CSE 163
Data & Privacy

Adrian Salguero
Spring 2026

💭 Icebreaker (discuss with neighbors):
Humans no longer need sleep to function! How would you spend that time?
Add to our Slido!

slido.com
#cse163

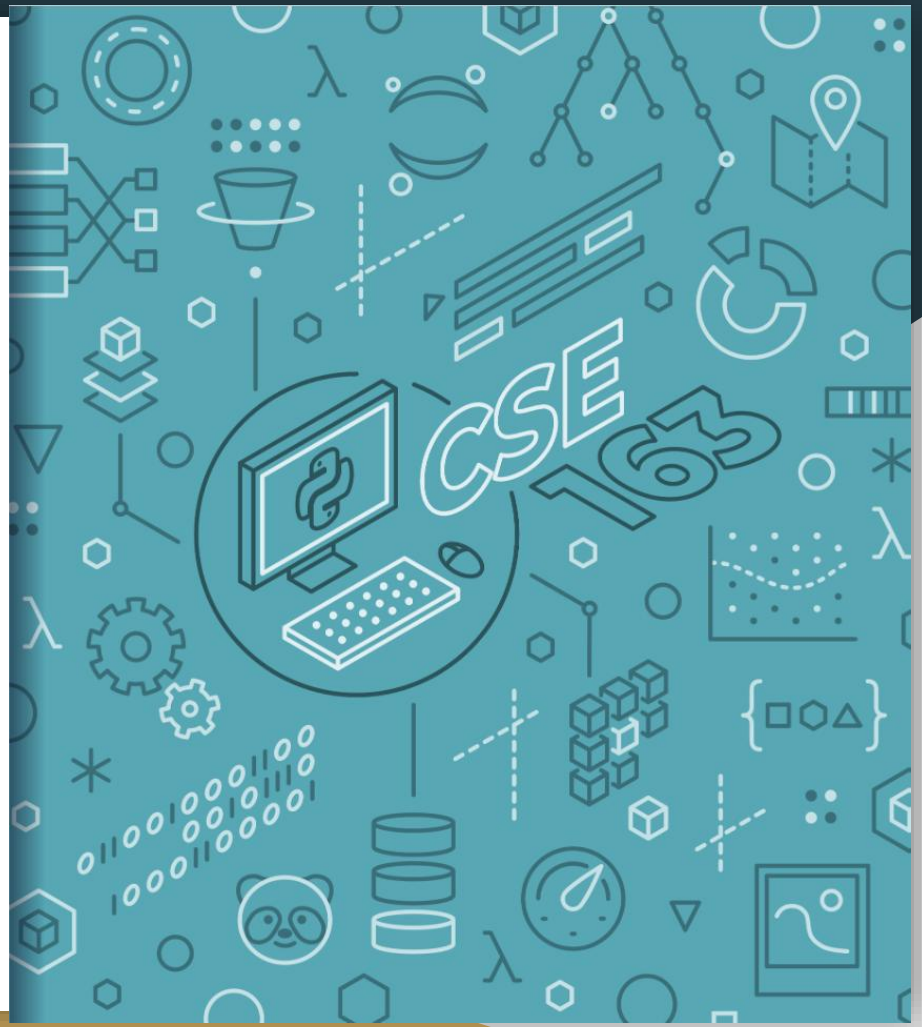

Announcements

- **Take Home Assessment 5: Mapping** due Tuesday May 26th at 11:59pm
- **Reading Assignment 5** due Tuesday, May 26th at 11:59pm!
- **Lesson 23 Canvas Quiz** due tonight at 11:59pm!
- **Project Part 3** due June 1st at 11:59pm on Gradescope!

Anonymous Data Isn't

- In the mid-1990s, an insurance group in Massachusetts published anonymous records of hospital visits with attributes like name, address, social security removed but left in demographic information (e.g. zip code, gender)
- Turns out this data release was not so anonymous!
  - Latanya Sweeney (an MIT PhD student) was able to link demographic information in hospital data with voter rolls. Found which hospital record corresponded to the governor (and mailed his medical records to his office)
- Sweeney estimates 87% of the US is uniquely identified by knowing three pieces of information: date of birth, sex, and zip code

k-anonymity

- K-anonymity: A first definition of privacy by Sweeney that requires every query results in at least k people in the dataset
  - Achieved by removing columns or fuzzing values

Name	Age	Gender	Zip Code	Smoker	Diagnosis
*	60–70	Male	191**	Y	Heart disease
*	60–70	Female	191**	N	Arthritis
*	60–70	Male	191**	Y	Lung cancer
*	60–70	Female	191**	N	Crohn's disease
*	60–70	Male	191**	Y	Lung cancer
*	***50–60***	***Female***	191**	N	HIV
*	50–60	Male	191**	Y	Lyme disease
*	50–60	Male	191**	Y	Seasonal allergies
*	***50–60***	***Female***	191**	N	Ulcerative colitis

Differential Privacy

- A stronger notion of privacy that guarantees how much information you can learn about a person.
- Consider two worlds:
  - Individual A participates in a study
  - Individual A does not participate in a study
- If the results of the study are similar, we say it respects differential privacy

