Foundations of Privacy

Lecture 3

Catuscia Palamidessi

Utility of a mechanism

Utility

Let us start with an example. Suppose we have a medical database, and we want to use it to do research about a certain disease.

For instance, we want to ask queries like:

1. How many people in the DB have the disease?
2. What is the average age of the people with the disease?

Suppose we know that :

- there are 1000 people in the DB
- the maximum age is 120
- both queries are sanitised with DP

age	disease
41	no
45	yes
37	no
50	yes
...	...
20	no

Loss function

How to measure the quality of the reported answer?

Consider the first query: $f(x) =$ number of people with the disease.
Let $y = f(x)$ be the true answer, and z the reported answer.
Which of the following loss functions is better?

1. $\ell(y, z) = |z - y|$
2. $\ell(y, z) = (z - y)^2$
3. $\ell(y, z) = \begin{cases} 0 & \text{if } z = y \\ 1 & \text{if } z \neq y \end{cases}$
4. $\ell(y, z) = 0$
5. $\ell(y, z) = z + y$

Loss function

How to measure the quality of the reported answer?

Consider the first query: $f(x) =$ number of people with the disease.
Let $y = f(x)$ be the true answer, and z the reported answer.
Which of the following loss functions is better?

1. $\ell(y, z) = |z - y|$
2. $\ell(y, z) = (z - y)^2$
3. $\ell(y, z) = \begin{cases} 0 & \text{if } z = y \\ 1 & \text{if } z \neq y \end{cases}$
4. $\ell(y, z) = 0$
5. $\ell(y, z) = z + y$

(1), (2) and (3) are all reasonable loss functions, they all measure the "precision" of the answer. Which one is more suitable for our purposes depends on what we want to do.
On the other hand, (4) does not measure anything, and (5) does not make sense.

Monotonicity of the loss

In general, if $\mathcal{Y} \subseteq \mathcal{Z}$ and the domain $\mathcal{Z}$ is equipped with a notion of distance d, we want the loss to be *monotonic* w.r.t. d. Namely:

$$\ell(y, z) \leq \ell(y', z') \quad \Leftrightarrow \quad |z - y| \leq |z' - y'|$$

Utility as expected loss

Since there are many possible true answers, and even for the same true answer we have many possible reported answer, it is reasonable to define the utility as expectation.

Let π be the prior on $\mathcal{Y}$ (the true answers) and p the probability associated to the mechanism. We could define:

$$\begin{aligned} \mathcal{U}(\mathcal{K}, \pi) &= \mathbb{E}_{\pi,p}\ell(y,z) \\ &= \sum_{y,z} \pi(y)\, p(z|y)\, \ell(y,z) \end{aligned}$$

Utility as expected loss

Since there are many possible true answers, and even for the same true answer we have many possible reported answer, it is reasonable to define the utility as expectation.

Let π be the prior on $\mathcal{Y}$ (the true answers) and p the probability associated to the mechanism. We could define:

$$\begin{aligned} \mathcal{U}(\mathcal{K}, \pi) &= \mathbb{E}_{\pi,p}\ell(y,z) \\ &= \sum_{y,z} \pi(y)\, p(z|y)\, \ell(y,z) \end{aligned}$$

Are we happy with this definition?

Utility as expected loss

Since there are many possible true answers, and even for the same true answer we have many possible reported answer, it is reasonable to define the utility as expectation.

Let π be the prior on $\mathcal{Y}$ (the true answers) and p the probability associated to the mechanism. We could define:

$$\begin{aligned} \mathcal{U}(\mathcal{K}, \pi) &= \mathbb{E}_{\pi,p}\ell(y,z) \\ &= \sum_{y,z} \pi(y)\, p(z|y)\, \ell(y,z) \end{aligned}$$

Are we happy with this definition?

What if we get a negative answer? Or an answer greater than 1000, the number of people in the DB? (it could happen, for instance, with the geometric mechanism).

Utility as expected loss

Since there are many possible true answers, and even for the same true answer we have many possible reported answer, it is reasonable to define the utility as expectation.

Let π be the prior on $\mathcal{Y}$ (the true answers) and p the probability associated to the mechanism. We could define:

$$\begin{aligned} \mathcal{U}(\mathcal{K}, \pi) &= \mathbb{E}_{\pi,p}\ell(y,z) \\ &= \sum_{y,z} \pi(y)\, p(z|y)\, \ell(y,z) \end{aligned}$$

Are we happy with this definition?

What if we get a negative answer? Or an answer greater than 1000, the number of people in the DB? (it could happen, for instance, with the geometric mechanism).

We are not going to believe these answers, so we could remap them in more likely values. For instance we could remap the negative values into 0, and those greater than 1000 into 1000

Remapping

We could use a remapping function defined as:

$$r(z) = \begin{cases} 0 & \text{if } z < 0 \\ z & \text{if } 0 \leq z \leq 1000 \\ 1000 & \text{if } z > 1000 \end{cases}$$

and define

$$\mathcal{U}(\mathcal{K}, \pi) \;=\; \sum_{y,z} \pi(y) p(z|y)\, \ell(y, r(z))$$

Remapping

We could use a remapping function defined as:

$$r(z) = \begin{cases} 0 & \text{if } z < 0 \\ z & \text{if } 0 \leq z \leq 1000 \\ 1000 & \text{if } z > 1000 \end{cases}$$

and define

$$\mathcal{U}(\mathcal{K}, \pi) \;=\; \sum_{y,z} \pi(y) p(z|y)\, \ell(y, r(z))$$

More in general, we assume that we exploit the prior knowledge, and the knowledge of the mechanism, to define and use the best possible remapping function:

$$\mathcal{U}(\mathcal{K}, \pi) \;=\; \min_{r} \sum_{y,z} \pi(y) p(z|y)\, \ell(y, r(z))$$

Notes about utility

- We saw a definition for discrete mechanisms. For continuous ones, like the Laplace, the definition is analogous except that the expectation has to be computed via integration

- The expected loss is not the only definition of utility that has been considered in the literature. There are others, for instance the worst-case loss, the expected ratio of ``good'' answers, etc. For the next results, however, we will assume that utility is defined as expected loss.

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

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 in the DB 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 and the truncated geometric mechanisms are universally optimal for counting queries and any monotonic loss function

Open question: can we extend this result to the continuous case?

Privacy vs utility: two fundamental results

2. [Brenner and Nissim, STOC 2010] On a discrete domain, 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 $(\mathcal{X}, \sim)$

