Bridging the Gap between Computer Science and Legal Approaches to Privacy

Alexandra Wood

Berkman Klein Center for Internet & Society at Harvard University

Computer Science 208: Applied Privacy for Data Science

April 22, 2019

An Interdisciplinary Collaboration

This work is the product of an *interdisciplinary working group* bringing together computer scientists, information scientists, and legal scholars

Computer Science

Kobbi Nissim, Aaron Bembenek, Mark Bun, Marco Gaboardi, Thomas Steinke, Salil Vadhan

CRCS Center for Research on Computation and Society
at Harvard John A. Paulson School of Engineering and Applied Sciences

Georgetown University

Law & Policy

Urs Gasser, David O'Brien, Alexandra Wood

BERKMAN KLEIN CENTER
FOR INTERNET & SOCIETY
AT HARVARD UNIVERSITY

Social Science & Information Science

Micah Altman

MIT Libraries

The Data Privacy Problem

Data Privacy: The Problem

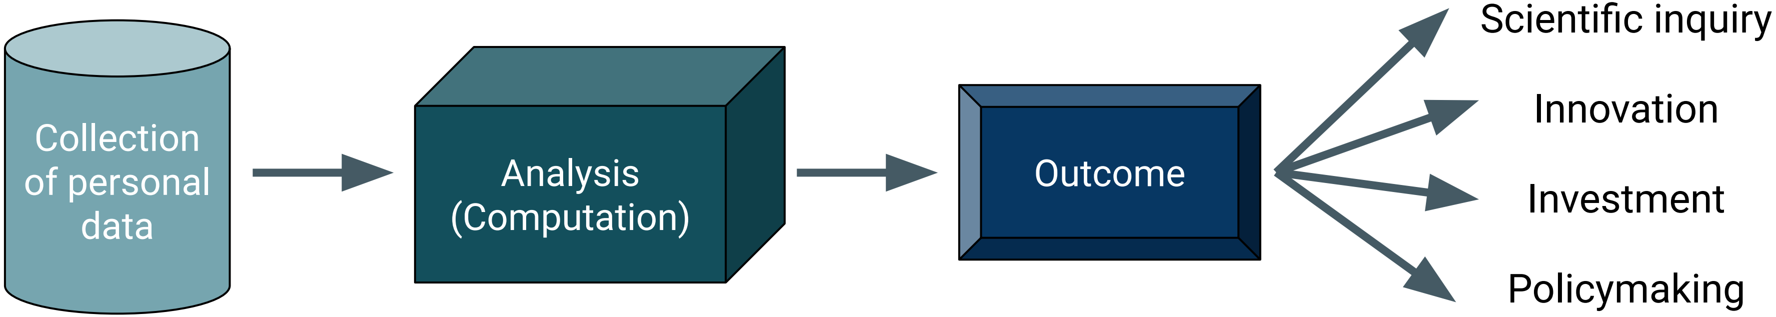

How can personal data can be analyzed and shared, while ensuring the privacy of the individuals in the data will be protected?

