


Foundations of Privacy

Lecture 4, Part I

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

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

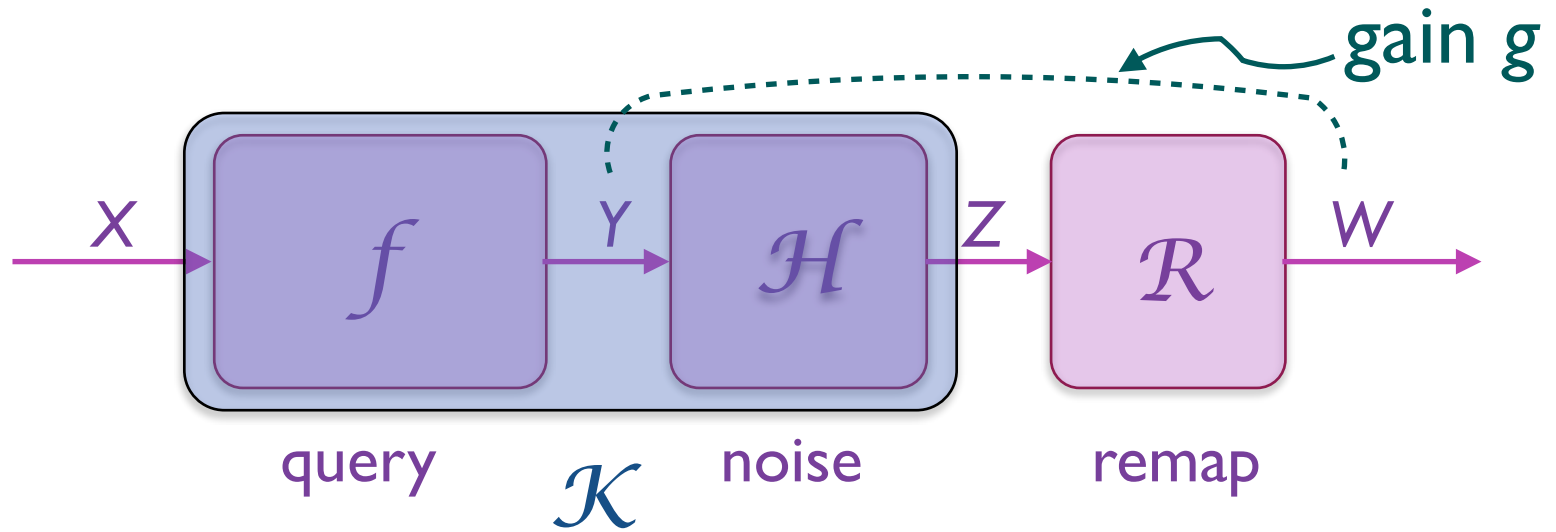
that represents the usefulness of the guess w when the true answer is y .

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

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

- 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

- 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 exist, 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 approximation 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 $p(z | y)$ 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




Privacy vs utility: two fundamental results

[Ghosh et al., STOC 2009]

1. The geometric mechanism and the truncated geometric mechanism are **universally optimal** for counting queries and any anti-monotonic gain function.
2. Every universally optimal mechanism can be remapped to the truncated geometric mechanism

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{V}, \sim) induced by (\mathcal{X}, \sim)



Exercises

1. Compute the utility of the geometric mechanism for a counting query, with privacy degree ϵ , on the uniform prior distribution, with the gain function defined as the identity relation.
2. Same exercise, but with the gain function defined as the converse of the distance.
3. Find a mechanism for the same counting query, with the same degree of privacy, but lower utility
4. We saw that post-processing cannot decrease privacy. Can it decrease the utility? Motivate your answer