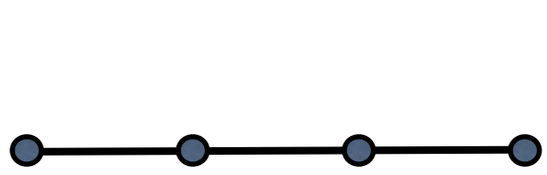

ok

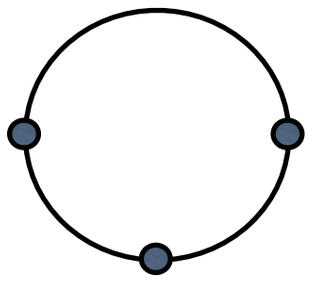

not ok

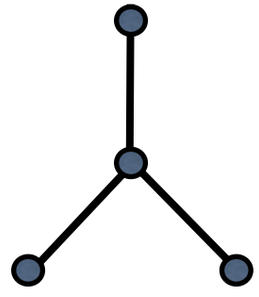

not ok

The Local Model

The Global Model

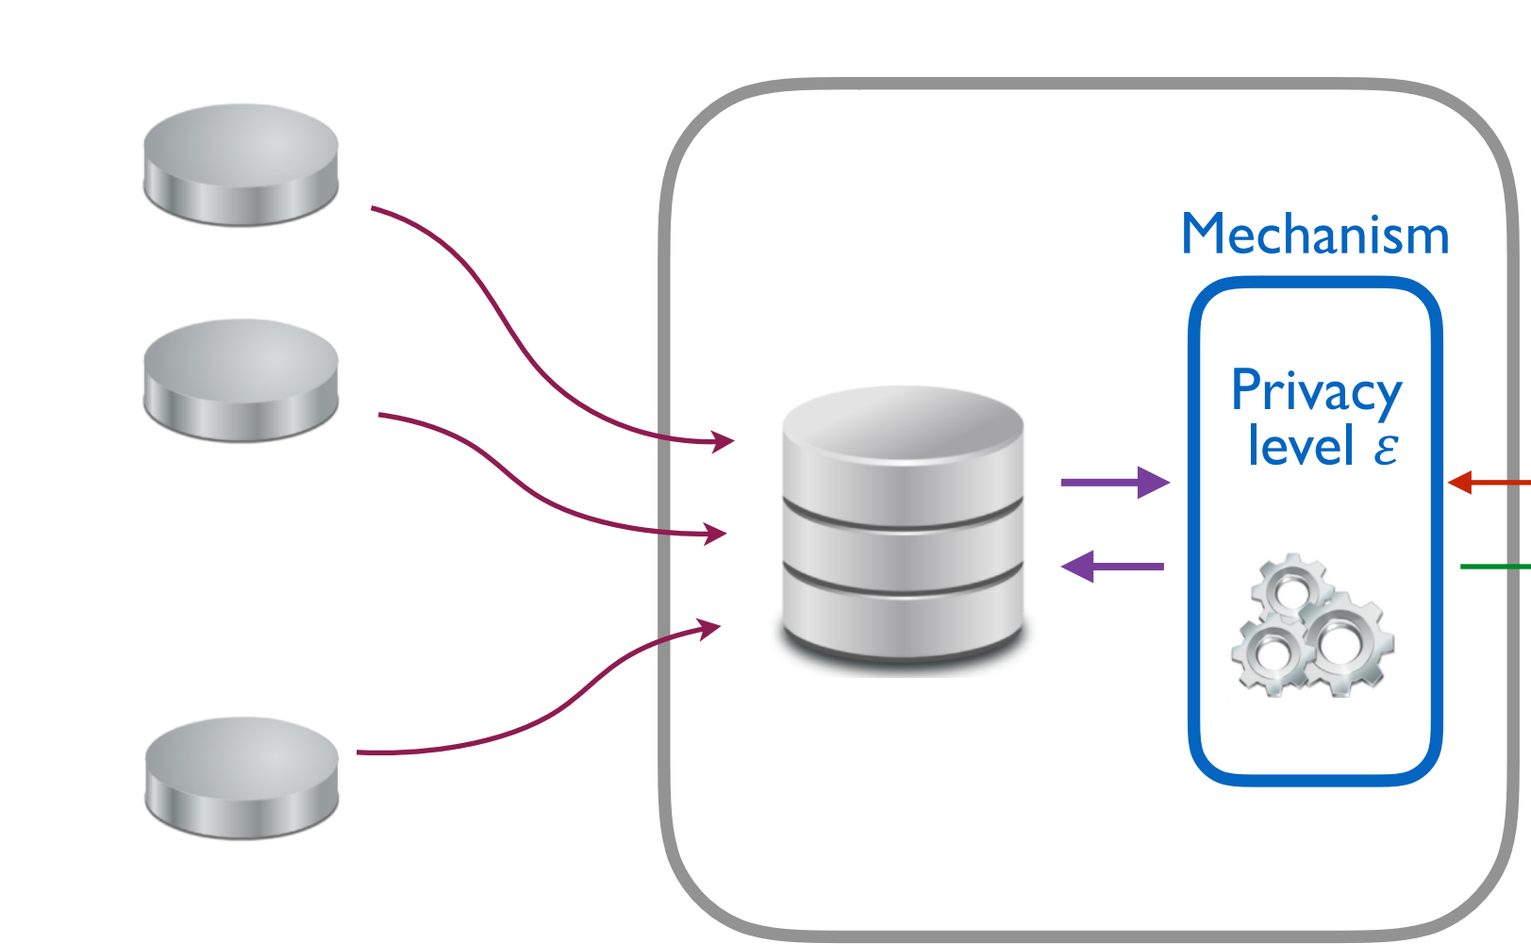

The Global Model

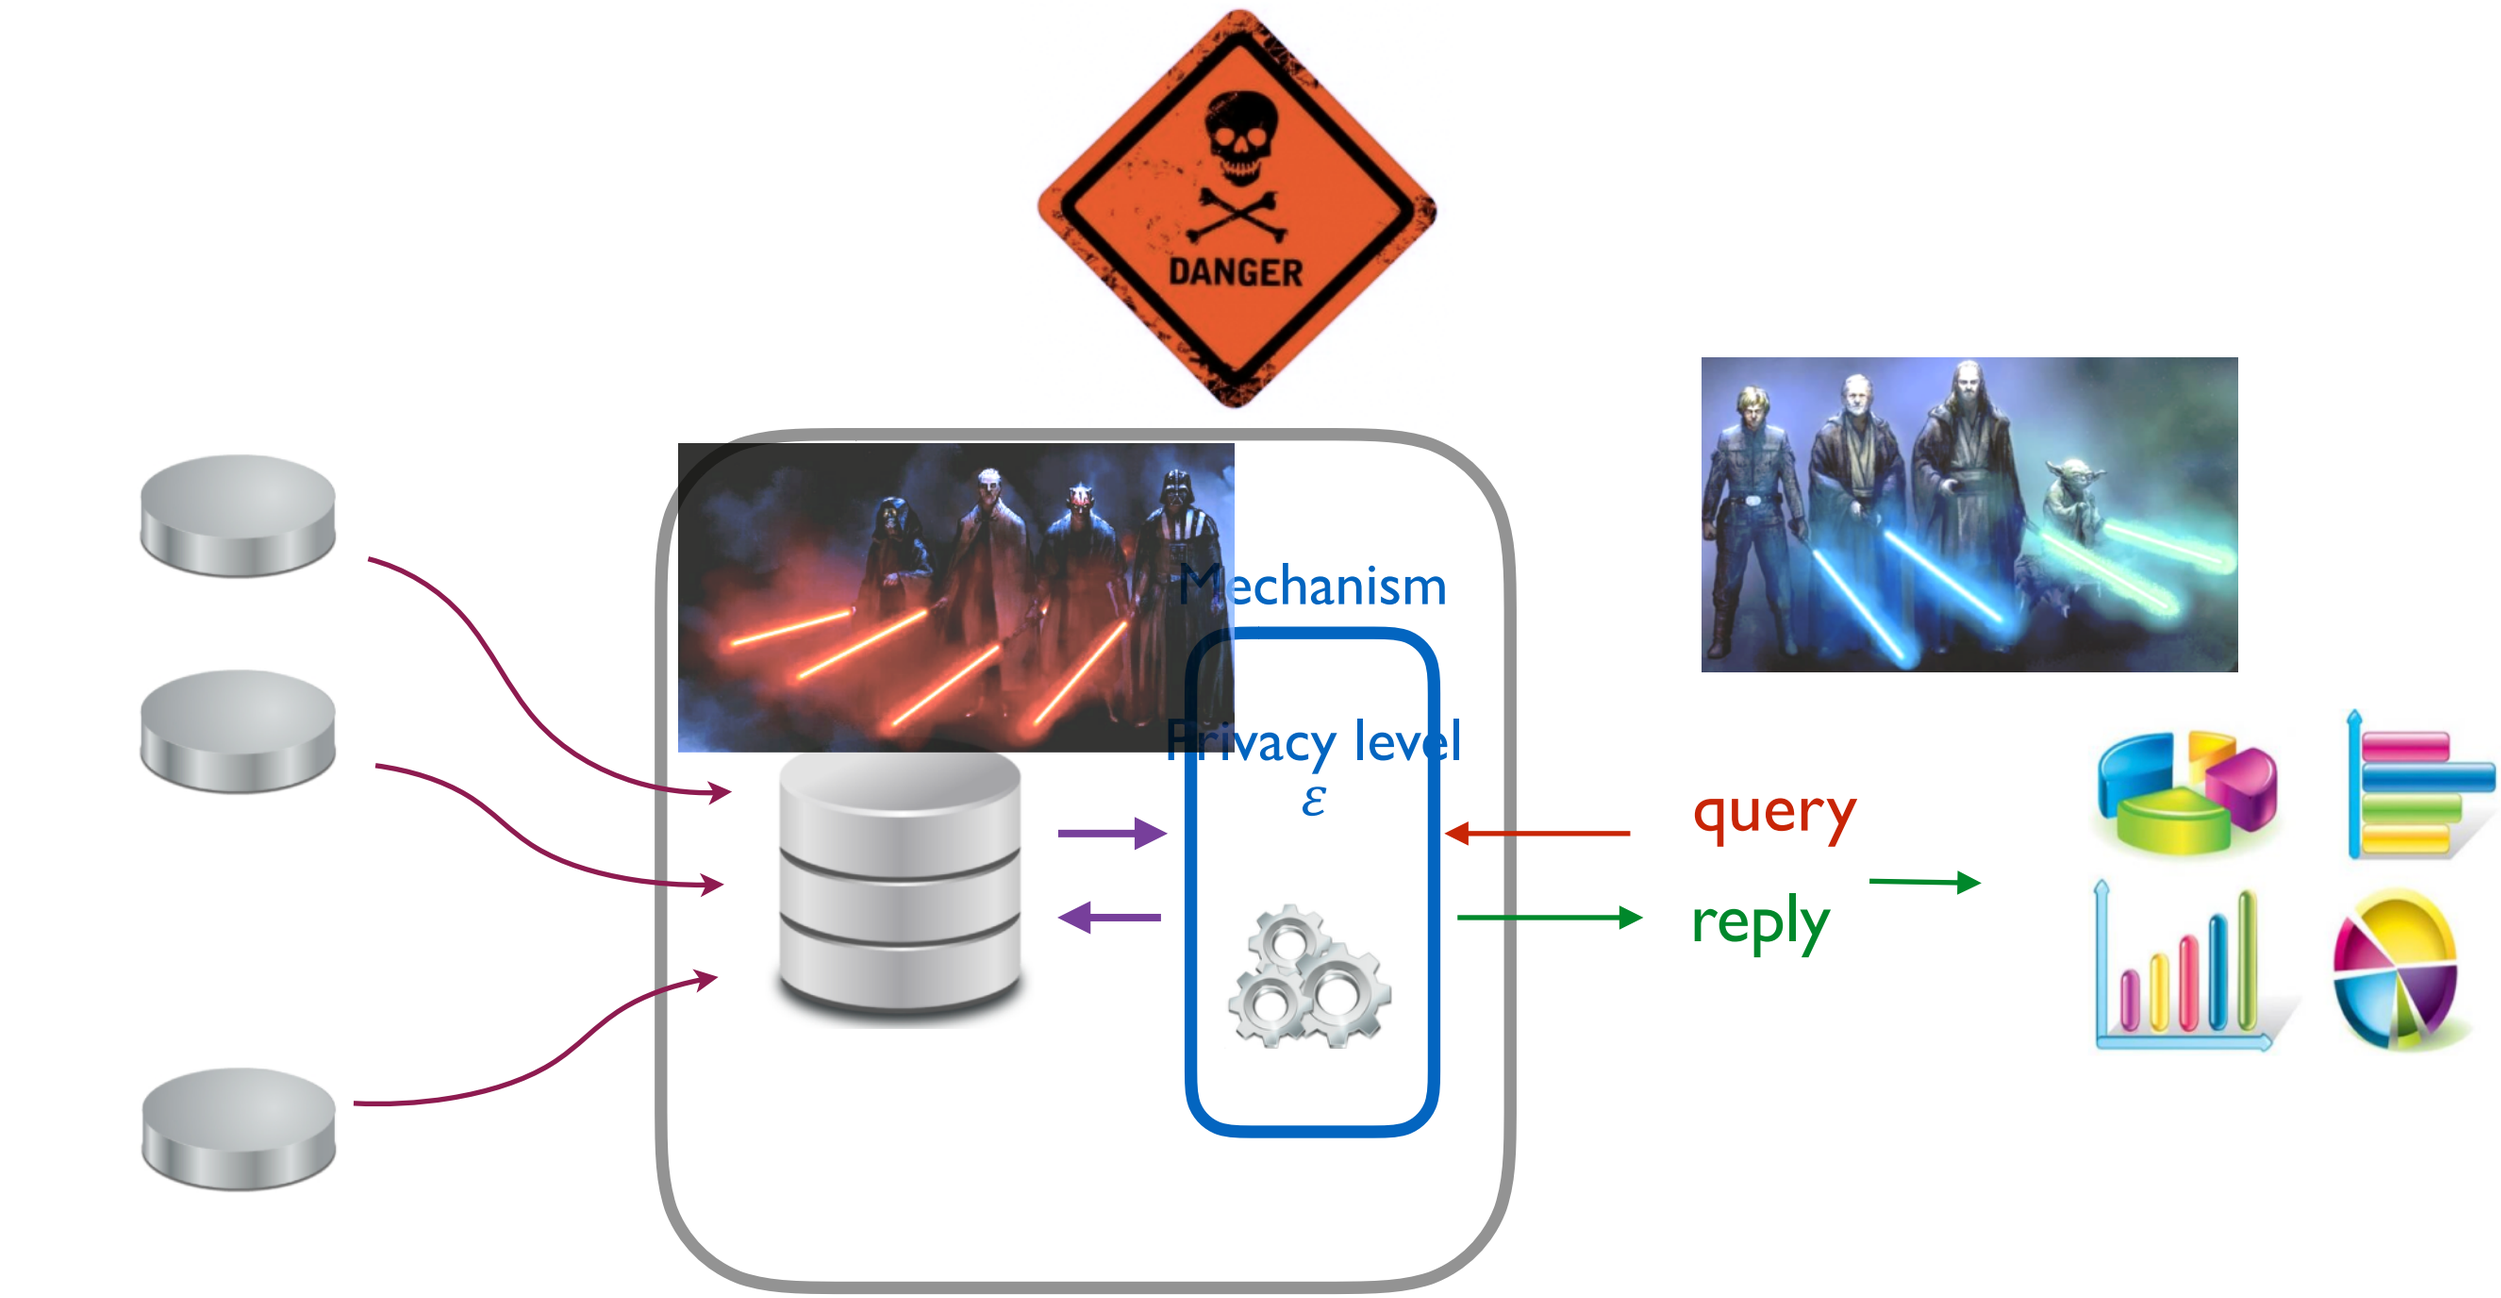

The Local Model

Google

Apple

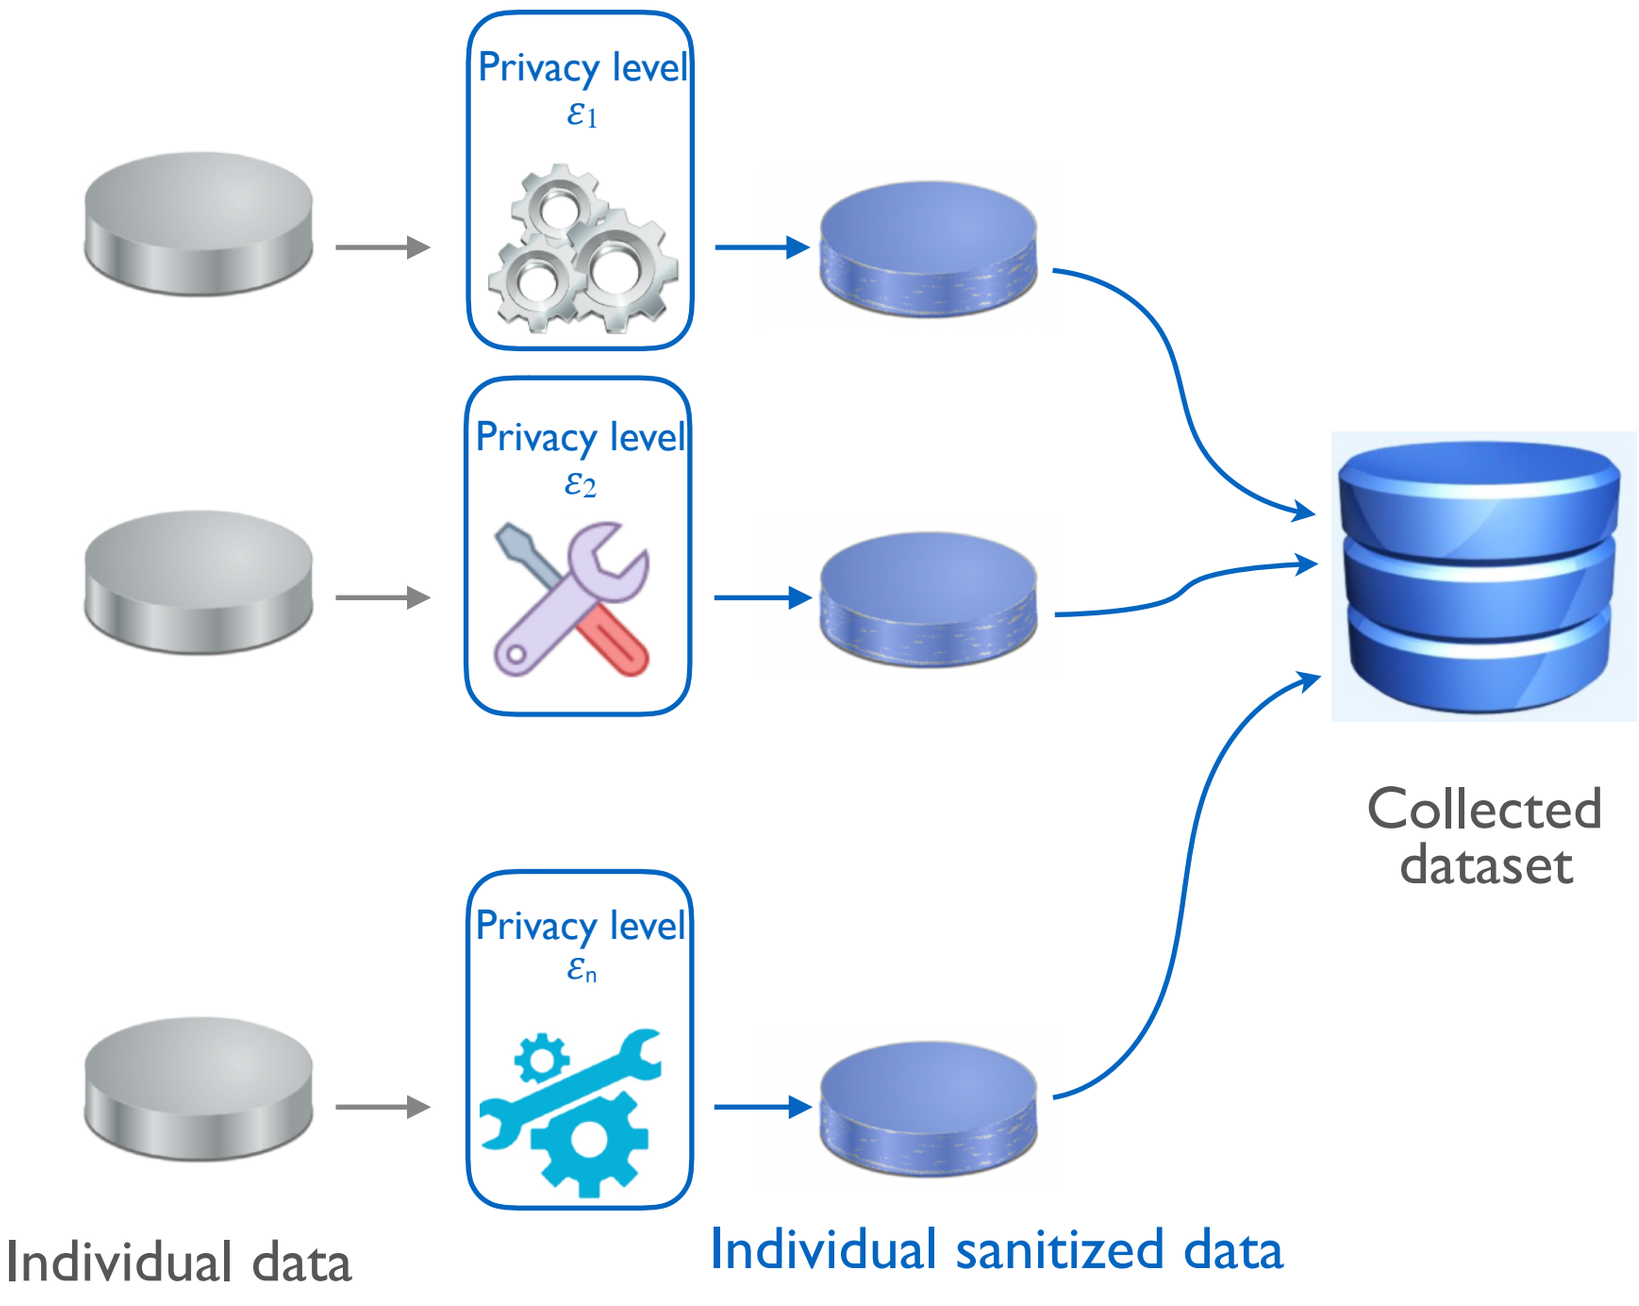

The Local Model

Google

Apple

Two notions of utility

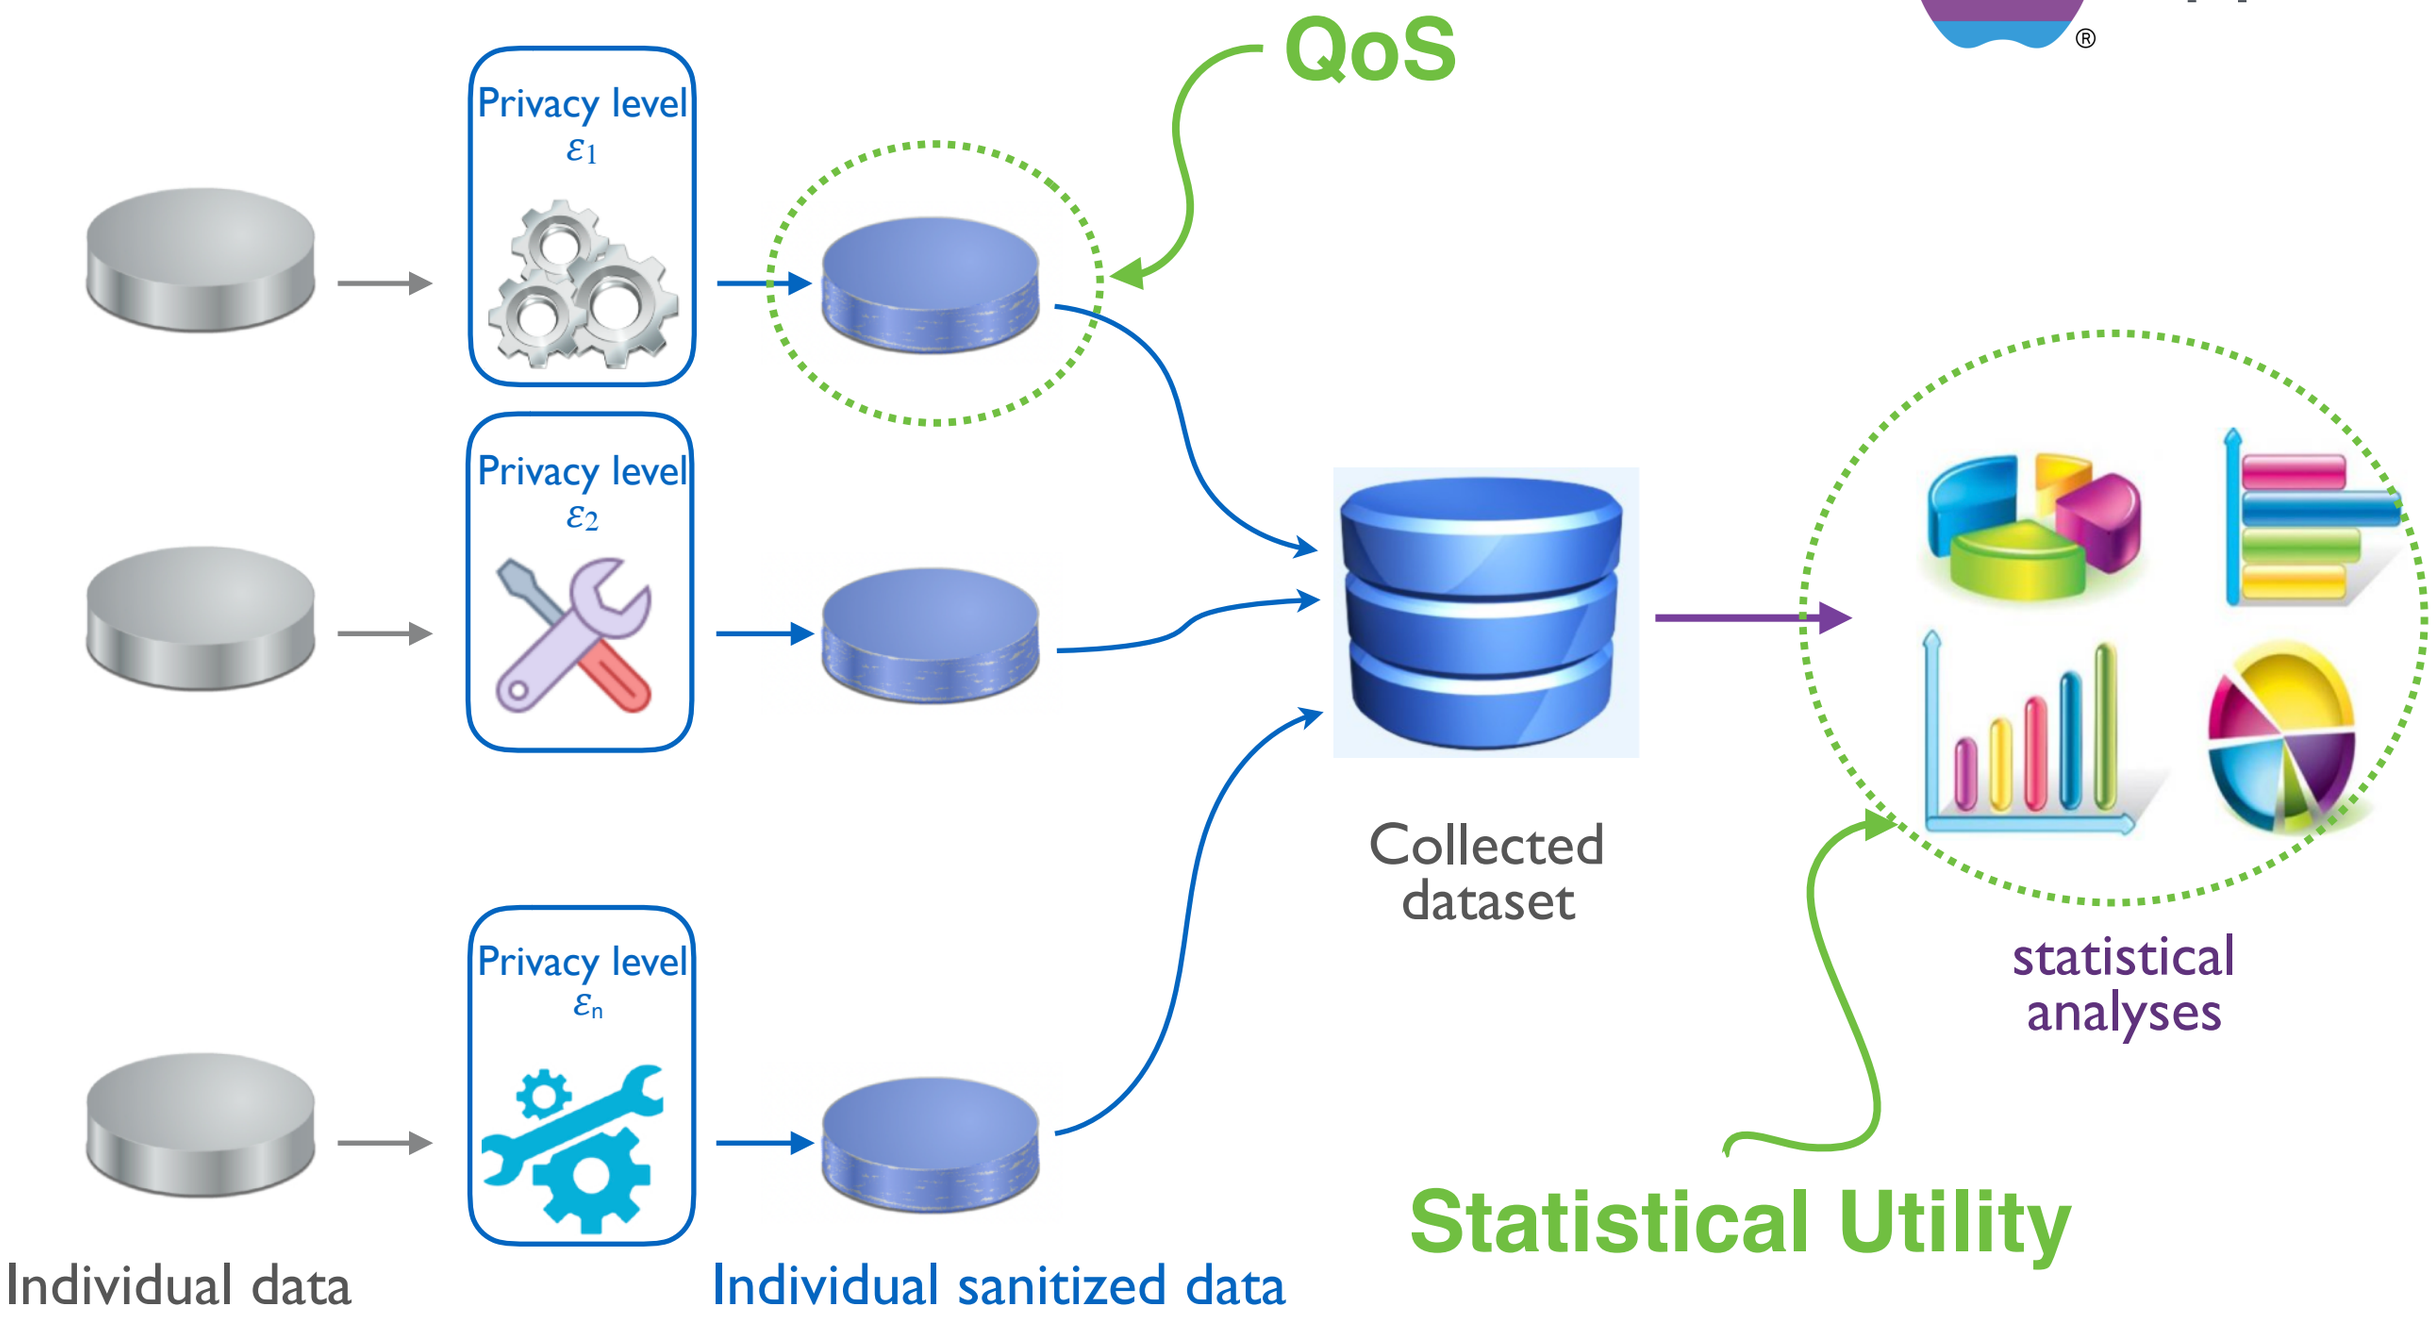

Local Differential Privacy

[Jordan &Wainwright '13]

One of the most popular definitions of privacy in the local model is Local Differential Privacy (LPD)

Definition Let $\mathcal{X}$ be a set of possible values and $\mathcal{Y}$ the set of noisy values. A mechanism $\mathcal{K}$ is ε-locally differentially private (ε-LDP) if for all $x_1, x_2 \in \mathcal{X}$ and for all $y \in \mathcal{Y}$

$$P[\mathcal{K}(x) = y] \leq e^{\varepsilon}\ P[\mathcal{K}(x') = y]$$

or equivalently, using the conditional probability notation:

$$p(y \mid x) \leq e^{\varepsilon}\ p(y \mid x')$$

Example:
The Randomized Response protocol

Suppose that I want to find out how many of you find my lectures boring.

If I ask you directly you will probaly answer "no" regardless of the truth.

So, I use the following following protocol instead:
You toss a coin, without showing me the result.
If the result is head, then you answer the truth, otherwise you toss the coin again, and answer according to the result ("Yes" if Head, "No" if Tail).

This protocol is called **Randomized Response**

Example:
The Randomized Response protocol

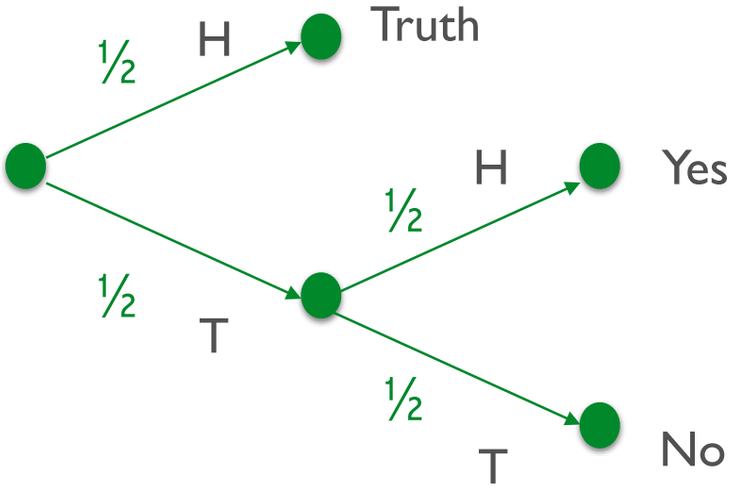

x \ y	yes	no
yes	¾	¼
no	¼	¾

Question: is the The Randomized Response protocol locally differentially private?

Yes, it is (log 3)-LPD

The k-RR mechanism (aka the flat m.)

[Kairouz et al, '16]

The k-RR is the extension of Randomized Response
to a secret's domain of k elements

The flat mechanism is the simplest way to implement LPD.
It is defined as follows:

$$p(y|x) = \begin{cases} c\,e^{\varepsilon} & \text{if } x = y \\ c & otherwise \end{cases}$$

