Journée Aléatoire 2022

Commune à la SFdS, la SMAI et la SMF

29 septembre 2022, Institut Henri Poincaré, Paris

Differential Privacy for Data Analysis

How to learn while respecting individual privacy?

Aurélien Garivier

École Normale Supérieure de Lyon, UMPA & LIP

UNIVERSITÉ DE LYON

ENS DE LYON

Outline

Differential Privacy

The Exponential Mechanism

Laplace Mechanism

The Statistical Price of Differential Privacy

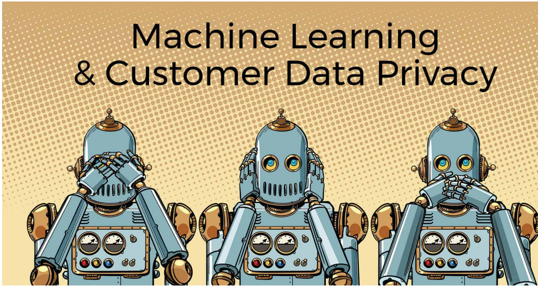

Src: https://www.actian.com/company/blog/

Outline

UNIVERSITÉ DE LYON

ENS DE LYON

Data

Record $x_i \in \mathcal{X}$ for individual i

Data $\mathbf{x} = (x_1, \ldots, x_n) \in \mathcal{X}^n$

Variables

Individuals	Gender (M/F)	Age	Weight (lbs.)	Height (in.)	Smoking (0=No, 1=Yes)	Race
Patient #1	M	59	175	69	0	White
Patient #2	F	67	140	62	1	Black
Patient #3	F	73	155	59	0	Asian
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
Patient #75	M	48	190	72	0	White

Src: `https://statacumen.com/`

Statistical model

The records are iid draws of an unknown probability law $P_\theta \in \mathfrak{M}_1(\mathcal{X})$: under $\mathbb{P}_\theta$, $\mathbf{X} = (X_1, \ldots, X_n) \sim P_\theta^{\otimes n}, \theta \in \Theta$

(Statistical) Data Analysis

Randomized algorithm: for $\mathbf{x} \in \mathcal{X}^n$, outputs $\psi(\mathbf{x}, U) \sim Q_{\mathbf{x}} \in \mathfrak{M}_1(\mathcal{T})$

[*Database*] Target = $f(\mathbf{x})$
[*Statistics*] Target = some functional of P_θ, while $\psi(\mathbf{X}, U) \sim \mathbb{P}_\theta^U$

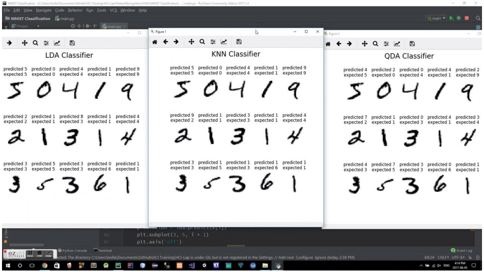

Example: image classification, parameter estimation, prediction rule, etc.

UNIVERSITÉ DE LYON

ENS DE LYON

Framework

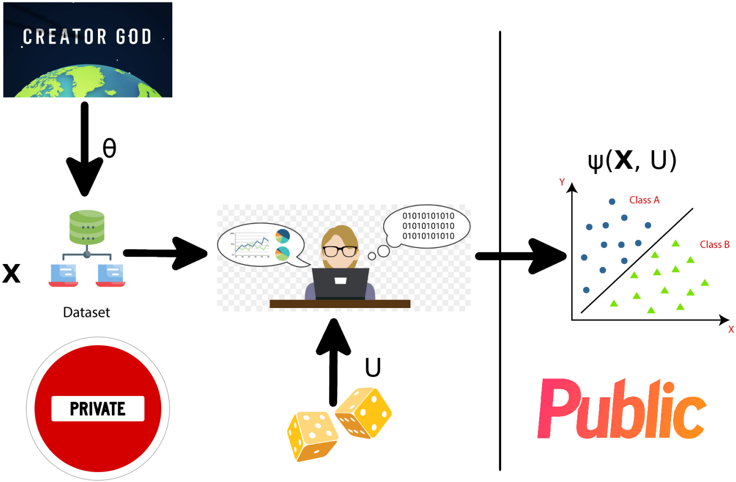

CREATOR GOD
θ
X
Dataset
PRIVATE
01010101010
01010101010
01010101010
U
ψ(X, U)
Y
Class A
Class B
X
Public

Information leakage

Membership attack

Src: https://www.arxiv-vanity.com/papers/1904.05506/

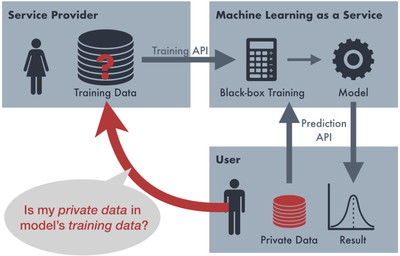

Model inversion attack [Fredrikson et al. '2015]

Figure 1: An image recovered using a new model inversion attack (left) and a training set image of the victim (right). The attacker is given only the person's name and access to a facial recognition system that returns a class confidence score.

See https://arxiv.org/abs/1610.05820 for more information: *Membership Inference Attacks against Machine Learning Models* by Reza Shokri, Marco Stronati, Congzheng Song, Vitaly Shmatikov

Anonymization is not the solution

Linkage attack

[*Simple Demographics Often Identify People Uniquely*, by Latanya Sweeney] showed that gender, date of birth, and zip code are sufficient to uniquely identify the vast majority of Americans.

$\Longrightarrow$ By linking these attributes in a supposedly anonymized healthcare database to public voter records, she was able to identify the individual health record of the Governor of Massachussetts.

Differencing attack

Imposing request on many lines is not the solution

Example from [Dwork & Roth]:

- How many people in the database have the sickle cell trait?
- How many people, not named Z, in the database have the sickle cell trait?

Differential Privacy

DP: attackers can learn virtually nothing *more* about an individual than they would understand if that individual's record were absent from the dataset.

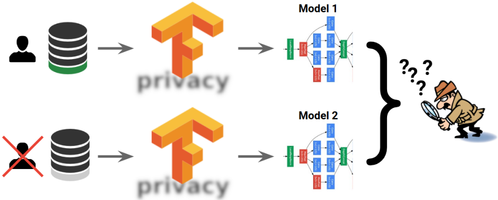

Smoker example

If an individual is openly "smoking" but wants privacy on her medical status,

- a medical study will prove the risk associated with smoking (whether she participates or not)
- a *DP* study will make it impossible to know if she indeed participated or not, even to someone who would have all the remaining information

Survey on triathletes: "do you use doping?"

Triathletes doping status $X_i \overset{iid}{\sim} \mathcal{B}(p)$
but they may lie: answer $\tilde{X}_i \in \{0, 1\}$

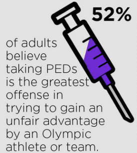

Survey on triathletes: ”do you use doping?”

Triathletes doping status $X_i \overset{iid}{\sim} \mathcal{B}(p)$
but they may lie: answer $\tilde{X}_i \in \{0, 1\}$

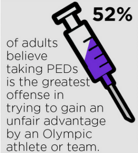

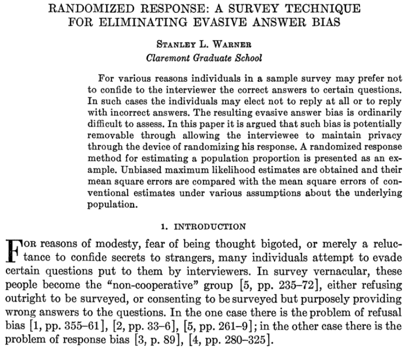

RANDOMIZED RESPONSE: A SURVEY TECHNIQUE FOR ELIMINATING EVASIVE ANSWER BIAS

STANLEY L. WARNER
Claremont Graduate School

For various reasons individuals in a sample survey may prefer not to confide to the interviewer the correct answers to certain questions. In such cases the individuals may elect not to reply at all or to reply with incorrect answers. The resulting evasive answer bias is ordinarily difficult to assess. In this paper it is argued that such bias is potentially removable through allowing the interviewee to maintain privacy through the device of randomizing his response. A randomized response method for estimating a population proportion is presented as an example. Unbiased maximum likelihood estimates are obtained and their mean square errors are compared with the mean square errors of conventional estimates under various assumptions about the underlying population.

