Differential privacy

Some material taken from here, here

Final project checkins

- Schedule a 45 min meeting with me next week
- Look on CMU Google calendar & send me an invite!

Secure multiparty computation

- Allows multiple non-colluding parties to compute on their joint data, without revealing their input data or any intermediate results
- One big drawback: *final result* usually needs to be revealed to all parties
- What to do if you must release some result to some untrusted entities?

Differential privacy

- How to publish data about someone *while protecting their privacy?*
- Cynthia Dwork et al. Calibrating Noise to Sensitivity in Private Data Analysis (TCC 2006)
- Idea: add randomness/noise to the final result so that the exact value is not revealed

DP intuition

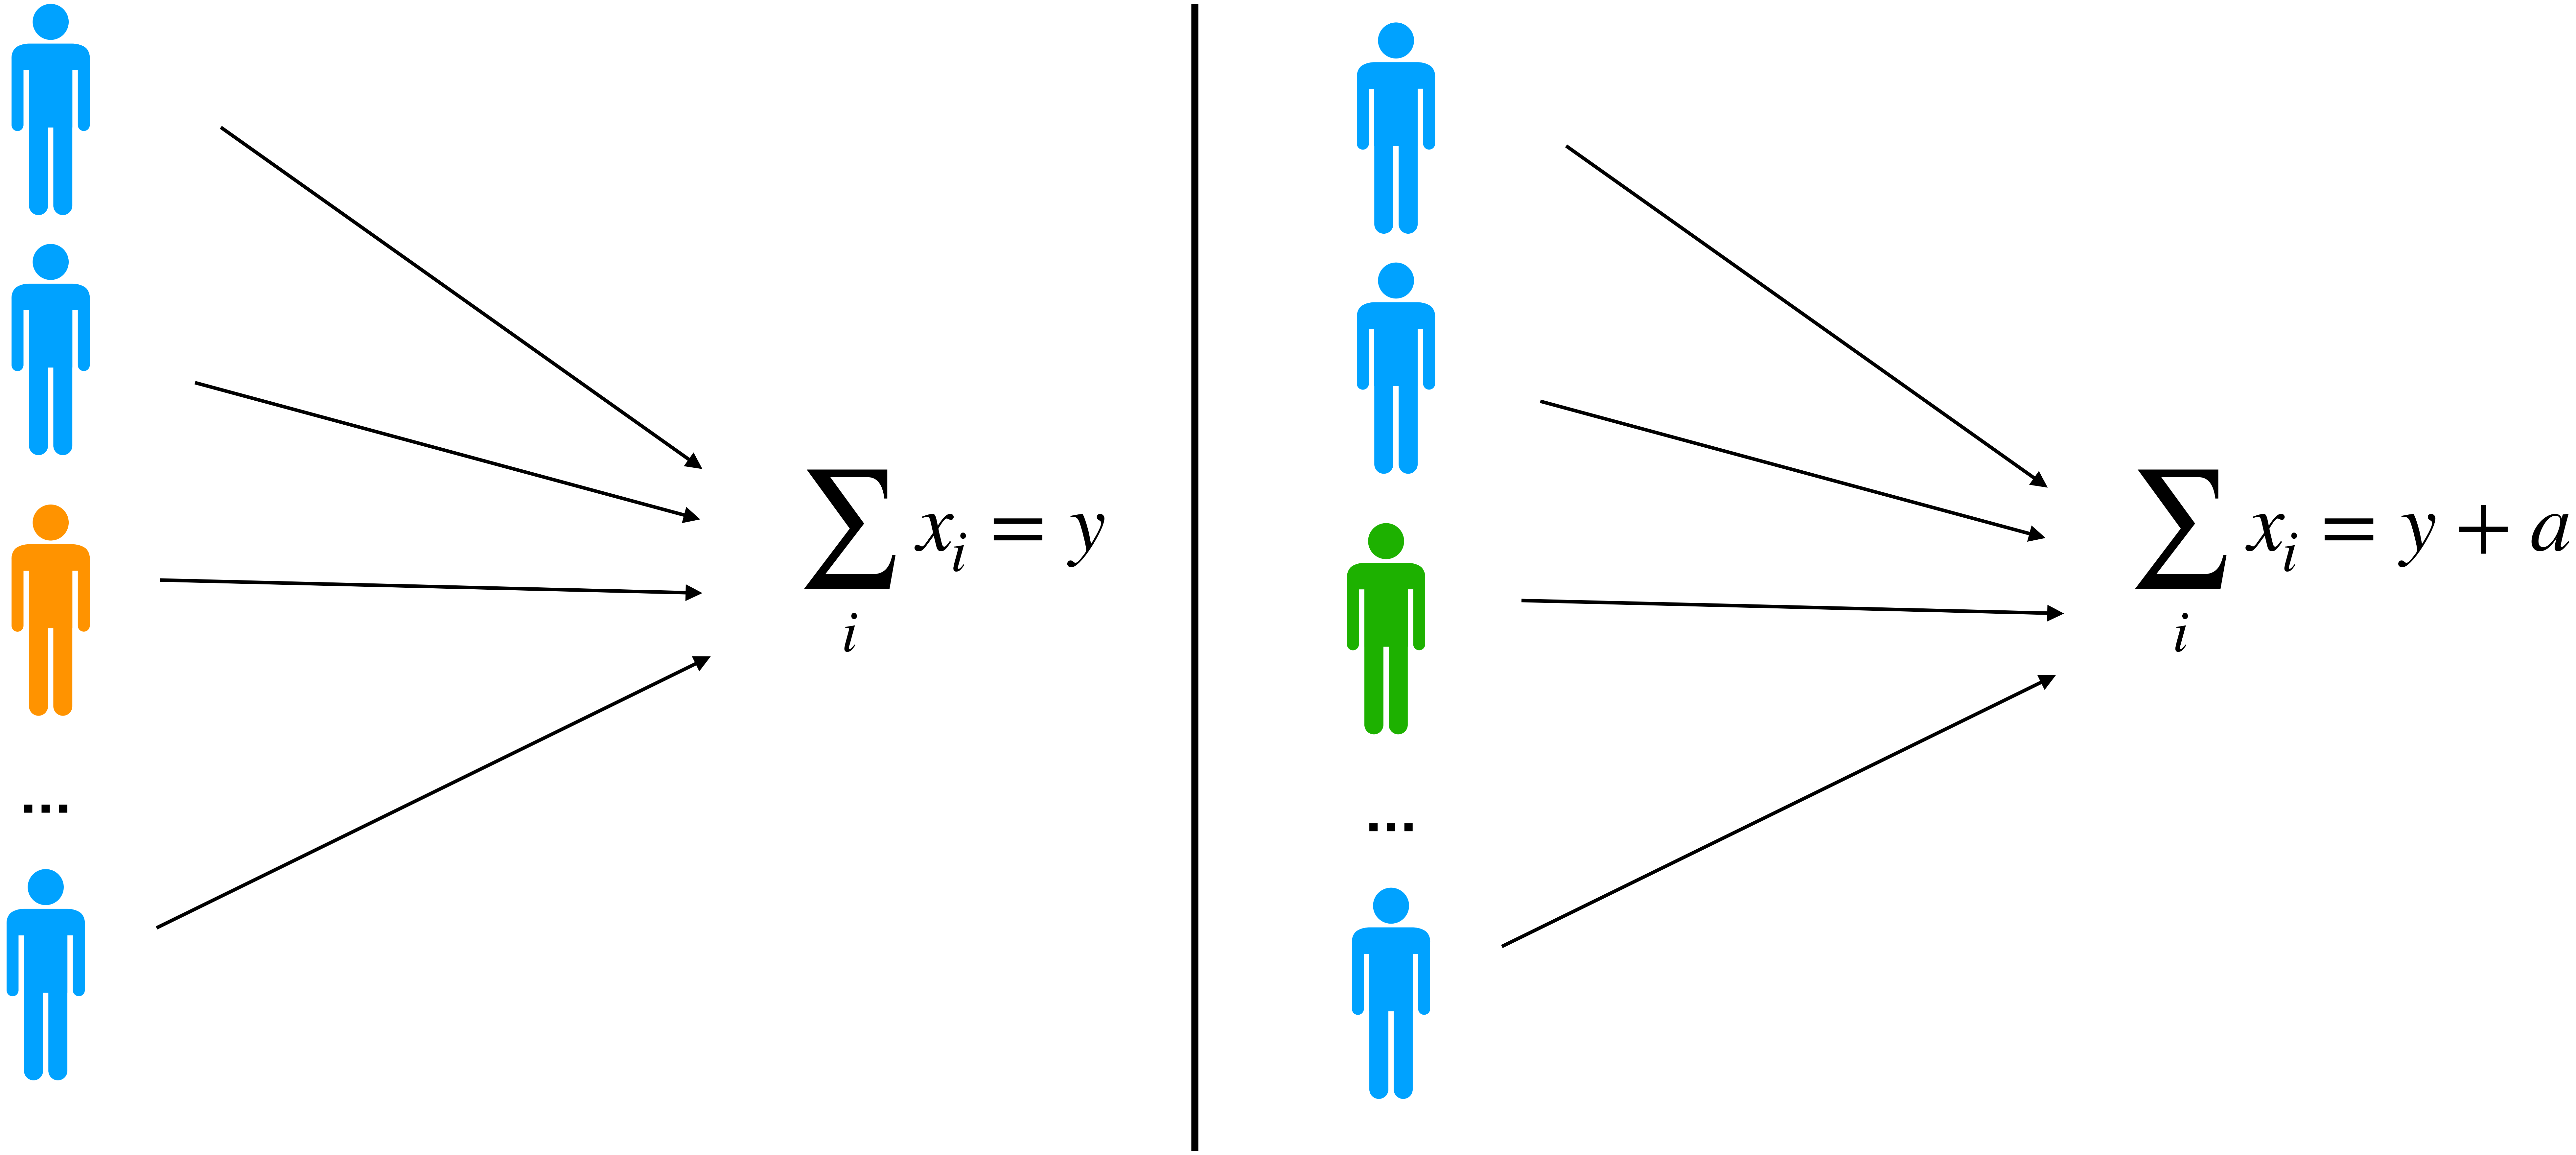

Anyone who sees the output (e.g., data scientist)
will learn about an individual user's data

DP definitions

- DP promises to protect an individual from any additional harm they might face due to their data being included in a database
  - Might still be harmed by the final result, but their decision to participate shouldn't increase the harm
  - Smoking causes cancer: result will increase insurance cost of Alice, but impact on Alice is independent of whether she participated in the study
