



**HAL**  
open science

# A unified view on differential privacy and robustness to adversarial examples

Rafael Pinot, Florian Yger, Cedric Gouy-Pailler, Jamal Atif

► **To cite this version:**

Rafael Pinot, Florian Yger, Cedric Gouy-Pailler, Jamal Atif. A unified view on differential privacy and robustness to adversarial examples. Workshop on Machine Learning for CyberSecurity at ECMLP-KDD 2019, Sep 2019, Wurzburg, Germany. hal-02892170

**HAL Id: hal-02892170**

**<https://hal.science/hal-02892170>**

Submitted on 7 Jul 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A unified view on differential privacy and robustness to adversarial examples

Rafael Pinot<sup>1,2</sup>, Florian Yger<sup>1</sup>, Cédric Gouy-Pailler<sup>2</sup>, and Jamal Atif<sup>1</sup>

<sup>1</sup> Université Paris-Dauphine, PSL Research University, LAMSADE, Paris, France

<sup>2</sup> Institut LIST, CEA, Université Paris-Saclay, LADIS, Palaiseau, France  
`rafael.pinot@dauphine.fr`

**Abstract.** This short note highlights some links between two lines of research within the emerging topic of trustworthy machine learning: differential privacy and robustness to adversarial examples. By abstracting the definitions of both notions, we show that they build upon the same theoretical ground and hence results obtained so far in one domain can be transferred to the other. More precisely, our analysis is based on two key elements: probabilistic mappings (also called randomized algorithms in the differential privacy community), and the Renyi divergence which subsumes a large family of divergences. We first generalize the definition of robustness against adversarial examples to encompass probabilistic mappings. Then we observe that Renyi-differential privacy (a generalization of differential privacy recently proposed in [10]) and our definition of robustness share several similarities. We finally discuss how can both communities benefit from this connection to transfer technical tools from one research field to the other.

**Keywords:** Differential Privacy · Adversarial Examples · Renyi divergence

## 1 Introduction

With the large adoption of machine learning techniques in several domains (including critical ones), researchers and practitioners are observing growing concerns on the security and privacy of the tools they develop. A primary concern is to guarantee that sensitive information from the used databases are not leaked, accidentally disclosed, or inferred from the sole release of the model (*privacy preserving algorithms*). Beyond preserving privacy, a crucial issue of recent machine learning approaches is to protect the methods against malicious users targeting their weaknesses (e.g. *adversarial examples*, or poisoning attacks).

*Privacy preserving algorithms:* Several definitions have been introduced to characterize privacy preserving algorithms in the context of machine learning and data publishing. Among them, differential privacy has become the dominant standard by providing a formal and adaptive conception of privacy preserving data-analysis. It has been broadly investigated in numerous frameworks

and applications (see [5] for a complete overview of the field). An algorithm is said to be differentially private if, given two close databases, it produces statistically indistinguishable outputs. Highly correlated to the notion of "closeness" both in the input and output spaces, most frameworks [5,2,7,6] rely on divergences/pseudo-metrics between probability measures to characterize this notion. Recently, Mironov [10] proposed to use the well-known Renyi divergence to obtain a more general definition of privacy. This notion is well defined, and it exhibits principled theoretical advantages over previous definitions, which makes it the most general formulation of differential privacy introduced so far.

*Adversarial examples attacks:* Modern neural networks achieve state of the art performances in a variety of domains. However, it has been shown that such neural networks can be vulnerable to adversarial examples, i.e. imperceptible variations of legitimate examples crafted to deliberately mislead a machine learning algorithm [15]. Since then, attacks and defenses are developed in a tight back-and-forth (see [1] for a complete overview of the field). Most past defenses were deterministic (see e.g [9,14]), but recently, the idea of using randomization in the learning process to ensure robustness against adversarial examples attacks is gaining in interest [8,3,12].

*Outline of the paper:* We first recall the key notions of probabilistic mapping and Renyi divergence in Section 2. Then we introduce the notion of differential privacy and present its generalization called Renyi-differential privacy in Section 3. Section 4 presents the problem of adversarial examples and our generalized definition of robustness to these attacks. Finally we discuss in Section 5 the similarity between the two concepts, and an application to image classification in which we transfer tools from differential privacy to make algorithms robust to adversarial examples.

## 2 Preliminaries

Let us consider two arbitrary metric spaces  $(\mathcal{X}, d_{\mathcal{X}})$ , and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , let  $\sigma(\mathcal{Y})$  be a  $\sigma$ -algebra over  $\mathcal{Y}$  and  $\mathcal{P}(\mathcal{Y})$  be the set of probability measures over  $(\mathcal{Y}, \sigma(\mathcal{Y}))$ . The notion of probabilistic mapping is the central concept used in differential privacy, we recall it below.

