A Statistical Viewpoint on Privacy: From Hypothesis Testing to Blackwell's Theorem

Weijie Su

University of Pennsylvania

Big Brother is watching you! [*1984*, George Orwell]

Does anonymization preserve privacy?

The Netflix competition

Does anonymization preserve privacy?

The Netflix competition

- In 2006, Narayanan and Shmatikov demonstrated that

$$\text{Netflix ratings} + \text{IMDb} = \text{De-anonymization!}$$

Does anonymization preserve privacy?

The Netflix competition

- In 2006, Narayanan and Shmatikov demonstrated that

$$\text{Netflix ratings} + \text{IMDb} = \text{De-anonymization!}$$

- The second Netflix competition was canceled

Releasing summary statistics?

Genomic research often releases minor allele frequencies (MAFs), i.e., sample mean

In 2008, Homer et al shocked the genetics community by showing that MAFs are *not* private

OPEN ACCESS Freely available online PLoS GENETICS

Resolving Individuals Contributing Trace Amounts of DNA to Highly Complex Mixtures Using High-Density SNP Genotyping Microarrays

Nils Homer[1,2], Szabolcs Szelinger[1], Margot Redman[1], David Duggan[1], Waibhav Tembe[1], Jill Muehling[1], John V. Pearson[1], Dietrich A. Stephan[1], Stanley F. Nelson[2], David W. Craig[1*]

1 Translational Genomics Research Institute (TGen), Phoenix, Arizona, United States of America, 2 University of California Los Angeles, Los Angeles, California, United States of America

Is this our future?

Can we give up privacy? [WSJ '13]

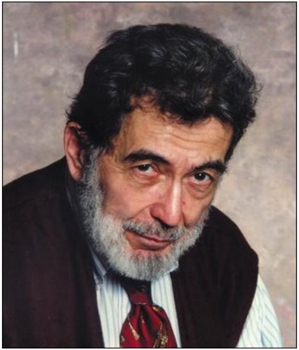

Peggy Noonan: *A loss of privacy is a loss of something personal and intimate*

Nat Hentoff: *Privacy is an American constitutional liberty right*

From hypothesis testing to privacy

In 2006, Dwork, McSherry, Nissim, and Smith introduced **differential privacy**

From hypothesis testing to privacy

In 2006, Dwork, McSherry, Nissim, and Smith introduced **differential privacy**

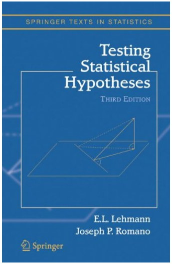

In 2010, Wasserman and Zhou related it to **hypothesis testing**

- Hypothesis testing serves as a convenient tool

From hypothesis testing to privacy

In 2006, Dwork, McSherry, Nissim, and Smith introduced **differential privacy**

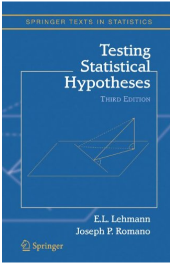

In 2010, Wasserman and Zhou related it to **hypothesis testing**

- Hypothesis testing serves as a convenient tool
- However, is it the *optimal* language for reasoning about differential privacy?

The intuition behind differential privacy

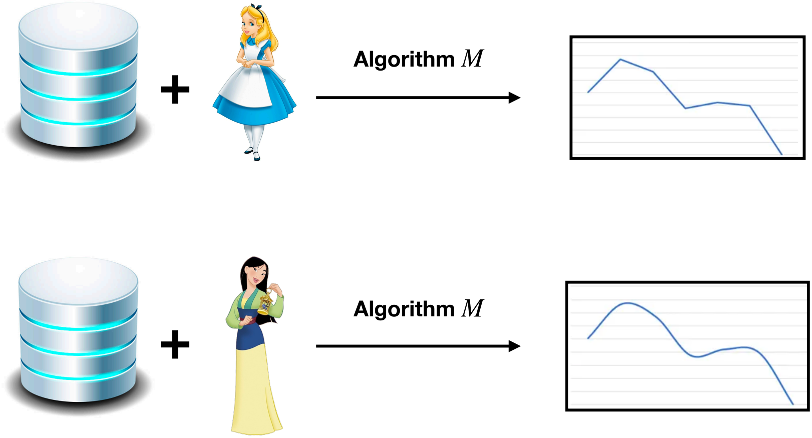

The intuition behind differential privacy

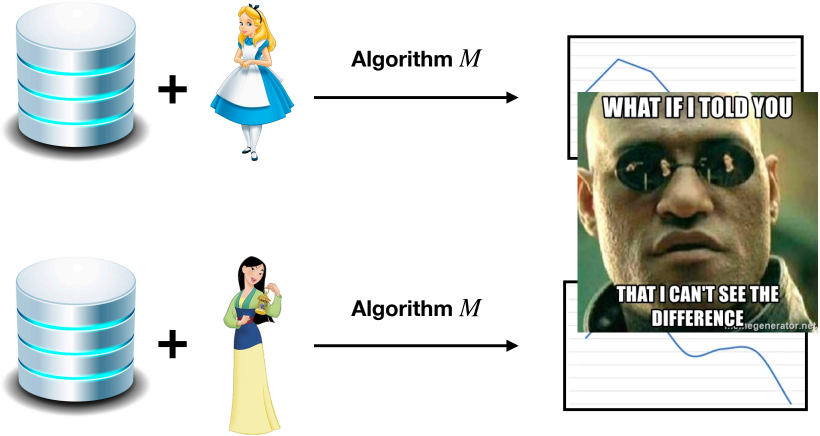

Setup for differential privacy

An example of a dataset S

	Gender	Age	Salary
Alice	F	25	\$75,000
Bob	M	20	\$45,000
Charlie	M	30	\$50,000
Dave	M	35	\$80,000
...	...	...	...
...	...	...	...

Setup for differential privacy

An example of a dataset S

	Gender	Age	Salary
Alice	F	25	\$75,000
Bob	M	20	\$45,000
Charlie	M	30	\$50,000
Dave	M	35	\$80,000
...	...	...	...
...	...	...	...

An example of a mechanism/algorithm

$$M(S) = \text{Average Salary} + \text{noise}$$

Interpreting differential privacy via hypothesis testing

Two *neighboring* datasets

$$S = \{\text{Alice}, \text{Bob}, \text{Charlie}, \text{Dave}\} \quad \text{and} \quad S' = \{\text{Anne}, \text{Bob}, \text{Charlie}, \text{Dave}\}$$

Based on output of algorithm M, perform hypothesis testing

$$H_0: \text{ true dataset is } S \qquad \mathbf{vs} \qquad H_1: \text{ true dataset is } S'$$

Interpreting differential privacy via hypothesis testing

Two *neighboring* datasets

$$S = \{\text{Alice}, \text{Bob}, \text{Charlie}, \text{Dave}\} \quad \text{and} \quad S' = \{\text{Anne}, \text{Bob}, \text{Charlie}, \text{Dave}\}$$

Based on output of algorithm M, perform hypothesis testing

$$H_0 : \text{ Alice in the dataset} \qquad \mathbf{vs} \qquad H_1 : \text{ Anne in the dataset}$$

- Intuitively, preserves privacy of Alice and Anne if hypothesis testing is difficult

Interpreting differential privacy via hypothesis testing

Two *neighboring* datasets

$$S = \{\text{Alice}, \text{Bob}, \text{Charlie}, \text{Dave}\} \quad \text{and} \quad S' = \{\text{Anne}, \text{Bob}, \text{Charlie}, \text{Dave}\}$$

Based on output of algorithm M, perform hypothesis testing

$$H_0: \text{ Alice in the dataset} \qquad \mathbf{vs} \qquad H_1: \text{ Anne in the dataset}$$

- Intuitively, preserves privacy of Alice and Anne if hypothesis testing is difficult
- Essence in *differential privacy* (DP)

The impact of differential privacy

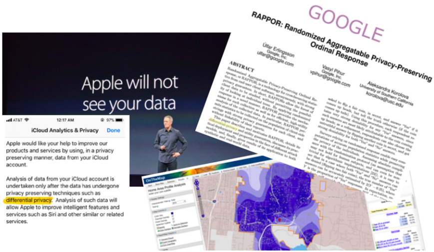

Google (Chrome), Apple (iOS 10+), Microsoft, U.S. Census Bureau [Dwork, Roth '14; Erlingsson et al '14; Apple DP team '17; Ding et al '17; Abowd '16]

Test of time: 2017 Gödel prize

The impact of differential privacy

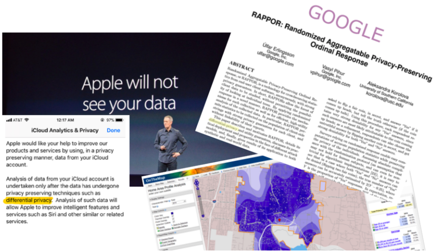

Google (Chrome), Apple (iOS 10+), Microsoft, U.S. Census Bureau [Dwork, Roth '14; Erlingsson et al '14; Apple DP team '17; Ding et al '17; Abowd '16]

Test of time: 2017 Gödel prize
Turing Award?

What's new in this talk?

