

Steven Wu

1 Approximate Differential Privacy

One of the notable features of differential privacy is the *multiplicative* notion of similarity between distributions, that is, we require

$$\mathbb{P}(A(\mathbf{x}) \in E) \leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') \in E) \quad (1)$$

for every event E . Intuitively this requirement seems overly stringent. For example, suppose there is an event E such that $\mathbb{P}(A(\mathbf{x}) \in E) = 0$. Then, for any finite ϵ , differential privacy says that $\mathbb{P}(A(\mathbf{x}') \in E) = 0$ for every neighboring dataset. Is this really necessary? Yes, if $\mathbb{P}(A(\mathbf{x}') \in E) = 10^{-80}$, and we see an outcome in E , then we know with certainty that the input was \mathbf{x}' and not \mathbf{x} . But that doesn't seem like a big problem given that E only occurs with probability one-over-the-number-of-atoms-in-the-universe.

In this lecture, we'll explore a slightly more permissive variant of differential privacy that captures the intuition that these sorts of highly disclosive, but extremely low probability events should be allowed. Of course, there's no reason to consider more permissive variants unless we get something in return, so, not surprisingly, we'll see how we can use this relaxation to get better utility.

1.1 Definition and Properties

Specifically, we will modify the definition of differential privacy to allow for a hybrid additive-multiplicative definition of closeness.

Definition 1.1 ((ϵ, δ) -DP with fixed-size data sets). A randomized algorithm $A : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially private for size n data sets if, for every pair of neighboring data sets \mathbf{x}, \mathbf{x}' , for all $E \subseteq \mathcal{Y}$,

$$\mathbb{P}(A(\mathbf{x}) \in E) \leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') \in E) + \delta \quad (2)$$

Note that $(\epsilon, 0)$ -differential privacy is equivalent to the standard definition of ϵ -differential privacy, but we will typically write $(\epsilon, 0)$ -DP to avoid confusion. We often call (ϵ, δ) -DP *approximate DP*, and $(\epsilon, 0)$ -DP *pure DP*.¹

Intuitively, we think of δ as the probability of a “total privacy failure,” so we want δ to be extremely small, like 2^{-20} or ideally more like 2^{-128} . When δ is small enough, this definition has very similar “semantics” to the standard definition of differential privacy, although with some technical subtleties that we won't focus on for now. The relaxed definition also satisfies many of the same useful properties that are true for $(\epsilon, 0)$ -differential privacy, with a suitable change of parameters to account for δ .

Lemma 1.2. *Approximate differential privacy satisfies the following properties:*

1. *Closure under Post-processing: for any $A : \mathcal{X}^n \rightarrow \mathcal{Y}$ that is (ϵ, δ) -DP, and any post-processing map $B : \mathcal{Y} \rightarrow \mathcal{Y}'$, $B(A(\cdot))$ is also (ϵ, δ) -DP.*

¹Jon doesn't like these names, since they suggest there is something substandard about (ϵ, δ) -DP, but they are convenient.

2. (ϵ, δ) -differential privacy satisfies (adaptive) composition. Running one mechanism satisfying (ϵ_1, δ_1) -DP followed by another mechanism satisfying (ϵ_2, δ_2) -DP satisfies $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -DP.

Exercise 1.3. Prove Lemma 1.2.

Note that, similar to pure DP, approximate DP satisfies a composition property where compositing T mechanisms, each with (ϵ, δ) -DP gives us a combined mechanism that is $(\epsilon T, \delta T)$ -DP. However, unlike pure DP, this composition bound can actually be improved considerably, and we can prove a bound more like $(\epsilon\sqrt{T}, \delta T)$ -DP, although the exact parameters are a bit weaker. This “strong composition” property is one of the most useful things about the approximate version of DP, and the next lecture is devoted to exploring this phenomenon in more detail.