**Definition 1 (probabilistic mapping).** *A probabilistic mapping from  $\mathcal{X}$  to  $\mathcal{Y}$  is a mapping  $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$ . Given  $x$ ,  $\mathcal{M}$  outputs a probability measure  $\mathcal{M}(x)$ . To get a numerical output  $y$  out of  $\mathcal{M}$  for  $x$ , one needs to sample  $y \sim \mathcal{M}(x)$ .*

Informally, a probabilistic mapping  $\mathcal{M}$  is said to be differentially private, if given  $x$  and  $x'$  two close inputs (i.e  $d_{\mathcal{X}}(x, x')$  is small enough) it outputs two close measures  $\mathcal{M}(x)$ , and  $\mathcal{M}(x')$ . To evaluate the closeness between this two probability measures in the formal definition of differential privacy, Dwork et. al [5] uses the maximum divergence, which is a special case of the more general Renyi divergence defined as follows:

**Definition 2 (Renyi divergence of order  $\lambda$  [13]).** *Let us consider  $\mu_1, \mu_2 \in \mathcal{P}(\mathcal{Y})$  two probability measures, both dominated by a third measure  $\nu$ . The Renyi divergence of order  $\lambda$  between  $\mu_1$  and  $\mu_2$  writes*

$$D_\lambda(\mu_1, \mu_2) := \frac{1}{\lambda - 1} \log \int_{\mathcal{Y}} g_2(y) \left( \frac{g_1(y)}{g_2(y)} \right)^\lambda d\nu(y).$$

Where  $g_1$  and  $g_2$  are the probability density of  $\mu_1$ , and  $\mu_2$  with respect to  $\nu$ .

The Renyi divergence (see [16] for more details) is defined for  $\lambda \in (1, \infty)$ . It equals the Kullback-Leibler divergence when  $\lambda \rightarrow 1$ , and the maximum divergence (denoted  $D_\infty$ ) when  $\lambda \rightarrow \infty$ . It also has the very special property of being non decreasing with respect to  $\lambda$ . This divergence is very common in machine learning (especially the Kullback-Leibler divergence), statistics, and information theory. Using this notion of closeness between distributions, one can define both differential privacy (with  $D_\infty$ ), and Renyi-differential privacy (with  $D_\lambda$ ).

### 3 Differential privacy and its generalization

We now present the definition of differential privacy, and its Renyi generalization.

**Definition 3 (Classical differential privacy [5]).** *Let  $\mathcal{X}$  be a space of databases,  $\mathcal{Y}$  an output space, and " $\sim_h$ " denoting the that two databases from  $\mathcal{X}$  only differ from one row. A probabilistic mapping  $\mathcal{M}$  from  $\mathcal{X}$  to  $\mathcal{Y}$  is called differentially private if for any  $x, x' \in \mathcal{X}$  s.t.  $x \sim_h x'$  and for any  $Y \in \sigma(\mathcal{Y})$  on has  $\mathcal{M}(x)(Y) \leq \exp(\epsilon) \mathcal{M}(x')(Y)$ .*

**Definition 4 (Metric differential privacy [2]).** *Let  $\epsilon > 0$ ,  $(\mathcal{X}, d_{\mathcal{X}})$  an arbitrary (input) metric space, and  $\mathcal{Y}$  an output space. A probabilistic mapping  $\mathcal{M}$  from  $\mathcal{X}$  to  $\mathcal{Y}$  is called  $(\epsilon, \alpha)$ - $d_{\mathcal{X}}$  private if for any  $x, x'$  s.t  $d_{\mathcal{X}}(x, x') \leq \alpha$ , one has  $D_\infty(\mathcal{M}(x), \mathcal{M}(x')) \leq \epsilon$ .*

Classical differential privacy is a particular case of Metric differential privacy where  $\mathcal{X}$  is a set of tabular databases,  $d_{\mathcal{X}}$  is the hamming distance, and  $\alpha = 1$ <sup>3</sup>. We finally introduce a general form of privacy definition that complies both with classical, and metric differential privacy, namely Renyi-differential privacy.

