Differential Privacy in Communications

Matthew Hale

Department of Mechanical and Aerospace Engineering
University of Florida

AFOSR Center of Excellence Kickoff
May 14, 2019

Encryption Can Sometimes Be Restrictive

- Sometimes we want to share *some* information

Encryption Can Sometimes Be Restrictive

- Sometimes we want to share *some* information
- Example: coalitions may wish to share approximate locations

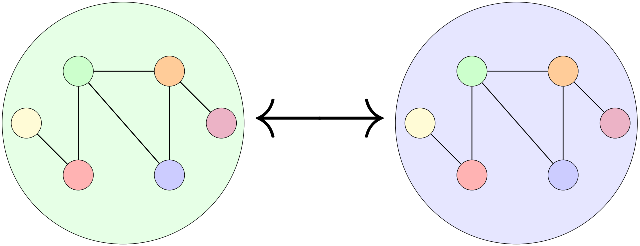

How should we make sensitive data private?

Fundamental Question

How can we share information but keep secrets in contested environments?

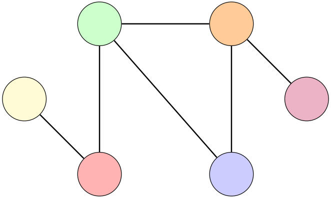

How should we make sensitive data private?

Fundamental Question

How can we share information but keep secrets in contested environments?

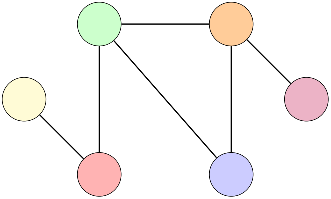

Goal

Develop theoretical tools for protecting data while sharing it.

Agents' Dynamics Generate Trajectories to Share

- Example: Agents in a coalition want to share their states with another coalition

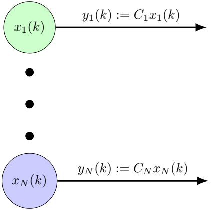

Agents' Dynamics Generate Trajectories to Share

- Example: Agents in a coalition want to share their states with another coalition

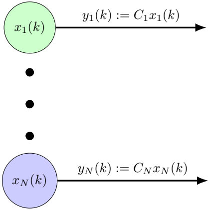

- No guarantee that the recipient only knows $y_i(k)$ at time k

Data Can be Aggregated and Processed

- We lose control of our data after sharing it

Data Can be Aggregated and Processed

- We lose control of our data after sharing it
- We cannot know what an adversary will do with what they receive
  - Aggregate it over time?
  - Filter it?

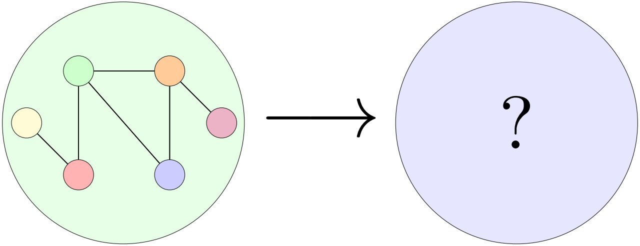

Data Can be Aggregated and Processed

- We lose control of our data after sharing it
- We cannot know what an adversary will do with what they receive
  - Aggregate it over time?
  - Filter it?

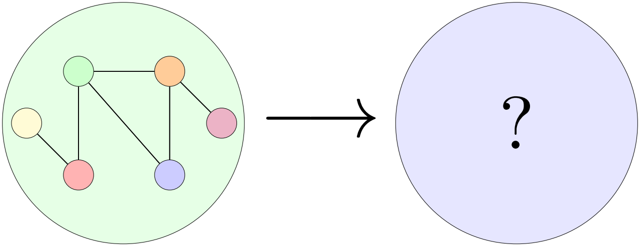

- Privacy must (somehow) account for this

How should we provide privacy?

How should we provide privacy?

Differential Privacy (DP)

DP is a privacy framework with a several key features:

- It offers a formal definition of “privacy”

How should we provide privacy?

Differential Privacy (DP)

DP is a privacy framework with a several key features:

- It offers a formal definition of "privacy"
- It is immune to post-processing
  - x private $\Rightarrow f(x)$ private for all f

How should we provide privacy?

