Canonical Noise and Private Hypothesis Tests

Jordan Awan
Purdue University
jawan@purdue.edu

Salil Vadhan
Harvard University
salil_vadhan@harvard.edu

Abstract

In the setting of f-DP, we propose the concept *canonical noise distribution* (CND) which captures whether an additive privacy mechanism is tailored for a given f, and give a construction of a CND for an arbitrary tradeoff function f. We show that private hypothesis tests are intimately related to CNDs, allowing for the release of private p-values at no additional privacy cost as well as the construction of uniformly most powerful (UMP) tests for binary data. We apply our techniques to difference of proportions testing.

CCS Concepts: • **Security and privacy** → **Privacy protections**; **Usability in security and privacy**; Social aspects of security and privacy.

Keywords: differential privacy, p-values, frequentist inference, optimal mechanism

ACM Reference Format:
Jordan Awan and Salil Vadhan. 2021. Canonical Noise and Private Hypothesis Tests. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

In this paper, we study two basic and fundamental privacy questions in the framework of f-DP. The first is based on the basic mechanism of adding independent noise to a real-valued statistic, and the second is about the nature of hypothesis tests under DP. We show that in fact, the two problems are intricately related, where the "canonical" additive noise distribution enables private p-values "for free," and gives a closed form construction of certain optimal hypothesis tests.

One of the most basic and fundamental privacy mechanisms is an additive mechanism, where independent noise is added to a real-valued statistic. The Laplace mechanism appeared along with the original definition of DP, and the Gaussian mechanism was another one of the earliest privacy mechanism designed for approximate DP. A natural question is what noise distributions are "optimal" or "canonical" for a given definition of privacy. The geometric mechanism/discrete Laplace mechanism is optimal for count data in terms of maximizing Bayesian utility [8], the staircase mechanism is optimal for ϵ-DP in terms of ℓ_2-error [7], and the truncated-uniform-Laplace (Tulap) distribution generalizes both the discrete Laplace and staircase mechanisms and is optimal for (ϵ, δ)-DP in terms of generating uniformly most powerful (UMP) hypothesis tests for Bernoulli data [1, 2]. In this paper, we give a formal definition of a *canonical noise distribution* (CND) which is applicable to any f-DP notion of privacy. We show that the Gaussian distribution is canonical for Gaussian differential privacy (GDP), and the Tulap distribution is canonical for (ϵ, δ)-DP. We prove that a CND always exists for any symmetric tradeoff function f, and give a construction to generate a CND given an arbitrary tradeoff function f. In fact, this construction results in the Tulap distribution in the case of (ϵ, δ)-DP.

Another basic privacy question is on the nature of DP hypothesis tests. Awan and Slavković [1] and Awan and Slavković [2] showed that for independent Bernoulli data, there exists uniformly most powerful (UMP) (ϵ, δ)-DP tests which are based on the Tulap distribution, enabling "free" private p-values, at no additional cost to privacy. We show that in general, given any f-DP test, a free private p-value can always be generated in terms of a CND for f. We also expand the main results of Awan and Slavković [1] from (ϵ, δ)-DP to f-DP as well as from i.i.d. Bernoulli variables to exchangeable binary data. This expansion shows that the CND is the proper analogue of the Tulap distribution, and gives an explicit construction of the most powerful f-DP test for binary data. Finally, we apply our results to private difference of proportions testing, available in the full paper.

2 Differential privacy

All of the major variants of DP state that given a randomized algorithm M, for any two adjacent databases X, X', the distributions of $M(X)$ and $M(X')$ should be "similar." While many DP variants measure similarity in terms of divergences, recently Dong et al. [5] proposed f-DP, which formalizes similarity in terms of constraints on hypothesis tests. We say that X and X' are *adjacent* if $H(X, X') \leq 1$, where H is the *Hamming metric*.

