A multiplicative version of the Kantorovich bisimulation metric to verify differential privacy

Kostas Chatzikokolakis, Catuscia Palamidessi, Lili Xu

INRIA & Ecole Polytechnique of Paris

A bit of history

- Our interest for differential privacy derives from our research on quantitative information flow, which started thanks to a collaboration with Prakash in 2005.
  "Go information theory!"

- The collaboration on QIF was also the basis of Kostas' PhD thesis

- Our interest for the Kantorovich metric derives from the seminal works by Prakash (and Desharnais, Jagadeesan, Gupta, CONCUR'99 , LICS'02)

Plan of the talk

- Motivations (statistical databases)
- Randomized mechanisms for queries
- Differential privacy
- Generalization to arbitrary metrics
- A multiplicative variant of the Kantorovich metric
- Properties

Differential Privacy

- Differential privacy [Dwork et al.,2006] is a notion of privacy originated from the area of **Statistical Databases**

- **The problem:** we want to use databases to get statistical information (aka aggregated information), but without violating the privacy of the people in the database

The problem

- The statistical queries should not reveal private information.

- Example: in a database meant to study a certain disease, we may want to ask queries that reveal the correlation between the disease and the age, but we should not be able to derive from this info whether a certain person has the disease.

name	age	disease
Alice	30	no
Bob	30	no
Don	40	yes
Ellie	50	no
Frank	50	yes

Query:
What is the youngest age of a person with the disease?

Answer:
40

Problem:
The adversary may know that Don is the only person in the database with age 40

The problem

- The statistical queries should not unveil private information.

- Example: in a database meant to study a certain disease, we may want to ask queries that reveal the correlation between the disease and the age, but we should not be able to derive from this info whether a certain person has the disease.

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

k-anonymity: the answer always partition the space in groups of at least k elements

Alice	Bob
Carl	Don
Ellie	Frank

The problem

Unfortunately, k-anonymity is very fragile under composition:

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

Alice	Bob
Carl	Don
Ellie	Frank

The problem of composition

Consider the query:

What is the minimal weight of a person with the disease?

Answer: 100

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

The problem of composition

Combine with the two queries:

minimal weight and the minimal age of a person with the disease

Answers: 40, 100

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

Solution

Introduce some noise on the answer, so that the answers of minimal age and minimal weight can be given also by other people with different age and weight

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers

minimal age:

40 with probability 1/2

30 with probability 1/4

50 with probability 1/4

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers

minimal weight:
100 with prob. 4/7
90 with prob. 2/7
60 with prob. 1/7

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers

Combination of the answers

The adversary cannot tell for sure whether a certain person has the disease

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers: a typical mechanism

- Randomized mechanism for a query $f: \mathcal{X} \to \mathcal{Y}$.
  Instead of the exact answer to the query, the curator gives a randomized answer $\mathcal{K}: \mathcal{X} \to \mathcal{Z}$ ($\mathcal{Z}$ may be different from $\mathcal{Y}$)
- The principle: little noise in global info produces large noise in individual info
- A typical randomized method: **the Laplacian noise.** If the exact answer is y, the reported answer is z, with a probability density function defined as:

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

where Δf is the *sensitivity* of f:

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

and c is a normalization factor:

