The Role of Interactivity in Local Differential Privacy

Matthew Joseph

Jieming Mao

Seth Neel

Aaron Roth

Q: How much does interaction matter in local differential privacy?

Q: How much does interaction matter in local differential privacy?

A: It depends.

Outline

1. Differential Privacy

2. Local Differential Privacy
   a. Result 1: Limits of full interaction
   b. Result 2: Power of full interaction

Outline

1. Differential Privacy

2. Local Differential Privacy
    a. Result 1: Limits of full interaction
    b. Result 2: Power of full interaction

Differential Privacy [DMNS06] in Words

Property of a randomized algorithm $\mathcal{A}$

Small changes in input ⇒ small changes in output

Add noise to output to obscure any small changes in input

Differential Privacy in Math

Definition: Two databases X and X' are *neighbors* if they differ in at most one entry. Randomized algorithm $A: \mathcal{X} \to \mathcal{Y}$ is (ε, δ)-*differentially private* (DP) if, for all neighbors X and X', and for all $Y \subset \mathcal{Y}$,

$$P[A(X) \text{ in } Y] \leq e^{\varepsilon} P[A(X') \text{ in } Y] + \delta.$$

Why is Differential Privacy "Private"?

Think of as X and X' "database with your data" and "database without your data"

If A is DP, then $A(X) \approx_{(\varepsilon, \delta)} A(X')$, so the computation is (almost) agnostic to your presence

Central DP Learning From Data

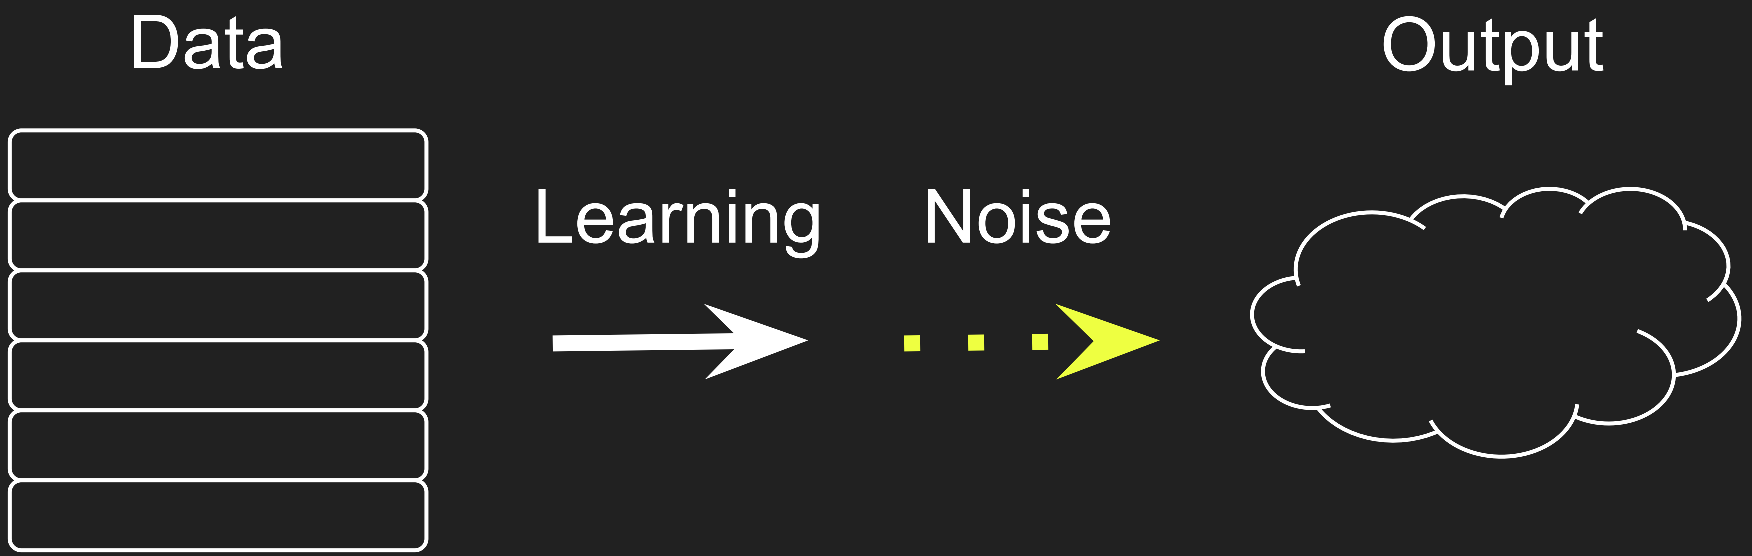

Useful DP Properties

Composition: For $A = (A_1, \ldots, A_k)$ where each A_i is $(\varepsilon_i, \delta_i)$-DP, A is $(\Sigma_i\varepsilon_i, \Sigma_i\delta_i)$-DP.

