
Towards Assessment of Randomized Smoothing Mechanisms for Certifying Adversarial Robustness

Tianhang Zheng,¹* Di Wang,^{2,3} Baochun Li,¹ Jinhui Xu²

¹ Department of Electrical and Computer Engineering
University of Toronto, Toronto, ON

² Department of Computer Science and Engineering
State University of New York at Buffalo, Buffalo, NY

³ King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

Abstract

As a certified defensive technique, randomized smoothing has received considerable attention due to its scalability to large datasets and neural networks. However, several important questions remain unanswered, such as (i) whether the Gaussian mechanism is an appropriate option for certifying ℓ_2 -norm robustness, and (ii) whether there is an appropriate randomized (smoothing) mechanism to certify ℓ_∞ -norm robustness. To shed light on these questions, we argue that the main difficulty is how to assess the appropriateness of each randomized mechanism. In this paper, we propose a generic framework that connects the existing frameworks in [1, 2], to assess randomized mechanisms. Under our framework, for a mechanism which can certify a certain extent of robustness, we define the magnitude (*i.e.*, the expected ℓ_∞ norm) of the random noise it adds as the metric for assessing its appropriateness. We also derive lower bounds on the metric for the ℓ_2 -norm and ℓ_∞ -norm cases as the criteria for assessment. Based on our framework, we assess the Gaussian and Exponential mechanisms by comparing the magnitude of noise added by these mechanisms and the corresponding criteria. We first conclude that the Gaussian mechanism is indeed an appropriate option to certify ℓ_2 -norm robustness. Surprisingly, we show that the Gaussian mechanism is also an appropriate option for certifying ℓ_∞ -norm robustness, instead of the Exponential mechanism. Finally, we generalize our framework to ℓ_p -norm for any $p \geq 2$. Our theoretical results are verified by evaluations on CIFAR10 and ImageNet.

1 Introduction

The past decade has witnessed tremendous success of deep learning in handling various learning tasks like image classification [3], natural language processing [4], and game playing [5]. Nevertheless, a major unresolved issue of deep learning is its vulnerability to adversarial samples, which are almost indistinguishable from natural samples to humans but can mislead deep neural networks (DNNs) to make wrong predictions with high confidence [6, 7]. This phenomenon, referred to as adversarial attack, is considered to be one of the biggest threats to the deployment of many deep learning systems. Thus, a great deal of effort has been devoted to developing defensive techniques against it. However, the majority of the existing defenses are of heuristic nature (*i.e.*, without any theoretical guarantees), implying that they may be ineffective against stronger attacks. Recent work [8, 9, 10] has confirmed this concern by showing that most of those heuristic defenses actually fail to defend strong adaptive attacks. This forces us to shift our attention to certifiable defenses as they can classify all the samples in a predefined neighborhood of the natural samples with a theoretically-guaranteed error bound. Among all the existing certifiable defensive techniques, randomized smoothing is

*The first two authors contributed equally to this work.

becoming increasingly popular due to its scalability to large datasets and arbitrary networks. [1] first relates adversarial robustness to differential privacy, and proves that adding noise is a certifiable defense against adversarial examples. [2] connects adversarial robustness with the concept of Rényi divergence, and improves the estimate on the lower bounds of the robust radii. Recently, [11] successfully certifies 49% accuracy on the original ImageNet dataset under adversarial perturbations with ℓ_2 norm less than 0.5.

Despite these successes, there are still several unanswered questions regarding randomized (smoothing) mechanisms[†]. One such question is, why should we use the Gaussian mechanism for randomized smoothing to certify ℓ_2 -norm robustness, or is there any more appropriate mechanism than the Gaussian mechanism? Another important question is regarding the ability of this method to certify ℓ_∞ -norm robustness. If randomized smoothing can certify ℓ_∞ -norm robustness, what mechanism is an appropriate choice? All these questions motivate us to develop a framework to assess the appropriateness of a randomized mechanism for certifying ℓ_p -norm robustness.

In this paper, we take a promising step towards answering the above questions by proposing a generic and self-contained framework, *which applies to different norms and connects the existing frameworks in [1, 2]*, for assessing randomized mechanisms. Our framework employs the Maximal Relative Rényi (MR) divergence as the probability distance measurement, and thus, the definition of robustness under this measurement is named as D_{MR} robustness. Under our framework, *we define the magnitude* (i.e., *expected ℓ_∞ norm*) *of the noise required by a mechanism to certify a certain extent of robustness as the metric for assessing the appropriateness of the mechanism*. To be specific, a more “appropriate” randomized mechanism under this definition refers to a mechanism that can certify a certain extent of robustness with “less” noise. Given this definition, it is natural to define the assessment criteria as the lower bounds on the magnitude of the noise required to certify ℓ_p -norm (D_{MR}) robustness, in that we can judge whether a mechanism is an appropriate option based on the gap between the magnitude of noise needed by the mechanism and the lower bounds.

Inspired by the theory regarding the lower bounds of one-way marginals in differential privacy models, we prove lower bounds on the noise required for certifying ℓ_2 -norm and ℓ_∞ -norm D_{MR} robustness for arbitrary randomized mechanisms. We demonstrate that the Gaussian mechanism is an appropriate option by showing the gaps between the required Gaussian noise and the lower bounds are only $O(\sqrt{\log d})$ for both ℓ_2 -norm and ℓ_∞ -norm, where d is the dimensionality of the data. Note that this gap is small for datasets like CIFAR-10 and ImageNet, which indicates that the Gaussian mechanism is an appropriate option for certifying ℓ_2 -norm or ℓ_∞ -norm robustness. Moreover, we also show that the Exponential mechanism is not an appropriate option for certifying ℓ_∞ -norm robustness since the gap scales in $O(\sqrt{d})$. To summarize, our contribution is three-fold:

- We propose a generic and self-contained framework for the assessment of randomized mechanisms induced by a new definition of robustness, namely D_{MR} robustness, which connects the existing frameworks such as [1] and [2].
- We define a metric for assessing randomized mechanisms, *i.e.*, the magnitude of the noise required to certify robustness, and we derive the lower bounds on the magnitude of the noise required by arbitrary randomized mechanisms to certify ℓ_2 -norm and ℓ_∞ -norm robustness as the criteria for the assessment. Also, we extend this framework to ℓ_p -norm for any $p \geq 2$.
- We assess the Gaussian mechanism and the Exponential mechanism based on the metric and the lower bounds (*i.e.*, the criteria). We first conclude that the Gaussian mechanism is an appropriate option to certify ℓ_2 -norm robustness. We also show that the Gaussian mechanism is an appropriate option for certifying ℓ_∞ -norm robustness, instead of the Exponential mechanism.

Due to the space limit, all the omitted proofs and some experimental results are included in the appendix in the Supplementary Material.

2 Related Work

To our knowledge, there are three approaches to certify adversarial robustness standing out recently. The first approach formulates the task of adversarial verification as an optimization problem and solves it by tools like convex relaxations and duality [12, 13, 14]. Given a convex set (usually an ℓ_p

[†]In this paper, “randomized mechanism” is an abbreviation for “randomized smoothing mechanism”.

ball) as input, the second approach maintains a convex outer approximation of all the possible outputs at each layer by various techniques, such as interval bound propagation, hybrid zonotope, abstract interpretations, and etc. [15, 16, 17, 18, 19]. The third approach uses randomized smoothing to certify robustness, which is the main focus of this paper. Randomized smoothing for certifying robustness becomes increasingly popular due to its strong scalability to large datasets and arbitrary networks [1, 20, 11, 12, 21]. For this approach, [1] first proves that randomized smoothing can certify the ℓ_2 and ℓ_1 -norm robustness using the differential privacy theory. [20] derives a tighter lower bound on the ℓ_2 -norm robust radius based on a lemma on Rényi divergence. [11] further obtains a tight guarantee on the ℓ_2 -norm robustness using the Neyman-Pearson lemma. [22] proposes a new framework based on f-divergence that applies to different measures. [21] combines [11] with adversarial training, and [23] extends the method in [11] to the top-k classification setting. We note that, compared with [11], the frameworks proposed in [1, 2] are more general. In the following, we briefly review the basic definitions and theorems in the frameworks of [1, 2], which helps us demonstrate the inherent connections between our framework and those two frameworks.

3 Preliminaries

In this section, we first introduce several basic definitions and notations. In general, we denote any randomized mechanism by $\mathcal{M}(\cdot)$, which outputs a random variable depending on the input. We represent any deterministic classifier that outputs a prediction label by $f(\cdot)$. A commonly-used randomized classifier can be constructed by $g(\cdot) = f(\mathcal{M}(\cdot))$. We denote a data sample and its label by \mathbf{x} and y , respectively. An ℓ_p -norm ball centered at \mathbf{x} with radius r is represented by $\mathbb{B}_p(\mathbf{x}, r)$. We say a data sample \mathbf{x}' is in the $\mathbb{B}_p(\mathbf{x}, r)$ iff $\|\mathbf{x}' - \mathbf{x}\|_p \leq r$. Next, we can detail the frameworks in [1] and [2], *i.e.*, PixelDP and Rényi-Divergence-based Bound.

PixelDP To the best of our knowledge, PixelDP [1] is the first framework to prove that randomized smoothing is a certified defense by connecting the concepts of adversarial robustness and differential privacy. The definition of adversarial robustness in the framework of PixelDP can be stated as follows:

Definition 1 (PixelDP [1]) For any \mathbf{x} , and \mathbf{x}' in the $\mathbb{B}_p(\mathbf{x}, r)$, if a randomized mechanism $\mathcal{M}(\cdot)$ satisfies

$$\forall S \subseteq \mathcal{O}, P(\mathcal{M}(\mathbf{x}) \in S) \leq e^\epsilon P(\mathcal{M}(\mathbf{x}') \in S) + \delta, \quad (1)$$

where \mathcal{O} denotes the output space of $\mathcal{M}(\cdot)$. Then we can say $\mathcal{M}(\cdot)$ is (ϵ, δ) -PixelDP.

