Differential Privacy

Beyond Global Sensitivity

Marco Gaboardi
Boston University

The opinions expressed in this course are mine and they do not not reflect those of the National Science Foundation or the US. Census Bureau.

Differential privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

Global Sensitivity

Definition 1.8 (Global sensitivity). The *global sensitivity* of a function $q : \mathcal{X}^n \to \mathbb{R}$ is:

$$\Delta q = \max \left\{ |q(D) - q(D')| \;\middle|\; D \sim_1 D' \in \mathcal{X}^n \right\}$$

Global Sensitivity

Definition 1.8 (Global sensitivity). The *global sensitivity* of a function $q : \mathcal{X}^n \to \mathbb{R}$ is:

$$\Delta q = \max\left\{ |q(D) - q(D')| \;\middle|\; D \sim_1 D' \in \mathcal{X}^n \right\}$$

Global Sensitivity

Definition 1.8 (Global sensitivity)**.** The *global sensitivity* of a function $q : \mathcal{X}^n \to \mathbb{R}$ is:

$$\Delta q = \max \left\{ |q(D) - q(D')| \;\middle|\; D \sim_1 D' \in \mathcal{X}^n \right\}$$

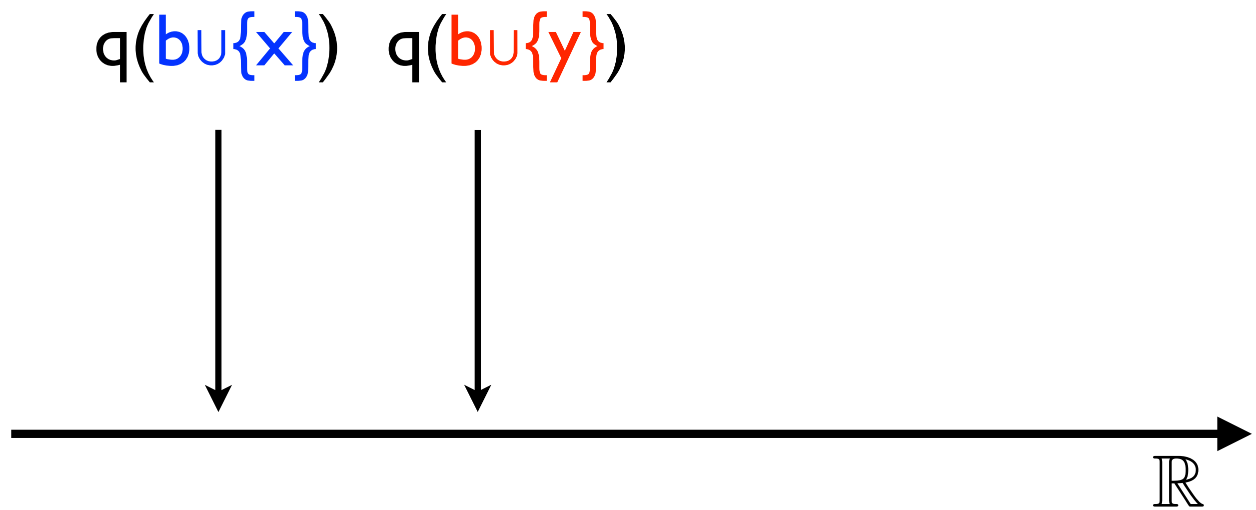

Global Sensitivity

Definition 1.8 (Global sensitivity)**.** The *global sensitivity* of a function $q : \mathcal{X}^n \to \mathbb{R}$ is:

$$\Delta q = \max \left\{ |q(D) - q(D')| \;\middle|\; D \sim_1 D' \in \mathcal{X}^n \right\}$$

