Foundations of Privacy

Lecture 3

Differential Privacy

- We are now ready to define **Differential Privacy** for a randomized mechanism $\mathcal{K}$.

- Let us first consider the discrete case. Namely, $\mathcal{K}(x)$ is discrete, for every database x.

- **Definition (Differential Privacy)** $\mathcal{K}$ is ε-differentially private if for every pair of databases $x_1, x_2 \in \mathcal{X}$ such that $x_1 \sim x_2$, and for every $z \in Z$, we have:

  $$p(Z = z | X = x_1) \leq e^{\varepsilon} p(Z = z | X = x_2)$$

  where $p(Z = z | X = x)$ represents the conditional probability of z given x, namely the probability that on the database x the mechanism reports the answer z

- This definition therefore means that the value (or the presence) of an individual does not affect significantly the probability of getting a certain reported value.

Differential Privacy

- Let us now consider the continuous case. Namely, $\mathcal{K}(x)$ is a probability density function on $\mathcal{Z}$. The only thing that changes is that we consider a measurable subset $\mathcal{S}$ of $\mathcal{Z}$ instead than a single z:

- **Definition (Differential Privacy)** $\mathcal{K}$ is ε-differentially private if for every pair of databases $x_1, x_2 \in \mathcal{X}$ such that $x_1 \sim x_2$, and for every measurable $\mathcal{S} \subseteq Z$, we have:

$$p(Z \in \mathcal{S} | X = x_1) \leq e^{\varepsilon} p(Z \in \mathcal{S} | X = x_2)$$

  where $p(Z \in \mathcal{S} | X = x)$ represents the probability that on the database x the mechanism reports an answer in $\mathcal{S}$

- This definition therefore means that the value (or the presence) of an individual does not affect significantly the probability that the reported value satisfy a certain property.

Oblivious Mechanisms

- Given $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $\mathcal{K} : \mathcal{X} \rightarrow \mathcal{Z}$, we say that $\mathcal{K}$ is oblivious if it depends only on $\mathcal{Y}$ (not on $\mathcal{X}$)

- If $\mathcal{K}$ is oblivious, it can be seen as the composition of f and a randomized mechanism $\mathcal{H}$ (noise) defined on the exact answers $\mathcal{K} = f$ x $\mathcal{H}$

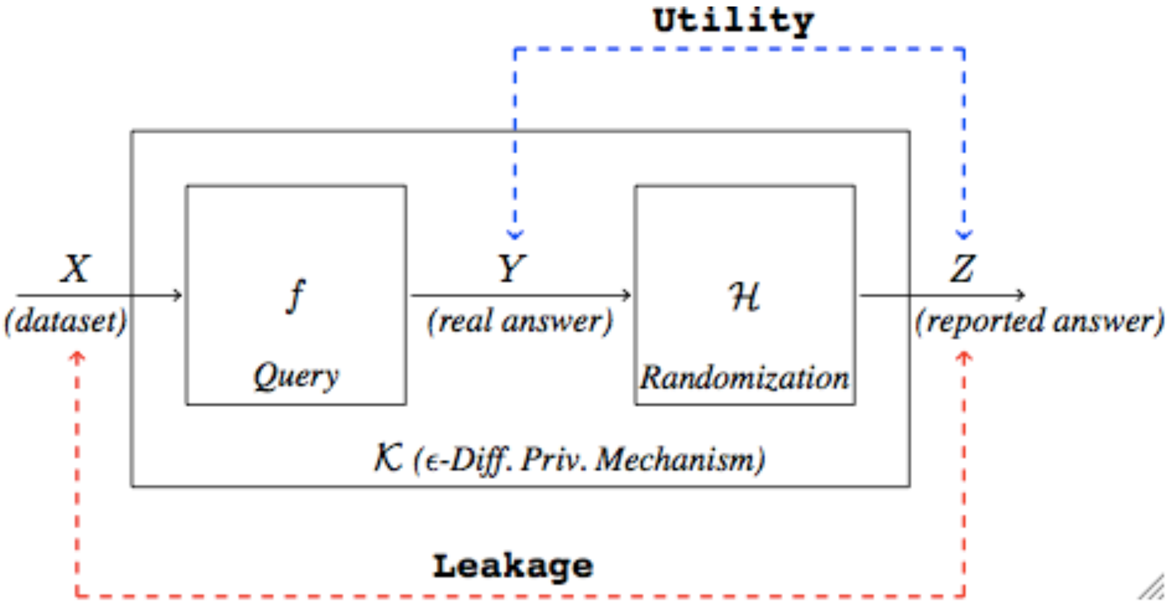

- Privacy concerns the information flow between the databases and the reported answers, while utility concerns the information flow between the correct answer and the reported answer