[1] connects PixelDP with adversarial robustness by the following lemma.

Lemma 1 (Robustness Condition [1]) Suppose $\mathbf{G}(\cdot)$ is randomized K -class classifier defined by $\mathbf{G}(\mathbf{x}) = (\mathbf{G}_1(\mathbf{x}), \dots, \mathbf{G}_K(\mathbf{x}))$ that satisfies (ϵ, δ) -PixelDP (in $\mathbb{B}_p(\mathbf{x}, r)$). For the ground-truth class y , if

$$\mathbb{E}(\mathbf{G}_y(\mathbf{x})) > e^{2\epsilon} \max_{i:i \neq y} \mathbb{E}(\mathbf{G}_i(\mathbf{x})) + (1 + e^\epsilon)\delta, \quad (2)$$

then the classification result $(\mathbb{E}(\mathbf{G}_1(\mathbf{x})), \dots, \mathbb{E}(\mathbf{G}_K(\mathbf{x})))$ is robust in $\mathbb{B}_p(\mathbf{x}, r)$, *i.e.*, $\forall \mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r)$, $\operatorname{argmax}_i \mathbb{E}(\mathbf{G}_i(\mathbf{x}')) = y$.

Note that the definition of the randomized classifier $\mathbf{G}(\cdot)$ is different from the definition of $g(\cdot)$ since the output of $g(\cdot)$ is a scalar not a vector (prediction label). $g(\cdot)$ is more popular in the follow-up works such as [2, 11]. [1] mainly utilizes two mechanisms, *i.e.*, Laplace mechanism and Gaussian mechanism, to guarantee PixelDP. Specifically, adding Laplace noise (*i.e.*, $p(z) = \frac{1}{2b} \exp(-\frac{|z|}{b})$) to the data samples can certify $(\epsilon, 0)$ -PixelDP in $\mathbb{B}_1(\mathbf{x}, b\epsilon)$ for any \mathbf{x} , and adding Gaussian noise (*i.e.*, $p(z) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-\frac{z^2}{2\sigma^2})$) can certify (ϵ, δ) -PixelDP in $\mathbb{B}_2(\mathbf{x}, \frac{\sigma\epsilon}{\sqrt{2 \log 1.25/\delta}})$ for any \mathbf{x} .

Rényi Divergence-based Bound [2] proves a tighter estimate (compared with [1]) on the robust radii based on the following lemma.

Lemma 2 (Rényi Divergence Lemma [2]) Let $P = (p_1, p_2, \dots, p_k)$ and $Q = (q_1, q_2, \dots, q_k)$ be two multinomial distributions. If the indices of the largest probabilities **do not** match on P and Q , then

the Rényi divergence between P and Q , i.e., $D_\alpha(P\|Q)$,[‡] satisfies

$$D_\alpha(P\|Q) \geq -\log(1 - p_{(1)} - p_{(2)}) + 2\left(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha})\right)^{\frac{1}{1-\alpha}}.$$

where $p_{(1)}$ and $p_{(2)}$ refer to the largest and the second largest probabilities in $\{p_i\}$, respectively.

If the Gaussian mechanism is applied to certify ℓ_2 -norm robustness, then we have the following bound of the robust radii.

Lemma 3 ([2]) *Let f be any deterministic classifier and $g(\mathbf{x}) = f(\mathbf{x} + \mathbf{z})$ be its corresponding randomized classifier for samples $\mathbf{x} \in \mathbb{R}^d$, where $\mathbf{z} \sim \mathcal{N}(0, \sigma^2 I_d)$. Then $\forall \mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r)$, $\operatorname{argmax}_y P(g(\mathbf{x}) = y) = \operatorname{argmax}_{y'} P(g(\mathbf{x}') = y')$, i.e., $g(\cdot)$ is robust in $\mathbb{B}_p(\mathbf{x}, r)$, and the ℓ_2 robust radii r that could be certified is given by*

$$r^2 \leq \sup_{\alpha > 1} -\frac{2\sigma^2}{\alpha} \log(1 - p_{(1)} - p_{(2)}) + 2\left(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha})\right)^{\frac{1}{1-\alpha}}. \quad (3)$$

$p_{(1)}$ and $p_{(2)}$ refer to the largest and the second largest probabilities in $\{p_i\}$, where p_i is the probability that $g(\mathbf{x})$ returns the i -th class, i.e., $p_i = P(g(\mathbf{x}) = i)$.

4 Framework Overview

In this section, we present a generic framework based on the Definition 2, 3, and 4, for assessing randomized mechanisms. According to Definition 3, our framework applies to any ℓ_p -norm. Moreover, we show that our proposed framework connects the existing general frameworks in [1, 2] by Theorem 4.1 & 4.2. Also, we note that it is difficult to involve the framework in [11] since [11] restricts the additive noise of the randomized mechanism to be isotropic such as Gaussian noise, while in our framework, we do not need to specify the type of additive noise.

4.1 Main Definitions

Under our framework, the definition of adversarial robustness is induced by maximal relative Rényi divergence (MR divergence), namely D_{MR} robustness, so we start from introducing the definition of MR divergence.

Definition 2 (Maximal Relative Rényi Divergence) *The Maximal Relative Rényi Divergence $D_{MR}(P\|Q)$ of distributions P and Q is defined as*

$$D_{MR}(P\|Q) = \max_{\alpha \in (1, \infty)} \frac{D_\alpha(P\|Q)}{\alpha}, \quad (4)$$

where $D_\alpha(P\|Q)$ is the Rényi Divergence between P and Q . Using D_{MR} as the probability measure, we can define adversarial robustness as follows:

Definition 3 (D_{MR} Robustness) *We say a randomized (smoothing) mechanism $\mathcal{M}(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust if for any \mathbf{x}*

$$\forall \mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r), \max\{D_{MR}(\mathcal{M}(\mathbf{x}), \mathcal{M}(\mathbf{x}')), D_{MR}(\mathcal{M}(\mathbf{x}'), \mathcal{M}(\mathbf{x}))\} \leq \epsilon. \quad (5)$$

If a randomized smoothing classifier $g(\cdot)$ satisfies the above equation, we can say it is a $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust classifier or it certifies $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robustness.

A property of D_{MR} robustness we use throughout this paper is its postprocessing property, which can be stated as follows:

Corollary 1 (Postprocessing Property) *Let $g(\mathbf{x}) = f(\mathcal{M}(\mathbf{x}))$ be a randomized classifier, where $f(\cdot)$ is any deterministic function (classifier). $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust if $\mathcal{M}(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust.*

This postprocessing property can be easily proved by $D_\alpha(f(\mathcal{M}(\mathbf{x}))\|f(\mathcal{M}(\mathbf{x}')))) \leq D_\alpha(\mathcal{M}(\mathbf{x})\|\mathcal{M}(\mathbf{x}'))$ for any $\alpha \in (1, \infty)$ [24]. This property allows us to only concentrate on the randomized mechanism $\mathcal{M}(\cdot)$ without considering the specific form of the deterministic classifier $f(\cdot)$, and therefore makes the framework applicable to an arbitrary neural network.

[‡]For $\alpha \in (1, \infty)$, $D_\alpha(P\|Q)$ is defined as $D_\alpha(P\|Q) = \frac{1}{\alpha-1} \log \mathbb{E}_{x \sim Q} \left(\frac{P(x)}{Q(x)}\right)^\alpha$.

4.2 Connections between D_{MR} robustness and the existing frameworks

The framework defined by Definition 2 & 3 is generic since it is closely connected with the existing ones [1, 2]. Here we demonstrate the connections by the following two theorems.

Theorem 4.1 (D_{MR} Robustness & PixelDP) *If $\mathcal{M}(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust, then $\mathcal{M}(\cdot)$ is also $(\epsilon + 2\sqrt{\log(1/\delta)}\epsilon, \delta)$ -PixelDP.*

We note that the opposite of Theorem 4.1 holds only when $\delta = 0$, which indicates our framework is a relaxed version of the PixelDP framework. But this should not be a surprise since most of the following frameworks [2, 11, 22] can somehow be considered more relaxed than the PixelDP framework and thus yield tighter certified bounds. Similarly, our framework can provide the same bound on the robust radius as in [2], which is tighter than the bound in [1] (Theorem 4.2).

Theorem 4.2 (D_{MR} Robustness & Rényi Divergence-based Bound) *If a randomized classifier $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust, then we have $\forall \mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r)$, $\operatorname{argmax}_y P(g(\mathbf{x}) = y) = \operatorname{argmax}_{y'} P(g(\mathbf{x}') = y')$ as long as*

$$\epsilon \leq \sup_{\alpha > 1} -\frac{1}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}}), \quad (6)$$

where $p_{(1)}$ and $p_{(2)}$ also refer to the largest and the second largest probabilities in $\{p_i\}$, and p_i is the probability that $g(\mathbf{x})$ returns the i -th class, i.e., $p_i = P(g(\mathbf{x}) = i)$. Based on the above theorem, we can derive the same ℓ_2 robust radius as in Lemma 3 [2]. We will detail how to derive the ℓ_2 robust radius in Section 5.

An interpretation of Theorem 4.1 and 4.2 is that, as long as we can use a randomized mechanism with a certain amount of noise to certify D_{MR} robustness, we can use the same mechanism with the same amount of noise to certify PixelDP and the Rényi Divergence-based bound. Thus, Theorem 4.1 and 4.2 indicate the assessment results based on the metric defined in Section 4.3 is very likely to generalize to the other frameworks.

