

# Random Differential Privacy

Rob Hall\*, Alessandro Rinaldo†, and Larry Wasserman‡

## 1 Introduction

*Differential privacy* (DP) ([6]) is a type of privacy guarantee that has become quite popular in the computer science literature. The advantage of differential privacy is that it gives a strong and mathematically rigorous guarantee. The disadvantage is that the strong privacy guarantee often comes at the expense of the statistical utility of the released information. We propose a weaker notion of privacy, called “random differential privacy” (RDP), under which it is possible to achieve better accuracy.

The privacy guarantee provided by RDP represents a radical weakening of the ordinary differential privacy. This could be a cause for concern for those who want very strong privacy guarantees. Indeed, *we are not suggesting the RDP should replace ordinary differential privacy*. However, as we shall show in this paper (and has been observed many times in the past), differential privacy can lead to large information losses in some cases (see e.g., [10, 18, 4]). Thus, we feel there is great value in exploring weakened versions of differential privacy. In other words, we are proposing a new privacy definition as a way of exploring the privacy/accuracy tradeoff.

We begin by introducing ordinary differential privacy and setting up some notation. We then explore the lower limits for accuracy of differentially private techniques in the context of histograms. We introduce a concept which parallels minimaxity in statistics, and identify the minimax risk for a differentially private histogram. We describe an important subset of these minimax differentially private histograms which we show to have risk which is uniformly lower bounded at a rate which is linear in the dimension of the histogram. We then introduce our proposed relaxation to differential privacy, under which our technique enjoys the same minimax risk, but with a lower bound which depends only on the size of the support of the histogram (namely, the number of non-zero cells). Thus we show that in the context of sparse histograms, the relaxation allows for a strictly better data release. We also demonstrate some important properties of our relaxation, such as an analog of the composition lemma.

---

\*Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, <mailto:rjhall@cs.cmu.edu>

†Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, <mailto:arinaldo@cmu.edu>

‡Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, <mailto:larry@stat.cmu.edu>

## 2 Differential Privacy (DP)

### 2.1 Definition

Let  $X = (X_1, \dots, X_n) \in \mathcal{X}^n$  be an input database with  $n$  observations where  $X_i \in \mathcal{X}$ . The goal is to produce some output  $Z \in \mathcal{Z}$ . For example, the inputs may consist of database rows in which each column is a measurement of an individual, and the output is the number of individuals having some property. Let  $Q_n(\cdot | X)$  be a conditional distribution for  $Z$  given  $X$ . Write  $X \sim X'$  if  $X, X' \in \mathcal{X}^n$ , and  $X$  and  $X'$  differ in one coordinate. We say that  $X$  and  $X'$  are *neighboring databases*.<sup>1</sup>

We say  $Q_n$  satisfies  $\alpha$  differential privacy if, for all measurable  $B \subset \mathcal{Z}$  and all  $X \sim X' \in \mathcal{X}^n$ ,

$$e^{-\alpha} \leq \frac{Q_n(Z \in B | X)}{Q_n(Z \in B | X')} \leq e^{\alpha}. \quad (1)$$

The intuition is that, for small  $\alpha > 0$ , the value of one individual's data has a small effect on the output. We consider any DP algorithm to be a family of distributions  $Q_n$  over the output space  $\mathcal{Z}$ . We index a family of distributions by  $n$  to show the size of the dataset.

It has been shown by researchers in privacy that differential privacy provides a very strong guarantee. Essentially it means that whether or not one particular individual is entered in the database has negligible effect on the output. The research in differential privacy is vast. A few key references are [6], [9], [2], [7], [3], and references therein.

### 2.2 Noninteractive Privacy and Histograms

Much research on differential privacy focuses on the case where  $Z$  is a response to some query such as “what is the mean of the data.” A simple way to achieve differential privacy in that case is to add some noise to the mean of  $X$  where the noise has a Laplace distribution. The user may send a sequence of such queries. This is called *interactive privacy*. We instead focus on the *noninteractive privacy*, where the goal is to output a whole database (or a “synthetic dataset”)  $Z = (Z_1, \dots, Z_N)$ . Then the user is not restricted to a small number of queries.

One way to release a private database is to first release a privatized histogram. We can then draw an arbitrarily large sample  $Z = (Z_1, \dots, Z_N)$  from the histogram. It is easy to show that if the histogram satisfies DP then  $Z$  also satisfies DP. Hence, in the rest of the paper, we focus on constructing a private histogram.

We consider privatization mechanisms which are permutation invariant with respect to their inputs (i.e., those distributions which treat the values  $x_i$  as a set rather than a vector); in the context of histograms this appears to be a very mild restriction.

---

<sup>1</sup>In some papers, the definition is changed so that one sample is a strict subset of the other, having exactly one less element. Although this definition is perhaps slightly stronger, we do not use it and remark that the approaches we present below may all be fit into this framework if so desired.

