Local Differential Privacy

Giulia Fanti

Slides based in part on material by Ananth Raghunathan

Fall 2019

Administrative

- HW3 out
  - Differential privacy and deanonymization
- Recitation on Friday
  - Local differential privacy (Sruti)
- Interesting talks
  - Today @ 5.30 Posner160, "Facebook Data Privacy. + Design"
  - Thursday 10/10 @ noon, **Hamburg Hall 1002,** "Next Generation Privacy Reviews", Dhanuja Shaji, SNAP
- Project budget
  - If you need money for your project (e.g. for datasets) send me an email with the amount you need and link to purchase

Canvas quiz

- 10 minutes

Different models

- Global (database) differential privacy

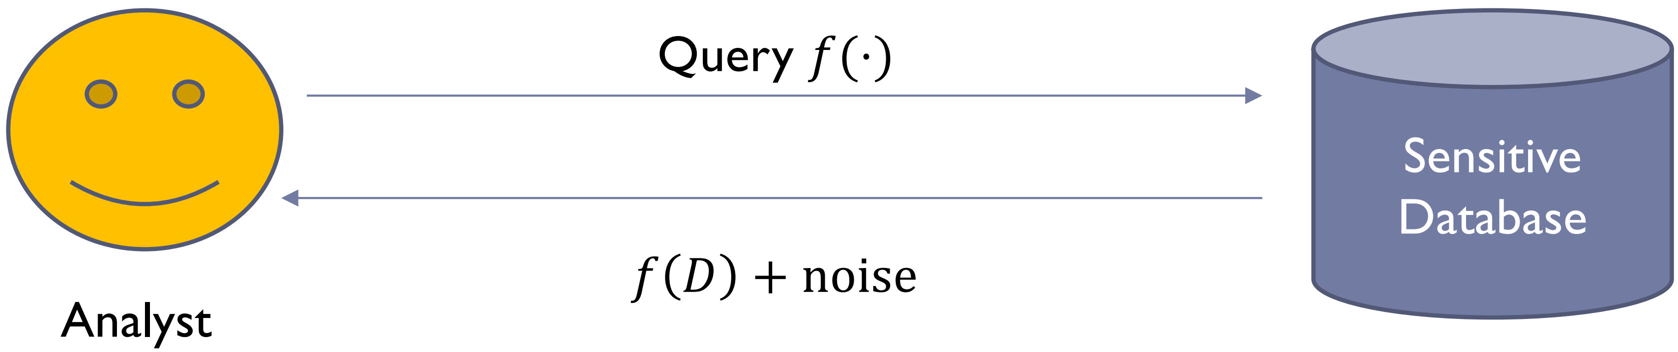

- Local differential privacy

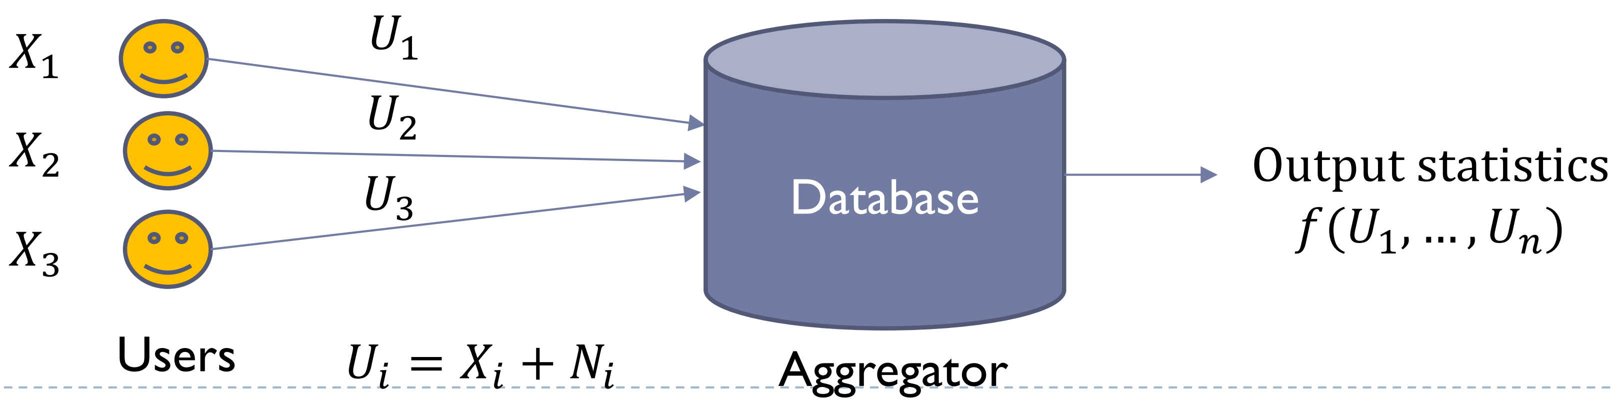

Local Differential Privacy

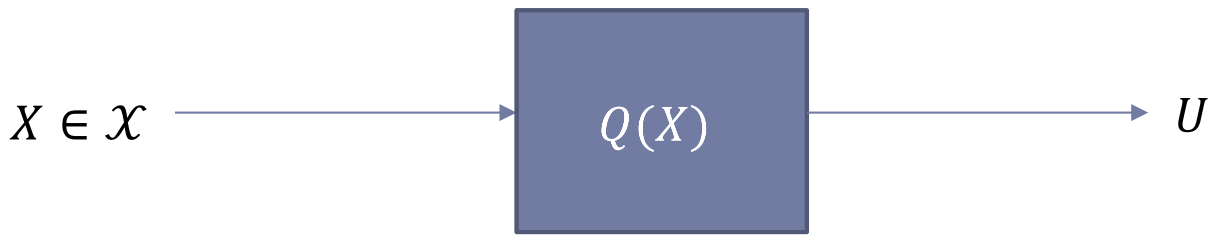