4.3 Assessment of Randomized Mechanisms

Since there are infinite randomized mechanisms, a natural problem is to determine whether a certain randomized mechanism is an appropriate option to certify adversarial robustness. However, we note that all the previous work [2, 11, 21] overlook this problem and assume the Gaussian mechanism to be an appropriate mechanism for certifying ℓ_2 -norm robustness without sufficient assessment. While in this paper, we attempt to provide a solution to this problem under our proposed framework. Specifically, we define a metric to assess randomized mechanisms as follows:

Definition 4 *Specify a p -norm, a robust radius r , and an epsilon ϵ , the magnitude (expected ℓ_∞ -norm) of the noise required by a randomized mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ to certify $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robustness is defined as the metric to assess the appropriateness of $\mathcal{M}(\cdot)$.*

We define this metric for assessing randomized mechanisms because *the accuracy of neural networks tends to decrease as the magnitude of the noise added to the inputs increases*. Note that if the magnitude of the noise required for a randomized classifier is too large, the accuracy of its predictions on clean samples will be very low, then robustness will be useless[§]. Given the above metric, we also need criteria to assess the (relative) appropriateness of a randomized mechanism. In this paper, we employ the lower bounds on the magnitude of the noise required by any randomized mechanism to certify certain $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robustness as the criteria. *We consider a randomized mechanism as an appropriate option if the gap between the magnitude of the additive noise required by this mechanism and the corresponding lower bound is small*. In the following, we will provide the lower bounds for ℓ_2 -norm and ℓ_∞ -norm, i.e., *the two most popular norms*, and assess the appropriateness of the Gaussian and Exponential mechanisms for certifying ℓ_2 -norm and ℓ_∞ -norm robustness. In Appendix, we generalize our framework to ℓ_p -norm for any $p \geq 2$.

5 Assessing Mechanisms for Certifying ℓ_2 -norm Robustness

In this section, we first elaborate on how the Gaussian mechanism certifies D_{MR} robustness, and then provide the lower bound on the magnitude of the additive noise required by any randomized

[§]Certified robustness only guarantees the predictions of the perturbed samples and the predictions of their clean samples are the same.

mechanism ($\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$) to certify ℓ_2 -norm robustness. By comparing the magnitude of the additive noise required by the Gaussian mechanism with the lower bound, we conclude that the Gaussian mechanism is an appropriate option to certify ℓ_2 -norm robustness.

Theorem 5.1 (Gaussian Mechanism for Certifying ℓ_2 -norm robustness) *Let f be any deterministic classifier and $g(\mathbf{x}) = f(\mathcal{M}(\mathbf{x}))$ be its corresponding randomized classifier for the sample $\mathbf{x} \in \mathbb{R}^d$, where $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \sigma^2 I_d)$. Then, $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_2, \frac{r^2}{2\sigma^2})$ -robust.*

According to Theorem 4.2, if we substitute ϵ with $\frac{r^2}{2\sigma^2}$, r can be given by $r^2 \leq \sup_{\alpha > 1} -\frac{2\sigma^2}{\alpha} \log(1 - p_1 - p_2 + 2(\frac{1}{2}(p_1^{1-\alpha} + p_2^{1-\alpha}))^{\frac{1}{1-\alpha}})$, which is same as the bound in [2] (Lemma 3). To provide a criterion for assessment of randomized mechanisms in the ℓ_2 -norm case, we prove a lower bound on the magnitude of the noise \mathbf{z} required by any randomized mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ to ensure that $\mathcal{M}(\mathbf{x})$ (as well as $f(\mathcal{M}(\mathbf{x}))$) is $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust. As mentioned in Section 4.3, if the magnitude of the additive Gaussian noise is close to the lower bound, then the Gaussian mechanism is considered as an appropriate option. The lower bound is given by the following theorem.

Theorem 5.2 (ℓ_2 -norm Criterion for Assessment) *For any $\epsilon \leq O(1)$, if there is an $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : \mathbb{R}^d \mapsto \mathbb{R}^d$ such that,*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}_{\mathcal{M}}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \leq \alpha, \quad (7)$$

for some $\alpha \leq O(1)$, then it must be true that $\alpha \geq \Omega(\frac{r}{\sqrt{\epsilon}})$. In another word, $\Omega(\frac{r}{\sqrt{\epsilon}})$ is the lower bound of the (expected) magnitude of the random noise, i.e., the ℓ_2 -norm criterion.

Note that proving this theorem on $\mathbb{R}^d \mapsto \mathbb{R}^d$ is non-trivial, which is detailed in Appendix. Theorem 5.2 indicates that the (expected) magnitude of the additive noise should be at least $\Omega(\frac{r}{\sqrt{\epsilon}})$ to certify $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robustness. For the Gaussian mechanism, the required magnitude of the additive noise is $O(\sigma\sqrt{\log d})$ according to [25], which is $O(\frac{r}{\sqrt{\epsilon}}\sqrt{\log d})$ to guarantee $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robustness, according to Theorem 5.1. This means that the gap between the Gaussian mechanism and the potential optimal mechanism is bounded by $O(\sqrt{\log d})$.

Remark 1 *We say Gaussian mechanism is an appropriate option because $O(\sqrt{\log d})$ is small for most commonly-used datasets. For instance, for CIFAR-10 ($d = 3072$), $\sqrt{\log_e d} \approx 2.83$, and for ImageNet ($d = 150528$), $\sqrt{\log_e d} \approx 3.45$.*

Equivalently, if we fix the expected ℓ_∞ -norm of the additive noise as α , the largest radius r that can be certified by any $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized mechanisms is upper bounded by $O(\alpha\sqrt{\epsilon})$, and the robust radius certified by Gaussian mechanism is $O(\frac{\alpha\sqrt{\epsilon}}{\sqrt{\log d}})$. The upper bound $O(\alpha\sqrt{\epsilon})$ can be simply derived by transforming $\alpha \geq \Omega(\frac{r}{\sqrt{\epsilon}})$ in Theorem 5.2. For the Gaussian mechanism, since $\alpha = O(\sigma\sqrt{\log d})$ and $\epsilon = \frac{r^2}{2\sigma^2}$ according to Theorem 5.1, $r = O(\frac{\alpha\sqrt{\epsilon}}{\sqrt{\log d}})$ [¶].

6 Assessing Mechanisms for Certifying ℓ_∞ -norm Robustness

In this section, we first discuss the possibility of using the Exponential mechanism, an analogue of the Gaussian mechanism in the ℓ_∞ -norm case, to certify ℓ_∞ -norm robustness. Then, we prove the lower bound on the magnitude of the additive noise required by any randomized mechanism to certify ℓ_∞ -norm robustness. By comparing the magnitude of the noise required by the Exponential mechanism with the lower bound, we conclude that the Exponential mechanism is not an appropriate option to certify ℓ_∞ -norm robustness. Surprisingly, we find that the Gaussian mechanism is a more appropriate option than the Exponential mechanism to certify ℓ_∞ -norm robustness.

We first recall the form of the density function of Gaussian noise: $p(\mathbf{z}) \propto \exp(-\frac{\|\mathbf{z}\|_2^2}{\sigma^2})$. Based on this, we conjecture that, to certify ℓ_∞ -norm robustness, we can sample the noise using the Exponential mechanism, an analogue of the Gaussian mechanism in the ℓ_∞ -norm case:

$$p(\mathbf{z}) \propto \exp(-\frac{\|\mathbf{z}\|_\infty}{\sigma}). \quad (8)$$

[¶]The theoretical results of the robust radii are verified by experiments.

We show in the following theorem that randomized smoothing using the Exponential mechanism can certify $(r, D_{MR}, \|\cdot\|_\infty, \frac{r^2}{2\sigma^2})$ -robustness, which is seemingly an extension of the ℓ_2 -norm case. However, the certified radius r is $O(1/d)$, which implies that it is unscalable to high-dimensional data, *i.e.*, The Exponential mechanism should not be an appropriate mechanism to certify ℓ_∞ -norm robustness. This conclusion is further verified by our assessment method, which will be detailed later.

Theorem 6.1 (Exponential Mechanism for Certifying ℓ_∞ -norm Robustness) *Let f be any deterministic classifier and $g(\mathbf{x}) = f(\mathcal{M}(\mathbf{x}))$ be its corresponding randomized classifier for sample $\mathbf{x} \in \mathbb{R}^d$, where $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with \mathbf{z} sampled from Exponential mechanism. Then, $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_\infty, \frac{r}{\sigma})$ -robust and also $(r, D_{MR}, \|\cdot\|_\infty, \frac{r^2}{2\sigma^2})$ -robust.*

According to Theorem 4.2, if we substitute ϵ with $\frac{r}{\sigma}$ or $\frac{r^2}{2\sigma^2}$, then r can be given by $r \leq \sup_{\alpha>1} -\frac{\sigma}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}})$, or $r^2 \leq \sup_{\alpha>1} -\frac{2\sigma^2}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}})$. Comparing this result and Theorem 5.1, we can see that randomized smoothing via the Exponential mechanism certifies a region with (almost) the same radius as that certified by the Gaussian mechanism in the ℓ_2 -norm case, indicating similarity in their robustness guarantees. However, the following corollary shows that the magnitude of the noise required by the Exponential mechanism is much larger than that of the Gaussian mechanism in the ℓ_2 -norm case.

Corollary 2 *For the Exponential mechanism that can guarantee Theorem 6.1, $\mathbb{E}_{\mathbf{z}}[\|\mathbf{z}\|_\infty] = d\sigma$.*

Equivalently, if we fix the expected ℓ_∞ -norm of the additive noise as α , according to Theorem 6.1, the robust radius r certified by Exponential mechanism is $\max\{O(\frac{\epsilon\alpha}{d}), O(\frac{\sqrt{2\epsilon\alpha}}{d})\}$ [¶], which is very small for high-dimensional datasets. The following theorem further shows that there is a huge gap between the noise required by the Exponential mechanism and the lower bound, indicating that the Exponential mechanism is probably not an appropriate option for certifying ℓ_∞ -norm robustness.