We partition the sample space  $\mathcal{X}$  into  $k$  cells (or bins)  $\{B_j\}_{j=1}^k$ .<sup>2</sup> We consider the input to be a lattice point in the  $k$ -simplex, by taking the function:  $\theta^n(x_1, \dots, x_n) = (\theta_1, \dots, \theta_k)$ ,  $\theta_j = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{x_i \in B_j\}$ . The image of this mapping  $\Theta = \theta^n(\mathcal{X}^n)$  is the set of lattice points in the simplex which correspond to histograms of  $n$  observations in  $k$  bins. Note that this is in essence a “normalized histogram” since the elements sum to one. This set depends on  $k$  although we suppress this notation. For the remainder of this paper we consider the output space  $\mathcal{Z}$  to be the same as the input space (i.e., a normalized histogram).

Now we give a concrete example of a  $Q_n$  which achieves differential privacy. Define  $z_j = \theta_j + 2L_j/(n\alpha)$  where  $L_1, \dots, L_k$  are independent draws from a Laplace distribution with mean zero and rate one. Then  $(z_1, \dots, z_k)$  satisfy DP (see e.g., [6]). However, the  $z_i$  themselves do not represent a histogram, because they can be negative and they do not necessarily sum to one. Hence we may take, for example:

$$\delta(z) = \arg \min_{\theta \in \Theta} \|z - \theta\|_1 \quad (2)$$

where we use the  $\ell_1$  norm:  $\|x\|_1 = \sum_j |x_j|$ . Hence this procedure results in a valid histogram. Note that  $\delta(z)$  satisfies the differential privacy since each subset of values it may take clearly corresponds to a measurable subset of  $\mathbb{R}^k$ . Since the differential privacy held for the real vector then it also holds for the projection (see e.g., [17]). We will refer to this as the *histogram perturbation method* (see e.g., [17]). There are other methods for generating differentially private histograms, and our results below concern hold over a large subset of all the possible techniques available (to be made precise after Proposition 3.2). Hence our results apply to more than the above concrete scheme.

### 3 Lower Bounds for Accuracy with Differential Privacy

To motivate the need for relaxed versions of differential privacy, we consider here the accuracy of differentially private histograms. We evaluate a differentially private procedure in terms of its “risk” which is a natural measure of accuracy taken from statistics. We consider the  $\ell_1$  loss function, and the associated risk:

$$R(\theta, Q_n) = \int_{\Theta} \|\hat{\theta} - \theta\|_1 dQ_n(\hat{\theta}|\theta). \quad (3)$$

where  $\hat{\theta}$  is the output of the differentially private algorithm,  $\theta$  is the input histogram, and the distribution  $Q_n$  is the one induced by the randomized algorithm. Typically this risk will be a non-constant function of the parameter  $\theta$  and of the distribution  $Q_n$ . Therefore we consider the “minimax risk” which is the smallest achievable worst-case risk, and gives a measure of the hardness of the problem which does not depend on a

---

<sup>2</sup> In this paper,  $k$  is taken as a given integer. The problem of choosing an optimal  $k$  in a private matter is the subject of future work.

particular choice of procedure:

$$R^* = \inf_{Q_n} \sup_{\theta \in \Theta} R(\theta, Q_n). \quad (4)$$

We next describe the minimax risk of the best fully differentially private mechanism  $Q_n$ .

**Proposition 3.1.**

$$R^* \geq c_0 \frac{k-1}{\alpha n}$$

*Proof.* The proof uses a standard method for deriving minimax lower bounds in statistical estimation. Consider the  $k-1$ -dimensional hypercube

$$\left\{ \left( \frac{\sigma_1 \tau}{n}, \dots, \frac{\sigma_{k-1} \tau}{n}, \frac{(n - \sum_{i=1}^{k-1} \sigma_i) \tau}{n} \right) : \sigma_i \in \{0, 1\} \right\}.$$

Take  $\theta, \theta'$  to be neighboring corners of this hypercube (namely two elements which differ in exactly one coordinate  $\sigma_i$ ). Take the KL divergence between the conditional distributions at these corners to be:

$$KL(Q_n(\cdot|\theta) \| Q_n(\cdot|\theta')) = \int_{\Theta} \log \frac{Q_n(\hat{\theta}|\theta)}{Q_n(\hat{\theta}|\theta')} dQ_n(\hat{\theta}|\theta).$$

By considering a sequence of points corresponding to neighboring inputs, we find the ratio of densities to have the upper bound:  $\frac{Q_n(\hat{\theta}|\theta)}{Q_n(\hat{\theta}|\theta')} \leq e^{\alpha \tau}$  since  $\tau$  elements of the input have to change to move from  $\theta$  to  $\theta'$ , and the ratio at each step is bounded by  $e^\alpha$ . Therefore the KL divergence obeys  $KL(Q_n(\cdot|\theta) \| Q_n(\cdot|\theta')) \leq \alpha \tau$ . The ‘‘affinity’’ between the two distributions is:

$$\|Q_n(\cdot|\theta) \wedge Q_n(\cdot|\theta')\| = \int_{\Theta} \min \{Q_n(\hat{\theta}|\theta), Q_n(\hat{\theta}|\theta')\} d\hat{\theta}.$$

