Modeling Differential Privacy

Presenter: Sun Jiaxuan

Huazhong University of Sci & Tech

Econ & CS Summer Reading Group

Related Literature

- Dwork, C., and Roth, A. (2014):"The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, 9(3-4). Pages 5-18
- Eilat, R., K. Eliaz, and X. Mu. (2021):"Bayesian Privacy," Theoretical Economics.

Natural Approach

Privacy-preserving data analysis:

- Data cannot be FULLY anonymized and remain useful.
- Re-identification could be harmful.

Key trade-off for privacy protection:

- Detailed and accurate statistics
- Privacy and confidentiality

Requirement: analysts know no more about any individual in the dataset after analysis than before.

Toy model: Randomized Disclosure

Asking participants having property P or not. Participants' state: $\omega \in \{P, \neg P\}$. Denote probability of P: $\pi = ¶(\omega = P)$. Participants report signal s following randomized rule:

1. Flip a coin.
2. If **tails**, respond truthfully.
3. If **heads**, then flip a second coin and respond "P" if heads and "$\neg P$" if tails.

$$Tails : ¶(\omega = s) = 1$$

$$Heads : ¶(s = P) = ¶(s = \neg P) = \frac{1}{2}$$

$¶(s = P)$: $\frac{\pi}{2} + \frac{1}{4} \Rightarrow$ Estimation of $\pi : 2¶(s = P) - \frac{1}{2}$

Formalizing DP

A randomized algorithm with domain A and range B will be associated with a mapping from A to probability simplex over B, denoted by ΔB.

Definition 1

Given a discrete set B, the probability simplex over B, ΔB is

$$\Delta B = \{x \in \mathbb{R}^{|B|} : x_i \geq 0 \text{ for all } i \text{ and } \sum_{i=1}^{|B|} x_i = 1\}$$

A randomized algorithm $\mathcal{M}$, with domain A and range B is associated with a mapping $M : A \to \Delta B$. On input $a \in A$, $\mathcal{M}$ outputs $\mathcal{M}(a) = b$ with probability $(M(a))_b$ for each $b \in B$.

Database Distances

We think of databases x as collections of records from universe $\mathcal{X}$. $x \in \mathbb{N}^{|\mathcal{X}|}$. x_i represents number of elements in x of type $i \in \mathcal{X}$.

A measure between two databases x and y will be l_1 distance: $\|x - y\|_1$ where $\|x\|_1 = \sum_{i=1}^{|\mathcal{X}|} |x_i|$.

$\|x\|_1$ is a measure of the SIZE of database x, and $\|x - y\|_1$ is a measure of how many records DIFFER between x and y.

Differential Privacy

Definition 2

A randomized algorithm $\mathcal{M}$ with domain $\mathcal{N}^{|\mathcal{X}|}$ is (ε, δ)-differentially private if for all $\mathcal{S} \subseteq \mathit{Range}(\mathcal{M})$ and for all $x, y \in \mathcal{N}^{|\mathcal{X}|}$ such that $\|x - y\|_1 \leq 1$

$$Pr[\mathcal{M}(x) \in \mathcal{S}] \leq \exp(\varepsilon) Pr[\mathcal{M}(y) \in \mathcal{S}] + \delta$$

If $\delta = 0$, we say that $\mathcal{M}$ is ε-differentially private.

Privacy Loss

For two outputs $\mathcal{M}(x), \mathcal{M}(y)$, KL-divergence between them is

$$D(\mathcal{M}(x) \| \mathcal{M}(y)) = \mathbb{E}_{z \sim \mathcal{M}(x)}[\ln \frac{Pr[\mathcal{M}(x) = z]}{Pr[\mathcal{M}(y) = z]}]$$

Max divergence between them is

$$\begin{aligned} D_\infty(\mathcal{M}(x) \| \mathcal{M}(y)) &= \max_{S \subset \text{supp}(\mathcal{M}(x))} [\ln \frac{Pr[\mathcal{M}(x) \in S]}{Pr[\mathcal{M}(y) \in S]}] \\ &= \max_{z \in \mathcal{M}(x)} [\ln \frac{Pr[\mathcal{M}(x) = z]}{Pr[\mathcal{M}(y) = z]}] \\ &\leq \varepsilon \end{aligned}$$

We define ε as maximum differential privacy loss.

Mechanism Design Approach (Eilat, Eliaz, and Mu '21)

A bayesian measure of privacy loss

How much the principal learns about agent types through observing what they choose in the mechanism. Specifically,

1. Principal has prior belief F about agent types t
2. He offers a general mechanism $\mathbb{M}$ specifying message set M and how messages are mapped to outcomes
3. Agents play a Bayesian equilibrium
4. Given message m, principal forms posterior belief $F(\cdot|m)$
5. Privacy loss defined as expected KL-divergence between posterior and prior beliefs: $I(\mathbb{M}) = \mathbb{E}_m[D(F(\cdot|m)\|F)]$
6. Principal constrianed by $I(\mathbb{M}) \leq \varepsilon$ exogenously given.

Screening Environment

Monopolistic screening of Mussa and Rosen ('78)

- A seller sells some quality $q \geq 0$ to a buyer for payment p
- Buyer type $\theta \in [\underline{\theta}, \overline{\theta}]$ distributed as F with positive density
- Profit net of production cost $p - \frac{q^2}{2}$. Buyer utility $q \cdot \theta - p$.
- Seller maximizes profit subject to privacy. That is,

$$\max \mathbb{E}_m[p(m) - c(q(m))] \quad s.t. \quad \mathbb{E}_m[D(F(\cdot|m)\|F)] \leq \varepsilon$$

Coarse Revelation

Theorem 3

Given $0 < \varepsilon < \infty$, *There exists optimal privacy-constrained mechanism* $\mathbb{M}$, *where the set of types* $[\underline{\theta}, \overline{\theta}]$ *is partitioned into finitely many intervals, and in equilibrium each type truthfully reports its interval.*

Other properties:

- privacy constraint binds in any optimal mechanism
- it ε small, exactly two intervals used

Welfare Analysis

Comparative Statics w.r.t. ε:

- Profit from a ε-constrained optimal mechanism increases in ε.
- Buyer surplus is maximized with full privacy, minimized with no privacy

Recap and Forward

- Differential privacy: trade-off between statistical accuracy and confidentiality
- Characterization: ε privacy loss
- DP in econ: Bayesian measure of privacy loss in mechanism design framework

Questions to ask:

- Regulators play strategically to decide ε
- Data Erasure right
- Dynamic environment