A *new* privacy notion and the *old* one

f-differential privacy: **this talk**

- Interpreting privacy via hypothesis testing

(ϵ, δ)-differential privacy: **Dwork et al.**

- Interpreting privacy via hypothesis testing

A *new* privacy notion and the *old* one

f-differential privacy: **this talk**

- Interpreting privacy via hypothesis testing
- Privacy measure: type I and II errors *trade-off*

(ϵ, δ)-differential privacy: **Dwork et al.**

- Interpreting privacy via hypothesis testing
- Privacy measure: *worst-case* likelihood ratio

A *new* privacy notion and the *old* one

f-differential privacy: **this talk**

- Interpreting privacy via hypothesis testing
- Privacy measure: type I and II errors *trade-off*
- Privacy *functional* parameter: $f : [0,1] \to [0,1]$

(ϵ, δ)-differential privacy: **Dwork et al.**

- Interpreting privacy via hypothesis testing
- Privacy measure: *worst-case* likelihood ratio
- Privacy parameters: $\epsilon \geqslant 0, 0 \leqslant \delta < 1$

A *new* privacy notion and the *old* one

f-differential privacy: **this talk**

- Interpreting privacy via hypothesis testing
- Privacy measure: type I and II errors *trade-off*
- Privacy *functional* parameter: $f : [0,1] \to [0,1]$
- How to achieve: adding *Gaussian* noise

(ϵ, δ)-differential privacy: **Dwork et al.**

- Interpreting privacy via hypothesis testing
- Privacy measure: *worst-case* likelihood ratio
- Privacy parameters: $\epsilon \geqslant 0, 0 \leqslant \delta < 1$
- How to achieve: adding Laplace noise

Outline

Paper

Gaussian Differential Privacy

Journal of the Royal Statistical Society: Series B (with discussion), 2022

- Jinshuo Dong (Penn/Northwestern/Tsinghua)
- Aaron Roth (Penn)

Trade-off functions

$$H_0: \text{ true dataset is } S \qquad \mathbf{vs} \qquad H_1: \text{ true dataset is } S'$$

Trade-off functions

$$H_0 : P \qquad \textbf{vs} \qquad H_1 : Q$$

For rejection rule $\phi \in [0, 1]$, denote by

$$\begin{aligned} &\text{type I error} \quad && \alpha_\phi = \mathbb{E}_P[\phi] \\ &\text{type II error} \quad && \beta_\phi = 1 - \mathbb{E}_Q[\phi] \end{aligned}$$

Trade-off functions

$$H_0 : P \qquad \textbf{vs} \qquad H_1 : Q$$

For rejection rule $\phi \in [0, 1]$, denote by

type I error $\quad \alpha_\phi = \mathbb{E}_P[\phi]$

type II error $\quad \beta_\phi = 1 - \mathbb{E}_Q[\phi]$

Definition

For two probability distributions P and Q, define the trade-off function $T(P, Q) : [0, 1] \to [0, 1]$ as

$$T(P, Q)(\alpha) = \inf_\phi \{\beta_\phi : \alpha_\phi \leqslant \alpha\}$$

Trade-off functions

$$H_0 : P \qquad \textbf{vs} \qquad H_1 : Q$$

For rejection rule $\phi \in [0, 1]$, denote by

$$\begin{aligned} &\text{type I error} \quad && \alpha_\phi = \mathbb{E}_P[\phi] \\ &\text{type II error} \quad && \beta_\phi = 1 - \mathbb{E}_Q[\phi] \end{aligned}$$

Definition

For two probability distributions P and Q, define the trade-off function $T(P, Q) : [0, 1] \to [0, 1]$ as

$$T(P, Q)(\alpha) = \inf_\phi \{\beta_\phi : \alpha_\phi \leqslant \alpha\}$$

- Neyman–Pearson lemma
- f is trade-off if and only if f is convex, continuous, non-increasing, and $f(\alpha) \leqslant 1 - \alpha$ for $\alpha \in [0, 1]$

Definition of f-DP

Definition (DRS)

A (randomized) algorithm M is said to be f-differentially private if

$$T\big(M(S), M(S')\big) \geqslant f$$

for all neighboring datasets S and S'

- Randomness of $M(S), M(S')$ is from the algorithm M
- Telling apart Alice and Anne is *no* easier than P and Q if $f = T(P, Q)$
- Related to hypothesis testing region [Kairouz et al '17]

(ϵ, δ)-DP is a special instance of f-DP

Definition of (ϵ, δ)-DP

$$\mathrm{e}^{-\epsilon}\, \mathbb{P}(M(S') \in E) - \mathrm{e}^{-\epsilon}\delta \leqslant \mathbb{P}(M(S) \in E) \leqslant \mathrm{e}^{\epsilon}\, \mathbb{P}(M(S') \in E) + \delta$$

(ϵ, δ)-DP is a special instance of f-DP

Definition of (ϵ, δ)-DP

$$\mathrm{e}^{-\epsilon}\,\mathbb{P}(M(S') \in E) - \mathrm{e}^{-\epsilon}\delta \leqslant \mathbb{P}(M(S) \in E) \leqslant \mathrm{e}^{\epsilon}\,\mathbb{P}(M(S') \in E) + \delta$$

Adapted from [Wasserman, Zhou '10]

An algorithm M is (ϵ, δ)-DP if and only if it is $f_{\epsilon,\delta}$-DP

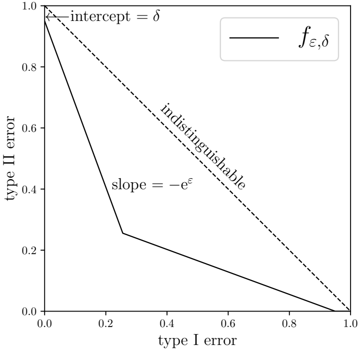

(ϵ, δ)-DP is a special instance of f-DP

Definition of (ϵ, δ)-DP

$$e^{-\epsilon}\, \mathbb{P}(M(S') \in E) - e^{-\epsilon}\delta \leqslant \mathbb{P}(M(S) \in E) \leqslant e^{\epsilon}\, \mathbb{P}(M(S') \in E) + \delta$$

Adapted from [Wasserman, Zhou '10]