Theorem 6.2 (ℓ_∞ -norm Criterion for Assessment) *For $\epsilon \leq O(1)$, if there is an $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robust mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : \mathbb{R}^d \mapsto \mathbb{R}^d$ that satisfies*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}_{\mathcal{M}}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \leq \alpha,$$

*it must be true that $\alpha \geq \Omega(\frac{r\sqrt{d}}{\sqrt{\epsilon}})$. In another word, $\Omega(\frac{r\sqrt{d}}{\sqrt{\epsilon}})$ is the lower bound of the expected magnitude of the required noise, *i.e.*, the criterion.*

According to Corollary 2 and Theorem 6.1, for Exponential mechanism, its required magnitude of noise is $O(\frac{dr}{\sqrt{\epsilon}})$ to guarantee $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robustness. Compared with Theorem 6.2, we can see that the gap between the magnitude of the noise required by the Exponential mechanism and the lower bound is $O(\sqrt{d})$, which can be very large for high-dimensional datasets. Therefore, we can conclude that the Exponential mechanism is probably not an appropriate mechanism for certifying ℓ_∞ -norm robustness. Surprisingly, the following theorem shows that the Gaussian mechanism is an appropriate choice for certifying $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robustness.

Theorem 6.3 (Gaussian Mechanism for Certifying ℓ_∞ -norm robustness) *Let $r, \epsilon > 0$ be some fixed number and $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \frac{dr^2}{2\epsilon} I_d)$. Then, $\mathcal{M}(\cdot)$ is $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robust, and $\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}_{\mathcal{M}}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty$ is upper bounded by $O(\frac{r\sqrt{d \log d}}{\sqrt{\epsilon}})$.*

From Theorem 6.2 and 6.3, we can see that the gap between the magnitude of the noise required by the Gaussian mechanism and the lower bound is also $O(\sqrt{\log d})$. Thus, we can say the Gaussian mechanism is a more appropriate option to certify ℓ_∞ -norm robustness (see Remark 1). Equivalently, if we fix the expected ℓ_∞ -norm of the additive noise as α , the robust radius is $O(\frac{\sqrt{\epsilon\alpha}}{\sqrt{d \log d}})$ [¶].

Remark 2 *Note that in the previous sections we only consider ℓ_2 -norm and ℓ_∞ -norm robustness and the corresponding mechanisms because ℓ_2 -norm and ℓ_∞ -norm are the two most important norms. But actually, we can extend our framework to ℓ_p -norm for any $p \geq 2$. See Section D in Appendix.*

7 Experiments

Datasets and Models Our theories are verified on two widely-used datasets, *i.e.*, CIFAR10 and ImageNet^{||}. We follow [11, 21] to use a 110-layer residual network and a ResNet-50 as the base models for CIFAR10 and ImageNet. The certified accuracy for radius R is defined as the fraction of the test set whose certified radii are larger than R , and predictions are correct. We note that the lower bounds (criteria) are not verifiable by experiments since we are still not sure if there exist any randomized mechanism that can achieve those lower bounds. *So in the following, we mainly verify the theoretical results regarding the Gaussian mechanism and the Exponential mechanism. We provide more details about the numerical method and more experimental results compared to the other frameworks in Appendix in the supplementary material.*

Empirical Results In the following, we verify our framework by comparing our theoretical results of the ℓ_2/ℓ_∞ robust radii with the ℓ_2/ℓ_∞ radii at which the Gaussian/Exponential mechanism can certify 40 ~ 60% accuracy in the experiments. Note that in the previous literature, 40 ~ 60% robust accuracy is considered as a reasonably good performance [26, 11]. Besides, selecting another reasonable accuracy does not affect the verification results too much because what our theories characterize are the asymptotic behaviors rather than the exact values of the robust radii.

In Fig. 1, we demonstrate the results of the Gaussian mechanism for certifying ℓ_2 -norm robustness. The red dashed lines show that the Gaussian mechanism can certify 40 ~ 60% accuracy at ℓ_2 radius = 0.34 (CIFAR-10, $d = 3072$) and ℓ_2 radius = 0.29 (ImageNet, $d = 150568$), *i.e.*, approximately $1/\sqrt{\log d}$. These results verify that the ℓ_2 radius certified by the Gaussian mechanism is $O(\frac{\alpha\sqrt{\epsilon}}{\sqrt{\log d}})^{**}$. We also argue that, $O(\frac{\alpha\sqrt{\epsilon}}{\sqrt{\log d}})$ is the scale of the largest certified ℓ_2 radius (*i.e.*, $\frac{\sigma}{2}(\Phi^{-1}(p_{(1)}) - \Phi^{-1}(p_{(2)}))$) in the previous literature since the ℓ_∞ -norm of the Gaussian noise α is $O(\sigma\sqrt{\log d})$. This argument is verified by Fig. 3 & 4 in Appendix.

Fig. 2 shows that the Gaussian mechanism certifies 40 ~ 60% accuracy at ℓ_∞ radius = $6e - 3$ on CIFAR-10 ($d = 3072$) and ℓ_∞ radius = $1.1e - 3$ on ImageNet ($d = 150568$), *i.e.*, approximately $O(1/\sqrt{d\log d})$. These results verify that the ℓ_∞ radius certified by the Gaussian mechanism is $O(\frac{\sqrt{\epsilon}\alpha}{\sqrt{d\log d}})$. Fig. 2 also shows that the Exponential mechanism certifies approximately 40 ~ 60% accuracy at ℓ_∞ radius = $1.5e - 4$ on CIFAR-10 and ℓ_∞ radius = $7e - 6$ on ImageNet, *i.e.*, approximately $O(1/d)$. These results verify that the ℓ_∞ robust radius certified by the Exponential mechanism scales in $O(\frac{\alpha\sqrt{\epsilon}}{d})$ or $O(\frac{\sqrt{2\epsilon}\alpha}{d})$. If we compare the performance of the Gaussian mechanism and the Exponential mechanism in Fig. 2, we can see that the Gaussian mechanism is a much more appropriate option than for certifying ℓ_∞ -norm robustness. It is worth noting that the performance of the Gaussian mechanism can be better with the bound proved in [11], which is comparable to the other state-of-the-art approaches introduced in Section 2. We detail some results regarding the comparison in Appendix.

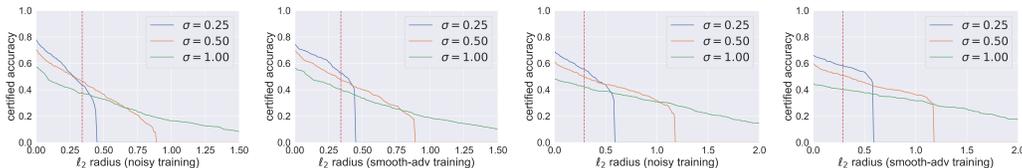


Figure 1: Certify ℓ_2 -norm robustness by the Gaussian mechanism: certified accuracy for CIFAR-10 (left two) and ImageNet (right two). Models: noisy training [11] & smooth-adv training [21].

8 Conclusion

In this paper, we present a generic and self-contained framework, which connects the existing frameworks such as [1, 2], for assessing randomized mechanisms. Under our framework, we define the magnitude of the noise required by a randomized mechanism to certify a certain extent of robustness as the metric for assessing this mechanism. We also provide the lower bounds on the magnitudes of the required noise as the assessment criteria. Comparing the noise required by the Gaussian and Exponential mechanism and the criteria, we conclude that (i) The Gaussian mechanism

^{||} Pixel value range is [0.0, 1.0]
^{**} $\alpha \leq O(1)$, and $\epsilon \leq O(1)$.

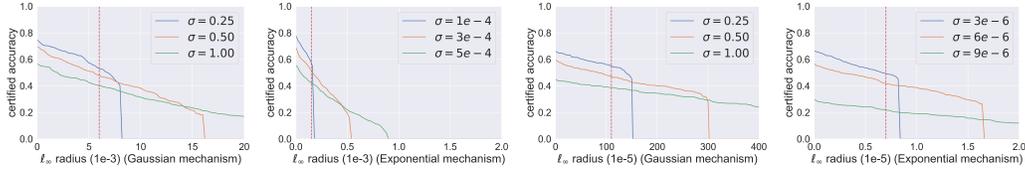


Figure 2: Certify ℓ_∞ -norm robustness by the Gaussian/Exponential mechanism: certified accuracy for CIFAR-10 (left two) and ImageNet (right two). Model: smooth-adv training [21].

is an appropriate option to certify ℓ_2 -norm and ℓ_∞ -norm robustness. (ii) The Exponential mechanism is not an appropriate mechanism to certify ℓ_∞ -norm robustness, although it seems an analogue of the Gaussian mechanism in the ℓ_∞ -norm case.

Broader Impact Statement

Since this work is a study on the core theory & framework to understand the randomized smoothing method and model robustness, we do not figure out any *substantial* ethic issue regarding this work. However, it might lead to several societal consequences. On the positive side, it can help the industry to assess a randomized smoothing mechanism before using them in a deep learning model for improving model robustness. Besides, it can motivate researchers to understand the ability and limitations of the randomized smoothing method. On the negative side, although the randomized smoothing method achieves state-of-the-art performance among certified defensive techniques, some of our theoretical results (*e.g.*, the lower bounds of the additive noise to certify robustness) indicate its limitations. These limitations might give the potential adversaries more confidence to attack deep learning models with adversarial examples.

References

- [1] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. *arXiv preprint arXiv:1802.03471*, 2018.
- [2] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Certified adversarial robustness with additive noise. In *Advances in Neural Information Processing Systems*, pages 9459–9469, 2019.
- [3] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [4] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [5] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [6] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [7] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [8] Warren He, James Wei, Xinyun Chen, Nicholas Carlini, and Dawn Song. Adversarial example defense: Ensembles of weak defenses are not strong. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, 2017.
- [9] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.