$$c = \frac{1}{2\,\Delta f}$$

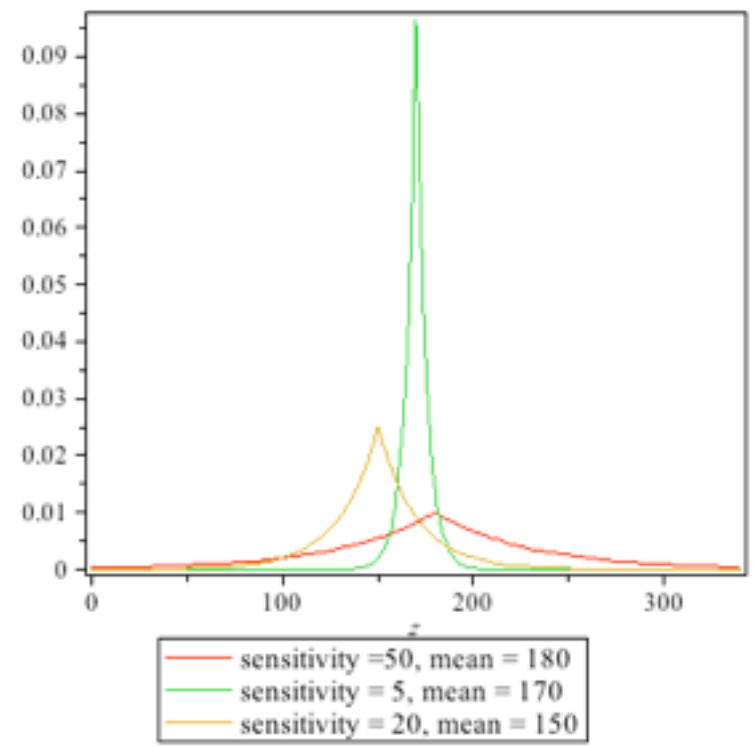

Privacy and Utility

- The two main criteria by which we judge a randomized mechanism:
  - **Privacy:** how good is the protection against leakage of private information
  - **Utility:** how useful is the reported answer

- Clearly there is a trade-off between privacy and utility, but they are not the exact opposites: privacy is about the individual data, while utility is about the aggregate data.

Differential Privacy

- There have been various attempts to formalize the notion of privacy, but the most successful one is the notion of Differential Privacy, recently introduced by Dwork
- **Differential Privacy** [Dwork 2006]: a randomized function $\mathcal{K}$ provides ε-differential privacy if for all adjacent databases x, x', and for all $z \in \mathcal{Z}$, we have

$$\frac{p(K = z|X = x)}{p(K = z|X = x')} \leq e^{\epsilon}$$

- The idea is that the likelihoods of x and x' are not too far apart, for every S
- Differential privacy is robust with respect to composition of queries
- The definition of differential privacy is independent from the prior (but this does not mean that the prior doesn't help in breaching privacy!)

Differential Privacy: alternative definition

- Perhaps the notion of differential privacy is easier to understand under the following equivalent characterization.
- In the following, X_i is the random variable representing the value of the individual i, and $X_{\neq i}$ is the random variable representing the value of all the other individuals in the database

- **Differential Privacy, alternative characterization:** a randomized function $\mathcal{K}$ provides ε-differential privacy if:

$$\text{for all } x \in \mathcal{X}, z \in \mathcal{Z}, p_i(\cdot)$$

$$\frac{1}{e^{\epsilon}} \leq \frac{p(X_i = x_i | X_{\neq i} = x_{\neq i})}{p(X_i = x_i | X_{\neq i} = x_{\neq i} \wedge K = z)} \leq e^{\epsilon}$$

Common misinterpretation and importance of the side knowledge

Generalization

Equivalent definition of Differential Privacy:

A mechanism is ϵ-differentially private iff for every pair of databases x, x' we have:

$$p(Z = z|X = x) \leq e^{\epsilon\, d_H(x,x')} p(Z = z|X = x')$$

where d_H is the Hamming distance between databases:
$d_H(x, x') =$ number of individuals in which x and x' differ.

