Optimal and Differentially Private Data Acquisition: Central and Local Mechanisms

Alireza Fallah*, Ali Makhdoumi†, Azarakhsh Malekian‡, Asuman Ozdaglar*

*: MIT †: Duke University ‡: University of Toronto

Data Acquisition Problem

- Advances in AI → demand for data is increasing!
- Companies collect data in various ways:
  - offering payments, e.g., Nielsen.
  - collecting it as a byproduct of their services, e.g., Netflix, Google, or Facebook.

- A common concern: **Privacy**!
- As a user's data is harnessed, more & more information about her behavior & preferences are uncovered.

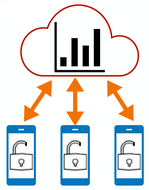

- **Differential Privacy (DP)** [Dwork et al. 2006]: a popular framework for characterizing privacy losses:
  - Widely used by tech companies, including Apple, Google, and Microsoft.

How much privacy an individual will obtain?

- Two individuals have similar data, but one is more privacy conscious.
- We want to provide different privacy levels to them.
- How should we do this? How should we elicit their privacy sensitivities?

- **Model:**
- A platform is interested in estimating a parameter θ by collecting data from users $1,\cdots,n$.
- Data of user i: $x_i = \theta + Z_i \in \mathcal{X}$.
- Z_i's: i.i.d. zero mean with variance VAR and $|Z_i| \le 1/2$.

Central & Local Differential Privacy

- In the central setting, users trust the platform and the platform releases a private estimate.
- In the local setting, users make their data private from the beginning: user i shares $\mathscr{C}_i(x_i)$ with the platform:

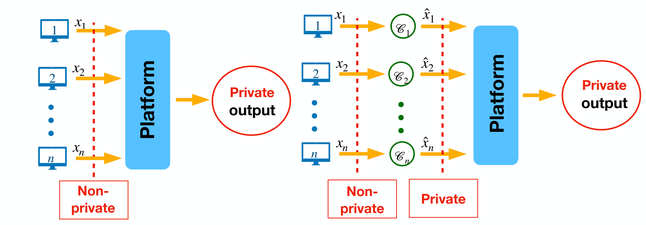

Central Differential Privacy **Local Differential Privacy**

Central & Local Differential Privacy

- An algorithm $\mathscr{A}$ is $(\varepsilon_i)_{i=1}^n$-**centrally DP** if:

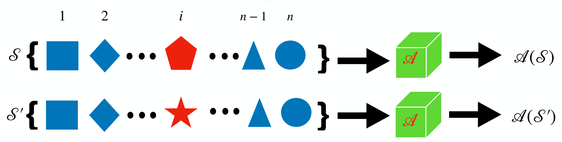

$$\mathbb{P}(\mathscr{A}(\mathcal{S}) \in \mathscr{W}) \le e^{\varepsilon_i}\, \mathbb{P}(\mathscr{A}(\mathcal{S}') \in \mathscr{W})$$

- ε_i: maximum **privacy loss level** of user i.
- Smaller ε_i = user i's data has less impact on the output of Algorithm → higher privacy guarantee for user i.

- An algorithm $\mathscr{A}$ is $(\varepsilon_i)_{i=1}^n$-**locally DP** if:

$$\mathbb{P}(\mathscr{C}_i(x) \in \mathscr{W}) \le e^{\varepsilon_i}\, \mathbb{P}(\mathscr{C}_i(x') \in \mathscr{W}), \quad \forall x, x' \in \mathcal{X}.$$

- **Laplace mechanism:**
  - Adding Laplace noise to ensure privacy.
  - **Lemma**: Linear estimator $\sum_{i=1}^n w_i x_i + \text{Laplace}(1/\eta)$ is $(w_i\eta)_{i=1}^n$ centrally DP.

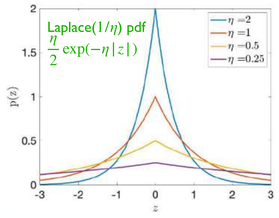

Question

- Assume $(\varepsilon_1)_{i=1}^n$ are given.
- Which estimator is *optimal* w.r.t. mean square error?