where c is a normalization constant.

namely $c = \dfrac{1}{k - 1 + e^{\varepsilon}}$ where k is the size of the domain

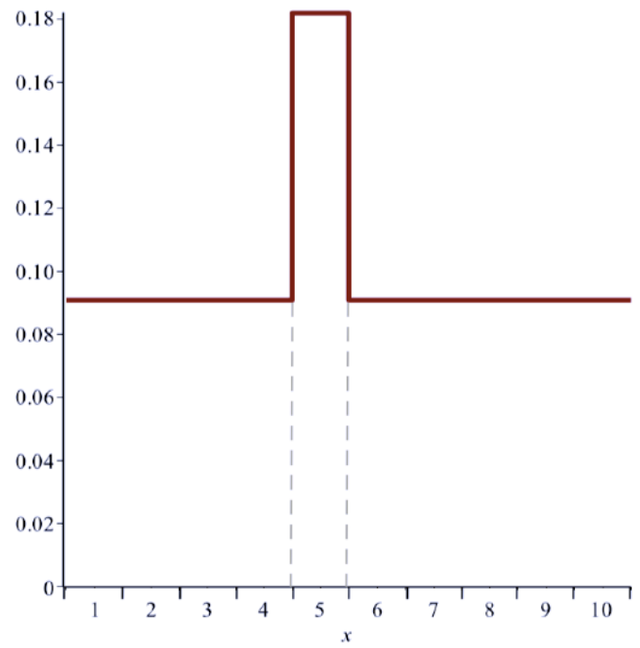

Privacy Properties:

- Compositionality
  (if we combine two LPD mechanisms, the resulting mechanism is still LPD)
- Independence from the side knowledge of the adversary

Another definition of privacy in the local model:

d-privacy

d-privacy: a generalization of DP and LDP

[Chatzikokolakis et al., '13]

d-privacy

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

$$\forall x, x' \in \mathcal{X}, \forall z \quad \frac{p(z \mid x)}{p(z \mid x')} \leq e^{\varepsilon\, d(x,x')}$$

generalizes

Differential Privacy

- x, x' are databases
- *d* is the Hamming distance

Local Differential Privacy

- *d* is the discrete distance

Intuition

d-privacy protects the *precision* of the secret. For instance, it allows to distinguish whetehr I am in Paris or London, but not where precisely I am in Paris

Properties

- Like LDP, it can be applied at the user side
- Like DP and LDP, it is compositional and independent from side knowledge of the adversary

Typical *d*-private mechanisms

Laplace, Geometric, and their Planar versions

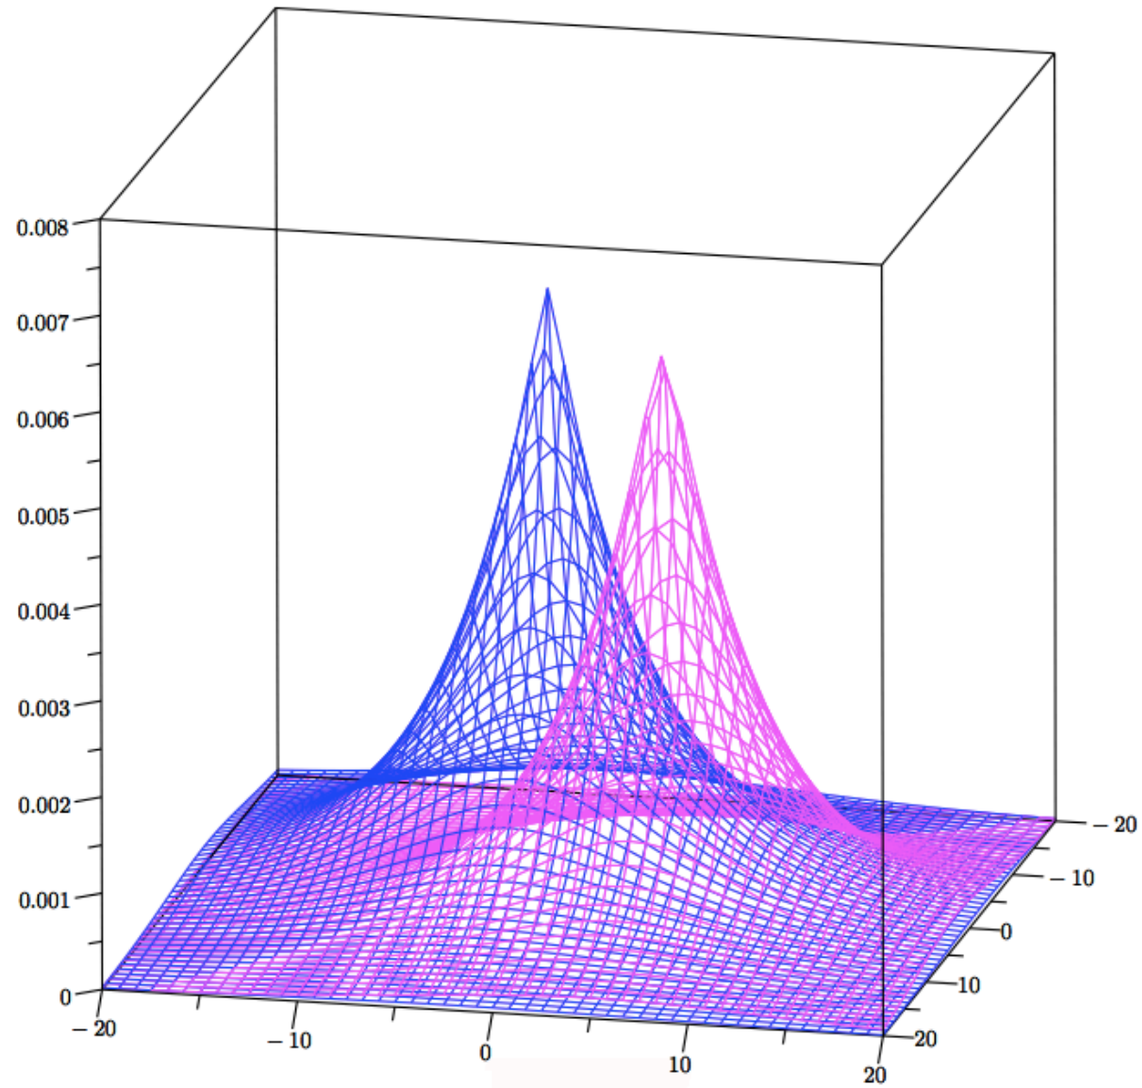

Planar Laplace

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

Used especially for location privacy

Example of application of *d*-privacy: Location Privacy for Location Based Services

- Example of LBS: find the restaurants near the user
- Revealing the exact location may be dangerous: profiling, inference of sensitive information, etc.
- Revealing an approximate location is usually ok
- The QoS decreases with the expected distance between the real location and the noisy one, so there is a trade-off.

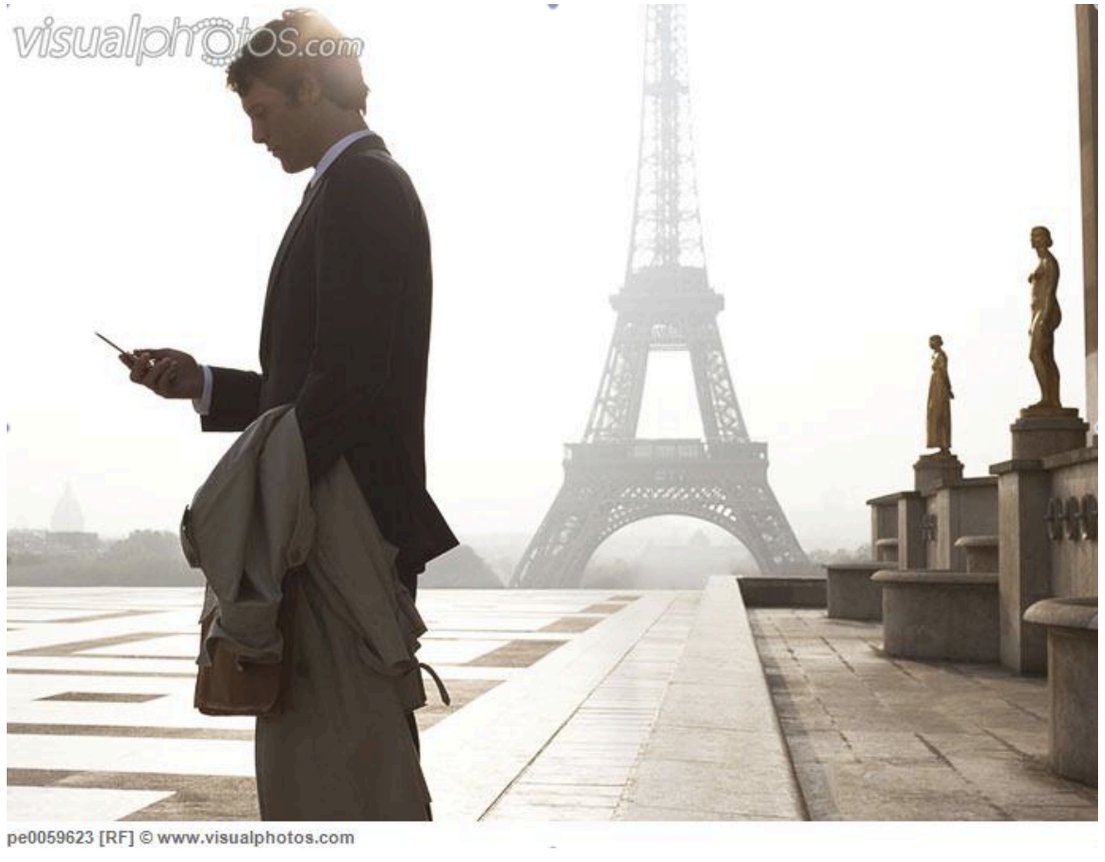

Geo-indistinguishability

In the case of location privacy, d-privacy is called **geo-indistinguishability**

d : the Euclidean distance

x : the exact location

z : the reported location

$d-$ privacy

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

where r is the distance between x and x'

