

# Heterogeneous Differential Privacy

Mohammad Nabil Alaggan<sup>1</sup>, Sébastien Gambs<sup>2</sup>, Anne-Marie Kermarrec<sup>3</sup>

<sup>1</sup> Helwan University, Egypt (work down while at IRISA/University of Rennes 1, France)

<sup>2</sup> Université de Rennes 1 – INRIA/IRISA, Rennes, France


<sup>3</sup> INRIA Rennes Bretagne-Atlantique, Rennes, France

The Theory and Practice of Differential Privacy

London, UK – April 18, 2015



- Personalization through user-based collaborative filtering
  - Users who share similar interests like similar items
- P2P system:
  - Users host their own data
  - They meet at random and find those who share their interests
    - By computing a similarity metric (probably based on the inner product)
  - They construct a  $k$ -Nearest Neighbor overlay network<sup>1</sup>



<sup>1</sup>The Gossple Anonymous Social Network [BFGKL 2010]


- There are  $n$  items (URLs, videos, images, news pieces, ...)
- Each user is associated with a private profile representing his interest with respect to these items (like/dislike, etc. ...)
- Profiles are binary vectors  $\in \{0, 1\}^n$ 
  - 1 means like
  - 0 means either dislike, or has never been rated

Example:

	wikipedia.com	gnu.org	debian.org	archlinux.org
Alice's profile	0	1	0	1
Bob's profile	1	0	1	1


## User's interests (population)

- To have any utility at all, some information has to leak about the users interests




- The user only has a sample of his interests (his profile)
- We assume a profile is sampled independently
  - Otherwise correlations may violate privacy [KM11]
  - Some solution (not part of this work): preprocess the profiles in a DP manner to remove correlations, for instance, by using dimensionality reduction or SVD

[KM11] Kifer and Machanavajjhala. No free lunch in data privacy. SIGMOD'11



- A function of the profile, satisfying differential privacy, is released



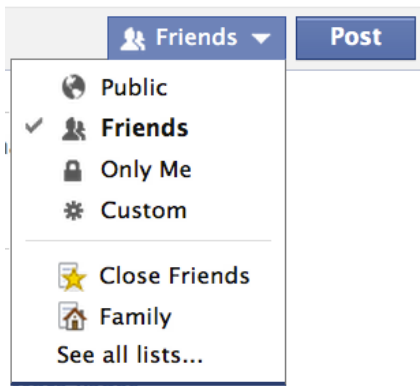
- Learning aggregate information about the user's interests (such as the similarity to another user) is not considered a privacy violation
- Inferring items of the profile is considered a privacy violation (evidence of participation of a particular item)
  - Obstruct linkage/deanonymization/reconstruction attacks
  - Plausible deniability
  - Decreases incentives to *not* add *an* item to the profile

- Using the Laplacian mechanism for a function  $f$  with global sensitivity  $\Delta f$
- Release

$$f(\mathbf{d}) + \text{Lap}(\Delta f / \epsilon) . \quad (1)$$

- Problem: assumes same privacy budget ( $\epsilon$ ) for all users and all items





Different people have different privacy expectations and some items may be considered more private than others

## Definition $((\varepsilon, v_1, \dots, v_n)$ -Differential Privacy (HDP))

A randomized algorithm  $\mathcal{A}$  satisfies  $(\varepsilon, v_1, \dots, v_n)$ -differential privacy if **for all items  $i$**  and for all neighboring profiles  $\mathbf{d}, \mathbf{d}^{(i)}$ , and  $S \subset \text{Range}(\mathcal{A})$ :

$$\Pr[\mathcal{A}(\mathbf{d}) \in S] \leq \exp(\varepsilon v_i) \Pr[\mathcal{A}(\mathbf{d}^{(i)}) \in S] , \quad (2)$$

in which  $\mathbf{d}$  and  $\mathbf{d}^{(i)}$  differ on *at most* item  $i$ , and each privacy weight  $v_i$  is in the range  $[0, 1]$ , where 0 means absolute privacy, and 1 means traditional  $\varepsilon$ -differential privacy.

- HDP is composable

## Definition (Modular Global Sensitivity)

Modular Global sensitivity of a function  $f$  with respect to  $i$  is

$$\Delta_i f = \max_{\mathbf{d} \sim \mathbf{d}^{(i)}} \left| f(\mathbf{d}) - f(\mathbf{d}^{(i)}) \right| . \quad (3)$$

The traditional global sensitivity is then

$$\Delta f = \max_i \Delta_i f . \quad (4)$$

# Key Observation

Let  $\lambda = \Delta f / \varepsilon$  and  $N = \text{Lap}(\lambda)$ . From the PDF of the Laplacian distribution:

$$\frac{\Pr[f(\mathbf{d}) + N \in S]}{\Pr[f(\mathbf{d}^{(i)}) + N \in S]} \leq \exp(\Delta_i(f)/\lambda) \quad (5)$$

$$= \exp(\underbrace{c_i \Delta(f)/\lambda}_{\varepsilon_i}) \quad (6)$$

$$\leq \exp(\varepsilon) \quad (7)$$

for some  $0 \leq c_i \leq 1$ .