For two probability distributions P and Q, the *tradeoff function* $T(P, Q) : [0, 1] \to [0, 1]$ is defined as $T(P, Q)(\alpha) =$

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

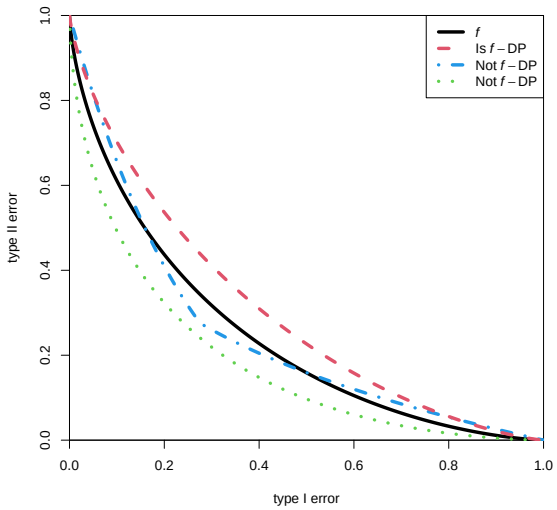

Figure 1. A plot of three examples of $T(M(D), M(D'))$. Only the red, dashed tradeoff curve satisfies f-DP.

$\inf\{1 - \mathbb{E}_Q\phi \mid \mathbb{E}_P(\phi) \leq \alpha\}$, where the infinimum is over all measurable tests ϕ. The tradeoff function can be interpreted as follows: If $T(P, Q)(\alpha) = \beta$, then the most powerful test ϕ which is trying to distinguish between $H_0 = \{P\}$ and $H_1 : \{Q\}$ at type I error $\leq \alpha$ has type II error β. A larger tradeoff function means that it is harder to distinguish between P and Q. A function $f : [0, 1] \to [0, 1]$ is a tradeoff function if and only if f is convex, continuous, decreasing, and $f(x) \leq 1 - x$ for all $x \in [0, 1]$ [5, Proposition 2.2]. We say that a tradeoff function f is *nontrivial* if f is not identically equal to $1 - \alpha$.

Definition 2.1 (f-DP)**.** Let f be a tradeoff function. A mechanism M satisfies f-DP if

$$T(M(D), M(D')) \geq f$$

for all $D, D' \in \mathcal{X}^n$ such that $H(D, D') \leq 1$.

See Figure 1 for examples of tradeoff functions which do and do not satisfy f-DP for a particular f. In the above definition, the inequality $T(M(D), M(D')) \geq f$ is shorthand for $T(M(D), M(D'))(\alpha) \geq f(\alpha)$ for all $\alpha \in [0, 1]$. Without loss of generality we can assume that f is symmetric: $f(\alpha) = f^{-1}(\alpha)$, where $f^{-1}(\alpha) = \inf\{t \in [0, 1] \mid f(t) \leq \alpha\}$ [5, Proposition 2.4].

Wasserman and Zhou [14] and Kairouz et al. [11] both showed that (ϵ, δ)-DP can be expressed in terms of hypothesis testing, and in fact Dong et al. [5] showed that (ϵ, δ)-DP can be expressed as a special case of f-DP.

Definition 2.2 ((ϵ, δ)-DP)**.** Let $\epsilon > 0$ and $\delta \geq 0$, and define $f_{\epsilon,\delta}(\alpha) = \max\{0, 1 - \delta - \exp(\epsilon)\alpha, \exp(-\epsilon)(1 - \delta - \alpha)\}$. A mechanism M satisfies (ϵ, δ)*-DP* if it satisfies $f_{\epsilon,\delta}$-DP.

Another notable special case of f-DP is Gaussian DP.

Definition 2.3 (Gaussian differential privacy)**.** For $\mu > 0$,

$$G_\mu(\alpha) := T(N(0, 1), N(\mu, 1))(\alpha) = \Phi(\Phi^{-1}(1 - \alpha) - \mu),$$

where Φ is the cdf of $N(0, 1)$. A mechanism M satisfies μ*-Gaussian differential privacy* (μ-GDP) if it is G_μ-DP.

3 Canonical noise distributions

One of the most basic techniques of designing a privacy mechanism is through adding data-independent noise. The earliest DP mechanisms add either Laplace or Gaussian noise, and there have since been several works developing optimal additive mechanisms including the geometric (discrete Laplace) [8], truncated-uniform-Laplace (Tulap) [1, 2], and staircase mechanisms [7]. There have also been several works exploring multivariate and infinite-dimensional additive mechanisms such as K-norm [3, 10], elliptical perturbations [13], and Gaussian processes [9, 12].

