Privacy Pinch Points for Applied ML

Bailey Kacsmar

PUPS Research Lab
UNIVERSITY OF ALBERTA
amii

This session:

Privacy Pinch Points for Applied ML

Not this session:

Privacy Pinch Points for **Applied ML**

Pinch Points?

Image source: https://www.constructionsafety.co.za/ems/pinch-points/

Pinch Points?

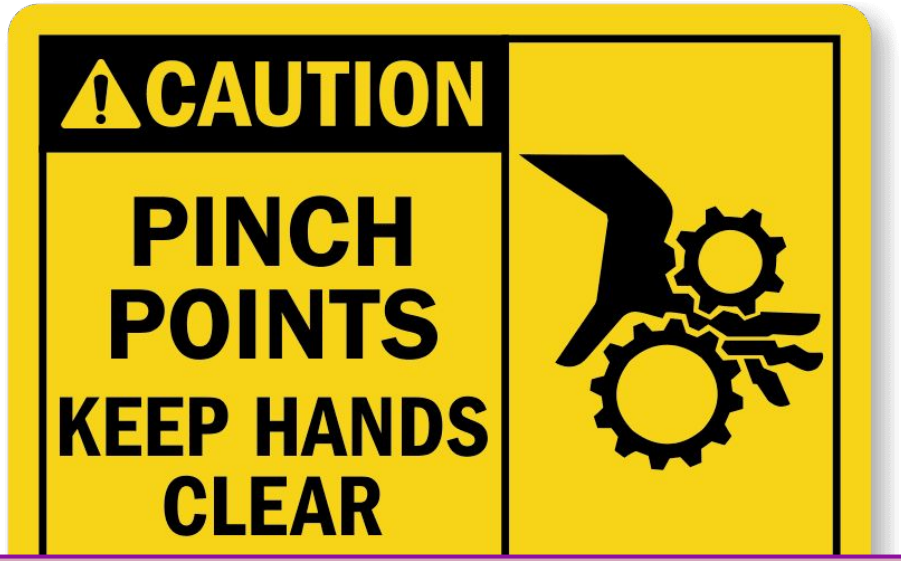

Def: When objects come together and there is a possibility that a person could be caught or injured

Image source: https://www.constructionsafety.co.za/ems/pinch-points/

Common Causes of Pinch Points?

- Lack of attention…
- Mobility (of equipment)
- Poor maintenance
- Lack of proper safe work procedures
- Reaching into moving points…

Common Causes of Pinch Points?

- Lack of attention…
- Mobility (of equipment)
- Poor maintenance
- Lack of proper safe work procedures
- Reaching into moving points…

Privacy? Applied ML?

Privacy Pinch Points

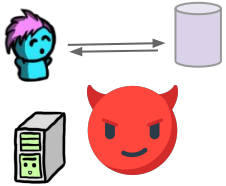

Technical Privacy

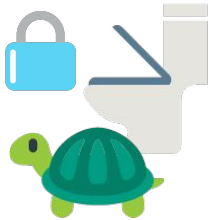

Conceptual Privacy

Legal Privacy

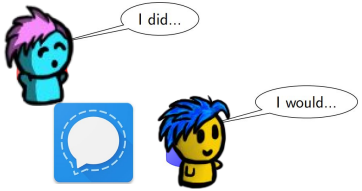

Usable Privacy

Challenge:

Why Privacy and ML?

What makes this hard? What's the risk?

Privacy and Data

Google and Mastercard Cut a Secret Ad Deal to Track Retail Sales

Google found the perfect way to link online ads to store purchases: credit card data

By Mark Bergen and Jennifer Surane
August 30, 2018, 3:43 PM EDT *Updated on August 31, 2018, 12:40 PM EDT*

washingtonpost.com

Now for sale: Data on your mental health

Drew Harwell

These retailers share customer data with Facebook's owner. Customers may not have been told | CBC News

Thomas Daigle · CBC News · Posted: Feb 07, 2023 4:00 AM EST | Last

Home Depot didn't get customer consent before sharing data with Facebook's owner, privacy watchdog finds | CBC News

Catharine Tunney · CBC News · Posted: Jan 26, 2023 9:53 AM
Updated: January 27

Double-double tracking: How Tim Hortons knows where you sleep, work and vacation

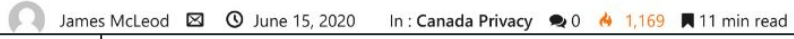

Why Privacy and ML?

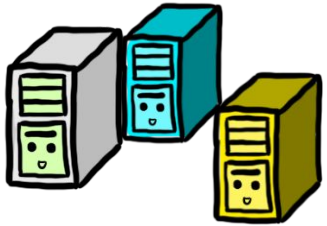

A company wants to analyze data

But the data has privacy implications for the data subjects

Researchers develop technical solutions

Privacy versus Security

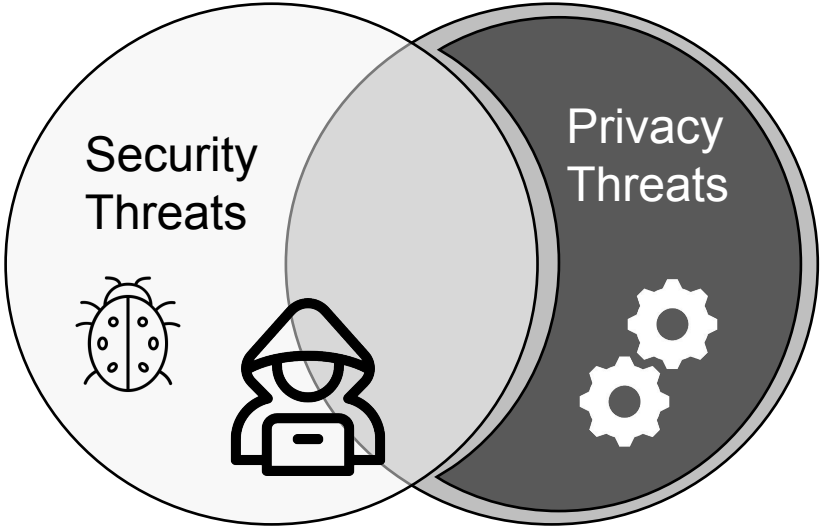

Westin's (1967)

An entity's **ability to control** how, when, and to what extent personal information about it is communicated to others

For privacy, focus on the **harms** (consequences) caused by privacy violations.