1.1.1 Proving Mechanism’s Satisfy Approximate DP

Approximate DP differs from pure DP in some critical ways though. For example, in proving pure DP, we focused on events E that were just singletons, $E = \{y\}$, and relied on the equivalence,

$$\forall y \in \mathcal{Y} \quad \mathbb{P}(A(\mathbf{x}) = y) \leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') = y) \quad (3)$$

$$\iff \forall E \subseteq \mathcal{Y} \quad \mathbb{P}(A(\mathbf{x}) \in E) \leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') \in E) \quad (4)$$

For approximate DP the analogous statement is *not* true, and we have to consider all sets $E \subseteq \mathcal{Y}$ to establish approximate DP.

However, to prove approximate-DP, it is enough to prove that if we draw y from $A(\mathbf{x})$, then with high probability we will have $\mathbb{P}(A(\mathbf{x}) = y) \leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') = y)$. We can capture this idea with the following useful lemma, that we will often use when we want to prove that a mechanism is (ϵ, δ) -DP.

Lemma 1.4. For a mechanism $A : \mathcal{X}^n \rightarrow \mathcal{Y}$, a pair of datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$, and any $\epsilon > 0$, define the sets

$$\text{Good}_{\mathbf{x}, \mathbf{x}'} = \left\{ y \in \mathcal{Y} : \frac{\mathbb{P}(A(\mathbf{x}) = y)}{\mathbb{P}(A(\mathbf{x}') = y)} \leq e^\epsilon \right\} \quad \text{Bad}_{\mathbf{x}, \mathbf{x}'} = \overline{\text{Good}_{\mathbf{x}, \mathbf{x}'}} \quad (5)$$

If $\mathbb{P}(A(\mathbf{x}) \in \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \leq \delta$ for every pair of neighboring datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$, then A satisfies (ϵ, δ) -DP.

Proof. To prove the statement, fix an arbitrary pair of neighboring datasets $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^n$ and an arbitrary event $E \subseteq \mathcal{Y}$. Then we can calculate:

$$\mathbb{P}(A(\mathbf{x}) \in E) = \mathbb{P}(A(\mathbf{x}) \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}) + \mathbb{P}(A(\mathbf{x}) \in E \cap \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \quad (6)$$

$$\leq \mathbb{P}(A(\mathbf{x}) \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}) + \mathbb{P}(A(\mathbf{x}) \in \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \quad (7)$$

$$\leq \mathbb{P}(A(\mathbf{x}) \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}) + \delta \quad (8)$$

$$= \sum_{y \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}} \mathbb{P}(A(\mathbf{x}) = y) + \delta \quad (9)$$

$$\leq \sum_{y \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}} e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') = y) + \delta \quad (10)$$

$$= e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') \in E \cap \text{Good}_{\mathbf{x}, \mathbf{x}'}) + \delta \quad (11)$$

$$\leq e^\epsilon \cdot \mathbb{P}(A(\mathbf{x}') \in E) + \delta \quad (12)$$

which is what we needed to show. □

One interesting thing to note is that (ϵ, δ) -DP *does not* imply that $\mathbb{P}(A(\mathbf{x}) \in \text{Bad}) \leq \delta$, as you might expect, although a version of this equivalence does hold but with slightly weaker parameters.

1.2 A First Example: Truncated Laplace

We wouldn't bother considering this definition if it didn't allow us to do something useful. In the next section we'll see a much more general feature of approximate DP, but let's start with a very simple example called the *truncated Laplace mechanism*.

Our canonical example of an $(\epsilon, 0)$ -differentially private mechanism for answering some real-valued statistic $f : \mathcal{X}^n \rightarrow \mathbb{R}$ was to compute

$$A(\mathbf{x}) = f(\mathbf{x}) + \frac{\Delta}{\epsilon} \cdot \text{Lap}(1) \quad (13)$$

where Δ is the global sensitivity of f . The noise distribution has standard deviation $O(\Delta/\epsilon)$, but has the undesirable property that the noise can be arbitrarily large. That is, there is some probability that we receive an answer that is complete nonsense. Although we haven't seen the tools to prove it yet, this issue is *inherent* for pure differentially private algorithms, as any $(\epsilon, 0)$ -DP algorithm for answering a single count has to have some small probability of giving an extremely bad answer.