A typical oblivious differentially private mechanism: Laplacian noise

- Randomized mechanism for a query $f : \mathcal{X} \rightarrow \mathcal{Y}$.
- A typical randomized method: **add Laplacian noise.** If the exact answer is y, the reported answer is z, with a probability density function defined as:

$$dP_y(z) = c\, e^{-\frac{|z-y|}{\Delta f}\varepsilon}$$

where Δf is the *sensitivity* of f:

$$\Delta f = \max_{x \sim x' \in \mathcal{X}} |f(x) - f(x')|$$

($x \sim x'$ means x and x' are adjacent, i.e., they differ only for one record)

and c is a normalization factor:

$$c = \frac{\varepsilon}{2\,\Delta f}$$

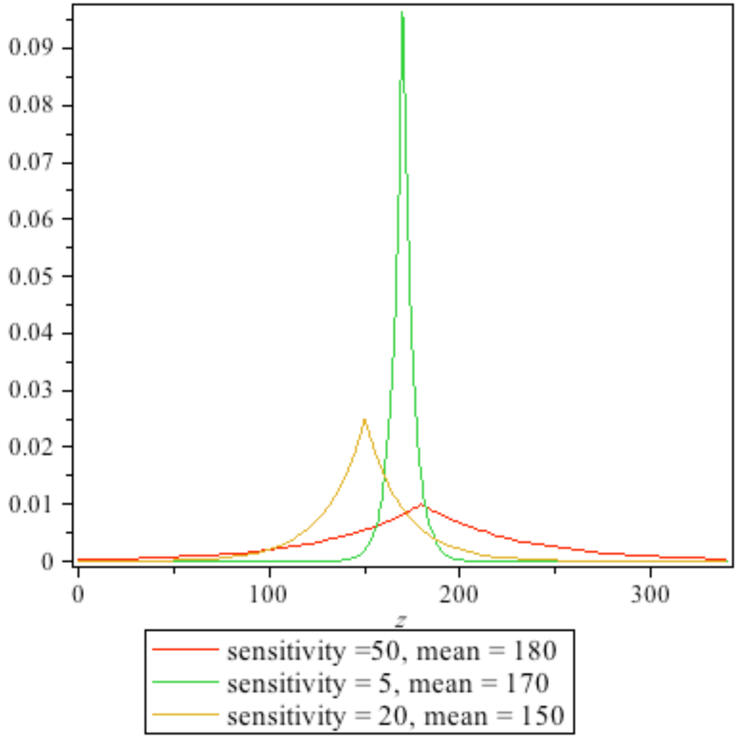

The geometric mechanism

- The Laplacian noise is typically used in the case that $\mathcal{Y}$ (the set of true answers of the query) is a **dense** numerical set, like the Reals or the Rationals.

- If $\mathcal{Y}$ is a **discrete** numerical set, like the Integers, then the typical mechanism used in this case is the **geometric mechanism**, which is a sort of discrete Laplacian.

- In the geometric mechanism, the probability distribution of the noise is:

$$p(z|y) = c\, e^{-\frac{|z-y|}{\Delta f}\varepsilon}$$

  - In this expression, c is a normalization factor, defined so to obtain a probability distribution,

  - Δf is the sensitivity of query f

Normalization constant in a geometric mechanism

- In the geometric mechanism, the probability distribution of the noise is:

$$p(z|y) = c\, e^{-\frac{|z-y|}{\Delta f}\varepsilon}$$

As usual, we can compute c (the normalization factor) by imposing that the sum of the probability on all Z is 1. It turns out that

$$c = \frac{1-\alpha}{1+\alpha} \qquad \text{where} \qquad \alpha = e^{-\frac{\varepsilon}{\Delta f}}$$

$$\text{hence} \qquad p(z|y) = \frac{1-\alpha}{1+\alpha}\, \alpha^{|z-y|}$$

- **Examples:** Compute the geometric mechanism for the following queries:
  - “ How many diabetic people weight more than 100 kilos ? ”
  - “ What is the max weight (in kilos) of a diabetic person ? ”

Gaussian noise

The formula for gaussian noise is

$$c\ e^{-\frac{(y-z)^2}{\sigma}\varepsilon}$$

where c is a normalization factor and σ is a suitable constant.

