CSE660
Differential Privacy

September 27, 2017

Marco Gaboardi
Room: 338-B
gaboardi@buffalo.edu
http://www.buffalo.edu/~gaboardi

Outline of the class

Week 1

Introduction, motivation and privacy limitations. Definition of Differential Privacy and the curator model.

Week 2

Basic mechanisms: Randomized Response, Laplace Mechanism,

Week 3

Basic properties following from the definition, Exponential Mechanism and comparison with the other basic mechanisms.

Week 4

The Report Noisy max algorithm. The Sparse Vector technique.

Week 5

Formalizing privacy proofs using an approximate probabilistic coupling argument.

Releasing Many Counting Queries with Correlated Noise. The smallDB algorithm.

Week 6

The MWEM algorithm. The DualQuery algorithm.

Outline of the class

Week 8

Studying the experimental accuracy. Adaptivity and adaptive MWEM.

Week 7

Variations on differential privacy: Renyi DP, zero-concentrated DP

Week 9

PAC learning and private PAC learning

Week 10

The local model for differential privacy.

Week 11

More algorithms for the local model.

Week 12

Differentially Private Hypothesis Testing

Week 13

Differential Privacy and Generalization in Adaptive Data Analysis

Week 14

Project presentations

(ε,δ)-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 $S \subseteq R$:

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

Composition

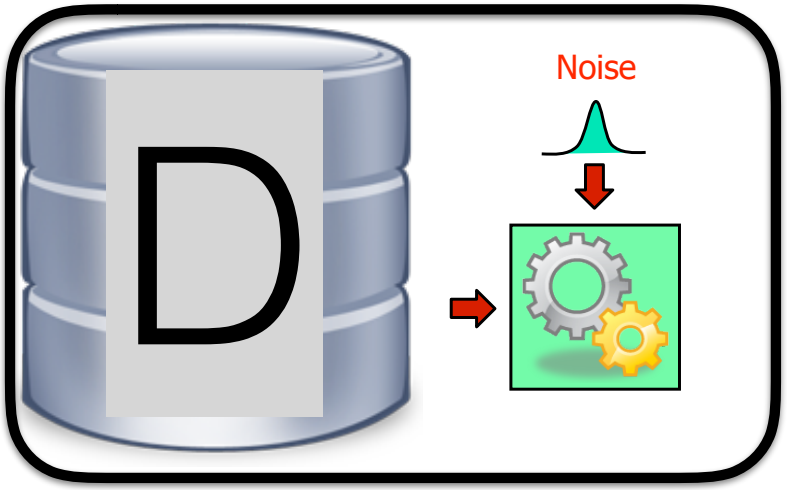

M_1 is ε_1-DP

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

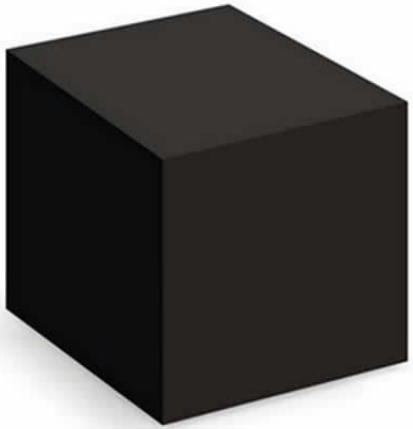

We always need to think before applying composition to whether we have other options!

Multiple queries

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

Reconstruction attack with polynomial adversary

Let $M:\{0,1\}^n \rightarrow R$ be a privacy mechanism adding noise within **$E=o(\sqrt{n})$** perturbation. Then we can show M blatantly non-private against an adversary A running in polynomial time and **answering n queries.**

[DinurNissim'02, DworkYekhanin'08]

Multiple queries

Question: how much perturbation do we have if we want to answer n counting queries under $\varepsilon_{\text{global}}$-DP?

We can split the privacy budget uniformly:

$$\epsilon = \frac{\epsilon_{\text{global}}}{n}$$

Laplace accuracy: with high probability we have:

$$\left|q(D) - r\right| \leq O\left(\frac{1}{\epsilon n}\right)$$

Multiple queries

Question: how much perturbation do we have if we want to answer n counting queries under $\varepsilon_{\text{global}}$-DP?