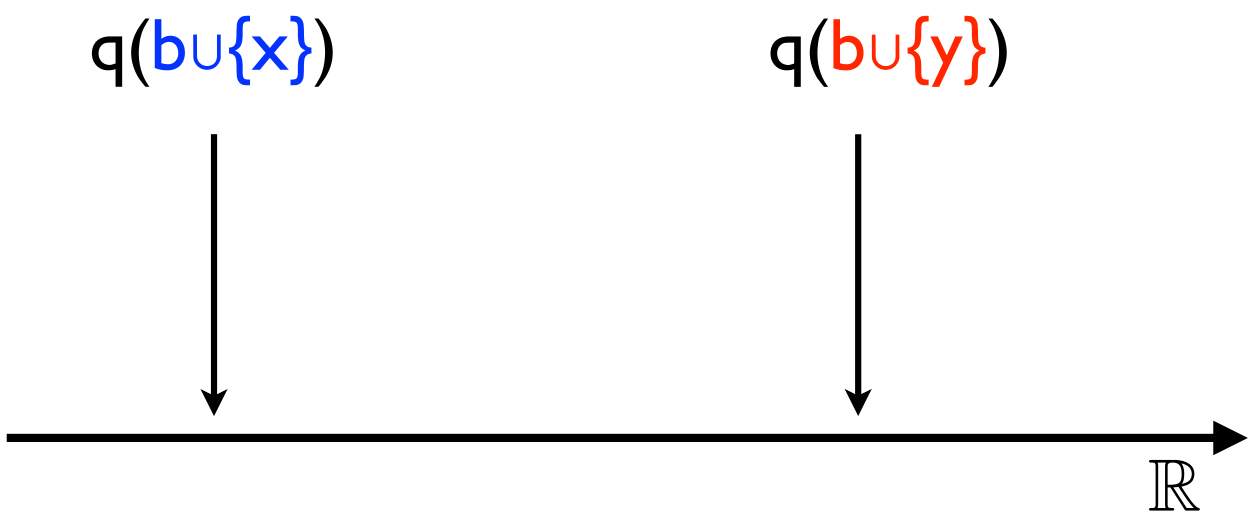

Laplace Mechanism

Algorithm 2 Pseudo-code for the Laplace Mechanism

1: **function** LAPMECH(D, q, ϵ)
2: $\quad Y \xleftarrow{\$} \mathsf{Lap}(\frac{\Delta q}{\epsilon})(0)$
3: $\quad$ **return** $q(D) + Y$
4: **end function**

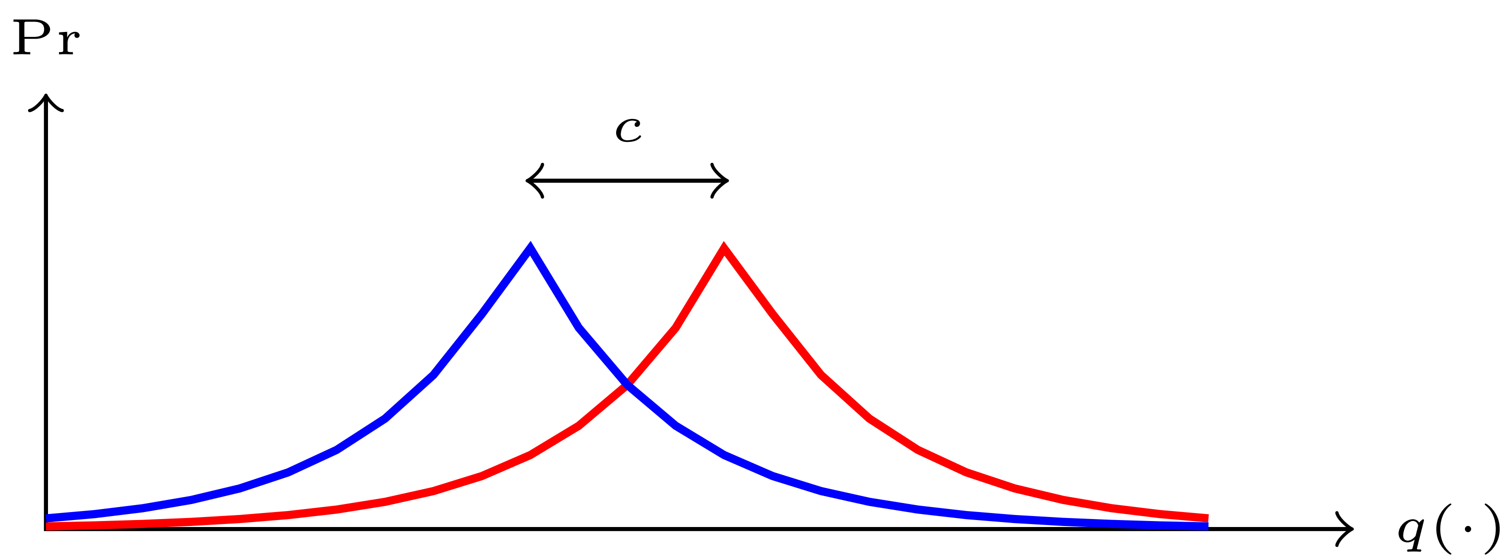

Laplace Mechanism

Accuracy Theorem: let $r = \mathsf{LapMech}(D, q, \epsilon)$

$$\Pr\left[|q(D) - r| \geq \left(\frac{\Delta q}{\epsilon}\right)\ln\left(\frac{1}{\beta}\right)\right] = \beta$$