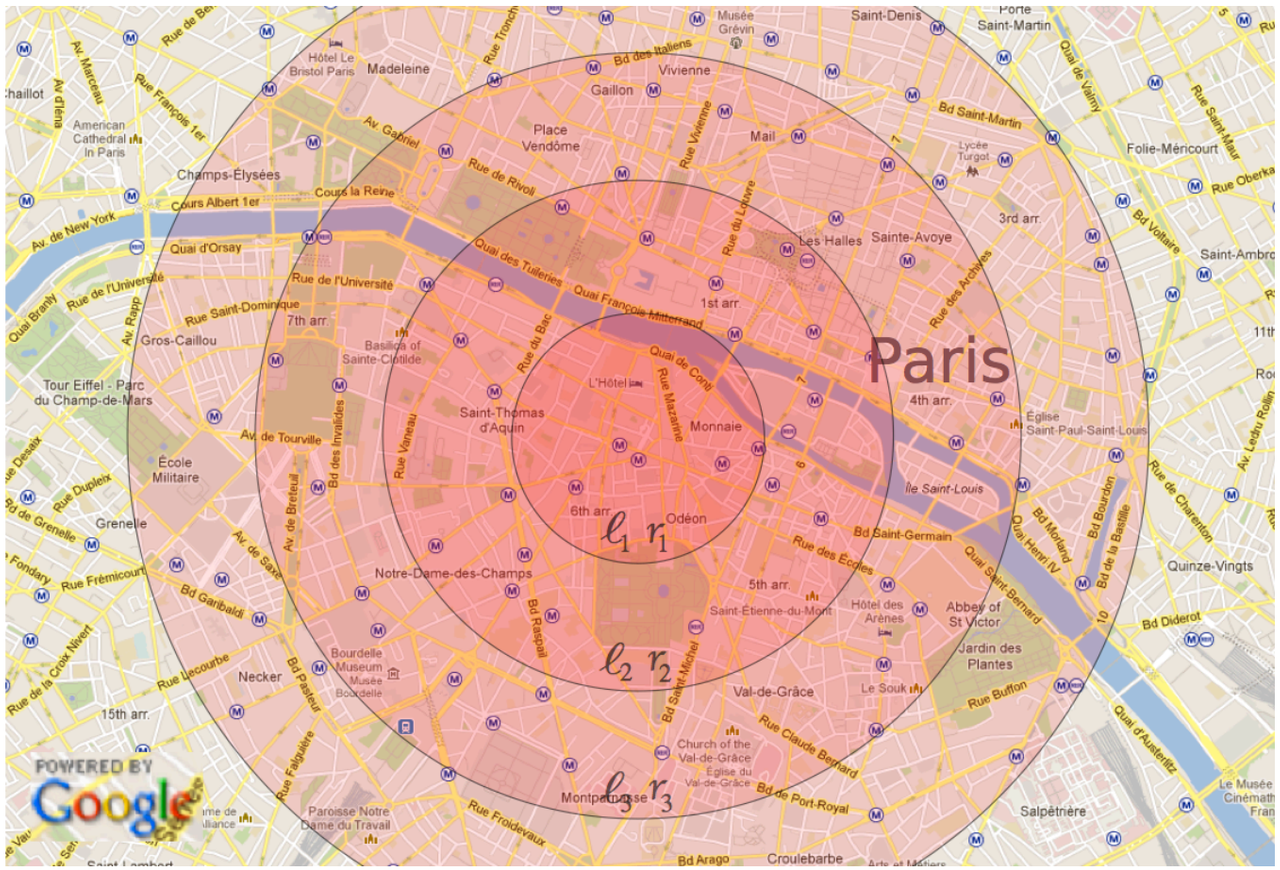

Like d-privacy, geo-indistinguishability is:

1) independent from the prior,
2) compositional

Meaning of geo-indistinguishability

$$\forall \pi. \quad \frac{P(x|z)}{P(x'|z)} \leq e^{\epsilon d(x,x')} \frac{\pi(x)}{\pi(x')}$$

The closer two points are,
the more they are indistinguishable

The level of distinguishability
also depends on the prior

We want to be unable to tell whether the user is in rue Pigalle or at Notre Dame, but it is ok to disclose that he is in Paris and not in London

Tool: "Location Guard"

http://www.lix.polytechnique.fr/~kostas/software.html

Extension for Firefox, Chrome, and Opera. It has been released about two yeas ago, and nowadays it has about 60,000 active users.