- We say mechanism Q is ϵ-locally differentially private if

$$\sup_{S,x,x' \in \mathcal{X}} \frac{Q(S|X = x)}{Q(S|X = x')} \leq e^{\epsilon},$$

Randomized Response

- **"Are you now, or have you ever been, a member of the communist party?"**

- Flip a coin, in private
- If the coin comes up heads, respond "Yes"
- Otherwise, tell the truth

- Estimate true "yes" ratio with

  # of "Yes" responses – 0.5

Real-World Application: RAPPOR

- Google wanted to detect hijacking of browser settings
  - Measure proportion of homepages
  - … without collecting everyone's data in plaintext

- RAPPOR
  - First internet-scale deployment of differential privacy
  - Open-source

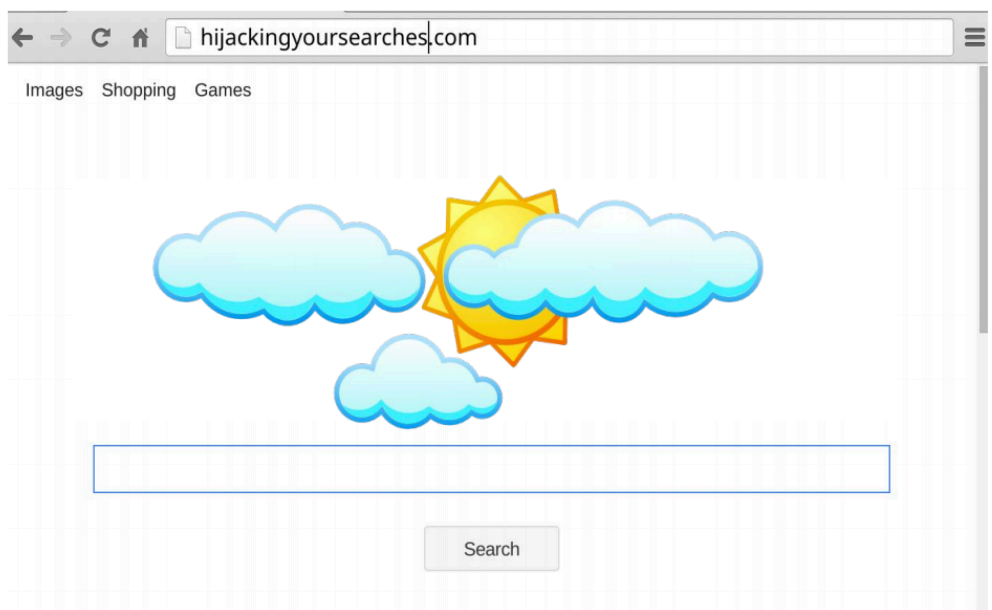

Traditional best practices

- Collect user data
- Scrub IP addresses, timestamps, etc.