Multidimensional Output

What can we do when we have a multidimensional output?

$$q : \mathscr{X}^n \to \mathbb{R}^m$$

Global L1 Sensitivity

We can generalize the notion of global sensitivity:

$$\Delta_1 q = \max \left\{ || q(D) - q(D') ||_1 \mid D \sim_1 D' \right\}$$

Where

$$|| \overrightarrow{v} ||_1 = \sum_{i=0} |v_i|$$

Global L1 Sensitivity

What is the L1 sensitivity of m counting query seen all together?

$$q : X^n \to \mathbb{R}^m \qquad q(D) = (q_1(D), \ldots, q_m(D))$$

Global L1 Sensitivity

What is the L1 sensitivity of m counting query seen all together?

$$q : X^n \to \mathbb{R}^m \qquad q(D) = (q_1(D), \ldots, q_m(D))$$

The L1 global sensitivity is

$$\frac{m}{n}$$

Laplace Mechanism

When $q : \mathcal{X}^n \rightarrow \mathbb{R}^m$

$$\texttt{LapMech}(D, q, \epsilon) = q(D) + (Y_1, \ldots, Y_m)$$

where $Y_i \sim_{i.i.d.} Lap(\frac{\Delta_1 q}{\epsilon}, 0)$

This mechanism is (eps,0)-DP

Accuracy revisited

Accuracy Theorem (for m counting queries together):

$$\Pr\left[||q(D) - r||_\infty \geq \left(\frac{n}{m\epsilon}\right)\ln\left(\frac{m}{\beta}\right)\right] \leq \beta$$

Where

$$||\overrightarrow{v}||_\infty = \max_{i=0} |v_i|$$

Global L2 Sensitivity

We can have another notion of global sensitivity:

$$\Delta_2 q = \max \left\{ ||\, q(D) - q(D') \,||_2 \;|\, D \sim_1 D' \right\}$$

Where

$$||\, \overrightarrow{v} \,||_2 = \sqrt{\sum_{i=0} v_i^2}$$

Gaussian Mechanism

Algorithm 14 Pseudo-code for the Gaussian Mechanism

1: **function** GaussMech(D, q, ϵ)
2: $\quad Y \xleftarrow{\$} \mathsf{Gauss}(0, \frac{2\ln(\frac{1.25}{\delta})(\Delta_2 q)^2}{\epsilon^2})$
3: $\quad$ **return** $q(D) + Y$
4: **end function**