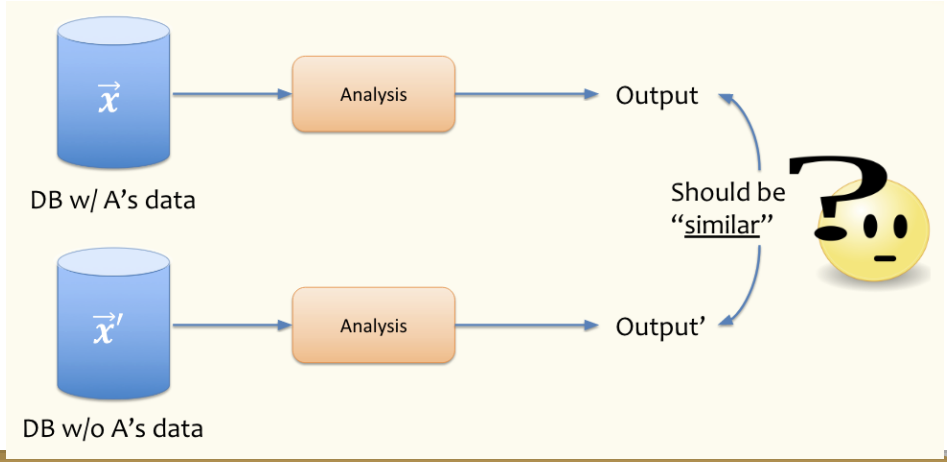

Differential Privacy

- Say an algorithm or analysis is ε-differentially private if results with or without any single person in the dataset are “at most ε” off.

- Defining how close results are is a little complex, but is a statement of probabilities
- If $\varepsilon = 0$, require results to be exactly the same
- If ε is small, require results to be very similar
- If ε is large, require more deviation in results (less privacy)

- Two methods for commonly achieving ε-differential privacy
  - Jittering Result (Laplace Mechanism)
  - Randomized Response

Jittering

- Take a result of analysis and add a small amount of random noise to the result
  - Example: Report average age of census but add a small random number to it

- Specifically if you add noise that follows a Laplace distribution with parameter ε, you can achieve ε-differential privacy.
  - See below for ε=0.5 (red dashes), ε=1 (blue solid), ε=2 (purple dots)

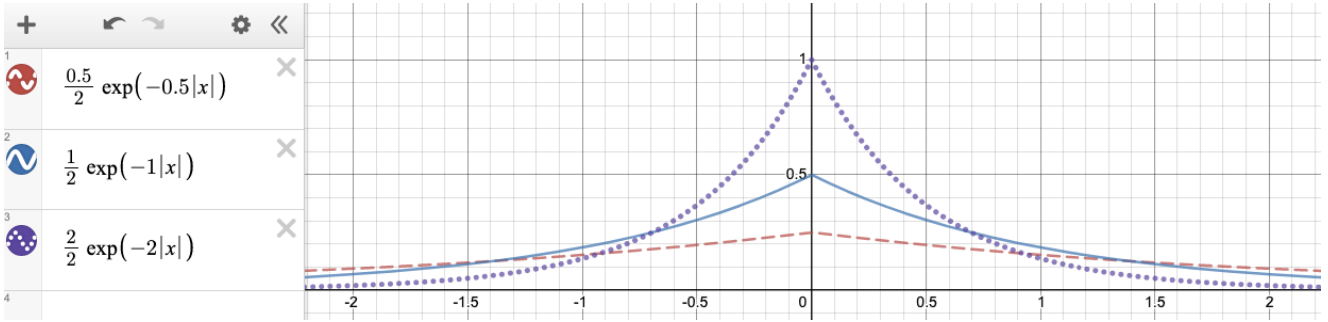

Randomized Response

- What if we don't trust the data collector with our data?
  - Even with differentially private statistics published, they still have access to the raw data. What if they get hacked?

- Change the differential privacy mechanism to be done locally rather than centrally!

- Differentially Private Polling Procedure
  - Call up a person. Ask them to flip a coin (don't say the result)
  - If Heads, tell us their honest answer to question ("Yes" or "No")
  - If Tails, flip the coin again
    - If Heads, report "Yes"
    - If Tails, report "No"

- ***Key idea***: Can learn aggregate trends without knowing true result of the individual

Randomized Response Analysis

- ***Key property***: People tell the truth ¾ of the time and lie ¼ of the time. ½ of the time they are honest, and then half of the time they tell us a random answer that lines up with the truth.

- To see why this work, suppose we know the answer is “Yes” for ⅓ of people. How many “Yes” responses would we expect in this procedure?
  - ⅓ of the population has true answer “Yes”. ¾ of them will tell us the truth so we will get a total of ¼ of responses being honest “Yes”es
  - ⅔ of the population has the true answer “No”, but ¼ of the time they will randomly tell us “Yes”. This means we would expect ⅙ of the population to lie and tell us “Yes”
  - Total of “Yes” received (on average): ¼ + ⅙ = 5/12
- In general, work backwards to solve for underlying probability

Case Studies Exercise

- With the people around you: choose the COVID-19 or the “Rides of Glory” case study from the lesson.

- Create a slide in the Google Slides document and share your responses to the “Food for Thought” questions in the lesson

- Finished early? Additional discussion questions to you can discuss and respond to:
  - Both case studies involve the tracking of location data. What other kinds of data would you be concerned about regarding privacy?
  - Both case studies report group patterns and aggregates rather than individuals. Does this change how you view the privacy concerns?