However, (ϵ, δ) -DP can actually be achieved while adding noise from a distribution of bounded support. Specifically, for $\lambda, \tau > 0$, the *truncated Laplace distribution* $\text{Lap}(\lambda, \tau)$ is defined by the probability density function

$$p(y) = \begin{cases} \frac{1}{Z} \cdot e^{-|y|/\lambda} & |y| \leq \tau \\ 0 & |y| > \tau \end{cases} \quad (14)$$

where $Z = \int_{-\tau}^{\tau} e^{-|y|/\lambda} dy$ is a normalizing constant. One can prove that, for some choice of $\tau = O(\log(1/\delta))$, the mechanism

$$A(\mathbf{x}) = f(\mathbf{x}) + \frac{\Delta}{\epsilon} \cdot \text{Lap}(1, \tau) \quad (15)$$

satisfies (ϵ, δ) -differential privacy. This mechanism has the essentially the same standard deviation, but also has a guaranteed worst-case bound on the magnitude of the noise! Intuitively, we're trading a small probability of getting an extremely inaccurate answer for a small probability of getting an answer that compromises privacy.

Exercise 1.5. Prove that the truncated Laplace mechanism in (15) is (ϵ, δ) -differentially private.

2 The Gaussian Mechanism

In this section we'll see that adding noise from a *Gaussian* distribution, rather than a Laplace distribution, satisfies approximate DP, and will in fact can give sometimes much better accuracy for functions that output high-dimensional vectors.

2.1 Univariate Gaussian Noise

Recall that the Gaussian distribution with mean μ and standard deviation σ , denoted $\mathcal{N}(\mu, \sigma^2)$, is defined by the probability density function

$$p_{\mu, \sigma^2}(y) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(y-\mu)^2}{2\sigma^2}} \quad (16)$$

We'll start by showing that adding Gaussian noise with mean 0 and variance $O(\Delta^2 \log(1/\delta)/\epsilon^2)$ satisfies (ϵ, δ) -differential privacy. Notice that this standard deviation is actually *larger* than the variance

for the Laplace mechanism by a factor of $O(\log(1/\delta))$, but analyzing the univariate case will be an important warmup for analyzing the case of functions that output high-dimensional vectors.

For a function $f : \mathcal{X}^n \rightarrow \mathbb{R}$, we'll define the *Gaussian mechanism* as

$$A(\mathbf{x}) = f(\mathbf{x}) + \mathcal{N}\left(0, \frac{2\Delta^2 \log(2/\delta)}{\varepsilon^2}\right) \quad (17)$$

and we will prove that this mechanism satisfies approximate differential privacy.

Theorem 2.1. *For any $\varepsilon \leq 1$ and $\delta \in (0, 1)$, the Gaussian mechanism (17) satisfies (ε, δ) -differential privacy.²*

Proof. To start, let's fix two neighboring datasets \mathbf{x}, \mathbf{x}' . Without loss of generality, we will assume $f(\mathbf{x}) = 0$, which implies $|f(\mathbf{x}')| \leq \Delta$. Recall that to use Lemma 1.4 we want to show that $\mathbb{P}(A(\mathbf{x}) \in \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \leq \delta$, so to this end let's study the following probability ratio. To clean up the calculations, we will look at the log of the probability ratio, and we'll write $\sigma^2 = 2\Delta^2 \log(2/\delta)/\varepsilon^2$ to be the variance of the Gaussian noise. Then we have

$$\ln\left(\frac{\mathbb{P}(A(\mathbf{x}) = y)}{\mathbb{P}(A(\mathbf{x}') = y)}\right) = \ln\left(\frac{\exp(-\frac{y^2}{2\sigma^2})}{\exp(-\frac{(y-f(\mathbf{x}'))^2}{2\sigma^2})}\right) = \frac{(y-f(\mathbf{x}'))^2}{2\sigma^2} - \frac{y^2}{2\sigma^2} \quad (18)$$

$$= \frac{(y-f(\mathbf{x}'))^2 - y^2}{2\sigma^2} \quad (19)$$

$$= \frac{-2f(\mathbf{x}')y + f(\mathbf{x}')^2}{2\sigma^2} \quad (20)$$

$$= \frac{-2f(\mathbf{x}')y + f(\mathbf{x}')^2}{4\Delta^2 \log(2/\delta)/\varepsilon^2} \quad (21)$$

$$\leq \frac{-y\varepsilon^2}{2\Delta \log(2/\delta)} + \frac{\varepsilon}{4} \quad (22)$$

where the last two lines use the fact that $|f(\mathbf{x}')| \leq \Delta$, $\varepsilon \leq 1$, and $\log(2/\delta) \geq 1$ and does some simplification.