An algorithm M is (ϵ, δ)-DP if and only if it is $f_{\epsilon,\delta}$-DP

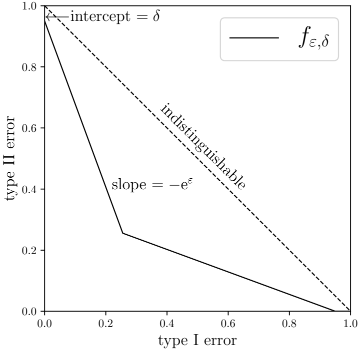

Issues with (ϵ, δ)-DP

- 4 segments. A bit *ad hoc*?
- w.p. δ, very bad events can happen

A primal-dual perspective on the relationship between f-DP and (ϵ, δ)-DP

From dual to primal

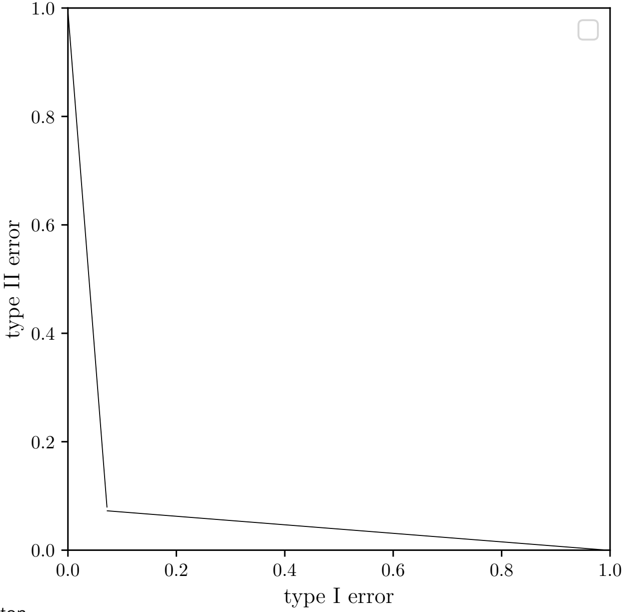

From dual to primal

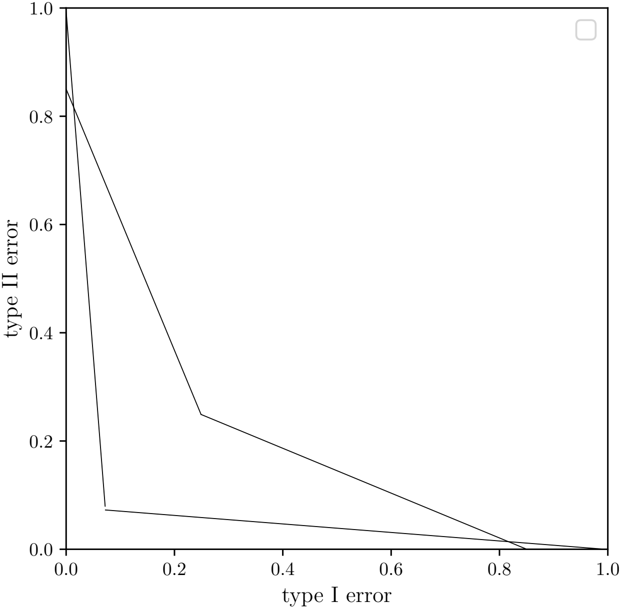

From dual to primal

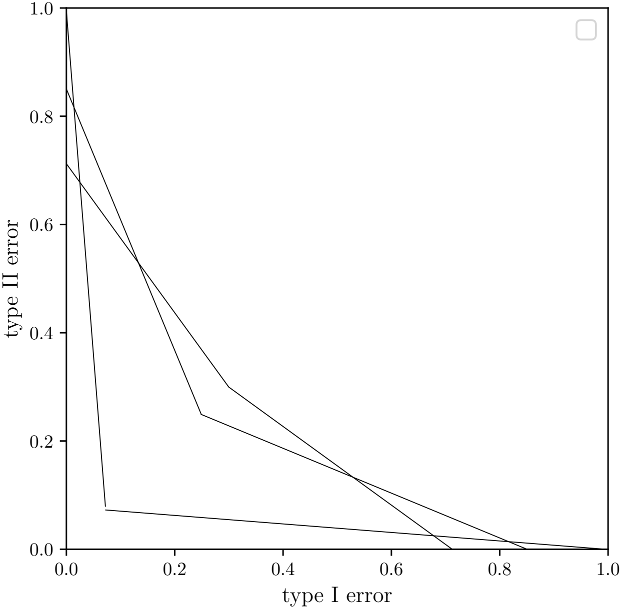

From dual to primal

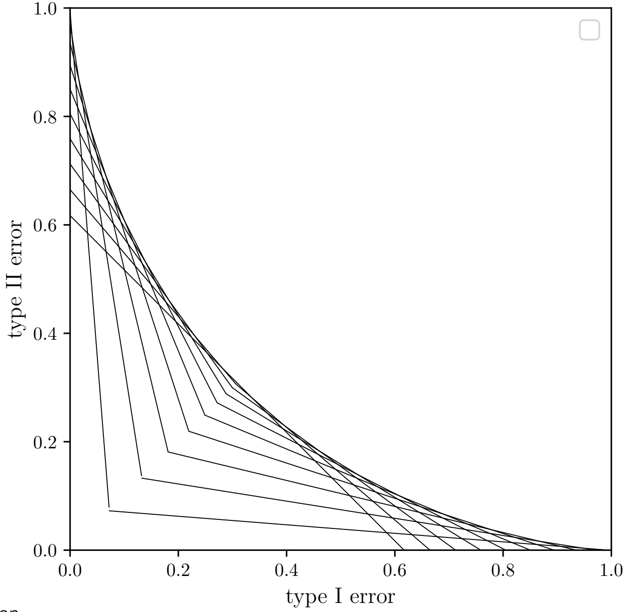

From dual to primal

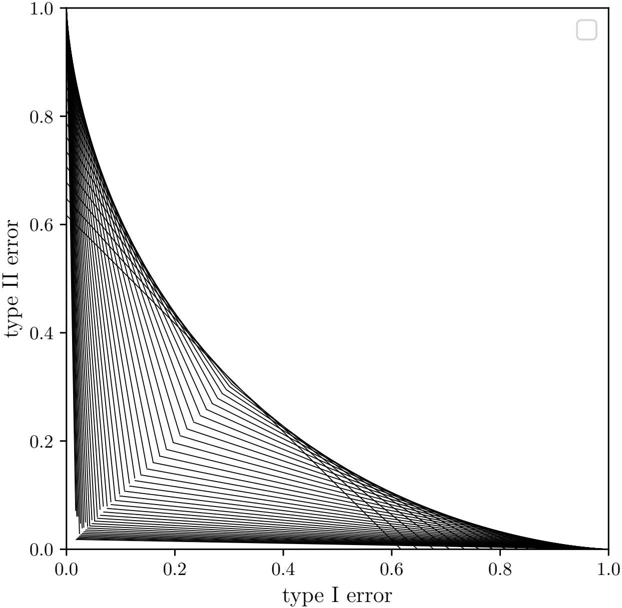

From dual to primal

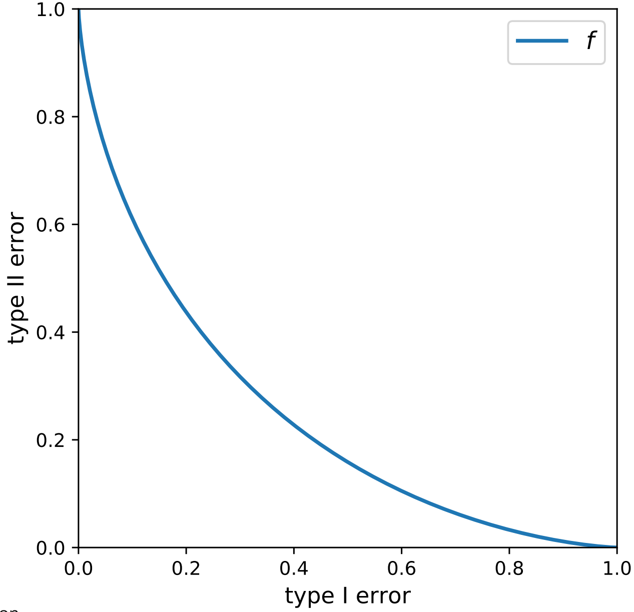

From primal to dual

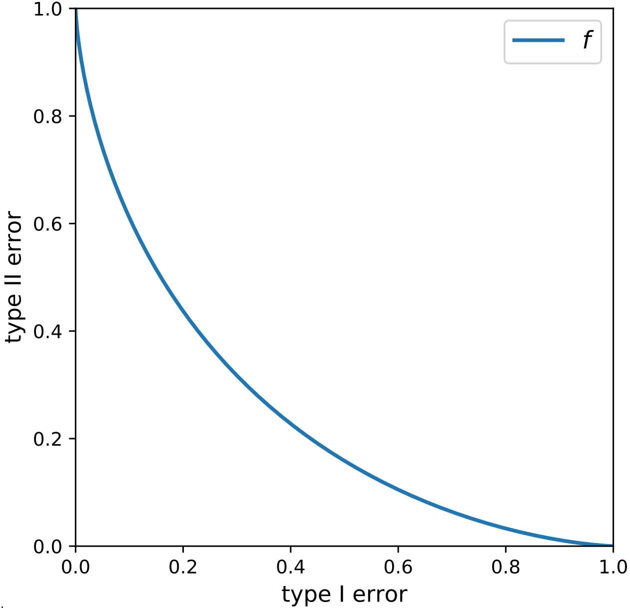

From primal to dual

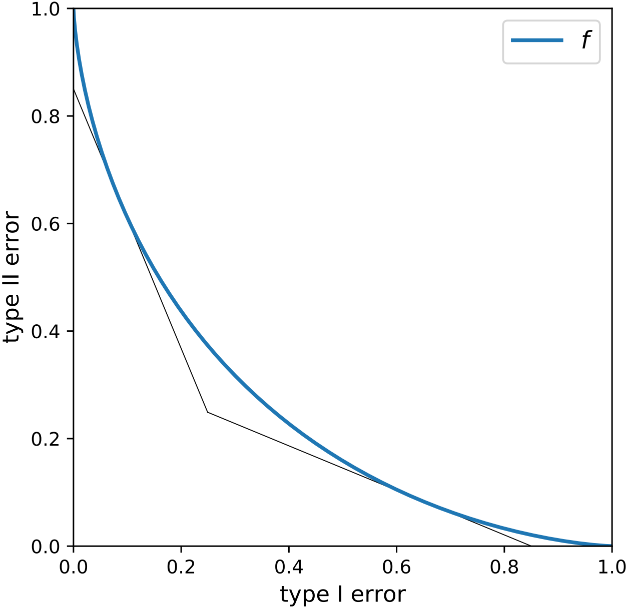

From primal to dual

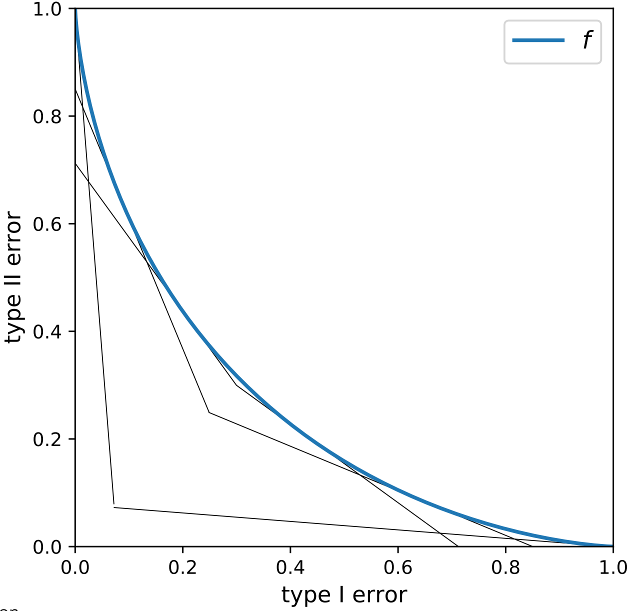

From primal to dual

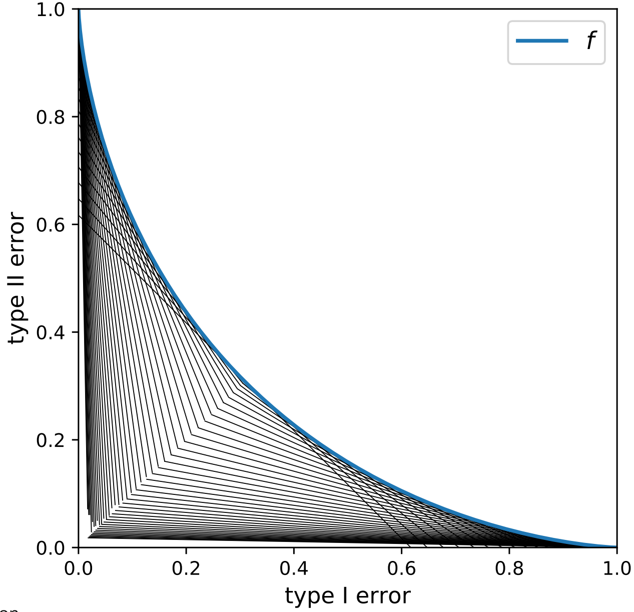

Is f too general? Let's focus!

Gaussian differential privacy (GDP)

Consider Gaussian trade-off function

$$G_\mu := T\big(\mathcal{N}(0,1), \mathcal{N}(\mu,1)\big)$$

for $\mu \geqslant 0$. Explicitly, $G_\mu(\alpha) = \Phi\big(\Phi^{-1}(1-\alpha) - \mu\big)$

Definition (DRS)

An algorithm M is said to be μ-GDP if

$$T\big(M(S), M(S')\big) \geqslant G_\mu$$

for all neighboring datasets S and S'

Gaussian differential privacy (GDP)

Consider Gaussian trade-off function

$$G_\mu := T\big(\mathcal{N}(0,1), \mathcal{N}(\mu,1)\big)$$

for $\mu \geqslant 0$. Explicitly, $G_\mu(\alpha) = \Phi\big(\Phi^{-1}(1-\alpha) - \mu\big)$

Definition (DRS)

An algorithm M is said to be μ-GDP if

$$T\big(M(S), M(S')\big) \geqslant G_\mu$$

for all neighboring datasets S and S'

- A single-parameter family (related to LDA)
- Focal to f-DP (a central limit theorem phenomenon)

How to interpret μ in GDP?