Useful DP Properties

Composition: For $A = (A_1, \ldots, A_k)$ where each A_i is $(\varepsilon_i, \delta_i)$-DP, A is $(\Sigma_i \varepsilon_i, \Sigma_i \delta_i)$-DP.

Robust to Post-Processing: If A is (ε, δ)-DP, then for any function f, $f(A)$ is also (ε, δ)-DP.

Key Takeaways About Differential Privacy

DP algorithms map similar databases to similar output distributions

Add randomness somewhere for privacy

Modular, can cut and paste

Outline

1. Differential Privacy

2. Local Differential Privacy
   a. Result 1: Limits of full interaction
   b. Result 2: Power of full interaction

Central DP Learning From Data

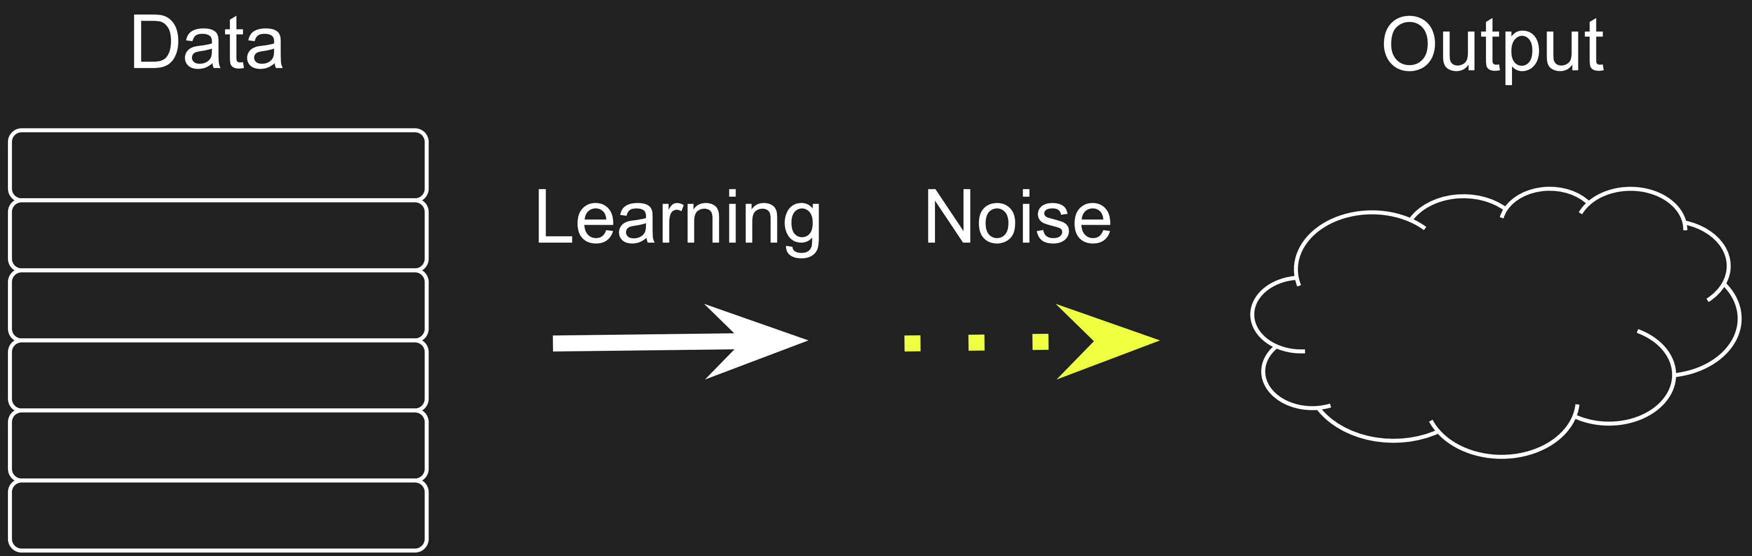

Local DP [DMNS06] Learning From Data

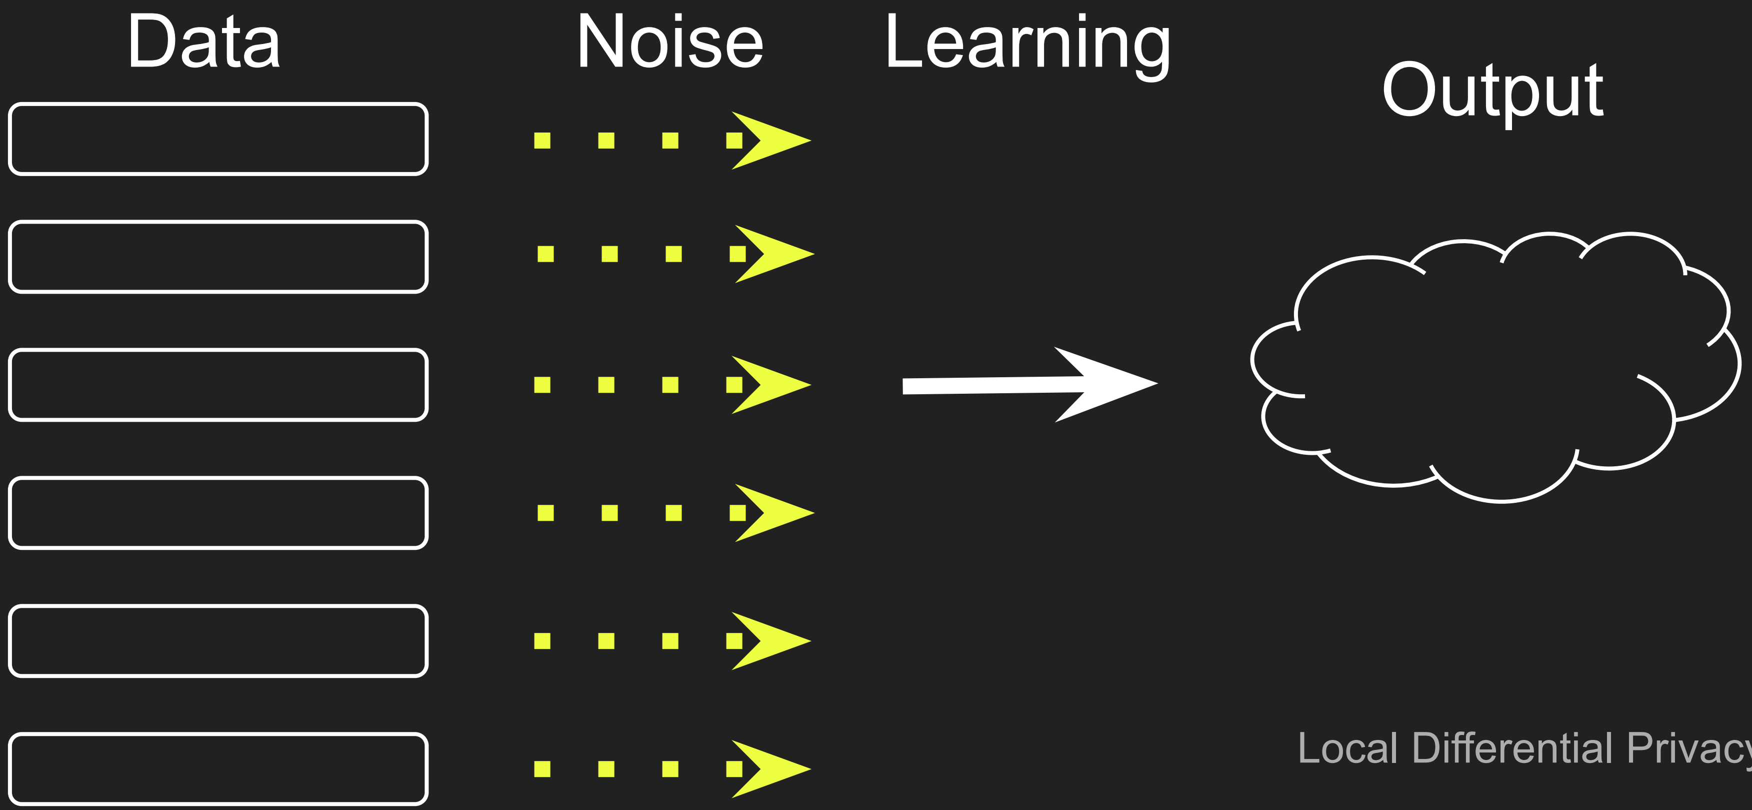

Local DP in Words

No more central database, users keep their data

Protocol $\mathcal{A}$ learns about the data through public communication with users

Users send responses through *randomizers* $\mathcal{R}$

Local DP in Math

Definition: Protocol $\mathcal{A}$ is **(ε,δ)**-locally differentially private (LDP) if the transcript of communications it generates is an **(ε,δ)**-DP function of the user data.