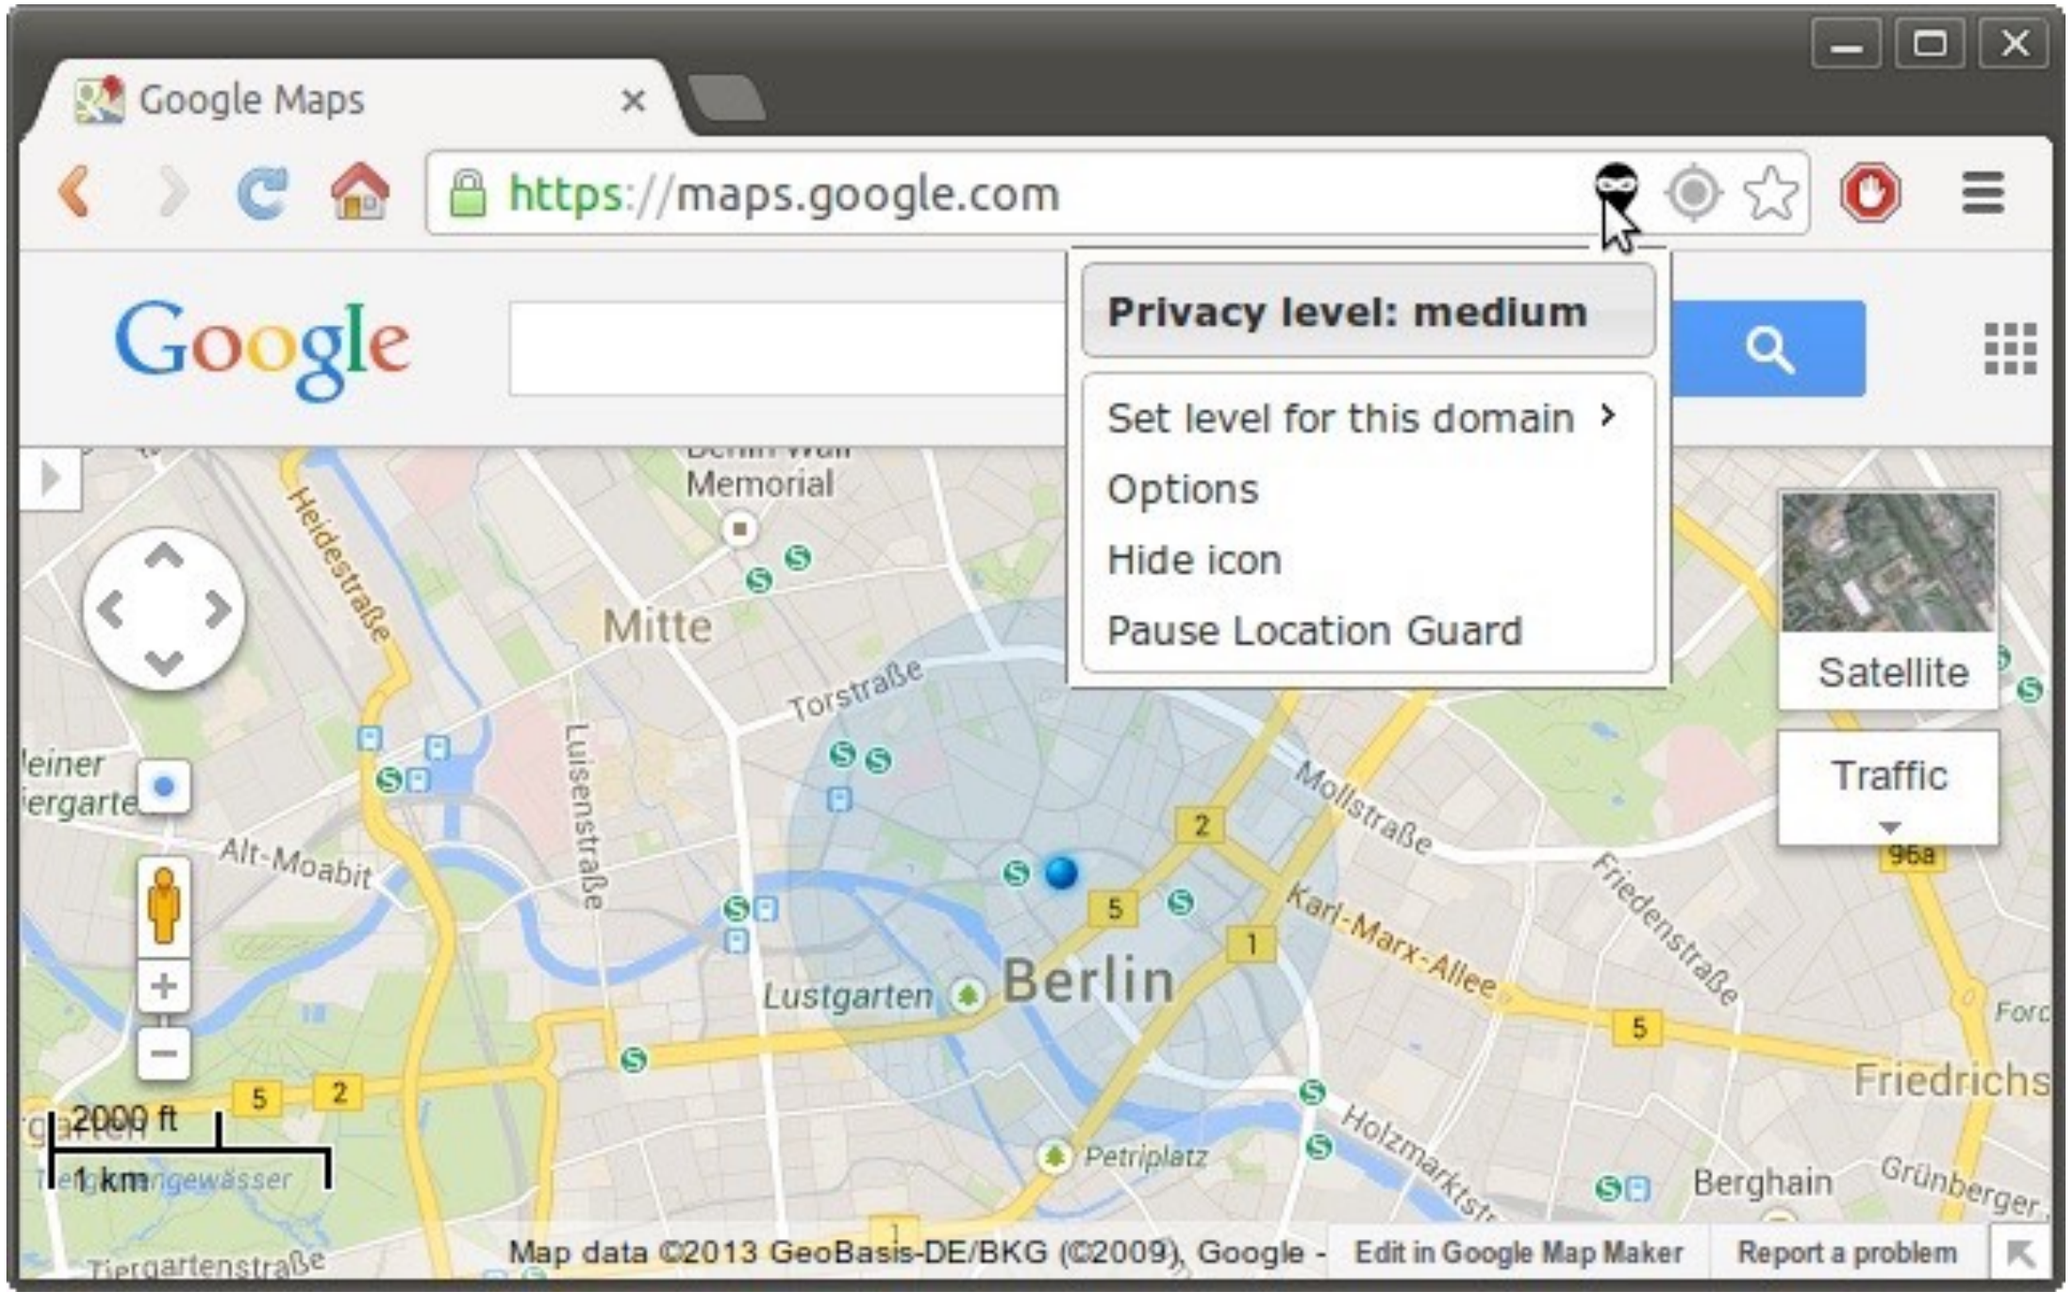

How it works

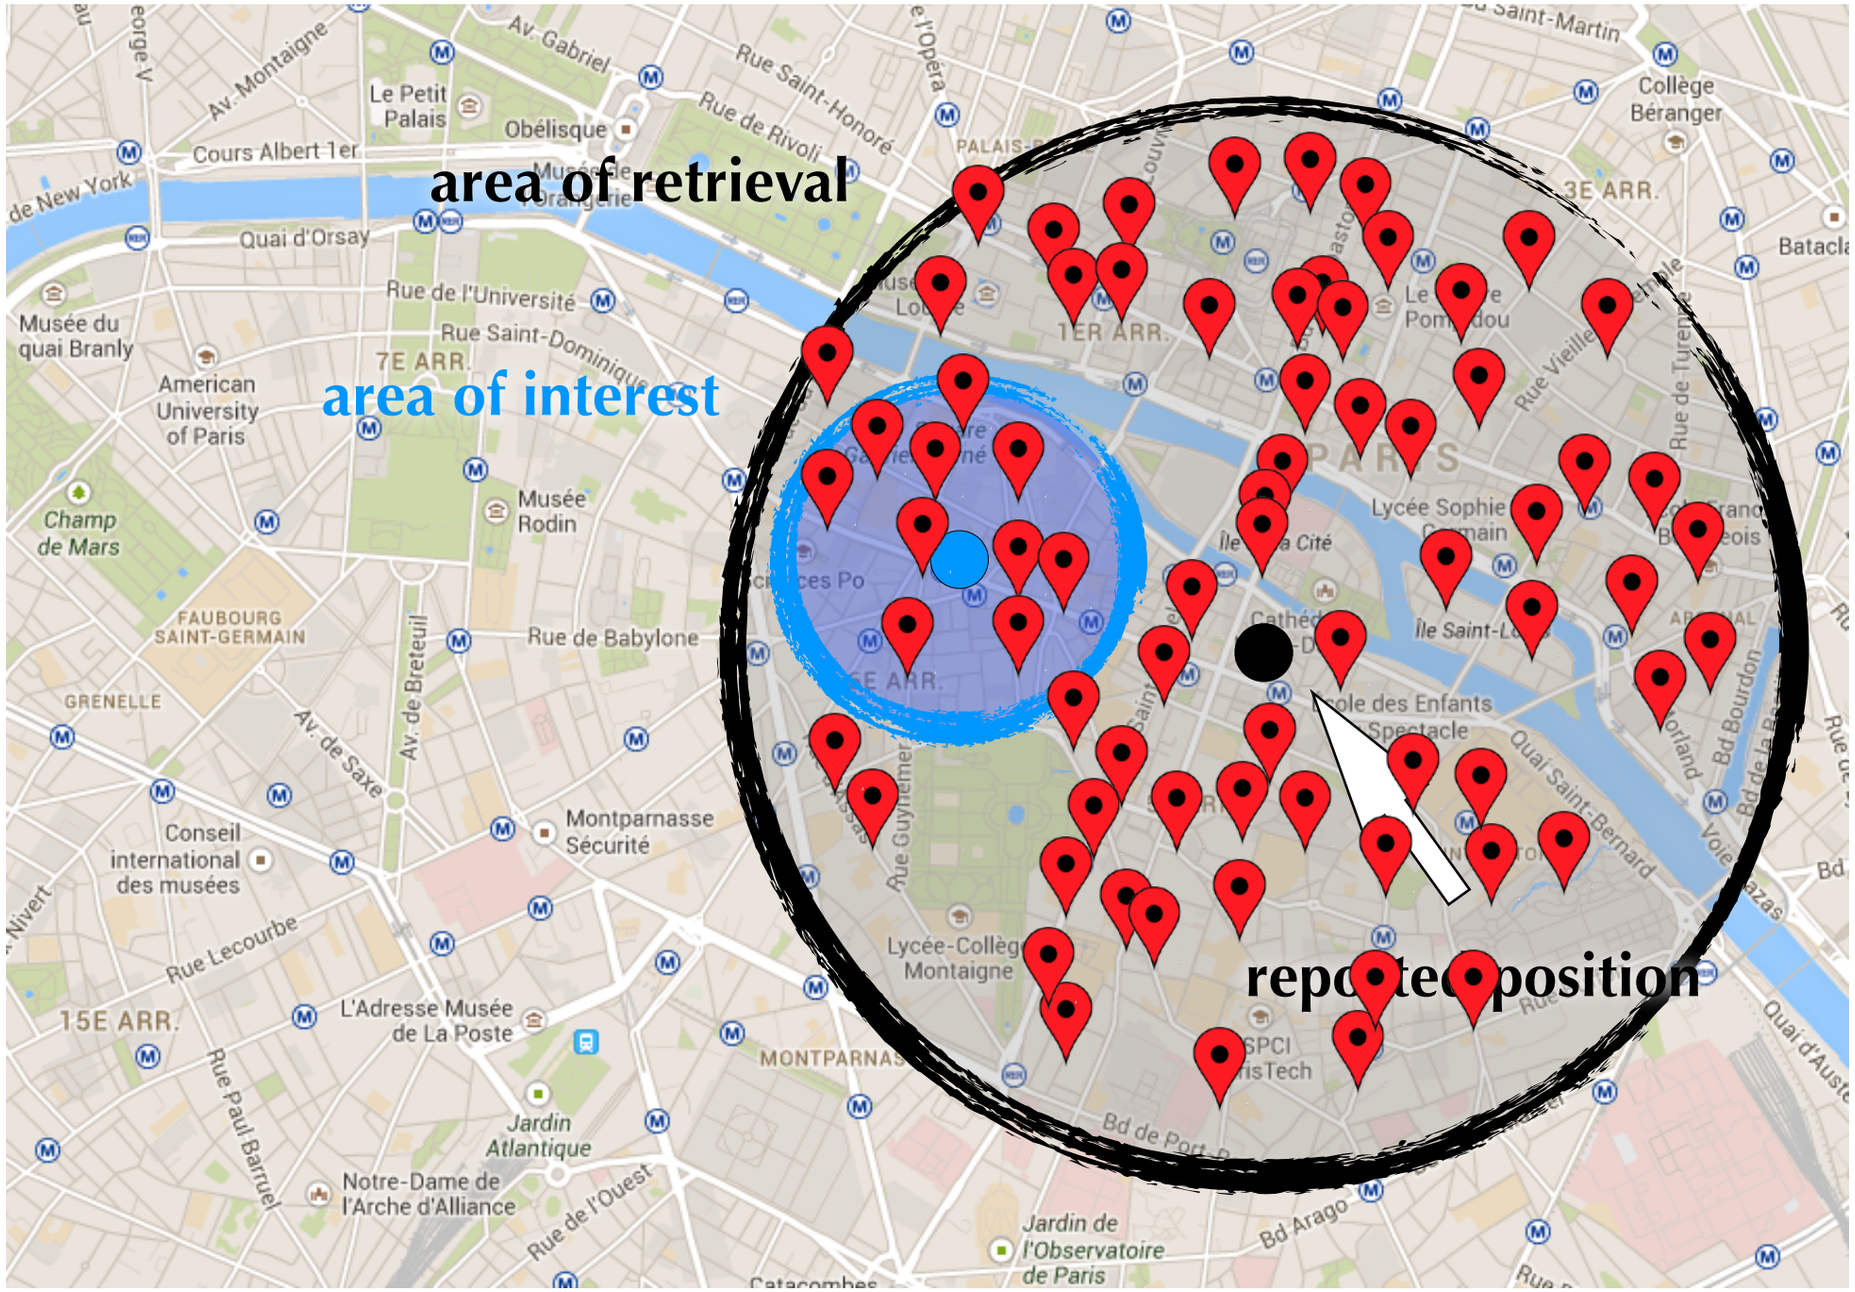
area of retrieval
area of interest
reported position
Av. Montaigne
Av. Gabriel
Rue Saint-Honoré
Rue de Rivoli
Le Petit Palais
Obélisque
Cours Albert 1er
Quai d'Orsay
Rue de l'Université
Rue Saint-Dominique
Musée du quai Branly
American University of Paris
7E ARR.
Musée Rodin
Champ de Mars
FAUBOURG SAINT-GERMAIN
Rue de Babylone
GRENELLE
Av. de Breteuil
Av. de Saxe
Montparnasse Sécurité
Conseil international des musées
Rue Lecourbe
15E ARR.
L'Adresse Musée de La Poste
Rue Paul Barruel
Jardin Atlantique
Notre-Dame de l'Arche d'Alliance
Rue de l'Ouest
MONTPARNAS
Jardin de l'Observatoire de Paris
Bd Arago
Lycée-Collège Montaigne
1ER ARR.
3E ARR.
Collège Béranger
Lycée Sophie Germain
Île de la Cité
Île Saint-Louis
Cathédrale
Rue Vieille du Temple
Rue de Turenne
Quai Saint-Bernard
Bd Bourdon
Rue Saint-Martin

Trade-off privacy-utility:

Utility as Quality of Service (QoS)

Trade-off privacy-QoS

Comparison with other methods for location privacy

Four mechanisms:

- Planar Laplacian
- Cloaking (report a region)
- Optimal by [Shokri et al. 2012] for uniform prior
- Optimal by [Shokri et al. 2012] for a given prior

(a)

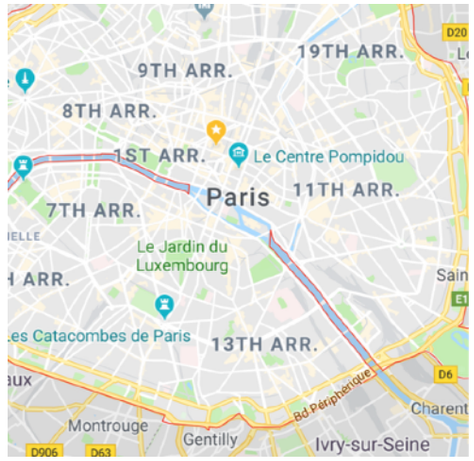

No need to compare with k-RR: it has a very bad QoS

(b)

Evaluation:

- Gowalla dataset, various towns, divided in a grid 10 x 10
- The levels of privacy are calibrated so that all methods offers the same level of privacy according to the definition of privacy of Shokri et al (Bayesian adversary)

(c)

Privacy versus QoS: evaluation

The four mechanisms:

- Cloaking,
- Optimal by [Shokri et al. CCS 2012] generated assuming uniform prior
- Planar Laplacian
- Optimal by [Shokri et al. CCS 2012] generated assuming the given prior

Based on linear optimization: high complexity

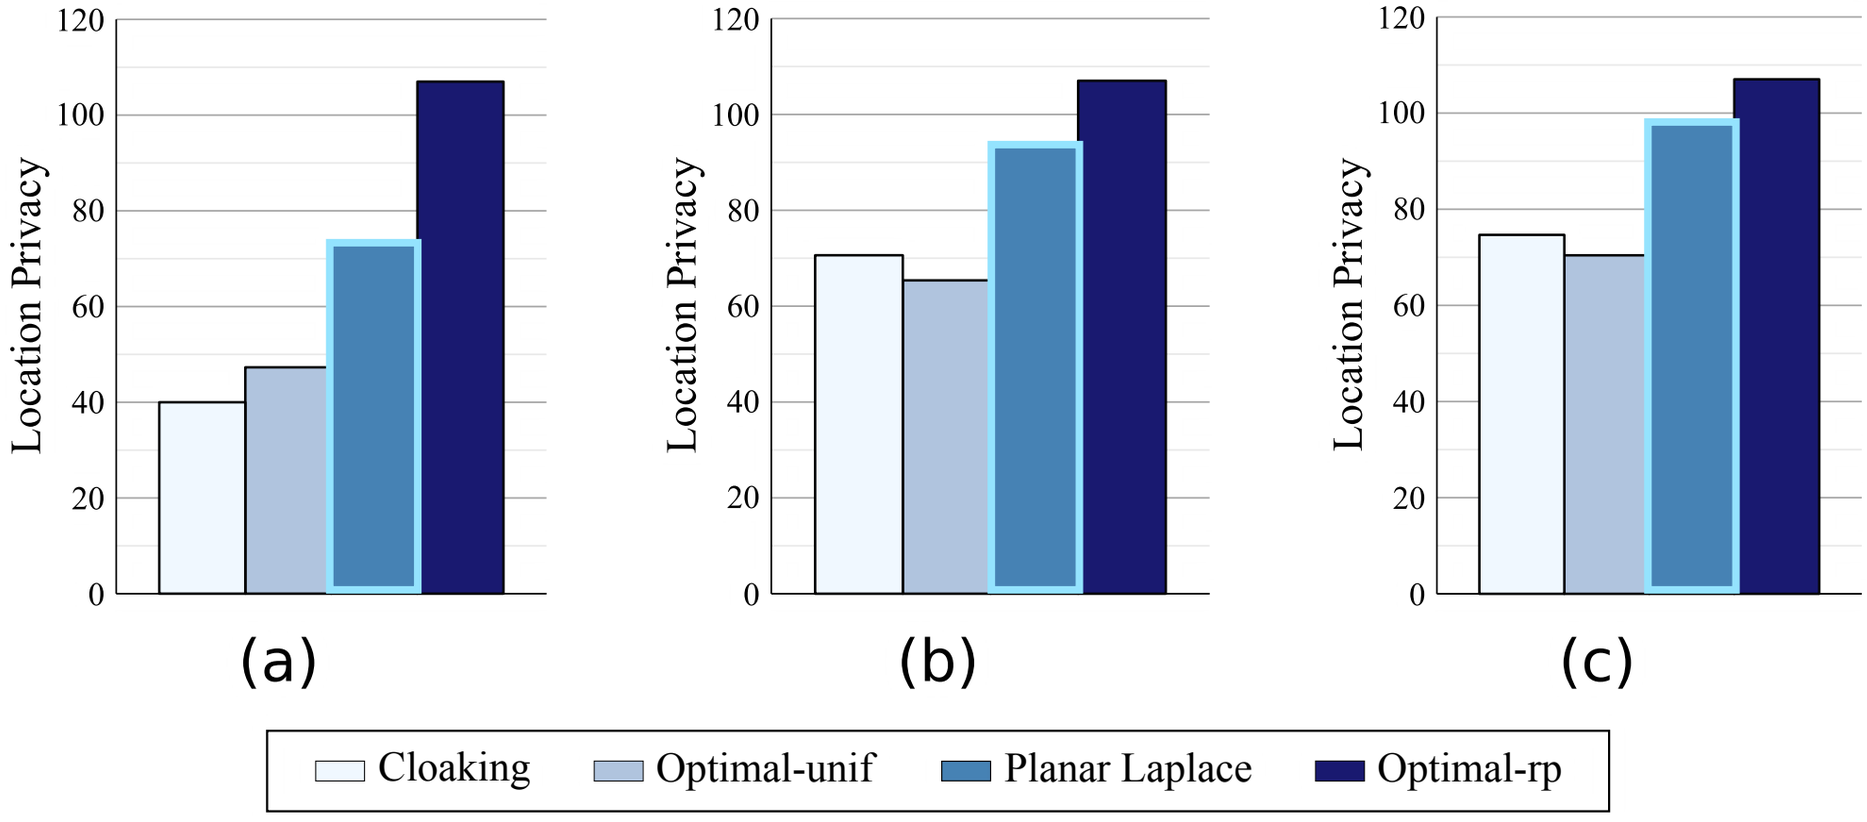

Trade-off privacy-utility:

Statistical Utility

Statistical utility: The problem

Consider again the Randomized Response mechanism

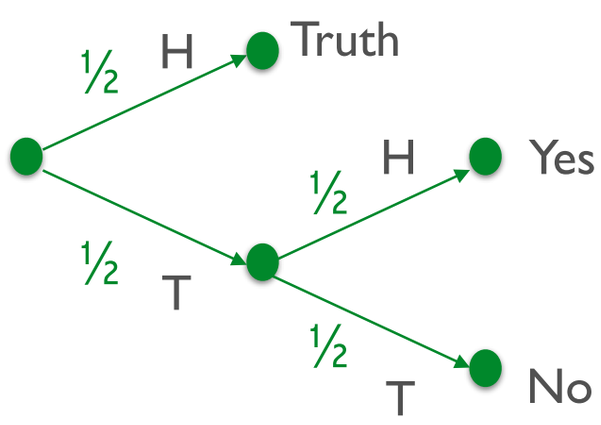

x \ y	yes	no
yes	¾	¼
no	¼	¾

Suppose that I get 60% answer "Yes" and 40% "no"

Do these figures represent the real percentages ?

Statistical utility: The matrix inversion method

[Kairouz et al, '16]

- Let C be the stochastic matrix associated to the mechanism
- Let q be the empirical distribution (derived from the noisy data reported by the mecahnism).
- Compute the approximation of the true distribution as $r = q\, C^{-1}$

Statistical utility: The matrix inversion method
[Kairouz et al, '16]

- Let C be the stochastic matrix associated to the mechanism
- Let q be the empirical distribution (derived from the noisy data reported by the mecahnism).
- Compute the approximation of the true distribution as $r = q\,C^{-1}$

Example: Randomized Response

Example Assume $q(\,Yes) = \frac{6}{10}$ and $q(No) = \frac{4}{10}$. Then:

$$\begin{aligned} \tfrac{3}{4}\,p(\,Yes) \;+\; \tfrac{1}{4}\,p(No) &= \tfrac{6}{10} \\ \tfrac{1}{4}\,p(\,Yes) \;+\; \tfrac{3}{4}\,p(No) &= \tfrac{4}{10} \end{aligned}$$

From which we derive $p(\,Yes) = \frac{7}{10}$ and $p(No) = \frac{3}{10}$

y

x	yes	no
yes	¾	¼
no	¼	¾

Statistical utility: The matrix inversion method

The mathix inversion method is simple and easy to analize.
Howerer it has two problems:

- Problem 1: C must be invertible, and not all mechanisms are
- Problem 2: The result may not be a distribution

Example: Consider again the Randomize Response and assume that $q(Yes) = \frac{4}{5}$ and $q(No) = \frac{1}{5}$. Then:

$$\begin{aligned} \tfrac{3}{4}\, p(Yes) &+ \tfrac{1}{4}\, p(No) &= \tfrac{4}{5} \\ \tfrac{1}{4}\, p(Yes) &+ \tfrac{3}{4}\, p(No) &= \tfrac{1}{5} \end{aligned}$$

y

	yes	no
yes	¾	¼
no	¼	¾

x

From which we derive $p(Yes) = \frac{11}{10}$ and $p(No) = -\frac{1}{10}$

Statistical utility: The matrix inversion method

$r = q\,C^{-1}$ may not be a distribution because it may contain negative elements. In order to try to obtain the true distribution **π** we can either:

- set to 0 all the negative elements, and renormalize (INV-N),

  or

- project r on the simplex (INV-P).

The resulting distribution however usually is not the best approximation of the original distribution.