Real-World Example of a Privacy Attack

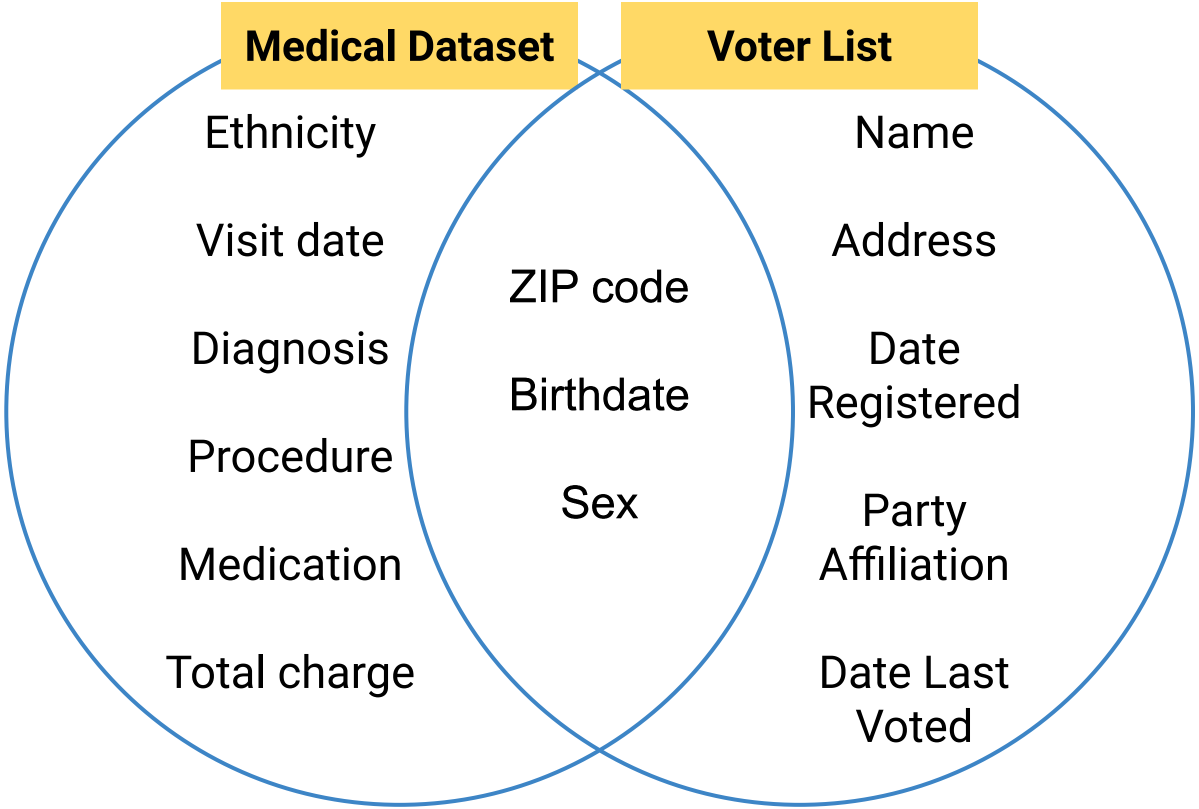

Medical Dataset

Ethnicity

Visit date

Diagnosis

Procedure

Medication

Total charge

ZIP code

Birthdate

Sex

Voter List

Name

Address

Date Registered

Party Affiliation

Date Last Voted

Source: Latanya Sweeney (1997)

HIPAA Privacy Rule (2000)

Safe Harbor Method for de-identifying protected health records:

> (i) Categories of information from a list of 18 identifiers (e.g., names, geographic units containing 20,000 or fewer people, dates (except year), telephone numbers, Social Security numbers, etc.) are removed, and
>
> (ii) The covered entity does not have actual knowledge that the information could be used alone or in combination with other information to identify an individual who is a subject of the information.

45 C.F.R. § 164.514

Database Reconstruction Example

- **US Census Bureau's dual goals**
  - Collecting and publishing data necessary for democracy
  - Protecting the privacy of individuals to ensure trust and prevent harm
    - 13 U.S.C. § 9 prohibits Census Bureau employees from

      > "mak[ing] any publication whereby the data furnished by any particular establishment or individual under this title can be identified" (13 U.S.C. § 9(a)(2))

- **These goals are in conflict**
  - Fundamental law of information recovery: "overly accurate" estimates of "too many" statistics completely destroys privacy (Cynthia Dwork)
  - Approximately 7.8 billion statistics released from 2010 Census

Database Reconstruction Example

- The Census Bureau’s confidentiality mechanisms have evolved over time.
    - E.g., suppression, data swapping, and synthetic data generation
- In 2019, Census Bureau researchers published results from a simulated data reconstruction attack on the 2010 Census data publications.
    - Performed a database reconstruction for all 308,745,538 people in 2010 Census
    - Linked reconstructed records to commercial databases from 2010
    - Block, sex, age, race, ethnicity reconstructed exactly for 46% of population, and 71% of population allowing age +/- one year
    - Linked to commercial PII and confirmed exact re-identifications for 17% of population

Attacks on Privacy: Key Takeaways

- Lack of rigor leads to unanticipated privacy failures.
    - New attack modes emerge as research progresses.
    - Redaction of identifiers, release of aggregates, etc. is insufficient.
    - Auxiliary information must be taken into consideration.
- Any useful analysis of personal data must leak some information about individuals.
- Information leakages accumulate with multiple analyses/releases.