Notice that the ratio depends on y itself, and is largest when y is negative with large magnitude. Since y is random and drawn from $\mathcal{N}(0, \sigma^2)$, it shouldn't take extreme negative values too often. In particular, we have

$$\text{Bad}_{\mathbf{x}, \mathbf{x}'} \subseteq \left\{y : |y| > \frac{\sqrt{2}\Delta \log(2/\delta)}{\varepsilon}\right\} \quad (23)$$

Using bounds on the tails of Gaussians, namely that $\mathbb{P}(|\mathcal{N}(0, \sigma^2)| > t\sigma) \leq 2e^{-t^2/2}$, we have

$$\mathbb{P}(A(\mathbf{x}) \in \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \leq \mathbb{P}\left(A(\mathbf{x}) \in \left\{y : |y| > \frac{\sqrt{2}\Delta \log(2/\delta)}{\varepsilon}\right\}\right) \leq \delta \quad (24)$$

Thus we have the conditions to apply Lemma 1.4 to conclude that the Gaussian mechanism satisfies (ε, δ) -DP. \square

²Note that, unlike the Laplace mechanism, which works for any value of $\varepsilon > 0$, for our analysis of the Gaussian mechanism we need ε to be somewhat small, although we could still prove that the Gaussian mechanism is private, with slightly different types of bounds, when ε is larger.

2.2 Multivariate Gaussian Noise and ℓ_2 -Sensitivity

So far we've seen two different mechanisms that satisfy approximate DP but not pure DP, but no killer application where (ϵ, δ) -DP allows us to obtain asymptotically lower error for any problem. In this section we'll show that the *multivariate* version of the Gaussian mechanism can do just that for when we want to approximate some function $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ where f has low global sensitivity in the Euclidean norm (ℓ_2 -norm) rather than in the ℓ_1 -norm.

2.2.1 ℓ_2 -Sensitivity

Definition 2.2 (Global ℓ_2 Sensitivity). For a function $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$, we define the global ℓ_2 -sensitivity to be

$$\Delta_2 = \max_{\text{neighboring } \mathbf{x}, \mathbf{x}'} \|f(\mathbf{x}) - f(\mathbf{x}')\|_2 \quad (25)$$

where $\|v\|_2 = (\sum_i v_i^2)^{1/2}$ is the Euclidean norm. Note that we will sometimes denote the ℓ_1 -sensitivity as Δ_1 to distinguish it from Δ_2 .

One important thing to remember is that the ℓ_2 -norm is never more than the ℓ_1 norm. In fact, we have $\Delta_2 \leq \Delta_1 \leq \sqrt{k}\Delta_2$, which indicates that the ℓ_2 -sensitivity is never more than the ℓ_1 -sensitivity but can actually be much lower when k is large. For example, suppose $\mathcal{X} = \{0, 1\}^k$ and the statistic f just computes the sum of the datapoints $f(\mathbf{x}) = \sum_{i=1}^n x_i$. Then we have $\Delta_1 = k$ and $\Delta_2 = \sqrt{k}$.

Another useful example to keep in mind comes from the query release problem. If we have a set of k statistics f_1, \dots, f_k where each has the form

$$f_j(\mathbf{x}) = \sum_{i=1}^n \varphi(x_i) \text{ for some } \varphi : \mathcal{X} \rightarrow \{0, 1\} \quad (26)$$

and we write

$$f(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_k(\mathbf{x})) \quad (27)$$

then f has $\Delta_1 = k$ and $\Delta_2 = \sqrt{k}$.

As we'll see the useful thing about the Gaussian mechanism as opposed to the Laplace mechanism is that the magnitude of the noise can be proportional to Δ_2 instead of Δ_1 .