Minimax error

$$\mathscr{L}(\Theta, \mathscr{P}) := \inf_{\hat{\theta} \in \Theta} \sup_{P \in \mathscr{P}} \mathbb{E}_{(x_i \sim P)_{i=1}^n, \hat{\theta}} \left[\left| \hat{\theta}(x_{1:n}) - \theta(P) \right|^2 \right]$$

Theorem I

- Assume $\varepsilon_1 \le \ldots \varepsilon_n \le 1$. Let $\mathscr{P}$ be family of distributions such that $|X| \le 1/2$ a.s.
- For Θ ≡ family of $(\varepsilon_i)_{i=1}^n$-centrally DP estimators

$$\mathscr{L}(\Theta, \mathscr{P}) \ge \mathcal{O}(1) \left(\max_{k \in \{0,1,\cdots,n\}} \frac{1}{n - k + (\sum_{i=1}^k \varepsilon_i)^2} \wedge 1 \right).$$

- For Θ ≡ family of $(\varepsilon_i)_{i=1}^n$-locally DP estimators

$$\mathscr{L}(\Theta, \mathscr{P}) \ge \mathcal{O}(1) \left(\frac{1}{\sum_{i=1}^n \varepsilon_i^2} \wedge 1 \right)$$

- Achievable (up to log factor) by **linear estimators.**

Private Data Acquisition Mechanism

- How $(\varepsilon_i)_{i=1}^n$ are endogenized?
- A user's **privacy sensitivity**: per unit cost of privacy loss.

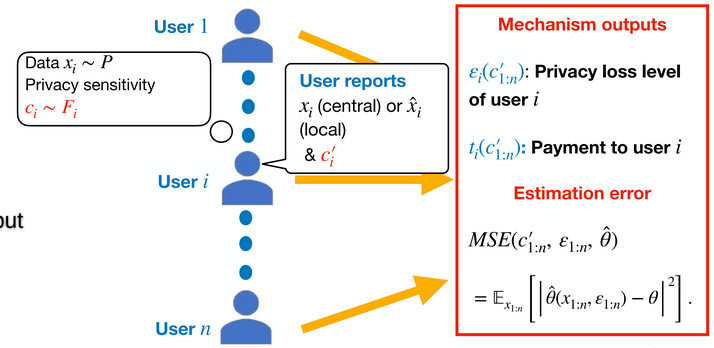

- The cost of user i with privacy sensitive c_i who reports c_i':

$$\text{Cost}_i(c_i', c_i; \varepsilon_{1:n}, t_i, \hat{\theta}) = \mathbb{E}\left[MSE(c_i', c_{-i}, \varepsilon_{1:n}, \hat{\theta}) + c_i\, \varepsilon_i(c_i', c_{-i}) - t_i(c_i', c_{-i}) \right]$$

- The cost of a nonparticipating user: VAR (her *best* estimate given her data alone).

Platform's problem

$$\min_{\hat{\theta}, \varepsilon_{1:n}(\cdot), t_{1:n}(\cdot)} \mathbb{E}_{c_{1:n}} \left[MSE(c_{1:n}, \varepsilon_{1:n}, \hat{\theta}) + \sum_{i=1}^n t_i(c_{1:n}) \right]$$

- Incentive Compatibility (IC): user i has no incentive to misrepresent her privacy sensitivity when others report truthfully:

$$Cost_i(c_i, c_i; \varepsilon_{1:n}, t_i, \hat{\theta}) \le Cost_i(c_i', c_i; \varepsilon_{1:n}, t_i, \hat{\theta}).$$

- Individual Rationality (IR): the platform does not make users worse off by participating in the mechanism:

$$Cost_i(c_i, c_i; \varepsilon_{1:n}, t_i, \hat{\theta}) \le VAR\,.$$

Proof Ideas [for the centrally DP case]