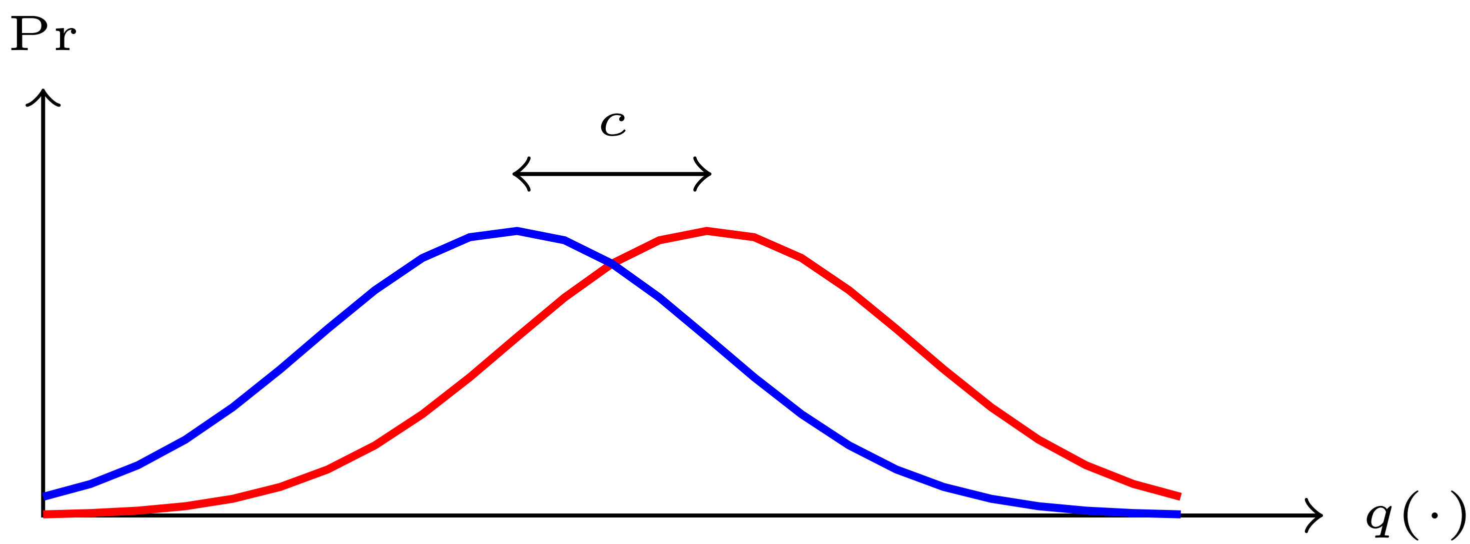

Gaussian Mechanism

Algorithm 14 Pseudo-code for the Gaussian Mechanism

function $\textsc{GaussMech}(D, q, \epsilon)$
    $Y \xleftarrow{\$} \mathsf{Gauss}(0, \frac{2\ln(\frac{1.25}{\delta})(\Delta_2 q)^2}{\epsilon^2})$
    return $q(D) + Y$
end function

Gaussian Mechanism

Algorithm 14 Pseudo-code for the Gaussian Mechanism

function GaussMech(D, q, ϵ)
    $Y \xleftarrow{\$} \mathsf{Gauss}(0, \frac{2\ln(\frac{1.25}{\delta})(\Delta_2 q)^2}{\epsilon^2})$
    return $q(D) + Y$
end function

Theorem (Privacy of the Gaussian Mechanism)

The Gaussian mechanism is (ε,δ)-differentially private.

Gaussian Mechanism

Theorem (Privacy of the Gaussian Mechanism)
The Gaussian mechanism is (ε,δ)-differentially private.

