the following inequality holds.

$$\Pr[\mathcal{M}(\mathbf{X}) \in S] \leq \Pr[\mathcal{M}(\tilde{\mathbf{X}}) \in S] \exp\{\epsilon\} + \delta \quad (1)$$

We will refer to the parameters ϵ and δ as privacy budget throughout the paper. Next definition reminds the reader of the L_2 norm global sensitivity.

Definition 2 (L_2 norm sensitivity). L_2 norm sensitivity denoted s refers to the smallest possible upper bound on the L_2 distance between the images of a query $q : D \rightarrow \mathbb{R}^k$ when applied to two neighboring datasets \mathbf{X} and $\tilde{\mathbf{X}}$ as

$$s = \sup_{d(\mathbf{X}, \tilde{\mathbf{X}})=1} \|q(\mathbf{X}) - q(\tilde{\mathbf{X}})\|_2. \quad (2)$$

Application of Gaussian noise results in a more relaxed privacy guarantee, that is (ϵ, δ) -DP contrary to Laplace mechanism, which brings about $(\epsilon, 0)$ -DP. (ϵ, δ) -DP is achieved by calibrating the noise variance as a function of the privacy budget and query sensitivity as given by the next definition.

Definition 3. Gaussian mechanism [9] is defined for a function (or a query) $q : D \rightarrow \mathbb{R}^k$ as follows

$$\mathcal{M}(\mathbf{X}, q(\cdot), \epsilon, \delta) = q(\mathbf{X}) + (Z_1, \dots, Z_k) \quad (3)$$

where $Z_i \sim \mathcal{N}(0, \sigma^2)$, $i = 1, \dots, k$ denote independent and identically distributed (i.i.d.) Gaussian random variables with the variance $\sigma_z^2 = \frac{2s^2 \log(1.25/\delta)}{\epsilon^2}$.

Lastly, we revisit the so-called Kullback-Leibler (KL) DP definition of [10].

Definition 4 (KL-DP). For a randomized mechanism $P_{Y|X}$ that guarantees ϵ -KL-DP, the following inequality holds for all its neighboring datasets x and \tilde{x} .

$$D(P_{Y|X=x} || P_{Y|X=\tilde{x}}) \leq \exp\{\epsilon\} \quad (4)$$

B. Problem Definition and Performance measures

We define the original dataset in the following form $\mathbf{X} = X^n = \{X_1, \dots, X_n\}$, where X_i are assumed to be i.i.d following the Gaussian distribution with the parameters $\mathcal{N}(0, \sigma_{X_i}^2)$. The query function takes the aggregation of this dataset as $q(\mathbf{X}) = \sum_i^n X_i$ and the DP-mechanism adds Gaussian noise Z on the query output leading to the noisy output in the following form $\mathcal{M}(\mathbf{X}, q(\cdot), \epsilon, \delta) = Y = \sum_i^n X_i + Z$. An adversary adds a single record denoted X_a to this dataset. The modified output of the DP-mechanism becomes $\sum_i^n X_i + X_a + Z$.

Statistical approach: In our first approach, we employ hypothesis testing to determine whether or not the defender fails to detect the attack. Accordingly, we set the following hypotheses where the null and alternative hypotheses are respectively translated into DP noise distribution with and without the bias induced by the attacker.

$$\begin{aligned} H_0 &: \text{defender fails to detect } X_a \\ H_1 &: \text{defender detects } X_a \end{aligned} \quad (5)$$

False alarm refers to the event when the defender detects the attack when in fact there was no attack with the corresponding probability denoted by α . Similarly, mis-detection is failing to detect an actual attack with the probability of occurrence denoted by β . This first part presents a trade-off between the shift due to the additional adversarial data, the privacy budget, the sensitivity of the query and the probability of false alarm by using the following likelihood ratio function $\Lambda = \frac{\mathcal{L}(p_1)}{\mathcal{L}(p_0)} \stackrel{H_0}{\leq} \stackrel{H_1}{\geq} k$ where p_1 and p_0 denote the noise distributions for the i^{th} hypothesis with the corresponding location parameter μ_i for $i = 0, 1$. The impact of the attack is denoted by $\Delta\mu$ and is equal to $\mu_1 - \mu_0$.

