NIST SP800-226

Guidelines for Evaluating Differential Privacy Guarantees

Joseph P. Near
University of Vermont
https://www.uvm.edu/~jnear/

David Darais
Galois, Inc.
https://david.darais.com/

with Naomi Lefkovitz & Gary Howarth
NIST

Outline

SP800-226: Guidelines for Evaluating Differential Privacy Guarantees

- Planned for public comment later this month
- Looking for your feedback!

This talk:

1. Intro to differential privacy
2. Goals of SP800-226
3. Examples of differential privacy hazards

The Differential Privacy Guarantee

What is Privacy?

Data Privacy:

An analysis is *privacy preserving* if:

- It reveals useful information about the population (**utility**)
- It does *not* reveal new information about individuals (**privacy**)

The Differential Privacy Guarantee

Differential privacy:

Analysis outcome is equally likely, **whether or not I contribute my data**

Implication #1: privacy harm following analysis ***would have happened anyway***

Implication #2: “off-grid cabin world” ≈ “real world”

I live in a cabin off-grid

≈

I live in the real world

Differential Privacy: a Scale for Privacy

Superpower #1:

Differential privacy ***measures privacy***

Proposed data release or analysis

Upper bound on potential privacy harm

Differential Privacy: Compositional

Superpower #2:

Differential privacy is ***compositional***

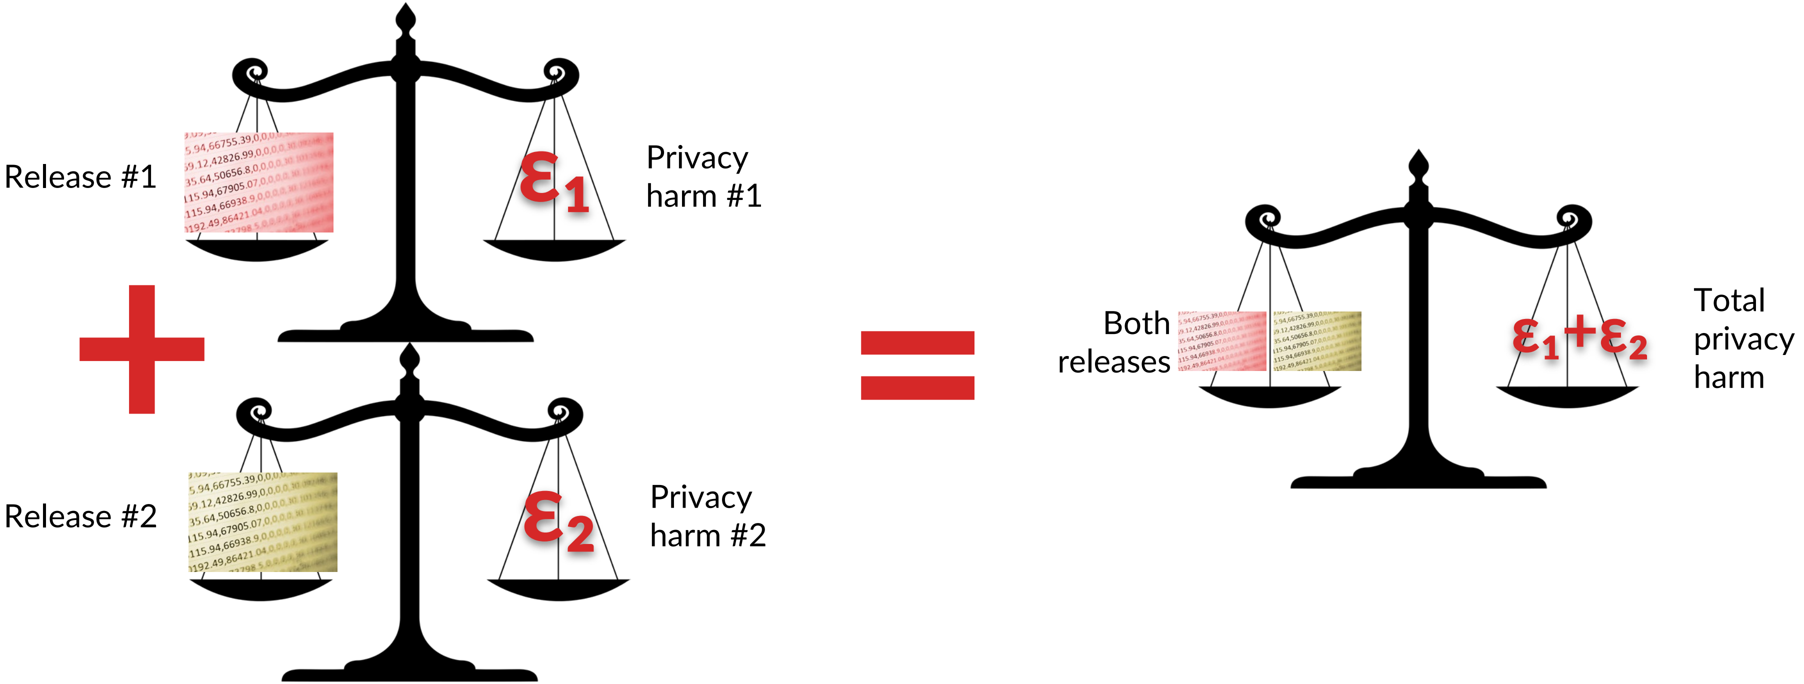

Achieving Differential Privacy

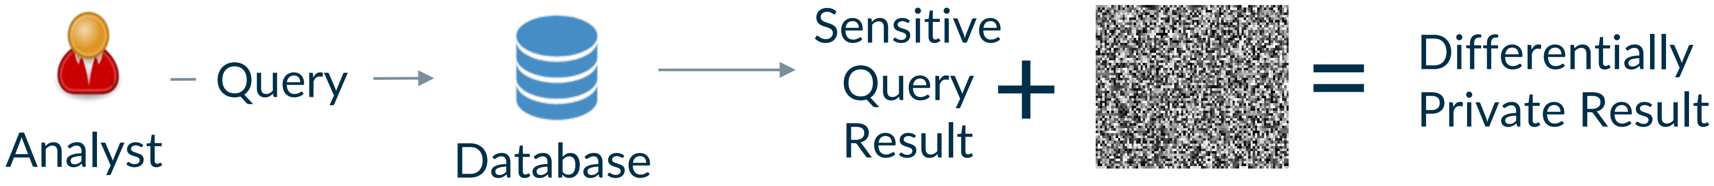

Prototypical solution: **add noise** to results

More noise = **more privacy**

Privacy tuned by ***privacy parameter*** ε

Impact of the Privacy Parameter

Smaller ε
More noise
More privacy
Less accuracy

Larger ε
Less noise
Less privacy
More accuracy

Goals of SP800-226

Goals of SP800-226

- Introduce differential privacy
- Summarize the aspects of a differential privacy guarantee
- Describe how to evaluate and compare guarantees
- Highlight important privacy hazards

Out of Scope

- Describe the math of differential privacy
- Teach how to implement differential privacy
- Compare differential privacy to other techniques

Target audience: practitioners

- Managers
- Software engineers
- Policymakers
- Data scientists

Structure of the Document

- Each section describes one aspect of the guarantee
- Evaluating a guarantee requires considering all aspects