Proof: Intuitively

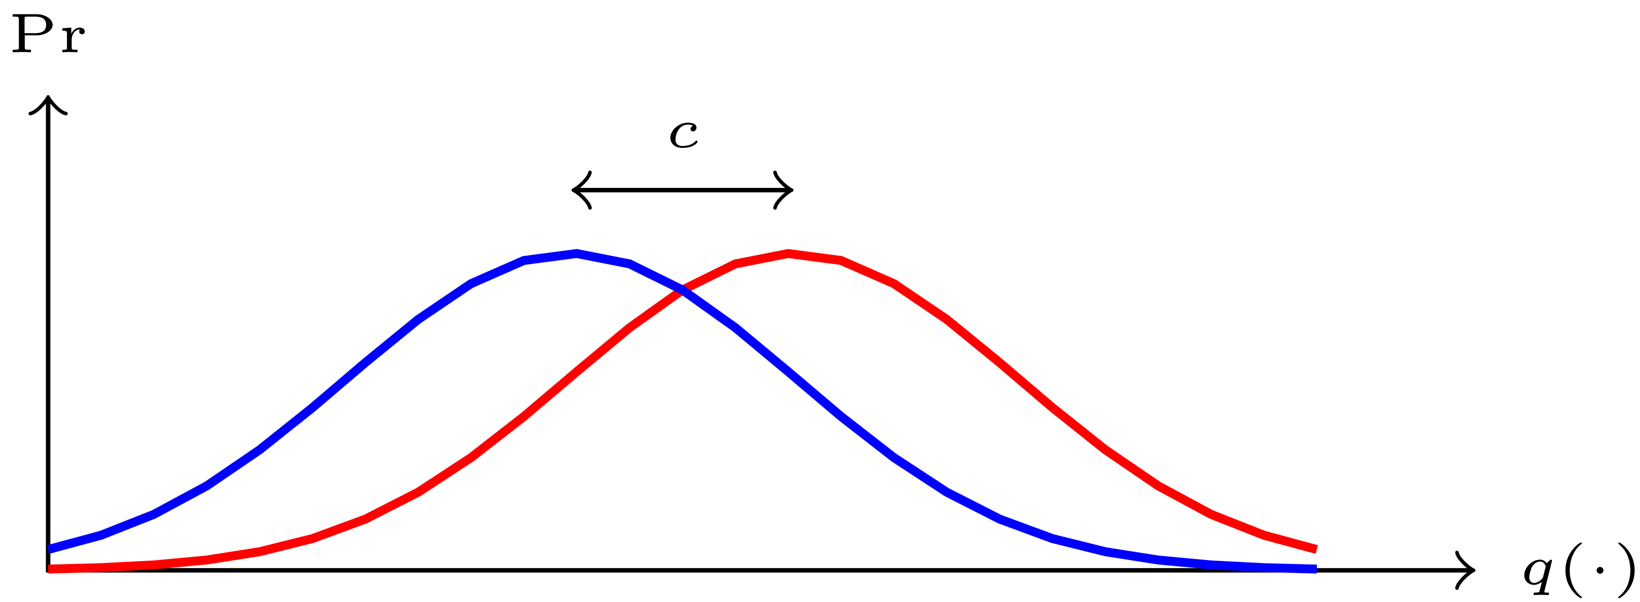

Gaussian Mechanism

Theorem (Privacy of the Gaussian Mechanism)
The Gaussian mechanism is (ε,δ)-differentially private.

Proof: Intuitively

We need δ to account for bigger differences in the tail

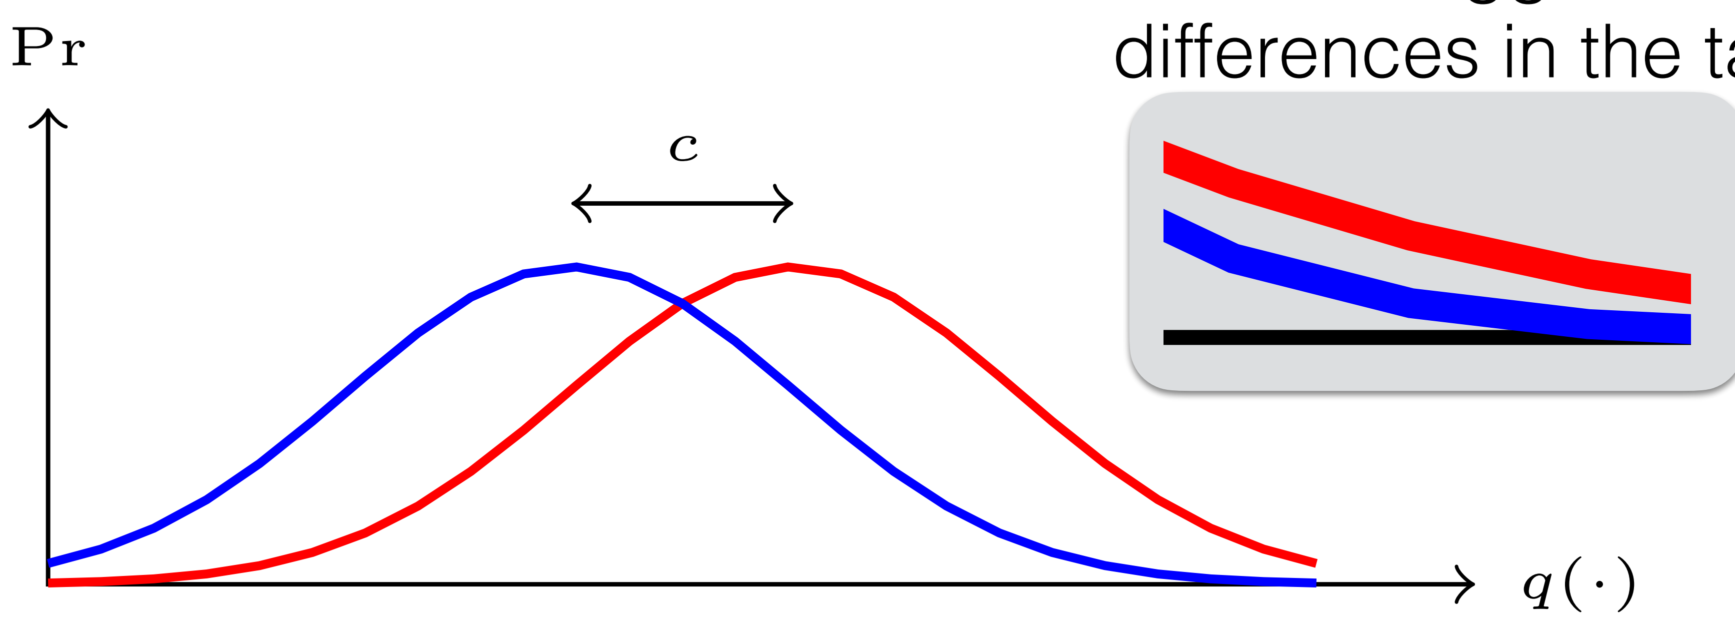

Gaussian Mechanism

Accuracy Theorem (for m counting queries together)

$$\Pr\left[||q(D) - r||_\infty \geq \frac{2\Delta_2 q}{\epsilon}\sqrt{\ln(\frac{1.25}{\delta})\ln\frac{m}{\beta}}\right] \leq \beta$$