Privacy Pinch Points = Risk of Harms

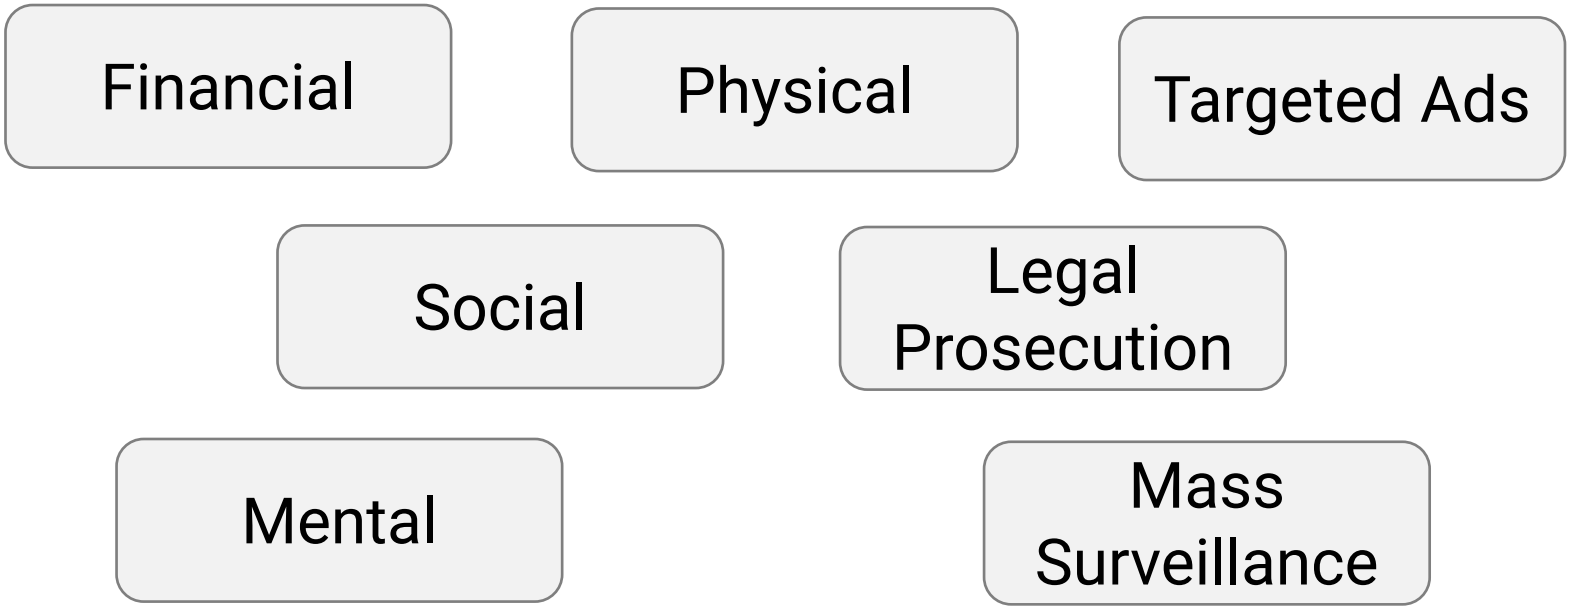

Privacy Violation Life-Cycle

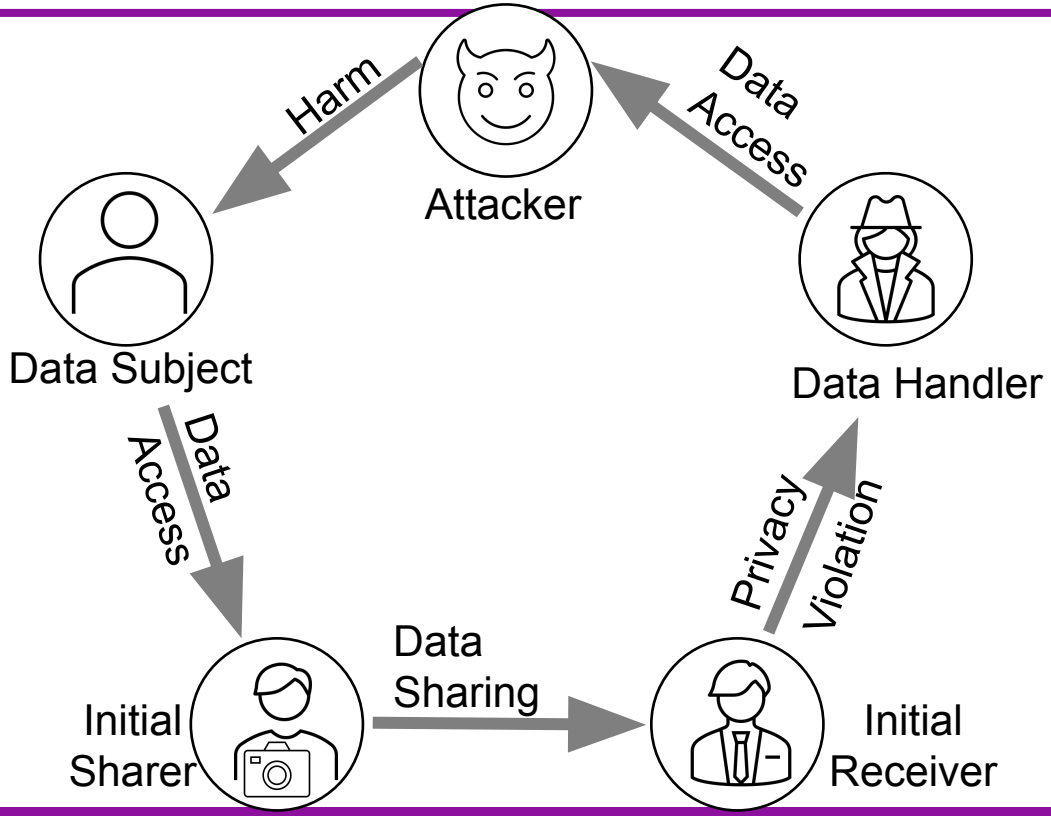

Adversarial Thinking

- Think like an adversary to understand the ***vulnerabilities*** of a system and develop ***protection techniques***.
- When designing inference attacks, we also apply **Kerckhoff's principle** (or Shannon's maxim), adapted to privacy

Adversarial Thinking

- Think like an adversary to understand the ***vulnerabilities*** of a system and develop ***protection techniques***.
- When designing inference attacks, we also apply **Kerckhoff's principle** (or Shannon's maxim), adapted to privacy

Assume the adversary knows how the system works

- There are **no hidden parameters** other than the users' data
- The adversary can **even know some rough distribution**

What are inference attacks?

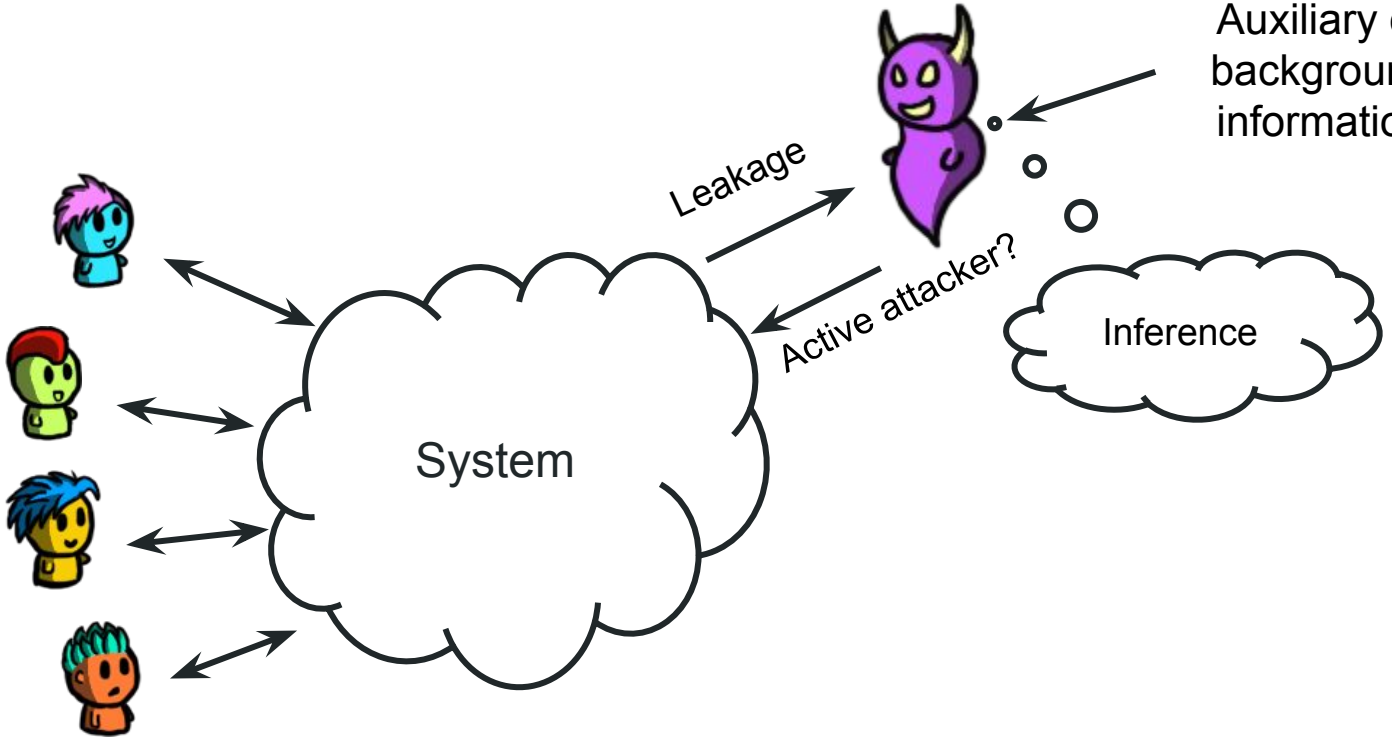

What are inference attacks?