Question: does an oblivious mechanism based on this noise function satisfy ε-differential privacy, for some suitable value of σ ?

Truncated geometric mechanism

- Often $\mathcal{Y}$ (the set of the true answers) does not coincide with the whole integers, but it is just subset, for instance an interval [a,b].
- With the geometric (resp. Laplacian) mechanism, however, the set of reported answers $\mathcal{Z}$ is always the whole set of integers $\mathbb{Z}$ (resp. $\mathbb{R}$)
- It is often inconvenient to report answers outside $\mathcal{Y}$. We can make the set of reported answers coincide with that of the true answers by **truncating** the mechanism. In particular, if $\mathcal{Y}$ is the interval [a,b], we can set $\mathcal{Z}$ = [a,b], and transfer on the extremes a and b all the probability that (according to the original mechanism) would fall outside the interval: The probability that would fall to the left of a is transferred into a, and probability that would fall to the right of b is transferred into b.
- Exercise: Compute the truncated geometric mechanism for a counting query if the set of true answers are the integers in the interval [0,100]

Post-processing

- Post-processing a mechanism $\mathcal{K}$ consists in composing $\mathcal{K}$ with another function $\mathcal{F}$
  - $\mathcal{F}$ can be probabilistic or deterministic
  - $\mathcal{K}$ can be oblivious or not — it does not matter for the theorem below

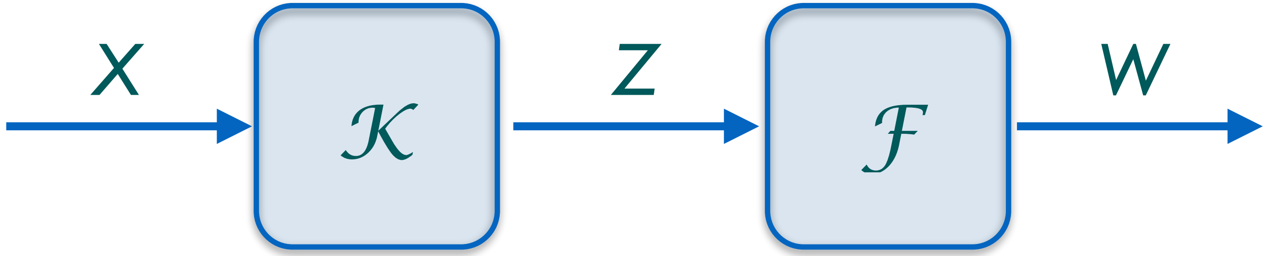

Theorem: Post processing does not harm privacy. Namely, if $\mathcal{K}$ is ε-differentially private, then also $\mathcal{F} \circ \mathcal{K}$ is ε-differentially private

Truncation

- Truncation is a typical example of post-processing
- In fact, assume that the true answer is in the interval [a,b]. Then truncation can be defined as follows: If the reported is smaller than a, then it gets remapped into a, and if it is greater than b, then it gets remapped into b.
- Because of the above theorem, truncation does not decrease the level of privacy.

Utility

- When a user sees the reported value z of the mechanism, he may take z as it is, or, based on his prior knowledge, he may guess another value w. We say that the user remaps z into w.
  Summarizing, we have:
- $\mathcal{X}$, the set of databases, with associated random variable X
- $\mathcal{Y}$, the set of true answers to the query f. Associated random variable Y
- $\mathcal{Z}$, the set of reported answers to the query f (after we apply the noise). Associated random variable Z
- $\mathcal{W}$, the set of guesses. Associated random variable W. $\mathcal{W}$ often coincides with $\mathcal{Y}$, but W usually does not coincide with Y.