LDP: Pros and Cons

Pros:

- ✓ Data never leaves user device, only DP outputs
- ✓ Don’t have to store any private data

LDP: Pros and Cons

Pros:

- ✓ Data never leaves user device, only DP outputs
- ✓ Don’t have to store any private data

Cons:

- ✗ More noise → worse utility
- ✗ Don’t get to store any private data

Q: How much does interaction matter for local differential privacy?

A: It depends.

Types of LDP Interactivity

Definition: Protocol $\mathcal{A}$ is *noninteractive* if all users speak once, simultaneously and independently.

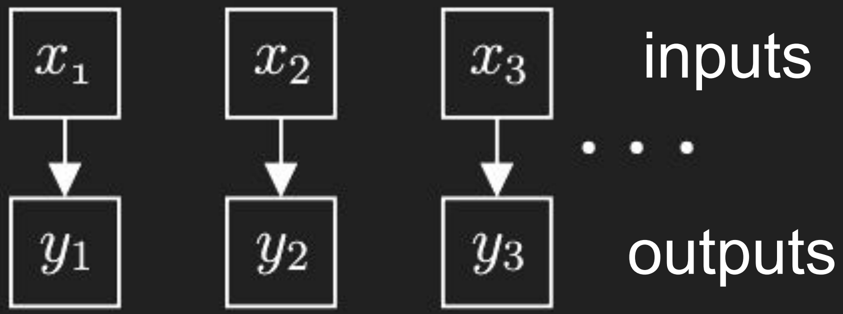

Types of LDP Interactivity

Definition: Protocol A is *sequentially interactive* [DJW13] if all users speak once (possibly in multiple rounds).

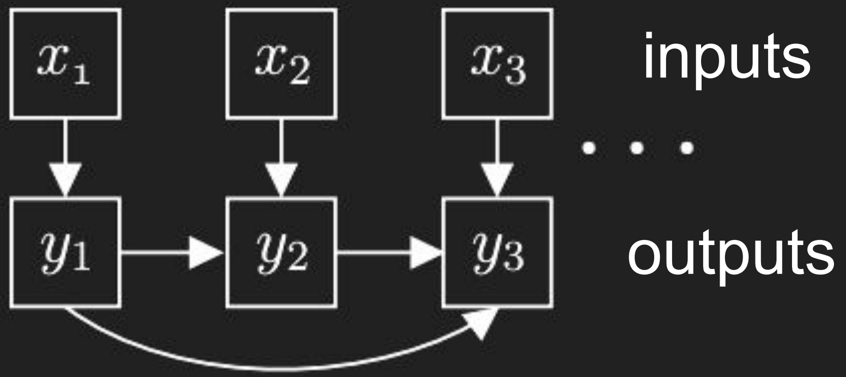

Types of LDP Interactivity

Definition: Protocol A is *fully interactive* if users may interact arbitrarily (possibly speak multiple times, in multiple rounds).