Mathematical facts, not matters of policy

When traditional techniques for protecting privacy can fail spectacularly, how do you proceed

… as a data publisher like the Census Bureau?
… as a regulator?

Emergence of a New Privacy-Enhancing Technology

Emergence of Differential Privacy

- A new line of privacy work in theoretical computer science (beginning ~2003)
- Yields a new concept: **Differential privacy** (2006)
  - Supported by rich theory
  - In its first stages of implementation and real-world use
    - US Census, Google, Apple, Uber, …
    - Disclosure avoidance mechanism for the 2020 Census

Relevance to Data Analysis and Sharing

- Various legal provisions restrict disclosures of identifiable or sensitive information about individuals.
  - e.g., FERPA generally prohibits the disclosure of personally identifiable information from education records, except with consent or pursuant to one of several narrow exceptions to the consent requirement. Notably, FERPA permits the disclosure of de-identified information.
- However, there is a lack of certainty around the use of terms like *personally identifiable information* and *de-identified information*, especially as the understanding of privacy risks continues to evolve over time.

A Practical Challenge

Formal privacy models like differential privacy offer a solution for providing wide access to statistical information with guarantees that individual-level information will not be leaked inadvertently or due to an attack.

- ➢ Formal mathematical privacy concept that addresses weaknesses of traditional schemes (and more).
- ➢ Supported by a rich theoretical literature and now in initial stages of implementation and testing by industry and statistical agencies.

🚫 However, these tools cannot be used to share sensitive data with the general public unless they satisfy legal standards with some certainty.

What is the relationship between differential privacy and existing regulatory standards for privacy protection?

Gaps between Differential Privacy and Legal Standards for Privacy

Challenges for Formal Privacy Models

Demonstrating that formal privacy models satisfy applicable legal requirements is challenging due to the conceptual gaps between legal and technical approaches to defining privacy.

Notably, information privacy laws are generally:

- context-specific,
- subject to interpretation,
- allow for some degree of flexibility, and
- rely on traditional, often heuristic, conceptions of privacy,

which creates uncertainty for the implementation of more formal approaches.

Gaps between Technical & Legal Concepts

1. Generality of protection afforded

- Regulatory requirements vary according to industry sector, jurisdiction, institution, types of information, and other contextual factors
  - e.g., FERPA, MA data security regulation, Privacy Act
- Challenges: In practice, privacy risks are not limited to the information categories and contexts contemplated by regulations.
  - Further, analysts may combine information from different contexts.
- In contrast, formal privacy models like differential privacy can be applied wherever statistical or machine learning analysis is performed, regardless of context, and protect all information specific to an individual.

Gaps between Technical & Legal Concepts

2. Scope of attacks contemplated

- Regulations often contemplate a limited set of specific attacks and failures.
    - e.g., **record linkage** (the re-identification of one or more records in a de-identified dataset by uniquely linking those records with identified records in a publicly available dataset) is often the primary or sole failure mode.
- Challenges: Other privacy attacks are identified over time.
    - e.g., reconstruction attacks, confirming an individual's presence in a dataset, singling out an individual (even if not fully identified), inferring information specific to an individual with less than absolute certainty.
- Formal privacy models provide protection against a wide collection of privacy attacks, even those that are not currently known.

Gaps between Technical & Legal Concepts

3. Expectations vs. the scientific understanding

- Regulations that rely on the concept of de-identification or anonymization are often not in agreement with the current scientific understanding of privacy.
    - They may be limited in scope, may not provide an adequate level of privacy in practice, and may not withstand rigorous, formal mathematical scrutiny.
    - e.g., HIPAA Privacy Rule safe harbor method
        - Redaction of identifiers can fail to protect privacy, especially when applied to detailed information such as medical records.
    - Any information, even information not traditionally considered identifying, has the potential to leak information specific to individuals.

Gaps between Technical & Legal Concepts

3. Expectations vs. the scientific understanding

- Statutes may be interpreted to require something that is not technically feasible (i.e., absolute privacy protection when sharing personal data).
  - e.g., Title 13, U.S. Code protects the confidentiality of respondent information protected by the US Census Bureau, prohibiting any publication whereby the data furnished by an individual "can be identified."
  - If this concept were interpreted very conservatively, Title 13 would disallow any leakage of information about individuals.