Information-theoretic approach: Our second approach is inspired by rate-distortion theory, where we employ the biggest possible difference between the images of the query for the datasets with and without the additional data X_a (i.e. neighboring inputs) as the fidelity criterion (Definition 2). Accordingly, we derive the mutual information between the original dataset

and its neighbor in order to bound the additional data's second-order statistics so that the defender fails to detect the attack. We assume that X_a follows a normal distribution with the variance $\sigma_{X_a}^2$. To simplify our derivations, we also assume that the original dataset $X^n = \{X_1, X_2, \dots, X_i, \dots, X_n\}$ and its neighbor $\tilde{X}^n = \{X_1, X_2, \dots, X_i + X_a, \dots, X_n\}$ have the same dimension n . Alternatively, the attack would change the size of the dataset as $n + 1$ where the additional data is not added to either of the X_i 's.

III. THRESHOLDS TO REMAIN UNDETECTED

In this part, we firstly present a statistical trade-off between the probability of false alarm, privacy budget and the impact of the attack via hypothesis testing. Additionally, we derive an information-theoretic upper bound on the second-order statistics of the additional data by employing a lossy source-coding approach to adversarial classification.

A. A Statistical Threshold to Avoid Detection

We present our main result by the following theorem.

Theorem 1. *The threshold of the best critical region of size α for the Gaussian mechanism with the largest possible power of the test for positive bias $\Delta\mu$ yields*

$$k = \exp \left\{ \frac{\Delta\mu}{\sigma_z} \left(Q^{-1}(\alpha) - \frac{\Delta\mu}{2\sigma_z} \right) \right\} \quad (6)$$

where $\sigma_z^2 = \frac{s^2 \log(1.25/\delta)^2}{\epsilon^2}$. The defender fails to detect the attack if $Y < \tilde{k} + q(\mathbf{X})$, where $q(\cdot)$ is the noiseless query output. By analogy, for $\Delta\mu < 0$, the attack is not detected if the DP output exceeds $\tilde{k} + q(\mathbf{X})$ where $\tilde{k} = \exp \left\{ \frac{\Delta\mu}{\sigma_z} \left(Q^{-1}(\bar{\alpha}) - \frac{\Delta\mu}{2\sigma_z} \right) \right\}$ for $\bar{\alpha} = 1 - \alpha$.

Proof. Likelihood ratio function Λ to choose between $Y - \sum_i^n X_i$ and $Y - \sum_i^n X_i - X_a$ results in $z > \tilde{k}$ where $\tilde{k} = \frac{\sigma_z \log k}{\Delta\mu} + \frac{\mu_1 + \mu_0}{2}$ by setting p_0 and p_1 as Gaussian distributions with respective location parameters μ_0 and μ_1 and the mutual scale parameter σ_z . Probability of false alarm is derived using this condition as

$$\alpha = \Pr[z > \tilde{k} | H_0 \text{ is true}] = Q \left(\frac{\sigma_z \log k}{\Delta\mu} + \frac{\Delta\mu}{2\sigma_z} \right) \quad (7)$$

where $Q(\cdot)$ denotes the Gaussian Q-function defined as $\Pr[Z > z]$ for standard Gaussian random variables. The threshold of the critical region k for $\Delta\mu > 0$ is obtained as a function of the probability of false-alarm as

$$k = \exp \left\{ \frac{\Delta\mu}{\sigma_z} \left(Q^{-1}(\alpha) - \Delta\mu/2\sigma_z \right) \right\} \quad (8)$$