- [10] Jonathan Uesato, Brendan O’Donoghue, Pushmeet Kohli, and Aaron Oord. Adversarial risk and the dangers of evaluating against weak attacks. In *International Conference on Machine Learning*, pages 5032–5041, 2018.
- [11] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. *arXiv preprint arXiv:1902.02918*, 2019.
- [12] Krishnamurthy Dvijotham, Robert Stanforth, Sven Gowal, Timothy A Mann, and Pushmeet Kohli. A dual approach to scalable verification of deep networks. In *UAI*, pages 550–559, 2018.
- [13] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018.
- [14] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pages 5283–5292, 2018.
- [15] Matthew Mirman, Timon Gehr, and Martin Vechev. Differentiable abstract interpretation for provably robust neural networks. In *International Conference on Machine Learning*, pages 3575–3583, 2018.
- [16] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Efficient formal safety analysis of neural networks. In *Advances in Neural Information Processing Systems*, pages 6367–6377, 2018.
- [17] Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Timothy Mann, and Pushmeet Kohli. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*, 2018.
- [18] Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Duane Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*, 2019.
- [19] Mislav Balunovic and Martin Vechev. Adversarial training and provable defenses: Bridging the gap. In *International Conference on Learning Representations*, 2020.
- [20] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Second-order adversarial attack and certifiable robustness. *arXiv preprint arXiv:1809.03113*, 2018.
- [21] Hadi Salman, Greg Yang, Jerry Li, Pengchuan Zhang, Huan Zhang, Ilya Razenshteyn, and Sebastien Bubeck. Provably robust deep learning via adversarially trained smoothed classifiers. *arXiv preprint arXiv:1906.04584*, 2019.
- [22] KD Dvijotham, J Hayes, B Balle, Z Kolter, C Qin, A Gyorgy, K Xiao, S Gowal, and P Kohli. A framework for robustness certification of smoothed classifiers using f-divergences. In *International Conference on Learning Representations*, 2020.
- [23] Jinyuan Jia, Xiaoyu Cao, Binghui Wang, and Neil Zhenqiang Gong. Certified robustness for top-k predictions against adversarial perturbations via randomized smoothing. *arXiv preprint arXiv:1912.09899*, 2019.
- [24] Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- [25] Francesco Orabona and Dávid Pál. Optimal non-asymptotic lower bound on the minimax regret of learning with expert advice. *arXiv preprint arXiv:1511.02176*, 2015.
- [26] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [27] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [28] Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- [29] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), 2016.
- [30] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.

- [31] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *arXiv preprint arXiv:1501.06095*, 2015.

To make the paper more readable, we first review some definitions about differential privacy [27].

Definition 5 Given a data universe X , we say that two datasets $D, D' \subset X$ are neighbors if they differ by only one entry, which is denoted by $D \sim D'$. A randomized algorithm \mathcal{M} is (ϵ, δ) -differentially private (DP) if for all neighboring datasets D, D' and all events S the following holds

$$P(\mathcal{M}(D) \in S) \leq e^\epsilon P(\mathcal{M}(D') \in S) + \delta.$$

Definition 6 A randomized algorithm \mathcal{M} is (α, ϵ) -Rényi differentially private (DP) if for all neighboring datasets D, D' the following holds

$$D_\alpha(\mathcal{M}(D) \parallel \mathcal{M}(D')) \leq \epsilon.$$

A Omitted Proofs in Section 4

Proof [Proof of Theorem 4.1] According to Definition 2, for a fixed \mathbf{x} , we have $\forall \mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r)$ and any $\alpha \in (1, \infty)$,

$$D_\alpha(\mathcal{M}(\mathbf{x}) \parallel \mathcal{M}(\mathbf{x}')) < \alpha\epsilon.$$

Therefore, $\mathcal{M}(\cdot)$ satisfies $(\alpha, \alpha\epsilon)$ -Rényi DP. According to the following lemma, i.e.,

Lemma 4 ([28]) If a randomized mechanism is (α, ϵ) -Rényi DP, then it is $(\epsilon + \frac{\log(1/\delta)}{\alpha-1}, \delta)$ -DP for any $\delta > 0$,

we have $\mathcal{M}(\cdot)$ is $(\alpha\epsilon + \frac{\log(1/\delta)}{1-\alpha}, \delta)$ -DP, for all $\alpha \in (1, +\infty)$. Since

$$\begin{aligned} \min_{\alpha \in (1, +\infty)} \left\{ \alpha\epsilon + \frac{\log(1/\delta)}{\alpha-1} \right\} &\stackrel{\beta=\alpha-1}{=} \min_{\beta \in (0, \infty)} \left\{ \epsilon(1 + \beta + \frac{\log(1/\delta)}{\epsilon\beta}) \right\} \\ &= \epsilon + 2\sqrt{\log(1/\delta)\epsilon}. \end{aligned}$$

Thus, by the definition of Approximate Differential Privacy (Definition 5), in total we have for any \mathbf{x} , $\mathbf{x}' \in \mathbb{B}_p(\mathbf{x}, r)$, and any event S

$$P(\mathcal{M}(\mathbf{x}') \in S) \leq e^{\epsilon + 2\sqrt{\log(1/\delta)\epsilon}} P(\mathcal{M}(\mathbf{x}) \in S) + \delta.$$

Thus by Definition 1, $\mathcal{M}(\cdot)$ is $(\epsilon + 2\sqrt{\log(1/\delta)\epsilon}, \delta)$ Pixel-DP. ■

Proof [Proof of Theorem 4.2] Recall that Lemma 2 indicates that we have $\operatorname{argmax}_y P(g(\mathbf{x}) = y) = \operatorname{argmax}_{y'} P(g(\mathbf{x}') = y')$ as long as

$$D_\alpha(g(\mathbf{x}) \parallel g(\mathbf{x}')) < \sup_{\alpha > 1} -\log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}}).$$

Thus we just need to prove that the above condition holds. Since $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust, for any \mathbf{x} and $\|\mathbf{x}' - \mathbf{x}\|_p \leq r$ we have

$$D_\alpha(g(\mathbf{x}) \parallel g(\mathbf{x}')) < \alpha\epsilon.$$

If we also have the additional condition:

$$\epsilon \leq \sup_{\alpha > 1} -\frac{1}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}}),$$

then $D_\alpha(g(\mathbf{x}) \parallel g(\mathbf{x}')) < \sup_{\alpha > 1} -\log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}})$. Thus, the additional condition to guarantee $\operatorname{argmax}_y P(g(\mathbf{x}) = y) = \operatorname{argmax}_{y'} P(g(\mathbf{x}') = y')$ can be stated as $\epsilon \leq \sup_{\alpha > 1} -\frac{1}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}})$. ■

B Omitted Proofs in Section 5

Proof [Proof of Theorem 5.1] By the postprocessing property we just need to show $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ is $(r, D_{MR}, \|\cdot\|_2, \frac{r^2}{2\sigma^2})$ robust.

Fix any x , we have for any $\mathbf{x}' \in \mathbb{B}_2(\mathbf{x}, r)$ and $\alpha \in (1, \infty)$

$$\begin{aligned} D_\alpha(\mathcal{M}(\mathbf{x})\|\mathcal{M}(\mathbf{x}')) &= D_\alpha(\mathcal{N}(\mathbf{x}, \sigma^2 I_d)\|\mathcal{N}(\mathbf{x}', \sigma^2 I_d)) \\ &= \frac{\alpha \|\mathbf{x}' - \mathbf{x}\|_2^2}{2\sigma^2} \leq \frac{r^2}{2\sigma^2}. \end{aligned}$$

■

Proof [Proof of Theorem 5.2] We first show that, in order to prove Theorem 5.2, we only need to prove Theorem B.1. Then we show that, to prove Theorem B.1, we only need to prove Theorem B.2. Finally, we give a formal proof of Theorem B.2.

Theorem B.1 *For some $\epsilon = O(1)$ and $d \leq O(r^2)$, if there is a $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ randomized (smoothing) mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : [0, \frac{r}{2\sqrt{d}}]^d \mapsto \mathbb{R}^d$ such that for any $\mathbf{x} \in [0, \frac{r}{2\sqrt{d}}]^d$,*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \alpha$$

for some constant $\alpha = O(1)$. Then it must be true that $\alpha \geq \Omega(\frac{r}{\sqrt{\epsilon}})$.

For any $\mathcal{M}(\mathbf{x}) : \mathbb{R}^d \mapsto \mathbb{R}^d$, in Theorem B.1, we only consider the expected ℓ_∞ -norm of the noise added by $\mathcal{M}(\mathbf{x})$ on $\mathbf{x} \in [0, \frac{r}{2\sqrt{d}}]^d$. Thus, the α in Theorem B.1 should be less than or equal to the α in Theorem 5.2 (on $\mathbf{x} \in \mathbb{R}^d$). Therefore, the lower bound for the α in Theorem B.1 (i.e., $\Omega(\frac{r}{\sqrt{\epsilon}})$) is also a lower bound for the α in Theorem 5.2. That is to say, if Theorem B.1 holds, then Theorem 5.2 also holds true.

Next, we show that if Theorem B.2 holds, then Theorem B.1 also holds.

Theorem B.2 *For some $\epsilon = O(1)$ and $d \leq O(r^2)$, if there is a $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized (smoothing) mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : [0, \frac{r}{2\sqrt{d}}]^d \mapsto [0, \frac{r}{2\sqrt{d}}]^d$ such that for any $\mathbf{x} \in [0, \frac{r}{2\sqrt{d}}]^d$*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \alpha$$

for some constant $\alpha = O(1)$. Then it must be true that $\alpha \geq \Omega(\frac{r}{\sqrt{\epsilon}})$.