Differential Privacy (DP)

DP is a privacy framework with a several key features:

- It offers a formal definition of "privacy"
- It is immune to post-processing
  - x private $\Rightarrow f(x)$ private for all f
- It is robust to side information

How should we provide privacy?

Differential Privacy (DP)

DP is a privacy framework with a several key features:

- It offers a formal definition of “privacy”
- It is immune to post-processing
  - x private $\Rightarrow f(x)$ private for all f
- It is robust to side information

- Used by:

Apple

Google

Uber

UBER

How should we provide privacy?

Differential Privacy (DP)

DP is a privacy framework with a several key features:

- It offers a formal definition of "privacy"
- It is immune to post-processing
  - x private $\Rightarrow f(x)$ private for all f
- It is robust to side information

- Used by: Apple Google Uber

UBER

DP Idea

Make "adjacent" state trajectories produce "similar" outputs

Fundamental Definitions in Differential Privacy

Adjacent trajectories in ℓ_p-spaces

We fix a constant $b > 0$ and define $\mathsf{Adj}_b : \ell_p^n \times \ell_p^n \to \{0, 1\}$ as

$$\mathsf{Adj}_b(x_1, x_2) = 1 \iff \|x_1 - x_2\|_{\ell_p} \leq b.$$

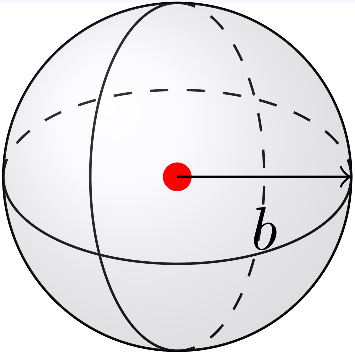

Similarity of Outputs

Fundamental Inequality of Differential Privacy

For adjacent state trajectories x_1 and x_2, we want the outputs y_1, y_2 to satisfy

$$\mathbb{P}(y_2) \leq e^{\epsilon}\mathbb{P}(y_1) + \delta,$$

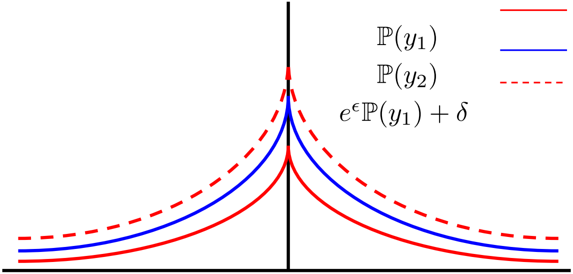

Similarity of Outputs

Fundamental Inequality of Differential Privacy

For adjacent state trajectories x_1 and x_2, we want the outputs y_1, y_2 to satisfy

$$\mathbb{P}(y_2) \leq e^{\epsilon}\mathbb{P}(y_1) + \delta,$$

This is the definition of (ϵ, δ)-differential privacy.

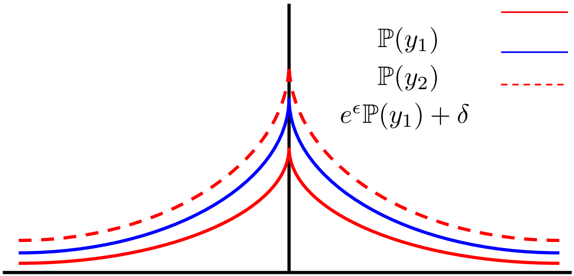

- Fix a probability space $(\Omega, \Sigma, \mathbb{P})$. Differential privacy is enforced by a *mechanism* of the form

$$M : \ell_p^n \times \Omega \to \ell_q^r.$$

Mechanisms for Differential Privacy

- Fix a probability space $(\Omega, \Sigma, \mathbb{P})$. Differential privacy is enforced by a *mechanism* of the form