**Definition 5 (Renyi differential privacy [10]).** *Let  $\epsilon > 0$ ,  $(\mathcal{X}, d_{\mathcal{X}})$  an arbitrary (input) metric space, and  $\mathcal{Y}$  the output space. A probabilistic mapping  $\mathcal{M}$  from  $\mathcal{X}$  to  $\mathcal{Y}$  is called  $(\lambda, \epsilon, \alpha)$ - $d_{\mathcal{X}}$  Renyi-private if for any  $x, x'$  s.t  $d_{\mathcal{X}}(x, x') \leq \alpha$ , one has  $D_\lambda(\mathcal{M}(x), \mathcal{M}(x')) \leq \epsilon$*

<sup>3</sup> Classical definitions set  $\alpha = 1$ , and argue that one can always scale  $d_{\mathcal{X}}$  such that  $d_{\mathcal{X}} \leq 1$  fits the notion of "close enough". We rather keep  $d_{\mathcal{X}}$  unchanged and take an arbitrary  $\alpha$  instead. Both definitions are equivalent.

According to Definition 5, it is clear that both Metric, and differential privacy are included in Renyi-differential privacy. Moreover, note that the definition above is based on arbitrary spaces, and metrics  $(\mathcal{X}, d_{\mathcal{X}}$ , and  $\mathcal{Y})$ . Hence, one can define Renyi-privacy for an arbitrary learning task, even if preserving privacy in this task has no clear semantic. In the following, we present robustness against adversarial examples, and how robustness and privacy are formally similar.

## 4 Robustness to adversarial examples

Let us now consider a classification task over  $\mathcal{X}$  (i.e  $\mathcal{Y} = [N]$ ). Let us denote  $\mathcal{D}$  the ground-truth distribution one tries to learn, and  $h$  the classifier at hand (trained over some subset of  $\mathcal{X} \times \mathcal{Y}$ ). An adversarial example attack for  $x$  is a small perturbation of  $x$  that fools the results of  $h$ . For instance, for image classification, the changes from the initial image to the perturbed one are visually imperceptible, but images are classified with two different labels. The problem of generating an adversarial example from an input  $x$  writes

$$\min d_{\mathcal{X}}(x, x + \tau), \text{ where } \tau \in \mathcal{X}, \text{ and } h(x + \tau) \neq h(x) \quad (1)$$

Even if adversarial examples are intensively studied, a broadly accepted definition of robustness against adversarial attacks does not seem to exist. We settle that the notion of prediction-change risk initially formalized in [4], and implicitly used in e.g [15] is a suitable start-point. Given a classifier  $h$ , it is defined as

$$\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}} [\exists x' \in B(x, \alpha) \text{ s.t } h(x') \neq h(x)].$$

Where  $B(x, \alpha) = \{x' \in \mathcal{X} \text{ s.t } d_{\mathcal{X}}(x, x') \leq \alpha\}$ , and  $\mathcal{D}_{\mathcal{X}}$  is the marginal distribution of  $\mathcal{D}$  with respect to  $\mathcal{X}$ . From this we can derive a definition of robustness to adversarial attacks.

**Definition 6 (Adversarial robustness).** *A classifier  $h$  is said to be  $(\alpha, \gamma)$ -robust if  $\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}} [\exists x' \in B(x, \alpha) \text{ s.t } h(x') \neq h(x)] \leq \gamma$ .*

Regarding [17,11], probabilistic mappings seem to be good candidates to defend against adversarial example attacks. The following definition gives a generalized notion of robustness against adversarial examples attacks complying with probabilistic mappings.

**Definition 7 (Generalized adversarial robustness).** *Let  $D_{\mathcal{P}(\mathcal{Y})}$  be a metric/divergence on  $\mathcal{P}(\mathcal{Y})$ . A randomized classifier  $\mathcal{M}$  is said to be  $D_{\mathcal{P}(\mathcal{Y})}$ - $(\alpha, \epsilon, \gamma)$ -robust if  $\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}} [\exists x' \in B(x, \alpha) \text{ s.t } D_{\mathcal{P}(\mathcal{Y})}(\mathcal{M}(x'), \mathcal{M}(x)) > \epsilon] \leq \gamma$ .*

Definition 7 is fully general, and depends on the metric/divergence  $D_{\mathcal{P}(\mathcal{Y})}$  one chooses to consider. In particular, if one restricts the study of randomized classifiers to Dirac measures, and sets  $D_{\mathcal{P}(\mathcal{Y})}$  to be the trivial distance (which takes the value 0 where the measures are equal and 1 elsewhere), definitions 6 and 7 match. One can refer to [12] for more details on definition 7 and proof on the interest of choosing  $D_{\mathcal{P}(\mathcal{Y})}$  to be the Renyi divergence  $D_{\lambda}$ .

## 5 Links between differential privacy and robustness to adversarial attacks

The starting point to highlight the similarities between both notions are Definitions 5 and 7. A first observation is that in an abstract way (i.e. without instantiating the spaces), and by considering the Renyi divergence both definitions are strictly equivalent. This suggests the following claim, the proof of which is straightforward since it follows from the definitions.