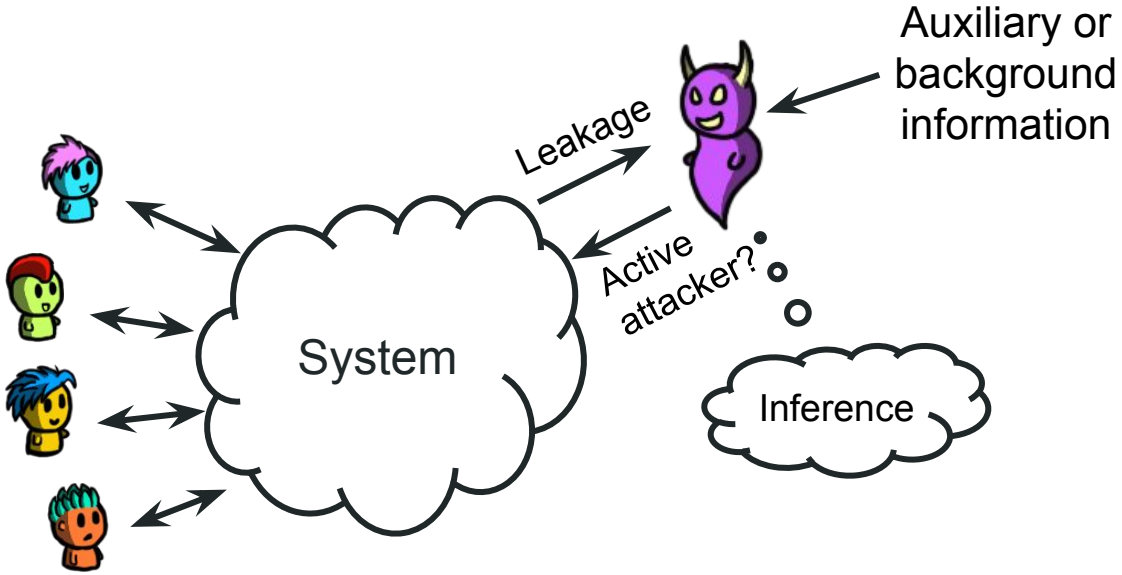

Goal: Learn something (non-trivial) and privacy sensitive from the system

What are inference attacks?

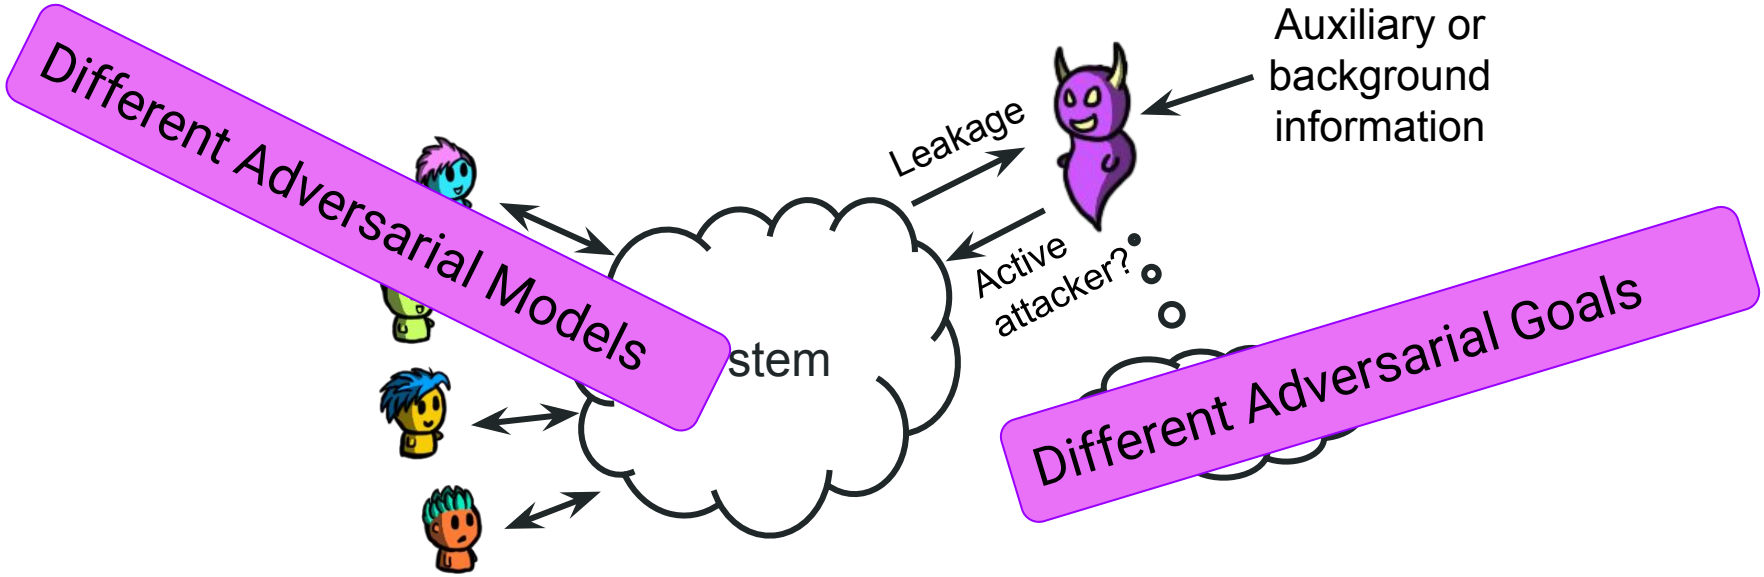

Goal: Learn something (non-trivial) and privacy sensitive from the system

Inference Attacks: Goals and Abilities

- **Goals**:
  - Infer data
  - Infer a property of the data
  - Infer the presence (membership) of some data
  - Infer the behavior of a user
  - Infer some attributes of a data sample
  - Infer dependencies among the data
  - …

Inference Attacks: Goals and Abilities

- **Goals**:
  - Infer data
  - Infer a property of the data
  - Infer the presence (membership) of some data
  - Infer the behavior of a user
  - Infer some attributes of a data sample
  - Infer dependencies among the data
  - …
- **Abilities**:
  - Statistical tools (estimation theory, detection theory, maximum likelihood, Bayesian inference…)
  - Combinatorics
  - Heuristics
  - Machine learning
  - …

Designing a System Aware of Inference Attacks

For any system that relies on users' data, there are two goals:

- **Utility:** Design a system that provides benefits to its users and the service provider
- **Privacy:** Design a system that provides protection against inference attacks

Q: What are "utility" and "privacy"? How do we "measure" them?

