CSE660
Differential Privacy

October 25, 2017

Marco Gaboardi

Room: 338-B

gaboardi@buffalo.edu

http://www.buffalo.edu/~gaboardi

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 $S \subseteq R$:

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

Blatantly non-privacy

The privacy mechanism $M:X^n \rightarrow R$ is blatantly non-private if an adversary can build a candidate database $D' \in X^n$, that agrees with the real database D in all but $o(n)$ entries:

$$d_H(D,D') \in o(n)$$

M is blatantly non private if we can reconstruct:

$n-f(n)$ entries for f sublinear.

This corresponds to require that the mechanism reconstructs a fraction $1-f(n)/n$ of the rows, and when $n \rightarrow \infty$ we have $1-f(n)/n \rightarrow 1$.

Differential privacy prevents blatantly non-privacy

Consider a uniformly random dataset $D \in X^n$.
Suppose $Q: X^n \rightarrow R$ is (ε,δ)-differentially private.
Then the the expected fraction of rows that any adversary can reconstruct is at most:

$$\frac{e^{\epsilon}}{|X|} + \delta$$

Let's try now to reason about the privacy parameters:

$\varepsilon=.01$ $\delta=10^{-15}$ $n=1000$ $|X|=1000$

$\varepsilon=.01$ $\delta=10^{-1}$ $n=1000$ $|X|=1000$

$\varepsilon=.1$ $\delta=10^{-15}$ $n=10000$ $|X|=100$

Multiple queries

Question: how much perturbation do we have if we want to answer n counting queries with Laplace under ε-DP?

Multiple queries

Question: how much perturbation do we have if we want to answer n counting queries with Laplace under ε-DP?

Using standard composition we have as a max error

$$O\left(\frac{n}{\epsilon_{\text{global}} n}\right) = O\left(\frac{1}{\epsilon_{\text{global}}}\right)$$

Notice that if we don't renormalize this is of the order of

$$O\left(\frac{n}{\epsilon_{\text{global}}}\right)$$

bigger than the sample error.

Composition

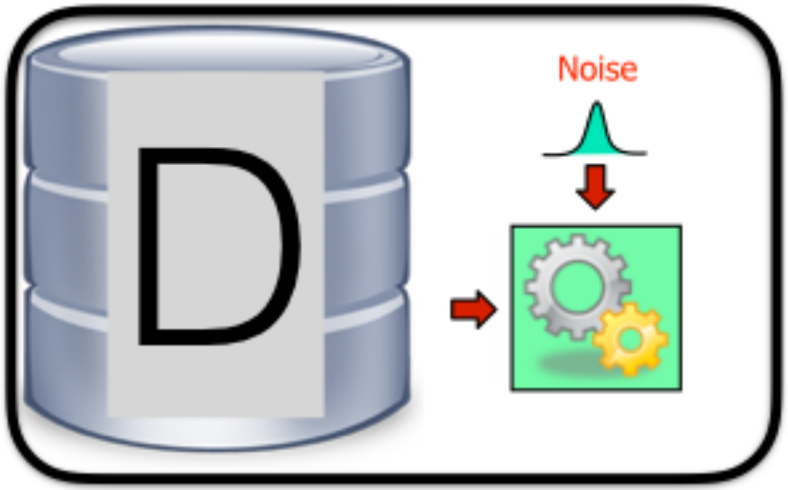

M_1 is ε_1-DP

M_2 is ε_2-DP

…

M_n is ε_n-DP

The overall process is $(\varepsilon_1+\varepsilon_2+\ldots+\varepsilon_n)$-DP

Composition

Theorem 1.18 (Standard composition for ϵ-differential privacy)**.** Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be ϵ_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 $\sum_{i=1}^{k} \epsilon_i$-differentially private.

Proof. Fix any pair of adjacent datasets $D \sim_1 D'$. Fix also an output $\vec{r} = (r_1, \ldots, r_k) \in R_1 \times \cdots \times R_k$. We have:

$$\frac{\Pr[\mathcal{M}(D) = \vec{r}]}{\Pr[\mathcal{M}(D') = \vec{r}]} = \frac{(\Pr[\mathcal{M}_1(D), \ldots, \mathcal{M}_k(D)) = (r_1, \ldots, r_k)]}{(\Pr[\mathcal{M}_1(D'), \ldots, \mathcal{M}_k(D')) = (r_1, \ldots, r_k)]}$$

$$= \frac{\Pr[\mathcal{M}_1(D) = r_1] \cdots \Pr[\mathcal{M}_k(D) = r_k]}{\Pr[\mathcal{M}_1(D') = r_1] \cdots \Pr[\mathcal{M}_k(D') = r_k]}$$

$$= \left(\frac{\Pr[\mathcal{M}_1(D) = r_1]}{\Pr[\mathcal{M}_1(D') = r_1]}\right) \cdots \left(\frac{\Pr[\mathcal{M}_k(D) = r_k]}{\Pr[\mathcal{M}_k(D') = r_k]}\right)$$

$$\leq \exp(\epsilon_1) \cdots \exp(\epsilon_k) = \exp\left(\sum_{i=1}^{k} \epsilon_i\right).$$