Utility

- When a user sees the reported value z of the mechanism, he may take z as it is, or, based on his prior knowledge, he may guess another value w. We say that the user remaps z into w.
  Summarizing, we have:

- $\mathcal{X}$, the set of databases, with associated random variable X

- $\mathcal{Y}$, the set of true answers to the query f. Associated random variable Y

- $\mathcal{Z}$, the set of reported answers to the query f (after we apply the noise). Associated random variable Z

- $\mathcal{W}$, the set of guesses. Associated random variable W. $\mathcal{W}$ often coincides with $\mathcal{Y}$, but W usually does not coincide with Y.

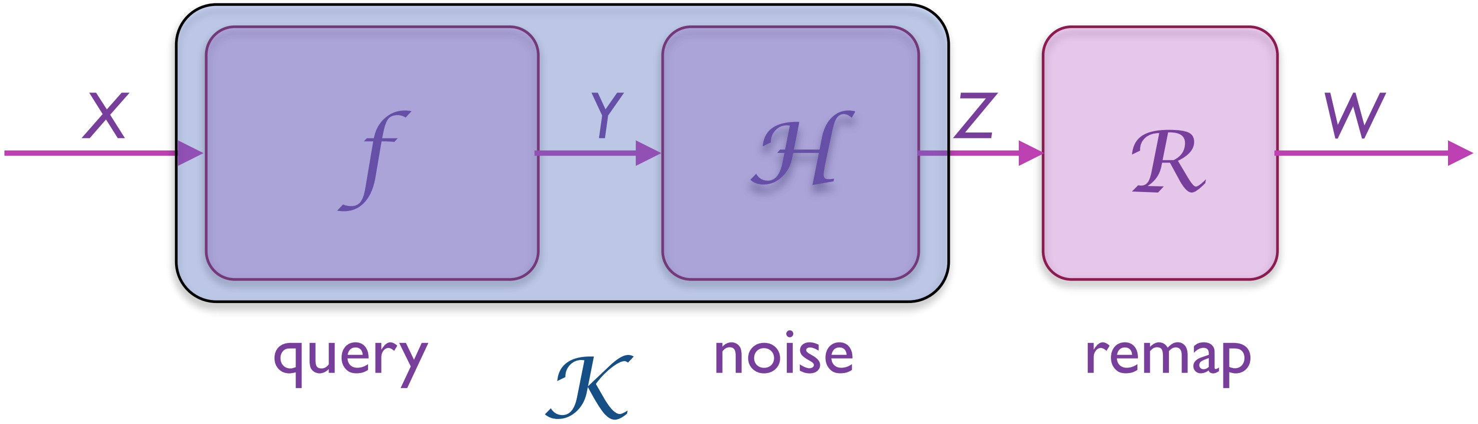

Schema for an oblivious mechanism. In a non-oblivious one Z depend also on X.

Utility

- A gain function is a function

$$g : \mathcal{W} \times \mathcal{Y} \to \mathbb{R}$$

  that represents the usefulness of the guess w when the true answer is y.

- Often there is a notion of distance d between w and y, representing how well w approximates y. Formally:

$$d : \mathcal{W} \times \mathcal{Y} \to \mathbb{R}$$

- The gain g is usually assumed to be anti-monotonic with respect to d. Namely:

$$\text{if } d(w,y) \leq d(w',y), \text{ then } g(w,y) \geq g(w',y)$$

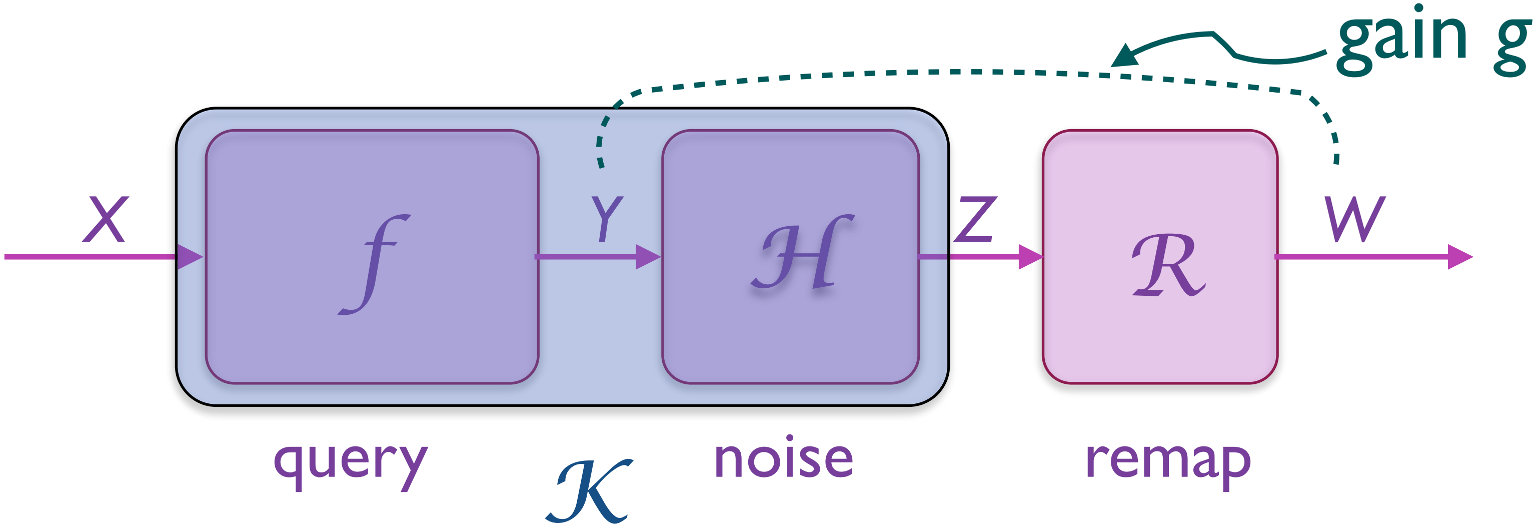

Schema for an oblivious mechanism. In a non-oblivious one *Z* depend also on *X*.

Utility

- Given a database x, consider the expected gain over all possible reported answers, for a certain remapping r. For an oblivious mechanism this is given by the formula:

$$\sum_z p_{\mathcal{H}}(z|f(x))g(r(z), f(x))$$

- For a generic (possibly non oblivious) mechanism, this is given by:

$$\sum_z p_{\mathcal{K}}(z|x)g(r(z), f(x))$$

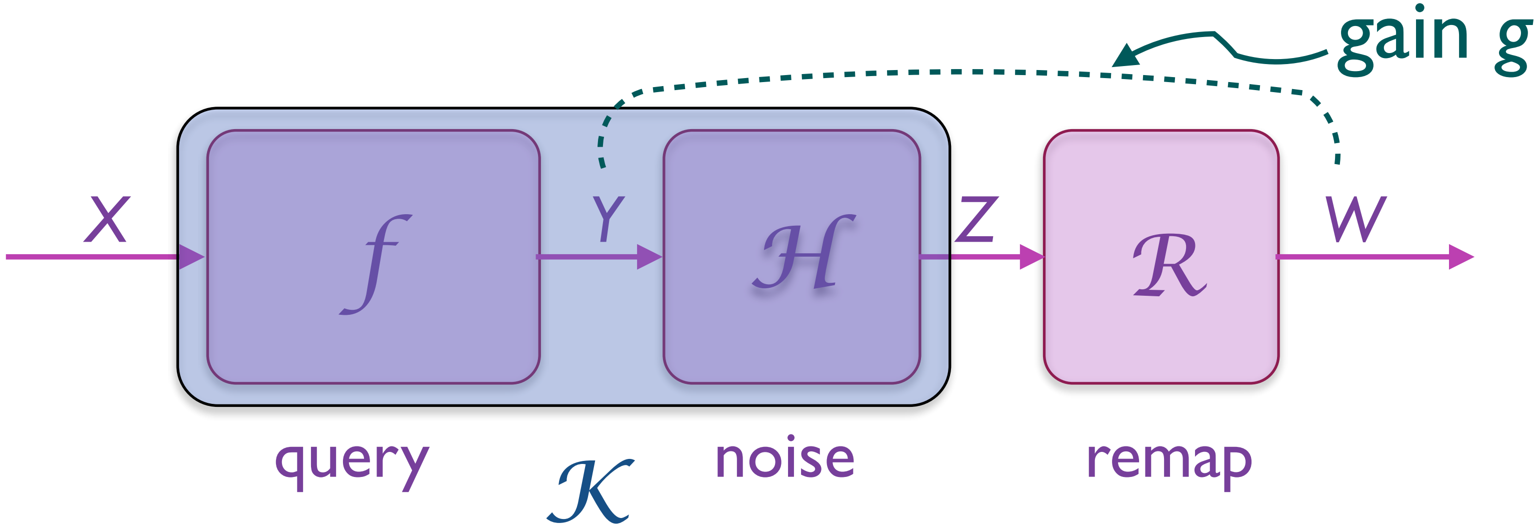

Schema for an oblivious mechanism. In a non-oblivious one *Z* depend also on *X*.

