Extremal Mechanisms for Local Differential Privacy

Peter Kairouz, Sewoong Oh, and Pramod Viswanath
E-mails: {kairouz2,swoh,pramodv}@illinois.edu
University of Illinois at Urbana Champaign, USA

Local Differential Privacy

Local Privacy Model:

- Warner's 1965 randomized response

Have you ever used illegal drugs?

say yes — answer truthfully

- tension between the need to ***share data*** and the need to ***protect privacy***
- data providers do not trust data collectors (analysts)

Local Differential Privacy:

- Q is a privatization mechanism that maps $X \in \mathcal{X}$ stochastically to $Y \in \mathcal{Y}$
- for a non-negative ε, we say that Q is ***ε-locally differentially private*** if

$$e^{-\varepsilon} \leq \frac{Q(Y=y|X=x)}{Q(Y=y|X=x')} \leq e^{\varepsilon}$$

Operational Interpretation of Differential Privacy:

- for any $A, B \subset \mathcal{X}$ such that $A \cap B = \emptyset$, form the following hypothesis test

$$H_0: \text{ original data } X \in A$$
$$H_1: \text{ original data } X \in B$$

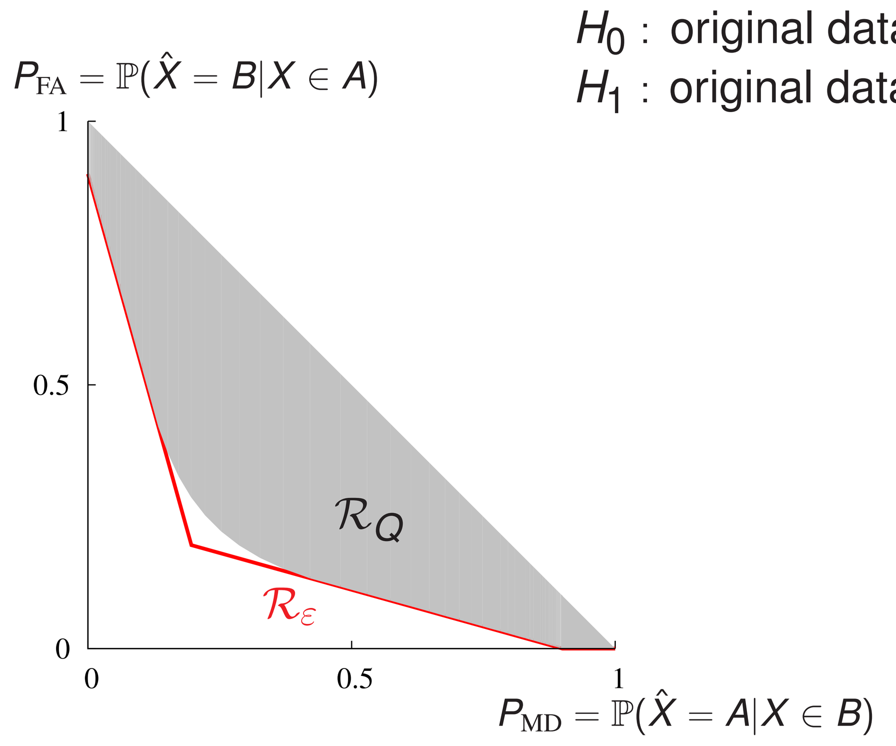

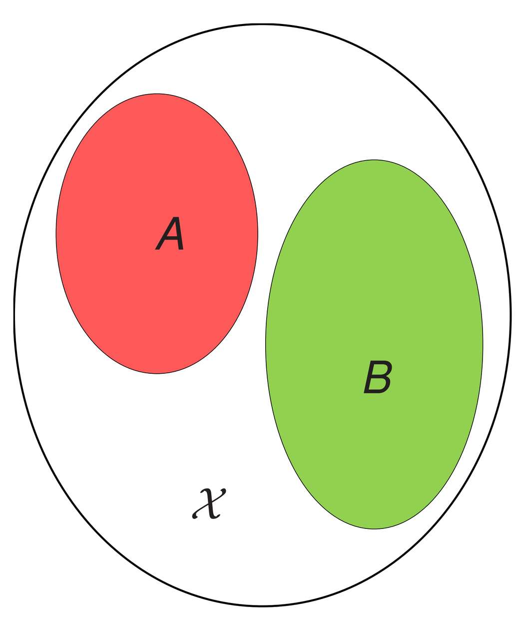

Operational Definition of Differential Privacy

Q is ε-locally differentially private $\iff$ $\mathcal{R}_Q \subseteq \mathcal{R}_\varepsilon$

$$P_{\rm FA} + e^{\varepsilon} P_{\rm MD} \geq 1$$
$$e^{\varepsilon} P_{\rm FA} + P_{\rm MD} \geq 1$$