Designing a System Aware of Inference Attacks

For any system that relies on users' data, there are two goals:

- **Utility:** Design a system that provides benefits to its users and the service provider
- **Privacy:** Design a system that provides protection against inference attacks

Q: What are "utility" and "privacy"? How do we "meas

It's complicated...

Privacy Mitigations? Private Computation?

Defenses!!

Private Computation

Private Computations

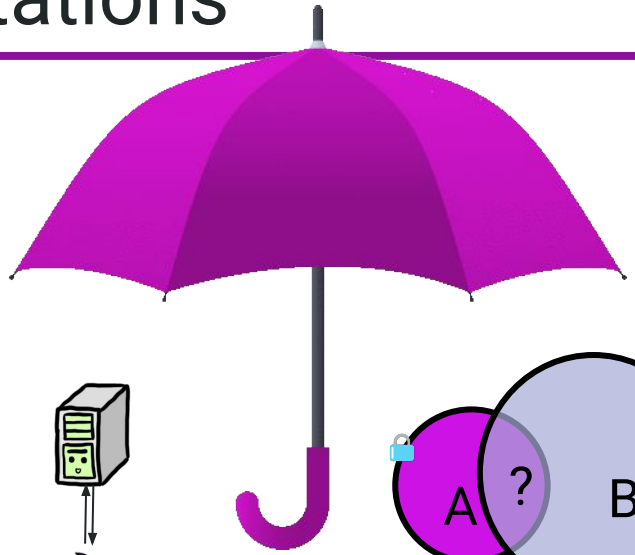

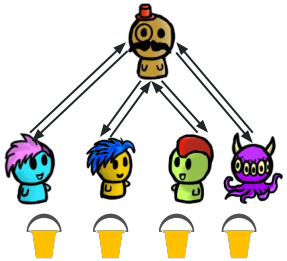

Private Machine Learning

Private Query Processing

Private Set Intersection

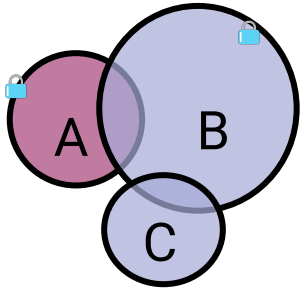

Multiparty Computations

Private Computations Class

Define, **what** is being protected, from **whom**,
and under what **conditions** this protection will hold.

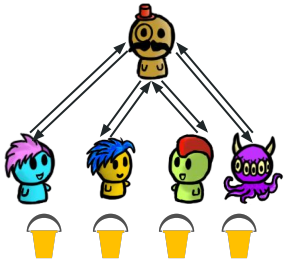

Private Machine Learning

Private Query Processing

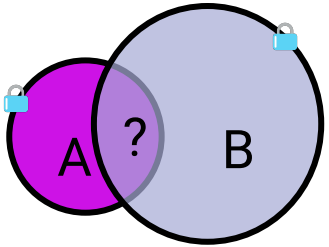

Private Set Intersection

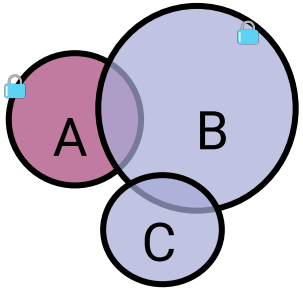

Multiparty Computations

Technical Guarantees Types

- Statistical
- Computational
- Information Theoretical

Quantifying Privacy: Theoretical Notions

- **Syntactic** notions of privacy: these are computed on the leaked or released data. They are data dependent
  - K-anonymity, l-diversity, t-closeness, etc

Quantifying Privacy: Theoretical Notions

- **Syntactic** notions of privacy: these are computed on the leaked or released data. They are data dependent
  - K-anonymity, l-diversity, t-closeness, etc
- **Semantic** notions of privacy: these are computed on the data release mechanism itself, and they hold regardless of the data (data independent)
  - Mostly Differential Privacy

Quantifying Privacy: Empirical Notions

- The performance of an **inference attack** e.g., the attacker error, accuracy, true positive rate, false positive rate, etc
- Can provide an **upper bound** on privacy

The Privacy-Utility trade-off

- Given any metric for privacy and for utility, they are usually at odds:

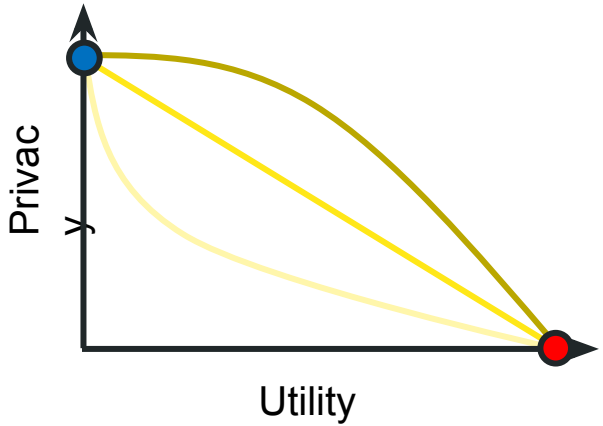

- **Q:** How do you design a system that provides maximum utility?
- **Q:** How do you design a system that provides maximum privacy?
- Designing a system that provides a good privacy-utility trade-off is hard!

The Privacy-Utility trade-off

- Given any metric for privacy and for utility, they are usually at odds:

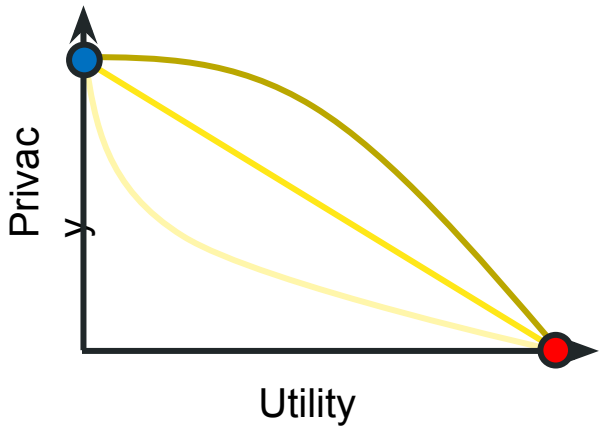

- How do you design a system that provides maximum utility?
  - You design it without privacy in mind
- How do you design a system that provides maximum privacy?
  - ..?
- Designing a system that provides a good privacy-utility trade-off is hard!

The Privacy-Utility trade-off

- Given any metric for privacy and for utility, they are usually at odds:

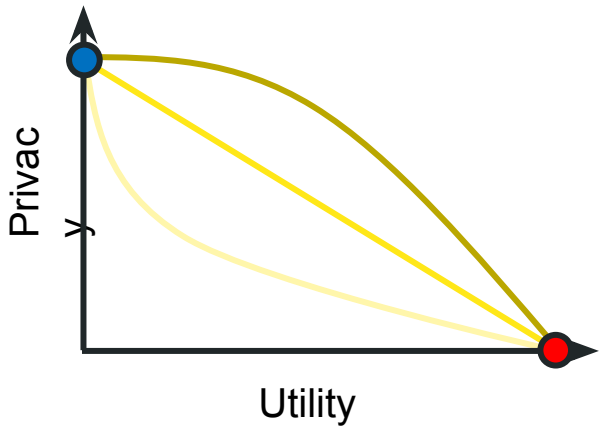

- How do you design a system that provides maximum utility?
  - You design it without privacy in mind
- How do you design a system that provides maximum privacy?
  - You don't design it
- Designing a system that provides a good privacy-utility trade-off is hard!

Setting:

What are we protecting and how?

A Private Computation? Cryptography!