- A mechanism A is ϵ-differentially private if for all databases D_1 and D_2 which differ in one individual, $P[A(D_1) = O] \leq e^{\epsilon} \cdot P[A(D_2) = O]$, for all values of O
  - Output of the process should be similar if you change one individual
  - ϵ captures the privacy-utility tradeoff

Randomized response

- How to survey how many people are illegal drug users?
    - Ask them directly, they might lie to you
- Why not add noise to their responses? Each person will
    - Flip a coin
    - If heads, then flip a second coin and respond “Yes” if heads, “No” if tails
    - If tails, respond truthfully
- Provides plausible deniability - even if answer is “Yes”, it might have been offered because first & second coins were both heads

Randomized response

- **Theorem:** The randomized response described is ϵ-differentially private, where $e^{\epsilon} = 3$ ($\epsilon \approx 1.1$).

- **Proof:** Fix a single individual. With 50% probability they will answer truthfully, and 50% probability they will answer randomly

    - Drug user: 75% chance to answer “Yes”; non-drug user: 25% chance to answer “Yes”

    - $P[A(Yes) = Yes]/P[A(No) = Yes] = 0.75/0.25 = 3$

Some properties of DP

- Privacy loss is quantified
    - $\epsilon = 0$ means perfect privacy; lower = better privacy but worse utility
- DP is immune to post processing: a data analyst without additional knowledge about the private database cannot compute a function of the output of a DP algorithm and make it less private
- Composition
    - Suppose two algorithms A and B are ϵ-differentially private
    - Publishing the result of both is 2ϵ-differentially private (two algorithms are independent so they have their own randomness)
    - Privacy budget
        - One individual is in C statistics
        - If each release is ϵ_i-differentially private, then overall release is ϵ-differentially private if $\sum_i \epsilon_i = \epsilon$

DP in the real world

- Apple uses DP for collecting telemetry data on iOS and MacOS
  - 2 submissions per day per user
  - $\epsilon = 8$ for each submission
- Microsoft collects number of minutes that Windows 10 users use in each app
  - $\epsilon = 0.7$, once every 6 hours

Today: RAPPOR

Next time: DP in ML system

- So far we've seen randomized response mechanism, which is a *locally differently private* mechanism
- Another model: *global differential privacy*
  - A trusted entity has collected all of the data, and adds noise before releasing the result
  - Census
  - Google trains and release an ML model