The Kullback-Csiszar-Kemperman inequality [19] yields a lower bound on the affinity between these distributions:

$$\|Q_n(\cdot|\theta) \wedge Q_n(\cdot|\theta')\| \geq 1 - \sqrt{\frac{\alpha \tau}{2}}.$$

Assouad’s lemma (see [19] again) thus gives the lower bound:

$$R^* \geq (k-1) \frac{\tau}{2n} \left( 1 - \sqrt{\frac{\alpha \tau}{2}} \right).$$

Taking  $\tau = t/\alpha$  gives

$$R^* \geq (k-1) \frac{t}{2\alpha n} \left(1 - \sqrt{\frac{t}{2}}\right).$$

For  $\alpha < 1$  we may take  $t < 1$ , which results in the parenthetical expression being positive.  $\square$

**Remark 1.** *The previous result demonstrates that the minimax risk of the differentially private histogram is of the order  $O\left(\frac{k}{\alpha n}\right)$ .*

**Remark 2.** *Hardt and Talwar [11] have a similar result although their setting is somewhat different. In particular, they do not restrict to the space of histograms based on  $n$  observations.*

The above results demonstrate that for every differentially private scheme, there is at least one input for which the risk is growing in the order shown (in fact, at least one point in every hypercube of side length  $\tau/n$ ). However, the prospect exists that at many other inputs the risk is much lower. We now demonstrate that this is not the case when  $k = 2$ , by presenting a uniform lower bound for the risk among all minimax schemes. In the case of  $k = 2$  the output may be regarded as a single number  $\frac{a}{n}$  where  $a \in \{0, \dots, n\}$ , which gives the proportion of the data points in the first bin. Our result will show that in a sense, the minimax differential privacy schemes are similar to “equalizer rules” in the sense that the risk is on the same order for every input.

**Proposition 3.2.** *For  $k = 2$ , for any  $Q_n$  which achieves  $\sup_{\theta} R(\theta, Q_n) \leq \frac{c_0}{\alpha n}$ , we have that  $\inf_{\theta} R(\theta, Q_n) \geq \frac{c_1}{\alpha n}$*

*Proof.* Note that for any  $\theta_1$  and  $c > c_0$ , due to the uniform upper bound on the risk, Markov’s inequality gives

$$\int_{\mathcal{Z}} \mathbf{1}\{|\hat{\theta} - \theta_1| \leq \frac{c}{\alpha n}\} dQ_n(\hat{\theta}|\theta_1) \geq 1 - \frac{c_0}{c}.$$

Therefore, due to the constraint of differential privacy, we have that, for any  $\theta_0$ ,

$$\int_{\mathcal{Z}} \mathbf{1}\{|\hat{\theta} - \theta_1| \leq \frac{c}{\alpha n}\} dQ_n(\hat{\theta}|\theta_0) \geq \left(1 - \frac{c_0}{c}\right) \exp\left\{-\frac{\alpha n}{2}\|\theta_0 - \theta_1\|_1\right\}.$$

Since  $\frac{n}{2}\|\theta_0 - \theta_1\|_1$ , elements of the input change to move from  $\theta_0$  to  $\theta_1$ . Therefore taking  $\theta_1$  to give  $\|\theta_0 - \theta_1\|_1 = \frac{2c}{\alpha n}$  gives

$$R(\theta_0, Q_n) \geq \frac{c}{\alpha n} \left(1 - \frac{c_0}{c}\right) e^{-c} = \frac{c_1}{\alpha n}.$$

As  $\theta_0$  is arbitrary, this gives a uniform lower bound under the conditions above.  $\square$

For the relaxation of differential privacy given in Definition 2.2 of [11], the above result remains intact for large enough  $n$ . The relaxation is:

$$Q_n(z|X) \leq Q_n(z|X')e^\alpha + \eta(n)$$

where  $\eta(n)$  is negligible (i.e., tending to zero faster than any inverse polynomial in  $n$ ). Thus via the same technique as above, we have

$$R(\theta_0, \delta, \mathbb{Q}_n) \geq \frac{c}{\alpha n} \left( (1 - \frac{c_0}{c})e^{-c} - c_2\eta(n) \right) = \frac{c_1 - \eta(n)}{\alpha n}.$$

For large enough  $n$  this latter term is bounded from below by  $\frac{c_3}{\alpha n}$ . This indicates that the above relaxation of differential privacy will not be useful in achieving higher accuracy.