1. INTRODUCTION

FOR reasons of modesty, fear of being thought bigoted, or merely a reluctance to confide secrets to strangers, many individuals attempt to evade certain questions put to them by interviewers. In survey vernacular, these people become the "non-cooperative" group [5, pp. 235–72], either refusing outright to be surveyed, or consenting to be surveyed but purposely providing wrong answers to the questions. In the one case there is the problem of refusal bias [1, pp. 355–61], [2, pp. 33–6], [5, pp. 261–9]; in the other case there is the problem of response bias [3, p. 89], [4, pp. 280–325].

Journal of the American Statistical Association, Mar. 1965, Vol.60, No.309, pp. 63-69

See also Chong, Chun Yin Andy & Chu, Amanda & So, Mike & Chung, Ray. (2019). *Asking Sensitive Questions Using the Randomized Response Approach in Public Health Research: An Empirical Study on the Factors of Illegal Waste Disposal*. International Journal of Environmental Research and Public Health.

Survey on triathletes: ”do you use doping?”

Triathletes doping status $X_i \overset{iid}{\sim} \mathcal{B}(p)$
but they may lie: answer $\tilde{X}_i \in \{0, 1\}$

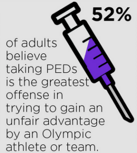

Randomized Response [Warner’65]

Flip a coin, then:

$\rightarrow$ if tails, answer according to another coin flip

$\rightarrow$ if heads, give the right answer

$$\mathbb{P}(\tilde{X}_i = 1 | X_i = x_i) = 1/4 + x_i/2 \qquad \frac{\mathbb{P}(\tilde{X}_i = 1 | X_i = 1)}{\mathbb{P}(\tilde{X}_i = 1 | X_i = 0)} = 3$$

- No triathlete can be prosecuted   one cannot condemn 1/4th of the innocent triathletes!
- But still permits to estimate the proportion of dopers by $\hat{p}_n = 2n^{-1} \sum_{i=1}^{n} \tilde{X}_i - 1$.

Cost: for the same precision, requires $\approx 4x$ more data   or even more if $x(1-x) \ll 1$

Survey on triathletes: "do you use doping?"

Triathletes doping status $X_i \overset{iid}{\sim} \mathcal{B}(p)$
but they may lie: answer $\tilde{X}_i \in \{0, 1\}$

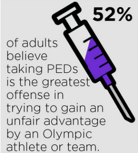

Randomized Response [Warner'65]

Flip a coin, then:

→ if tails, answer according to another coin flip

→ if heads, give the right answer

$$\mathbb{P}(\tilde{X}_i = 1 | X_i = x_i) = 1/4 + x_i/2 \qquad \frac{\mathbb{P}(\tilde{X}_i = 1 | X_i = 1)}{\mathbb{P}(\tilde{X}_i = 1 | X_i = 0)} = 3$$

- No triathlete can be prosecuted — one cannot condemn 1/4th of the innocent triathletes!
- But still permits to estimate the proportion of dopers by $\hat{p}_n = 2n^{-1} \sum_{i=1}^{n} \tilde{X}_i - 1$.

Cost: for the same precision, requires $\approx 4x$ more data — or even more if $x(1-x) \ll 1$

"smoker example": if $\hat{p}_n = 98\%$,
a lot of information on each triathlete
BUT no more than if she had not participated in the study

Formal Definition

Randomized algorithm $\mathcal{A}(\mathbf{x}) = \psi(\mathbf{x}, U)$ = random variable on $\mathcal{T}$

Def: Neighboring databases $\mathbf{x} \sim \mathbf{x}'$ if $\exists i \in \{1, \ldots, n\}, \forall j \neq i, x_j = x'_j$

Differential Privacy

[« Calibrating Noise to Sensitivity », TCC'2006, by Cynthia Dwork, Frank McSherry, Kobbi Nissim et Adam Smith $\implies$ Gödel Prize 2017]

ψ is ε-DP if for all $\mathbf{x} \sim \mathbf{x}'$ and all $\mathcal{S} \subset \mathcal{T}$

$$\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) \in S\big) \leq e^{\varepsilon}\ \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') \in S\big)$$

A person's privacy cannot be compromised by a statistical release if their data are not in the database. Therefore, with differential privacy, the goal is to give each individual roughly the same privacy that would result from having their data removed $\implies$ the statistical functions run on the database should not overly depend on the data of any one individual.

UNIVERSITÉ DE LYON

ENS DE LYON

Formal Definition

Randomized algorithm $\mathcal{A}(\mathbf{x}) = \psi(\mathbf{x}, U)$ = random variable on $\mathcal{T}$

Def: Neighboring databases $\mathbf{x} \sim \mathbf{x}'$ if $\exists i \in \{1, \ldots, n\}, \forall j \neq i, x_j = x'_j$

Differential Privacy

[« Calibrating Noise to Sensitivity », TCC'2006, by Cynthia Dwork, Frank McSherry, Kobbi Nissim et Adam Smith $\Longrightarrow$ Gödel Prize 2017]

ψ is ε-DP if for all $\mathbf{x} \sim \mathbf{x}'$ and all $\mathcal{S} \subset \mathcal{T}$

$$\mathbb{P}^U\Big(\mathcal{A}(\mathbf{x}) \in S\Big) \leq e^{\varepsilon}\, \mathbb{P}^U\Big(\mathcal{A}(\mathbf{x}') \in S\Big)$$

Equivalently,

- if $\mathcal{A}(x)$ is discrete, $\quad -\varepsilon \leq \ln \frac{\mathbb{P}^U\big(\mathcal{A}(x)=t\big)}{\mathbb{P}^U\big(\mathcal{A}(x')=t\big)} \leq \varepsilon \quad$ for all $t \in \mathcal{T}$
- if $\mathcal{A}(x)$ has density $f(\cdot|x)$, $\quad -\varepsilon \leq \ln \frac{f(t|x)}{f(t|x')} \leq \varepsilon \quad$ for all $t \in \mathcal{T}$

Formal Definition

Randomized algorithm $\mathcal{A}(\mathbf{x}) = \psi(\mathbf{x}, U)$ = random variable on $\mathcal{T}$

Def: Neighboring databases $\mathbf{x} \sim \mathbf{x}'$ if $\exists i \in \{1, \ldots, n\}, \forall j \neq i, x_j = x'_j$

Differential Privacy

[« Calibrating Noise to Sensitivity », TCC'2006, by Cynthia Dwork, Frank McSherry, Kobbi Nissim et Adam Smith $\Longrightarrow$ Gödel Prize 2017]

ψ is ε-DP if for all $\mathbf{x} \sim \mathbf{x}'$ and all $\mathcal{S} \subset \mathcal{T}$

$$\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) \in S\big) \leq e^{\varepsilon}\, \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') \in S\big)$$

Differential privacy mathematically guarantees that anyone seeing the result of a differentially private analysis will essentially make the same inference about any individual's private information, whether or not that individual's private information is included in the input to the analysis.