Utility

- The utility $\mathcal{U}$ of a mechanism is the maximum expected gain over all possible databases. The maximum is over all possible remappings: It is assumed that the user is rational and therefore makes the guesses that are the most useful to him. Note that $\mathcal{U}$ depends also on the prior π over $\mathcal{X}$ Formally, let us denote by r a remapping function. For an oblivious mechanism we have:

$$\mathcal{U}(\mathcal{K}, \pi, g) = \max_r \sum_x \pi(x) \sum_z p_{\mathcal{H}}(z|f(x)) g(r(z), f(x))$$

For a general (possibly non-oblivious) mechanism, we have:

$$\mathcal{U}(\mathcal{K}, \pi, g) = \max_r \sum_x \pi(x) \sum_z p_{\mathcal{K}}(z|x) g(r(z), f(x))$$

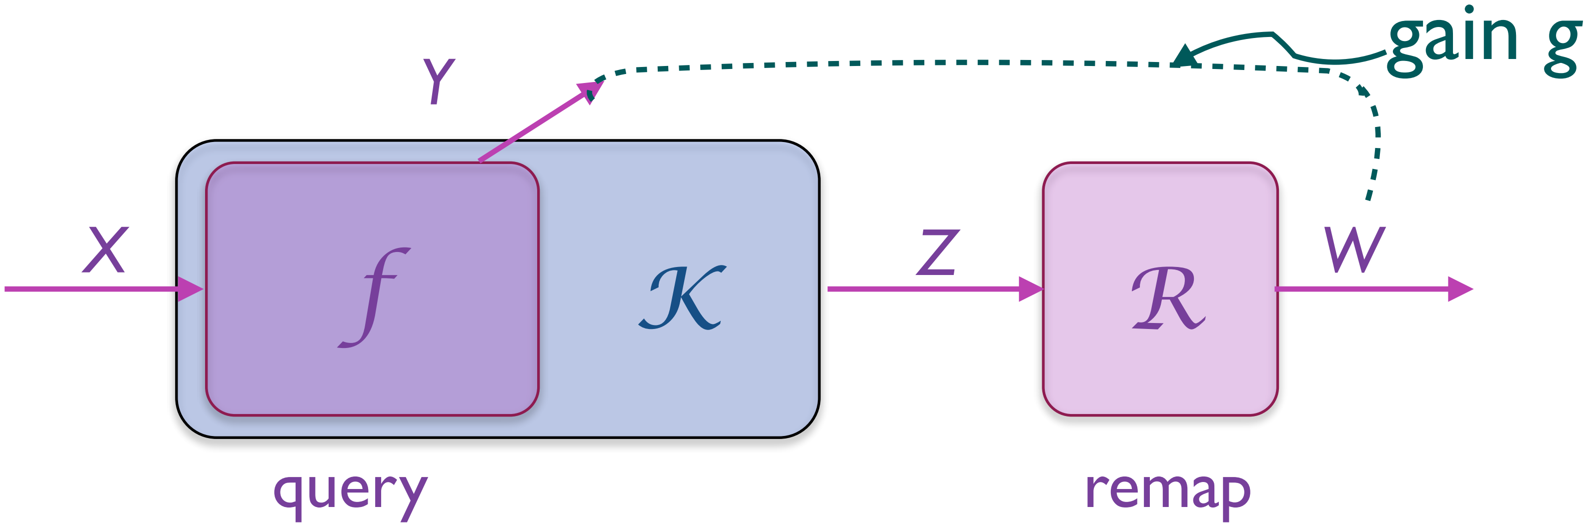

Example

The simplest gain function is the identity relation:

$$g(w,x) = \begin{cases} 1 & w = x \\ 0 & w \neq x \end{cases}$$

It represents the situation in which we are happy only if we guess the true answer.