By analogy, for $\Delta\mu < 0$, the probability of raising a false-alarm and the power of the test yield

$$\alpha = 1 - Q \left(\frac{\sigma_z \log k}{\Delta\mu} + \Delta\mu/(2\sigma_z) \right) \quad (9)$$

Rewriting (9), we obtain the second threshold given by Theorem 1 for negative bias. \square

In order to compute the receiver operating characteristic (ROC) curves, to visualize the effect of the privacy budget and the impact of the attack on the accuracy of correctly detecting the attacker, we derive the power of the test for both cases obtained through $\Pr[z > \tilde{k} | H_1 \text{ is true}]$ as follows

$$\bar{\beta} = \begin{cases} Q(Q^{-1}(\alpha) - \Delta\mu/\sigma_z), & \text{for } \Delta\mu > 0, \\ 1 - Q(Q^{-1}(\bar{\alpha}) - \Delta\mu/\sigma_z) & \text{for } \Delta\mu < 0 \end{cases} \quad (10)$$

where $\bar{\beta} = 1 - \beta$. Numerical evaluation results of Theorem 1 are presented in Section V.

B. Privacy-Distortion Trade-off for Adversarial Classification

The idea in this part is to render the problem of adversarial classification under differential privacy as a lossy source-coding problem. Instead of using the mutual information between the input and output (or the input's estimate obtained by using the output) of the mechanism, for this problem we derive the mutual information between the (neighboring) datasets before and after the attack, according to the adversary's conflicting goals as maximizing the induced bias while remaining undetected. The first expansion proceeds as follows considering the neighbor that includes X_a has now $(n + 1)$ entries over n rows.

$$I(X^n; \tilde{X}^n) = h(\tilde{X}^n) - h(\tilde{X}^n | X^n) \quad (11)$$

$$= \frac{1}{2} \log(2\pi e)^n \left(\sum_{i=1}^n \sigma_{X_i}^2 + \sigma_{X_a}^2 \right) - \frac{1}{2} \log((2\pi e)^n \sigma_{X_a}^2) \quad (12)$$

$$= \frac{1}{2} \log \left(1 + \sum_i^n \sigma_{X_i}^2 / \sigma_{X_a}^2 \right) \quad (13)$$

Due to the adversary's attack, in the first term of (12), we add up the variances of $(n + 1)$ X_i 's including X_a . For the second expansion, we have

$$I(X^n; \tilde{X}^n) = h(X^n) - h(X^n | \tilde{X}^n) \quad (14)$$

$$\geq h(X^n) - h(f(X^n) - \tilde{X}^n | \tilde{X}^n) \quad (15)$$

$$= h(X^n) - h(f(X^n) - f(\tilde{X}^n) | \tilde{X}^n) \quad (16)$$

$$\geq h(X^n) - h(f(X^n) - f(\tilde{X}^n)) \quad (17)$$

$$\geq \frac{1}{2} \log \left((2\pi e)^n \sum_i^n \sigma_{X_i}^2 \right) - \frac{1}{2} \log((2\pi e)s^2) \quad (18)$$

$$= \frac{1}{2} \log \left((2\pi e)^{n-1} \sum_i^n \sigma_{X_i}^2 / s^2 \right) \quad (19)$$

In (15), we apply the following property due to concavity of entropy function, $h(g(x)) \leq h(x)$ for any function $g(\cdot)$. In (17), conditioning reduces entropy and in (18), we plug in Definition 2 into the second term. Since (13) \geq (19), global sensitivity is bounded as follows in terms of the second-order statistics of the original data and those of the additional data X_a .

$$s \geq (2\pi e)^{\frac{n-1}{2}} \left(\frac{\sum_i^n \sigma_{X_i}^2 \cdot \sigma_{X_a}}{\sum_i^n \sigma_{X_i}^2 + \sigma_{X_a}^2} \right)^{1/2} \quad (20)$$

Alternatively, the lower bound on the sensitivity of the Gaussian mechanism can be used as an upper bound on $\sigma_{X_a}^2$ to yield a threshold in terms of the additional data's variance as a function of the privacy budget and the original data's statistics to guarantee that the adversary avoids being detected. Accordingly, we obtain the following upper bound

$$\sigma_{X_a}^2 \leq \frac{\sum_i^n \sigma_{X_i}^2}{\left(\sum_i^n \sigma_{X_i}^2 (2\pi e)^{n-1} / s^2\right) - 1} \quad (21)$$

where $s^2 = \frac{\sigma_0^2 \epsilon^2}{2 \log(1.25/\delta)}$ due to Definition 3.

