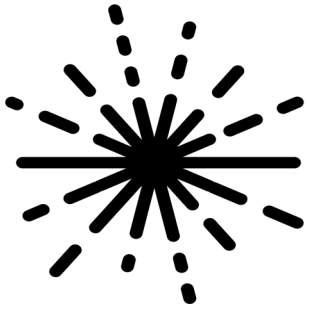

CS2080: Applied Privacy for Data Science
Conclusions

James Honaker, Priyanka Nanayakkara, Salil Vadhan
School of Engineering & Applied Sciences
Harvard University

April 30, 2025

Announcements

- Final paper drafts due: today!
- Poster session and party: Thu 5/8, 9am-12pm, SEC
  - Plan 1min presentation from each group member (on different aspects of the project)
- Final paper revision due: Thu 5/8

Privacy Risks

- Deidentified data can often be reidentified.
- Naïve query systems are subject to differencing-style attacks.
- Releasing too many aggregate statistics allows for reconstruction or membership attacks (Census, Diffix).
- AI models can memorize their training data

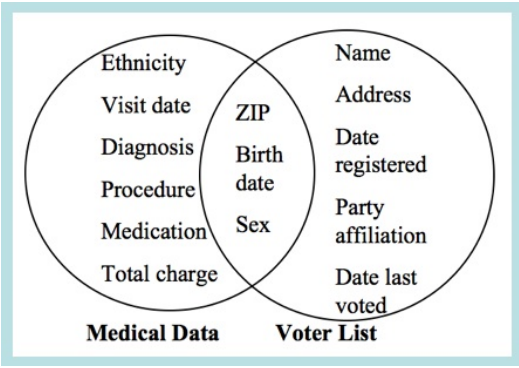

[Sweeney `97]

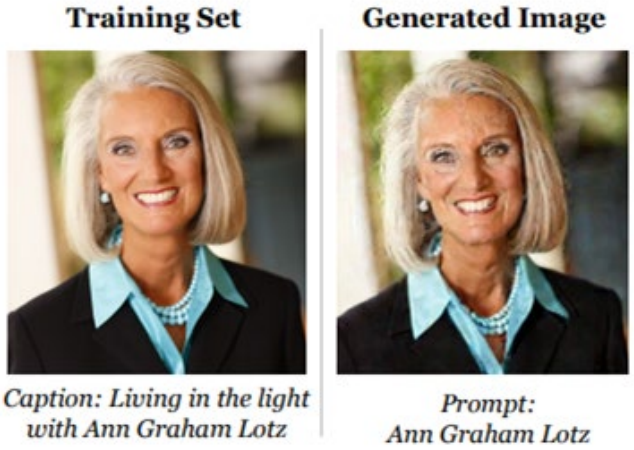

[Carlini, Hayes, Nasr et al. 2023]

Definition of Differential Privacy

- Strong privacy definition.
- Compatible with many statistical analyses.
- Ensures that “individual-level information” does not leak.
- Applies regardless of adversary’s auxiliary information.
- Adversary can be external “analysts” (centralized DP) or aggregator (local DP) or in between (federated and shuffle).

But:

- Adversary may still infer sensitive attributes.
- Not applicable when utility requires individual-level data.
- “Privacy” has many other meanings beyond what DP captures.

Core Properties

DP is closed under post-processing:

"No adversary can break the privacy guarantee"

Differentially private mechanisms compose:

"The epsilons add up"

Group privacy:

"The level of privacy degrades linearly with group size"

Composition of DP

- DP and variants (pure, approximate, zCDP, moments accountant) satisfy composition thms.
- Leads to tradeoff of # queries vs. accuracy (vs. privacy)