- Keep central database of scrubbed data (e.g., 2 weeks)
  - Keep only aggregates of older data

- Report aggregates of data over threshold (e.g., 10 users)

- Can be the best approach for opt-in, low-sensitivity data

RAPPOR

- Learn statistics with differential privacy

RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response

Úlfar Erlingsson
Google, Inc.
ulfar@google.com

Vasyl Pihur
Google, Inc.
vpihur@google.com

Aleksandra Korolova
University of Southern California
korolova@usc.edu

- Pros:
  - Strong privacy guarantees
  - Robust to hackers, subpoenas, etc.
- Cons:
  - How do you collect string-valued data with LDP?

Bloom Filters

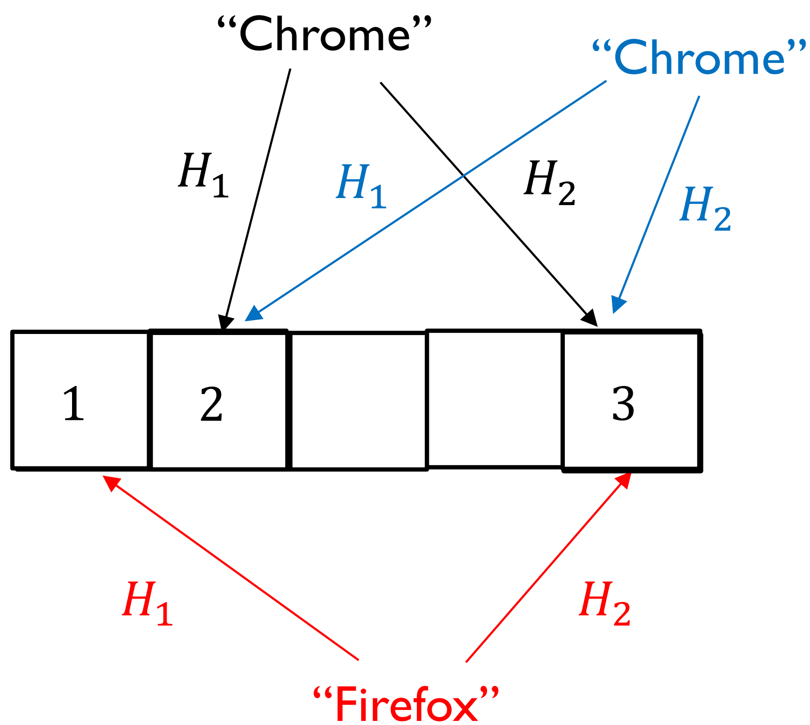

We use k hash functions.