Formal Definition

Randomized algorithm $\mathcal{A}(\mathbf{x}) = \psi(\mathbf{x}, U)$ = random variable on $\mathcal{T}$

Def: Neighboring databases $\mathbf{x} \sim \mathbf{x}'$ if $\exists i \in \{1, \dots, n\}, \forall j \neq i, x_j = x'_j$

Differential Privacy

[« Calibrating Noise to Sensitivity », TCC'2006, by Cynthia Dwork, Frank McSherry, Kobbi Nissim et Adam Smith $\Longrightarrow$ Gödel Prize 2017]

ψ is ε-DP if for all $\mathbf{x} \sim \mathbf{x}'$ and all $\mathcal{S} \subset \mathcal{T}$

$$\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) \in S\big) \leq e^{\varepsilon}\, \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') \in S\big)$$

In the previous example on the DP survey, algorithm $\mathcal{A}(\mathbf{X}) = (\tilde{X}_1, \dots, \tilde{X}_n)$ is $\ln(3)$-DP.

Note that it outputs an entire (differentially private), which is unusual: more often, we just want the answer to a query.

Properties

Post-processing

If $\mathcal{A} : \mathcal{X}^n \to \mathfrak{M}_1(\mathcal{T})$ is ε-DP, then for every $f : \mathcal{T} \to \mathcal{T}'$ algorithm $f \circ \mathcal{A}$ is also ε-DP

Group privacy

If $\mathbf{x} \sim \mathbf{x}^2 \sim \cdots \sim \mathbf{x}^k$, then for all $\mathcal{S} \subset \mathcal{T}$, $\quad \mathbb{P}\big(\mathcal{A}(\mathbf{x}) \in S\big) \leq e^{k\varepsilon}\, \mathbb{P}\big(\mathcal{A}(\mathbf{x}^k) \in S\big)$

"Composition"

If $\mathcal{A}_1 : \mathcal{X}^n \to \mathfrak{M}_1(\mathcal{T})$ is ε-DP and if $\mathcal{A}_2 : \mathcal{X}^n \to \mathfrak{M}_1(\mathcal{T}')$ is ε'-DP, then $\mathbf{x} \mapsto \big(\mathcal{A}_1(\mathbf{x}), \mathcal{A}_2(\mathbf{x})\big)$ is $(\varepsilon + \varepsilon')$-DP

DP defines privacy not as a binary notion of "was the data of individual exposed or not", but rather a matter of accumulative risk

Outline

Example: Majority of Binary Observations

$\mathcal{X} = \{0,1\}$, $n = 2k+1$ target $f(\mathbf{x}) = \mathbb{1}\{\sum x_i \geq n/2\} = \text{median}(x)$

- $\mathcal{A}(\mathbf{x})$ depends only on $s = \sum x_i \implies \mathbb{P}^U(\mathcal{A}(\mathbf{x}) = 1) =: p(s)$
- By symmetry $p(n-s) = 1 - p(s)$
- DP: $p(k+1) \leq e^{\varepsilon} p(k) = e^{\varepsilon}(1 - p(k+1)) \implies p(k+1) \leq \frac{1}{1+e^{-\varepsilon}}$
- More generally, for all $s > n/2$, $p(s) \leq \dfrac{1}{1 + e^{-(2s-n)\varepsilon}}$ and $p(n) \leq \dfrac{1}{1+e^{-(n+2)\varepsilon}}$
- In fact, $p(s) = \dfrac{1}{1 + e^{-(2s-n)\varepsilon/2}}$ is ε-DP (see next slide)

  Better: $p(k+r) = \dfrac{1}{1 + e^{-r\varepsilon}}$ is ε-DP: $\frac{p(k+r+1)}{p(k+r)} = e^{\varepsilon} \frac{1+e^{-r\varepsilon}}{e^{\varepsilon}+e^{-r\varepsilon}} \leq e^{\varepsilon}$

  and similarly for $\dfrac{p(k+1)}{p(k)}$ and $\dfrac{1 - p(k+r+1)}{1 - p(k+r)}$

Example: Majority of Binary Observations

$\mathcal{X} = \{0,1\}$, $n = 2k+1$ target $f(\mathbf{x}) = \mathbb{1}\{\sum x_i \geq n/2\} = \text{median}(x)$

- $\mathcal{A}(\mathbf{x})$ depends only on $s = \sum x_i \implies \mathbb{P}^U(\mathcal{A}(\mathbf{x}) = 1) =: p(s)$
- By symmetry $p(n-s) = 1 - p(s)$
- DP: $p(k+1) \leq e^\varepsilon p(k) = e^\varepsilon(1 - p(k+1)) \implies p(k+1) \leq \frac{1}{1+e^{-\varepsilon}}$
- More generally, for all $s > n/2$, $p(s) \leq \dfrac{1}{1 + e^{-(2s-n)\varepsilon}}$ and $p(n) \leq \dfrac{1}{1+e^{-(n+2)\varepsilon}}$
- In fact, $p(s) = \dfrac{1}{1 + e^{-(2s-n)\varepsilon/2}}$ *is* ε-DP (see next slide)
- Requires $n \gg 1/\varepsilon$
- If $|s - n/2| \geq 3/\varepsilon$, the answer is correct with probability $\geq 95\%$
- But if $|s - n/2| \leq \sqrt{n}$, the chances are high that the majority in the sample is not the majority in the population
- $\implies$ if $\varepsilon \geq 3/\sqrt{n} \iff n \geq 9/\varepsilon^2$, ε-DP does not really cost any precision!

More generally: Exponential Mechanism

If $\mathcal{T}$ is discrete, one wants $\mathcal{A}$ to assign a probability to each possible outcome $t \in \mathcal{T}$ that depends on its utility $u(\mathbf{x}, t)$ on the data $\mathbf{x}$

The sensibility of u is defined as $\Delta u = \max\limits_{t \in \mathcal{T}} \max\limits_{\mathbf{x} \sim \mathbf{x}'} \left| u(\mathbf{x}, t) - u(\mathbf{x}', t) \right|$

Exponential Mechanism

The algorithm $\mathcal{A}$ defined by $\mathbb{P}^U\left(\mathcal{A}(\mathbf{x}) = t\right) = \dfrac{\exp\left(\frac{\varepsilon u(\mathbf{x},t)}{2\Delta u}\right)}{\sum_{t' \in \mathcal{T}} \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)}$ is ε-DP

Previous example: for $u(\mathbf{x}, t) = (2t - 1)\left(s - \dfrac{n}{2}\right) = -u(\mathbf{x}, 1 - t)$,

$$\mathbb{P}^U\left(\mathcal{A}(x) = 1\right) = \frac{\exp\left(\frac{\left(s - \frac{n}{2}\right)\varepsilon}{2}\right)}{\exp\left(\frac{\left(s - \frac{n}{2}\right)\varepsilon}{2}\right) + \exp\left(-\frac{\left(s - \frac{n}{2}\right)\varepsilon}{2}\right)} = \frac{1}{1 + \exp\left(-\left(s - \frac{n}{2}\right)\varepsilon\right)}$$

Proof

For every $t \in \mathcal{T}$ and $\mathbf{x} \sim \mathbf{x}'$,

$$
\begin{aligned}
\frac{\mathbb{P}^U(\mathcal{A}(\mathbf{x}) = t)}{\mathbb{P}^U(\mathcal{A}(\mathbf{x}') = t)} &= \frac{\exp\left(\frac{\varepsilon u(\mathbf{x},t)}{2\Delta u}\right)}{\sum_{t' \in \mathcal{T}} \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)} \Bigg/ \frac{\exp\left(\frac{\varepsilon u(\mathbf{x}',t)}{2\Delta u}\right)}{\sum_{t' \in \mathcal{T}} \exp\left(\frac{u(\mathbf{x}',t')\varepsilon}{2\Delta u}\right)} \\
&= \exp\left(\frac{\varepsilon\left(u(\mathbf{x},t) - u(\mathbf{x}',t)\right)}{2\Delta u}\right) \frac{\sum_{t' \in \mathcal{T}} \exp\left(\frac{\left(u(\mathbf{x}',t') - u(\mathbf{x},t)\right)\varepsilon}{2\Delta u}\right) \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)}{\sum_{t' \in \mathcal{T}} \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)} \\
&\leq \exp\left(\frac{\varepsilon}{2}\right) \frac{\sum_{t' \in \mathcal{T}} \exp\left(\frac{\varepsilon}{2}\right) \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)}{\sum_{t' \in \mathcal{T}} \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right)} = \exp(\varepsilon)
\end{aligned}
$$