- **Lower bound:** Le Cam's method
1. Replace $\sup_{p \in \mathscr{P}}$ by average over arbitrary $P_1, P_2 \in \mathscr{P}$.
2. We show $\mathscr{L}(\Theta, \mathscr{P}) \ge \frac{(\mu_{P_1} - \mu_{P_2})^2}{8}(1 - \|Q_1 - Q_2\|_{TV})$ where Q_i is the distribution of $\hat{\theta}(x_{1:n})$ when $x_i \sim P_i$.
3. Then, we show that for any k

$$\|Q_1 - Q_2\|_{TV} \le 2\|P_1 - P_2\|_{TV} \sum_{i=1}^k (e^{\varepsilon_i} - 1) + \sqrt{\frac{n-k}{2} D_{KL}(P_1, P_2)}$$

4. Finally, choose P_1 & P_2 as two Bernoullis and optimize over their means' difference.

- **Upper bound** idea: w_i grows proportional to ε_i up to some k and then **remains constant!**

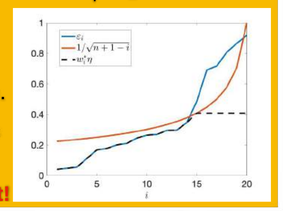

From Mechanism Design to Optimization

- We characterize the payment that satisfies IC & IR.
- Plugging it back, we simplify the optimization problem:

$$\min_{\varepsilon_{1:n}(\cdot)} \mathbb{E}_{c_{1:n}} \left[(n+1) MSE(c_{1:n}, \varepsilon_{1:n}, \hat{\theta}) + \sum_{i=1}^n \varepsilon_i(c_{1:n}) \psi_i(c_i) \right] - nVAR$$

$$\text{s.t. } \mathbb{E}_{\mathbf{c}_{-i}} \left[\varepsilon_i(z, \mathbf{c}_{-i}) \right] \text{ is decreasing in } z \text{ for all } i$$

- Assume **virtual cost** $\psi_i(c) = c + \frac{F_i(c)}{f_i(c)}$ is increasing.

Theorem II

- Let us focus on **linear estimators.**
- For central DP, the optimization problem reduces to

$$\min_{\varepsilon_{1:n}} \frac{n+1}{\left(\sum_{j=1}^n \varepsilon_j\right)^2} \left(2 + \sum_{i=1}^n \varepsilon_i^2\, VAR \right) + \sum_{i=1}^n \psi_i(c_i)\varepsilon_i$$

- We develop a score-based algorithm to solve this problem in $\mathcal{O}(n\log(n))$.
- For local DP, the optimization problem reduces to

$$\min_{\varepsilon_{1:n}} \frac{n+1}{\sum_{i=1}^n \frac{1}{VAR + 2/\varepsilon_i^2}} + \sum_{i=1}^n \psi_i(c_i)\varepsilon_i$$

- We develop an algorithm to find a δ-approximate solution in time $\text{poly}(n, 1/\delta)$.
- Proofs rely on using KKT conditions, defining auxiliary variables, and characterizing structures of the solution.

Central vs. Local: Platform's objective

- We establish that the platform's optimal cost under central differential privacy setting is always **weakly smaller.**
- Example with two users with c_i uniform over $[1,2]$:

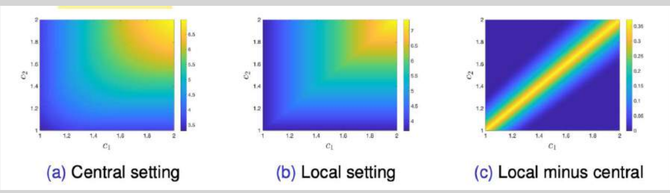

(a) Central setting (b) Local setting (c) Local minus central

- However, user's privacy loss can be smaller in the local setting.

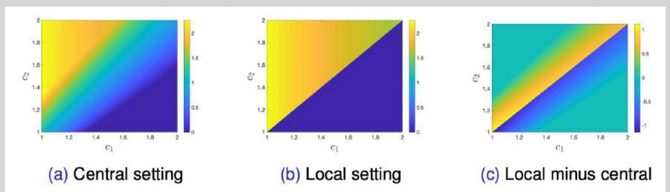

(a) Central setting (b) Local setting (c) Local minus central