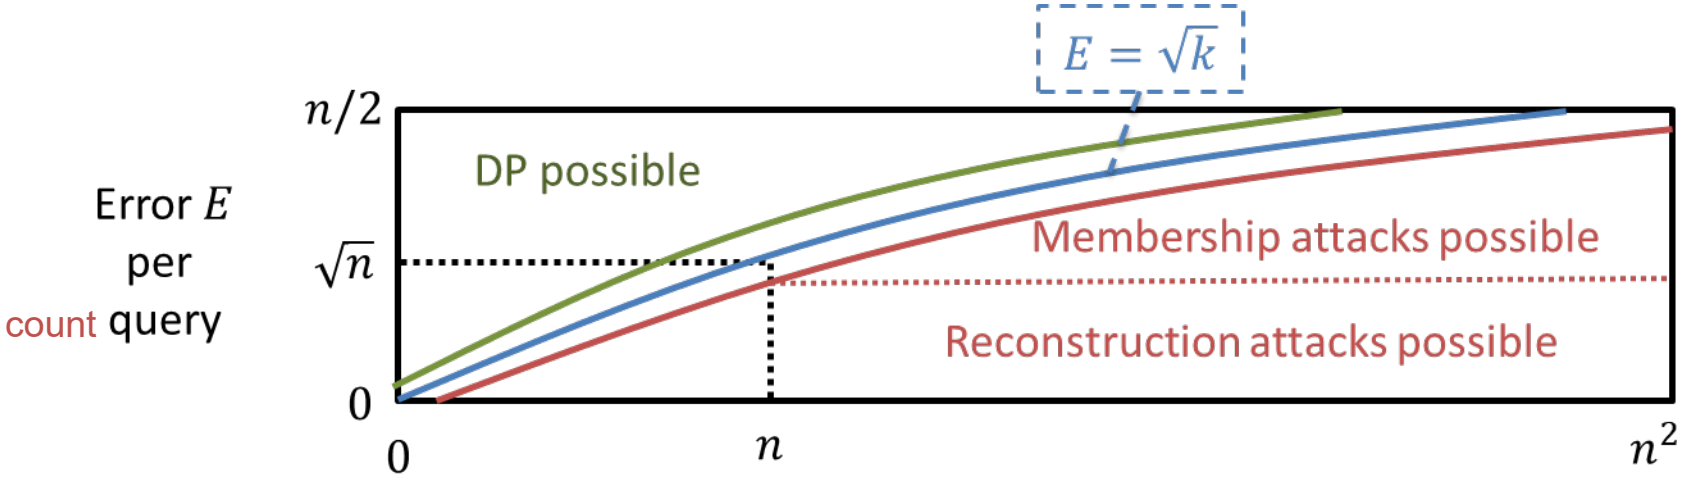

- Tradeoff is worse in local model.
- Allows for modular design of DP algorithms (w/post-processing).

Value of Rigorous Thinking in Privacy & Security

- Break cycle of attack-defense-attack-defense-…
- Separates goal from solutions.
  - Can evaluate privacy/security definition on its own.
  - Opens design space for solutions.
- Makes assumptions about adversary and implementation explicit, evaluable.
- Allows for study of tradeoffs (e.g., privacy vs. utility) and limits (impossibility, hardness).

Designing DP Algorithms

Every non-trivial DP mechanism must be randomized!

How to add noise and how much noise?

Scale to the query sensitivity (possibly after clipping)

- Means
- Variance
- Medians/Quantiles/Ranges
- Histograms
- Regression (OLS & Logistic)
- Synthetic Data Generation

- Empirical Risk Minimization
- Deep Learning/SGD
- Confidence intervals
- Causal inference (difference of means)
- Hypothesis tests

Core Components

A small number of primitives form the building blocks of some of the most complicated models, including:

- Clipping/Clamping
- Laplace and Gaussian Mechanisms
- Exponential Mechanism
- Randomized Response
- Composition
- Binning, One-hot encoding
- Histogram method

As well as some core recurring ideas:

- Post-processing
- Stable transformations
- Subsampling

DP Programming Paradigms & Challenges

- Measurements, Transformations, Combinators (PINQ, OpenDP)
- Multi-relational databases (Flex, PrivateSQL, GoogleDP)
- User interfaces (PSI/DP Wizard, ViP, etc.)
- General optimization/ERM/SGD (Opacus)
- Combining DP and other PETs (Private RCT)
- Formal verification of DP
- Finite-precision arithmetic
- Timing and other side-channel attacks