Our approach: Iterative Bayesian Update

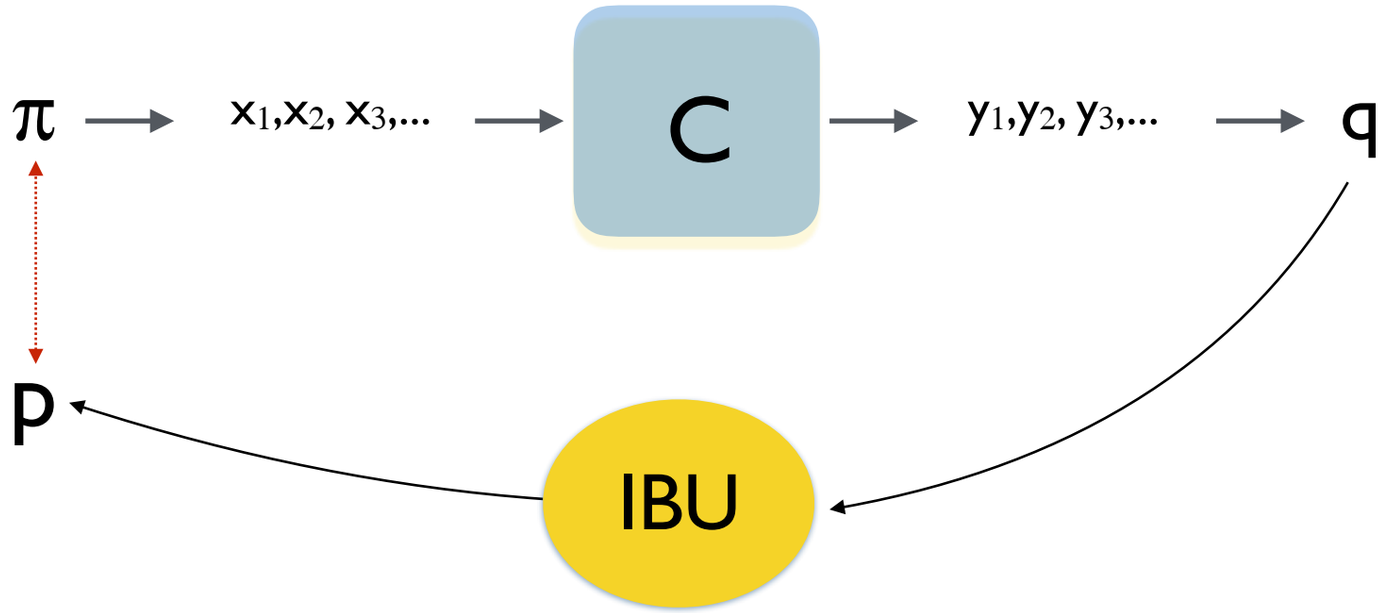

The IBU:

- is based on the **Maximization-Expectation** method
- produces a **Maximum Likelihood Estimator** p of the true distribution π
- If C is invertible, the MLE is unique and as the number of samples grows it converges to π

The Iterative Bayesian Update

- Define $p^{(0)}$ = any fully supported distribution (for example the uniform distribution)
- Repeat: Define $p^{(n+1)}$ as the Bayesian update of $p^{(n)}$ weighted on the corresponding element of q, namely:

$$p_x^{(n+1)} = \sum_y q_y \frac{p_x^{(n)} C_{xy}}{\sum_z p_z^{(n)} C_{zy}}$$

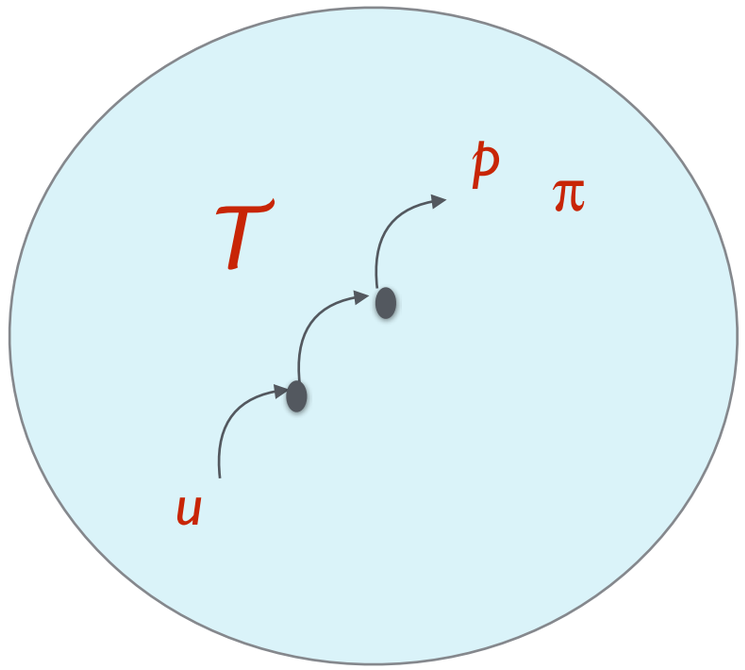

- Note that $p^{(n+1)} = T(p^{(n)})$
- When C is invertible, T has unique fix point (the MLE)
- Open problem: in some cases (with few samples) the MLE may not be the best estimation of the true distribution. We are trying to devise corrective methods.

Comparison between

the matrix inversion method

and

IBU

Comparison between Matrix Inversion and IBU:
Mechanism: Laplace $\epsilon = 0.1$ Data domain: {0,1,...,99}

Reconstruction Method \ Original Distribution	Gaussian	Uniform on an interval
INV-N	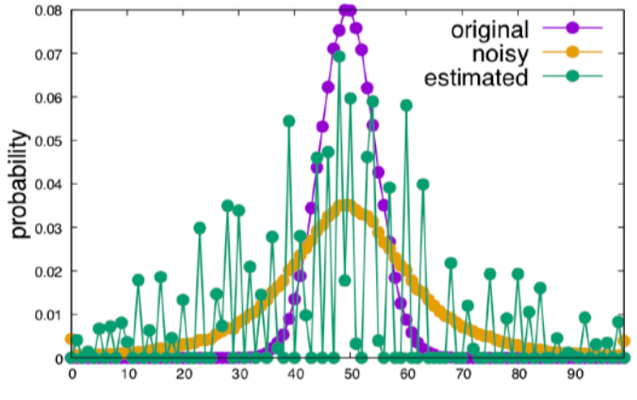	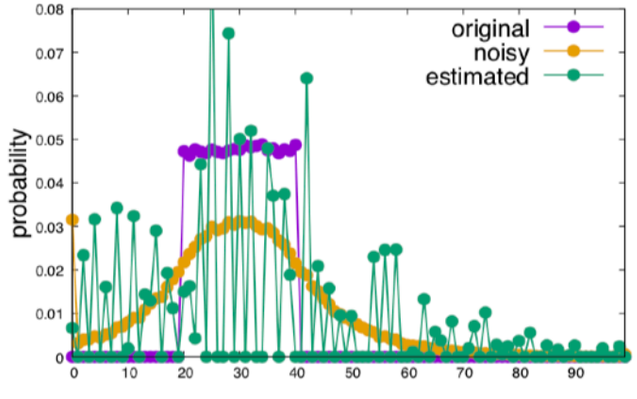
INV-P	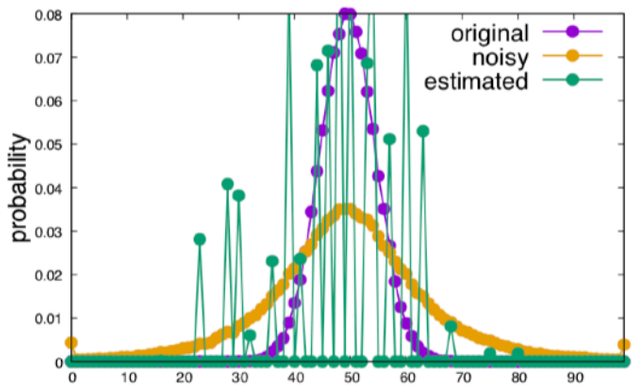	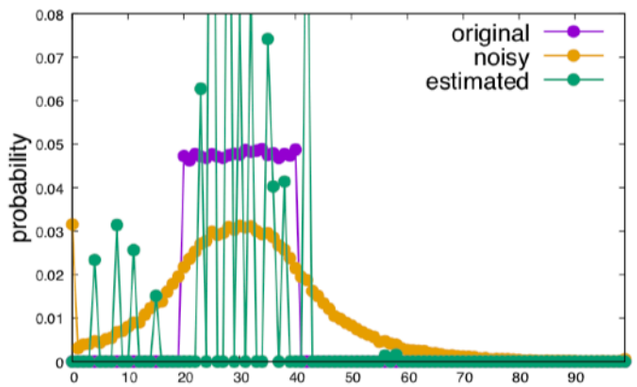
IBU	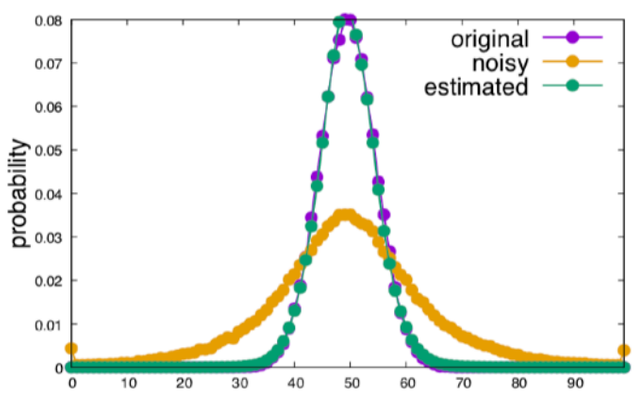	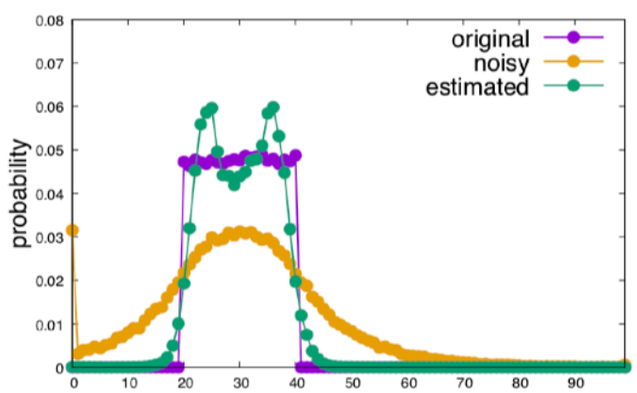

Comparison between Matrix Inversion and IBU:
Mechanism: Planar Laplace $\epsilon = 1$
Data domain: Gowalla Location Data in S. Francisco

planar Laplace noise

San Francisco area

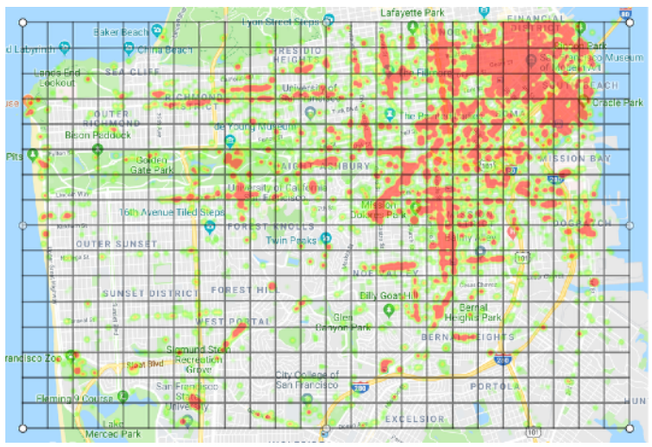

Original

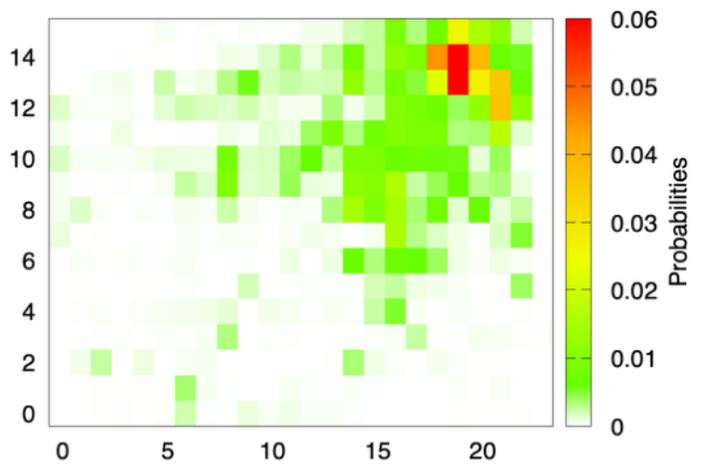

Empirical from noisy data

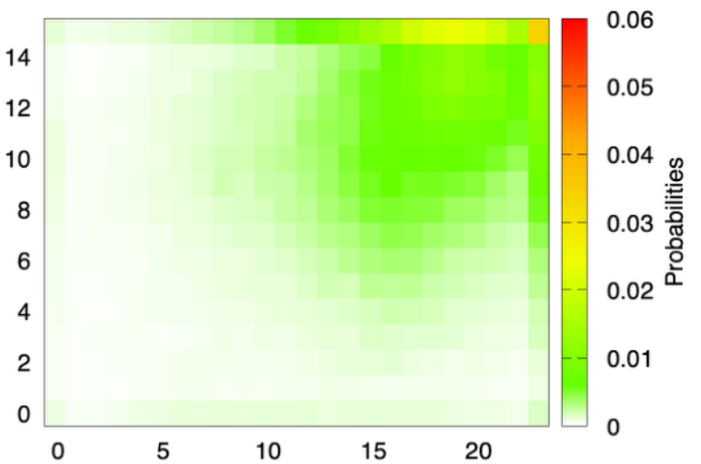

IBU

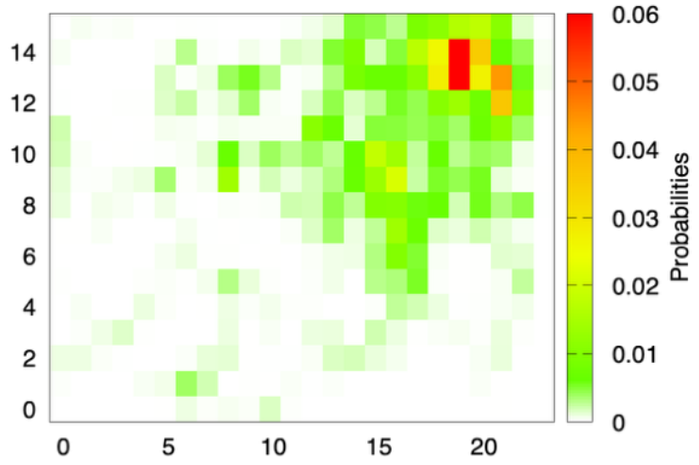

INV-N

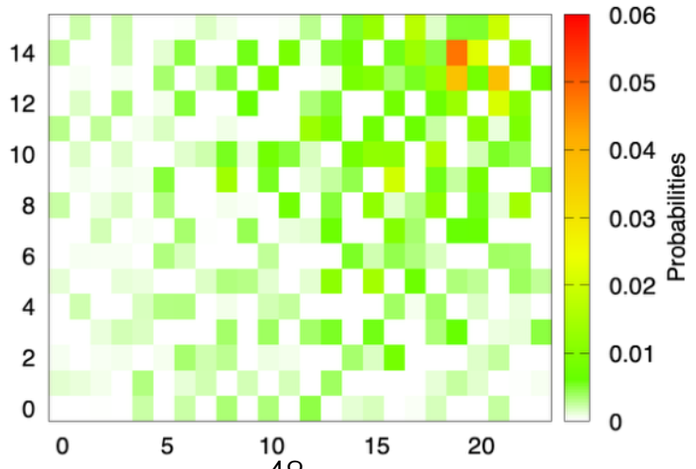

INV-P

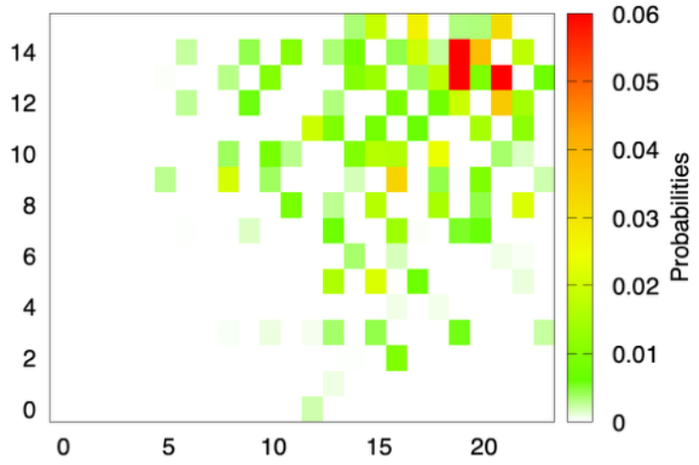

Comparison wrt statistical utility (using IBU) between different definitions of privacy and their typical obfuscation mechanisms

LPD / k-RR

and

***d*-privacy / Laplace**

Trade-off between utility and statistical privacy
Comparison between Laplace and k-RR

Problem: Both K-RR and the geometric / laplace mechanisms are parametrized by ε, but it has a different meaning.