## Core Idea

Manipulating  $\Delta_i$  controls  $\varepsilon_i$ :  
The lower  $\Delta_i$  is, the more private the item  $i$

- For a privacy vector  $\mathbf{v}$ , any  $(\min_i v_i)$ -differentially private mechanism satisfies HDP
- Sacrifices a lot of utility:
  - What if  $\min_i v_i$  is much lower than other weights, or zero?
- Would like to use heterogeneity to achieve better utility

Given a profile  $\mathbf{d}$  and privacy weights  $\mathbf{v} \in [0, 1]^n$ , find a weight vector  $\mathbf{w} \in [0, 1]^n$  such that the modular global sensitivities for all  $i$  of

$$\hat{f} = f(w_1 d_1, \dots, w_n d_n) \quad (8)$$

satisfy

$$\Delta_i(\hat{f}) \leq v_i \Delta f \quad (9)$$

then invoke the standard Laplacian mechanism  $\hat{f} + N$

## Theorem

The stretching mechanism achieves  $(\varepsilon, v_1, \dots, v_n)$ -differential privacy

## Theorem (Distortion)

$$\left| f(\mathbf{d}) - \hat{f}(\mathbf{d}) \right| \leq (1 - m) \|\nabla f(B \cdot \mathbf{d})\| \|\mathbf{d}\|, \quad (10)$$

*in which*

- $m = \min_i v_i$
- $B = cI + (1 - c)T$  for some  $0 \leq c \leq 1$
- $T = \text{diag}(\mathbf{w})$

## Distortion for the Inner Product

- Additive distortion (for inner product) is  $\tilde{O}(\sqrt{n})$  when the privacy weights are  $1 - \tilde{O}(1/n^{3/2})$
- The privacy weights should approach 1 as  $n$  grows, otherwise the distortion will grow beyond the desired bound
- This distortion does not hurt the utility since the lower bound on the noise for the two-party inner product is  $\tilde{\Omega}(\sqrt{n})$  [MMPRTV]<sup>2</sup>

---

<sup>2</sup>[MMPRTV] McGregor, Mironov, Pitassi, Reingold, Talwar, and Vadhan. The limits of two-party differential privacy. FOCS'10




- The privacy weights themselves should be private
- If user  $A$  gives a high privacy weight to item  $j$ , it indicates that this item as a special interest for him
- Our stretching mechanism guarantees that the privacy weights **themselves** are protected with the standard  $\varepsilon$ -differential privacy: for all profiles  $\mathbf{d}$  and neighboring privacy vectors  $\mathbf{v}, \mathbf{v}'$



$$\max_{\mathbf{v} \sim \mathbf{v}'} |\hat{f}_{\mathbf{v}}(\mathbf{d}) - \hat{f}_{\mathbf{v}'}(\mathbf{d})| \leq \Delta f \quad (11)$$

# Experimental Results

## Digg



## Delicious



## JYC15<sup>3</sup>

- Parallel line of work, with the following differences:
- There is one privacy weight per user
- Privacy weights are public
  - Can be determined only based on public data; like the user's occupation
  - Cannot be determined by the user's preference if his preference reflects private information

## ESS15<sup>4</sup>

- Previous Talk!

## PGM14<sup>5</sup>

- Uses rescaling as a mathematical tool to make query transformations (like Join) stable
  - To avoid worst-case sensitivity, for the purpose of better utility
- Incomparable to our work: Unclear how it could be applied to ensure different levels of privacy

---

<sup>3</sup>[JYC] Jorgensen, Yu and Cormode. ICDE'15

<sup>4</sup>[ESS15] Ebadi, Sands and Schneider. POPL'15

<sup>5</sup>[PGM14] Proserpio, Goldberg and McSherry. PVLDB 7(8): 637–648 (2014)

## Summary

- Proposed a mechanism to ensure heterogeneous differential privacy
- Demonstrated that it protects both the items and the privacy weights
- Some functions do not lend themselves to heterogeneous differential privacy
  - The  $\ell_0$  function (the support of the input vector)
  - The minimum function (semantics of the function not preserved)

- Investigate solutions for profiles with correlated items
  - Perhaps through dimensionality reduction
- Propose mechanisms with less distortions
  - Investigate pre/post processing techniques that can reduce distortion
    - May depend on the weight vector; must take care
- Find lower bounds on distortion for families of functions
- Investigate non-interactive mechanisms for HDP

Thank you!