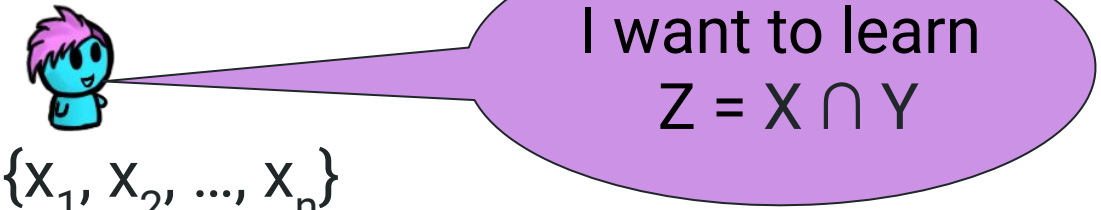

$X = \{x_1, x_2, \ldots, x_n\}$

$Y = \{y_1, y_2, \ldots, y_m\}$

Private Set Intersection

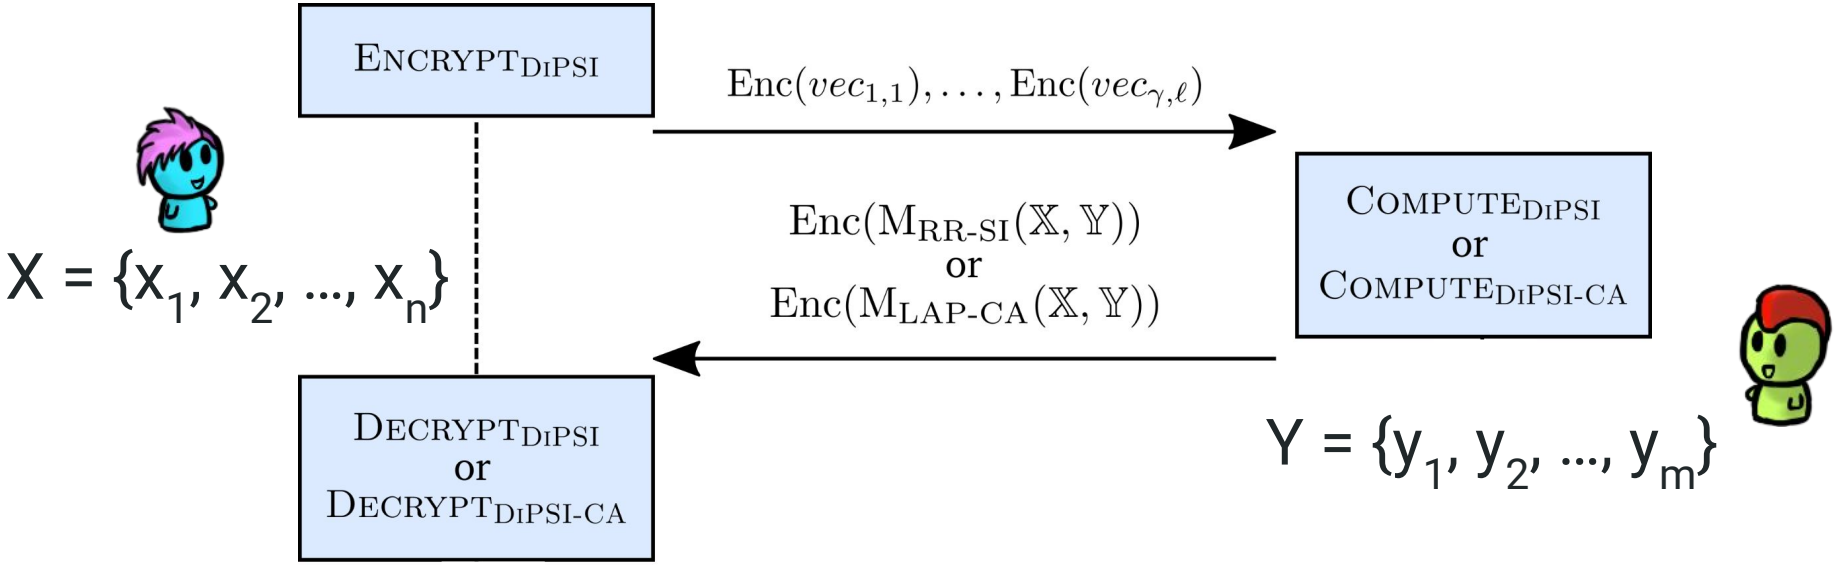

B. Kacsmar, B. Khurram, N. Lukas, A. Norton, et al. "Differentially private two-party set operations." In 2020 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 390-404. IEEE, 2020.

Private Computation and Machine Learning?

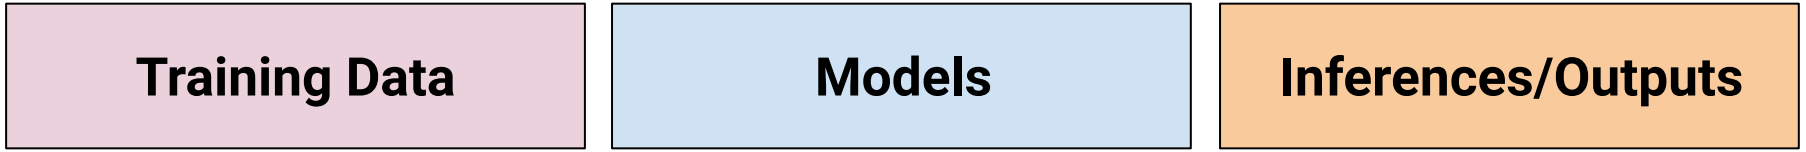

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Private Computation and Machine Learning?

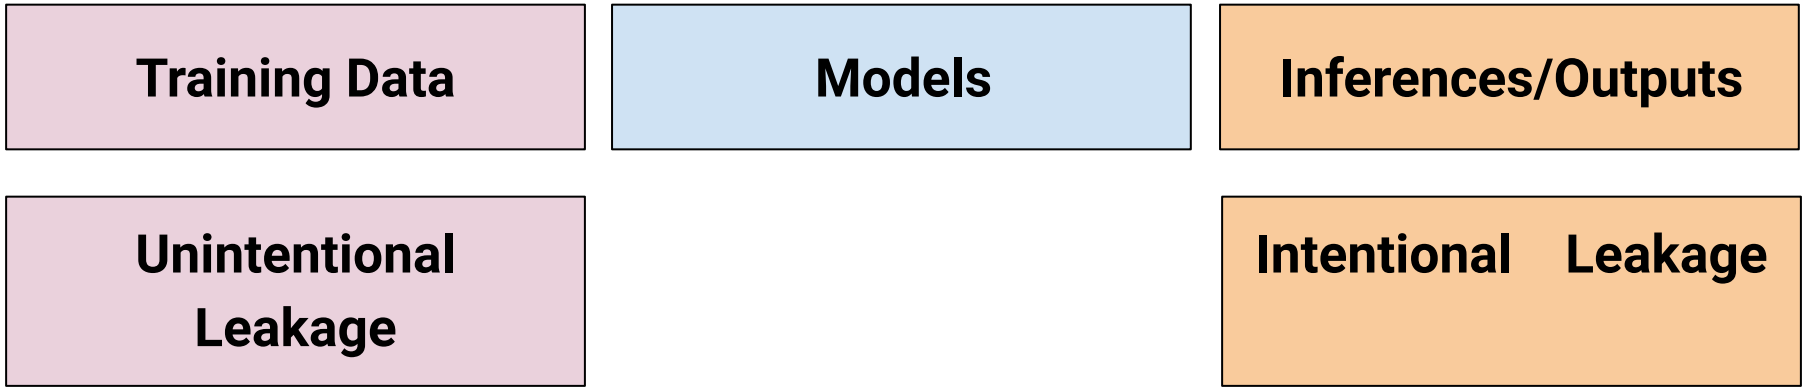

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Private Computation and Machine Learning?