2.2.2 The Multivariate Gaussian Mechanism

The *spherical multivariate Gaussian distribution* in \mathbb{R}^k with mean $\vec{\mu}$ and variance σ^2 is denoted $\mathcal{N}(\vec{\mu}, \sigma^2 \mathbb{I}_{k \times k})$ is defined by the density function

$$p_{\vec{\mu}, \sigma^2}(\mathbf{y}) = \frac{1}{(2\pi\sigma^2)^{k/2}} e^{-\frac{\|\mathbf{y} - \vec{\mu}\|_2^2}{2\sigma^2}} \quad (28)$$

In more operational terms, it is the random variable $Z = (Z_1, \dots, Z_k)$ where each Z_j is sampled independently from the univariate Gaussian distribution $\mathcal{N}(\mu_j, \sigma^2)$. A few basic facts about this random variable will be useful to know:

Lemma 2.3. *If Z is drawn from the spherical multivariate Gaussian distribution $\mathcal{N}(\vec{0}, \sigma^2 \mathbb{I}_{k \times k})$, then*

1. $\mathbb{E}(\|Z\|_2^2) = \sigma^2 k$
2. $\mathbb{E}(\|Z\|_2) \leq \sigma \sqrt{k}$

3. $\mathbb{E}(\max\{|Z_1|, \dots, |Z_k|\}) \leq \sigma \cdot O(\sqrt{\ln k})$

4. For any vector $v \in \mathbb{R}^k$ the dot product $Z \cdot v$ is distributed as $\mathcal{N}(0, \sigma^2 \|v\|_2^2)$.

Now we're ready to define the Gaussian mechanism. Given a function $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ with ℓ_2 -sensitivity Δ_2 , we can define the Gaussian mechanism as

$$A(\mathbf{x}) = f(\mathbf{x}) + \mathcal{N}\left(\vec{0}, \frac{2\Delta_2^2 \log(2/\delta)}{\varepsilon^2} \cdot \mathbb{I}_{k \times k}\right) \quad (29)$$

Unsurprisingly, we will prove that this mechanism satisfies (ε, δ) -differential privacy

Theorem 2.4. For any $\varepsilon \leq 1$ and $\delta \in (0, 1)$, the Gaussian mechanism (29) satisfies (ε, δ) -differential privacy.

Proof. The proof follows the exact same approach as the univariate Gaussian mechanism, but we have to calculate the probability ratio for the multivariate Gaussian. As before, fix two neighboring datasets \mathbf{x}, \mathbf{x}' and assume without loss of generality that $f(\mathbf{x}) = \vec{0}$ so that $\|f(\mathbf{x}')\|_2 \leq \Delta_2$. We will study the log-probability-ratio again

$$\ln\left(\frac{\mathbb{P}(A(\mathbf{x}) = \vec{y})}{\mathbb{P}(A(\mathbf{x}') = \vec{y})}\right) = \ln\left(\frac{\exp(-\frac{\|\vec{y}\|_2^2}{2\sigma^2})}{\exp(-\frac{\|\vec{y}-f(\mathbf{x}')\|_2^2}{2\sigma^2})}\right) = \frac{\|\vec{y}-f(\mathbf{x}')\|_2^2 - \|\vec{y}\|_2^2}{2\sigma^2} \quad (30)$$

To analyze this expression precisely, we want to use the fact that for any pair of vectors $u, v \in \mathbb{R}^k$, $\|u - v\|_2^2 = \|u\|_2^2 + \|v\|_2^2 - 2u \cdot v$. Thus we have

$$\frac{\|\vec{y}-f(\mathbf{x}')\|_2^2 - \|\vec{y}\|_2^2}{2\sigma^2} = \frac{-2f(\mathbf{x}') \cdot \vec{y} + \|f(\mathbf{x}')\|_2^2}{2\sigma^2} \quad (31)$$

$$= \frac{-2f(\mathbf{x}') \cdot \vec{y} + \|f(\mathbf{x}')\|_2^2}{4\Delta_2^2 \log(2/\delta)/\varepsilon^2} \quad (32)$$

$$\leq \frac{-(f(\mathbf{x}') \cdot \vec{y}) \cdot \varepsilon^2}{2\Delta_2^2 \log(2/\delta)} + \frac{\varepsilon}{4} \quad (33)$$