On a generic domain $\mathcal{X}$ provided with a metric d:

$$p(Z = z|X = x) \leq e^{\epsilon\, d(x,x')} p(Z = z|X = x')$$

d-privacy

Example: Location Based Services

- Use an LBS to find a restaurant
- Without revealing the exact location
- Revealing an approximate location is ok

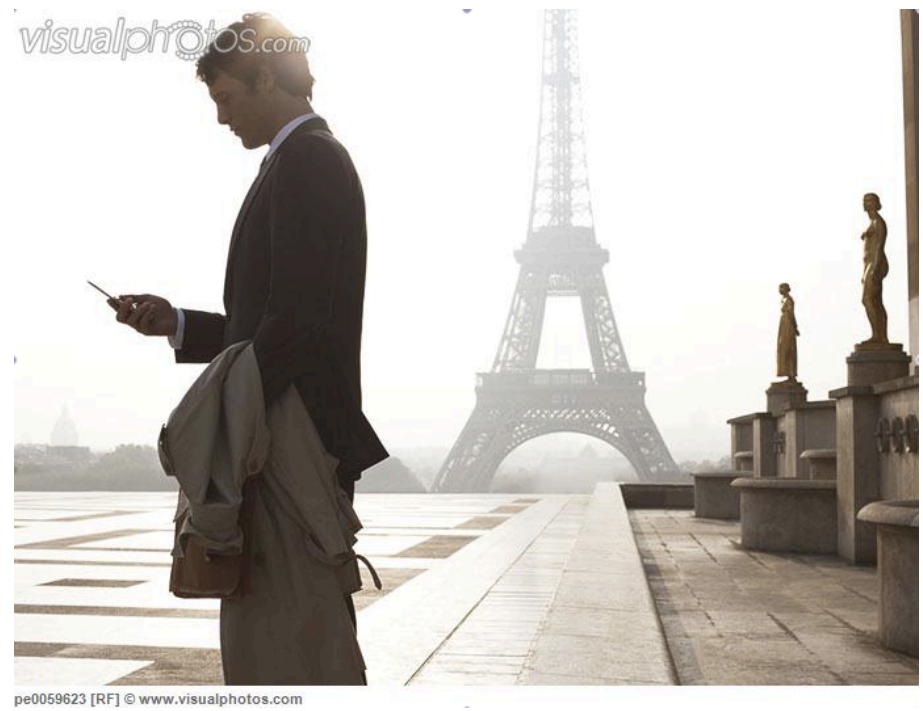

Example: Location Based Services

d : the Euclidean distance

x : the exact location

z : the reported location

d-privacy:

$$\frac{p(x|z)}{p(x'|z)} \leq e^{\epsilon r} \frac{p(x)}{p(x')}$$

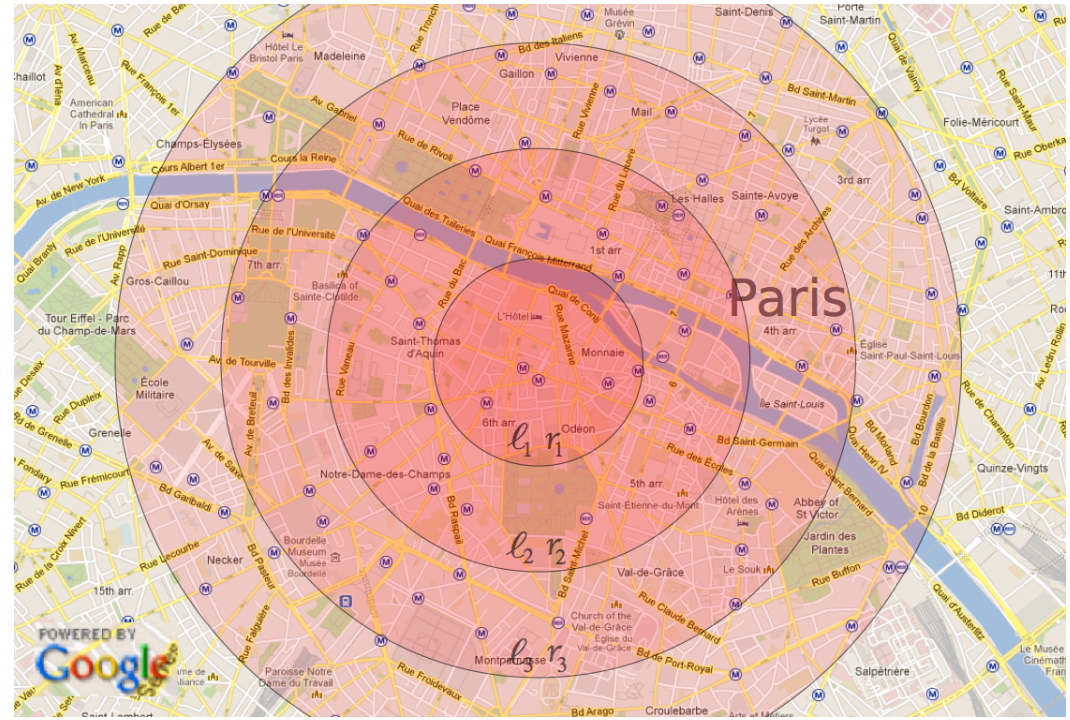

geo-indistinguishability

A *d*-private mechanism for LBS

Bivariate Laplacian

$$dp_x(z) = \frac{\epsilon^2}{2\pi} e^{\epsilon\, d(x,z)}$$

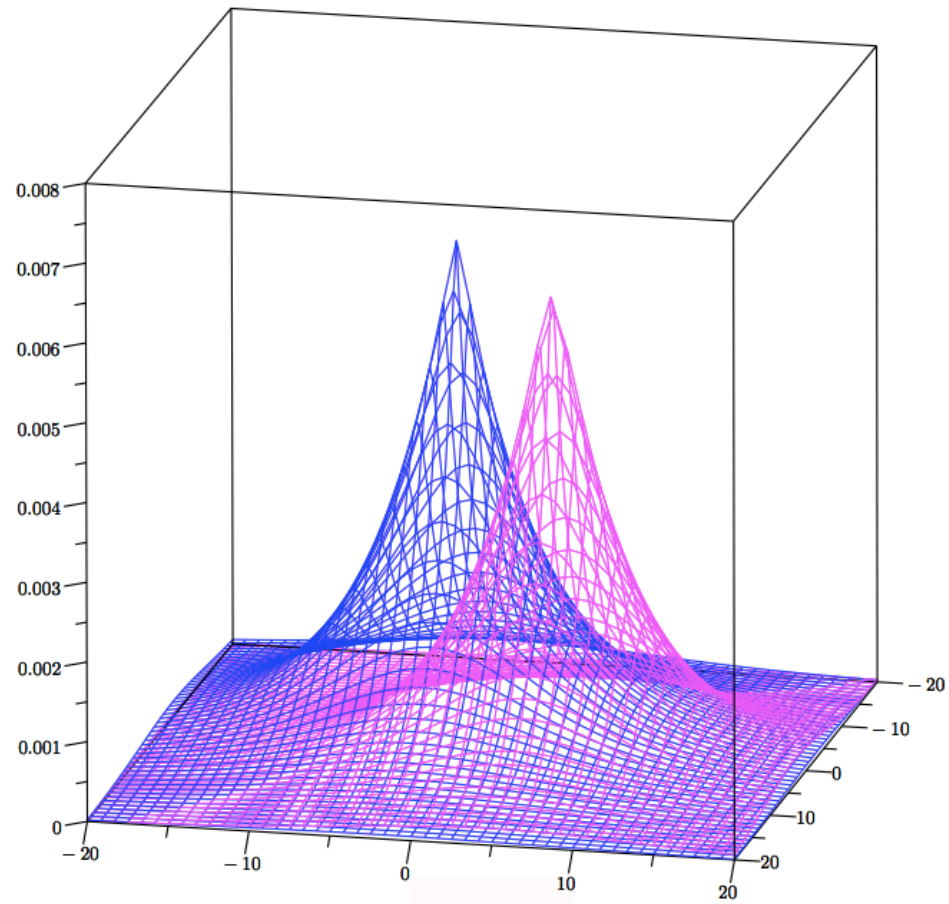

General Laplacian: $dp_x(z) = c\, e^{\epsilon\, d(x,z)}$ is d-private for any d

DP on probabilistic processes

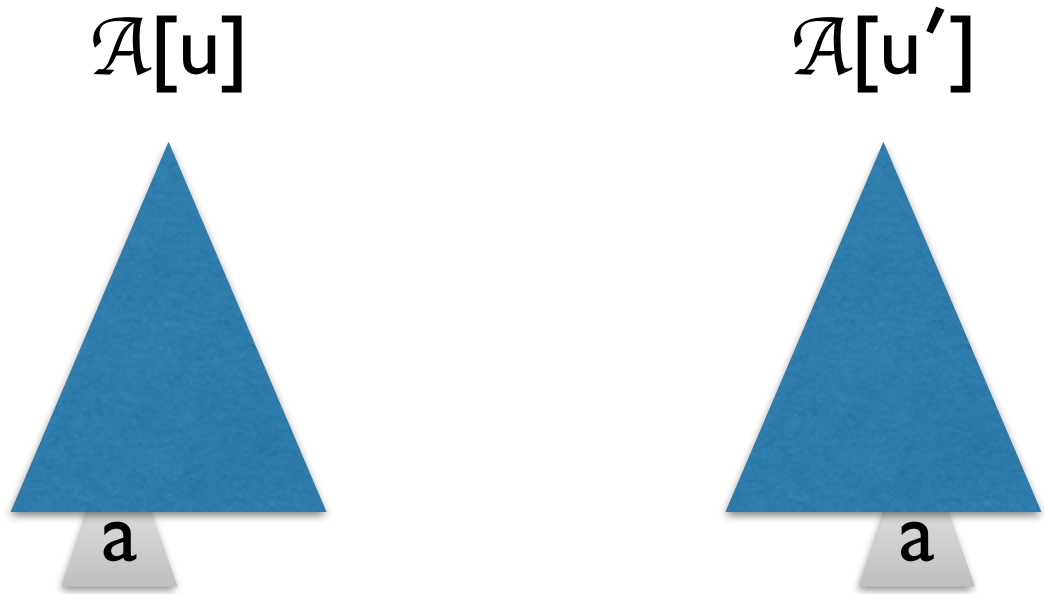

$$\frac{p(a|u)}{p(a|u')} \leq e^{d(u,u')}$$

The Kantorovich bisimulation metric and a multiplicative variant

The basic ingredient of the construction is the following transformation on metrics

	Standard [Prakash et al. van Breughel et al.]	Multiplicative variant
Primal	maximize $\sum_i(\mu(s_i) - \mu'(s_i))x_i$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i,j.\ x_i - x_j \leq m(s_i,s_j)$	maximize $\lvert \ln \frac{\sum_i \mu(s_i)x_i}{\sum_i \mu'(s_i)x_i} \rvert$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i,j.\ x_i \leq e^{m(s_i,s_j)}x_j$
Dual	minimize $\sum_{i,j} l_{ij}m(s_i,s_j)$ subject to $\forall i.\ \sum_j l_{ij} = \mu(s_i)$ $\forall j.\ \sum_i l_{ij} = \mu'(s_j)$ $\forall i,j.\ l_{ij} \geq 0$	minimize $\ln z$ subject to $\forall i.\ \sum_j l_{ij} - r_i = \mu(s_i)$ $\forall j.\ \sum_i l_{ij}e^{m(s_i,s_j)} - r_j \leq z \cdot \mu'(s_j)$ $\forall i,j.\ l_{ij}, r_i \geq 0$

Properties of the multiplicative variant of the bisimulation metric

- Let m be the greatest fixpoint of the transformation (the smallest metric among those for which the transformation is not expansive)

- If $\mathbf{m}(s, s') = 0$ then s, s' are (probabilistically) bisimilar

- If for all u,u' $\mathbf{m}(\mathcal{A}[u],\mathcal{A}[u']) \leq d(u,u')$ then $\mathcal{A}$ is differentially private

Thank you !