While there are many choices of additive mechanisms to achieve f-DP, we are interested in adding the least noise necessary in order to maximize the utility of the output. Rather than measuring the amount of noise by its variance or entropy, we focus on whether the privacy guarantee is tight.

In this section, we introduce the concept *canonical noise distribution* (CND), which captures whether a real-valued distribution is perfectly tailored to satisfy f-DP. We formalize this in Definition 3.1. We then show that for any symmetric f, we can always construct a CND, where the construction is given in Definition 3.4 and proved to be a CND in Theorem 3.5. We will see in Section 4 that CNDs are fundamental for understanding the nature of f-DP hypothesis tests, for constructing "free" DP p-values, and for the design of uniformly most powerful f-DP tests for binary data.

Before we define canonical noise distribution, we must introduce the *sensitivity* of a statistic, a central concept of DP [6]. A statistic $T : \mathcal{X}^n \to \mathbb{R}$ has *sensitivity* $\Delta > 0$ if $|T(X) - T(X')| \leq \Delta$ for all $H(X, X') \leq 1$. As the sensitivity measures how much a statistic can change when one person's data is modified, additive noise must be scaled proportionally to the sensitivity in order to protect privacy.

Definition 3.1. Let f be a symmetric tradeoff function. A cdf F is a *canonical noise distribution* (CND) for f if

1. given a statistic $S(X)$ with sensitivity $\Delta > 0$, and $N \sim F(\cdot)$, the mechanism $S(X) + \Delta N$ satisfies f-DP. Equivalently, for any $m \in (0, 1)$, $T(F(\cdot), F(\cdot - m)) \geq f$,
2. $f(\alpha) = T(F(\cdot), F(\cdot - 1))(\alpha)$ for all α,
3. $T(F(\cdot), F(\cdot - 1))(\alpha) = F(F^{-1}(1 - \alpha) - 1)$ for all α,
4. $F(x) = 1 - F(-x)$ for all $x \in \mathbb{R}$; that is, F is the cdf of a random variable which is symmetric about zero.

The most important conditions of Definition 3.1 are 1 and 2, which state that the distribution can be used to satisfy f-DP and that the privacy bound is tight. Condition 3 of Definition 3.1 gives a closed form for the tradeoff function,

and is equivalent to requiring that the optimal rejection set for discerning between $F(\cdot)$ and $F(\cdot-1)$ is of the form (x,∞) for some $x \in \mathbb{R}$. The last condition of Definition 3.1 enforces symmetry of the distribution, which makes CNDs much easier to work with.

It is easy to show that $\Phi(\mu\cdot)$, the cdf of $N(0,1/\mu^2)$ is a CND for G_μ.

Proposition 3.2. *Let f be a symmetric tradeoff function. Let F be a CND for f, and G be another cdf such that $T(G(\cdot), G(\cdot-1)) \geq f$. Let $N \sim F$ and $M \sim G$. Then there exists a randomized function* $\text{Proc} : \mathbb{R} \to \mathbb{R}$ *which satisfies* $\text{Proc}(N) \stackrel{d}{=} M$ *and* $\text{Proc}(N+1) \stackrel{d}{=} M+1$, *where "$\stackrel{d}{=}$" means* equal in distribution.

Proposition 3.2 follows from property 2 in Definition 3.1 along with Dong et al. [5, Theorem 2.10], which was originally a result of Blackwell [4]. Proposition 3.2 shows that if we add noise from a CND, we can post-process to obtain the same result as if we added noise from another distribution. This shows in a very general sense that a CND adds the least noise necessary to achieve f-DP.

In the remainder of this section, we show that given any tradeoff function f, we can always construct a canonical noise distribution (CND), but that a CND need not be unique.