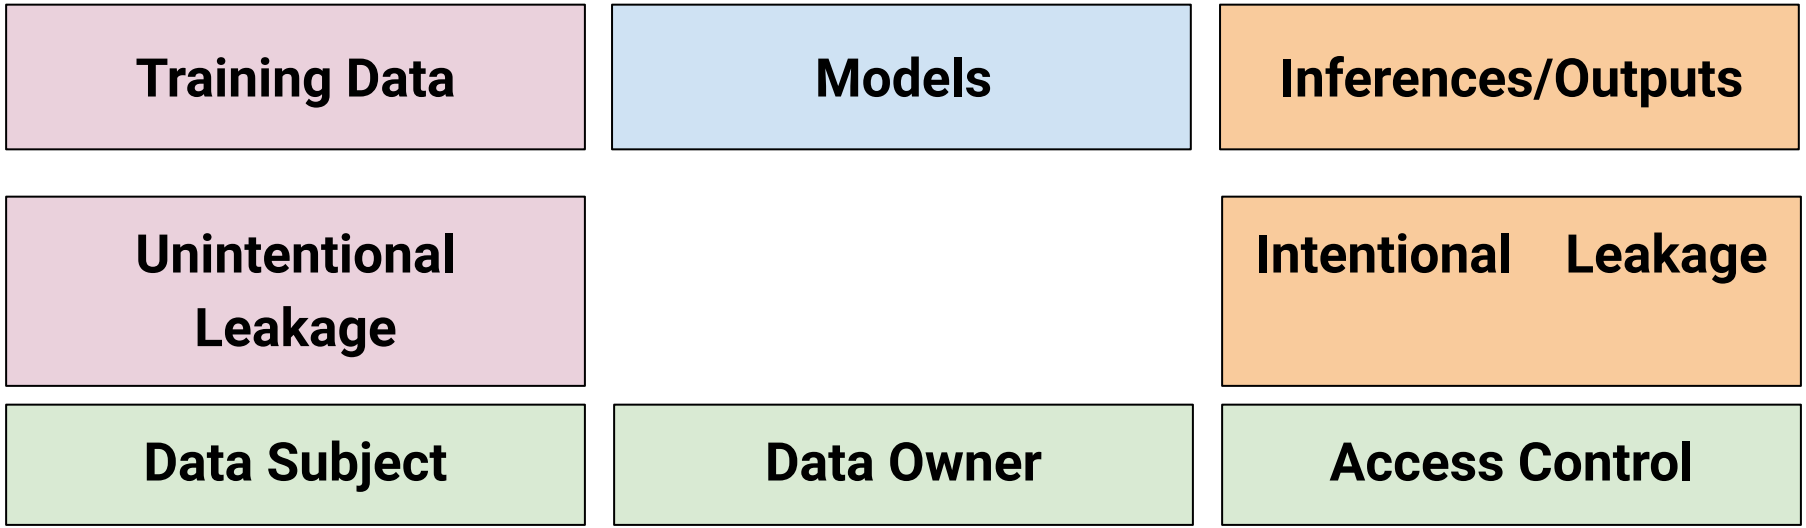

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Design:

A Bit on Privacy Mitigation Techniques

Towards Privacy by Design, Core Tenets

- User centric
- Embedding privacy into the design
- Having privacy as the default configuration
- Ensuring privacy across the whole software life-cycle

Technical Privacy: Differential Privacy Intuition

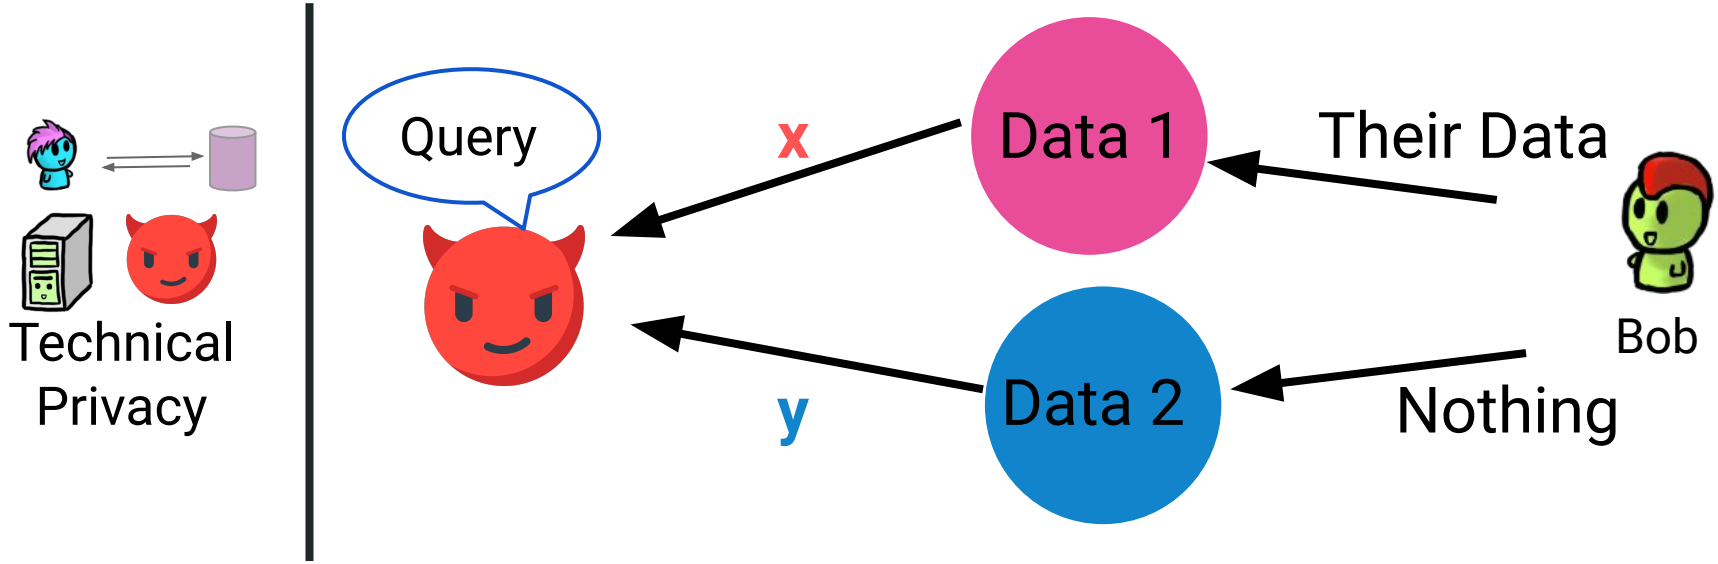

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Differential Privacy and Machine Learning

- DP-SGD
- Individualized Differential Privacy (PATE)
- More…

However, still **require expertise** for deployment

Distribution of Trust

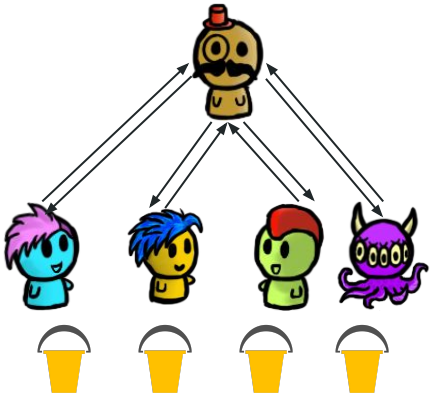

Federated Learning
<u>PLUS</u> something

- Distribution alone is not private
- SMPC is…expensive
- But…

Not putting all the eggs in one basket, will always have appeal.