$$M : \ell_p^n \times \Omega \to \ell_q^r.$$

- For us this will take the form

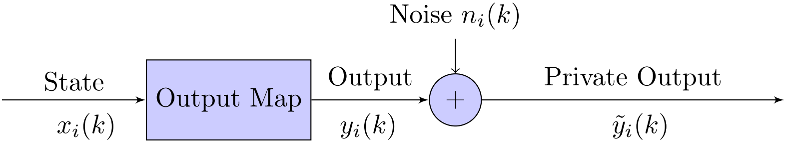

Privacy Noise is Calibrated by What We Share

Sensitivity

The p-norm sensitivity of a mapping $\mathcal{F}$ is

$$\Delta_p \mathcal{F} = \sup_{x_1, x_2 : \mathsf{Adj}_B(x_1, x_2)} \|\mathcal{F}(x_1) - \mathcal{F}(x_2)\|_{\ell_p}.$$

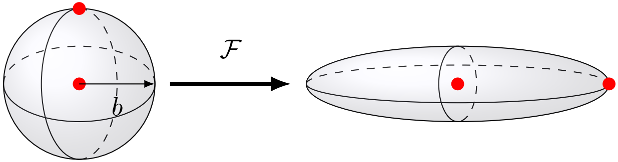

Privacy Noise is Calibrated by What We Share

Sensitivity

The p-norm sensitivity of a mapping $\mathcal{F}$ is

$$\Delta_p \mathcal{F} = \sup_{x_1, x_2 : \mathsf{Adj}_B(x_1, x_2)} \left\| \mathcal{F}(x_1) - \mathcal{F}(x_2) \right\|_{\ell_p}.$$

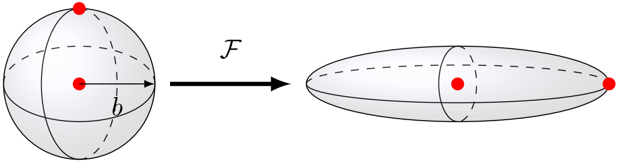

- For an agent sharing $y_i(k) := C_i x_i(k)$: $\Delta_p \mathcal{F} = s_1(C_i) b$

Privacy Noise is Calibrated by What We Share

Sensitivity

The p-norm sensitivity of a mapping $\mathcal{F}$ is

$$\Delta_p \mathcal{F} = \sup_{x_1, x_2 : \mathsf{Adj}_B(x_1, x_2)} \|\mathcal{F}(x_1) - \mathcal{F}(x_2)\|_{\ell_p}.$$

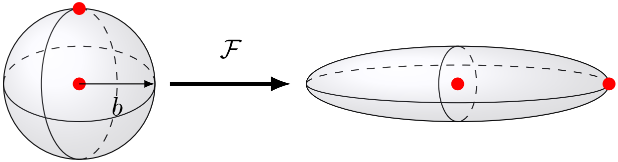

- For an agent sharing $y_i(k) := C_i x_i(k)$: $\Delta_p \mathcal{F} = s_1(C_i) b$
- We make it differentially private by adding noise $w(k) \sim \mathcal{N}\big(0, s_1(C_i) b \cdot \kappa(\epsilon, \delta)\big)$

Differential Privacy in Control

- Differential privacy has been applied to:
  - Kalman filtering

Differential Privacy in Control

- Differential privacy has been applied to:
  - Kalman filtering
  - Distributed linear-quadratic control

Differential Privacy in Control

- Differential privacy has been applied to:
  - Kalman filtering
  - Distributed linear-quadratic control
  - Consensus problems

Differential Privacy in Control

- Differential privacy has been applied to:
  - Kalman filtering
  - Distributed linear-quadratic control
  - Consensus problems
  - Optimization in several forms

Differential Privacy in Control

- Differential privacy has been applied to:
  - Kalman filtering
  - Distributed linear-quadratic control
  - Consensus problems
  - Optimization in several forms
- Always involves introducing randomness

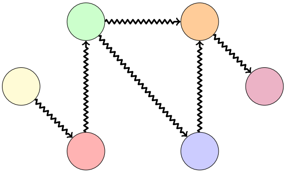

Next Steps

- Contested environments have asynchronous communications
- How can we use asynchronous private information?

Next Steps

- Contested environments have asynchronous communications
- How can we use asynchronous private information?
- How can we privatize new data types, such as sets?