Remark 1. *The second expansion of the mutual information between neighboring datasets derived in (19), can be related to the well-known **rate-distortion function of the Gaussian source** which, originally, provides the minimum possible transmission rate for a given distortion balancing (mostly for the Gaussian case) the squared-error distortion with the source variance. This is contradicting the adversary's goal in our setting, where the adversary aims to **maximize** the damage that s/he inflicts on the DP-mechanism. But at the same time to avoid being detected, the attack is calibrated according to the sensitivity which here replaces the distortion. Thus, unlike the classical rate-distortion theory, here the mutual information between the neighbors should be maximized for a given sensitivity to simultaneously satisfy adversary's conflicting goals for the problem of adversarial classification under Gaussian DP-mechanism. Also note that the additional factor $2\pi e$ appears in our bounds as opposed to the original rate-distortion function and the corresponding lower bound on squared-error distortion due to the choice of the query function that aggregates the entire dataset and reduces the dimension.*

IV. CHERNOFF DIFFERENTIAL PRIVACY

In the classical approach, the best error exponent in hypothesis testing for choosing between two probability distributions is the Kullback-Leibler divergence between these two distributions due to Stein's lemma [11]. In the Bayesian setting, however, assigning prior probabilities to each of the hypotheses in a binary hypothesis testing problem the best error exponent when the weighted sum probability of error, i.e. $\pi = a\alpha + b\beta$ for $b = 1 - a$ and $a \in (0, 1)$, is minimized, corresponds to the *Chernoff information/divergence*. The Chernoff information between two probability distributions f_0 and f_1 with prior probabilities a and b is defined as

$$C_a(f_0||f_1) = \log \int_x f_0(x)^a f_1^b(x) dx \quad (22)$$

The Renyi divergence denoted $D_a(f_0||f_1)$ between two Gaussian distributions with parameters $\mathcal{N}(\mu_0, \sigma_0^2)$ and $\mathcal{N}(\mu_1, \sigma_1^2)$ is given in [12] by

$$D_a(f_0||f_1) = \ln \frac{\sigma_1}{\sigma_0} + \frac{1}{2(a-1)} \ln \left(\frac{\sigma_1^2}{(\sigma^2)_a^*} \right) + \frac{1}{2} \frac{a(\mu_0 - \mu_1)^2}{(\sigma^2)_a^*} \quad (23)$$

where $(\sigma^2)_a^* = a\sigma_1^2 + b\sigma_0^2$. Using the following relation between Chernoff information and Renyi divergence

$D_a(f_0||f_1) = \frac{1}{1-a} C_a(f_0||f_1)$, we obtain the Gaussian univariate Chernoff information¹ with priors $a = (1 - a)$ and constant standard deviation $\sigma_0 = \sigma_1 = \sigma$ as follows.

$$C(f_0||f_1) = \frac{(\mu_0 - \mu_1)^2}{8\sigma^2}. \quad (24)$$

On the other hand, KL divergence between two Gaussian distributions denoted $D_{KL}(f_0||f_1)$ is derived as $\log \left(\frac{\sigma_1}{\sigma_0} \right) + \frac{1}{2} \frac{\sigma_0^2}{\sigma_1^2} + \frac{(\mu_1 - \mu_0)^2}{2\sigma_1^2} - \frac{1}{2}$. The next definition provides an adaptation of Chernoff information to quantify differential privacy guarantee as a stronger alternative to KL-DP of Definition 4 and (ϵ, δ) -DP for Gaussian mechanisms. We apply this to our problem setting for adversarial classification under Gaussian mechanisms, where the query output before and after the attack are $\sum_i^n X_i$ and $\sum_i^n X_i + X_a$, respectively. The corresponding distributions are considered as the DP noise with and without the induced value of X_a by the attacker as in our original hypothesis testing problem in (5).