Bound on the resulting utility

Theorem

For every database $\mathbf{x}$ the exponential mechanism satisfies:

$$\mathbb{P}^U\left(u\big(\mathbf{x},\mathcal{A}(\mathbf{x})\big) \leq u\big(\mathbf{x},f(\mathbf{x})\big) - \frac{2\Delta u}{\varepsilon}\ln\frac{|\mathcal{T}|}{\delta}\right) \leq \delta$$

Proof: for any $t \in \mathcal{T}$ such that $u\big(\mathbf{x},t\big) \leq u\big(\mathbf{x},f(\mathbf{x})\big) - 2\Delta u\varepsilon^{-1}\ln\left(|\mathcal{T}|/\delta\right)$,

$$\mathbb{P}^U\left(\mathcal{A}(\mathbf{x}) = t\right) \leq \frac{\exp\left(\frac{\varepsilon\left(u\left(\mathbf{x},f(\mathbf{x})\right) - \frac{2\Delta u}{\varepsilon}\ln\frac{|\mathcal{T}|}{\delta}\right)}{2\Delta u}\right)}{\exp\left(\frac{\varepsilon u(\mathbf{x},f(\mathbf{x}))}{2\Delta u}\right)} = \frac{\delta}{|\mathcal{T}|}$$

and there are at most $\mathcal{T}$ of them

Bound on the resulting utility

Theorem

For every database $\mathbf{x}$ the exponential mechanism satisfies:

$$\mathbb{P}^U\left(u\big(\mathbf{x}, \mathcal{A}(\mathbf{x})\big) \leq u\big(\mathbf{x}, f(\mathbf{x})\big) - \frac{2\Delta u}{\varepsilon}\ln\frac{|\mathcal{T}|}{\delta}\right) \leq \delta$$

Equivalently, for every $v > 0$

$$\mathbb{P}^U\left(u\big(x, \mathcal{A}(\mathbf{x})\big) \leq u\big(x, f(\mathbf{x})\big) - v\right) \leq |\mathcal{T}|e^{-\frac{\varepsilon v}{2\Delta_u}}$$

In the example:

$$\mathbb{P}^U\left(\mathcal{A}(\mathbf{x}) \neq f(x)\right) = \mathbb{P}^U\Big(u\big(x, \mathcal{A}(\mathbf{x})\big) \leq u\big(\mathbf{x}, f(\mathbf{x})\big) - 2u\big(\mathbf{x}, f(\mathbf{x})\big)\Big) \leq 2\exp\left(-\frac{u\big(\mathbf{x}, f(x)\big)\varepsilon}{\Delta u}\right) = 2e^{-\left|s-\frac{n}{2}\right|\varepsilon}$$

DB Lower Bound for Discrete Mechanisms

$\mathcal{A}$ (discrete) is said to be unbiased if for all $\mathbf{x} \in \mathcal{X}$ and $t \in \mathcal{T}$,

$$\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = t\big) \leq \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)$$

Inverse Sensibility

The inverse sensibility of function f on data $\mathbf{x}$ to output $t \in \mathcal{T}$ is defined as

$$\mathcal{D}_f(\mathbf{x}, t) = \min\left\{k : \exists \mathbf{x} \sim \mathbf{x}^1 \sim \ldots \sim \mathbf{x}^k \text{ and } f(\mathbf{x}^k) = t\right\}$$

Lower bound

For every unbiased ε-DP mechanism $\mathcal{A}$, $\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big) \leq \dfrac{1}{\sum_{t \in \mathcal{T}} e^{-2\mathcal{D}_f(\mathbf{x},t)\varepsilon}}$

Proof

Let $t \in \mathcal{T}$ and let $\mathbf{x}'$ be such that $h(\mathbf{x}, \mathbf{x}') \triangleq \sum_i \mathbb{1}\{x_1 \neq x'_i\} = D_f(\mathbf{x}, t)$ and $f(\mathbf{x}') = t$. DP implies

$$\frac{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = t\big)}{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)} = \frac{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = t\big)}{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') = t\big)} \frac{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') = t = f(\mathbf{x}')\big)}{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') = f(\mathbf{x})\big)} \frac{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}') = f(\mathbf{x})\big)}{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)} \geq e^{-\mathcal{D}_f(\mathbf{x},t)\varepsilon} \times 1 \times e^{-\mathcal{D}_f(\mathbf{x},t)\varepsilon}$$

and hence

$$1 = \sum_{t \in \mathcal{T}} \frac{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = t\big)}{\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)} \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big) \geq \sum_{t \in \mathcal{T}} e^{-2\mathcal{D}_f(\mathbf{x},t)\varepsilon} \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)$$

In the previous example, $\mathcal{D}_f(\mathbf{x}, 1 - f(\mathbf{x})) = \left|s - \frac{n}{2}\right| + \frac{1}{2}$ and this yields:

$$\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big) \leq \frac{1}{1 + e^{-\big(|2s-n|+1\big)\varepsilon}}$$

The Exponential Mechanism above is almost optimal: it has $\mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big) = \dfrac{1}{1 + e^{-\frac{|2s-n|\varepsilon}{2}}}$

ENS DE LYON

Inverse Sensitivity Mechanism

Inverse sensibility $\mathcal{D}_f$ = good candidate utility function for an exponential mechanism! $\Delta\mathcal{D}_f = 1 \implies \varepsilon$-DP Inverse Sensitivity Mechanism (ISM)

$$\mathbb{P}^U\left(\mathcal{A}(\mathbf{x}) = t\right) = \frac{e^{-\varepsilon\mathcal{D}_f(\mathbf{x},t)/2}}{\sum_{t'\in\mathcal{T}} e^{-\mathcal{D}_f(\mathbf{x},t')\varepsilon/2}}$$

Remark: if $|\mathcal{T}| = 2$ the denominator 2 is not needed:

$$\mathbb{P}^U\left(\mathcal{A}(\mathbf{x}) = t\right) = \frac{e^{-\varepsilon\mathcal{D}_f(\mathbf{x},t)}}{\sum_{t'\in\mathcal{T}} e^{-\mathcal{D}_f(\mathbf{x},t')\varepsilon}}$$

is ε-DP

Previous example: here $\mathcal{D}_f(\mathbf{x}, 1 - f(\mathbf{x})) = \left|s - \frac{n}{2}\right| + \frac{1}{2}$, the ISM $p(s) = \frac{1}{1 + e^{-\left(s-k-\mathbb{1}\{s\leq k\}\right)\varepsilon}}$ is an $\varepsilon-$DP mechanism slightly better than the exponential mechanism above

UNIVERSITÉ DE LYON

ENS DE LYON

DB Near-Optimality of the Inverse Senbility Mechanism

$1/4$-Optimality of the ISM

The ISM $\mathcal{A}$ is "more accurate" than any $\varepsilon/4$-DP algorithm $\mathcal{A}'$:

$$\mathbb{P}^U\big(\mathcal{A}'(\mathbf{x}) = f(\mathbf{x})\big) \leq \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big)$$

Proof: Since $\mathcal{D}_f(\mathbf{x}, f(\mathbf{x})) = 0$, $\quad \mathbb{P}^U\big(\mathcal{A}(\mathbf{x}) = f(\mathbf{x})\big) = 1/\sum_{t\in\mathcal{T}} e^{-\mathcal{D}_f(\mathbf{x},t)\varepsilon/2}$

Recall the lower bound: for every unbiased ε-DP mechanism $\mathcal{A}'$, $\quad \mathbb{P}^U\big(\mathcal{A}'(\mathbf{x}) = f(\mathbf{x})\big) \leq 1/\sum_{t\in\mathcal{T}} e^{-2\mathcal{D}_f(\mathbf{x},t)\varepsilon}$

Remark: if $|\mathcal{T}| = 2$, the ISM is $1/2$-optimal

UNIVERSITÉ DE LYON

ENS DE LYON

Continuous Exponential Mechanism

For a continuous $\mathcal{T}$, taking $\mathcal{A}(\mathbf{x})$ with density

$$f_{\mathbf{x}}(t) = \frac{\exp\left(\frac{\varepsilon u(\mathbf{x},t)}{2\Delta u}\right)}{\int_{\mathcal{T}} \exp\left(\frac{u(\mathbf{x},t')\varepsilon}{2\Delta u}\right) dt'}$$

also yields an ε-DP mechanism.