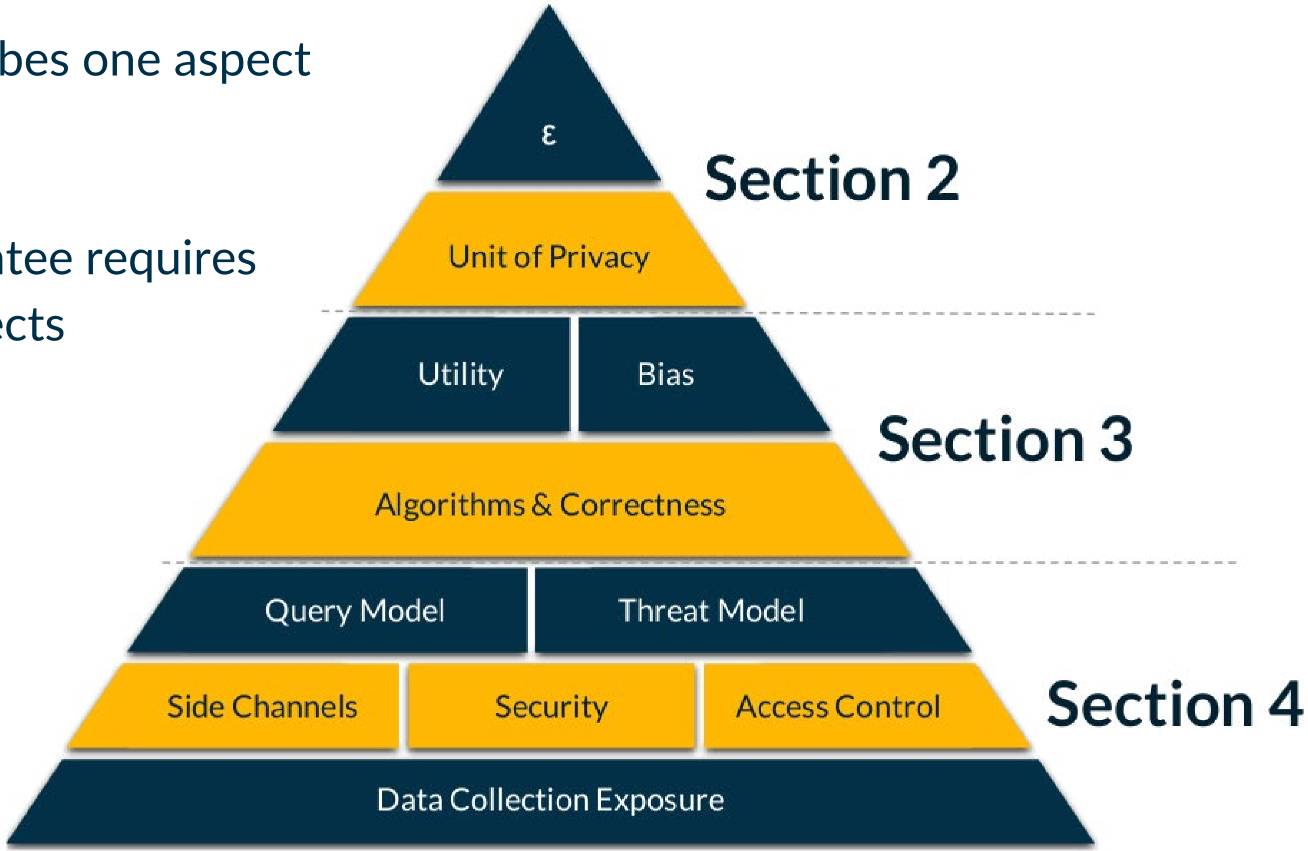

Differential Privacy Hazards

Hazard #1: Setting and Interpreting ε

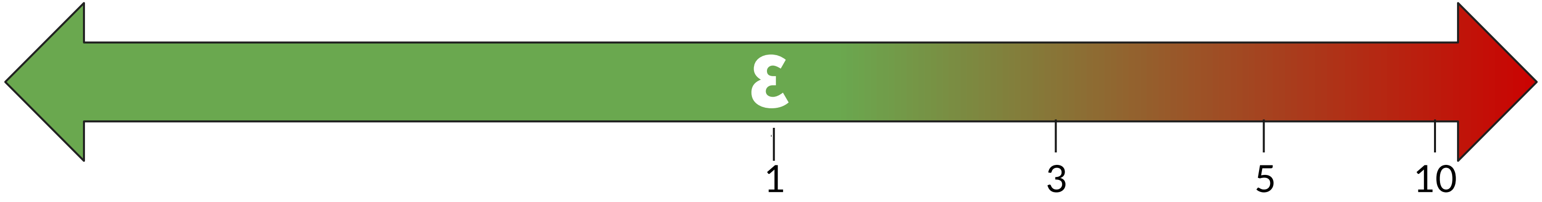

Traditional "rule of thumb": **ε ≤ 1 is best**

Hazard #2: Buggy Algorithms

Implementation bugs: easy to introduce, hard to detect!

Use well-tested libraries whenever possible

On Significance of the Least Significant Bits For Differential Privacy

Ilya Mironov

Abstract

We describe a new type of vulnerability present in many implementations of differentially private mechanisms. In particular, all four publicly available general purpose systems for differentially private computations are susceptible to our attack.

