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

Need for Privacy:

- increasing tension between the need to ***share data*** and the need to ***protect personal information***
- with big data comes big responsibilities
- individuals want ***plausible deniability***

Local Privacy Model:

- data providers do not trust data collectors (analysts)
- privacy is achieved by ***randomizing*** the data before releasing it
- this local privacy model dates back to Warner, 1965

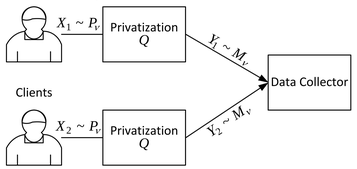

Local Differential Privacy:

- there are n data providers: user i owns $X_i \in \mathcal{X}$ (input alphabet)
- we study finite input alphabets $|\mathcal{X}| = k < \infty$
- Q is a privatization mechanism that maps X_i stochastically to $Y_i \in \mathcal{Y}$ (output alphabet)
- we allow $\mathcal{Y}$ to be (possibly) larger than $\mathcal{X}$
- for a non-negative ε, we say that Q is ***ε-locally differentially private*** if

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

- we cannot distinguish x from x' upon observing $y \in \mathcal{Y}$

Information Theoretic Utility Functions

Statistical Data Model:

- the X_i's are independently sampled from a distribution P_ν parameterized by $\nu \in \{0, 1\}$
- the data analyst is interested in the ***statistics*** of the data as opposed to individual samples
- the power to discriminate data generated from P_0 to data generated from P_1 depends on the 'distance' between the privatized marginals M_0 and M_1

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

Information Theoretic Utility:

- 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 $D_{kl}(M_0||M_1)$ and total variation $\|M_0 - M_1\|_{TV}$ are special cases
- f-divergences capture the quality of statistical inference: ***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
- can we design Q to maximize utility subject to local differential privacy?
- we solve the following problem

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

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

Binary Hypothesis Testing:

- given $\{Y_i\}_{i=1}^n$, the data analyst would like to detect whether $\nu = 0$ or $\nu = 1$
- Chernoff-Stein's lemma: the best type II error probability scales as $e^{-n D_{kl}(M_0||M_1)}$
- we show that when ε is sufficiently small, the effective sample size is reduced from n to $\varepsilon^2 n$

Staircase Mechanisms

Definition of Staircase Mechanisms:

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

$$\frac{Q(y|x)}{Q(y|x')} \in \left\{e^{-\varepsilon}, 1, e^{\varepsilon}\right\}$$

- examples of staircase mechanisms: binary and randomized response mechanisms

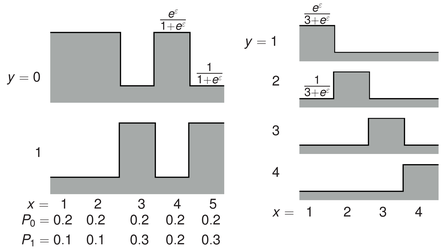

Main Result 1: Optimality of Staircase Mechanisms

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: $|\mathcal{Y}| \leq |\mathcal{X}|$.

Combinatorial Representation of Staircase Mechanisms:

- $S^{(k)} \in \{1, e^\varepsilon\}^{k \times 2^k}$ is called a staircase pattern matrix if its j-th column $S_j^{(k)} = (e^\varepsilon - 1)\, b_{j-1} + \mathbb{1}$
- b_j be the k-dimensional vector corresponding to the binary representation of the integer j.

$$S^{(3)} = \begin{bmatrix} 1 & 1 & 1 & 1 & e^\varepsilon & e^\varepsilon & e^\varepsilon & e^\varepsilon \\ 1 & 1 & e^\varepsilon & e^\varepsilon & 1 & 1 & e^\varepsilon & e^\varepsilon \\ 1 & e^\varepsilon & 1 & e^\varepsilon & 1 & e^\varepsilon & 1 & e^\varepsilon \end{bmatrix} .$$

- if Q is a staircase mechanism, then Q is 'equivalent' to $\tilde{Q} = S^{(k)}\Theta$ for some diagonal matrix $\Theta \in \mathbb{R}^{2^k \times 2^k}$

Main Result 2: Linear Program Formulation

For any ε, any P_0 and P_1, and any f-divergence, maximizing f-divergence subject to local differential privacy is equivalent to solving the following linear program

$$\begin{aligned} \underset{\Theta \in \mathbb{R}^{2^k \times 2^k}}{\text{maximize}} \;\; & \sum_{i=1}^{2^k} \mu_i \Theta_{ii} \\ \text{subject to} \;\; & S^{(k)}\Theta\mathbb{1} = \mathbb{1} \, , \\ & \Theta \text{ is a diagonal matrix} \, , \\ & \Theta \geq 0 \, , \end{aligned}$$

where $\mu_i = \left(\sum_{x \in \mathcal{X}} P_1(x) S_{xi}^{(k)}\right) f\left(\sum_{x \in \mathcal{X}} P_0(x) S_{xi}^{(k)} / \sum_{x \in \mathcal{X}} P_1(x) S_{xi}^{(k)}\right)$ and $D_f(M_0||M_1) = \sum_{i=1}^{2^k} \mu_i \Theta_{ii}$.

- the original maximization problem is now reduced to a ***finite dimensional linear program***
- however, solving this linear program might be ***computationally expensive*** if k is large

Binary Mechanisms

Definition of Binary Mechanisms:

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

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.

Binary Mechanisms Cnt'd

Optimality of Binary Mechanisms in the High Privacy Regime:

- the binary mechanism is universally optimal in the high privacy regime
- what about the other regimes?

Optimality of Binary Mechanisms for Total Variation Distances

For any P_0 and P_1, and any $\varepsilon \geq 0$, the binary mechanism maximizes total variation between the induced marginals M_0 and M_1 among all ε-locally differentially private mechanisms.

Near Optimality of Binary Mechanisms in the Moderate Privacy Regime

For any ε and any P_0 and P_1, the binary mechanism is an $1/(2(e^\varepsilon + 1)^2)$ approximation of the maximum KL-divergence between the induced marginals M_0 and M_1 among all ε-locally differentially private mechanisms.

- for $\varepsilon \leq 1$, $2(e^\varepsilon + 1)^2 \leq 32$

Randomized Response Mechanism

Definition of the Randomized Response Mechanism:

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

- 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

Global Differential Privacy:

- ***preserve the identity*** of participating individuals and not their data
- global differential privacy is used in ***data release*** applications
- most research in this area focuses on the high privacy regime
- in the high privacy regime, adding Laplace noise to data is order optimal
- exact optimality results are known only in few cases

Our Approach:

- local differential privacy is recent
- the local privacy model is particulary important in ***data collection*** applications
- we consider 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 find optimal privatization mechanisms in a setting where different individuals can collaborate interactively and each individual can be an analyst
- preprint available on arXiv

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

Future Directions:

- what if the X_i's are correlated?
- what if P_0 and P_1 are not known?
- what about m-ary hypothesis testing?

Neural Information Processing Systems (NIPS), December 2014