Laplace vs Gaussian Mechanism

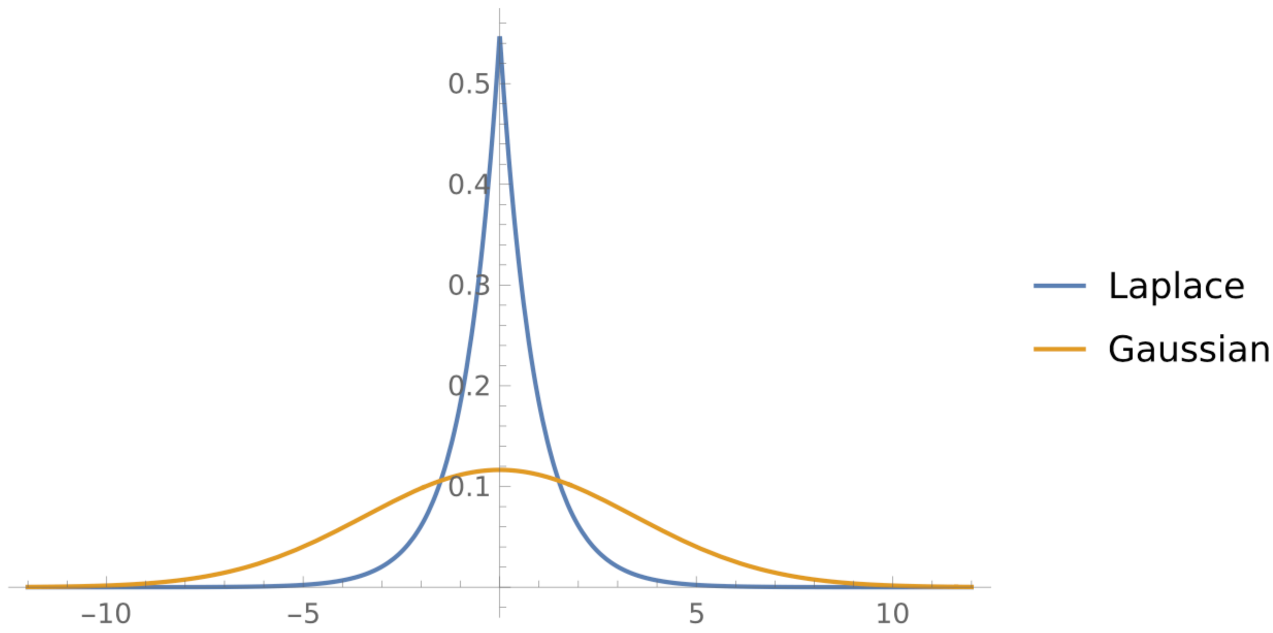

Global L2 Sensitivity

What is the L2 sensitivity of m counting query seen all together?

$$q : X^n \to \mathbb{R}^m \qquad q(D) = (q_1(D), \ldots, q_m(D))$$

Global L2 Sensitivity

What is the L2 sensitivity of m counting query seen all together?

$$q : X^n \to \mathbb{R}^m \qquad q(D) = (q_1(D), \ldots, q_m(D))$$

The L2 global sensitivity is

$$\frac{\sqrt{m}}{n}$$

Differential privacy

Definition

Given $\varepsilon, \delta \geq 0$, a probabilistic query $Q: X^n \to R$ is (ε, δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon) \Pr[Q(b_2) \in S] + \delta$$

Privacy Loss

In general we can think about the following quantity as the privacy loss incurred by observing r as output of $\mathcal{M}$ on the databases D and D'.

$$\mathcal{L}_{\mathcal{M}}^{D \to D'}(r) = \ln\left(\frac{\Pr[\mathcal{M}(D) = r]}{\Pr[\mathcal{M}(D') = r]}\right) = -\mathcal{L}_{\mathcal{M}}^{D' \to D}(r)$$

The $(\epsilon, 0)$-differential privacy requirement corresponds to requiring that for every r and every adjacent D, D' we have:

$$\left|\mathcal{L}_{\mathcal{M}}^{D \to D'}(r)\right| \leq \epsilon$$

(ε,δ)-Differential Privacy

This corresponds to a privacy loss of the form:

$$\mathcal{L}_{\mathcal{M}}^{D \to D'}(r) = \ln\left(\frac{\Pr[\mathcal{M}(D) = r|E]}{\Pr[\mathcal{M}(D') = r|E']}\right)$$

The (ϵ, δ)-differential privacy requirement corresponds to requiring that for every r and every adjacent D, D' we have:

$$\Pr\left[\left|\mathcal{L}_{\mathcal{M}}^{D \to D'}(r)\right| \leq \epsilon\right] \geq 1 - \delta$$

Composition for (ε,δ)-DP

Theorem 1.22 (Standard composition for (ϵ, δ)-differential privacy). Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ_i, δ_i)-differentially private algorithms (for $1 \leq i \leq k$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(\sum_{i=1}^k \epsilon_i, \sum_{i=1}^k \delta_i)$-differentially private.

