Four big ideas in privacy (as it relates to statistics)

John Duchi
Stanford University

Frejus 2025

The plan

We will discuss differential privacy and its extensions. After the basics, we will look at four big ideas

- Composition: if we use many private mechanisms, how do we lose privacy?
- Amplification: how can we improve privacy by simple methods?
- Advanced privacy mechanisms: stability, robustness, matrix mechanisms
- Optimality and lower bounds

Notation and setting

- Data X_i of individuals $i = 1, 2, \ldots, n$
- Data represented as $P_n = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{X_i}$
- $P_n \in \mathcal{P}_n$, the space of *empirical distributions*
- Wish to compute

  $$\theta(P_n) \in \Theta$$

  (e.g., mean, minimizer of loss)
- A *mechanism* M is a randomized mapping $M : \mathcal{P}_n \to \Theta$

The definition of privacy

Definition (Dwork et al. [5, 4])

A randomized mechanism M is (ε, δ)-*differentially private* if

$$\mathbb{P}(M(P_n) \in A) \leq e^{\varepsilon}\mathbb{P}(M(P_n') \in A) + \delta$$

for all neighboring P_n and $P_n' \in \mathcal{P}_n$

Observation (Bayesian perspective)

Cannot update a prior very much based on $M(P_n)$

A hypothesis testing perspective

- Adversary tests $H_0 : P_n$ against $H_1 : P'_n$
- Define errors

$$\alpha_0 = \mathbb{P}\big(\text{reject}(M(P_n))\big) \quad \text{and} \quad \alpha_1 = \mathbb{P}\big(\text{accept}(M(P'_n))\big)$$

Lemma (Wasserman and Zhou [8])

M is (ε, δ)-differentially private if and only if

$$\alpha_0 + e^{\varepsilon}\alpha_1 \geq 1 - \delta \quad \textit{and} \quad \alpha_1 + e^{\varepsilon}\alpha_0 \geq 1 - \delta$$

Local differential privacy

- when curator of data may be untrustworthy

Definition (Local differential privacy)

A mechanism $M : \mathcal{X} \to \mathcal{Z}$ is ε-*locally differentially private* if

$$\mathbb{P}(M(x) \in A) \leq e^{\varepsilon}\mathbb{P}(M(x') \in A)$$

Example (Randomized response, Warner [7])

Wish to release a sensitive answer $X \in \{0, 1\}$.

Basic mechanisms

Global sensitivity

Definition

function $f : \mathcal{P}_n \to \mathbb{R}$ has *global sensitivity*

$$\mathsf{GS}(f) := \sup\left\{\left|f(P_n) - f(P_n')\right| \mid d_{\mathsf{ham}}(P_n, P_n') \leq 1\right\}.$$

Examples:

- means with bounded data
- some optimization solutions

Laplace mechanism

- Laplace random variable $Z \sim \mathsf{Lap}(1)$ has density

$$p(z) = \frac{1}{2}\exp(-|z|)$$

- assume f has global sensitivity $\mathsf{GS}(f) < \infty$

Definition

The *Laplace mechanism* releases

$$M(P_n) = f(P_n) + \frac{\mathsf{GS}(f)}{\varepsilon} \cdot \mathsf{Lap}(1)$$

- it is ε-differentially private

Laplace mechanism (d-dimensions)

Definition

A function f has ℓ_p-global sensitivity

$$\mathsf{GS}_p(f) := \sup\left\{\left\|f(P_n) - f(P_n')\right\|_p \mid d_{\mathsf{ham}}(P_n, P_n') \leq 1\right\}$$

Definition

The *Laplace mechanism* releases

$$M(P_n) = f(P_n) + \frac{\mathsf{GS}_1(f)}{\varepsilon} \cdot W$$

where $W \in \mathbb{R}^d$ has $W_j \overset{\text{iid}}{\sim} \mathsf{Lap}(1)$

Mean estimation with Laplace mechanisms

Example

Assume data $X_i \in \mathbb{R}^d$ have $\|X_i\|_2 \le r$ and wish to estimate

$$f(P_n) := \mathbb{E}_{P_n}[X] = \overline{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i.$$

Laplace mechanism behavior on this?

Gaussian mechanism

Definition (see ref. [1])

The *Gaussian mechanism* releases

$$M(P_n) = f(P_n) + \mathsf{GS}_2(f) \cdot \mathsf{N}\left(0, \sigma^2(\varepsilon, \delta) I\right)$$

where

$$\sigma^2(\varepsilon, \delta) \leq O(1) \frac{\log \frac{1}{\delta}}{\varepsilon^2}$$

Privacy loss random variable

Definition

if Q_0, Q_1 are distributions of $M(P_n)$ and $M(P_n')$, *privacy loss*

$$L_M(z) := \log \frac{dQ_0(z)}{dQ_1(z)}$$

and *privacy loss random variable*

$$L_M := \log \frac{dQ_0(Z)}{dQ_1(Z)} \quad \text{for } Z \sim Q_0$$

Lemma (Dwork and Roth [3], Lemma 3.17 or Duchi [2], Lemma 8.2.10)

M is (ε, δ)-differentially private if and only if $\mathbb{P}(|L_M| \geq \varepsilon) \leq \delta$

Privacy of Gaussian mechanism

- Control the privacy loss random variable

Mean estimation with Gaussian mechanisms

Example

Assume data $X_i \in \mathbb{R}^d$ have $\|X_i\|_2 \le r$ and wish to estimate

$$f(P_n) := \mathbb{E}_{P_n}[X] = \overline{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i.$$

Gaussian mechanism behavior on this?