Here $k = 2$

Let's add differential privacy

- User side: Randomized response

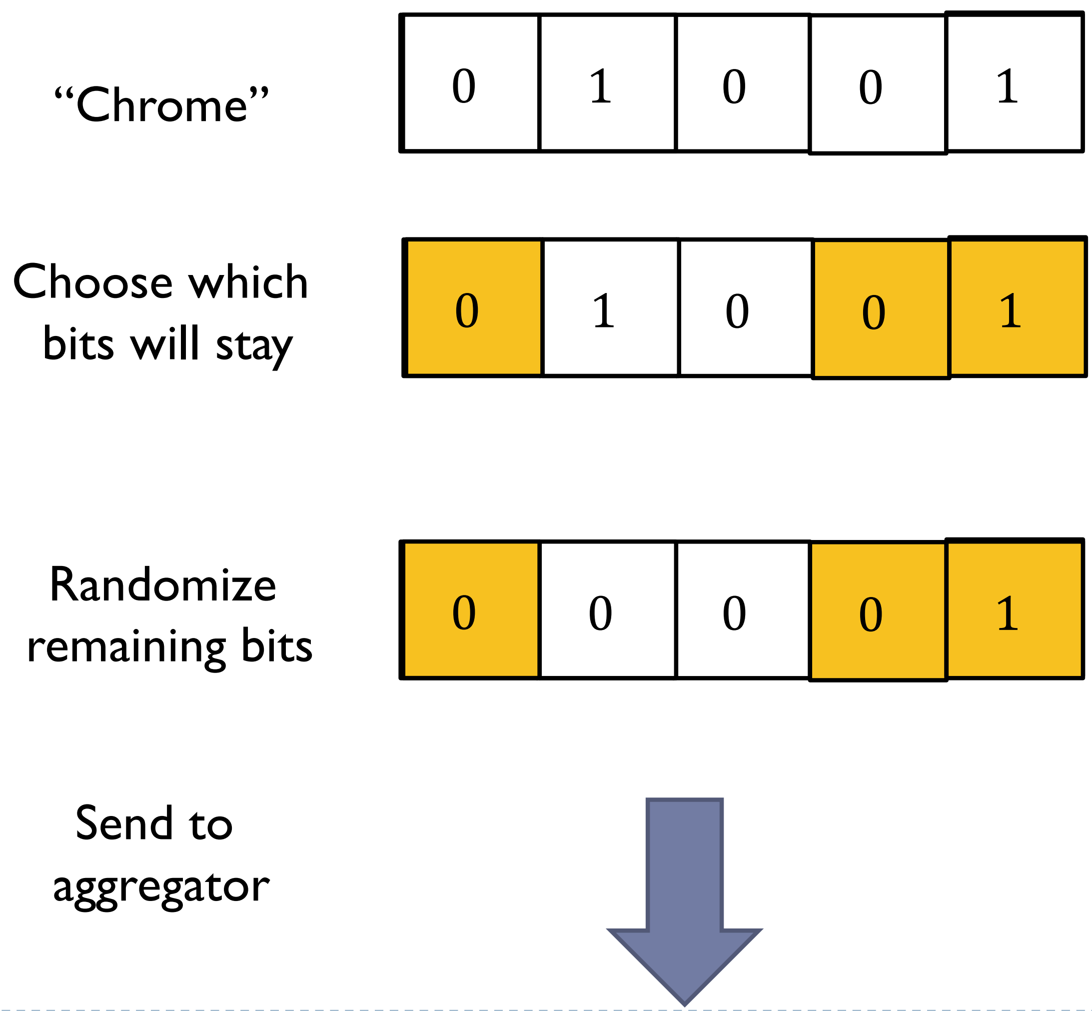

w.p. $1 - f$, report true bit
w.p. f, report random bit

e.g., let $f = \frac{1}{2}$

Let's add differential privacy

<u>Mechanism</u>
w.p. $1 - f$, report true bit
w.p. f, report random bit

- What privacy guarantee does this give you?