- The ISM is hence a very good candidate in theory.
- It is reminiscent of statistical physics ”Gibbs law” (thermodynamics).
- It can be hard to sample from.
- In fact, the discrete case is already computationally challenging when the output space $\mathcal{T}$ is ”big”.
- Research question: does approximate sampling preserve differential privacy?

WIP: Multi-quantile Estimation

Private Quantiles Estimation in the Presence of Atoms

Clément Lalanne[1], Aurélien Garivier[1,2], Rémi Gribonval[1,4], Clément Gastaud[3], and Nicolas Grislain[3]

[1]LIP Laboratory, École Normale Supérieure de Lyon, Lyon, France
[2]UMPA Laboratory, École Normale Supérieure de Lyon, Lyon, France
[3]Sarus Technologies, Paris, France
[4]INRIA, France

February 3, 2022

Abstract

We address the differentially private estimation of multiple quantiles (MQ) of a dataset, a key building block in modern data analysis. We apply the recent non-smoothed Inverse Sensitivity (IS) mechanism to this specific problem and establish that the resulting method is closely related to the current state-of-the-art, the JointExp algorithm, sharing in particular the same computational complexity and a similar efficiency. However, we demonstrate both theoretically and empirically that (non-smoothed) JointExp suffers from an important lack of performance in the case of peaked distributions, with a potentially catastrophic impact in the presence of atoms. While its smoothed version would allow to leverage the performance guarantees of IS, it remains an open challenge to implement. As a proxy to fix the problem we propose a simple and numerically efficient method called Heuristically Smoothed JointExp (HSJointExp), which is endowed with performance guarantees for a broad class of distributions and achieves results that are orders of magnitude better on problematic datasets.

Introduction

As more and more data is collected on individuals and as data science techniques become more powerful, threats to privacy have multiplied and serious concerns have emerged [NS06, BDK07, FJR15, DN03, HSR⁺08, LDM10, NS08, Swe00, WE18, Swe02]. Against this background, *differential privacy* [DR⁺14] has become the *gold standard* in privacy protection. By introducing randomness at a level calibrated to the *sensitivity*

Outline

Example: estimate the mean

Here $x \in \mathfrak{M}_{n,1}(\{0,1\})$ and $f(\mathbf{x}) = \bar{\mathbf{x}}$ $\quad \mathcal{T} = \{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, 1\}$
Inverse sensibility function: $\mathcal{D}_f(\mathbf{x}, t) = n|t - \bar{\mathbf{x}}|$

ISM: $\mathbb{P}(\mathcal{A}(\mathbf{x}) = t) = \dfrac{e^{-\varepsilon n|t-\bar{\mathbf{x}}|/2}}{\sum_{j=0}^{n} e^{-\varepsilon n|j/n-\bar{\mathbf{x}}|/2}}$

$$\mathbb{P}\left(\mathcal{A}(\mathbf{x}) \geq \bar{\mathbf{x}} + \frac{k}{n}\right) = \frac{\sum_{j=k}^{n} e^{-\varepsilon j/2}}{2\sum_{j=0}^{n} e^{-\varepsilon j/2}} = \frac{e^{-\varepsilon k/2}}{2} \quad \text{and} \quad \mathbb{P}\left(\mathcal{A}(x) \leq \bar{x} - \frac{k}{n}\right) = \frac{e^{-\varepsilon k/2}}{2}$$

$\implies$ up to the discretization, $\mathcal{A}(x)$ has the same distribution as

$$\mathcal{A}'(\mathbf{x}) = \bar{\mathbf{x}} + Y$$

where $Y \sim \text{Lap}\left(\frac{2}{n\varepsilon}\right)$ has density $\ell_{\frac{2}{n\varepsilon}}(\mathbf{x}) = \dfrac{n\varepsilon}{4} e^{-\frac{n\varepsilon|x|}{2}}$ since $\mathbb{P}(\mathcal{A}'(\mathbf{x}) \geq \bar{\mathbf{x}} + k/n) = \dfrac{e^{-\varepsilon k/2}}{2}$.

$\implies$ very simple mechanism: just add (well-calibrated) Laplace noise!

UNIVERSITÉ DE LYON

ENS DE LYON

Laplace mechanism

The L^1-sensitivity of a function $f : \mathcal{X}^n \to \mathbb{R}^k$ is defined as

$$\Delta f = \max_{\mathbf{x} \sim \mathbf{x}'} \|f(\mathbf{x}) - f(\mathbf{x}')\|_1$$

Example: if $f(\mathbf{x}) = n^{-1} \sum x_i$ and $x_i \in [0, 1]$, then $\Delta(f) = 1/n$

$\mathrm{Lap}(\sigma)$ has density $\ell_\sigma(\mathbf{x}) = \dfrac{1}{2\sigma} e^{-\frac{|x|}{\sigma}}$

Laplace Mechanism

The Laplace mechanism for $f : \mathcal{X}^n \to \mathbb{R}^k$ defined by

$$\mathcal{A}(\mathbf{x}) = f(\mathbf{x}) + (Y_1, \ldots, Y_k), \qquad Y_i \overset{iid}{\sim} \mathrm{Lap}\left(\frac{\Delta f}{\varepsilon}\right)$$

is ε-differentially private

Proof

For $t \in \mathbb{R}^k$ let

$$p_{\mathbf{x}}(t) = \prod_{j=1}^{k} \frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon\left|t_j - f(\mathbf{x})_j\right|}{\Delta f}\right) = \left(\frac{\varepsilon}{2\Delta f}\right)^k \exp\left(-\frac{\varepsilon\left\|t - f(\mathbf{x})\right\|_1}{\Delta f}\right)$$

be the density of $\mathcal{A}(\mathbf{x})$. Then for every $\mathbf{x} \sim \mathbf{x}'$ and every $t \in \mathbb{R}^k$,

$$\frac{p_{\mathbf{x}}(t)}{p_{\mathbf{x}'}(t)} = \frac{\exp\left(-\frac{\varepsilon\left\|t - f(\mathbf{x})\right\|_1}{\Delta f}\right)}{\exp\left(-\frac{\varepsilon\left\|t - f(\mathbf{x}')\right\|_1}{\Delta f}\right)} = \exp\left(-\frac{\varepsilon\left(\left\|t - f(\mathbf{x})\right\|_1 - \left\|t - f(\mathbf{x}')\right\|_1\right)}{\Delta f}\right)$$

$$\leq \exp\left(\frac{\varepsilon\left\|f(\mathbf{x}) - f(\mathbf{x}')\right\|_1}{\Delta f}\right) \leq \exp(\varepsilon)$$

since by definition $\|f(\mathbf{x}) - f(\mathbf{x}')\|_1 \leq \Delta f$

Examples

- Estimate averages (or sum)
- Counting queries   accuracy independent of *n*
- Histogram queries   accuracy independent of *n and of the number of bins!*
- Most common value of a variable
- Noisy max   accuracy depends on the number of values, but not on *n*
- Non-parametrics (estimate coefficients in trucated basis and add Laplace noise) [see Wasserman&Zhou '08]

⚠ **Finite precision arithmetic $\Longrightarrow$ possible privacy leaks**

For example, one can have $\mathbb{P}\big(\mathcal{A}(\mathbf{x}) = t\big) > 0$ while $\mathbb{P}\big(\mathcal{A}(\mathbf{x}') = t\big) = 0$ for two neighbors $\mathbf{x} \sim \mathbf{x}'$ because of rounding.