Therefore, in order to make a fair comparison, we need to calibrate ε, in such a way that the requested ratio is satisfied in the "area of interest" (area in which we want to be indistinguishable)

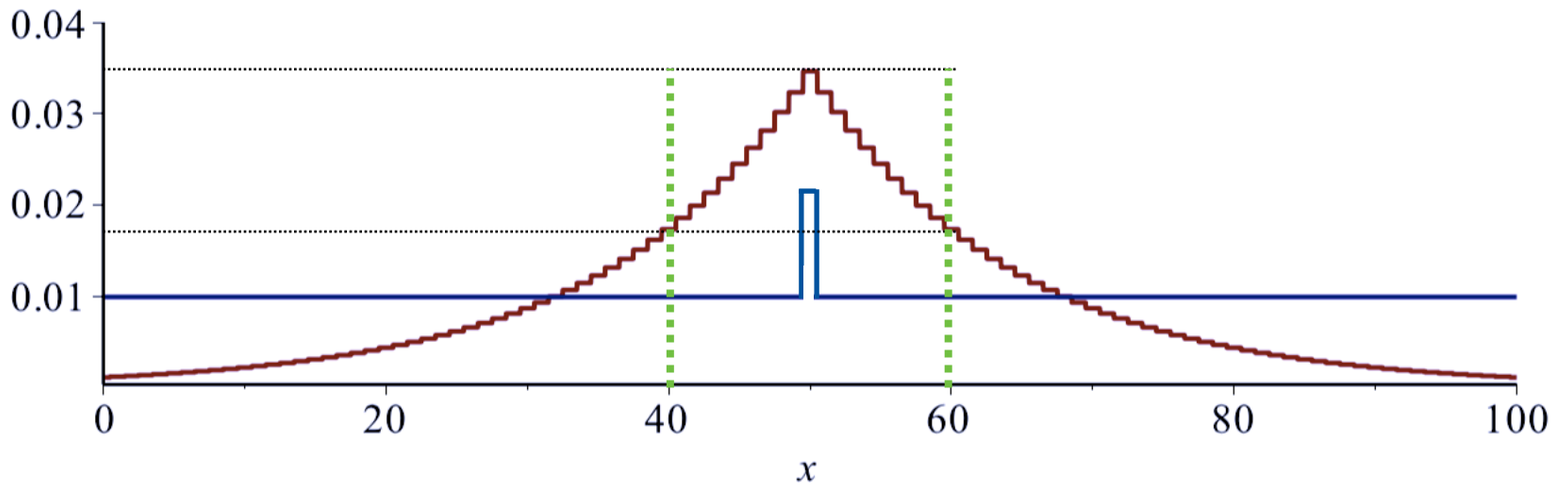

Comparison between LPD and *d*-privacy
Experiments on the Gowalla dataset

- Gowalla is a dataset of geographical checkins in several cities in the world
- We compare the statistical utility of kRR and Planar Laplace with the respective ε calibrated so to provide the same level of privacy within about 1 Km^2

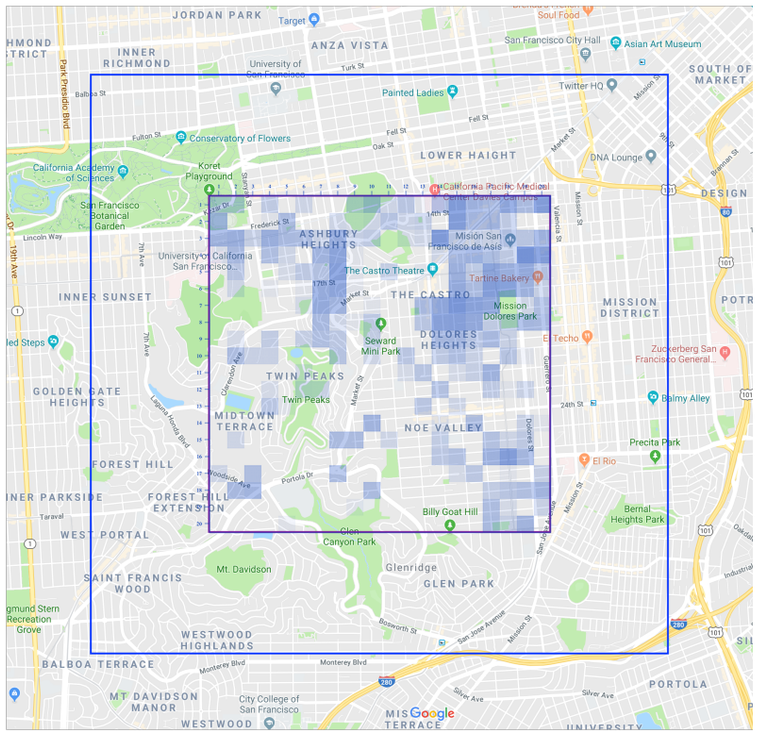

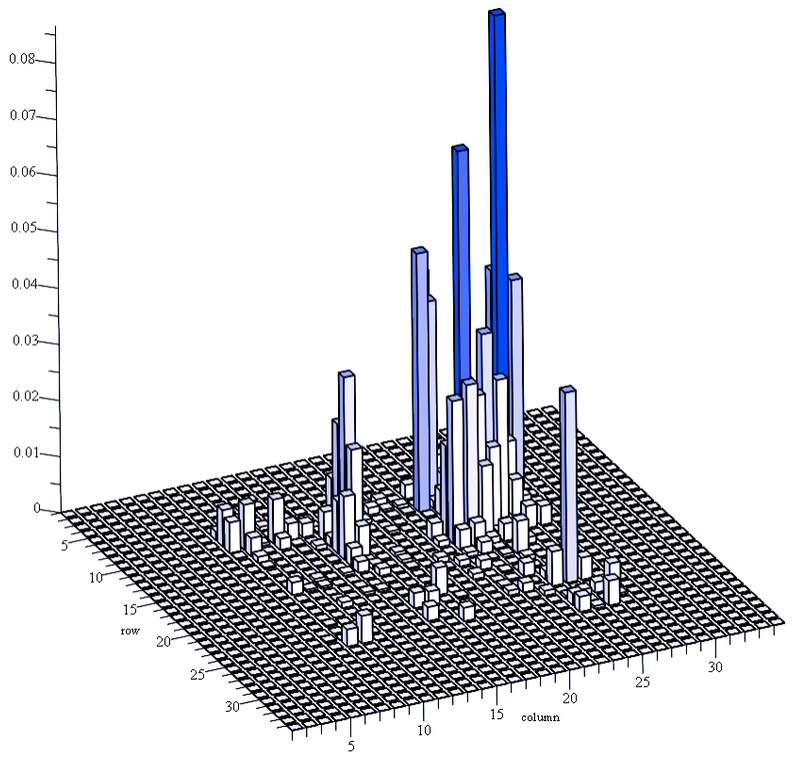

Gowalla checkins in an area of 3x3 km^2 in San Francisco downtown (about 10K checkins)

The Planar Laplace mechanism

$\varepsilon = \ln(2)$

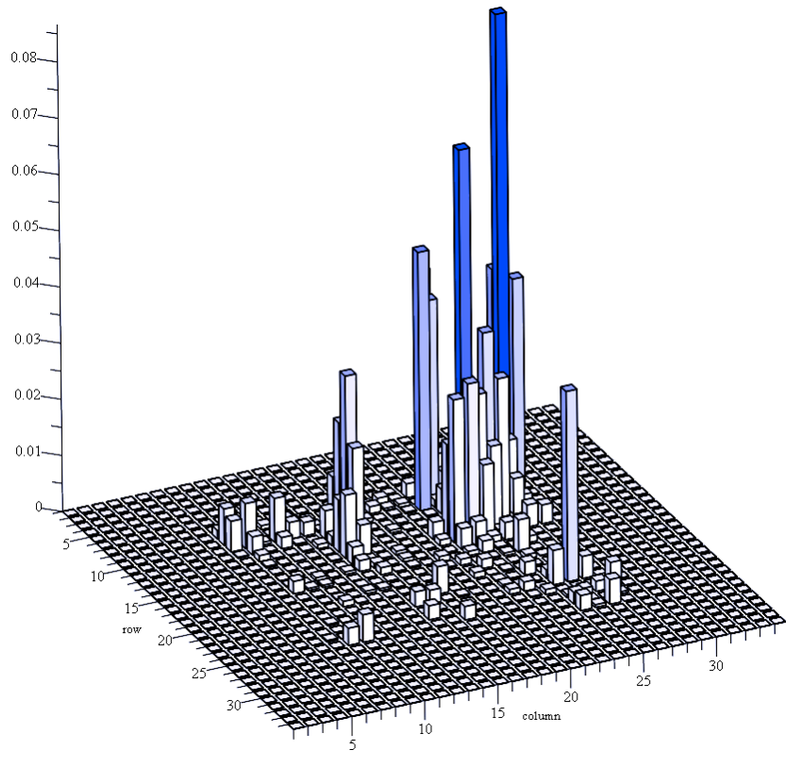

The real distribution

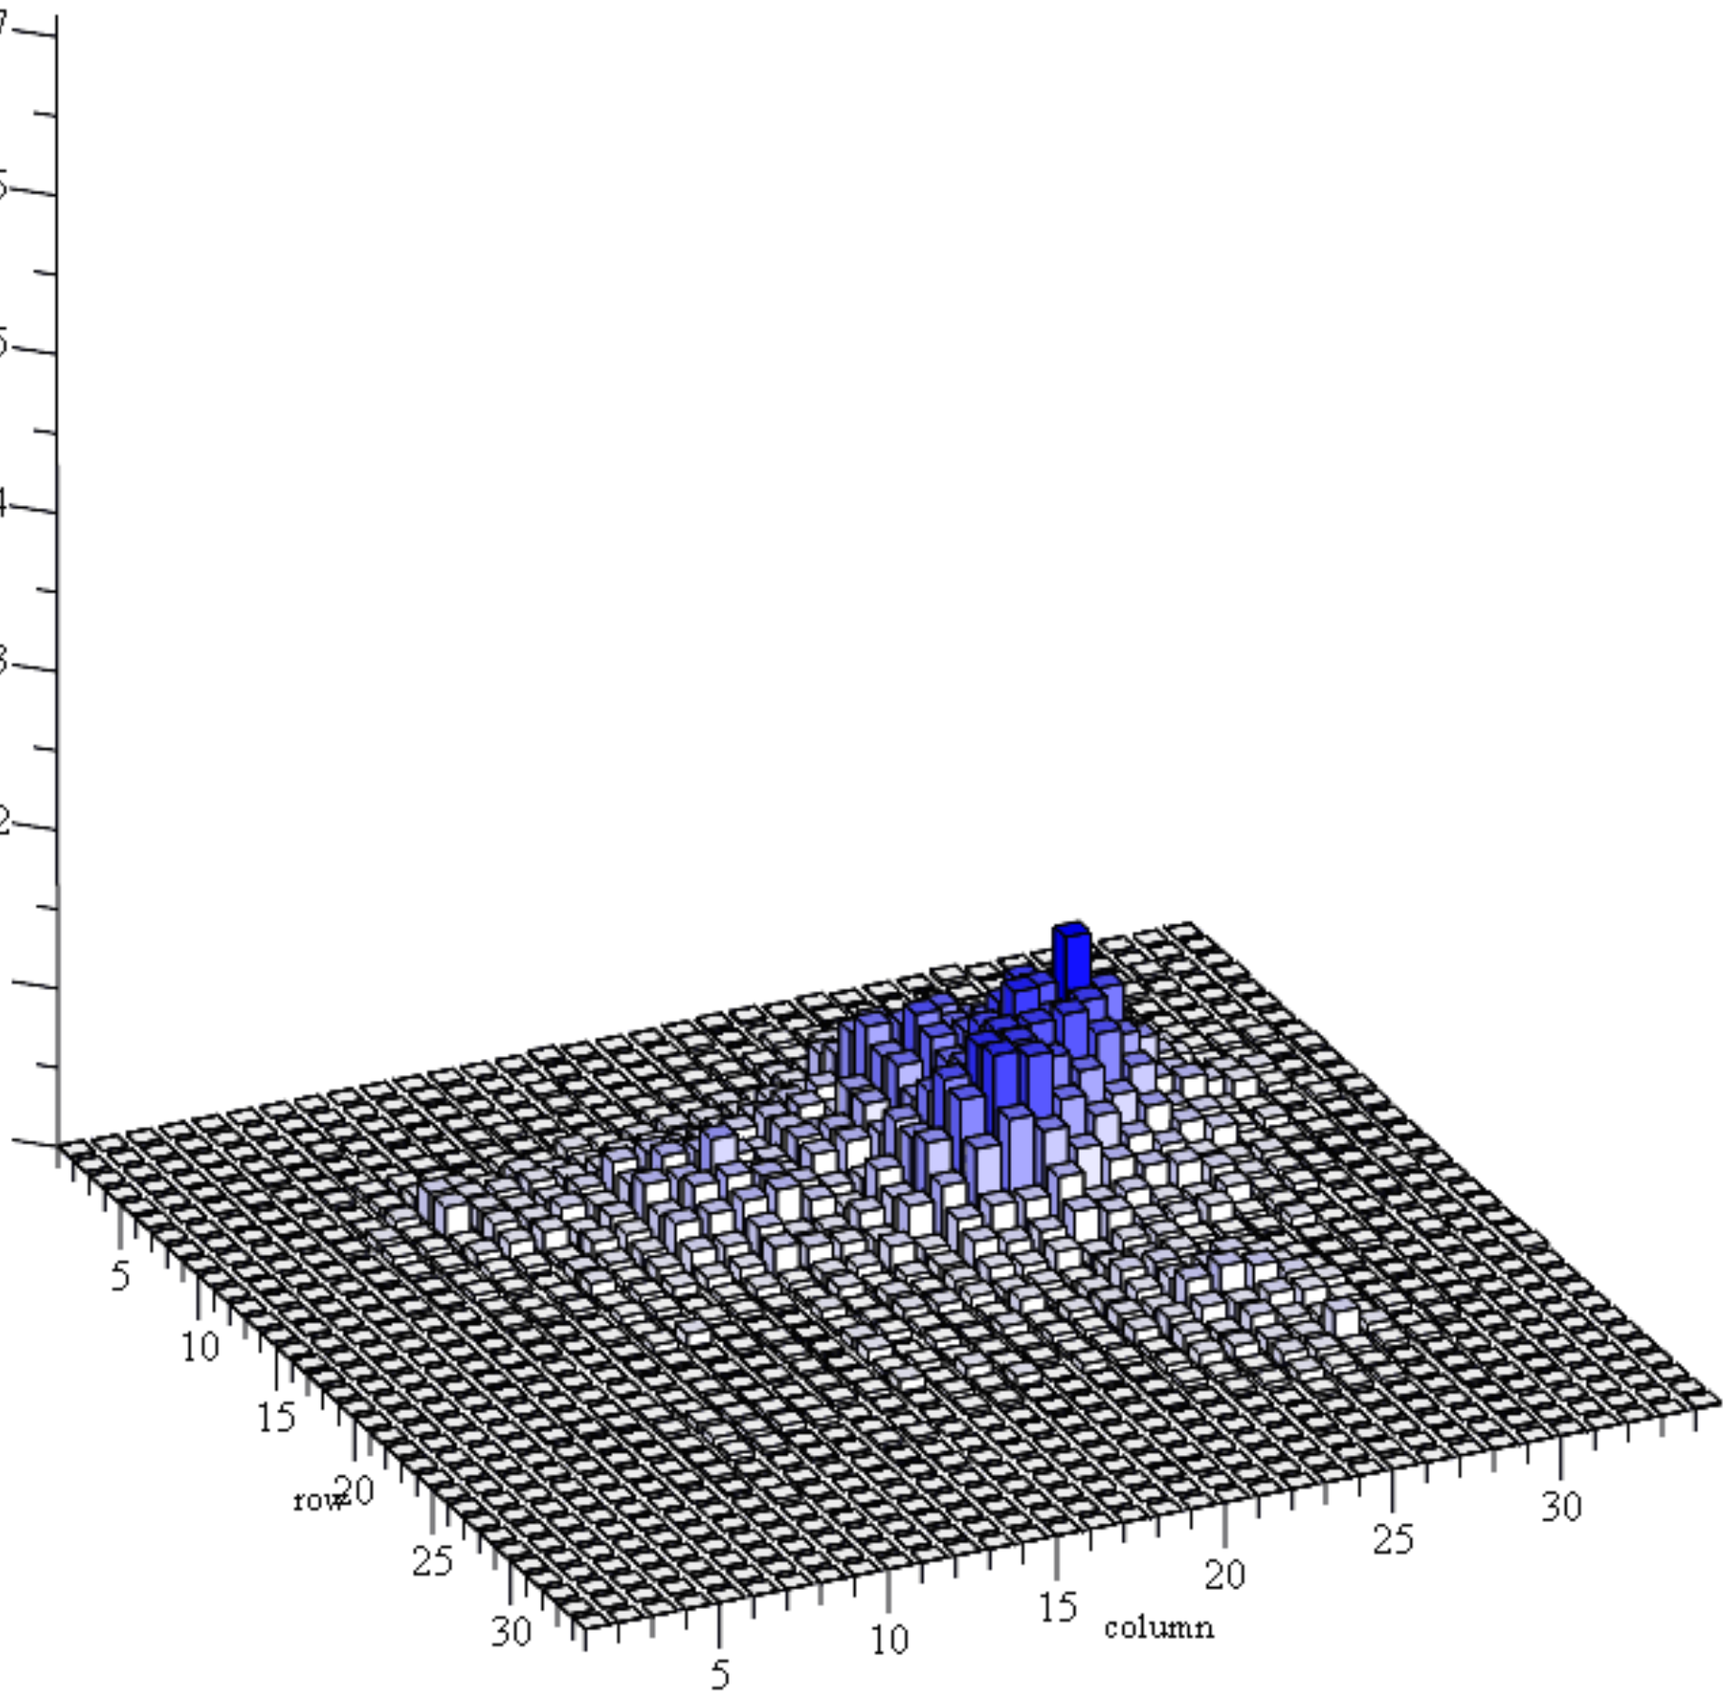

The noisy distribution and the result of the IBU (300 iterations)