For  $k > 2$ , we may write

$$R(\theta, Q_n) = \sum_{i=1}^k R_i(\theta, Q_n)$$

with

$$R_i(\theta, Q_n) \stackrel{\text{def}}{=} \int_{\mathcal{Z}} |\hat{\theta} - \theta_i| dQ_n(\hat{\theta}|\theta),$$

where the subscript means the  $i^{\text{th}}$  coordinate. Thus, whenever we have that  $R_i \leq \frac{c_0}{\alpha n}$  uniformly over  $i$ , we have that  $R(\theta, \delta, \mathbb{Q}_n) \geq \frac{c_1(k-1)}{\alpha n}$ . Therefore the only opportunity to improve upon the rate of  $\frac{k}{\alpha n}$  is when some  $\theta$  have some coordinate  $i$  at which the risk upper bound does not apply.

We conclude by remarking that we have demonstrated that for a certain class of differentially private algorithms which achieve the “minimax rate,” their risk is uniformly lower bounded at the same rate. The rate in question is linear in  $k$ , which is problematic when  $k$  is large relative to  $n$ . It remains an open question whether there are different techniques which achieve the minimax rate, yet do not have this property. Such a technique would have to lose the uniform upper bound on the coordinate-wise risk. Below, we present a weakening of differential privacy, which admits release mechanisms, which both keep the uniform upper bound on the coordinate-wise risk, and also have a minimax risk which is growing only in the support of the histogram (namely, the number of cells which contain observations).

## 4 Random Differential Privacy

In random differential privacy (RDP) we view the data  $X = (X_1, \dots, X_n)$  as random draws from an unknown distribution  $P$ . This is certainly the case in statistical sampling and of course it is the usual assumption in most learning theory. Let us denote the

observed values of the random variables  $X = (X_1, \dots, X_n)$  by  $x = (x_1, \dots, x_n)$ . Recall that under DP,  $Q(Z \in B|x_1, \dots, x_n)$  is not strongly affected if we replace some value  $x_i$  with another value  $x'_i$ . We continue to restrict to the case in which  $Q(Z \in B|x_1, \dots, x_n)$  is invariant to permutations of  $(x_1, \dots, x_n)$ . Thus we may restate DP by saying that  $Q(Z \in B|x_1, \dots, x_n)$  is not strongly affected if we replace  $x_n$  by some other arbitrary value  $x'_n$ . In RDP, we require instead that the distribution  $Q_n(\cdot|x_1, \dots, x_n)$  is not strongly affected if we replace  $x_n$  by some new  $x'_n$  which is also randomly drawn from  $P$ .

**Definition 1 (( $\alpha, \gamma$ )-Random Differential Privacy).** . We say that a randomized algorithm  $Q_n$  is ( $\alpha, \gamma$ )-Randomly Differentially Private when:

$$\mathbb{P}\left(\forall B \subseteq \mathcal{Z}, e^{-\alpha} \leq \frac{Q_n(Z \in B|X)}{Q_n(Z \in B|X')} \leq e^\alpha\right) \geq 1 - \gamma$$

where

$$X = (X_1, \dots, X_{n-1}, X_n), \quad X' = (X_1, \dots, X_{n-1}, X_{n+1})$$

(i.e.,  $X \sim X'$ ), and the probability is with respect to the  $n + 1$ -fold product measure  $P^{n+1}$  on the space  $\mathcal{X}^{n+1}$ , that is,  $X_1, \dots, X_{n+1} \stackrel{iid}{\sim} P$ .

We also give the “random” analog of the ( $\alpha, \delta$ )-Differential Privacy:

**Definition 2 (( $\alpha, \eta, \gamma$ )-Random Differential Privacy).** . We say that a randomized algorithm  $Q_n$  is ( $\alpha, \eta, \gamma$ )-Randomly Differentially Private when:

$$\mathbb{P}(\forall B \subseteq \mathcal{Z}, Q_n(Z \in B|X) \leq e^\alpha Q_n(Z \in B|X') + \eta(n)) \geq 1 - \gamma$$

where  $\eta$  is negligible (i.e., decreasing faster than any inverse polynomial).

We note that [13] also consider a probabilistic relaxation of DP. However, their relaxation is quite different than the one considered here. Namely, their relaxation bounds the probability that the differential privacy criteria is not met, but where the probability is taken with respect to the randomized algorithm itself. Our relaxation takes the probability with respect to the generation of the data itself. The following result is clear from the definition of random differential privacy.

**Proposition 4.1.** ( $\alpha, \gamma$ )-RDP is a strict relaxation of  $\alpha$ -DP. That is, if  $Q_n$  is DP then it is also RDP. However, there are RDP procedures that are not DP.

**Remark 3.** Although an  $\alpha$ -DP procedure fulfills the requirement of ( $\alpha, 0$ )-RDP, the converse is not true. The reason is that the latter requires that the condition (that the ratio of densities be bounded) holds almost everywhere with respect to the unknown measure, whereas DP requires that this condition holds uniformly everywhere in the space.

We next show an important property of the definition, namely, that RDP algorithms may be composed to give other RDP algorithms with different constants. The analogous

composition property for DP was considered to be important because it allowed rapid development of techniques which release multiple statistics, as well as techniques which allow interactive access to the data.