For any $\mathcal{M}(\mathbf{x}) : [0, \frac{r}{2\sqrt{d}}]^d \mapsto \mathbb{R}^d$ considered in Theorem B.1, there exists a $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized mechanism $\mathcal{M}''(\mathbf{x}) : [0, \frac{r}{2\sqrt{d}}]^d \mapsto [0, \frac{r}{2\sqrt{d}}]^d$ considered in Theorem B.2 such that for all $\mathbf{x} \in [0, \frac{r}{2\sqrt{d}}]^d$

$$\mathbb{E}[\|\mathcal{M}''(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty].$$

To prove the above statement, we first let $a = \frac{r}{2\sqrt{d}}$ and $\mathcal{M}'(\mathbf{x}) = \min\{\mathcal{M}(\mathbf{x}), a\}$, where min is a coordinate-wise operator. Now we fix the randomness of $\mathcal{M}(\mathbf{x})$ (that is $\mathcal{M}(\mathbf{x})$ is deterministic), and we assume that $\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty = |\mathcal{M}_j(\mathbf{x}) - x_j|$, $\|\mathcal{M}'(\mathbf{x}) - \mathbf{x}\|_\infty = |\mathcal{M}'_i(\mathbf{x}) - x_i|$. If $\mathcal{M}_i(\mathbf{x}) < a$, then by the definitions, we have $\|\mathcal{M}'(\mathbf{x}) - \mathbf{x}\|_\infty = |\mathcal{M}'_i(\mathbf{x}) - x_i| = |\mathcal{M}_i(\mathbf{x}) - x_i| \leq \|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty$. If $\mathcal{M}_i(\mathbf{x}) \geq a$, then we have $|\mathcal{M}'_i(\mathbf{x}) - x_i| = |a - x_i|$. Since $x_i \in [0, a]$ and $\mathcal{M}_i(\mathbf{x}) \geq a$, $|\mathcal{M}_i(\mathbf{x}) - x_i| \geq |a - x_i|$. $\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \geq |\mathcal{M}_i(\mathbf{x}) - x_i| \geq |a - x_i|$. Thus, $\mathbb{E}[\|\mathcal{M}'(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty]$.

Then, we let $\mathcal{M}''(\mathbf{x}) = \max\{\mathcal{M}'(\mathbf{x}), 0\}$ where max is also a coordinate-wise operator. We can use a similar method to prove that $\mathbb{E}[\|\mathcal{M}''(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \mathbb{E}[\|\mathcal{M}'(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty]$ i.e., the α in Theorem B.2 should be less than or equal to the α in Theorem B.1. Also we can see that $\mathcal{M}''(\mathbf{x}) = \max\{0, \min\{\mathcal{M}(\mathbf{x}), a\}\} = \mathbf{x} + \max\{-\mathbf{x}, \min\{\mathbf{z}, a - \mathbf{x}\}\}$, which means \mathcal{M}'' is also $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized mechanism due to the postprocessing property. Therefore, the lower bound for the α in Theorem B.2 (i.e., $\Omega(\frac{r}{\sqrt{\epsilon}})$) is also a lower bound for the α in Theorem B.1.

Finally, we give a proof of Theorem B.2.

We consider a special case: let $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2^d}\} = \{0, \frac{r}{2\sqrt{d}}\}^d$. Since for all $\mathbf{x}_i, \mathbf{x}_j, i, j \in [2^d]$, $\|\mathbf{x}_i - \mathbf{x}_j\|_2 \leq r$, \mathcal{M} is $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust on $\{0, \frac{r}{2\sqrt{d}}\}^d$. According to Theorem 4.1, \mathcal{M} is also $(\epsilon + 2\sqrt{\log(1/\delta)\epsilon}, \delta)$ -PixelDP. Thus, we have $\mathcal{M}(\cdot)$ is $(\epsilon + 2\sqrt{\log(1/\delta)\epsilon}, \delta)$ DP on X .

Let us first connect the lower bound of one-way marginal (*i.e.*, mean estimation) with the lower bound studied in Theorem B.2. Suppose an n -size dataset $X \in \mathbb{R}^{n \times d}$, the one-way marginal is $h(D) = \frac{1}{n} \sum_{i=1}^n X_i$, where X_i is the i -th row of X . In particular, when $n = 1$, one-way marginal is just the data point itself, and thus, the condition in Theorem B.2 can be rewritten as

$$\mathbb{E}_{\mathcal{M}} \|\mathcal{M}(D) - h(D)\|_{\infty} \leq \alpha. \quad (9)$$

Based on this connection, we first prove the case where $r = 2\sqrt{d}$, and then generalize it to any r . For $r = 2\sqrt{d}$, the conclusion reduces to $\alpha \geq \Omega(\sqrt{\frac{d}{\epsilon}})$. To prove this, we employ the following lemma, which provides a one-way margin estimation for all DP mechanisms.

Lemma 5 (Theorem 1.1 in [29]) *For every $\epsilon \leq O(1)$, every $2^{-\Omega(n)} \leq \delta \leq \frac{1}{n^{1+\Omega(1)}}$ and every $\alpha \leq \frac{1}{10}$, if $\mathcal{M} : (\{0, 1\}^d)^n \mapsto [0, 1]^d$ is (ϵ, δ) -DP and $\mathbb{E}[\|\mathcal{M}(D) - h(D)\|_{\infty}] \leq \alpha$, then we have $n \geq \Omega(\frac{\sqrt{d \log \frac{1}{\delta}}}{\epsilon \alpha})$.*

Setting $n = 1$, $\epsilon = \epsilon + 2\sqrt{\epsilon \log \frac{1}{\delta}}$ in Lemma 5, we can see that if $\mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_{\infty}] \leq \alpha$, then we must have

$$1 \geq \Omega\left(\frac{\sqrt{d \log \frac{1}{\delta}}}{(\epsilon + 2\sqrt{\epsilon \log \frac{1}{\delta}})\alpha}\right) \geq \Omega\left(\frac{\sqrt{d}}{\sqrt{\alpha^2 \epsilon}}\right),$$

where the last inequality is due to the fact that $\frac{\sqrt{\log \frac{1}{\delta}}}{\epsilon + 2\sqrt{\epsilon \log \frac{1}{\delta}}} \geq \Omega(\frac{1}{\sqrt{\epsilon}})$, since $\epsilon \leq O(1)$. Therefore, we have the following theorem,

Theorem B.3 *For some $\epsilon \leq O(1)$, if there is a $(2\sqrt{d}, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust randomized mechanism $\mathcal{M} : \{0, 1\}^d \mapsto [0, 1]^d$ satisfies that for all $\mathbf{x} \in \{0, 1\}^d$*

$$\mathbb{E}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_{\infty} \leq \alpha, \quad (10)$$

for some constant $\alpha = O(1)$. Then $1 \geq \Omega(\sqrt{\frac{d}{\epsilon \alpha^2}})$.

Now we come back to the proof for any r . Note that if a $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ -robust mechanism \mathcal{M} satisfies

$$\mathbb{E}_{\mathcal{M}} \|\mathcal{M}(\mathbf{x}_i) - \mathbf{x}_i\|_{\infty} \leq \alpha,$$

then we have

$$\mathbb{E}_{\mathcal{M}} \left\| \frac{2\sqrt{d}}{r} \mathcal{M}(\mathbf{x}_i) - \frac{2\sqrt{d}}{r} \mathbf{x}_i \right\|_{\infty} \leq \frac{2\sqrt{d}}{r} \alpha.$$

Since $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : [0, \frac{r}{2\sqrt{d}}]^d \mapsto [0, \frac{r}{2\sqrt{d}}]^d$ is $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ robust, thus $\frac{2\sqrt{d}}{r} \mathcal{M} : [0, 1]^d \mapsto [0, 1]^d$ is still $(r, D_{MR}, \|\cdot\|_2, \epsilon)$ robust. By Theorem B.3 with $\alpha = \frac{2\sqrt{d}}{r} \alpha \leq O(1)$, we have

$$1 \geq \Omega\left(\frac{r}{\sqrt{\epsilon \alpha^2}}\right), \quad \text{i.e., } \alpha \geq \Omega\left(\frac{r}{\sqrt{\epsilon}}\right). \quad (11)$$

■

C Omitted Proofs in Section 6

Proof [Proof of Theorem 6.1] We first prove that $D_\infty(g(\mathbf{x})\|g(\mathbf{x}')) \leq \frac{r}{\sigma}$ for all $\mathbf{x}' \in \mathbb{B}_\infty(\mathbf{x}, r)$. Since $\|\mathbf{x}' - \mathbf{x}\|_\infty \leq r$, for any \mathbf{y} ,

$$\frac{p(\mathbf{y} - \mathbf{x})}{p(\mathbf{y} - \mathbf{x}')} = \frac{\exp(-\frac{\|\mathbf{y} - \mathbf{x}\|_\infty}{\sigma})}{\exp(-\frac{\|\mathbf{y} - \mathbf{x}'\|_\infty}{\sigma})} \leq \exp\left(\frac{\|\mathbf{y} - \mathbf{x}'\|_\infty - \|\mathbf{y} - \mathbf{x}\|_\infty}{\sigma}\right) \leq \exp\left(\frac{\|\mathbf{x}' - \mathbf{x}\|_\infty}{\sigma}\right) \leq \exp\left(\frac{r}{\sigma}\right).$$

Since

$$D_\alpha(g(\mathbf{x})\|g(\mathbf{x}')) < D_\infty(g(\mathbf{x})\|g(\mathbf{x}')) = \mathbb{E}\left[\log \frac{p(\mathbf{x})}{p(\mathbf{x}')}\right] \leq \frac{r}{\sigma} < \frac{r}{\sigma}\alpha,$$

$\forall \alpha \in (1, +\infty)$, $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_\infty, \frac{r}{\sigma})$ -robust. Also, based on the following lemma,

Lemma 6 ([30]) *Let P and Q be two probability distributions satisfying $D_\infty(P\|Q) \leq \epsilon$ and $D_\infty(Q\|P) \leq \epsilon$. Then, $D_\alpha(P\|Q) \leq \frac{1}{2}\epsilon^2\alpha$,*

we have $D_\alpha(g(\mathbf{x})\|g(\mathbf{x}')) \leq \frac{1}{2}(\frac{r}{\sigma})^2\alpha$, i.e., $g(\cdot)$ is $(r, D_{MR}, \|\cdot\|_\infty, \frac{r^2}{2\sigma^2})$ -robust. \blacksquare

Proof [Proof of Corollary 2] Define the distribution D on $[0, \infty)$ to be $Z \sim D$, meaning $Z = \|\mathbf{z}\|_\infty$ for $\mathbf{z} \sim p(\mathbf{z})$, where $p(\mathbf{z})$ is defined in Eq.(8). The probability density function of D is given by

$$p_D(Z) \propto Z^{d-1} \exp\left(-\frac{Z}{\sigma}\right),$$

which is obtained by integrating the probability density function in Eq. (8) over the infinity ball of radius Z with surface area $d2^d Z^{d-1} \propto Z^{d-1}$. p_D is the Gamma distribution with shape d and mean σ , and thus $\mathbb{E}[\mathbf{z}] = d\sigma$. \blacksquare

Proof [Proof of Theorem 6.2] Similar to the proof of Theorem 5.2, in order to prove Theorem 6.2, we only need to prove the following theorem:

Theorem C.1 *For some $\epsilon = O(1)$ and $r = \Omega(1)$, if there is a $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ randomized (smoothing) mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : [0, \frac{r}{2}]^d \mapsto [0, \frac{r}{2}]^d$ such that for any $\mathbf{x} \in [0, \frac{r}{2}]^d$, the following holds*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \alpha$$

for some constant $\alpha = O(1)$. Then it must be true that $\alpha \geq \Omega(\frac{r\sqrt{d}}{\sqrt{\epsilon}})$.