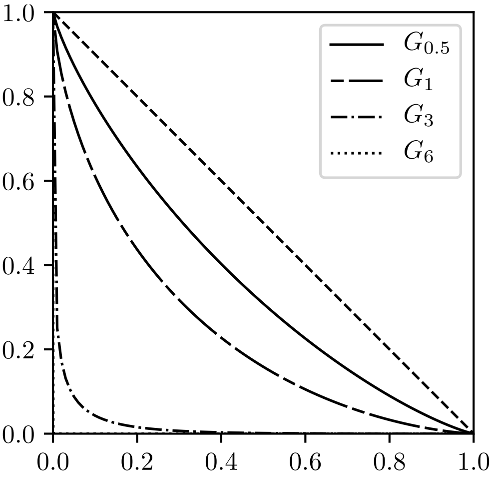

- Privacy amounts to distinguishing between $\mathcal{N}(0,1)$ and $\mathcal{N}(\mu,1)$
- $\mu \leqslant 1$: reasonably private. $\mu \geqslant 6$: blatantly non-private

HOW IS DIFFERENTIAL PRIVACY ACHIEVED?

A universal template: adding noise!

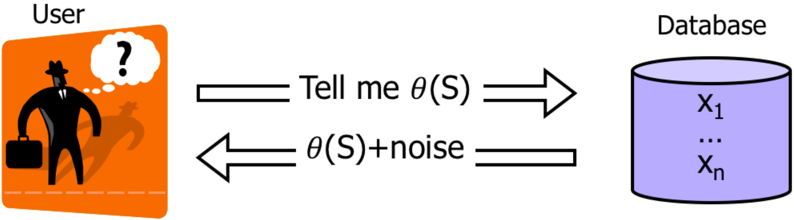

A universal template: adding noise!

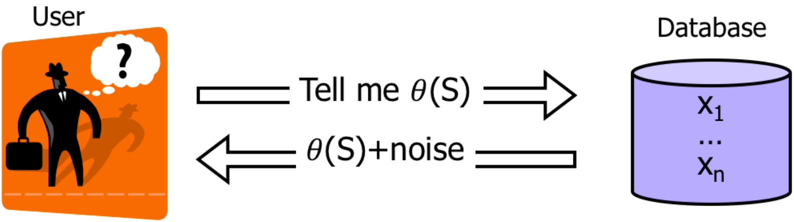

Sensitivity $\Delta\theta := \max_{S \sim S'} |\theta(S) - \theta(S')|$

Privacy guarantee

Consider the Gaussian mechanism $M(S) = \theta(S) + \mathcal{N}(0, \sigma^2)$. Then, M is μ-GDP with $\mu = \Delta\theta/\sigma$

A universal template: adding noise!

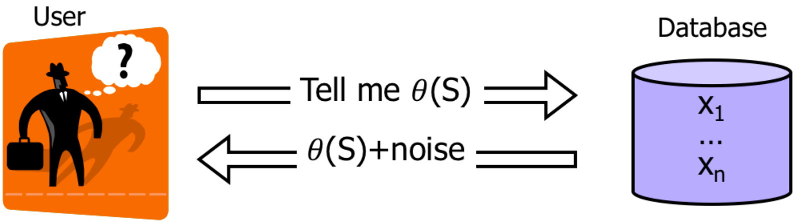

Sensitivity $\Delta\theta := \max_{S \sim S'} |\theta(S) - \theta(S')|$

Privacy guarantee

Consider the Gaussian mechanism $M(S) = \theta(S) + \mathcal{N}(0, \sigma^2)$. Then, M is μ-GDP with $\mu = \Delta\theta/\sigma$

- Gaussian mechanism is to GDP as Laplace mechanism is to $(\epsilon, 0)$-DP

Outline

Paper

A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem

Annual Review of Statistics and Its Application, 2025

Post-processing

A post-processing operation is a (randomized) algorithm that takes as input $M(S)$ and yields a new algorithm that we denote by $\texttt{Proc} \circ M$

- aka garbling

Post-processing

A post-processing operation is a (randomized) algorithm that takes as input $M(S)$ and yields a new algorithm that we denote by $\texttt{Proc} \circ M$

- aka garbling

Axiom

If an algorithm M is private, then its post-processing $\texttt{Proc} \circ M$ must also be private

Post-processing

A post-processing operation is a (randomized) algorithm that takes as input $M(S)$ and yields a new algorithm that we denote by $\texttt{Proc} \circ M$

- aka garbling

Axiom

If an algorithm M is private, then its post-processing $\texttt{Proc} \circ M$ must also be private

f-DP satisfies the axiom

f-DP satisfies the post-processing property because, for any P and Q,

$$T\big(\texttt{Proc}(P), \texttt{Proc}(Q)\big) \geqslant T(P, Q)$$

A representation theorem

Theorem (S)

Under the axiom, any DP definition must have its metric defined through the trade-off function:

$$D(P,Q) = d(T(P,Q))$$

A representation theorem

Theorem (S)

Under the axiom, any DP definition must have its metric defined through the trade-off function:

$$D(P,Q) = d(T(P,Q))$$

- Thus, f-DP is the most informative

A representation theorem

Theorem (S)

Under the axiom, any DP definition must have its metric defined through the trade-off function:

$$D(P, Q) = d(T(P, Q))$$

- Thus, f-DP is the most informative
- For ϵ-DP: $D(P, Q) := \sup_E \log \frac{P(E)}{Q(E)}$
- For (ϵ, δ)-DP: $D(P, Q) = \max_{E: P(E) \geqslant \delta} \log \frac{P(E) - \delta}{Q(E)}$

A representation theorem

Theorem (S)

Under the axiom, any DP definition must have its metric defined through the trade-off function:

$$D(P, Q) = d(T(P, Q))$$

- Thus, f-DP is the most informative
- For ϵ-DP: $D(P, Q) := \sup_E \log \frac{P(E)}{Q(E)}$
- For (ϵ, δ)-DP: $D(P, Q) = \max_{E: P(E) \geqslant \delta} \log \frac{P(E) - \delta}{Q(E)}$
- How to prove it?

Blackwell's informativeness theorem

[1]Greatest Of All Theorems in Slides

Blackwell's informativeness theorem

Lemma (Blackwell '51, GOATS[1])

Informativeness and post-processing are equivalent:

(a) $T(P', Q') \geqslant T(P, Q)$ (***informativeness***)

(b) (P', Q') *is Blackwell harder to distinguish than* (P, Q) (***post-processing/garbling***). *(That is,* $P' = \texttt{Proc}(P), Q' = \texttt{Proc}(Q)$*)*

[1] Greatest Of All Theorems in Slides

Blackwell's informativeness theorem

Lemma (Blackwell '51, GOATS[1])

Informativeness and post-processing are equivalent:

(a) $T(P', Q') \geqslant T(P, Q)$ (***informativeness***)

(b) (P', Q') *is Blackwell harder to distinguish than* (P, Q) (***post-processing/garbling***). *(That is,* $P' = \texttt{Proc}(P), Q' = \texttt{Proc}(Q)$*)*

COMPARISON OF EXPERIMENTS

DAVID BLACKWELL

HOWARD UNIVERSITY

1. Summary

Bohnenblust, Shapley, and Sherman [2] have introduced a method of comparing two sampling procedures or experiments; essentially their concept is that one experiment α is more informative than a second experiment β, $\alpha \supset \beta$, if, for every possible risk function, any risk attainable with β is also attainable with α. If α is a sufficient statistic for a procedure equivalent to β, $\alpha > \beta$, it is shown that $\alpha \supset \beta$. In the case of dichotomies, the converse is proved. Whether $>$ and $\supset$ are equivalent in general is not known. Various properties of $>$ and $\supset$ are obtained, such as the following: if $\alpha > \beta$ and γ is independent of both, then the combination $(\alpha, \gamma) > (\beta, \gamma)$. An application to a problem in 2×2 tables is discussed.

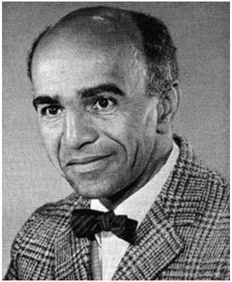

- Blackwell used terms: experiment & transformation

[1]Greatest Of All Theorems in Slides

Blackwell's informativeness theorem

Lemma (Blackwell '51, GOATS[1])

Informativeness and post-processing are equivalent:

(a) $T(P', Q') \geqslant T(P, Q)$ ***(informativeness)***

(b) (P', Q') *is Blackwell harder to distinguish than* (P, Q) ***(post-processing/garbling)****. (That is,* $P' = \texttt{Proc}(P), Q' = \texttt{Proc}(Q)$*)*

COMPARISON OF EXPERIMENTS

DAVID BLACKWELL
HOWARD UNIVERSITY

1. Summary

Bohnenblust, Shapley, and Sherman [2] have introduced a method of comparing two sampling procedures or experiments; essentially their concept is that one experiment α is more informative than a second experiment β, $\alpha \supset \beta$, if, for every possible risk function, any risk attainable with β is also attainable with α. If α is a sufficient statistic for a procedure equivalent to β, $\alpha \succ \beta$, it is shown that $\alpha \supset \beta$. In the case of dichotomies, the converse is proved. Whether $\succ$ and $\supset$ are equivalent in general is not known. Various properties of $\succ$ and $\supset$ are obtained, such as the following: if $\alpha \succ \beta$ and γ is independent of both, then the combination $(\alpha, \gamma) \succ (\beta, \gamma)$. An application to a problem in 2×2 tables is discussed.