**Proposition 4.2 (Composition).** . Suppose  $Q, Q'$  are distributions over  $\mathcal{Z}, \mathcal{Z}'$  which are  $(\alpha, \gamma)$ -RDP and  $(\alpha', \gamma')$ -RDP, respectively. The following distribution  $C$  over  $\mathcal{Z} \times \mathcal{Z}'$  is  $(\alpha + \alpha', \gamma + \gamma')$ -RDP:

$$C(Z, Z'|X) = Q(Z|X) \cdot Q'(Z'|X).$$

This result is simply an application of the union bound combined with the standard composition property of differential privacy. As an example, suppose it is required to release  $k$  different statistics of some data sample. If each one is released via a  $(\alpha/k, \gamma/k)$ -RDP procedure, then the overall release of all  $k$  statistics together achieves  $(\alpha, \gamma)$ -RDP. A similar result holds for the composition of  $(\alpha, \delta, \gamma)$ -RDP releases.

## 5 RDP Sparse Histograms

We first give a technique for the release of a histogram which works well in the case of a sparse histogram, and which satisfies the  $(\alpha, \gamma)$ -Random Differential Privacy. We then compare the accuracy of this method to a lower bound on the accuracy of a  $\alpha$ -Differentially Private approach.

The basic idea is to not add any noise to cells with low counts. This results in partitioning the space into two blocks and releasing a noise-free histogram in one block, and using a differentially private histogram in the other. The partition will depend on the data itself. For a sample  $x_1, \dots, x_n$ , we denote:  $S = S(x_1, \dots, x_n) = \{j : \theta_j = 0\}$ .

Then we consider the release mechanism:

$$z_j = \begin{cases} \theta_j & j \in S \text{ and } 2k \leq \gamma n \\ \theta_j + \frac{2}{n\alpha}L & \text{o/w.} \end{cases} \quad (5)$$

**Proposition 5.1.** The random vector  $Z = (z_1, \dots, z_k)$  as defined in (5) satisfies the  $(\alpha, \gamma)$ -RDP.

In demonstrating RDP, we take the sample  $x_1, \dots, x_n, x_{n+1}$  and denote:  $S = S(x_1, \dots, x_n)$  and  $S' = S(x_1, \dots, x_{n-1}, x_{n+1})$ . We consider the output distribution of our method when applied to each of the neighboring samples. The event that the ratio of densities fail to meet the requisite bound is a subset of the event where either  $x_{n+1} \in S$  or  $x_n \in S'$ , and when  $2k \leq \gamma n$ . In the complement of this event then the partitions are the same, and the differing samples both fall within the block which receives the Laplace noise, so the DP condition is achieved. In demonstrating the RDP, we simply bound the probability of the aforementioned event, conditional on the order statistics.

*Proof of proposition 5.1.* In the interest of space let the vector of order statistics be denoted  $T = (x_{(1)}, \dots, x_{(n+1)})$ . Let  $S^*(x_1, \dots, x_n, x_{n+1}) = \left\{ j : \sum_{i=1}^{n+1} \mathbf{1}\{x_i = j\} \leq 1 \right\}$ .



We have that  $S, S' \subseteq S^*$ . We thus have

$$\mathbb{P}(x_n \in S' \text{ or } x_{n+1} \in S|T) \leq \mathbb{P}(x_n \in S^* \text{ or } x_{n+1} \in S^*|T).$$

The latter probability is just the fraction of ways in which the order statistics may be rearranged so that  $x_n, x_{n+1}$  fall within  $S^*$ . Due to the condition  $2k \leq \gamma n$ , we have  $|S^*| \leq k \leq \frac{\gamma n}{2}$ . Therefore the number of rearrangements having at least one of  $x_n$  or  $x_{n+1}$  in  $S^*$  is bounded above

$$\mathbb{P}(x_n \in S^* \text{ or } x_{n+1} \in S^*|T) \leq \frac{2|S^*|}{n+1} < \gamma.$$

Therefore

$$\begin{aligned} \mathbb{P}(x_n \in S' \text{ or } x_{n+1} \in S) &\leq \int_{\mathcal{X}^{n+1}} \mathbb{P}(x_n \in S' \text{ or } x_{n+1} \in S|T) dP(T) \\ &\leq \int_{\mathcal{X}^{n+1}} \mathbb{P}(x_n \in S^* \text{ or } x_{n+1} \in S^*|T) dP(T) \\ &< \gamma \int_{\mathcal{X}^{n+1}} dP(T) \\ &= \gamma. \end{aligned}$$

Finally:

$$\begin{aligned} \mathbb{P}\left(\forall Z \subseteq \mathcal{Z}, e^{-\alpha} \leq \frac{Q_n(Z|X)}{Q_n(Z|X')} \leq e^\alpha\right) &= 1 - \mathbb{P}(x_n \in S' \text{ or } x_{n+1} \in S) \\ &> 1 - \gamma. \end{aligned}$$

□

## 5.1 Accuracy

Here we show that  $\delta(z)$  from (2) is close to  $\theta$  even when the histogram is sparse.