Lemma 3.3. *Let f be a symmetric tradeoff function and let F be a CND for f. Then $F(x) = 1-f(F(x-1))$ when $F(x-1) > 0$ and $F(x) = f(1-F(x+1))$ when $F(x+1) < 1$.*

In the Lemma 3.3, we see that a CND satisfies an interesting recurrence relation. If we know the value $F(x) = c$ for some $x \in \mathbb{R}$ and $c \in (0,1)$, then we know the value of $F(y)$ for all $y \in \mathbb{Z} + x$. This means that if we specify F on an interval of length 1, such as $[-1/2, 1/2]$, then F is completely determined by the recurrence relation. We will leverage this fact to construct a CND from scratch, by specifying a linear function on $[-1/2, 1/2]$. The remainder of this section is devoted to the construction of a CND and the proof that it has the properties of Definition 3.1.

Definition 3.4. Let f be a symmetric tradeoff function, and let $c \in [0,1]$ be the unique fixed point of f: $f(c) = c$. We $F_f : \mathbb{R} \to \mathbb{R}$ as

$$F_f(x) = \begin{cases} f(1-F_f(x+1)) & x < -1/2 \\ c(1/2-x) + (1-c)(x+1/2) & -1/2 \leq x \leq 1/2 \\ 1-f(F_f(x-1)) & x > 1/2. \end{cases}$$

Note that in Definition 3.4, on $[-1/2, 1/2]$ the CDF corresponds to a uniform random variable, but then due to the recursive nature of F_f and the fact that f is in general nonlinear, the CND of Definition 3.4 is in general not uniformly distributed on any other intervals. See Figure 2 for a plot of the pdf of the CND of Definition 3.4 corresponding to the tradeoff function G_1.

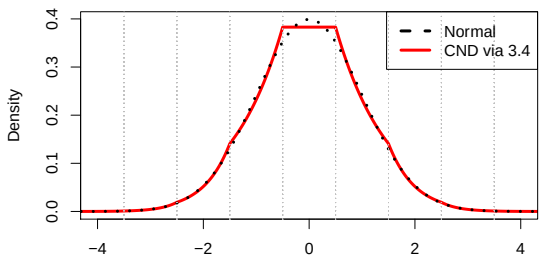

Figure 2. Density plots of $N(0,1)$ as well as the CND of Definition 3.4 for the tradeoff function G_1.

Theorem 3.5 below states that for any nontrivial tradeoff function, this construction yields a canonical noise distribution, which can be constructed as in Definition 3.1. This CND can be used to add perfectly calibrated noise to a statistic to achieve f-DP. As we will see later, the existence (and construction) of a CND will enable us to prove that any f-DP test can be post-processed from a test statistic, and that this implies that we can always obtain hypothesis testing p-values at no additional privacy cost, a generalization of the result of Awan and Slavković [1], which previously only held for (ϵ,δ)-DP and for Bernoulli data.

Theorem 3.5. *Let f be a nontrivial symmetric tradeoff function and let F_f be as in Definition 3.4. Then F_f is a canonical noise distribution for f.*

It turns out that the requirements of Definition 3.1 do not uniquely determine a distribution. For instance, Φ the cdf of a standard normal is a CND for 1-GDP, but Φ is different from the construction in Definition 3.4. See Figure 2 for the pdf of these two CNDs. Note that the CND of Definition 3.4 can be seen to be uniform in $[-1/2, 1/2]$ and has "kinks" at each half-integer value. On the other hand, the standard normal is smooth.

What we have developed in this section is a constructive and general method of generating canonical noise distributions for f-DP. In the special case of (ϵ,δ)-DP, the CND F_f is equal to the cdf of the Tulap distribution, proposed in Awan and Slavković [1], which is an extension of the Staircase mechanism Geng and Viswanath [7] from $(\epsilon, 0)$-DP to (ϵ,δ)-DP.

4 The nature of f-DP tests

A *test* is a function $\phi : \mathcal{X}^n \to [0,1]$, where $\phi(x)$ represents the probability of rejecting the null hypothesis given the observation x. The mechanism that implements this test releases a random value drawn as $\text{Bern}(\phi(x))$, where 1 represents "Reject" and 0 represents "Accept." We say that the

test ϕ satisfies f-DP if $\text{Bern}(\phi(x))$ satisfies f-DP. Lemma 4.1 shows that a test satisfies f-DP if for adjacent databases x and x', the values $\phi(x)$ and $\phi(x')$ are close in terms of an inequality based on f.