Information Theoretic Utility Functions

Hypothesis Testing and Classification:

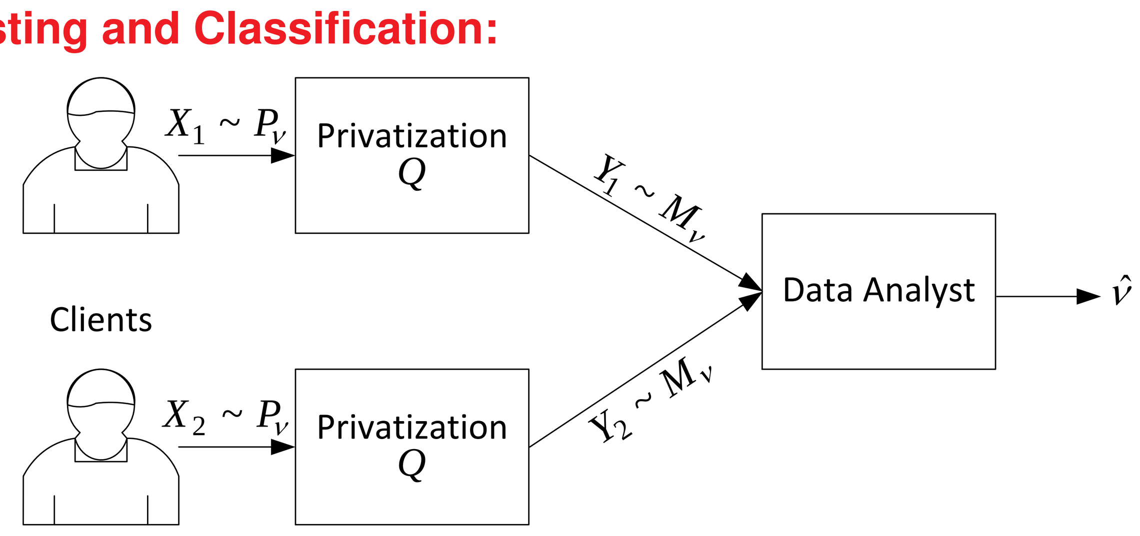

- the X_i's are sampled from a distribution P_ν parameterized by $\nu \in \{0, 1\}$
- given the Y_i's, the data analyst would like to detect whether $\nu = 0$ or $\nu = 1$
- performance is a function of distance between M_0 from M_1

$$M_\nu(S) = \int Q(S|x) dP_\nu(x) ,$$

- Chernoff-Stein's lemma: the best type II error probability scales as $e^{-n D_{\rm kl}(M_0 || M_1)}$
- **result:** when ε is sufficiently small, the effective sample size is reduced from n to $\varepsilon^2 n$

Information Theoretic Utilities:

- for some convex function f such that $f(1) = 0$, Csiszár's f-divergence is defined as

$$D_f(M_0 || M_1) = \int f\left(\frac{dM_0}{dM_1}\right) dM_1 ,$$

- KL divergence $D_{\rm kl}(M_0||M_1)$ and total variation $\|M_0 - M_1\|_{\rm TV}$ are special cases
- f-divergences capture: ***minimax rates*** and ***error exponents***

Fundamental Limits of Privacy:

- the **more** private you want to be, the **less** utility you get
- there is a ***fundamental trade-off*** between privacy and utility

$$\underset{Q}{\text{maximize}} \;\; D_f(M_0|M_1)$$
$$\text{subject to } \;\; Q \in \mathcal{D}_\varepsilon$$

- $\mathcal{D}_\varepsilon$ is the set of all ε-locally differentially private mechanisms
- this maximization problem is ***nonlinear, non-standard, and infinite dimensional***

Binary Data

Optimality of the Binary Randomized Response Mechanism:

When $|\mathcal{X}| = 2$, the following mechanism is optimal:

w.p. $\frac{1}{1+e^{\varepsilon}}$ **lie** — w.p. $\frac{e^{\varepsilon}}{1+e^{\varepsilon}}$ **answer truthfully**

Larger Alphabets

Definition of Staircase Mechanisms:

- a privatization mechanism is a ***staircase mechanism*** if

$$\frac{Q(Y=y|X=x)}{Q(Y=y|X=x')} \in \{e^{-\varepsilon}, 1, e^{\varepsilon}\}$$

- examples of staircase mechanisms: binary and randomized response mechanisms

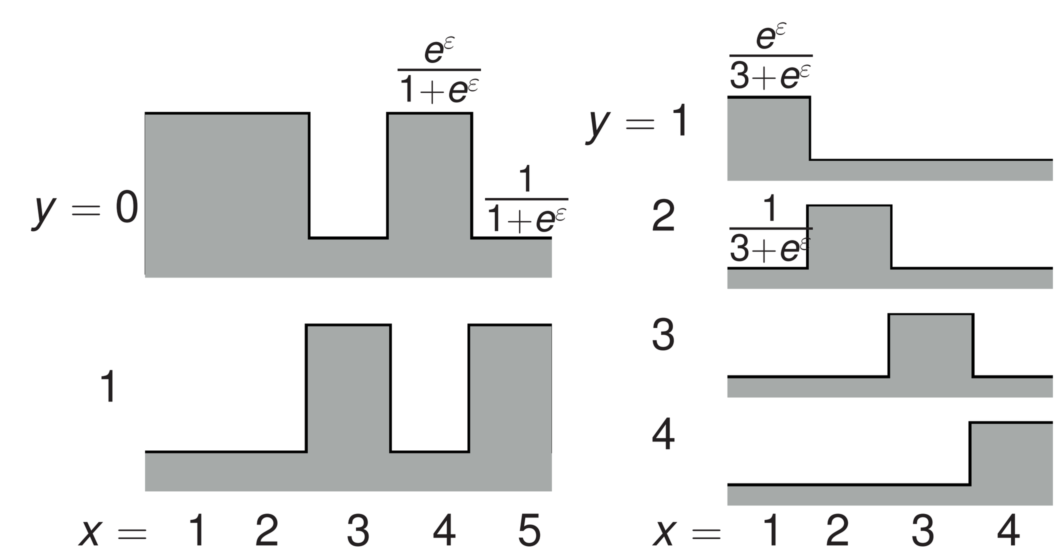

Optimality of Staircase Mechanisms

For any ε, any P_0 and P_1, and any f-divergence, there exists an optimal mechanism Q^* maximizing the f-divergence over all ε-locally differentially private mechanisms, such that Q^* is a staircase mechanism. Moreover, the output alphabet size is at most equal to the input alphabet size: $|\mathcal{Y}| \leq |\mathcal{X}|$.

Definition of Binary Mechanisms:

$$Q(Y=0|X=x) = \begin{cases} \frac{e^{\varepsilon}}{1+e^{\varepsilon}} & \text{if } P_0(x) \geq P_1(x) , \\ \frac{1}{1+e^{\varepsilon}} & \text{if } P_0(x) < P_1(x) . \end{cases} \quad Q(Y=1|X=x) = \begin{cases} \frac{e^{\varepsilon}}{1+e^{\varepsilon}} & \text{if } P_0(x) < P_1(x) , \\ \frac{1}{1+e^{\varepsilon}} & \text{if } P_0(x) \geq P_1(x) . \end{cases}$$

Optimality of Binary Mechanisms in the High Privacy Regime

For any P_0 and P_1, there exists a positive ε^* that depends on P_0 and P_1 such that for any f-divergences and all positive $\varepsilon \leq \varepsilon^*$, the binary mechanism maximizes $D_f(M_0||M_1)$ over all ε-local differentially private mechanisms.

Definition of the Randomized Response Mechanism:

$$Q(Y=y|X=x) = \begin{cases} \frac{e^{\varepsilon}}{|\mathcal{X}|-1+e^{\varepsilon}} & \text{if } y = x , \\ \frac{1}{|\mathcal{X}|-1+e^{\varepsilon}} & \text{if } y \neq x . \end{cases}$$

- can be viewed as a multiple choice generalization to Warner's randomized response
- observe that Q is independent of P_0 and P_1

Optimality of the Randomized Response Mechanism in the Low Privacy Regime

There exists a positive ε^* that depends on P_0 and P_1 such that for any P_0 and P_1, and all $\varepsilon \geq \varepsilon^*$, the randomized response mechanism maximizes the KL-divergence between the induced marginals over all ε-locally differentially private mechanisms.

Big Picture

Local Privacy:

- the local privacy model is particulary important in ***data collection*** applications
- we study a broad class of information theoretic utilities
- we provide ***explicit constructions*** of ***optimal mechanisms***

Our Methods Generalize:

- similar optimality results hold for a large class of convex utility functions
- our techniques can be generalized to private multi-party computation settings
- preprint available on arXiv:

"Differentially Private Multi-party Computation: Optimality of Non-Interactive Randomized Response
Peter Kairouz, Sewoong Oh, and Pramod Viswanath, 2014"