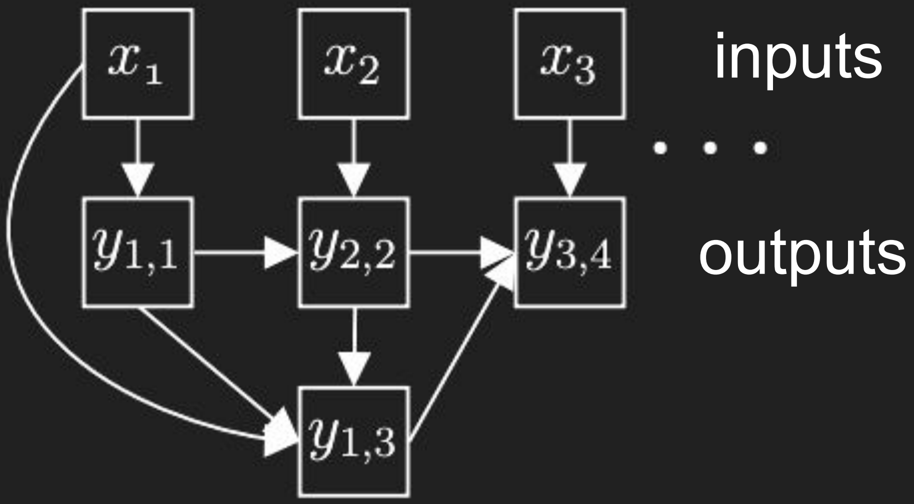

Types of LDP Interactivity

Noninteractive

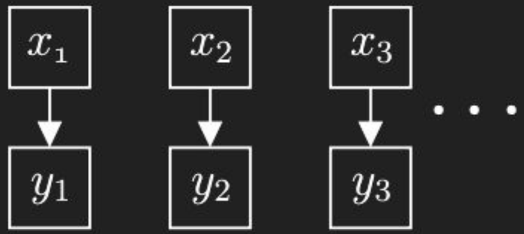

Sequentially Interactive

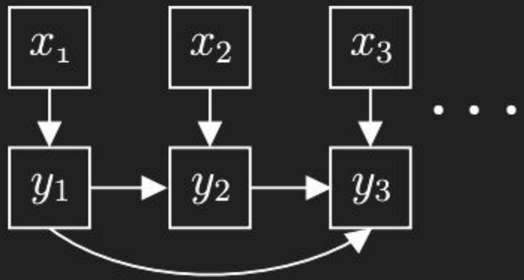

Fully Interactive

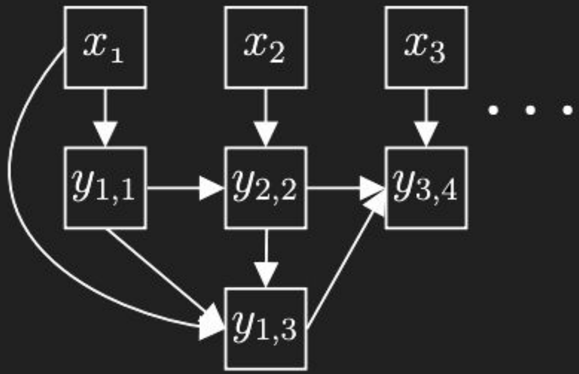

Types of LDP Interactivity

Noninteractive

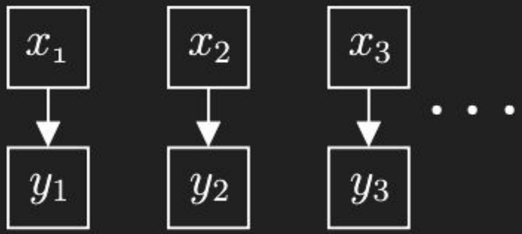

rounds = 1

Sequentially Interactive

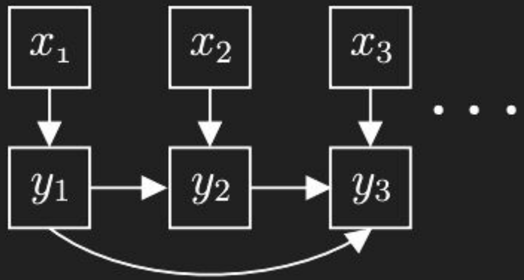

rounds
≤ # users

Fully Interactive

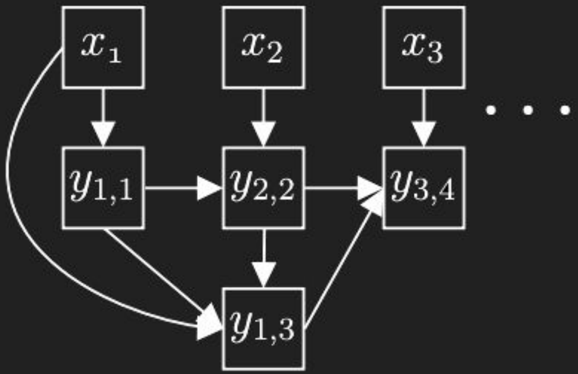

rounds = ???