Gaps between Technical & Legal Concepts

3. Expectations vs. the scientific understanding

- Binary view of privacy found in Title 13—whereby information is either identifiable or not—is common to many regulations.
- Issues with this approach:
  - (1) Information can never be made completely non-identifiable and
  - (2) Fails to recognize that privacy loss accumulates with successive releases of information about the same individuals (and can eventually amount to a significant disclosure of personal information).
- In contrast, formal privacy models bound the privacy leakage of each release, and bound the total privacy leakage across multiple releases.

Gaps between Technical & Legal Concepts

4. (In)-stability over time

- Notions of privacy embedded in regulations are continually evolving.
  - e.g., OMB guidance updated over time to address new ways de-identified data may be vulnerable to potential attacks.
    - ➢ New 2017 guidance advises that non-PII may become PII in the future.
  - As hard-wired techniques (e.g., HIPAA safe harbor) are shown to be inadequate to protect privacy, it is unclear how to satisfy regulatory standards that are out of step with best practice.
- In contrast, differential privacy is the subject of ongoing scientific research and, regardless of implementation, there is a strong assurance that it provides a sufficient level of privacy in a wide variety of settings.

Gaps between Technical & Legal Concepts

5. Relationship to normative expectations

- Comparing the protection afforded by a particular privacy technology to a normative standard must be done with care.
- E.g., it may be difficult to make a sufficiency claim with respect to differential privacy and regulatory requirements.
    - Formal privacy models are “privacy-first” definitions, and real-world uses of data may demand a compromise between privacy and accuracy.
    - Regulations express requirements that can be interpreted to exceed the protection provided by formal privacy models.
        - E.g., Title 13 (especially with respect to protecting establishments)

How can privacy practitioners and regulators overcome these gaps?

Opportunities for Bridging Privacy Definitions

Approach #1: Modeling Legal Requirements Formally

Is it possible to bridge these very different languages?

$M: X^n \to T$ satisfies ϵ-differential privacy if

$\forall x, x' \in X^n$ s.t. $dist_H(x, x') = 1$ $\forall S \subseteq T$,

$$\Pr_{\mathrm{M}}[M(x) \in S] \le e^{\epsilon} \Pr_{\mathrm{M}}[M(x') \in S].$$

Formal Modeling

We seek a methodology for rigorously arguing that a technological privacy solution satisfies the requirements of a particular law.

The proposed approach has two components:

1. Extraction of a formal mathematical requirement of privacy based on a legal standard found in an information privacy law, and
2. Construction of a rigorous mathematical proof for establishing that a technological privacy solution satisfies the mathematical requirement derived from the law.

Illustration: Formally Modeling FERPA

Goal: To extract a formal model of the Department of Education's privacy desiderata for FERPA, in the form of a game-based privacy definition:

- Provides a concise and fairly intuitive abstraction of the requirements in FERPA.
- Enables us to prove that if a formal model, such as differential privacy, satisfies the game-based definition, then we have a strong argument that it satisfies the requirements of FERPA.

Although FERPA is not written with a privacy game framework in mind, we claim (and demonstrate) that it is possible to extract a game that is based on its requirements.

FERPA: Family Educational Rights and Privacy Act

Protects personally identifiable information in education records maintained by educational agencies and institutions, including

“names, addresses, personal identifiers (e.g., SSNs, student numbers, biometric records), indirect identifiers (e.g., date of birth, place of birth, mother’s maiden name), other information that, alone or in combination, is linked or linkable to a specific student that would allow a reasonable person in the school community, who does not have personal knowledge of the relevant circumstances, to identify the student with reasonable certainty, or information requested by a person who the educational agency or institution reasonably believes knows the identity of the student [in the requested record].”

(20 C.F.R. § 99.3)

FERPA: Family Educational Rights and Privacy Act

Permits the release of **de-identified information**, without consent,

“after the removal of all personally identifiable information provided that the educational agency or institution or other party has made a reasonable determination that a student’s identity is not personally identifiable, whether through single or multiple releases, and taking into account other reasonably available information.”

20 C.F.R. § 99.31(b)(1)

FERPA: Family Educational Rights and Privacy Act

Permits the release of **directory information**, as long as the students (or, if minors, their parents) have received notice and an opportunity to opt out.

"'Directory information' means information contained in an education record of a student that would not generally be considered harmful or an invasion of privacy if disclosed."