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:

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

Composition

Theorem 1.18 (Standard composition for ϵ-differential privacy)**.** Let $\mathcal{M}_i : \mathcal{X}^n \to R_i$ be ϵ_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 $\sum_{i=1}^{k} \epsilon_i$-differentially private.

Proof. Fix any pair of adjacent datasets $D \sim_1 D'$. Fix also an output $\vec{r} = (r_1, \ldots, r_k) \in R_1 \times \cdots \times R_k$. We have:

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

(ε,δ)-Differential Privacy

We can also reformulate (ε,δ)-differential privacy in terms of the privacy loss. Informally, we would like that to be equivalent to:

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

(ε,δ)-Differential Privacy

Lemma 1.20. A mechanism $\mathcal{M} : \mathcal{X}^n \to \mathcal{R}$ is (ϵ, δ)-differentially private iff for every adjacent D, D' there exist events E (depending on $\mathcal{M}(D)$) and E' (depending on $\mathcal{M}(D')$) such that $\Pr[E] \geq 1 - \delta$, $\Pr[E'] \geq 1 - \delta$ and such that :

$$\Pr[\mathcal{M}(D) \in T|E] \leq e^{\epsilon} \Pr[\mathcal{M}(D') \in T|E']$$

and

$$\Pr[\mathcal{M}(D') \in T|E'] \leq e^{\epsilon} \Pr[\mathcal{M}(D) \in T|E]$$

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

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

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

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

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 $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D), \ldots, \mathcal{M}_k(D))$ 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

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 $\delta' > 0$.

Lemma 1.24. If M_i is ϵ-differentially private, where $\epsilon \leq 1$, then:

$$\mathbb{E}_{r_i \leftarrow \mathcal{M}_i(D)}[\mathcal{L}_{\mathcal{M}_i}^{D \to D'}(r_i)] \leq 2\epsilon^2$$

By linearity of expectation we have

$$\mathbb{E}_{r \leftarrow \mathcal{M}(D)}[\mathcal{L}_{\mathcal{M}}^{D \to D'}(r)] \leq kO(\epsilon^2)$$

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 $\delta' > 0$.

Applying the Chernoff Bound for random variables whose absolute value is bounded by ϵ, we get that with probability at least $1 - \delta'$ over $\vec{r} \leftarrow \mathcal{M}(D)$, we have

$$\mathcal{L}_{\mathcal{M}}^{D \to D'}(\vec{r}) \leq kO(\epsilon^2) + O(\sqrt{k\ln(1/\delta')})\epsilon \leq O(\sqrt{k\ln(1/\delta')})\epsilon$$

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 $\delta' > 0$.

So, let ϵ' be $O(\sqrt{k\ln(1/\delta')})\epsilon$. For every $T \subseteq R$ we have

$$
\begin{aligned}
Pr[\mathcal{M}(D) \in T] &\leq \Pr_{\vec{r} \leftarrow \mathcal{M}(D)}[\mathcal{L}_{\mathcal{M}}^{D \to D'}(\vec{r}) > \epsilon'] + \sum_{\vec{r} \in T : \mathcal{L}_{\mathcal{M}}^{D \to D'}(\vec{r}) \leq \epsilon'} \Pr[\mathcal{M}(D) = \vec{r}] \\
&\leq \delta' + \sum_{\vec{r} \in T : \mathcal{L}_{\mathcal{M}}^{D \to D'}(\vec{r}) \leq \epsilon'} e^{\epsilon'} \Pr[\mathcal{M}(D') = \vec{r}] \\
&\leq \delta' + e^{\epsilon'} \Pr[\mathcal{M}(D') \in T]
\end{aligned}
$$

Composition

We always need to think before applying composition to whether we have other options!

Answering multiple queries

We have seen several methods to answer a single query:

- Randomized Response
- Laplace Mechanism
- Exponential Mechanism

And methods to answer multiple queries with small error:

- Standard composition - we can answer $\sqrt{n}$ queries.
- Advanced composition - we can answer n queries.