By putting them together (hiding some details) 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.

Multiple queries

Question: how many counting queries can we answer with small error under $\varepsilon_{\text{global}}$-DP?

Let's now target an error similar to sample error. How many queries we can answer?
If we want a non-normalized error of:

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

we can answer at most $\sqrt{n}$ queries.

Multiple queries

Question: Can we do better?

Multiple queries

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

Advanced Composition

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

We have (by hiding many details) 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]

Mechanisms

Question: How about non-numeric data?

Exponential Mechanism

The Exponential Mechanism generalize this approach.

Suppose that we have a scoring function u(D,o) that to each pair (database, potential output) assign a score (a negative real number).

We want to output approximately the element with the max score.

Exponential Mechanism

Exponential Mechanism:

$\mathcal{M}_E(x, u, \mathcal{R})$

return $r \in \mathcal{R}$ with prob. $\dfrac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r' \in \mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}$

where

$$\Delta u = \max_{r \in \mathcal{R}} \max_{x \sim_1 y} \left| u(x, r) - u(y, r) \right|$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

The proof is very similar to the one for the Laplace Mechanism.

$$\frac{\Pr[\mathcal{M}_E(x,u,\mathcal{R})=r]}{\Pr[\mathcal{M}_E(y,u,\mathcal{R})=r]} = \frac{\left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)}{\left(\frac{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}\right)}$$

$$= \left(\frac{\exp(\frac{\varepsilon u(x,r)}{2\Delta u})}{\exp(\frac{\varepsilon u(y,r)}{2\Delta u})}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

$$= \exp\left(\frac{\varepsilon(u(x,r')-u(y,r'))}{2\Delta u}\right)\cdot\left(\frac{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}}\exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Exponential Mechanism

Privacy theorem:
The Exponential Mechanism is differentially private.

Continuing

$$= \exp\left(\frac{\varepsilon(u(x,r') - u(y,r'))}{2\Delta u}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(y,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

$$\leq \exp\left(\frac{\varepsilon}{2}\right) \cdot \exp\left(\frac{\varepsilon}{2}\right) \cdot \left(\frac{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}{\sum_{r'\in\mathcal{R}} \exp(\frac{\varepsilon u(x,r')}{2\Delta u})}\right)$$

Here we change y with x by paying exp(ε/2).

Exponential Mechanism

Exponential Mechanism Accuracy theorem:
Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

It follows from this lemma

$$\Pr\left[u(\mathcal{M}_E(x, u, \mathcal{R})) \leq \mathrm{OPT}_u(x) - \frac{2\Delta u}{\varepsilon}\left(\ln\left(\frac{|\mathcal{R}|}{|\mathcal{R}_{\mathrm{OPT}}|}\right) + t\right)\right] \leq e^{-t}$$

Proof.

$$\Pr[u(\mathcal{M}_E(x, u, \mathcal{R})) \leq c] \leq \frac{|\mathcal{R}| \exp(\varepsilon c / 2\Delta u)}{|\mathcal{R}_{\mathrm{OPT}}| \exp(\varepsilon \mathrm{OPT}_u(x) / 2\Delta u)}$$

$$= \frac{|\mathcal{R}|}{|\mathcal{R}_{\mathrm{OPT}}|} \exp\left(\frac{\varepsilon(c - \mathrm{OPT}_u(x))}{2\Delta u}\right).$$

Exponential Mech

Here we have a dependency on the size of the output space

Exponential Mechanism Accuracy the

Let $\mathrm{OPT}_u(x) = \max_{r \in \mathcal{R}} u(x, r)$. Then

$$\Pr\left[OPT_u(x) - u(x, \mathcal{M}_E(x, u, \mathcal{R})) \geq \left(\frac{2\Delta u}{\epsilon}\right) \ln\left(\frac{|\mathcal{R}|}{\beta}\right)\right] \leq \beta$$

Let's compare it with the accuracy of the Laplace Mechanism.

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

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

Exponential Mechanism

The Exponential Mechanism is a very general mechanism. It can actually be used as a kind of universal mechanism.

Unfortunately, when the output space is big it can be very costly to sample from it - the best option is to enumerate all the possibilities.

Moreover, when the output space is big also the accuracy get worse.