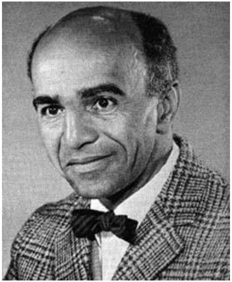

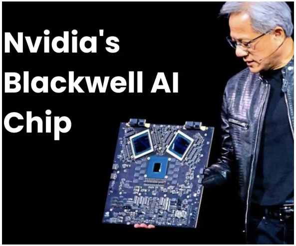

- Blackwell used terms: experiment & transformation

[1] Greatest Of All Theorems in Slides

Rényi divergence is not as informative

Rényi divergence of order γ

$$R_\gamma(P\|Q) := \frac{1}{\gamma - 1} \log \mathbb{E}_Q \left(\frac{\mathrm{d}P}{\mathrm{d}Q}\right)^\gamma$$

- Concentrated DP [Dwork, Rothblum '16], zero concentrated DP [Bun, Steinke '16], truncated concentrated DP [Bun, Dwork, Rothblum, Steinke '18], and Rényi DP [Mironov '17] are all defined via Rényi divergence

Rényi divergence is not as informative

Rényi divergence of order γ

$$R_\gamma(P\|Q) := \frac{1}{\gamma - 1} \log \mathbb{E}_Q \left(\frac{\mathrm{d}P}{\mathrm{d}Q}\right)^\gamma$$

- Concentrated DP [Dwork, Rothblum '16], zero concentrated DP [Bun, Steinke '16], truncated concentrated DP [Bun, Dwork, Rothblum, Steinke '18], and Rényi DP [Mironov '17] are all defined via Rényi divergence

Proposition (DRS)

Let $P_\epsilon = \mathrm{Bern}(\frac{\mathrm{e}^\epsilon}{1+\mathrm{e}^\epsilon}), Q_\epsilon = \mathrm{Bern}(\frac{1}{1+\mathrm{e}^\epsilon})$*. For* $0 < \epsilon < 4$*, the following are true:*

(a) *For all* $\gamma > 1$*,* $R_\gamma(P_\epsilon\|Q_\epsilon) < R_\gamma\big(\mathcal{N}(0,1)\|\mathcal{N}(\epsilon,1)\big)$

(b) *Using total variation,* $d_{\mathrm{TV}}(P_\epsilon, Q_\epsilon) > d_{\mathrm{TV}}\big(\mathcal{N}(0,1), \mathcal{N}(\epsilon,1)\big)$

Rényi divergence is not as informative

Rényi divergence of order γ

$$R_\gamma(P\|Q) := \frac{1}{\gamma-1}\log \mathbb{E}_Q\Big(\frac{\mathrm{d}P}{\mathrm{d}Q}\Big)^\gamma$$

- Concentrated DP [Dwork, Rothblum '16], zero concentrated DP [Bun, Steinke '16], truncated concentrated DP [Bun, Dwork, Rothblum, Steinke '18], and Rényi DP [Mironov '17] are all defined via Rényi divergence

Proposition (DRS)

Let $P_\epsilon = \mathrm{Bern}(\frac{\mathrm{e}^\epsilon}{1+\mathrm{e}^\epsilon}), Q_\epsilon = \mathrm{Bern}(\frac{1}{1+\mathrm{e}^\epsilon})$*. For* $0 < \epsilon < 4$*, the following are true:*

(a) *For all* $\gamma > 1$, $R_\gamma(P_\epsilon\|Q_\epsilon) < R_\gamma\big(\mathcal{N}(0,1)\|\mathcal{N}(\epsilon,1)\big)$

(b) *Using total variation,* $d_{\mathrm{TV}}(P_\epsilon, Q_\epsilon) > d_{\mathrm{TV}}\big(\mathcal{N}(0,1), \mathcal{N}(\epsilon,1)\big)$

- No such a phenomenon for trade-off functions
- Similar examples exist for (ϵ,δ)-DP

Properties f-DP

- Informative representation of privacy ✓

Outline

What is composition?

Sex	Blood	···	HIV
F	B	···	Y
M	A	···	N
M	O	···	N
M	O	···	Y
F	A	···	N
M	B	···	Y

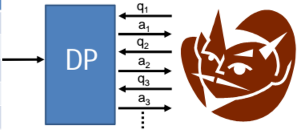

What is composition?

Sex	Blood	···	HIV
F	B	···	Y
M	A	···	N
M	O	···	N
M	O	···	Y
F	A	···	N
M	B	···	Y

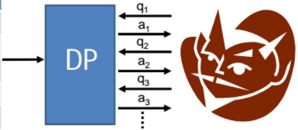

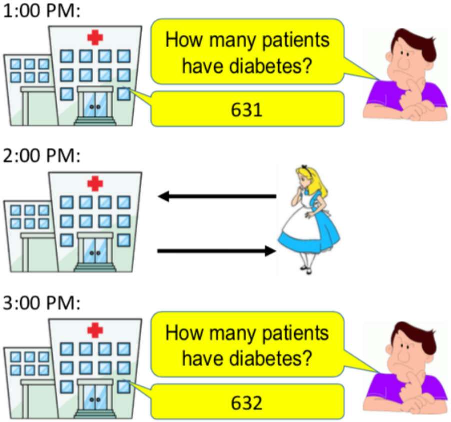

What is composition?

Sex	Blood	···	HIV
F	B	···	Y
M	A	···	N
M	O	···	N
M	O	···	Y
F	A	···	N
M	B	···	Y

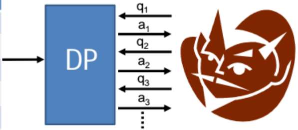

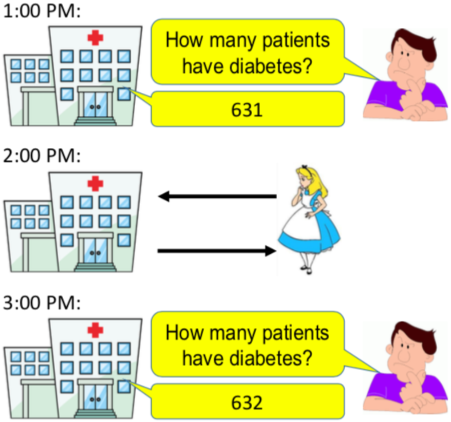

Composition surely leads to a privacy compromise. But how fast?

Definition of composition

Let $M_1 : X \to Y_1$ and $M_2 : X \times Y_1 \to Y_2$ be private algorithms. Define their composition $M : X \to Y_1 \times Y_2$ as

$$M(S) = (M_1(S), M_2(S, M_1(S)))$$

Given a sequence of algorithms $M_i : X \times Y_1 \times \cdots \times Y_{i-1} \to Y_i$ for $i \leqslant k$, recursively define the composition:

$$M : X \to Y_1 \times \cdots \times Y_k$$

Tensor product of trade-off functions

Definition

The tensor product of two trade-off functions $f = T(P, Q)$ and $g = T(P', Q')$ is defined as

$$f \otimes g := T(P \times P', Q \times Q')$$

- Well-defined
- The operator $\otimes$ is commutative and associative
- For GDP, $G_{\mu_1} \otimes G_{\mu_2} \otimes \cdots \otimes G_{\mu_k} = G_\mu$, where $\mu = \sqrt{\mu_1^2 + \cdots + \mu_k^2}$

Composition is an algebra

Proposition

Suppose $M_i(\cdot, y_1, \cdots, y_{i-1})$ *is* f_i*-DP for all* $y_1 \in Y_1, \ldots, y_{i-1} \in Y_{i-1}$. *Then the composition algorithm* $M: X \to Y_1 \times \cdots \times Y_k$ *is*

$$f_1 \otimes \cdots \otimes f_k\text{-DP}$$

Composition is an algebra

Proposition

Suppose $M_i(\cdot, y_1, \cdots, y_{i-1})$ is f_i-DP for all $y_1 \in Y_1, \ldots, y_{i-1} \in Y_{i-1}$. Then the composition algorithm $M : X \to Y_1 \times \cdots \times Y_k$ is

$$f_1 \otimes \cdots \otimes f_k\text{-DP}$$

- *Cannot* be improved in general
- Composition in f-DP is reduced to *algebra*
- k-step composition of μ-GDP algorithms is $\sqrt{k}\mu$-GDP

Central limit theorem for f-DP

Theorem (DRS)

Let $\{f_{ki} : 1 \leqslant i \leqslant k, k = 1, 2, \ldots\}$ *be a triangular array of trade-off functions, each being* $O(1/\sqrt{k})$ *close to perfect privacy. Then*

$$\lim_{k \to \infty} f_{k1} \otimes f_{k2} \otimes \cdots \otimes f_{kk} = G_\mu$$

- *The convergence is uniform on* $[0, 1]$
- μ *can be computed from* $\{f_{ki}\}$

Central limit theorem for f-DP

Theorem (DRS)

Let $\{f_{ki} : 1 \leqslant i \leqslant k, k = 1, 2, \ldots\}$ be a triangular array of trade-off functions, each being $O(1/\sqrt{k})$ close to perfect privacy. Then

$$\lim_{k\to\infty} f_{k1} \otimes f_{k2} \otimes \cdots \otimes f_{kk} = G_\mu$$

- *The convergence is uniform on $[0, 1]$*
- *μ can be computed from $\{f_{ki}\}$*

- If M_{ki} is f_{ki}-DP, their composition is approximately μ-GDP

Central limit theorem for f-DP

Theorem (DRS)

Let $\{f_{ki} : 1 \leqslant i \leqslant k, k = 1, 2, \ldots\}$ be a triangular array of trade-off functions, each being $O(1/\sqrt{k})$ close to perfect privacy. Then

$$\lim_{k \to \infty} f_{k1} \otimes f_{k2} \otimes \cdots \otimes f_{kk} = G_\mu$$

- *The convergence is uniform on $[0, 1]$*
- *μ can be computed from $\{f_{ki}\}$*

- If M_{ki} is f_{ki}-DP, their composition is approximately μ-GDP
- An effective approximation tool

Central limit theorem for f-DP

Theorem (DRS)

Let $\{f_{ki} : 1 \leqslant i \leqslant k, k = 1, 2, \ldots\}$ be a triangular array of trade-off functions, each being $O(1/\sqrt{k})$ close to perfect privacy. Then

$$\lim_{k \to \infty} f_{k1} \otimes f_{k2} \otimes \cdots \otimes f_{kk} = G_\mu$$

- *The convergence is uniform on $[0, 1]$*
- *μ can be computed from $\{f_{ki}\}$*

- If M_{ki} is f_{ki}-DP, their composition is approximately μ-GDP
- An effective approximation tool
- GDP is to f-DP as Gaussian variables (rvs) to general rvs

Central limit theorem for ϵ–DP

Theorem (DRS)