Assume that we have a set of data points $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2^d}\} = \{0, \frac{r}{2}\}^d$. Since $\|\mathbf{x}_i - \mathbf{x}_j\|_\infty \leq r$, thus we have \mathcal{M} is $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robust on $\{0, \frac{r}{2}\}^d$. According to Theorem 4.1, \mathcal{M} is also $(\epsilon + 2\sqrt{\log(1/\delta)\epsilon}, \delta)$ -PixelDP. Thus, it is $(\epsilon + 2\sqrt{\log(1/\delta)\epsilon}, \delta)$ -DP on \mathcal{X} . We first consider the case where $r = 2$. By Lemma 2 with $n = 1$ and $\epsilon = \epsilon + 2\sqrt{\epsilon \log 1/\delta}$ we have a similar result as in Theorem B.3:

Theorem C.2 *For some $\epsilon \leq O(1)$, if there is a $(2, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robust randomized mechanism $\mathcal{M} : \{0, 1\}^d \mapsto [0, 1]^d$ satisfies that for all $\mathbf{x} \in \{0, 1\}^d$*

$$\mathbb{E}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \leq \alpha, \tag{12}$$

for some constant $\alpha = O(1)$. Then $1 \geq \Omega(\sqrt{\frac{d}{\epsilon\alpha^2}})$.

For general r , if for any $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ robust randomized mechanism $\mathcal{M}(\cdot) : [0, \frac{r}{2}]^d \mapsto [0, \frac{r}{2}]^d$ we have

$$\mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \alpha,$$

that is

$$\mathbb{E}\left[\left\|\frac{2}{r}\mathcal{M}(\mathbf{x}) - \frac{2}{r}\mathbf{x}\right\|_\infty\right] \leq \frac{2}{r}\alpha.$$

This indicates that $\frac{2}{r}\mathcal{M}(\cdot) : [0, 1]^d \mapsto [0, 1]^d$ is $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$. Thus by Theorem C.2 with $\alpha = \frac{2}{r}\alpha$ we have

$$1 \geq \Omega\left(\frac{\sqrt{d}}{\sqrt{\epsilon(2/r\alpha)^2}}\right),$$

thus we have Theorem C.1. ■

Proof [Proof of Theorem 6.3] By simple calculation we have

$$D_\alpha(\mathcal{N}(\mathbf{x}, \frac{dr^2}{2\epsilon}I_d) \|\mathcal{N}(\mathbf{x}', \frac{dr^2}{2\epsilon}I_d) = \frac{\alpha\epsilon\|\mathbf{x} - \mathbf{x}'\|_2^2}{dr^2} \leq \frac{\alpha d\epsilon\|\mathbf{x} - \mathbf{x}'\|_\infty^2}{dr^2} \leq \alpha\epsilon.$$

Therefore, $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \frac{dr^2}{2\epsilon}I_d)$ is $(r, D_{MR}, \|\cdot\|_\infty, \epsilon)$ -robust. The bound of $\mathbb{E}[\|\mathbf{z}\|_\infty]$ can be easily proved by substituting σ in $O(\sigma\sqrt{\log d})$ [25] with $\sigma = \sqrt{\frac{dr^2}{2\epsilon}}$. ■

D Extension to ℓ_p -norm robustness for Any $p \in [2, \infty)$

In previous sections, we studied ℓ_2 -norm and ℓ_∞ -norm robustness. As we mentioned earlier, our framework can be applied to general norm. In this section, we will study the general ℓ_p -norm robustness with $p \geq 2$. Just as the previous sections, here we first investigate the ℓ_p -norm criteria for assessment.

Theorem D.1 *Given $p \geq 2$, for some $\epsilon = O(1)$ and $\sqrt[p]{d} \leq O(r)$, if there is a $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ randomized (smoothing) mechanism $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : \mathbb{R}^d \mapsto \mathbb{R}^d$ such that*

$$\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}[\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty] \leq \alpha$$

for some constant $\alpha = O(1)$. Then it must be true that $\alpha \geq \Omega(\frac{rd^{\frac{1}{2}-\frac{1}{p}}}{\sqrt{\epsilon}})$. When $p \rightarrow \infty$, according to Theorem 6.2, $\alpha \geq \Omega(\frac{rd^{\frac{1}{2}}}{\sqrt{\epsilon}})$

Proof [Proof of Theorem D.1] The proof is also almost the same as that of Theorem 5.2. Following the proof of Theorem 5.2, we can only constrain on the case where $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z} : [0, \frac{r}{2\sqrt[p]{d}}]^d \mapsto [0, \frac{r}{2\sqrt[p]{d}}]^d$.

Assume we have a set of data points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{2^d}\} = \{0, \frac{r}{2\sqrt[p]{d}}\}^d$. Since we have for any $i, j \in [2^d]$, $\|\mathbf{x}_i - \mathbf{x}_j\|_p \leq r$, thus by Theorem 4.1 we have \mathcal{M} is $(\epsilon + 2\sqrt{\epsilon \log 1/\delta}, \delta)$ PixelDP. And thus, it is $(\epsilon + 2\sqrt{\epsilon \log 1/\delta}, \delta)$ DP on \mathcal{X} .

We first consider the case where $r = 2\sqrt[p]{d}$, then we extend to the general case. When $r = 2\sqrt[p]{d}$, like Theorem B.3, we have

Theorem D.2 *For some $\epsilon \leq O(1)$, if a $(2\sqrt[p]{d}, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust randomized mechanism $\mathcal{M} : \{0, 1\}^d \mapsto [0, 1]^d$ satisfies that for all $\mathbf{x} \in \{0, 1\}^d$*

$$\mathbb{E}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \leq \alpha, \tag{13}$$

then $1 \geq \Omega(\sqrt{\frac{d}{\epsilon\alpha^2}})$. ■

Now for any $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ robust mechanism $\mathcal{M} : [0, \frac{r}{2\sqrt[p]{d}}]^d \mapsto [0, \frac{r}{2\sqrt[p]{d}}]^d$, if we have

$$\mathbb{E}\|\mathbf{z}\|_\infty = \mathbb{E}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty \leq \alpha,$$

then we have

$$\mathbb{E}\left[\left\|\frac{2\sqrt[p]{d}}{r}\mathcal{M}(\mathbf{x}) - \frac{2\sqrt[p]{d}}{r}\mathbf{x}\right\|_\infty\right] \leq \frac{2\sqrt[p]{d}}{r}\alpha.$$

Here we note that $\frac{2\sqrt[p]{d}}{r}\mathcal{M} : \{0, 1\}^d \mapsto [0, 1]^d$ and it is also $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ robust, and also $\frac{2\sqrt[p]{d}}{r}\alpha \leq O(1)$. Thus by Theorem D.1 we have

$$1 \geq \Omega\left(\sqrt{\frac{d}{\epsilon(2\sqrt[p]{d}/r\alpha)^2}}\right)$$

Thus we have $\alpha \geq \Omega\left(\frac{rd^{\frac{1}{2}-\frac{1}{p}}}{\sqrt{\epsilon}}\right)$.

Remark 3 First, we can see that when $p = 2$, Theorem D.1 is the same as Theorem 5.2. Thus, we can see it as a generalization of the previous theorem. Secondly, Theorem D.1 indicates that for a certain extent of robustness, the magnitude of the noise we add should be at least $\Omega(d^{\frac{1}{2}-\frac{1}{p}})$, which will quite large when in the high dimensional case. This means that for ℓ_p -norm robustness with $p > 2$, random smoothing mechanisms are not appropriate to high dimensional data. Or, we call it as the curse of dimensionality on random smoothing mechanism for certifying ℓ_p ($p \geq 2$) robustness.

Next, we show that based on the above criteria for ℓ_p -norm robustness, Gaussian mechanism is an appropriate option. Since the gap between the lower bound in the criteria and the magnitude of Gaussian mechanism is just $O(\sqrt{\log d})$.

Theorem D.3 (Gaussian Mechanism for Certifying ℓ_p -norm robustness) Let $r, \epsilon > 0$ be some fixed number and $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \frac{d^{1-\frac{2}{p}}r^2}{2\epsilon}I_d)$. Then, $\mathcal{M}(\cdot)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust, and $\mathbb{E}[\|\mathbf{z}\|_\infty] = \mathbb{E}_{\mathcal{M}}\|\mathcal{M}(\mathbf{x}) - \mathbf{x}\|_\infty$ is upper bounded by $O\left(\frac{rd^{\frac{1}{2}-\frac{1}{p}}\sqrt{\log d}}{\sqrt{\epsilon}}\right)$.