Definition 5 (Chernoff differential privacy). *For a randomized mechanism $P_{Y|X}$ guarantees ϵ -Chernoff-DP, if the following inequality holds for all its neighboring datasets x and \tilde{x}*

$$C_a(P_{Y|X=x}||P_{Y|X=\tilde{x}}) \leq \exp(\epsilon) \quad (25)$$

where $C_a(\cdot||\cdot)$ is defined by (22).

[10, Theorem 1] proves that KL-DP defined in Definition 4 is a stronger privacy metric than (ϵ, δ) -DP that is achieved by Gaussian mechanism. Accordingly, the following chain of inequalities are proven to hold for various definitions of differential privacy

$$\epsilon - \text{DP} \stackrel{a}{\succeq} \text{KL} - \text{DP} \stackrel{b}{\succeq} \text{MI} - \text{DP} \stackrel{c}{\succeq} \delta - \text{DP} \stackrel{d}{=} (\epsilon, \delta) - \text{DP}$$

where MI-DP refers to the mutual information DP defined by $\sup_{i, P_{X^n}} I(X_i; Y|X^{-i}) \leq \epsilon$ nats for a dataset $X^n = (X_1, \dots, X_n)$ with the corresponding output Y according to the randomized mechanism represented by $P_{Y|X^n}$ where X^{-i} denotes the dataset entries excluding X_i . δ -DP represents the case when $\epsilon = 0$ in (ϵ, δ) -DP.

Chernoff information based definition of differential privacy is a **stronger privacy metric** than KL-DP, and thus (ϵ, δ) -DP for the Gaussian mechanism due to prior probabilities. Such a comparison is presented numerically in Figure 1. Numerical evaluation also supports the same conclusion.

V. NUMERICAL EVALUATIONS AND CONCLUSION

Figure 1 depicts Chernoff DP and KL-DP for various levels of privacy and the impact of the attack which were set as a function of the global sensitivity. Accordingly, the attack is compared to the privacy constrained of Definition 5 that is referred as the upper bound in the legend. Due to prior probabilities, Chernoff information is tighter than KL divergence consequently, it provides a more strict privacy

¹An alternative method to derive Chernoff information is the use of Exponential families as shown in [13].

constraint. Figure 1 confirms that increasing the impact of the attack as a function of the sensitivity closes the gap with the upper bound for Chernoff-DP. Additionally, the KL-DP does not violate the upper bound of the privacy budget only in the high privacy regime (when ϵ is small) for the cases of $\Delta\mu = 2 \cdot s$ and $\Delta\mu = 4 \cdot s$.

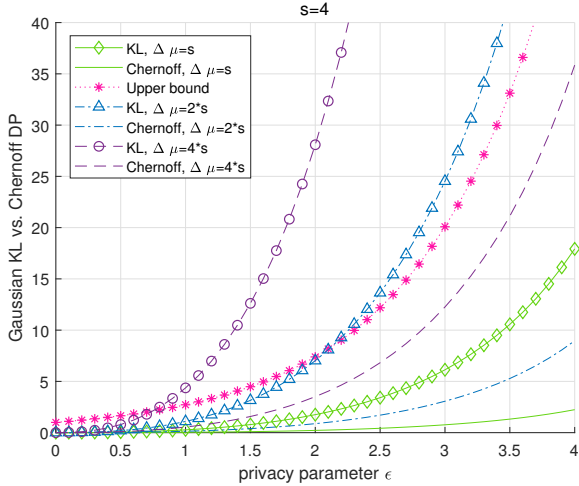


Fig. 1. KL-DP vs. Chernoff DP for various levels of privacy budget with global sensitivity $s = 4$