Fix $\mu > 0$ and assume $\epsilon = \sqrt{\mu/k}$. Then

$$G_\mu\left(\alpha + \frac{c}{k}\right) - \frac{c}{k} \leqslant f_{\epsilon,0}^{\otimes k}(\alpha) \leqslant G_\mu\left(\alpha - \frac{c}{k}\right) + \frac{c}{k}$$

Central limit theorem for ϵ-DP

Theorem (DRS)

Fix $\mu > 0$ *and assume* $\epsilon = \sqrt{\mu/k}$. *Then*

$$G_\mu\left(\alpha + \frac{c}{k}\right) - \frac{c}{k} \leqslant f_{\epsilon,0}^{\otimes k}(\alpha) \leqslant G_\mu\left(\alpha - \frac{c}{k}\right) + \frac{c}{k}$$

- Local computation is #P-complete [Murtagh, Vadhan '16]

Central limit theorem for ϵ-DP

Theorem (DRS)

Fix $\mu > 0$ and assume $\epsilon = \sqrt{\mu/k}$. Then

$$G_\mu\left(\alpha + \frac{c}{k}\right) - \frac{c}{k} \leqslant f_{\epsilon,0}^{\otimes k}(\alpha) \leqslant G_\mu\left(\alpha - \frac{c}{k}\right) + \frac{c}{k}$$

- Local computation is #P-complete [Murtagh, Vadhan '16]
- Sharper than the $O(1/\sqrt{k})$ bound in Berry–Esseen

Central limit theorem for ϵ-DP

Theorem (DRS)

Fix $\mu > 0$ and assume $\epsilon = \sqrt{\mu/k}$. Then

$$G_\mu\left(\alpha + \frac{c}{k}\right) - \frac{c}{k} \leqslant f_{\epsilon,0}^{\otimes k}(\alpha) \leqslant G_\mu\left(\alpha - \frac{c}{k}\right) + \frac{c}{k}$$

- Local computation is #P-complete [Murtagh, Vadhan '16]
- Sharper than the $O(1/\sqrt{k})$ bound in Berry–Esseen

Privacy CLT Beats Berry–Esseen for ϵ-DP! Why?

Central limit theorem for ϵ-DP

Theorem (DRS)

Fix $\mu > 0$ *and assume* $\epsilon = \sqrt{\mu/k}$. *Then*

$$G_\mu\left(\alpha + \frac{c}{k}\right) - \frac{c}{k} \leqslant f_{\epsilon,0}^{\otimes k}(\alpha) \leqslant G_\mu\left(\alpha - \frac{c}{k}\right) + \frac{c}{k}$$

- Local computation is #P-complete [Murtagh, Vadhan '16]
- Sharper than the $O(1/\sqrt{k})$ bound in Berry–Esseen

Privacy CLT Beats Berry–Esseen for ϵ-DP! Why?

- Due to *randomization* of rejection rules, leading to continuity of trade-off functions

A numerical example

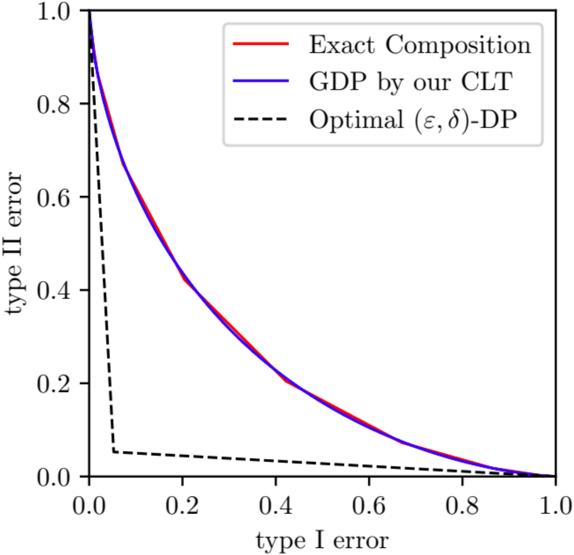

10-fold composition of $(1/\sqrt{10}, 0)$-DP. $\delta = 0.001$ in green curve

Properties of f-DP

- Informative representation of privacy ✓
- Algebraically convenient and tight composition operations ✓

Outline

What is subsampling for privacy?

Given dataset S, apply the algorithm M on a subsampled dataset $\mathrm{sub}(S)$, resulting a new algorithm $M \circ \mathrm{sub}(S)$

- Subsampling provides stronger privacy guarantees than when run on the whole dataset
- A frequently used tool for amplifying privacy

Subsampling theorem for f-DP

sub_m uniformly picks an m-sized subset from S. Let $p := m/n$

p-sampling operator C_p acting on trade-off functions

$$C_p(f) := \mathrm{Conv}\big(\min\{f_p, f_p^{-1}\}\big) = \min\{f_p, f_p^{-1}\}^{**}$$

- $f_p = pf + (1-p)\mathrm{Id}$, with $\mathrm{Id}(\alpha) = 1 - \alpha$
- $\min\{f_p, f_p^{-1}\}^{**}$ is double (convex) conjugate of $\min\{f_p, f_p^{-1}\}$ (the greatest convex lower bound)

Subsampling theorem for f-DP

sub_m uniformly picks an m-sized subset from S. Let $p := m/n$

p-sampling operator C_p acting on trade-off functions

$$C_p(f) := \mathrm{Conv}\big(\min\{f_p, f_p^{-1}\}\big) = \min\{f_p, f_p^{-1}\}^{**}$$

- $f_p = pf + (1-p)\mathrm{Id}$, with $\mathrm{Id}(\alpha) = 1 - \alpha$
- $\min\{f_p, f_p^{-1}\}^{**}$ is double (convex) conjugate of $\min\{f_p, f_p^{-1}\}$ (the greatest convex lower bound)

If M is f-DP, then $M \circ \mathrm{sub}_m$ is $C_p(f)$-DP, and it is tight

- The subsampling theorem for Rényi DP is complex [Wang, Balle, Kasiviswanathan '18]

Numerical examples

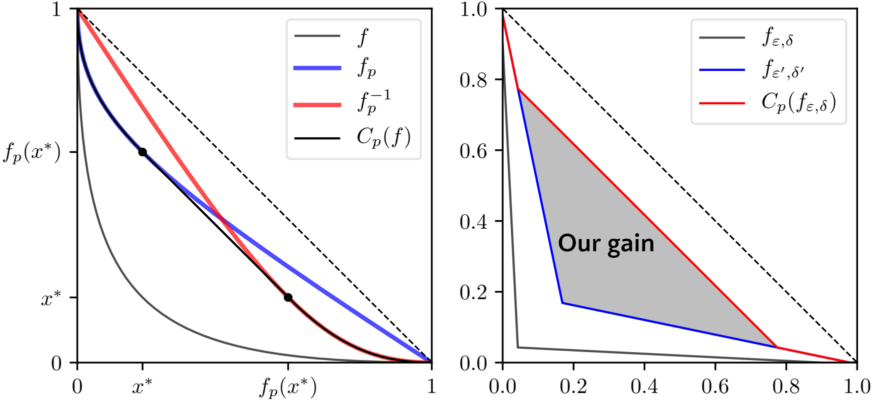

Left: $f = G_{1.8}, p = 0.35$. Right: $\epsilon = 3, \delta = 0.1, p = 0.2$

Properties of f-DP

- Informative representation of privacy ✓
- Algebraically convenient and tight composition operations ✓
- Sharp privacy amplification via subsampling ✓

Outline

Paper

Deep Learning with Gaussian Differential Privacy

Harvard Data Science Review, 2020

- Zhiqi Bu (Penn/Amazon)
- Jinshuo Dong (Penn/Northwestern/Tsinghua)
- Qi Long (Penn)

Privacy concerns in deep learning

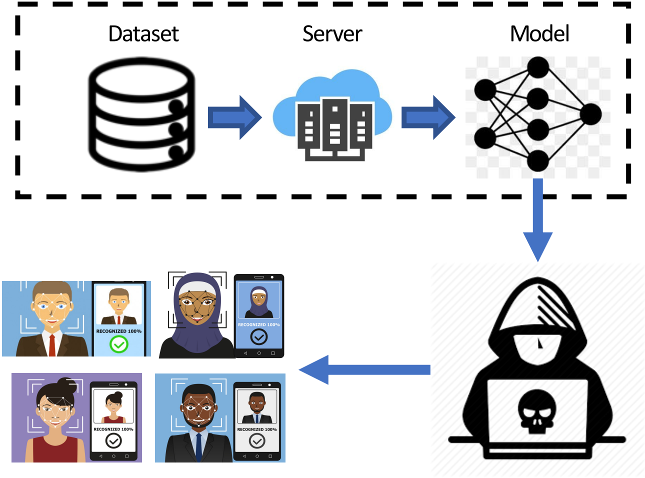

Privacy concerns in deep learning

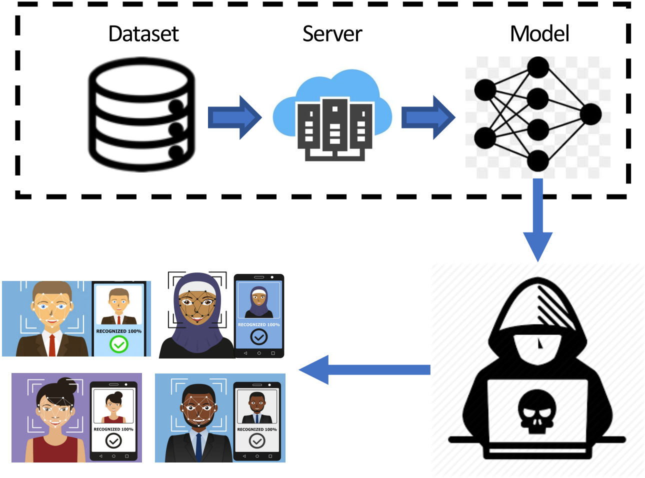

- Private deep learning by Google Brain [Abadi et al '16]

Private deep learning [Abadi et al '16]

Input: Dataset $S = \{x_1, \ldots, x_n\}$, loss function $\ell(\theta, x)$.
Parameters: initial weights θ_0, learning rate η_t, subsampling probability p, number of iterations T, noise scale σ, gradient norm bound R.

for $t = 0, \ldots, T-1$ **do**

Take a Poisson subsample $I_t \subseteq \{1, \ldots, n\}$ with subsampling probability p

for $i \in I_t$ **do**

$v_t^{(i)} \leftarrow \nabla_\theta \ell(\theta_t, x_i)$

$\bar{v}_t^{(i)} \leftarrow v_t^{(i)} / \max\left\{1, \|v_t^{(i)}\|_2 / R\right\}$ ▷ **Clip gradient**

$\theta_{t+1} \leftarrow \theta_t - \eta_t \cdot \frac{1}{|I_t|}\Big(\sum_{i \in I_t} \bar{v}_t^{(i)} + \sigma R \cdot \mathcal{N}(0, I)\Big)$ ▷ **Gaussian mechanism**

Output θ_T

- Moments accountant for (ϵ, δ)-DP [Abadi et al '16]
- Extends to noisy Adam

Can the f-DP framework improve privacy analysis?