20 C.F.R. § 99.3

Extracting a Formal Definition from FERPA

FERPA allows the release of de-identified information and directory information from education records.

De-identification can be thought of in terms of a computation; e.g., requiring the removal of identifying attributes can be seen as requiring a computation to redact those identifiers from the input data.

- ➢ This framing is useful for modeling a law’s requirements using the formal language used in computer science. This modeling allows us to extract a mathematical definition for determining whether a computation meets the FERPA privacy standard.

But how do we know whether a given computation provides a sufficient level of privacy protection to meet the requirements of FERPA?

Components of a FERPA Privacy Game

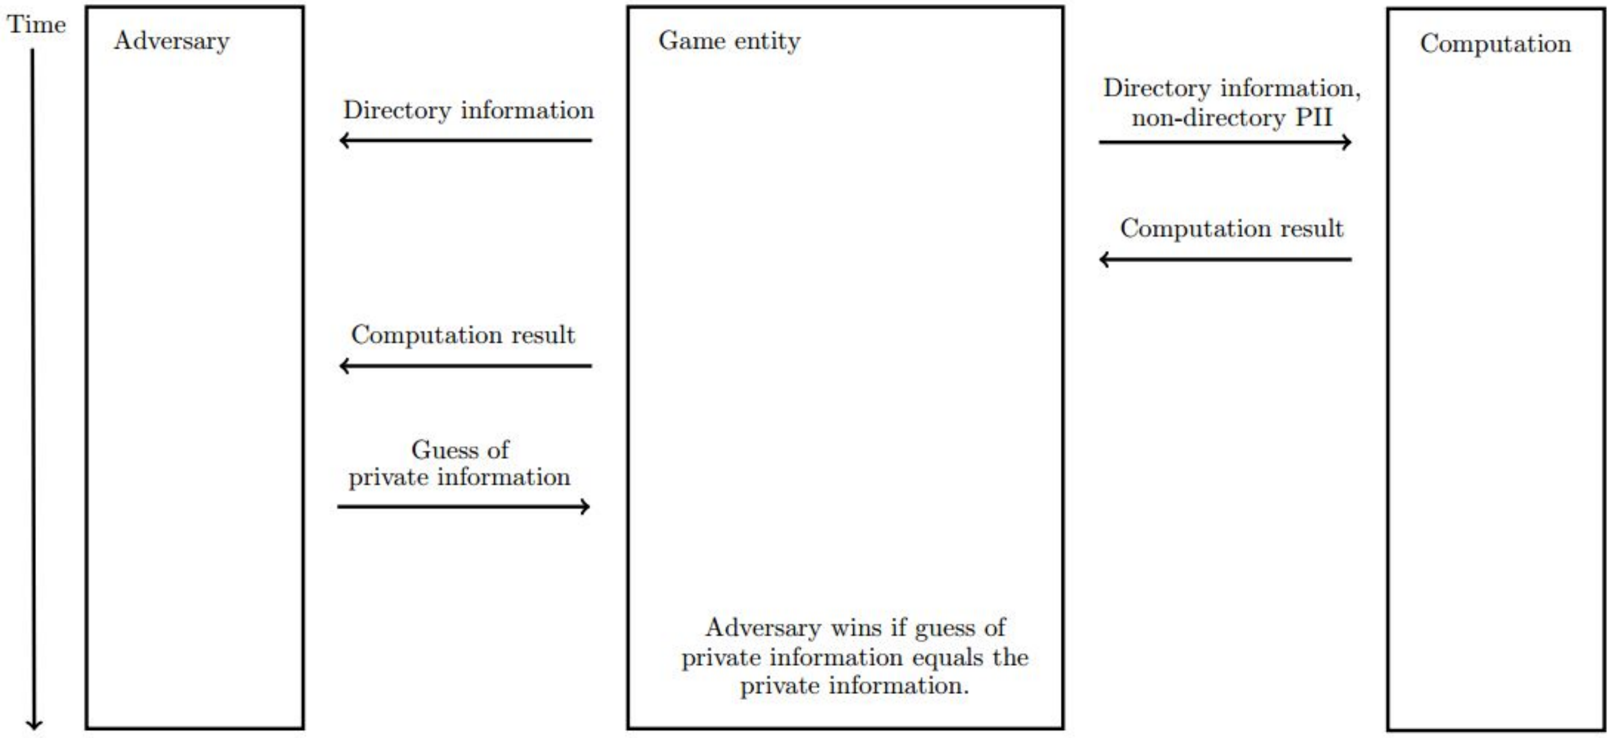

Modeling FERPA: Directory Information

The regulatory language is **ambiguous**, so we interpret the language as conservatively as reasonably possible. In other words, where there is ambiguity, we err on the side that is most beneficial for the adversary.