Proof. Fix any pair of adjacent datasets $D \sim_1 D'$. Fix also an output $\vec{r} = (r_1, \ldots, r_k) \in R_1 \times \cdots \times R_k$. Since each $\mathcal{M}_i : \mathcal{X}^n \to R_i$ is $(epsilon_i, \delta_i)$-differentially private, we have events E_i and E'_i such that $\Pr[E_i] \geq 1 - \delta_i$ and $\Pr[E'_i] \geq 1 - \delta_i$. We can then consider $E = E_1 \wedge \cdots \wedge E_k$ and $E' = E'_1 \wedge \cdots \wedge E'_k$.

Composition for (ε,δ)-DP

Theorem 1.22 (Standard composition for (ϵ, δ)-differential privacy)**.** Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ_i, δ_i)-differentially private algorithms (for $1 \leq i \leq k$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(\sum_{i=1}^k \epsilon_i, \sum_{i=1}^k \delta_i)$-differentially private.

We have:

$$\begin{aligned}
\mathcal{L}_{\mathcal{M}}^{D \to D'}(\vec{r}) &= \ln\left(\frac{\Pr[\mathcal{M}(D) = r|E]}{\Pr[\mathcal{M}(D') = r|E']}\right) \\
&= \ln\left(\frac{\Pr[\mathcal{M}_1(D) = r_1|E_1] \cdots \Pr[\mathcal{M}_k(D) = r_k|E_k]}{\Pr[\mathcal{M}_1(D') = r_1|E_1'] \cdots \Pr[\mathcal{M}_k(D') = r_k|E_k']}\right) \\
&= \ln\left(\frac{\Pr[\mathcal{M}_1(D) = r_1|E_1]}{\Pr[\mathcal{M}_1(D') = r_1|E_1']}\right) + \cdots + \ln\left(\frac{\Pr[\mathcal{M}_k(D) = r_k|E_k]}{\Pr[\mathcal{M}_k(D') = r_k|E_k']}\right) \\
&= \mathcal{L}_{\mathcal{M}_1}^{D \to D'}(r_1) + \cdots + \mathcal{L}_{\mathcal{M}_k}^{D \to D'}(r_k) \leq \epsilon_1 + \cdots \epsilon_k = \sum_{i=1}^k \epsilon_i.
\end{aligned}$$

Composition for (ε,δ)-DP

Theorem 1.22 (Standard composition for (ϵ, δ)-differential privacy). Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ_i, δ_i)-differentially private algorithms (for $1 \le i \le k$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(\sum_{i=1}^k \epsilon_i, \sum_{i=1}^k \delta_i)$-differentially private.

We still need to reason about the probability of E and E'. We know that for each E_i, E'_i we have $\Pr[E_i] \ge 1 - \delta_i$ and $\Pr[E'_i] \ge 1 - \delta_i$. So, by union bound we have $\Pr[E] \ge 1 - \sum_{i=1}^k \delta_i$ and $\Pr[E'] \ge 1 - \sum_{i=1}^k \delta_i$, and so we can conclude.

Advanced Composition

Question: how much perturbation do we have if we want to answer n queries under (ε,δ)-DP?

Using advanced composition we have as a max error

$$O\Big(\frac{1}{\epsilon_{\text{global}}\sqrt{n}}\Big)$$

If we don't renormalize this is of the order of

$$O\Big(\frac{\sqrt{n}}{\epsilon_{\text{global}}}\Big)$$

comparable to the sample error.

[DworkRothblumVadhan10, SteinkeUllman16]

Advanced Composition

Theorem 1.23 (Advanced composition)**.** Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ, δ)-differentially private algorithms (for $1 \leq i \leq k$ and $k < 1/\epsilon$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \dots, \mathcal{M}_k(D))$ is $(O(\sqrt{2k \ln(1/\delta')})\epsilon, k\delta + \delta')$-differentially private for every $\delta' > 0$.