Privacy analysis of deep learning

SGD equation

$$\theta_{t+1} = \mathtt{SGD} \circ \mathrm{sub}(S; \theta_t)$$

Observation

Deep Learning = Subsampling + Composition

Privacy analysis of deep learning

SGD equation

$$\theta_{t+1} = \texttt{SGD} \circ \mathrm{sub}(S; \theta_t)$$

Observation

Deep Learning = Subsampling + Composition

Thus, we get

Theorem (BDLS)

Private deep learning $M(S) = (\theta_1, \theta_2, \ldots, \theta_T)$ *is asymptotically* μ*-GDP with*

$$\mu = \frac{m}{n}\sqrt{T(\mathrm{e}^{1/\sigma^2} - 1)}$$

- m is the mini-batch size, and n is the total number of examples

f-DP gives tighter analysis on MNIST

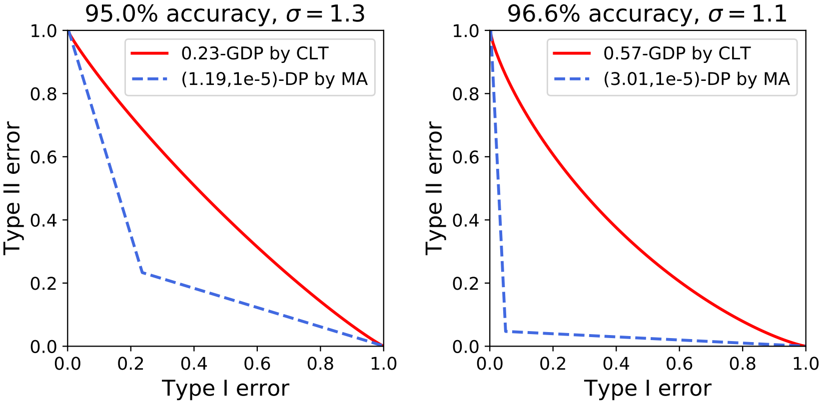

Solid red: our f-DP analysis. Dashed blue: moments accountant by Google Brain

f-DP gives tighter analysis on MNIST

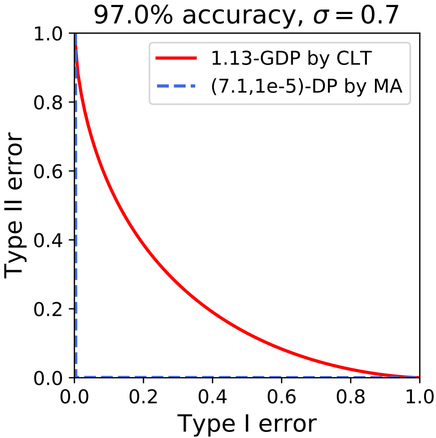

- Our f-DP interpretation is $\mathcal{N}(0,1)$ vs $\mathcal{N}(1.13,1)$; while MA gives $(7.1, 10^{-5})$-DP, noting $\mathrm{e}^{7.1} = 1212.0$

A Pareto improvement of privacy vs accuracy trade-off

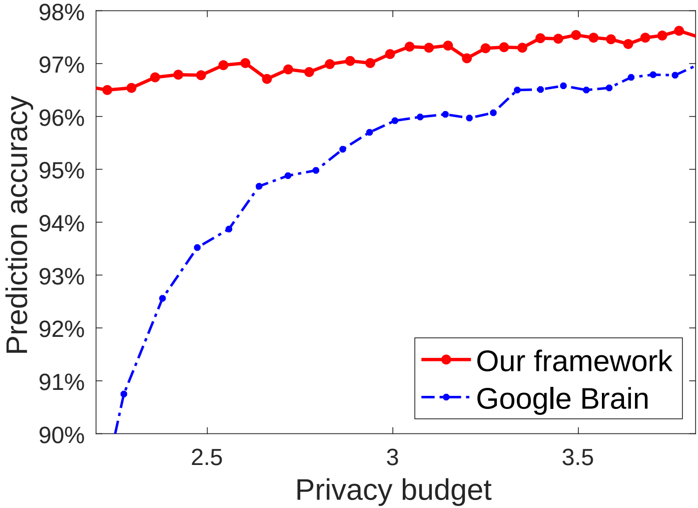

Fix $\delta = 10^{-5}$ but vary ϵ

Outline

1. Introduction to f-DP

2. Informative representation of privacy

3. Composition and central limit theorems

4. Amplifying privacy via subsampling

5. Application to deep learning

6. Application to 2020 United States Census

Paper

Revealing the Underestimated Privacy of the 2020 United States Census

Coming soon

- Buxin Su (Penn)
- Chendi Wang (Penn)

US Census Bureau adopted DP in 2020 decennial census

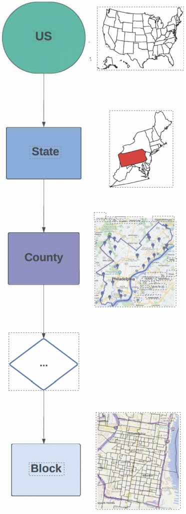

- Most queries take integer values, e.g.,

$$M(S) = \sum_{x \in \mathrm{NY}} \mathbf{1}\left[x \text{ is 18 or older}\right]$$

US Census Bureau adopted DP in 2020 decennial census

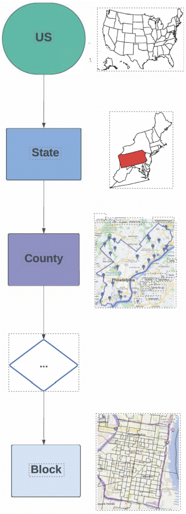

- Most queries take integer values, e.g.,

$$M(S) = \sum_{x \in \text{NY}} \mathbf{1}\left[x \text{ is 18 or older}\right]$$

- Add integer-valued noise to census microdata, with pdf

$$p_{\text{DG}}(x) = \frac{1}{Z(\mu, \sigma^2)} \mathrm{e}^{-\frac{(x-\mu)^2}{2\sigma^2}}, \qquad \text{for all } x \in \mathbb{Z}$$

US Census Bureau adopted DP in 2020 decennial census

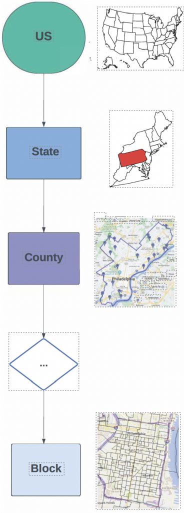

- Most queries take integer values, e.g.,

$$M(S) = \sum_{x \in \text{NY}} \mathbf{1}\left[x \text{ is 18 or older}\right]$$

- Add integer-valued noise to census microdata, with pdf

$$p_{\mathrm{DG}}(x) = \frac{1}{Z(\mu, \sigma^2)} \mathrm{e}^{-\frac{(x-\mu)^2}{2\sigma^2}}, \qquad \text{for all } x \in \mathbb{Z}$$

- Composition of 9 queries for each geographical level

Bureau hasn't fully used up privacy budget!

Theorem (SSW)

For any δ, f-DP yields a tighter ϵ privacy bound for census data than the Bureau's approach

Bureau hasn't fully used up privacy budget!

Theorem (SSW)

For any δ, f-DP yields a tighter ϵ privacy bound for census data than the Bureau's approach

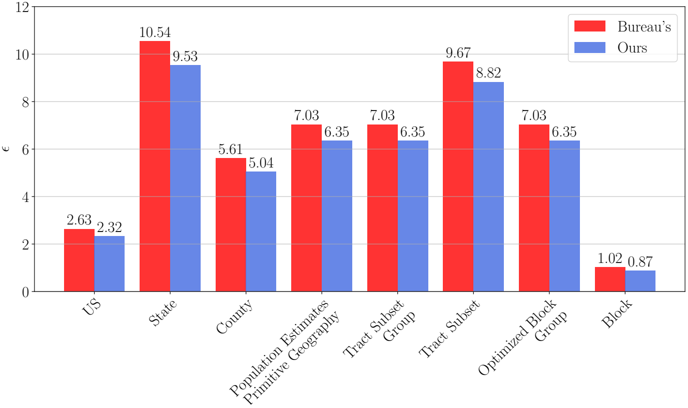

Bureau hasn't fully used up privacy budget!