**Theorem 5.2.** *Suppose that  $2k \leq \gamma n$ . Let  $\theta^n(x_1, \dots, x_n) = (\theta_1, \dots, \theta_r, 0, \dots, 0)$  for some  $1 \leq r < k$ . Then  $\|\theta - \delta(z)\|_1 = O_P(r/\alpha n)$ .*

*Proof.* Let  $L_1, \dots, L_r \sim \text{Laplace}$ . Let  $\mathcal{E}$  be the event that  $L_j > -\frac{n\alpha}{2}\theta_j$  for all  $1 \leq j \leq r$ . Then  $\mathcal{E}$  holds, except on a set of exponentially small probability. Suppose  $\mathcal{E}$  holds. Let  $W = \sum_{j=1}^r L_j = O_P(r)$ . For  $1 \leq j \leq r$ ,  $z_j = \left(\theta_j + (2L_j)/(n\alpha)\right)$ . For  $j > r$ ,  $z_j = \theta_j = 0$ . Hence  $\|z - \theta\|_1 = O_P(r/\alpha n)$ . Furthermore,  $\|\delta(z) - z\|_1 \leq \frac{r}{n} \leq \frac{r}{\alpha n}$ . Hence via the triangle inequality we have  $\|\delta(z) - \theta\|_1 = O_P(r/\alpha n)$ . □

We thus have a technique for which the risk is uniformly bounded above by  $O(k/\alpha n)$ , as with the DP technique, and which also enjoys the coordinate-wise upper bound on the risk. However in this regime, the risk is no longer uniformly lower bounded with a rate linear in  $k$ , since the upper bound is linear in  $r$  in the case of sparse vectors.

## 6 RDP via Sensitivity Analysis

We next demonstrate that RDP allows schemes for release of other kinds of statistics (besides histograms). A common technique used to establish a differentially private technique is to use Laplace noise with variance proportional to the “global sensitivity” of the function [8]. We show that there is an analog of this technique for RDP. We next demonstrate a method for the RDP release of an arbitrary function  $g_n(x_1, \dots, x_n) \in \mathbb{R}$ .

We consider the algorithm which samples the distribution

$$Q_n(z|x_1, \dots, x_n) \propto \exp \left\{ \frac{-\alpha |z - g_n(x_1, \dots, x_n)|}{s_n(x_1, \dots, x_n)} \right\}. \quad (6)$$

It is well known that when  $s_n$  is the constant function which gives an upper bound of the global sensitivity [8] of  $g_n$ , this method enjoys the  $\alpha$ -DP. As we allow  $s_n$  to depend on the data we may make use of the local sensitivity framework of [15]. There it is demonstrated that whenever:

$$\forall X \sim X' \quad s_n(X) \leq e^\beta s_n(X') \quad (7)$$

and

$$\forall X \quad \sup_{X' \sim X} |g_n(X) - g_n(X')| \leq s_n(X), \quad (8)$$

then (6) gives  $(2\alpha, \eta)$ -DP with:

$$\eta = e^{-\frac{\alpha}{2\beta}} \quad (9)$$

(see [15] Definition 2.1, Lemma 2.5 and Example 3). In moving from DP to RDP we may now require that conditions (7) and (8) hold only with the requisite probability  $1 - \gamma$ . Then (6) will achieve  $(2\alpha, \eta, \gamma)$ -RDP.

We consider a special subset of functions for which:

$$\sup_{X \sim X'} |g_n(X) - g_n(X')| = n^{-1} \sup_{x, x'} h(x, x').$$

Examples of functions satisfying this property are e.g., statistical point estimators [16] and regularized logistic regression estimates [5]. In particular in these cases it is assumed that  $\mathcal{X}$  is some compact subset of  $\mathbb{R}^d$  and then e.g.,  $\sup_{x, x'} h(x, x') = \|x - x'\|_2$  gives the diameter of this set.

We replace conditions (7) and (8) with:

$$P(s_n(X) \leq e^\beta s_n(X')) \geq 1 - \gamma_1 \quad (10)$$

and

$$P(n^{-1}h(x, x') \leq \min\{s_n(X), s_n(X')\}) \geq 1 - \gamma_2. \quad (11)$$

Note that  $x, x'$  are random draws from  $P$  which are independent of the random vectors  $X, X'$ . The first condition simply requires (7) to hold except on a set of measure  $\gamma_1$ . The second condition implies that both  $s_n(X)$  and  $s_n(X')$  give upper bounds to the local sensitivity, except on a set of measure  $\gamma_2$ . Putting these together along with the above considerations will yield a  $(2\alpha, \eta, \gamma_1 + \gamma_2)$ -RDP method. We note that we are essentially asking that  $s_n(X)$  and  $s_n(X')$  both give valid quantiles for the random variable  $h(x, x')$ , and that they give similar values with high probability.