➢ For example, the definition of **directory information** (i.e., information that can be disclosed because it is not considered harmful) is ambiguous (e.g., the definition varies between schools).

We could make assumptions in defining directory information in our model. However, new interpretations could call these assumptions into question.

➢ Instead, **we let the attacker choose** what constitutes directory information.

Modeling FERPA: The Adversary

Personally identifiable information: "information that, alone or in combination, is linked or linkable to a specific student that would allow a reasonable person in the school community, who does not have personal knowledge of the relevant circumstances, to identify the student with reasonable certainty."

This is FERPA's **implicit adversary**. Key points from guidance:

- We should not assume anything about the skill level of the adversary.
- Standard is based on the knowledge of a member of the school community, which is stronger than one based on the knowledge of any reasonable person.
- The adversary can have both high-level knowledge (e.g., demographics of school) and "insider" knowledge about specific individuals in local community.

Modeling FERPA: Adversary’s Knowledge

The adversary clearly has (potentially a lot of) knowledge, but by definition does not have “personal knowledge of the relevant circumstances.”

In our model, the adversary has access to any information that is publicly available, but has some uncertainty about private student information.

We model the adversary’s knowledge via probability distributions. Adversary associates with each student a probability distribution that represents her knowledge about the private information of that student. We allow the adversary to choose these statistics.

Example: If Alice comes from a school where 50% of the students failed the state math proficiency exam, then adversary might associate with Alice a distribution that has her failing the exam with a probability of 0.5.

Proving Differential Privacy Satisfies FERPA

Developing a formal definition of privacy protection based on the requirements of FERPA allows us to reason, with high confidence, about whether the use of a privacy technology satisfies FERPA.

- We can prove mathematically that any computation that is differentially private meets this definition, and (since the requirements of this definition are likely stricter than that of FERPA) thus satisfies the privacy requirements of FERPA.

Opportunities for Bridging Privacy Definitions

Approach #2: Integrating Modern Privacy Approaches Across the Information Lifecycle

Framework for Privacy-Aware Data Releases

Modeled on information security and lifecycle frameworks:

1. Developing a catalog of privacy controls
2. Identifying information uses, threats, and vulnerabilities
3. Designing data releases by aligning uses and risks with controls—at each stage of the information lifecycle

Catalog of Privacy Controls

Procedural, technical, educational, economic, and legal means for enhancing privacy—at each stage of the information lifecycle

	Procedural	Economic	Educational	Legal	Technical

Access/Release	Access controls;
	Consent;
	Expert panels; Individual privacy settings;
	Presumption of openness vs. privacy;
	Purpose specification;
	Registration;
	Restrictions on use by data controller;
	Risk assessments
		Access/Use fees (for data controller or subjects);
		Property rights assignment
			Data asset registers;
			Notice;
			Transparency
				Integrity and accuracy requirements; Data use agreements (contract with data recipient)/
				Terms of service
					Authentication;
					Computable policy;
					Differential privacy;
					Encryption (incl. Functional; Homomorphic);
					Interactive query systems;
					Secure multiparty computation