DB lower bound: continuous mechanisms

We consider a target function $f : \mathcal{X}^n \to \mathbb{R}$. For every $\mathbf{x} \in \mathcal{X}^n$, let $\Delta_f(\mathbf{x}, k) = \sup \left\{ |f(\mathbf{x}') - f(\mathbf{x})| : h(\mathbf{x}, \mathbf{x}') \leq k \right\}$ and $\Delta_f(\mathbf{x}) = \Delta_f(\mathbf{x}, 1)$

Lower bound

$\mathcal{A}$ is unbiased if $\mathbb{E}^U[\mathcal{A}(\mathbf{x})] = f(\mathbf{x})$ for every $\mathbf{x}$. Then, for every $\mathbf{x} \in \mathcal{X}^n$,

$$\mathbb{E}^U \left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x}) \right)^2 \right] \geq \sup_k \frac{\Delta_f(\mathbf{x}, k)^2 / 4}{1 + e^{2k\varepsilon}}$$

and in particular if $\Delta_f(\mathbf{x}, k) = k\Delta_f(\mathbf{x})$, then for $\varepsilon \leq 1/2$

$$\mathbb{E}^U \left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x}) \right)^2 \right] \geq \frac{\Delta_f(\mathbf{x})^2}{68\varepsilon^2}$$

For the average $f(\mathbf{x}) = \bar{\mathbf{x}}_n$, $\Delta_f(\mathbf{x}, k) = k\Delta_f(\mathbf{x}) = k/n$ and $\mathbb{E}^U \left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x}) \right)^2 \right] \geq \frac{1}{68n^2\varepsilon^2}$

Proof

Let $k \geq 1$ and let $\mathbf{x}'$ be such that $h(\mathbf{x}, \mathbf{x}') = k$ and $|f(\mathbf{x}') - f(\mathbf{x})| = \Delta_f(\mathbf{x}, k)$. By definition,

$$\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}') - t\right)^2\right] = \int_{\mathcal{T}} (s-t)^2 d\mathbb{P}^U_{\mathcal{A}(\mathbf{x}')}(s) \leq \int_{\mathcal{T}} (s-t)^2 e^{k\varepsilon} d\mathbb{P}^U_{\mathcal{A}(\mathbf{x})}(s) = e^{k\varepsilon}\ \mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - t\right)^2\right]$$

and hence

$$\frac{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}') - f(\mathbf{x}')\right)^2\right]}{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x})\right)^2\right]} = \frac{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}') - f(\mathbf{x}')\right)^2\right]}{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}') - f(\mathbf{x})\right)^2\right]} \frac{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}') - f(\mathbf{x})\right)^2\right]}{\mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x})\right)^2\right]} \leq 1 \times e^{k\varepsilon}$$

since $\mathcal{A}$ is unbiased. Therefore, by the Bienaymé-Chebishev inequality

$$\mathbb{P}^U\left(|\mathcal{A}(\mathbf{x}) - f(\mathbf{x})| \geq \frac{\Delta_f(\mathbf{x}, k)}{2}\right) \leq \frac{4\ \mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x})\right)^2\right]}{\Delta_f(\mathbf{x}, k)^2} \text{ and}$$

$$\mathbb{P}^U\left(|\mathcal{A}(\mathbf{x}) - f(\mathbf{x}')| \geq \frac{\Delta_f(\mathbf{x}, k)}{2}\right) \leq e^{k\varepsilon}\mathbb{P}^U\left(|\mathcal{A}(\mathbf{x}') - f(\mathbf{x}')| \geq \frac{\Delta_f(\mathbf{x}, k)}{2}\right) \leq \frac{4e^{2k\varepsilon}\ \mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x})\right)^2\right]}{\Delta_f(\mathbf{x}, k)^2}.$$

But since $|f(\mathbf{x}') - f(\mathbf{x})| = \Delta_f(\mathbf{x}, k)$,

$$1 \leq \mathbb{P}^U\left(|\mathcal{A}(\mathbf{x}) - f(\mathbf{x})| \geq \frac{\Delta_f(\mathbf{x}, k)}{2}\right) + \mathbb{P}^U\left(|\mathcal{A}(\mathbf{x}) - f(\mathbf{x}')| \geq \frac{\Delta_f(\mathbf{x}, k)}{2}\right) \leq \frac{4\left(1 + e^{2k\varepsilon}\right)\ \mathbb{E}^U\left[\left(\mathcal{A}(\mathbf{x}) - f(\mathbf{x}\right.\right.}{\Delta_f(\mathbf{x}, k)^2}$$

The second statement is obtained by the choice $k = \lceil 1/(2\varepsilon) \rceil$, noting that $1/(2\varepsilon) \leq k \leq 1/\varepsilon$

Outline

Le Cam's argument for Minimax Risk

For all $\psi : \mathcal{X}^n \times [0,1] \to \mathcal{Y}$, $\theta_1, \theta_2 \in \Theta$, let $S = \{t \in \theta : d(t, \theta_1) > d(t, \theta_2)\}$. Then

$$
\begin{aligned}
\max_{\theta \in \Theta} \mathbb{E}_\theta^U \left[d\left(\psi(\mathbf{X}, U), \theta\right)\right] &\geq \max \left\{\mathbb{E}_{\theta_1}^U \left[d\left(\psi(\mathbf{X}, U), \theta_1\right), \mathbb{E}_{\theta_2}^U \left[d\left(\psi(\mathbf{X}, U), \theta_2\right)\right]\right\} \\
&\geq \frac{1}{2}\left(\mathbb{E}_{\theta_1}^U \left[d\left(\psi(\mathbf{X}, U), \theta_1\right)\right] + \mathbb{E}_{\theta_2}^U \left[d\left(\psi(\mathbf{X}, U), \theta_2\right)\right]\right) \\
&\geq \frac{1}{2}\left(\mathbb{E}_{\theta_1}^U \left[d\left(\psi(\mathbf{X}, U), \theta_1\right)\mathbb{1}\left\{\psi(\mathbf{X}, U) \in S\right\}\right] + \mathbb{E}_{\theta_2}^U \left[d\left(\psi(\mathbf{X}, U), \theta_2\right)\mathbb{1}\left\{\psi(\mathbf{X}, U) \in \bar{S}\right\}\right]\right) \\
&\geq \frac{d(\theta_1, \theta_2)}{2 \times 2}\left(\mathbb{P}_{\theta_1}^U \left(\psi(\mathbf{X}, U) \in S\right) + \mathbb{P}_{\theta_2}^U \left(\psi(\mathbf{X}, U) \in \bar{S}\right)\right) \\
&= \frac{d(\theta_1, \theta_2)}{4}\mathbb{E}_{\theta_1, \theta_2}\left[Q_{\mathbf{X}_1}(S) + Q_{\mathbf{X}_2}(\bar{S}) \middle| \mathbf{X}_1, \mathbf{X}_2\right] \\
&\geq \frac{d(\theta_1, \theta_2)}{4}\mathbb{E}_{\theta_1, \theta_2}\left[\mathbb{1}\{\mathbf{X}_1 = \mathbf{X}_2\} \middle| \mathbf{X}_1, \mathbf{X}_2\right] \\
&= \frac{d(\theta_1, \theta_2)}{4}\left(1 - \mathrm{TV}\left(P_{\theta_1}^{\otimes n}, P_{\theta_2}^{\otimes n}\right)\right) \quad \text{by the coupling lemma} \\
&\geq \frac{d(\theta_1, \theta_2)}{4}\,\frac{e^{-\mathrm{KL}\left(P_{\theta_1}^{\otimes n}, P_{\theta_2}^{\otimes n}\right)}}{2} = \frac{d(\theta_1, \theta_2)}{8}e^{-n\,\mathrm{KL}\left(P_{\theta_1}, P_{\theta_2}\right)}
\end{aligned}
$$