Advanced Composition

Theorem 1.23 (Advanced composition). Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ, δ)-differentially private algorithms (for $1 \leq i \leq k$ and $k < 1/\epsilon$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(O(\sqrt{2k \ln(1/\delta')})\epsilon, k\delta + \delta')$-differentially private for every $\delta' > 0$.

Intuition: some of the outputs have positive privacy loss (i.e. give evidence for dataset D) and some have negative privacy loss (i.e. give evidence for dataset D'). The cancellations gives a smaller overall privacy loss.

Advanced Composition

Theorem 1.23 (Advanced composition). Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ, δ)-differentially private algorithms (for $1 \leq i \leq k$ and $k < 1/\epsilon$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(O(\sqrt{2k\ln(1/\delta')})\epsilon, k\delta + \delta')$-differentially private for every $\delta' > 0$.

Intuition: some of the outputs have positive privacy loss (i.e. give evidence for dataset D) and some have negative privacy loss (i.e. give evidence for dataset D'). The cancellations gives a smaller overall privacy loss.

Strategy:
1-considering the expected value of the privacy loss,
2-bound the privacy loss of all the variables together
3-compute the probability

The roles of δ

We have seen three roles that δ plays in DP

1. to account for the probability of failure in a DP computation
2. in the advanced composition theorem to have a better bound on the growth of ε when composing n queries,
3. to allow an analysis of the Gaussian Mechanism.

The point 3 (and 2) were the original motivations for introducing (ε,δ)-differential privacy while the point 1 is somehow undesirable.

The roles of δ

We have seen three roles that δ plays in DP

1. to account for the probability of failure in a DP computation
2. in the advanced composition theorem to have a better bound on the growth of ε when composing n queries,
3. to allow an analysis of the Gaussian Mechanism.

The point 3 (and 2) were the original motivations for introducing (ε,δ)-differential privacy while the point 1 is somehow undesirable.

Can we give other privacy definitions that behave well with respect to 3 and 2 and do not require 1?

Advanced Composition

Theorem 1.23 (Advanced composition). Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ, δ)-differentially private algorithms (for $1 \leq i \leq k$ and $k < 1/\epsilon$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(O(\sqrt{2k \ln(1/\delta')})\epsilon, k\delta + \delta')$-differentially private for every $\delta' > 0$.

Intuition: some of the outputs have positive privacy loss (i.e. give evidence for dataset D) and some have negative privacy loss (i.e. give evidence for dataset D'). The cancellations gives a smaller overall privacy loss.

Advanced Composition

Theorem 1.23 (Advanced composition)**.** Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be (ϵ, δ)-differentially private algorithms (for $1 \leq i \leq k$ and $k < 1/\epsilon$). Then, their composition defined to be $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ is $(O(\sqrt{2k \ln(1/\delta')})\epsilon, k\delta + \delta')$-differentially private for every $\delta' > 0$.

Intuition: some of the outputs have positive privacy loss (i.e. give evidence for dataset D) and some have negative privacy loss (i.e. give evidence for dataset D'). The cancellations gives a smaller overall privacy loss.

Strategy:
1-considering the expected value of the privacy loss,
2-bound the privacy loss of all the variables together
3-compute the probability

(ε,δ)-Differential Privacy

This corresponds to a privacy loss of the form:

$$\mathcal{L}_{\mathcal{M}}^{D \to D'}(r) = \ln\left(\frac{\Pr[\mathcal{M}(D) = r|E]}{\Pr[\mathcal{M}(D') = r|E']}\right)$$

The (ϵ, δ)-differential privacy requirement corresponds to requiring that for every r and every adjacent D, D' we have:

$$\Pr\left[\left|\mathcal{L}_{\mathcal{M}}^{D \to D'}(r)\right| \leq \epsilon\right] \geq 1 - \delta$$

Not exactly!

Bounding the moments

A random variable can be described using its moments.

$$\mu_n = \mathbb{E}[X^n]$$

Here we consider central moments. For instance, the first central moment is the mean, the second is the variance, the third is the skewness, etc.

Can we bound the moments of the privacy loss?

Moment generating function

The probability distribution of a random variable X can be described by its moment generating function:

$$\mathtt{m}_X(\alpha) = \mathbb{E}[e^{\alpha X}]$$

This function can be used to compute, or give upper bounds on the moments of the random variable X.

$$m_X(\alpha) = 1 + \alpha\mu_1 + \frac{\alpha^2\mu_2}{2!} + \ldots + \frac{\alpha^n\mu_n}{n!} + \ldots$$