Experimental Investigation

Monte Carlo simulation methods are a valuable tool for investigating utility and other performance measures of algorithms. We have used this underlying template repeatedly:

1. Simulate data from distribution with known properties (or bootstrap from large dataset as if a population).
2. Release DP estimate and compare to true estimand.
3. Repeat 1 & 2 to integrate over simulation error and summarize.
4. Repeat 3 over free parameters of interest.

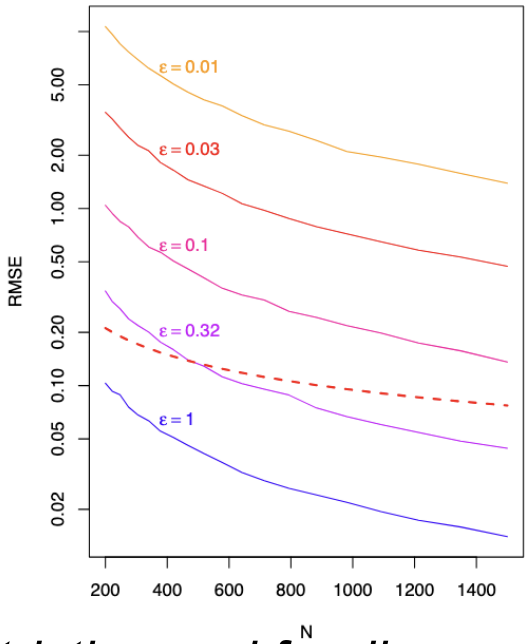

Real-world analyses of utility: impacts of DP on redistricting and funding allocation in the 2020 Census use case

Usable DP

- Explanation methods for conveying DP guarantees to data subjects
  - Text descriptions
  - Metaphors
  - Diagrams
  - Probability/risk-based explanations of epsilon
- Interactive interfaces for practitioners
  - PSI DP Wizard
  - ViP
  - Measure-Observe-Remeasure
  - …and more

Societal Perspectives

- Contextual Integrity
- Taxonomy of Privacy
- Surveillance and modulation
- Bridging technical-legal gaps
- Essentially contested concept
- Social construction of technology vs. politics of technology
- Handoff lens to identify value shifts

Deployments of DP

Census, Opportunity Atlas, Wikimedia, Apple, Google, Meta, Mozilla, …

Challenges and Open Problems:

- Getting both sufficient utility and satisfactory privacy.
- Managing privacy budget over many queries and analysts.
- Compatibility with stakeholder practices & expectations.
- Practical methods for generating synthetic data.
- Integration into existing workflows and pipelines.
- Enabling analysts to interpret noise, perform inference, measure uncertainty.
- Social, political, ethical, legal considerations
- Side channel attacks (e.g. randomness, timing).
- Vetted and general-purpose software tools.
- Combining DP with other PETs (MPC, TEE, FHE, ZKP) for end-to-end solutions.

To Pursue Further at Harvard

- Some final projects may lead to publishable papers.
- Attend Charles River Privacy Days
- Apply for an OpenDP internship and/or attend our regular DP seminar (write info@opendp.org)
- Explore annotated bibliography.
- Come discuss with us in office hours.
- Other Harvard faculty working on DP: Flavio Calmon (EE), Cynthia Dwork (CS), Gary King (Gov't), Xiao-Li Meng (Stats), Seth Neel (HBS)

To Pursue Further Elsewhere

- Apply for a job as a privacy engineer/data scientist/researcher.
  - Big & small tech companies
  - Privacy start-ups
  - Government agencies
  - Privacy non-profits and advocacy organizations
  - Industries grappling with data privacy (healthcare, finance, …)
  - Follow OpenDP slack #jobs channel
- Apply to graduate programs at places doing DP (we're happy to provide advice).