Mean estimation with randomized response

Example

Assume data $X_i \in \mathbb{R}$ have $X_i \in \{0, 1\}$. Randomized response:

$$Z_i = \begin{cases} X_i & \text{w.p. } e^{\varepsilon}/(1+e^{\varepsilon}) \\ 1 - X_i & \text{w.p. } 1/(1+e^{\varepsilon}) \end{cases}$$

Then for appropriate a, b, $\widehat{\theta}_n = a\overline{Z}_n + b$ satisfies

$$\mathbb{E}\left[(\widehat{\theta}_n - \mathbb{E}[X])^2\right] \lesssim \frac{1}{\varepsilon^2 \wedge 1} \cdot \frac{1}{n}.$$

What we want from a privacy definition

- Protection against side information

- No post-processing improvements
- Graceful privacy degradation *after multiple releases*

Composition of privacy algorithms

- for mechanisms $M_1 : \mathcal{P}_n \to \Theta_1$ and $M_2 : \mathcal{P}_n \times \Theta_1 \to \Theta_2$, their *composition*

$$M_1 \circ M_2(P_n) := \big(M_1(P_n), M_2(P_n, M_1(P_n))\big)$$

- k-fold *adaptive composition*

$$M_1 \circ M_2 \circ \cdots \circ M_k(P_n) := \big(M_1(P_n), \ldots, M_k(P_n, M_{k-1}, \ldots, M_1)\big)$$

big question: if each mechanism is private, is $M_1 \circ \cdots \circ M_k$ private?

Composition

Theorem

The k-fold adaptive composition of (ε, δ)-differentially private mechanisms is $(k\varepsilon, k\delta)$-differentially private and

$$\left(k\varepsilon(e^{\varepsilon}-1)+O(1)\sqrt{k\varepsilon^2\log\frac{1}{\delta}}, O(1)k\delta\right)\text{-differentially private}$$

Proof sketch of composition

- For q_i = density of M_i, define privacy loss

$$L_i := \log \frac{q_i(\theta_i \mid P_n, \theta_1^{i-1})}{q_i(\theta_i \mid P_n', \theta_1^{i-1})}$$

- Apply Azuma-Hoeffding inequality

Alternative definitions

- "play better" with composition
- admit cleaner analyses in some cases

Rényi-differential privacy

- Rényi α-divergence between P and Q is

$$D_\alpha\left(P\|Q\right) := \frac{1}{\alpha-1}\log\int\left(\frac{dP}{dQ}\right)^\alpha dQ$$

Definition (Mironov [6])

Mechanism M is (α,ε)-*Rényi differentially private* if induced measures $Q(\cdot \mid P_n)$ and $Q(\cdot \mid P_n')$ satisfy

$$D_\alpha\left(Q(\cdot \mid P_n)\|Q(\cdot \mid P_n')\right) \leq \varepsilon.$$

Composition in Rényi privacy

Proposition (Mironov [6])

Let $Z_0 = M_0(P_n)$*,* $Z_1 = M_1(P_n, Z_0)$ *be* (α, ε_0) *and* (α, ε_1)*-RDP. Then* (Z_0, Z_1) *is* $(\alpha, \varepsilon_0 + \varepsilon_1)$*-RDP.*

From Rényi privacy to differential privacy

Proposition (Mironov [6])

If M is (α, ε)-RDP, then it is $(\varepsilon + \frac{\log \frac{1}{\delta}}{\alpha - 1}, \delta)$-DP.

Lemma

For any event A, $P(A) \leq \left(\exp(D_\alpha\left(P \| Q\right)) \cdot Q(A)\right)^{\frac{\alpha-1}{\alpha}}$.

Part 2: case-by-case analysis of $P(M \in A)$ versus $P(M' \in A)$

Composition of Gaussian mechanisms

Corollary

Adaptively choose k *functions* $f_i(P_n)$*, each* $\mathsf{GS}(f) \le 1$*. Then* $M_i = f_i(P_n) + \mathsf{N}(0, \sigma^2)$ *with* $\sigma^2 = \frac{k \log \frac{1}{\delta}}{\varepsilon^2} + \frac{k}{\varepsilon}$ *is* $(2\varepsilon, \delta)$*-DP.*

[1] B. Balle and Y.-X. Wang. Improving the Gaussian mechanism for differential privacy: analytical calibration and optimal denoising. *arXiv:1805.06530*, 2018.

[2] J. C. Duchi. *Lecture Notes on Information Theory and Statistics*. 2024. URL `http://web.stanford.edu/class/stats311`.

[3] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3 & 4):211–407, 2014.

[4] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EUROCRYPT 2006)*, 2006.

[5] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, 2006.

[6] I. Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium (CSF)*, pages 263–275, 2017.

[7] S. Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[8] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489): 375–389, 2010.