Mironov, I., 2012, October. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security* (pp. 650-661).

Hazard #3: Systemic Bias

Noise has bigger impact on small groups

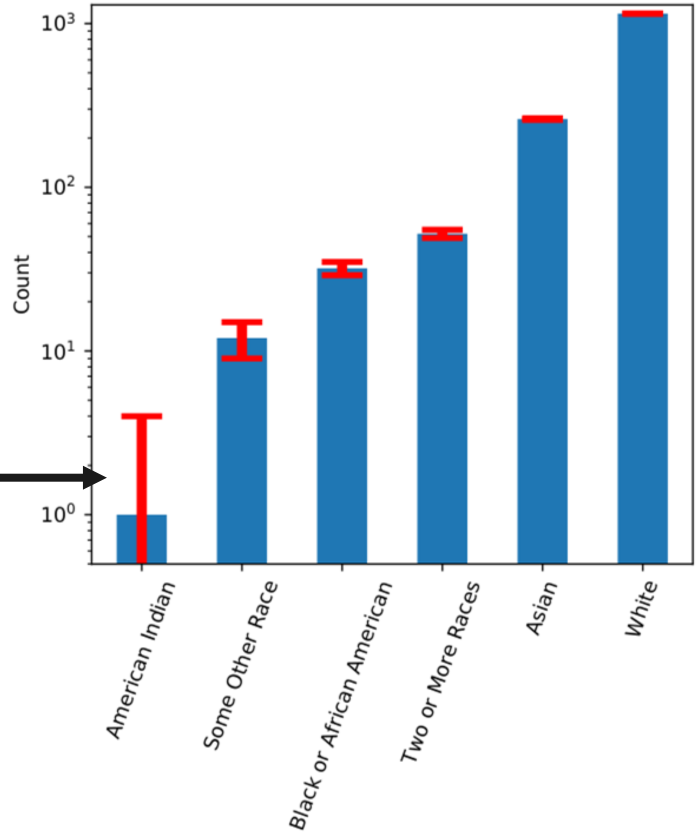

Differential privacy can create or amplify systemic bias

Hazard #4: Security

What if the server gets hacked?

Differential privacy doesn't necessarily protect data at rest

Thank you!

We need your help to improve the publication!

Seeking feedback on:

- Clarity & understandability
- Correctness & accuracy
- Missing info