We consider the empirical process based on  $h$  and the data sample  $X$  given by:

$$D(X, t) = \frac{2}{n} \sum_{i=1}^{n/2} \mathbf{1} \{h(x_i, x_{i+n/2}) \leq t\}.$$

This is exactly an empirical CDF for the distribution of  $h(x, x')$ , based on  $n/2$  independent samples of  $h(x, x')$ . We may anticipate that sample quantiles of this empirical CDF will be close to the quantiles from the true CDF, which we denote by  $H(t) = P(h \leq t)$ . This is made precise by the DKW inequality (see e.g., [14]), which in this case yields:

$$P\left(\sup_t |H(t) - D(X, t)| \geq \epsilon\right) \leq 2e^{-n\epsilon^2}. \quad (12)$$

Thus taking  $d_\delta(X)$  to be the smallest  $d$  with  $D(X, d) = 1 - \delta$ , and  $h_{\delta'}$  to give the  $1 - \delta'$  quantile of  $h$ , with  $\delta < \delta'$ , we have:

$$\begin{aligned} P(h(x, x') > d_\delta) &\leq \delta' + P(d_\delta(X) < h_{\delta'}) \\ &\leq \delta' + 2e^{-(\delta' - \delta)^2 n}. \end{aligned}$$

The second inequality comes from applying the monotone function  $D(X, \cdot)$  to both sides of the inequality statement in the probability, and then rearranging, to yield  $P(D(X, h_{\delta'}) - H(h_{\delta'}) > (\delta' - \delta))$ , which is bounded due to the DKW inequality (12). Thus for some appropriate choice of  $\delta, \delta'$  we may take  $s_n(X) = n^{-1}d_\delta(X)$ , and thus achieve (11).

Now to achieve (10) we turn to the Bahadur-Kiefer representation of sample quantiles (see [12]). We have that:

$$d_\delta(X) - h_\delta = \frac{D(X, h_\delta) - H(h_\delta)}{H'(h_\delta)} + O_p(n^{-3/4})$$

where  $H'$  is the derivative of  $H$  (namely the density). Hence we concentrate on the case when  $h$  is a continuous random variable. We find the ratio to be bounded in probability:

$$\frac{d_\delta(X)}{d_\delta(X')} \leq 1 + \frac{|d_\delta(X) - d_\delta(X')|}{d_\delta(X')} = 1 + \frac{O_p(n^{-1/2})}{h_\delta + O_p(n^{-1/2})}$$

where the final equality stems from using DKW to bound the  $D(X, h_\delta) - H(h_\delta)$  and along with the triangle inequality to bound  $|D(X, h_\delta) - D(X', h_\delta)|$ . This therefore demonstrates that:

$$\frac{d_\delta(X)}{d_\delta(X')} \leq 1 + O_p(n^{-1/2}) = O_p(e^{n^{-1/2}}).$$

This means that for large enough  $n$ , and some probability  $1 - \gamma_2$ , the ratio is bounded by  $e^\beta$  where  $\beta$  is polynomial in  $n^{-1/2}$ . Examining (9) we find  $\eta$  to be negligible for such a choice of  $\beta$ . Therefore the use of  $s_n(X) = n^{-1}d_\delta$  achieves the RDP as required.

We note that in principle this same approach would work, were we to replace  $D(X, t)$  with the U-statistic process:

$$U(X, t) = \frac{1}{\binom{n}{2}} \sum_{i>j} \mathbf{1}\{h(x_i, x_j) \leq t\}.$$

Though this is essentially another empirical CDF, it is based on non-independent samples since each  $x_i$  participates in  $n - 1$  of the evaluations of  $h$ . Nevertheless, an analog of the DKW inequality still applies to this process, and we still have the same behavior of the quantiles (see e.g., [1]).

## 7 Privacy Concerns

As stated above, we mainly use random differential privacy as a vehicle for a theoretical exploration of the boundaries of differential privacy. Although it is a conceptually reasonable weakening of differential privacy, whether it is appropriate to use in practice requires more attention. For example, if the hypothesized adversary (of e.g., [17] Theorem 2.4), really had access to a subset of  $n - 1$  of the data, and the one remaining element was the only inhabitant of its histogram cell, then this would be immediately revealed to the adversary. Whether this is a critical problem depends on the application.

## 8 Example

We present two examples in which the RDP technique and the DP techniques are compared on synthetic histogram data. In the first example the histogram has  $k = 25$  bins, all but two of which are empty, and  $n = 500$  points fall in to the other two.

Figure 1(a) shows the original data as well as the sanitized data due to differential privacy and RDP. Figure 1(b) shows the distribution of  $L_1$  loss from 100 simulations of both approaches. We see that the risk of the RDP histogram is typically much lower than that of the DP histogram, which occasionally has risk in excess of 0.5 (recall that the maximum possible loss is 2 in the case that the original and sanitized histograms had completely disjoint support).

We present an analogous two-dimensional example in Figure 2. Here the histogram has  $k = 400$  bins in which all but 16 are empty. In this example we see that the RDP technique has uniformly better loss than the DP technique.