Proof [Proof of Theorem D.3] By simple calculation we have

$$D_\alpha(\mathcal{N}(\mathbf{x}, \frac{d^{1-\frac{2}{p}}r^2}{2\epsilon}I_d) \|\mathcal{N}(\mathbf{x}', \frac{d^{1-\frac{2}{p}}r^2}{2\epsilon}I_d)) = \frac{\alpha\epsilon\|\mathbf{x} - \mathbf{x}'\|_2^2}{d^{1-\frac{2}{p}}r^2} \leq \frac{\alpha d^{1-\frac{2}{p}}\epsilon\|\mathbf{x} - \mathbf{x}'\|_p^2}{d^{1-\frac{2}{p}}r^2} \leq \alpha\epsilon.$$

Therefore, $\mathcal{M}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \frac{d^{1-\frac{2}{p}}r^2}{2\epsilon}I_d)$ is $(r, D_{MR}, \|\cdot\|_p, \epsilon)$ -robust. The bound of $\mathbb{E}[\|\mathbf{z}\|_\infty]$ can be easily proved by substituting σ in $O(\sigma\sqrt{\log d})$ [25] with $\sigma = \sqrt{\frac{d^{1-\frac{2}{p}}r^2}{2\epsilon}}$. ■

E Additional Details & Results

E.1 Numerical Method

We first detail the numerical method for the experiments in the following. The certification algorithm is first detailed in Alg. 1. Here we highlight the sampling method for the Exponential mechanism. Due to the high dimensionality of samples in real world applications, directly sampling $\mathbf{z} \sim p(\mathbf{z})$ as in Eq. 8 by the Markov Chain Monte Carlo (MCMC) algorithm requires a large number of random-walks that can incur high computational cost. To alleviate this issue, we adopt an efficient sampling method from [31] that first samples R from $\text{Gamma}(d+1, \sigma)$ and then samples \mathbf{z} from $[-R, R]^d$ uniformly. The complexity of this sampling algorithm is only $O(d)$.

E.2 Additional Experiment Results

ℓ_2 -norm Case In Fig. 3, we can see that, although [11] proves a tighter bound than ours, it also certifies approximately 40 ~ 60% accuracy at ℓ_2 radius = 0.34 (CIFAR-10, $d = 3072$) and ℓ_2 radius = 0.29 (ImageNet, $d = 150568$), i.e., $O(1/\sqrt{\log d})$. Even after using the advanced training method in [21], the scale of the robust radii is still $O(1/\sqrt{\log d})$, as shown in Fig. 4.

ℓ_∞ -norm Case Note that it seems obvious that the Gaussian mechanism is an appropriate mechanism to certify ℓ_2 -norm robustness since [11, 2, 21] have achieved the state-of-the-art certification results compared with the other methods in the ℓ_2 -norm case. However, in the ℓ_∞ -norm case, it is a little counterintuitive that the Gaussian mechanism is also an appropriate choice, which performs

Algorithm 1 Certifying ℓ_2/ℓ_∞ -norm Robustness

Require: Input \mathbf{x} , a classifier $f(\cdot)$, parameter $\sigma > 0$, number of samples for estimating confidence interval n .

Sample n samples from the Gaussian/Exponential mechanism $\{\mathbf{z}_i\}_{i=1\dots n}$

Output $c_i = f(\mathbf{x} + \mathbf{z}_i)$, and estimate the distribution of c_i , i.e., $p_j = \frac{\#\{c_i=j\}_{i=1\dots n}}{n}$

if Choose the Gaussian mechanism **then**

 Compute the robust radius by $r_2 = \sup_{\alpha>1} (-\frac{2\sigma^2}{\alpha} \log(1 - p_1 - p_2 + 2(\frac{1}{2}(p_1^{1-\alpha} + p_2^{1-\alpha}))^{\frac{1}{1-\alpha}})) \frac{1}{2}$

else if Choose the Exponential mechanism **then**

$r_a = \sup_{\alpha>1} -\frac{2\sigma^2}{\alpha} \log(1 - p_1 - p_2 + 2(\frac{1}{2}(p_1^{1-\alpha} + p_2^{1-\alpha}))^{\frac{1}{1-\alpha}})$

$r_b = \sup_{\alpha>1} (-\frac{2\sigma^2}{\alpha} \log(1 - p_{(1)} - p_{(2)} + 2(\frac{1}{2}(p_{(1)}^{1-\alpha} + p_{(2)}^{1-\alpha}))^{\frac{1}{1-\alpha}})) \frac{1}{2}$

$r_\infty = \max\{r_a, r_b\}$

end if

For the Gaussian mechanism, ℓ_2 robust radius is r_2 , and ℓ_∞ robust radius is $\sqrt{r_2^2/d}$. For the Exponential mechanism, ℓ_∞ robust radius is r_∞ .

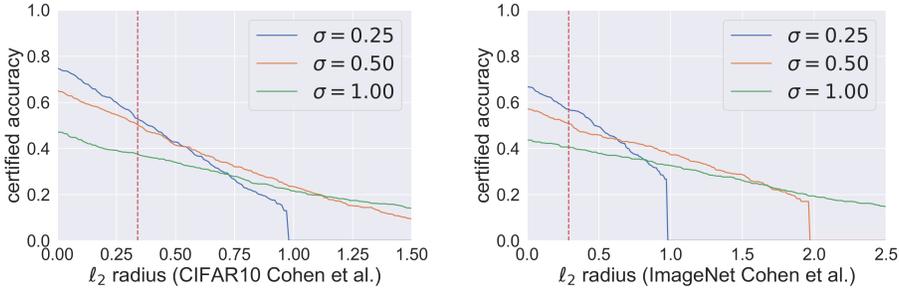


Figure 3: Certify ℓ_2 -norm robustness by the Gaussian mechanism [11]

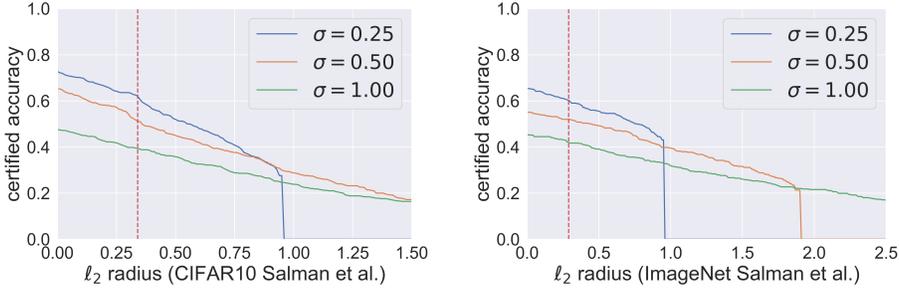


Figure 4: Certify ℓ_2 -norm robustness by the Gaussian mechanism and the adversarial training method in [21]

much better than the Exponential mechanism. In the Table 1, we compare the ℓ_∞ -norm certification results of the Gaussian mechanism and the other two representative approaches. Although [11] and the D_{MR} framework perform slightly worse than [14] or [17] on CIFAR10, they are more scalable to high-dimensional datasets like ImageNet. So we can say their ℓ_∞ -norm certification results are comparable. Besides, in Fig. 5 & 6, we show that the Gaussian mechanism certifies approximately 40 ~ 60% accuracy at ℓ_∞ radius = $6e - 3$ on CIFAR-10 and ℓ_∞ radius = $1.1e - 3$ on ImageNet, which are also approximately $O(1/\sqrt{d \log d})$ for both datasets.

All in all, the empirical results indicate the theorems proved under our framework are valid and very likely to generalize to the other frameworks.

Model	CIFAR-10		IMAGENET	
	ℓ_∞ Acc at 2/255	Standard Acc	ℓ_∞ Acc at 1/255	Standard Acc
Cohen et al. [11] (Gaussian)	47.0%	74.8% ($\sigma = 0.25$)	27.4%	57.2% ($\sigma = 0.5$)
D_{MR} Framework (Gaussian)	42.4%	69.6% ($\sigma = 0.5$)	24.4%	45.2% ($\sigma = 1.0$)
Wong et al. [14] (Single model)	53.9%	68.3%	-	-
IBP [17]	50.0%	70.2%	-	-

Table 1: Comparing the performance of the Gaussian mechanism with the other works in the ℓ_∞ case

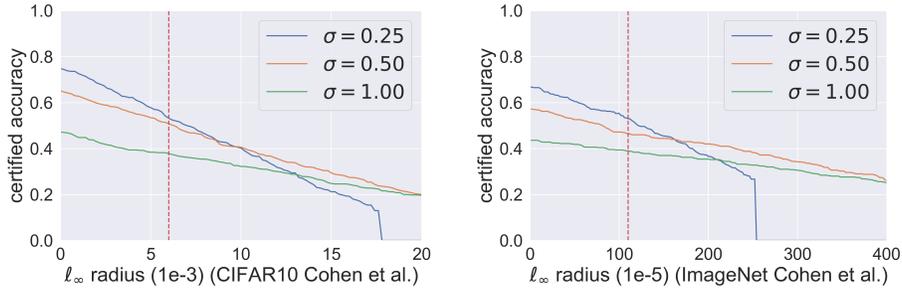


Figure 5: Certify ℓ_∞ -norm robustness by the Gaussian mechanism [11]

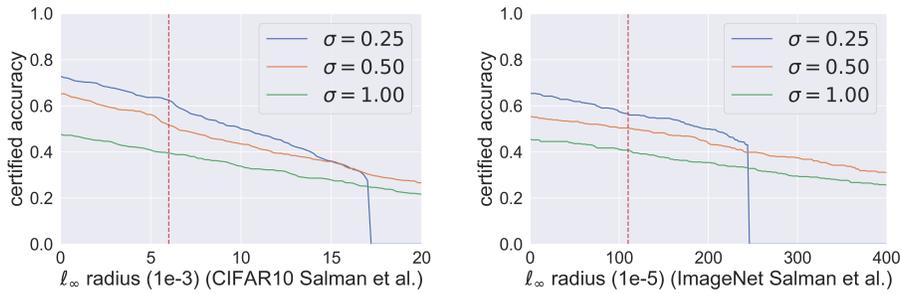


Figure 6: Certify ℓ_∞ -norm robustness by the Gaussian mechanism and the adversarial training method in [21]