Types of LDP Interactivity

Noninteractive

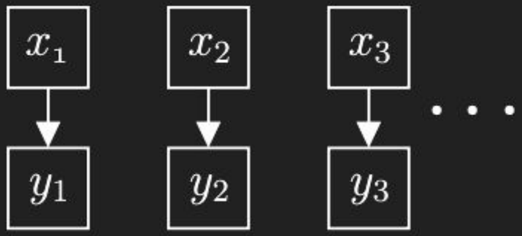

Sequentially Interactive

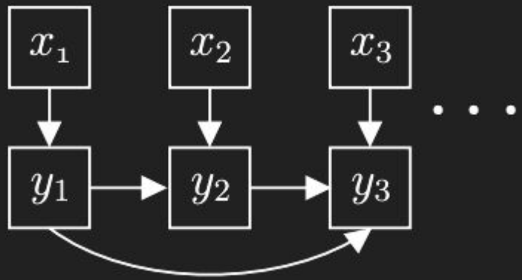

Fully Interactive

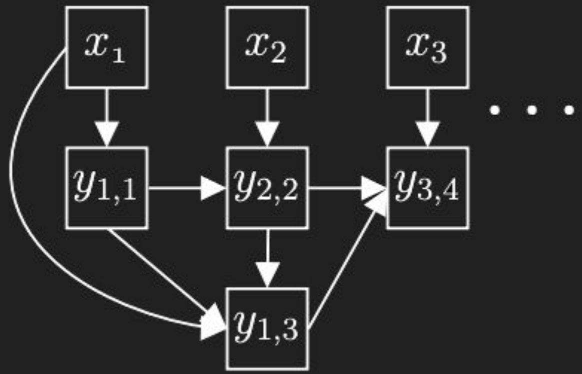

[KLNRS08]
[DF18]

Types of LDP Interactivity

Noninteractive

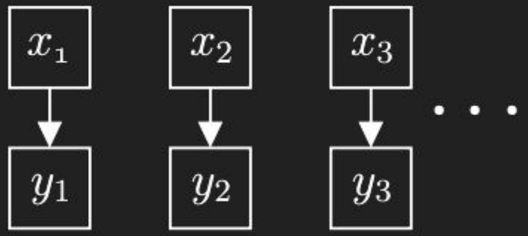

Sequentially Interactive

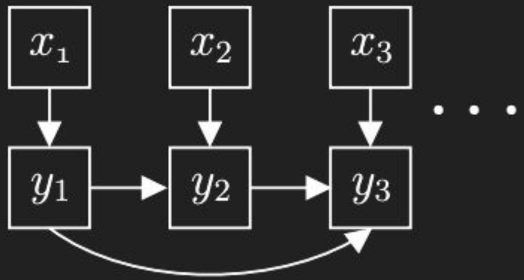

[KLNRS08]
[DF18]

Fully Interactive

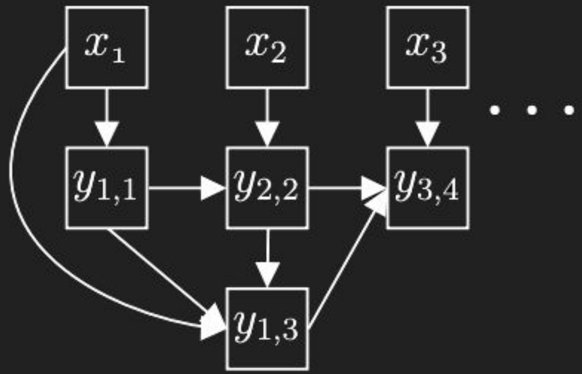

This Work

Outline

1. Differential Privacy

2. Local Differential Privacy
   a. Result 1: Limits of full interaction
   b. Result 2: Power of full interaction

Result 1: Limits of Full Interaction

Theorem (Informal): Any fully interactive protocol A_F can be converted into an identical sequentially interactive protocol A_S, with a controlled increase in sample complexity.

Result 1: Limits of Full Interaction

Theorem (Informal): Any fully interactive protocol A_F can be converted into an identical sequentially interactive protocol A_S, with a controlled increase in sample complexity.

Increase is sometimes small, sometimes large.
Depends on *compositionality*.

Compositionality

Composition: cut and paste randomizers together, privacy parameters add up

Any algorithm analyzed this way is 1-compositional

Not the only way to analyze!

Compositionality Example

Each user i has private datum $x_i \in \{1, 2, \ldots, k\}$, operator wants to compute counts

Protocol: each user outputs $y_i \in \{0,1\}^k$ where

- $y_i^j \sim \mathrm{Ber}(1/[e^\varepsilon+1])$ if $j \neq x_i$
- $y_i^j \sim \mathrm{Ber}(e^\varepsilon/[e^\varepsilon+1])$ otherwise

Compositionality Example

Each user i has private datum $x_i \in \{1, 2, \ldots, k\}$, operator wants to compute counts

If $x_i = 4$, $Y_i = \downarrow\ \downarrow\ \downarrow\ \uparrow\ \downarrow\ \downarrow\ \downarrow\ \cdots\ \downarrow\ \downarrow$