## 9 Conclusion

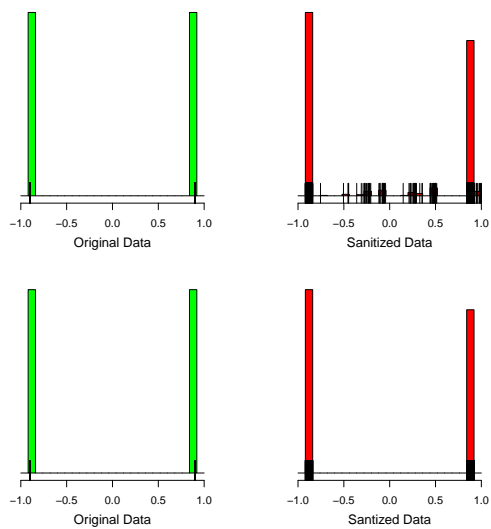
We have introduced a relaxed version of differential privacy—random differential privacy—shown how to apply it to histograms, and examined the accuracy of the resulting method. We also demonstrated some properties of our definition, and explained a basic construction for release of arbitrary functions of the data. As we mentioned in the introduction, we are not suggesting that differential privacy should be abandoned and replaced by random differential privacy. However, we do think it is fruitful to consider various relaxations of differential privacy to gain a deeper understanding of the trade-offs between the strength of the privacy guarantee and the accuracy of the data release mechanism.

In ongoing work we are extending this work to allow for data dependent choices of the number of bins and to allow for other density estimators besides histograms. We are also considering other relaxations of differential privacy. We will report on these results in future work.


## References

- [1] Arcones, M. (1996). The Bahadur-Kiefer representation for U-quantiles. *The Annals of Statistics*, 24(3):1400–1422.
- [2] Barak, B., Chaudhuri, K., Dwork, C., Kale, S., McSherry, F., and Talwar, K. (2007). Privacy, accuracy, and consistency too: A holistic solution to contingency table release. In *Proceedings of the 26<sup>th</sup> ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 273–282.
- [3] Blum, A., Dwork, C., McSherry, F., and Nissim, K. (2005). Practical privacy: the SuLQ framework. In *Proceedings of the 24<sup>th</sup> ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 128–138.
- [4] Charest, A.-S. (2012). Empirical evaluation of statistical inference from differentially-private contingency tables. *Privacy in Statistical Databases*, 257–272.
- [5] Chaudhuri, K. and Monteleoni, C. (2008). Privacy preserving logistic regression. In *Proceedings of the 22<sup>nd</sup> Annual Conference on Neural Information Processings Systems (NIPS 2008)*.
- [6] Dwork, C. (2006). Differential privacy. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*. 1–12.
- [7] Dwork, C. and Lei, J. (2009). Differential privacy and robust statistics. In *Proceedings of the 41<sup>st</sup> ACM Symposium on Theory of Computing*. 371–380.
- [8] Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3<sup>rd</sup> Theory of Cryptography Conference*. 265–284.
- [9] Dwork, C., McSherry, F., and Talwar, K. (2007). The price of privacy and the limits of LP decoding. In *Proceedings of the 39<sup>th</sup> Annual ACM symposium on Theory of computing*. 85–94.
- [10] Fienberg, S. E., Rinaldo, A., and Yang, X. (2010). Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. *Privacy in Statistical Databases*, 197–199.
- [11] Hardt, M. and Talwar, K. (2010). On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC '10)*. 705–714.
- [12] Kiefer, J. (1967). On Bahadur’s representation of sample quantiles. *The Annals of Mathematical Statistics*, 38(5):1323–1342.
- [13] Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., and Vilhuber, L. (2008). Privacy: Theory meets practice on the map. In *Proceedings of the 24<sup>th</sup> International Conference on Data Engineering*. 277–286.

- [14] Massart, P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, 18(3).
- [15] Nissim, K., Raskhodnikova, S., and Smith, A. (2007). Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39<sup>th</sup> Annual ACM Symposium on Theory of Computing*. 75–84.
- [16] Smith, A. (2011). Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the 43<sup>rd</sup> Annual ACM symposium on Theory of Computing (STOC '11)*. New York: ACM. 813–822.
- [17] Wasserman, L. and Zhou, S. (2010). A statistical framework for differential privacy. *The Journal of the American Statistical Association*, 105:375–389.
- [18] Yang, X., Fienberg, S. E., and Rinaldo, A. (2012). Differential privacy for protecting multi-dimensional contingency table data: Extensions and applications. *Journal of Privacy and Confidentiality*, 4(1): 101–125.
- [19] Yu, B. (1997). Assouad, fano, and le cam. In D. Pollard, E. Torgersen, and G. Yang (eds.), *Festschrift for Lucien Le Cam*. Springer. 423–435.



(a) Original and synthetic data for DP (top) and RDP (bottom)



(b) Empirical error distribution for DP (top) and RDP (bottom)

Figure 1: A one dimensional example.






Figure 2: Empirical error distributions for a two dimensional histogram, displayed in the top left.