Lemma 4.1. *A test $\phi : \mathcal{X}^n \to [0,1]$ satisfies f-DP if and only if $\phi(x) \le 1 - f(\phi(x'))$ for all $H(x,x') \le 1$.*

Lemma 4.1 simplifies the search for f-DP hypothesis tests and generalizes the bounds on private tests established in Awan and Slavković [1].

The result of Lemma 4.1 can also be expressed in terms of canonical noise distributions in Corollary 4.2, giving the elegant relation that $F^{-1}(\phi(x))$ and $F^{-1}(\phi(x'))$ differ by at most 1 when x and x' are adjacent.

Corollary 4.2 (Canonical Noise Distributions)**.** *Let f be a symmetric nontrivial tradeoff function and let F be a canonical noise distribution for f. Then a test ϕ satisfies f-DP if and only if $F^{-1}(\phi(x)) \le F^{-1}(\phi(x')) + 1$ for all $H(x,x') \le 1$.*

As we will see, Corollary 4.2 is also important for the construction of free DP p-values in Section 4.1.

4.1 Free f-DP p-values

In Awan and Slavković [1], it was shown that for Bernoulli data, the uniformly most powerful DP test could also be expressed as the post-processing of a private p-value, offering p-values at no additional privacy cost. We generalize this result using the concept of canonical noise distributions and show that for f-DP test can be expressed as a post-processing threshold test based on a private test statistic, and that the test statistic can also be used to give private p-values.

Theorem 4.3. *Let $\phi : \mathcal{X}^n \to [0,1]$ be an f-DP test. Let F be a CND for f, and draw $N \sim F$. Then*

1. *releasing $T = F^{-1}(\phi(x)) + N$ satisfies f-DP,*
2. *the variable $Z = I(T \ge 0)$, a post-processing of T is distributed as $Z \mid X = x \sim \text{Bern}(\phi(x))$,*
3. *the value $p = \sup_{\theta_0 \in H_0} \mathbb{E}_{X \sim \theta_0} F(F^{-1}(\phi(X)) - T)$ is also a post-processing of T and is a p-value for H_0,*
4. *if H_0 is a simple hypothesis and $\mathbb{E}_{H_0}\phi = \alpha$, then at type I error α, the p-value from part 3. is as powerful as ϕ at every alternative.*

We see from Theorem 4.3 that given an f-DP test ϕ, we can report both a summary statistic (namely, T) as well as a p-value (a post-processing of T) which contain strictly more information than only sampling $\text{Bern}(\phi(x))$. This shows that for simple null hypotheses, there is no general privacy amplification when post-processing a p-value or test statistic to a binary accept/reject decision.

Note that Theorem 4.3 starts with an f-DP test, and shows how to get a private summary statistic and p-values. However, constructing a private test ϕ is another matter. In Section 4.2, we show that for exchangeable binary data, we can construct a most powerful f-DP test in terms of a CND.

4.2 Most powerful tests for binary data

In this section, we extend the main result of Awan and Slavković [1], that of constructing most powerful DP tests, to general f-DP as well as exchangeable distributions on $\{0,1\}^n$. In contrast, the hypothesis tests of Awan and Slavković [1] were limited to (ϵ, δ)-DP and i.i.d. Bernoulli data. A distribution P on a set $\mathcal{X}^n$ is *exchangeable* if given $\underline{X} \sim P$ and a permutation π, $\underline{X} \overset{d}{=} \pi(\underline{X})$. Note that i.i.d. data are always exchangeable, but there are exchangeable distributions that are not i.i.d.

Theorem 4.4. *Let f be a symmetric nontrivial tradeoff function and let F be a CND of f. Let $\mathcal{X} = \{0,1\}$. Let P and Q be two exchangeable distributions on $\mathcal{X}^n$ such that $\frac{dQ}{dP}$ is a increasing function of $X = \sum_{i=1}^n X_i$. Let $\alpha \in (0,1)$. Then a most powerful f-DP test ϕ with level α for $H_0 : X \sim P$ versus $H_1 : X \sim Q$ can be expressed in any of the following forms:*