$$\epsilon = 2\ln\left(\frac{\left(1 - \frac{f}{2}\right)}{\frac{f}{2}}\right)$$

Aggregator

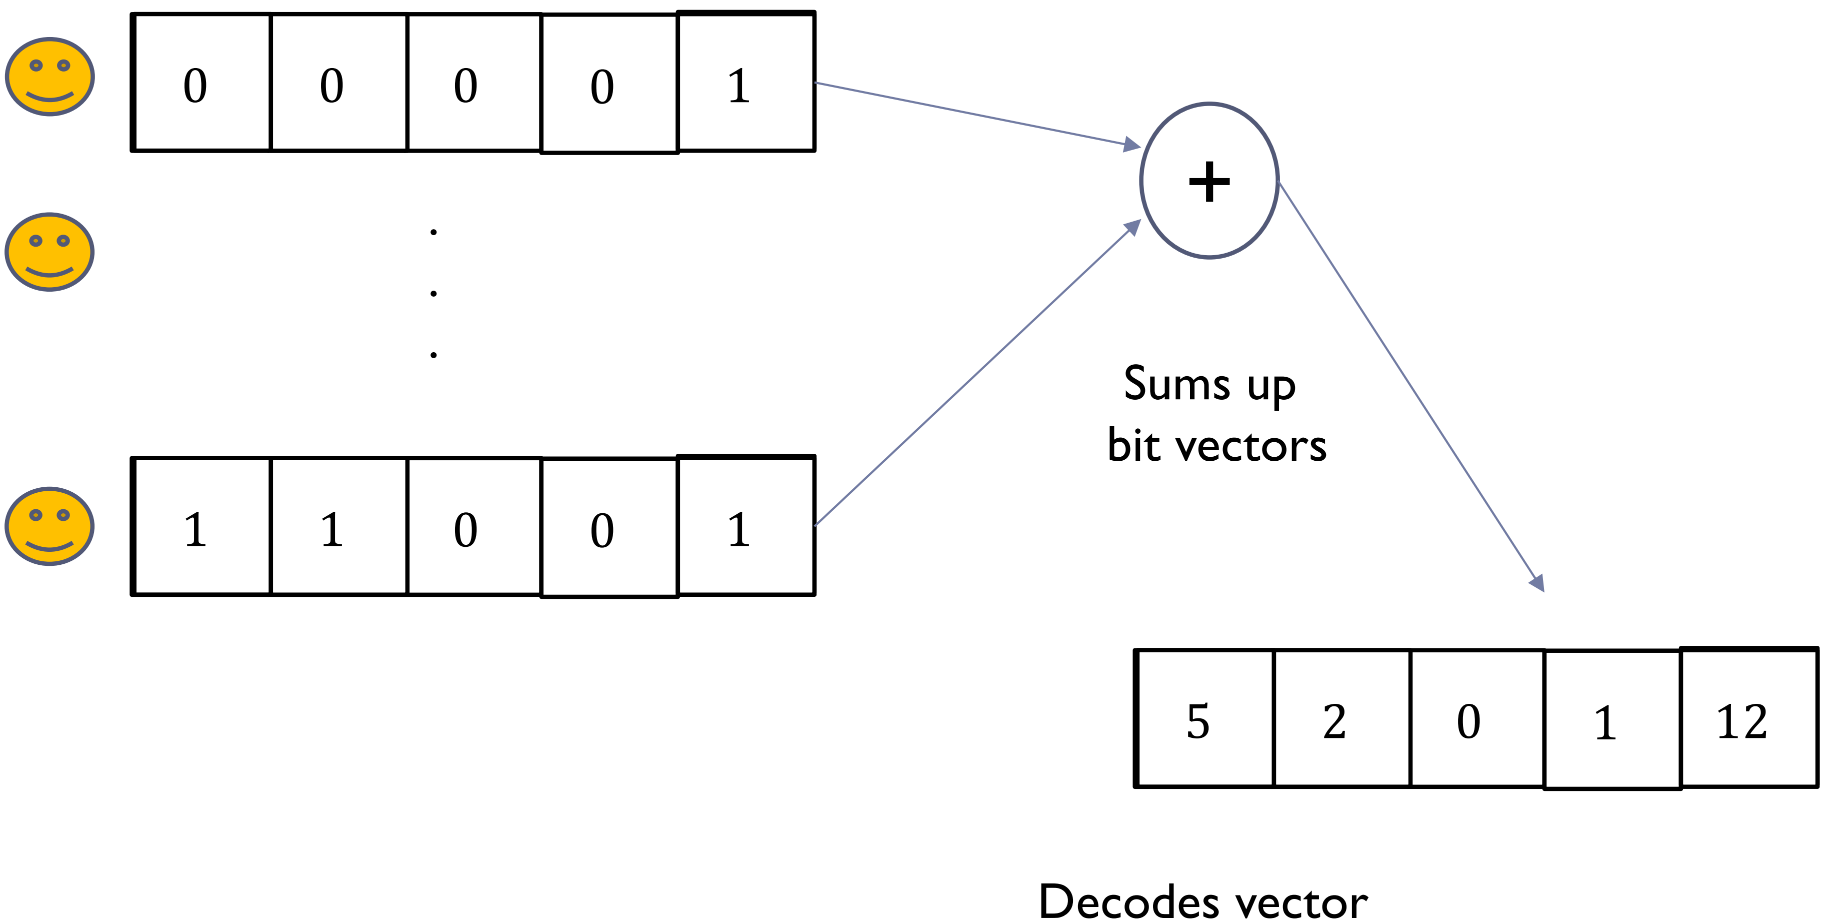

Decoding Bloom Filter

- Aggregator knows:
  - Mapping from words to bits

"Chrome"	0	1	0	0	1
"Firefox"	1	0	0	0	1
"Opera"	0	0	1	1	0

  - Aggregate sum of reported (noisy) vectors
  - Value of parameter f

In-Class Exercise

- Step 1: Go to https://forms.gle/vtsZaTv8CnqqyYsS6 and record your operating system
- Step 2: Create RAPPOR-randomized bits for your OS, and submit them at the same link.
- (wait for class to synchronize)
- Step 3: Form teams of 2-3 students. Try to recover the original distribution. (don't look at RAPPOR paper for this!) Submit your guess here (one per group!): https://forms.gle/CDWPkD6GVPpyYFMx7