Compositionality Example

Protocol: each user outputs $y_i \in \{0,1\}^k$ where

- $y_i^j \sim \mathrm{Ber}(1/[e^{\varepsilon}+1])$ if $j \neq x_i$
- $y_i^j \sim \mathrm{Ber}(e^{\varepsilon}/[e^{\varepsilon}+1])$ otherwise

Composition way: k total ε-randomizers

… so $k\varepsilon$-LDP

Compositionality Example

Protocol: each user outputs $y_i \in \{0,1\}^k$ where

- $y_i^j \sim \mathrm{Ber}(1/[e^{\varepsilon}+1])$ if $j \neq x_i$
- $y_i^j \sim \mathrm{Ber}(e^{\varepsilon}/[e^{\varepsilon}+1])$ otherwise

Direct way: $$\frac{P[y_i^j = y \mid x_i = x]}{P[y_i^j = y \mid x_i = x']} \leq \frac{e^{\varepsilon}/[e^{\varepsilon}+1]}{1/[e^{\varepsilon}+1]} = e^{\varepsilon}$$

… so ε-LDP. Took advantage of histogram data structure.

Compositionality

Definition: The *compositionality* of an LDP protocol is the multiplicative factor by which its minimal composition privacy guarantee exceeds its overall privacy guarantee.

Previous algorithm is k-compositional.

Result 1: Limits of Full Interaction

Theorem: Any fully interactive ε-LDP k-compositional protocol A_F can be converted into an identical 3ε-LDP sequentially interactive protocol A_S on, w.p. $1-\beta$, $O(e^{\varepsilon}(nk + \sqrt{nk}\log(\frac{1}{\beta})))$ samples.

Result 1: Limits of Full Interaction

Theorem: Any fully interactive ε-LDP k-compositional protocol A_F can be converted into an identical 3ε-LDP sequentially interactive protocol A_S on, w.p. $1-\beta$, $O(e^{\varepsilon}(nk + \sqrt{nk}\log(\frac{1}{\beta})))$ samples.

Is this tight?

Outline

1. Differential Privacy

2. Local Differential Privacy
    a. Result 1: Limits of full interaction
    b. Result 2: Power of full interaction

Result 2: Powers of Full Interaction

Yes! (up to log factors)

Theorem: There exists a fully interactive d-compositional ε-LDP protocol that solves *multi-party pointer jumping* in $\tilde{O}(d^2)$ samples, but any sequentially interactive (ε,δ)-LDP protocol requires $\tilde{\Omega}(d^3)$ samples.