Figure 2 presents ROC curves computed using the threshold of (6) for adversarial classification under Gaussian DP for three different scenarios where the impact of the attack is greater than, equal to and less than the L_2 norm global sensitivity (in this order) for various levels of privacy budget. We observe that in the low privacy regime (i.e. when ϵ is large) the accuracy of the test is high which comes at the expense of the privacy guarantee since as the privacy budget is decreased (higher privacy) the test is no longer accurate and the adversary cannot be correctly detected with high probability. Another observation can be made based on the effect of the relationship between the attack and sensitivity. Unsurprisingly, increasing the bias $\Delta\mu$ as opposed to s also increases the probability of correctly detecting the attacker.

Conclusion: We established statistical and information-theoretic trade-offs between the security of the Gaussian DP-mechanism and the adversary's advantage who aims to trick the classifier that detects anomalies/outliers. Firstly, we determined a statistical threshold that offsets the DP-mechanism's privacy budget against the impact of the adversary's attack to remain undetected. Secondly, we characterized the privacy-distortion trade-off of the Gaussian mechanism in a form of the well-known Gaussian rate-distortion function and bounded the impact of the adversary's modification on the original data in order to avoid detection. We introduced Chernoff DP and its application to adversarial classification which turned out to be a stronger privacy metric than KL-DP and (ϵ, δ) -DP for the Gaussian mechanism. Numerical evaluation shows that, the effect of increasing the impact of the attack closes the gap with the DP upper bound. Using information-theoretic quantities as




Fig. 2. Eqs. (9) and (10) for various values of ϵ , $\Delta\mu > 0$ and $\delta = \epsilon/20$.

privacy constraint is not fully exploited despite its practicality. Future work will focus on general solutions for different types of queries and attacks.

REFERENCES

- [1] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, 2006, pp. 1–12.
- [2] A. Ünsal and M. Önen, "A Statistical Threshold for Adversarial Classification in Laplace Mechanisms," in *IEEE Information Theory Workshop 2021*, Oct. 2021.
- [3] C. Liu, X. He, T. Chanyaswad, S. Wang, and P. Mittal, "Investigating Statistical Privacy Frameworks from the Perspective of Hypothesis Testing," in *PETS 2019 Proceedings on Privacy Enhancing Technologies*, 2019, pp. 233–254.
- [4] W. Wang, L. Ying, and J. Zhang, "On the relation between identifiability, differential privacy and mutual information privacy," *IEEE Transactions on Information Theory*, vol. 62, pp. 5018–5029, Sep. 2016.
- [5] A. Sarwate and L. Sankar, "A rate-distortion perspective on local differential privacy," in *Fiftieth Annual Allerton Conference*, Oct. 2014, pp. 903–908.
- [6] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Fiftieth Annual Allerton Conference*, Oct. 2012, pp. 1401–1408.
- [7] A. Pastore and M. Gastpar, "Locally differentially private randomized response for discrete distribution learning," *Journal on Machine Learning Research*, vol. 22, pp. 1–56, Jul. 2021.
- [8] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science 2014*, vol. 9, pp. 211–407, 2014.
- [9] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating Noise to Sensitivity in Private Data Analysis," in *Theory of Cryptography Conference*, 2006, pp. 265–284.
- [10] P. Cuff and L. Yu, "Differential Privacy as a Mutual Information Constraint," in *CCS 2016, Vienna, Austria*, Oct. 2016.
- [11] T. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley Series in Telecommunications, 1991.
- [12] M. Gil, F. Alajaji, and T. Linder, "Renyi divergence measures for commonly used univariate continuous distributions," *Information Sciences*, vol. 249, Nov. 2013.
- [13] F. Nielsen, "An information-geometric characterization of chernoff information," *IEEE Signal Processing Letters*, vol. 20, Mar. 2013.