$n = 0.$

The kRR mechanism

$\varepsilon = \ln(8)$

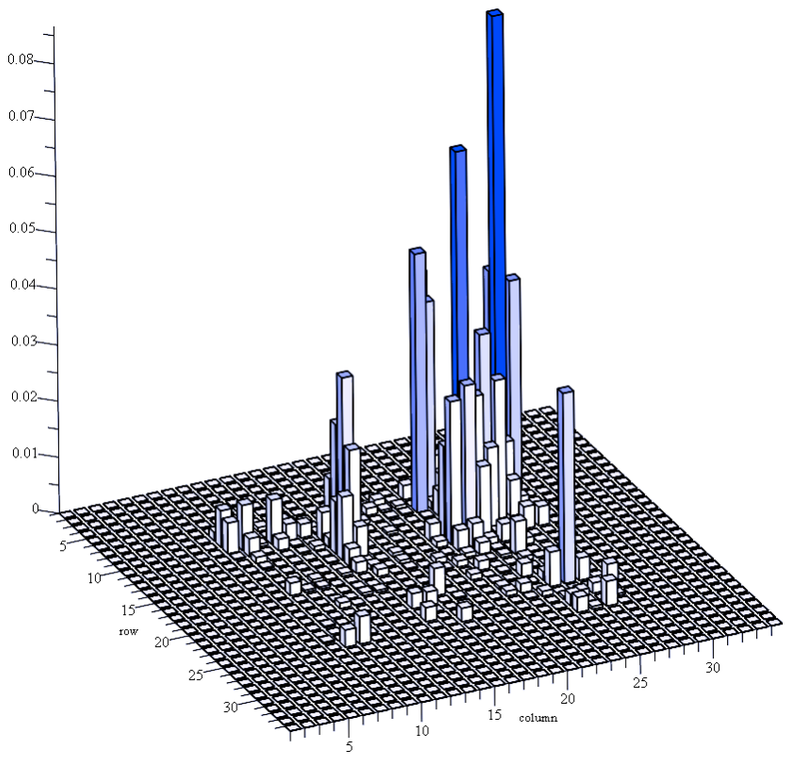

The real distribution

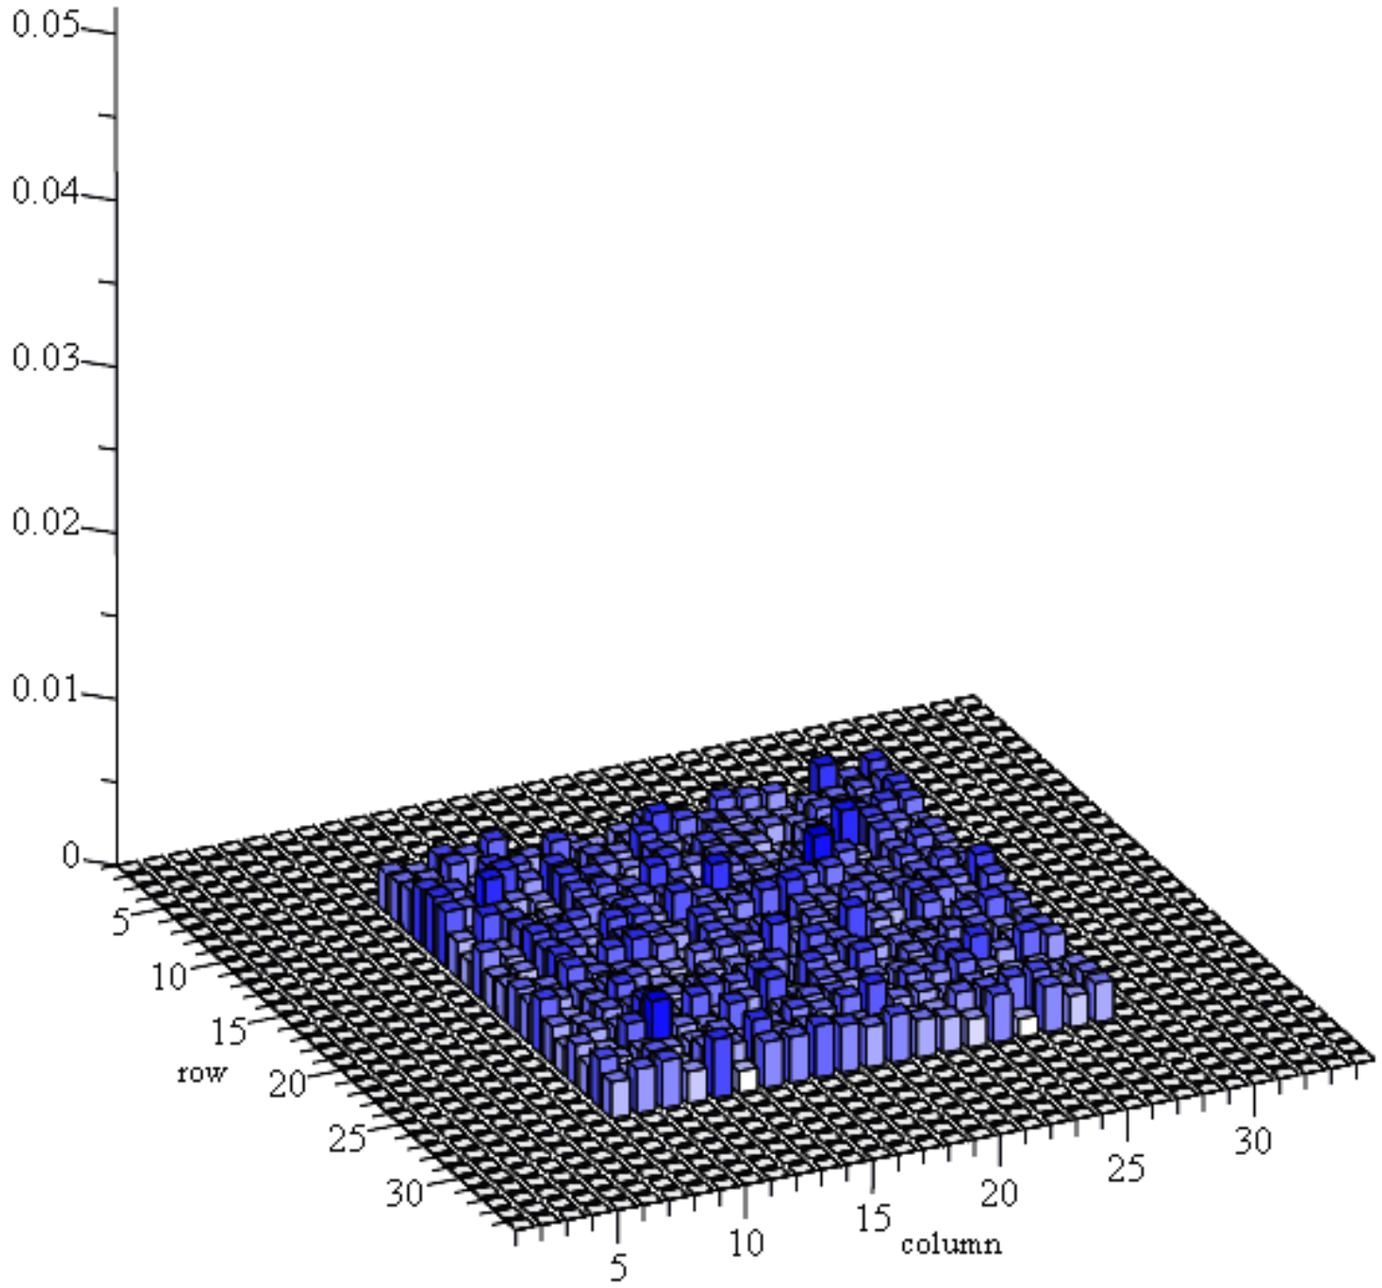

The noisy distribution and the result of the IBU (500 iterations)

Measuring the quality of the approximation

There are many measures of distance between distributions.
A typical metric is the total variation distance. If we are interested in statistic related to the ground distance, however, a more appropriate metric is the **Kantorovich distance** (aka Earth Movers distance).

- The Total Variation distance measures only the area between the two probability distributions

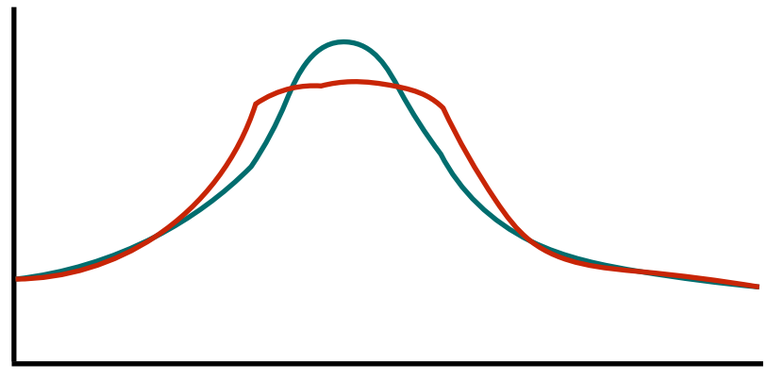

- The Kantorovich takes into account also the ground distance; it measures the "transportation effort" to make the two distributions equal. Cfr. "Earth moving distance"

- In these two examples the TV is the same, while the Kantorovich is larger in the second case

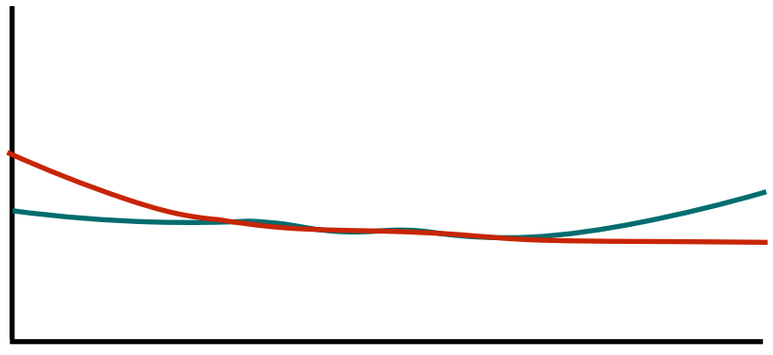

- The Kantorovich metric is particularly suitable when we are interested in statistics that are sensitive to the underlying distance.
Example: placement of hotspots.

$$K_d(\mu,\nu) = \sup_{f \in Lip} \left| \sum_x \mu_x f(x) - \sum_x \nu_x f(x) \right|$$

where Lip is the set of Lipshitz functions wrt d

Evaluation: San Francisco

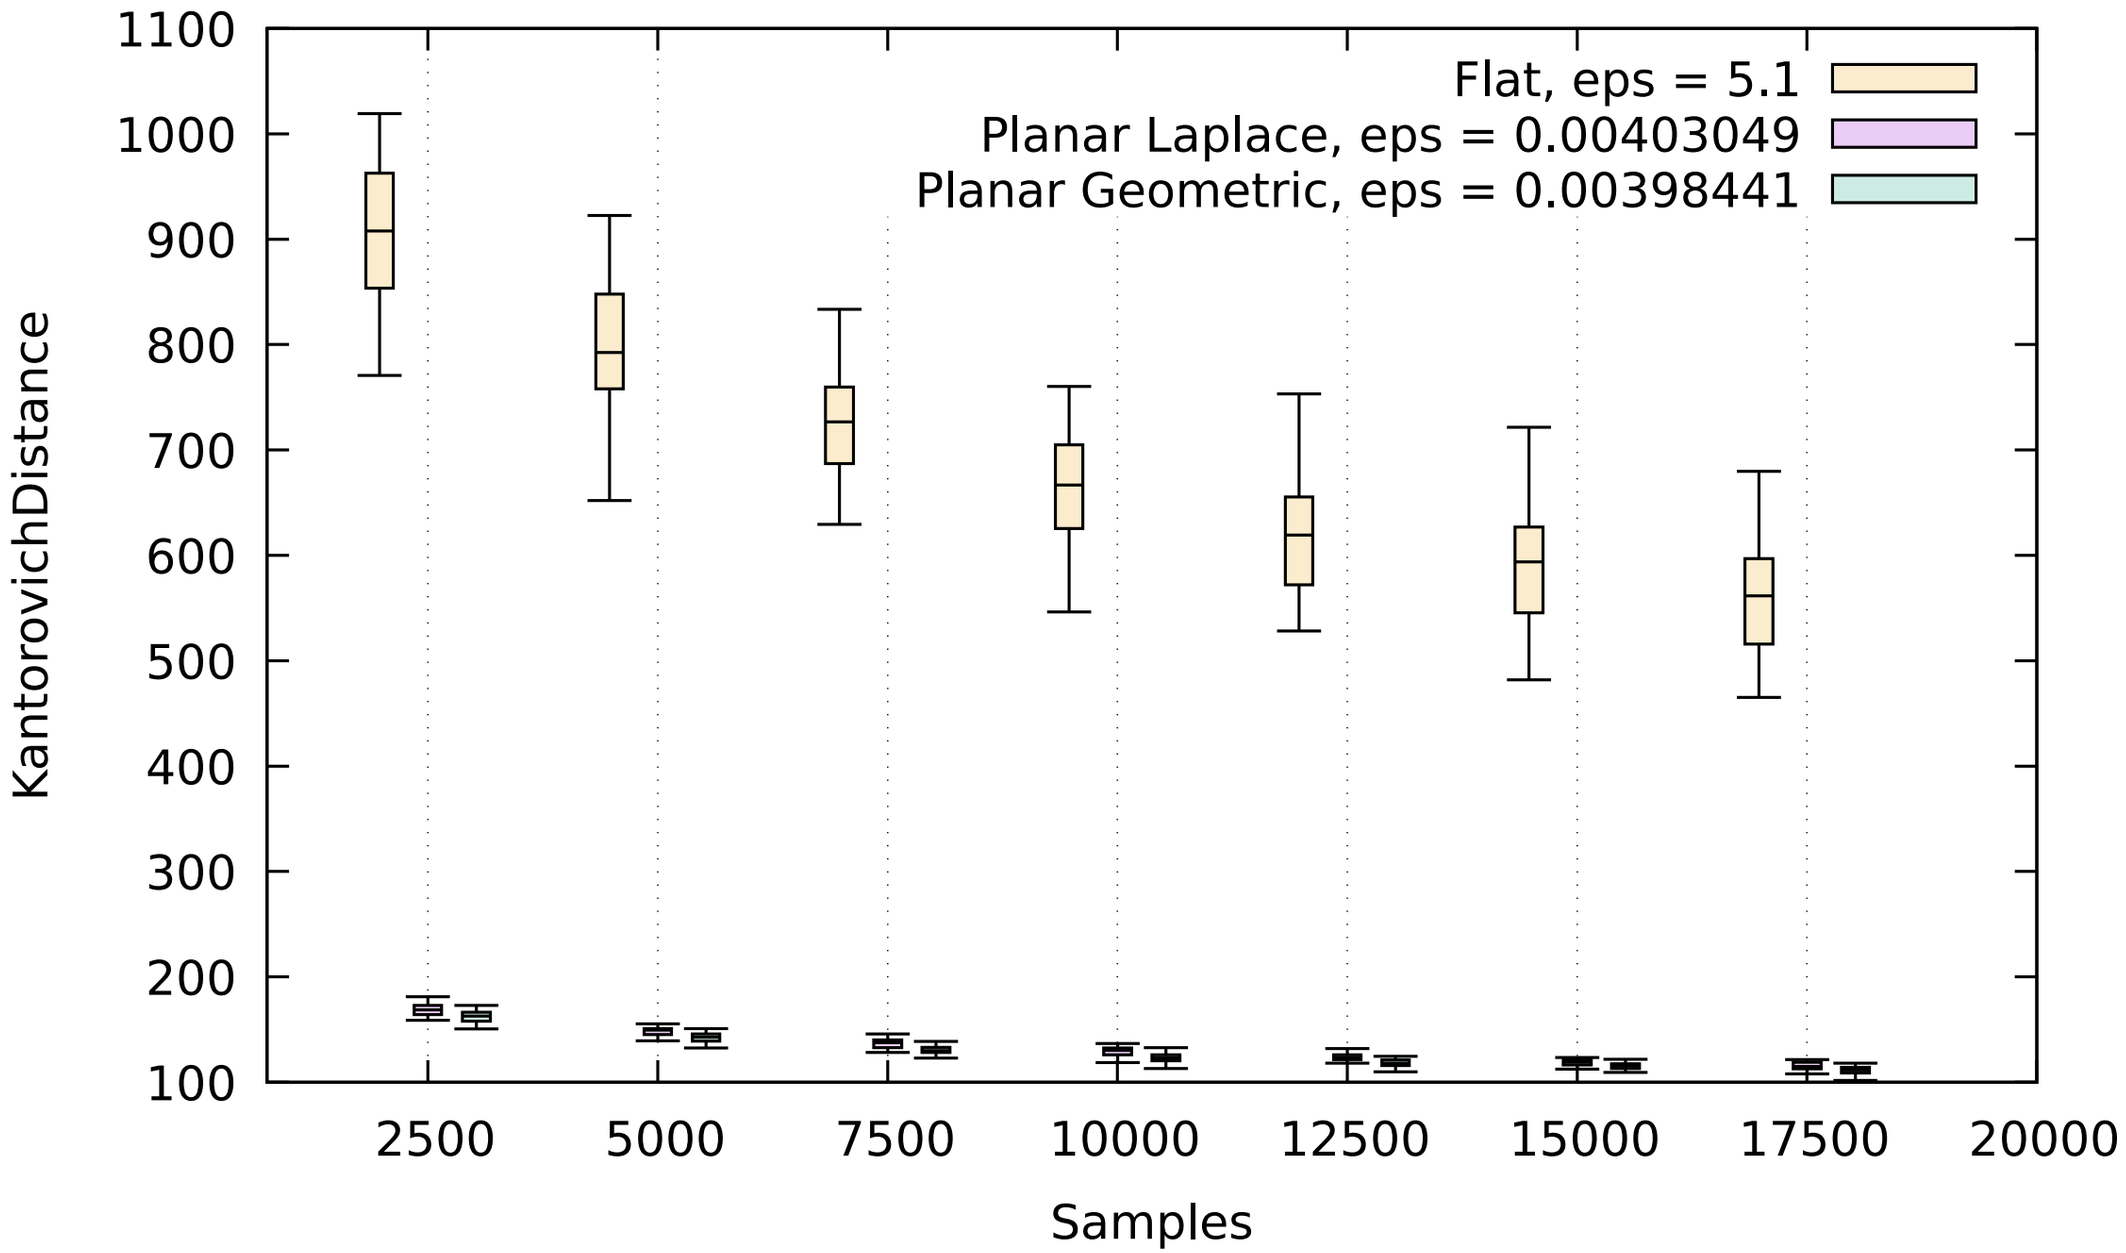

Evaluation: Paris