1. *There exists $y \in \{0,1,2,\dots,n\}$ and $c \in (0,1)$ such that for all $x \in \{0,1,2,\dots,n\}$,*

$$\phi(x) = \begin{cases} 0 & x < y, \\ c & x = y, \\ 1 - f(\phi(x-1)) & x > y, \end{cases}$$

   where if $y > 0$ then c satisfies $c \le 1 - f(0)$, and c and y are chosen s.t. $\mathbb{E}_P\phi(x) = \alpha$. If $f(0) = 1$, then $y = 0$.
2. *$\phi(x) = F(x - m)$, where m satisfies $\mathbb{E}_P\phi(x) = \alpha$.*
3. *Let $N \sim F$. The variable $T = X + N$ satisfies f-DP. Then $p = \mathbb{E}_{X \sim P} F(X - T)$ is a p-value and $I(p \le \alpha) \mid X = I(T \ge m) \mid X \sim \text{Bern}(\phi(X))$.*

While Theorem 4.3 took an f-DP test and produced "free" private p-values, Theorem 4.4 constructs an optimal test from scratch beginning only with a CND.

5 Difference of Proportions Tests

Testing two population proportions is a very basic and common hypothesis testing setting that in any when there are two groups with binary responses, such as A/B testing, clinical trials, and observational studies. While there does not exist a UMP unbiased f-DP test for this problem, using our results on DP most powerful tests for binary data, we show that there does exist a UMP unbiased "semi-private" test, which satisfies a weakened version of f-DP. While this test cannot be used for privacy purposes, it does provide an upper bound on the power of any f-DP test, and gives intuition on the structure of an optimal f-DP test for this problem. Using this as a benchmark we propose a private test, which allows for optimal inference for the two population parameters and is nearly as powerful as the semi-private UMPU. In the special case of $(\epsilon, 0)$-DP, we show empirically that our proposed test is more powerful than any $(\epsilon/\sqrt{2})$-DP test and has more accurate type I errors than the classic normal approximation test. The details and results are available in the full paper.

References

[1] Jordan Awan and Aleksandra Slavković. 2018. Differentially private uniformly most powerful tests for binomial data. *Advances in Neural Information Processing Systems* 31 (2018), 4208–4218.

[2] Jordan Awan and Aleksandra Slavković. 2020. Differentially Private Inference for Binomial Data. *Journal of Privacy and Confidentiality* 10, 1 (2020).

[3] Jordan Awan and Aleksandra Slavković. 2020. Structure and sensitivity in differential privacy: Comparing k-norm mechanisms. *J. Amer. Statist. Assoc.* (2020), 1–20.

[4] David Blackwell. 1950. *Comparison of experiments.* Technical Report. Howard University Washington United States.

[5] Jinshuo Dong, Aaron Roth, and Weijie J Su. 2019. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383* (2019).

[6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*. Springer, 265–284.

[7] Quan Geng and Pramod Viswanath. 2015. The optimal noise-adding mechanism in differential privacy. *IEEE Transactions on Information Theory* 62, 2 (2015), 925–951.

[8] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2012. Universally utility-maximizing privacy mechanisms. *SIAM J. Comput.* 41, 6 (2012), 1673–1693.

[9] Rob Hall, Alessandro Rinaldo, and Larry Wasserman. 2013. Differential privacy for functions and functional data. *The Journal of Machine Learning Research* 14, 1 (2013), 703–727.

[10] Moritz Hardt and Kunal Talwar. 2010. On the geometry of differential privacy. In *Proceedings of the forty-second ACM symposium on Theory of computing*. 705–714.

[11] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2017. The Composition Theorem for Differential Privacy. *IEEE Trans. Information Theory* 63, 6 (2017), 4037–4049. https://doi.org/10.1109/TIT.2017.2685505

[12] Ardalan Mirshani, Matthew Reimherr, and Aleksandra Slavković. 2019. Formal privacy for functional data with gaussian perturbations. In *International Conference on Machine Learning*. PMLR, 4595–4604.

[13] Matthew Reimherr and Jordan Awan. 2019. Elliptical Perturbations for Differential Privacy. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2019/file/b3dd760eb02d2e669c604f6b2f1e803f-Paper.pdf

[14] Larry Wasserman and Shuheng Zhou. 2010. A Statistical Framework for Differential Privacy. *J. Amer. Statist. Assoc.* 105:489 (2010), 375–389.