As with the univariate case, we have

$$\text{Bad}_{\mathbf{x}, \mathbf{x}'} \subseteq \left\{ \vec{y} : |f(\mathbf{x}') \cdot \vec{y}| > \frac{\sqrt{2}\Delta_2^2 \log(2/\delta)}{\varepsilon} \right\} \quad (34)$$

Now, since $f(\mathbf{x}') \cdot \vec{y}$ is distributed as the univariate Gaussian $\mathcal{N}(0, \|f(\mathbf{x}')\|_2^2 \sigma^2)$ where $\|f(\mathbf{x}')\|_2^2 \leq \Delta_2^2$, and $\sigma^2 = 2\Delta_2^2 \log(2/\delta)/\varepsilon^2$, we have

$$\mathbb{P}(A(\mathbf{x}) \in \text{Bad}_{\mathbf{x}, \mathbf{x}'}) \leq \mathbb{P}\left(A(\mathbf{x}) \in \left\{ \vec{y} : |f(\mathbf{x}') \cdot \vec{y}| > \frac{\sqrt{2}\Delta_2^2 \log(2/\delta)}{\varepsilon} \right\}\right) \leq \delta \quad (35)$$

□

2.2.3 Discussion

As promised, the Gaussian mechanism allows us to get better variance for the linear query release problem when the number of queries is relatively large. Specifically, if we want to solve query release with k queries, then the Laplace mechanism requires us to add noise to each query proportional to the ℓ_1 -sensitivity $\Delta_1 \leq k$, which gives an expected maximum error of

$$\mathbb{E} \left(\max_{j=1}^k |f_j(\mathbf{x}) - a_j| \right) \leq O \left(\frac{k \log k}{\epsilon} \right) \quad (36)$$

Although we haven't seen the tools to prove it yet, one can actually show that *any* $(\epsilon, 0)$ -DP algorithm for answering k arbitrary queries of this form must have expected maximum error $\Omega(k/\epsilon)$, so we cannot improve the Laplace mechanism too much without relaxing the definition of privacy. However, since $\Delta_2 \leq \sqrt{k}$, if we are willing to accept a small $\delta > 0$, the Gaussian mechanism will guarantee

$$\mathbb{E} \left(\max_{j=1}^k |f_j(\mathbf{x}) - a_j| \right) \leq O \left(\frac{\sqrt{k \log k \log(1/\delta)}}{\epsilon} \right) \quad (37)$$

Thus, if k is large enough so that $k \log k \geq \log(1/\delta)$ then we get strictly less error by relaxing the definition to approximate DP and using the Gaussian mechanism. One thing we will see in the next lecture or two is that this improvement is really a consequence of relaxing the definition of DP, and not of using Gaussian noise *per se*, and we can actually add Laplace noise of similar magnitude if we're willing to settle for approximate DP. However, Gaussian noise is nice for lots of reasons

Additional Reading and Watching

- Our analysis of the Gaussian mechanism isn't quite tight, both because we wanted to avoid some nasty calculations involving Gaussian density functions, and because we lose something by going through Lemma 1.4 rather than analyzing (ϵ, δ) -DP directly. A tighter analysis of the Gaussian mechanism for (ϵ, δ) -DP was done by Balle and Wang [BW18].
- There are also alternative variants of differential privacy that, in some sense, are tailored to the precise privacy properties of the Gaussian mechanism. These variants are called *concentrated DP* [DR16, BS16] and *Gaussian DP* [DRS19].

Acknowledgement This lecture note is built on the course material developed by Adam Smith and Jonathan Ullman.

References

- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference, TCC '16*, 2016. <https://arxiv.org/abs/1605.02065>.
- [BW18] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning, ICML '18*. PMLR, 2018. <https://arxiv.org/abs/1805.06530>.
- [DR16] Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016. <https://arxiv.org/abs/1603.01887>.

[DRS19] Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019. <https://arxiv.org/abs/1905.02383>.