Le Cam's argument for Minimax Risk

For all $\psi : \mathcal{X}^n \times [0,1] \to \mathcal{Y}$, $\theta_1, \theta_2 \in \Theta$,

$$\max_{\theta \in \Theta} \mathbb{E}_\theta^U \left[d\big(\psi(\mathbf{X}, U), \theta\big) \right] \geq \frac{d(\theta_1, \theta_2)}{8} e^{-n \,\mathrm{KL}\left(p_{\theta_1}, p_{\theta_2}\right)}$$

Hence,

Le Cam's bound

$$\min_{\psi} \max_{\theta \in \Theta} \mathbb{E}_\theta^U \left[d\big(\psi(\mathbf{X}, U), \theta\big) \right] \geq \max_{\mathrm{KL}\left(p_{\theta_1}, p_{\theta_2}\right) \leq \frac{1}{n}} \frac{d(\theta_1, \theta_2)}{8e} \overset{\text{typ.}}{=} \frac{C}{\sqrt{n}}$$

Tight in simple, low-dimensional models

Le Cam's argument for Minimax DP Risk

For all $\psi : \mathcal{X}^n \times [0,1] \to \mathcal{Y}$, $\theta_1, \theta_2 \in \Theta$, let $S = \{t \in \theta : d(t, \theta_1) > d(t, \theta_2)\}$. Then

$$
\begin{aligned}
\max_{\theta \in \Theta} \mathbb{E}^U_\theta \left[d(\psi(\mathbf{X}, U), \theta) \right] &\geq \max \left\{ \mathbb{E}^U_{\theta_1} \left[d(\psi(\mathbf{X}, U), \theta_1) \right], \mathbb{E}^U_{\theta_2} \left[d(\psi(\mathbf{X}, U), \theta_2) \right] \right\} \\
&\geq \frac{1}{2} \left(\mathbb{E}^U_{\theta_1} \left[d(\psi(\mathbf{X}, U), \theta_1) \right] + \mathbb{E}^U_{\theta_2} \left[d(\psi(\mathbf{X}, U), \theta_2) \right] \right) \\
&\geq \frac{1}{2} \left(\mathbb{E}^U_{\theta_1} \left[d(\psi(\mathbf{X}, U), \theta_1) \mathbb{1}\{\psi(\mathbf{X}, U) \in S\} \right] + \mathbb{E}^U_{\theta_2} \left[d(\psi(\mathbf{X}, U), \theta_2) \mathbb{1}\{\psi(\mathbf{X}, U) \in \bar{S}\} \right] \right) \\
&\geq \frac{d(\theta_1, \theta_2)}{2 \times 2} \left(\mathbb{P}^U_{\theta_1} \left(\psi(\mathbf{X}, U) \in S \right) + \mathbb{P}^U_{\theta_2} \left(\psi(\mathbf{X}, U) \in \bar{S} \right) \right) \\
&= \frac{d(\theta_1, \theta_2)}{4} \mathbb{E}_{\theta_1, \theta_2} \left[Q_{\mathbf{X}_1}(S) + Q_{\mathbf{X}_2}(\bar{S}) \middle| \mathbf{X}_1, \mathbf{X}_2 \right] \\
&\geq \frac{d(\theta_1, \theta_2)}{4} \mathbb{E}_{\theta_1, \theta_2} \left[Q_{\mathbf{X}_2}(S) e^{-\varepsilon h(\mathbf{X}_1, \mathbf{X}_2)} + Q_{\mathbf{X}_2}(\bar{S}) \middle| \mathbf{X}_1, \mathbf{X}_2 \right] \\
&\geq \frac{d(\theta_1, \theta_2)}{4} \mathbb{E}_{\theta_1, \theta_2} \left[e^{-\varepsilon h(\mathbf{X}_1, \mathbf{X}_2)} \right] \geq \frac{d(\theta_1, \theta_2)}{4} e^{-\varepsilon \mathbb{E}_{\theta_1, \theta_2} \left[h(\mathbf{X}_1, \mathbf{X}_2) \right]} \\
&= \frac{d(\theta_1, \theta_2)}{4} e^{-n\varepsilon \, \mathrm{TV}\left(P_{\theta_1}, P_{\theta_2}\right)} \quad \text{for the product coupling st } \forall i \in [n], \mathbb{P}_{\theta_1, \theta_2}(X_{1,i} \neq X_{2,i}) = \mathrm{TV}(P_{\theta_1}, P_{\theta_2})
\end{aligned}
$$

Le Cam's argument for Minimax DP Risk

For all $\psi : \mathcal{X}^n \times [0,1] \to \mathcal{Y}$, $\theta_1, \theta_2 \in \Theta$,

$$\max_{\theta \in \Theta} \mathbb{E}_\theta^U \left[d\left(\psi(\mathbf{X}, U), \theta\right) \right] \geq \frac{d(\theta_1, \theta_2)}{4} e^{-n\varepsilon \, \mathrm{TV}\left(P_{\theta_1}, P_{\theta_2}\right)}$$

Hence,

Le Cam's private bound

$$\min_{\psi} \max_{\theta \in \Theta} \mathbb{E}_\theta^U \left[d\left(\psi(\mathbf{X}, U), \theta\right) \right] \geq \max_{\mathrm{TV}\left(P_{\theta_1}, P_{\theta_2}\right) \leq \frac{1}{n\varepsilon}} \frac{d(\theta_1, \theta_2)}{4e} \overset{\text{typ.}}{=} \frac{C'}{n\varepsilon}$$

Tight for example for the estimation of a $1/n$-sensitive function by using the Laplace mechanism
$\implies$ for $\varepsilon \gg 1/\sqrt{n}$, no cost for privacy

Extensions

$\rightarrow$ better couplings?

$\rightarrow$ beyond Le Cam: Fano (high dimension, non-parametrics)

$\rightarrow$ better distances on image laws? ρ-zero differential privacy

$\rightarrow$ Sequential statistics?

References

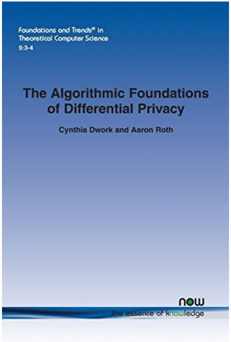

Near Instance-Optimality in Differential Privacy

Hilal Asi John C. Duchi
asi@stanford.edu jduchi@stanford.edu
Stanford University

Abstract

We develop two notions of instance optimality in differential privacy, inspired by classical statistical theory: one by defining a local minimax risk and the other by considering unbiased mechanisms and analogizing the Cramér-Rao bound, and we show that the local modulus of continuity of the estimand of interest completely determines these quantities. We also develop a complementary collection mechanisms, which we term the *inverse sensitivity* mechanisms, which are instance optimal (or nearly instance optimal) for a large class of estimands. Moreover, these mechanisms uniformly outperform the smooth sensitivity framework—on each instance—for several function classes of interest, including $\mathbb{R}$-valued continuous functions. We carefully present two instantiations of the mechanisms for median and robust regression estimation with corresponding experiments.

1 Introduction

We study instance-specific optimality for differentially private release of a function $f(\boldsymbol{x})$ of a dataset $\boldsymbol{x} \in \mathcal{X}^n$. In contrast to existing notions of optimality for private procedures, which measure mechanisms' worst case performance over all instances, we develop instance-specific notions to capture the difficulty of—and potential adaptivity of private mechanisms to—the given data x, rather than some potential worst case.