Always: Data Minimization

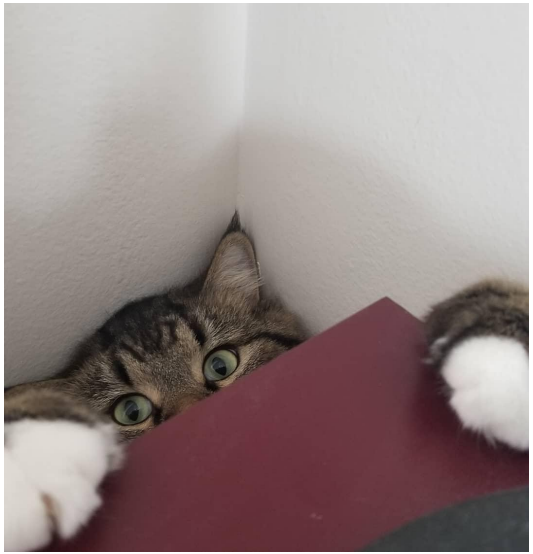

Challenge:

Is it enough? What about the other vectors…

Challenge:

Is it enough? What about the other vectors…

Consent and Communication

A Wider View of Technical Privacy

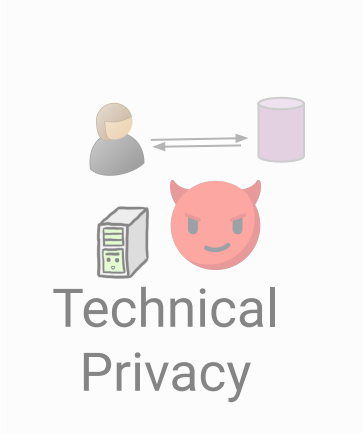

Technical Privacy

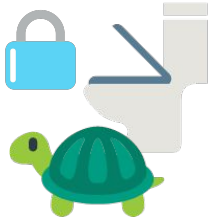

Conceptual Privacy

Legal Privacy

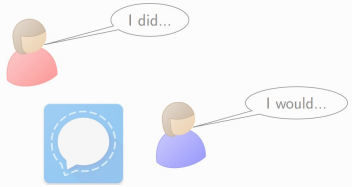

Usable Privacy

Understanding privacy notions and behaviours, **right to privacy**, and privacy expectations

M. Oates, et al. Turtles, locks, and bathrooms: Understanding mental models of privacy through illustration." Proceedings on Privacy Enhancing Technologies 2018.

Why Private Computation?

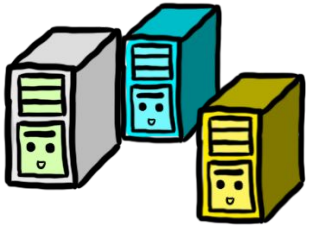

A company wants to analyze data

But the data has privacy implications for the data subjects

Researchers develop technical solutions

In what ways does private computation matter to people?

Perceptions and Expectations

- What do data subjects understand?
- How is a data subject's willingness to share impacted?
- How do data subjects perceive the risks?

What they "want"

What they "need"

Build towards those attributes

The Scenarios

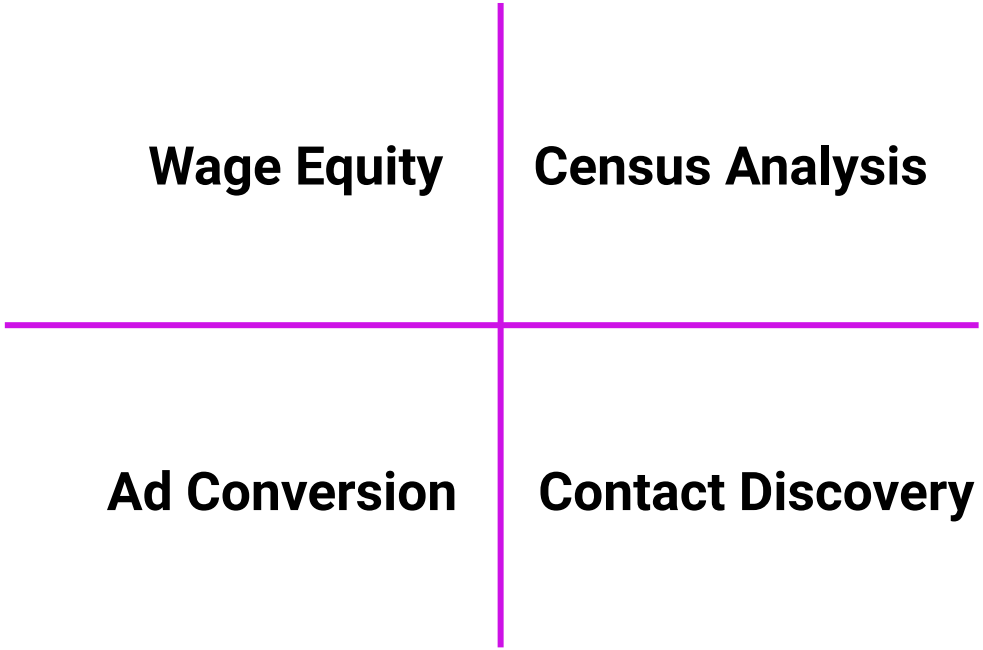

Contact Discovery Conceptual Example

The app wants to **determine the common contacts** between the new user and the existing users via…

1. …the new user shares all their contact information with the social media app.
2. … the new user shares **a modified version** of their contact information…**such that** the social media app does not learn non-users…thus, **this means**…