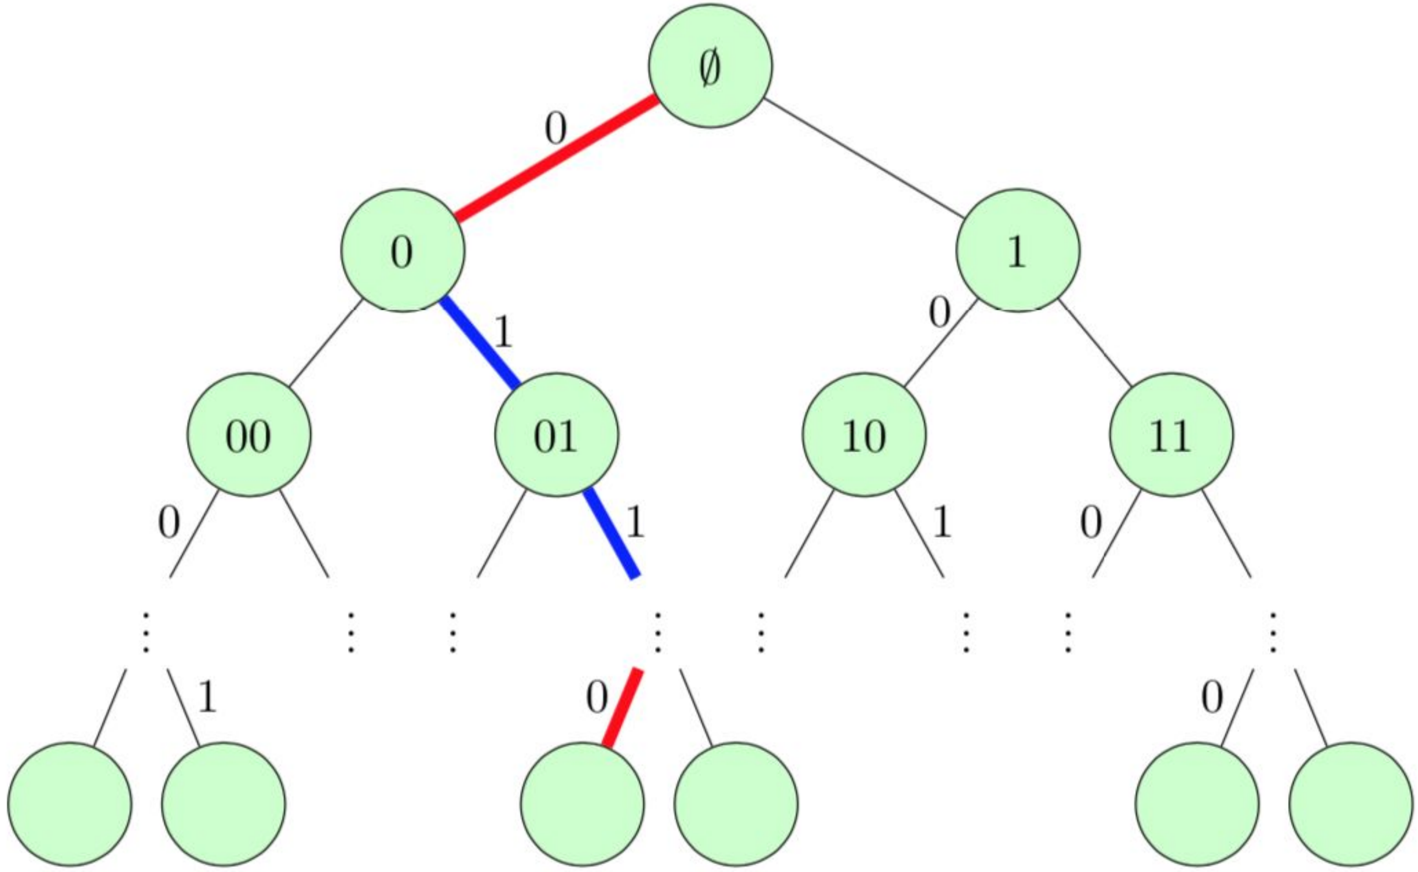
∅
0
1
0
1
00
01
10
11
0
1
1
0
1
0
0

Result 2: Powers of Full Interaction

Yes! (up to log factors)

Theorem: There exists a fully interactive d-compositional ε-LDP protocol that solves *multi-party pointer jumping* in $\tilde{O}(d^2)$ samples, but any sequentially interactive (ε,δ)-LDP protocol requires $\tilde{\Omega}(d^3)$ samples.

Can't avoid compositionality dependence.

Q: How much does interaction matter for local differential privacy?

A: It depends *on compositionality*.

Takeaways

- Can convert full to sequential, sample complexity blowup proportional to compositionality

Takeaways

- Can convert full to sequential, sample complexity blowup proportional to compositionality
  - Full interaction can only beat sequential interaction when the solution is highly compositional

Takeaways

- Can convert full to sequential, sample complexity blowup proportional to compositionality
  - Full interaction can only beat sequential interaction when the solution is highly compositional
- Unavoidably highly compositional (but also highly specific) problems exist

Takeaways

- Can convert full to sequential, sample complexity blowup proportional to compositionality
  - Full interaction can only beat sequential interaction when the solution is highly compositional
- Unavoidably highly compositional (but also highly specific) problems exist
- Didn’t mention: local-central separation for simple hypothesis testing

Takeaways

- Can convert full to sequential, sample complexity blowup proportional to compositionality
  - Full interaction can only beat sequential interaction when the solution is highly compositional
- Unavoidably highly compositional (but also highly specific) problems exist
- Didn't mention: local-central separation for simple hypothesis testing

arxiv.org/abs/1904.03564

References

1. [BBGN19] “The Privacy Blanket of the Shuffle Model”. Balle, Bell, Gascon, Nissim. CRYPTO.
2. [DF19] “Learning without Interaction Requires Separation”. Daniely and Feldman. NeurIPS.
3. [DJW13] “Local Privacy, Data Processing Inequalities, and Statistical Minimax Rates”. Duchi, Jordan, Wainwright. FOCS.
4. [DMNS06] “Calibrating Noise to Sensitivity in Private Data Analysis”. Dwork, Mcsherry, Nissim, Smith. TCC.
5. [KLNRS08] “What Can We Learn Privately?”. Kasiviswanathan, Lee, Nissim, Raskhodnikova, Smith. STOC.