With this gain function, the utility becomes (we give the formula for the oblivious case, the non-oblivious one is analogous):

$$\begin{aligned} \mathcal{U}(\mathcal{K},\pi,g) &= \max_r \sum_x \pi(x) \sum_z p_{\mathcal{H}}(z|f(x))\, g(r(z), f(x)) \\ &= \max_r \sum_y p_f(y) \sum_z p_{\mathcal{H}}(z|y)\, g(r(z), y) \\ &= \sum_z \max_y (p_f(y)\, p_{\mathcal{H}}(z|y)) \end{aligned}$$

This utility function essentially gives the expected probability of guessing the true answer. It is the converse of the Bayes risk

Example

Another typical gain function is the converse of the distance:

$$g(w, x) = D - d(w, x)$$

where D is the maximum possible distance between reported answers and true answers (it works well for truncated mechanisms). If such maximum does not exists, we can take $D = 0$. The only problem is that we get negative gains With this gain function, the utility is the expected distance between our best guess and the true answer. It gives a measure of how good is the approximated of the true answer that we can get with the mechanism.

Optimal mechanisms

- Given a prior π, and a privacy level ε, an ε-differentially private mechanism K is called optimal if it provides the best utility among all those which provide ε-differential privacy
- Note that the privacy does not depend on the prior, but the utility (in general) does.
- In the finite case the optimal mechanism can be computed with linear optimization techniques, where the variables are the conditional probabilities $p(z \mid y)$
  where y is the exact answer and z is the reported answer
- A mechanism is universally optimal if it is optimal for all priors π

Counting Queries

- Counting queries are typical examples of discrete queries. They are of the form: How many individuals in the database satisfy the property $\mathcal{P}$?
  - Examples:
    - How many individuals are affected by diabetes?
    - How many diabetic people are obese?
- Question: what is the sensitivity of a counting query?

Privacy vs utility: two fundamental results

1. [Ghosh et al., STOC 2009]
   The geometric mechanism and the truncated geometric mechanism are universally optimal for counting queries and any anti-monotonic gain function

Privacy vs utility: two fundamental results

2. [Brenner and Nissim, STOC 2010] The counting queries are the only kind of queries for which a universally optimal mechanism exists

   - This means that for other kind of queries one the optimal mechanism is relative to a specific user.
   - The precise characterization is given in terms of the graph $(\mathcal{Y},\sim)$ induced by $(\mathcal{X},\sim)$

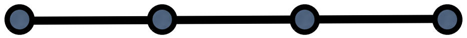

ok

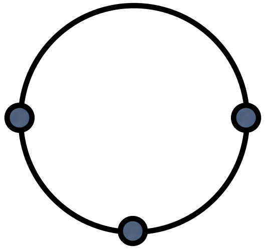

not ok

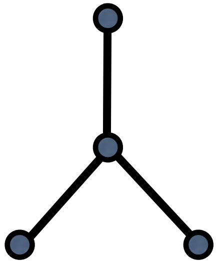

not ok

Exercises

1. Define the noise density function for the Laplacian mechanism for the query "What is the percentile of the people in the DB who earn more than 10K Euro a month", assuming that the database contains at least 1000 elements.

2. Define the truncated Laplacian mechanism for the above query. Note that $\mathcal{Y}$ is the interval [0,100].

3. Prove that ε-differential privacy can be equivalently defined as follows

$\mathcal{K}$ is ε-differentially private if for every pair of databases $x_1, x_2 \in \mathcal{X}$ (not necessarily adjacents), and for every $z \in \mathcal{Z}$, we have:

$$p(Z = z|X = x_1) \leq e^{\varepsilon h(x_1,x_2)} p(Z = z|X = x_2)$$

where $h(x_1, x_2)$ represents the Hamming distance between x_1 and x_2

Exercises

4. Compute the utility of the geometric mechanism for the counting query "how many people are in the database", with privacy degree ε, on the uniform prior distribution, with the gain function defined as the identity relation.

5. Find a mechanism for the same counting query, with the same degree of privacy, but lower utility.

6. We saw that post-processing cannot decrease privacy. Can it decrease the utility? Motivate your answer.

7. Show the graph on $\mathcal{Y}$ induced by the query "what is the average age of the people in the database"

8. Prove the result of Brenner and Nissim at Page 22 (hint: find two prior distributions on $\mathcal{Y}$ which have different optimal mechanisms)