The Interview

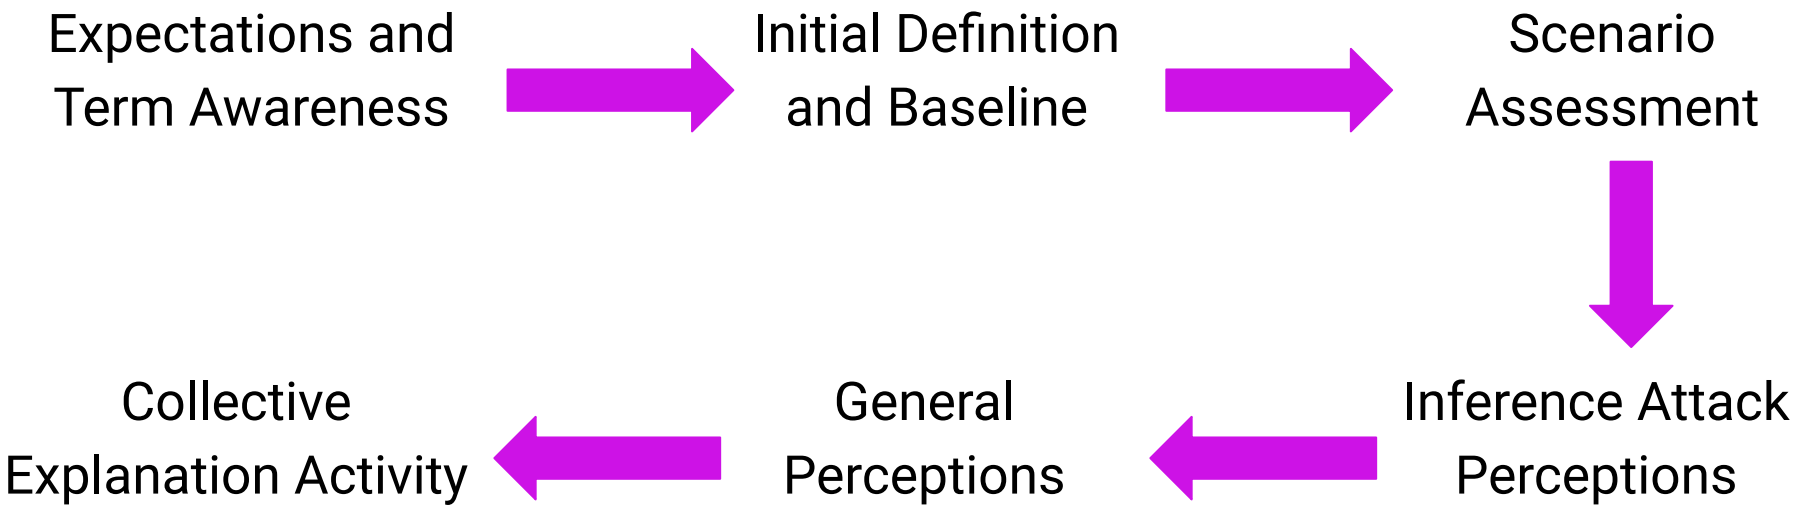

Participant Comprehension and Expectations

First Attempt

Second Attempt

Secure computation is a way that a company analyzes your data. The final analysis will be made public [at access location]. However, your specific data is protected and cannot be traced back to you nor can your specific data points be traced back to you. The analysis will be specifically [example], and this is being done because [purpose].

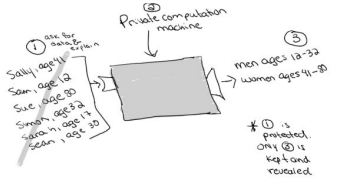

This is the information we're getting from you, but, rest assured, only Part Three will be shown. You can trust us to keep your information private. <If true>This information will only be used for this project and nothing else in the future.

Final Consensus

Participant Comprehension and Expectations

First Attempt

Second Attempt

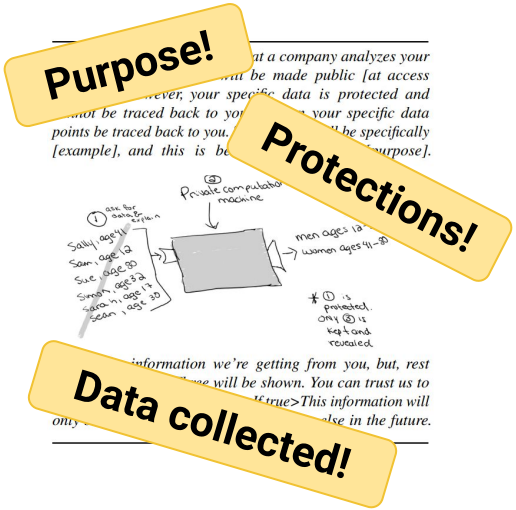

Final Explanation

Participant Comprehension and Expectations

First Attempt

Second Attempt

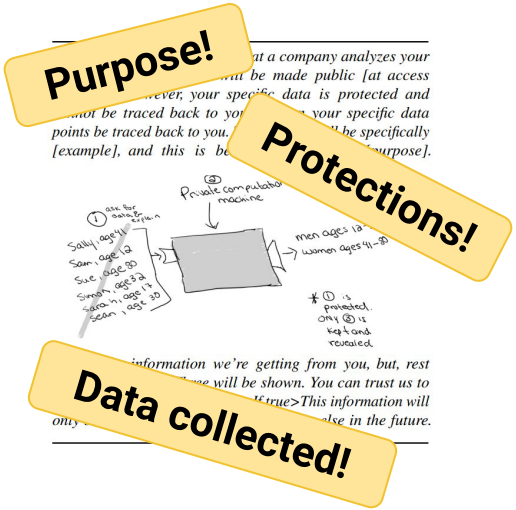

Final Explanation

Unconcerned with details of the mechanism, **impact** matters

Impact of Private Computation

"...they're trying to make it sound a little bit better" (P19).

"...it feels a little bit more protected that way" (P12)

Bounded Impact of Private Computation

Intentions Matter

Divulge the Details

Regulate the Restrictions

Consent Above All

“At the end of the day,
they’re still like learning specific things about me” (P7)

Awareness of Unique Threat Models

Alice

Joins Social App

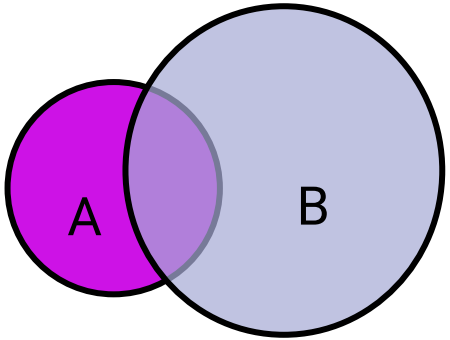

Contact Discovery

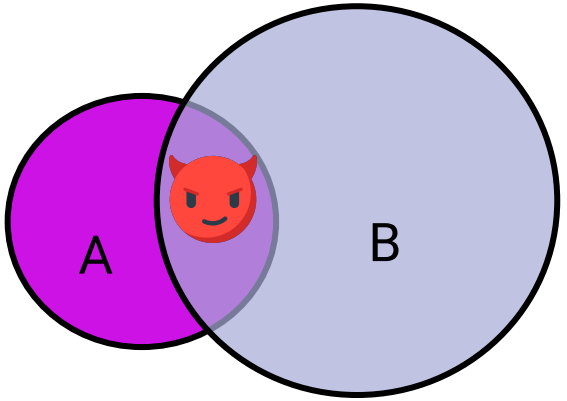

Real Identity Connected

Awareness of Unique Threat Models

Alice

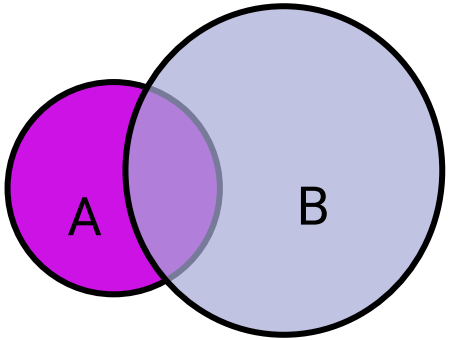

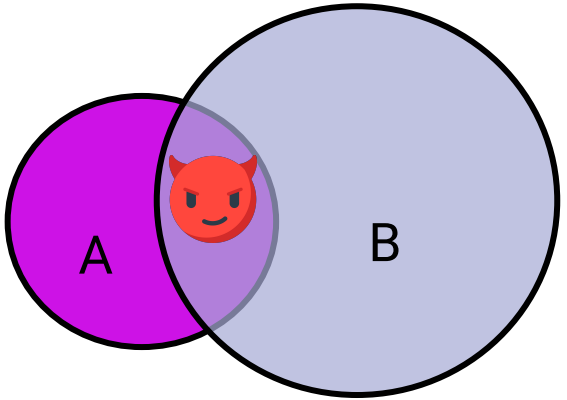

Joins Social App

Contact Discovery

Real Identity Connected

There exist, and will continue to exist risks that cannot be regulated by technology

Pinch Points. Many. But getting better…

Image source: https://www.constructionsafety.co.za/ems/pinch-points/

Takeaways

- **Protections provided by protocols** and constructions **do not encompass the full range of risks** experienced **by individuals** in society
- Privacy mitigation techniques are a treatment and not a cure for data privacy concerns
- People find private computation plausible, but they **care about the context, not the math**

Takeaways

- **Protections provided by protocols** and constructions **do not encompass the full range of risks** experienced **by individuals** in society
- Privacy mitigation techniques are a treatment and not a cure for data privacy concerns
- People find private computation plausible, but they **care about the context, not the math**

People can reason about private computation; let them.

Thanks!