*Claim (Renyi-DP  $\iff$   $D_\lambda$ -robustness).* An algorithm  $\mathcal{M}$  is  $D_\lambda$ - $(\alpha, \epsilon, 0)$ -robust if and only if  $\mathcal{M}$  is  $\mathcal{D}_\mathcal{X}$ -almost surely  $(\lambda, \epsilon, \alpha)$ - $d_\mathcal{X}$  Renyi-differentially private.

While this mathematical equivalence is important from a theoretical point of view, we will now go into deeper details to consider practical implications of this formulation. Without loss of generality, practical settings, in which privacy or robustness are needed, can be classified into three categories:

1. Differential privacy and adversarial robustness need to be ensured: prominent examples of this situation are image or voice classification. In this case, instead of considering two separate methodologies, both problems can be treated simultaneously, with the same tools.

2. Adversarial robustness has to be ensured but there are no special constraint on privacy: in this case, thanks to Claim 5, one could be able to design new defense mechanisms against adversarial examples attacks based on the extensive literature on differential privacy. Accordingly, one can design new defense mechanisms against adversarial example attacks based on the noise injection techniques traditionally used in the differential privacy literature proposed in e.g [12,8,3]. Note that, even though the formal connecting between differential privacy and robustness to adversarial examples is not identically stated in [8,3] and this note, both visions are not conflicting.

3. Differential privacy needs to be ensured but robustness to adversarial examples is not needed: while this setting does not seem intuitively natural, we advocate that a few emerging frameworks currently actively developed to test against adversarial robustness could also be used to evaluate differential privacy with minor adaptations.

Point 2 is currently being investigated. We however argue that the explicit link we just draw between differential privacy and robustness to adversarial examples might lead practitioners from both side to investigate further points 1 and 3.

## References

1. Carlini, N., Athalye, A., Papernot, N., Brendel, W., Rauber, J., Tsipras, D., Goodfellow, I.J., Madry, A., Kurakin, A.: On evaluating adversarial robustness. CoRR abs/1902.06705 (2019), <http://arxiv.org/abs/1902.06705>

2. Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the scope of differential privacy using metrics. In: De Cristofaro, E., Wright, M. (eds.) *Privacy Enhancing Technologies*. pp. 82–102. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
3. Cohen, J.M., Rosenfeld, E., Kolter, J.Z.: Certified adversarial robustness via randomized smoothing. *CoRR* **abs/1902.02918** (2019), <http://arxiv.org/abs/1902.02918>
4. Diochnos, D., Mahloujifar, S., Mahmoody, M.: Adversarial risk and robustness: General definitions and implications for the uniform distribution. In: *Advances in Neural Information Processing Systems*. pp. 10380–10389 (2018)
5. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* **9**(3-4), 211–407 (2013)
6. Dwork, C., Rothblum, G.N.: Concentrated differential privacy. *CoRR* **abs/1603.01887** (2016)
7. ElSalamouny, E., Chatzikokolakis, K., Palamidessi, C.: Generalized Differential Privacy: Regions of Priors That Admit Robust Optimal Mechanisms, pp. 292–318. Springer International Publishing, Cham (2014)
8. Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., Jana, S.: Certified robustness to adversarial examples with differential privacy. *arXiv preprint arXiv:1802.03471* (2018)
9. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. *CoRR* **abs/1706.06083** (2018)
10. Mironov, I.: Rényi differential privacy. In: *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*. pp. 263–275 (2017)
11. Moosavi-Dezfooli, S.M., Shrivastava, A., Tuzel, O.: Divide, denoise, and defend against adversarial attacks. *CoRR* **abs/1802.06806** (2018)
12. Pinot, R., Meunier, L., Araujo, A., Kashima, H., Yger, F., Gouy-Pailler, C., Atif, J.: Theoretical evidence for adversarial robustness through randomization: the case of the exponential family. *CoRR* **abs/1902.01148** (2019)
13. Rényi, A.: On measures of entropy and information. In: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. pp. 547–561. University of California Press, Berkeley, Calif. (1961)
14. Samangouei, P., Kabkab, M., Chellappa, R.: Defense-GAN: Protecting Classifiers Against Adversarial Attacks Using Generative Models. *ArXiv e-prints* (May 2018)
15. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I.J., Fergus, R.: Intriguing properties of neural networks. *CoRR* **abs/1312.6199** (2013)
16. van Erven, T., Harremoës, P.: Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory* **60**(7), 3797–3820 (July 2014)
17. Xie, C., Wang, J., Zhang, Z., Ren, Z., Yuille, A.L.: Mitigating adversarial effects through randomization. *CoRR* **abs/1711.01991** (2017)