Different Techniques

- Let
  - $Y \in R^d$ denote the observed Bloom filter
  - $A \in R^{d \times n}$ the matrix mapping words to initial (unnoised) bits
  - $X \in R^n$ the vector of all real word counts
- Linear regression:

$$\min_{X \in R^n} \|Y - AX\|_2$$

- LASSO

$$\min_{X \in R^n} \|Y - AX\|_2^2 + \lambda \|X\|_1$$

- Hybrid
  - Find support of X via LASSO
  - Solve linear regression to find weights

Chrome homepages estimated by RAPPOR

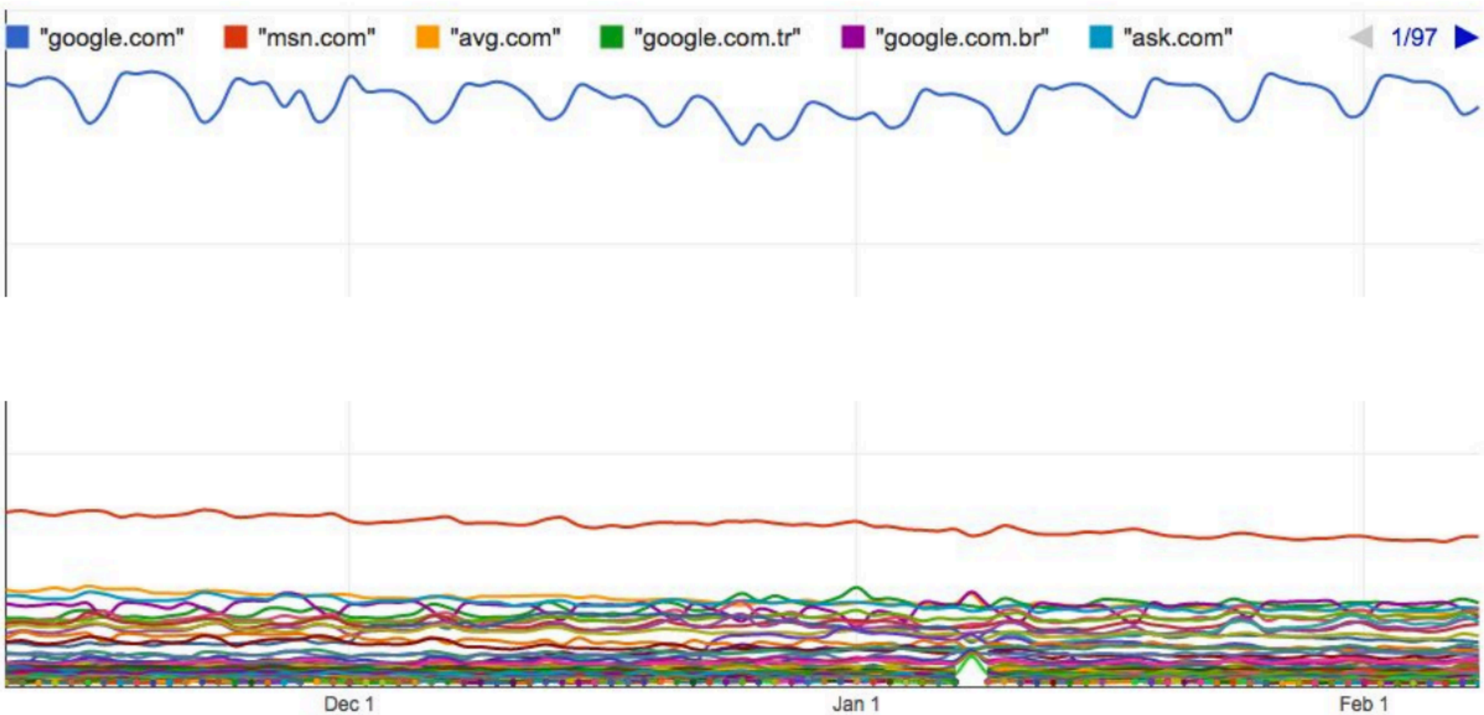

google
msn
avg
google tr
google br

What is the downside of LDP?

- Higher ϵ requires more data
  - Train models
  - Release statistics with given accuracy

- How much more?