Theorem (SSW)

For any δ, f-DP yields a tighter ϵ privacy bound for census data than the Bureau's approach

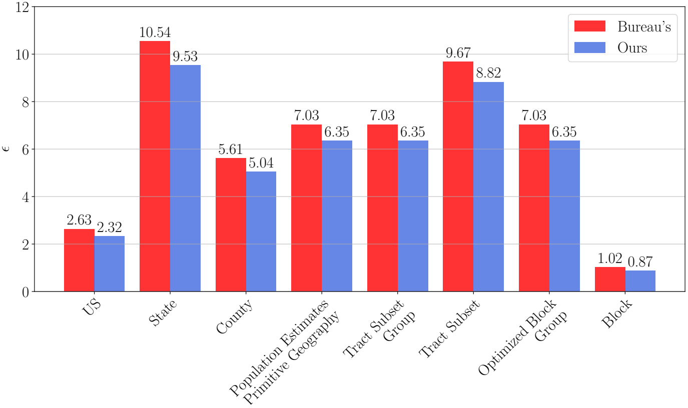

- Resolved an open question posed by the US Census Bureau [Kifer et al '22]

Less noise can be added to the census for the same privacy budget

Reduced noise with equivalent privacy bound

Comparison of variance: Bureau's approach vs. our f-DP based approach

Geographic Level	US	State	County	PEPG
Bureau's	69.40	5.00	16.07	10.47
Ours	54.74	4.25	13.21	8.71
Reduction	13.9%	15%	17.8%	16.8%

Geographic Level	Tract Subset Group	Tract Subset	Optimized Block Group	Block
Bureau's	10.47	5.77	10.47	451.13
Ours	8.71	4.89	8.71	338.28
Reduction	16.8%	15.3%	16.8%	25%

Reduced noise with equivalent privacy bound

Comparison of variance: Bureau's approach vs. our f-DP based approach

Geographic Level	US	State	County	PEPG
Bureau's	69.40	5.00	16.07	10.47
Ours	54.74	4.25	13.21	8.71
Reduction	13.9%	15%	17.8%	16.8%

Geographic Level	Tract Subset Group	Tract Subset	Optimized Block Group	Block
Bureau's	10.47	5.77	10.47	451.13
Ours	8.71	4.89	8.71	338.28
Reduction	16.8%	15.3%	16.8%	25%

- More accurate data for downstream applications of census

Concluding remarks

Privacy: a foundation for trustworthy data science

Summary

	Informativeness	Composition	Subsampling
ϵ-DP			
(ϵ, δ)-DP			
Divergence based DPs			
f-DP			

Summary

	Informativeness	Composition	Subsampling
ϵ-DP	✗		
(ϵ, δ)-DP	✗		
Divergence based DPs	✗		
f-DP	✓		

Gaussian differential privacy

- Trade-off functions are informative

Summary

	Informativeness	Composition	Subsampling
ϵ-DP	✗	✗	
(ϵ, δ)-DP	✗	✗	
Divergence based DPs	✗	✓	
f-DP	✓	✓	

Gaussian differential privacy

- Trade-off functions are informative
- Tight composition

Summary

	Informativeness	Composition	Subsampling
ϵ-DP	✗	✗	✓
(ϵ, δ)-DP	✗	✗	✓
Divergence based DPs	✗	✓	✗
f-DP	✓	✓	✓

Gaussian differential privacy

- Trade-off functions are informative
- Tight composition
- Sharp subsampling

Summary

	Informativeness	Composition	Subsampling
ϵ-DP	✗	✗	✓
(ϵ, δ)-DP	✗	✗	✓
Divergence based DPs	✗	✓	✗
f-DP	✓	✓	✓

Gaussian differential privacy

- Trade-off functions are informative
- Tight composition
- Sharp subsampling
- State-of-the-art applications to private deep learning and US Census

Take-home messages

Most Informative — Algebraically Nice

$$
\begin{tikzcd}
 & f\text{-DP} \arrow[dl] \arrow[dr] \arrow[dr, dashed, shift left, "\text{CLT}"] & \\
(\varepsilon, \delta)\text{-DP} \arrow[rr, dashed, "\text{CLT}"'] & & \text{GDP}
\end{tikzcd}
$$

Available in `TensorFlow Privacy`

1. *Gaussian Differential Privacy*
2. *A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem*
3. *Deep Learning with Gaussian Differential Privacy*
4. *Revealing the Underestimated Privacy of the 2020 United States Census*

Take-home messages

Most Informative   Algebraically Nice

$$\begin{tikzcd} & f\text{-DP} \arrow[dl] \arrow[dr] \arrow[dr, dashed, bend left, "\text{CLT}"] & \\ (\varepsilon, \delta)\text{-DP} \arrow[rr, dashed, "\text{CLT}"'] & & \text{GDP} \end{tikzcd}$$

Available in `TensorFlow Privacy`

1. *Gaussian Differential Privacy*
2. *A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem*
3. *Deep Learning with Gaussian Differential Privacy*
4. *Revealing the Underestimated Privacy of the 2020 United States Census*

Take-home messages

Most Informative f-DP Algebraically Nice

$$\begin{tikzcd} & f\text{-DP} \arrow[dl] \arrow[dr, "\text{CLT}"] \arrow[dr, dashed] & \\ (\varepsilon, \delta)\text{-DP} \arrow[rr, dashed, "\text{CLT}"'] & & \text{GDP} \end{tikzcd}$$

Available in `TensorFlow Privacy`

1. *Gaussian Differential Privacy*
2. *A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem*
3. *Deep Learning with Gaussian Differential Privacy*
4. *Revealing the Underestimated Privacy of the 2020 United States Census*

The Return of the King

Most Informative · Algebraically Nice

$$\begin{tikzcd} & f\text{-DP} \arrow[dl] \arrow[dr, "\mathrm{CLT}"] \arrow[dr, dashed] & \\ (\varepsilon, \delta)\text{-DP} \arrow[rr, dashed, "\mathrm{CLT}"'] & & \mathrm{GDP} \end{tikzcd}$$

Available in `TensorFlow Privacy`

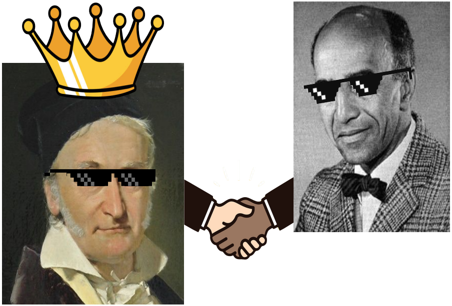

1. *Gaussian Differential Privacy*
2. *A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem*
3. *Deep Learning with Gaussian Differential Privacy*
4. *Revealing the Underestimated Privacy of the 2020 United States Census*

The Return of the King

Most Informative

Algebraically Nice

$$\begin{tikzcd} & f\text{-DP} \arrow[dl] \arrow[dr, "\text{CLT}"] & \\ (\varepsilon, \delta)\text{-DP} \arrow[rr, dashed, "\text{CLT}"'] & & \text{GDP} \end{tikzcd}$$

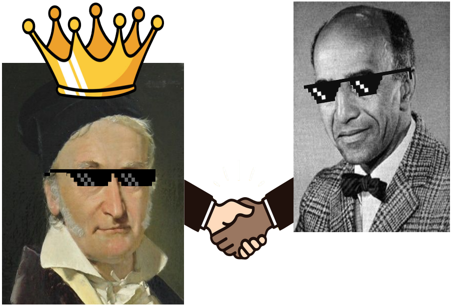

Available in `TensorFlow Privacy`

Learned from Blackwell & Gauss

Despite its origin in computer science, DP is fundamentally a statistical concept

1. *Gaussian Differential Privacy*
2. *A Statistical Viewpoint on Differential Privacy: Hypothesis Testing, Representation and Blackwell's Theorem*
3. *Deep Learning with Gaussian Differential Privacy*
4. *Revealing the Underestimated Privacy of the 2020 United States Census*

Proof sketch

Why $f_1 \otimes f_2 \otimes \cdots \otimes f_k \approx G_\mu$?

Let $f_i = T(P_i, Q_i)$. Test $H_0 : \boldsymbol{y} \sim P_1 \times \cdots \times P_k$ vs $H_1 : \boldsymbol{y} \sim Q_1 \times \cdots \times Q_k$

Proof sketch

Why $f_1 \otimes f_2 \otimes \cdots \otimes f_k \approx G_\mu$?

Let $f_i = T(P_i, Q_i)$. Test $H_0 : \boldsymbol{y} \sim P_1 \times \cdots \times P_k$ vs $H_1 : \boldsymbol{y} \sim Q_1 \times \cdots \times Q_k$

- Optimal test is

$$T := \frac{L - \mathbb{E}_P L}{\sqrt{\mathrm{Var}_P(L)}},$$

  where the log-likelihood ratio

$$L = \log \prod_{i=1}^{k} \frac{q_i(y_i)}{p_i(y_i)} = \sum_{i=1}^{k} \log \frac{q_i(y_i)}{p_i(y_i)} \equiv \sum_{i=1}^{k} L_i(y_i)$$

Proof sketch

Why $f_1 \otimes f_2 \otimes \cdots \otimes f_k \approx G_\mu$?

Let $f_i = T(P_i, Q_i)$. Test $H_0 : \boldsymbol{y} \sim P_1 \times \cdots \times P_k$ vs $H_1 : \boldsymbol{y} \sim Q_1 \times \cdots \times Q_k$

- Optimal test is

$$T := \frac{L - \mathbb{E}_P L}{\sqrt{\mathrm{Var}_P(L)}},$$

  where the log-likelihood ratio

$$L = \log \prod_{i=1}^{k} \frac{q_i(y_i)}{p_i(y_i)} = \sum_{i=1}^{k} \log \frac{q_i(y_i)}{p_i(y_i)} \equiv \sum_{i=1}^{k} L_i(y_i)$$

- Under H_0, T is approximately $\mathcal{N}(0,1)$; and under H_1, T is approximately $\mathcal{N}(\mu, 1)$

Proof sketch

Why $f_1 \otimes f_2 \otimes \cdots \otimes f_k \approx G_\mu$?

Let $f_i = T(P_i, Q_i)$. Test $H_0 : \boldsymbol{y} \sim P_1 \times \cdots \times P_k$ vs $H_1 : \boldsymbol{y} \sim Q_1 \times \cdots \times Q_k$

- Optimal test is

$$T := \frac{L - \mathbb{E}_P L}{\sqrt{\mathrm{Var}_P(L)}},$$

  where the log-likelihood ratio

$$L = \log \prod_{i=1}^{k} \frac{q_i(y_i)}{p_i(y_i)} = \sum_{i=1}^{k} \log \frac{q_i(y_i)}{p_i(y_i)} \equiv \sum_{i=1}^{k} L_i(y_i)$$

- Under H_0, T is approximately $\mathcal{N}(0,1)$; and under H_1, T is approximately $\mathcal{N}(\mu, 1)$
- Le Cam's third lemma