The trajectory of differential privacy research and private mechanisms reflects the desire to be adaptive both to the function f to be computed and dataset $\boldsymbol{x}$ at hand. Dwork et al.'s original perspective [21] targets the former, privatizing $f(x)$ by adding noise commensurate with the global sensitivity $\mathrm{GS}_f := \sup_{\boldsymbol{x},\boldsymbol{x}':d_{\mathrm{ham}}(\boldsymbol{x},\boldsymbol{x}')\le 1} |f(\boldsymbol{x}) - f(\boldsymbol{x}')|$ of f and adapting to the function f at hand; more recent work expands this function adaptive approach [6]. As the classical approach can be conservative—it does not reflect the sensitivity of the underlying data $\boldsymbol{x}$—a natural idea is to add noise that scales with the *local sensitivity* (or local modulus of continuity) $\mathrm{LS}_f(\boldsymbol{x}) := \sup_{\boldsymbol{x}':d_{\mathrm{ham}}(\boldsymbol{x},\boldsymbol{x}')\le 1} |f(\boldsymbol{x}) - f(\boldsymbol{x}')|$ of f at the dataset $\boldsymbol{x}$. Unfortunately, this fails to protect privacy, as the sensitivity itself may be compromising, leading Nissim et al. [35] to propose mechanisms that rely on smooth upper bounds to the local sensitivity. Yet these mechanisms are complex and, as our results show, may be conservative.

To understand these phenomena, we take a two-pronged approach, presenting both lower bounds on error and complementary (near) optimal mechanisms. We first consider the desiderata a lower bound should satisfy, following a program Cai and Low [11] develop (see also [16]):

A Statistical Framework for Differential Privacy

Larry Wasserman[*‡] Shuheng Zhou[†]

[*]Department of Statistics
[‡]Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213

[†]Seminar für Statistik
ETH Zürich, CH 8092

October 22, 2018

One goal of statistical privacy research is to construct a data release mechanism that protects individual privacy while preserving information content. An example is a *random mechanism* that takes an input database X and outputs a random database Z according to a distribution $Q_n(\cdot|X)$. *Differential privacy* is a particular privacy requirement developed by computer scientists in which $Q_n(\cdot|X)$ is required to be insensitive to changes in one data point in X. This makes it difficult to infer from Z whether a given individual is in the original database X. We consider differential privacy from a statistical perspective. We consider several data release mechanisms that satisfy the differential privacy requirement. We show that it is useful to compare these schemes by computing the rate of convergence of distributions and densities constructed from the released data. We study a general privacy method, called the exponential mechanism, introduced by McSherry and Talwar (2007). We show that the accuracy of this method is intimately linked to the rate at which the probability that the empirical distribution concentrates in a small ball around the true distribution.

1 Introduction

One goal of data privacy research is to derive a mechanism that takes an input database X and releases a transformed database Z such that individual privacy is protected yet information content is preserved. This is known as disclosure limitation. In this paper we will consider various methods

+ our contributions (here and to come):

- *Private Quantiles Estimation in the Presence of Atoms*. Clément Lalanne, Clément Gastaud, Nicolas Grislain, Aurélien Garivier, Rémi Gribonval.
- *On the Statistical Complexity of Estimation and Testing under Privacy Constraints*. Clément Lalanne, Aurélien Garivier, Rémi Gribonval
- *On Private Bandits*. Aymen Al Marjani, Aurélien Garivier, Emilie Kaufmann

The end


THANK YOU
GRACIAS
ARIGATO
SHUKURIA
DANKSCHEEN
TASHAKKUR ATU
YAQHANYELAY
SUKSAMA
EKHMET
GRAZIE
MEHRBANI
BOLZÏN
MERCI
BİYAN
SHUKRIA
TINGKI
GOZAIMASHITA
EFCHARISTO
JUSPAXAR
KOMAPSUMNIDA
MAAKE
PALDIES


Journée Aléatoire 2022
Commune à la SFdS, la SMAI et la SMF
29 septembre 2022, Institut Henri Poincaré, Paris

They already adopted DP

WIKIPEDIA
The Free Encyclopedia

Several uses of differential privacy in practice are known to date:

2008 U.S. Census Bureau, for showing commuting patterns.

2014 Google's RAPPOR, for telemetry such as learning statistics about unwanted software hijacking users' settings.

2015 Google, for sharing historical traffic statistics.

2016 Apple announced its intention to use differential privacy in iOS 10 to improve its Intelligent personal assistant technology.

2017 Microsoft, for telemetry in Windows.

2019 Privitar Lens is an API using differential privacy.

2019 Sarus provides ML with DP as a service.

2020 LinkedIn, for advertiser queries.

Abstraction

Computer = machine able to make a few elementary operations on data, that can be combined arbitrarily

Problem = description of the desired output for any given input

Input $\longrightarrow$  $\longrightarrow$ Output

Examples:

[3,2,5,1,4] $\longrightarrow$ [1,2,3,4,5]

le petit chat $\longrightarrow$ the little cat

 $\longrightarrow$ 5

 $\longrightarrow$ 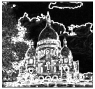

Two approaches

Classical approach : reduction = describe the sequence of elementary operations that permit to construct the output from the input
= computer programming (coding)

Artificial Intelligence = use a computer to *build itself* the program that will solve the task
= meta-programming

Machine Learning = feed the computer only with a (large number of) *example pairs* (input, output)
= search for the program that is supposed to work best on *new examples*

$\rightarrow$ for each problem, both approaches are possible
$\rightarrow$ but they are more or less efficient...

(Cliché) example: MNIST dataset

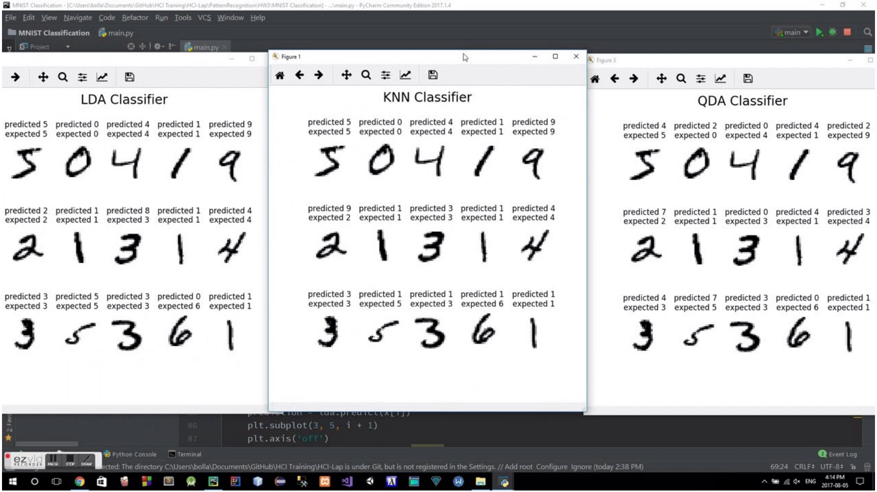

UNIVERSITÉ DE LYON

ENS DE LYON

Data in Machine Learning

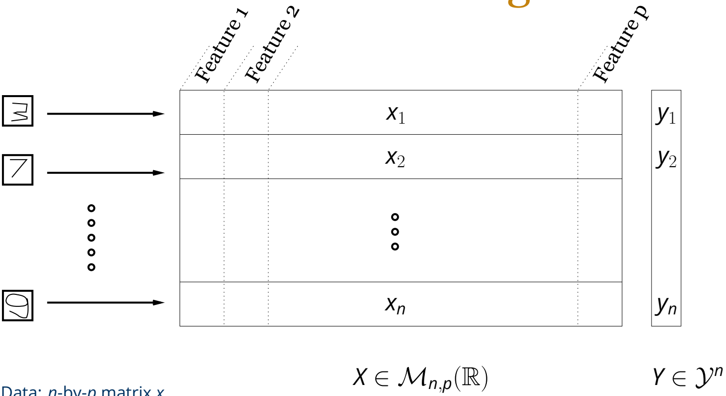

Data: n-by-p matrix x

- n examples = points of observations
- p features = characteristics measured for each example

Classifier $\mathcal{A}_n$

$h_n : \mathcal{X} \to \mathcal{Y}$

 $\mapsto$ 6