Guide to Selecting Appropriate Privacy Controls

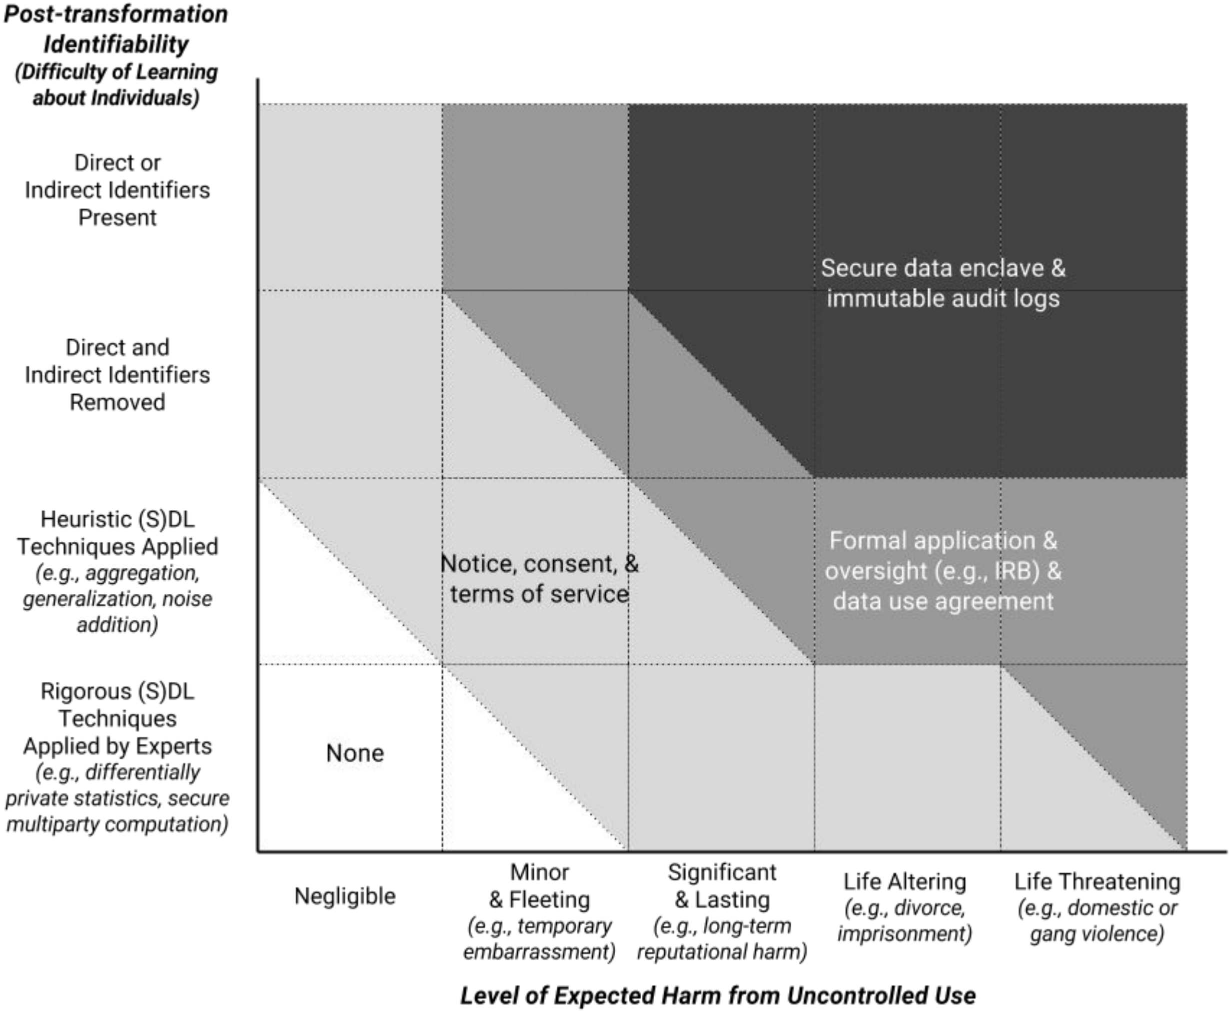

Post-transformation Identifiability
(Difficulty of Learning about Individuals)

Direct or Indirect Identifiers Present

Direct and Indirect Identifiers Removed

Heuristic (S)DL Techniques Applied *(e.g., aggregation, generalization, noise addition)*

Rigorous (S)DL Techniques Applied by Experts *(e.g., differentially private statistics, secure multiparty computation)*

Secure data enclave & immutable audit logs

Formal application & oversight (e.g., IRB) & data use agreement

Notice, consent, & terms of service

None

Negligible

Minor & Fleeting *(e.g., temporary embarrassment)*

Significant & Lasting *(e.g., long-term reputational harm)*

Life Altering *(e.g., divorce, imprisonment)*

Life Threatening *(e.g., domestic or gang violence)*

Level of Expected Harm from Uncontrolled Use

Example Tiered Access Model

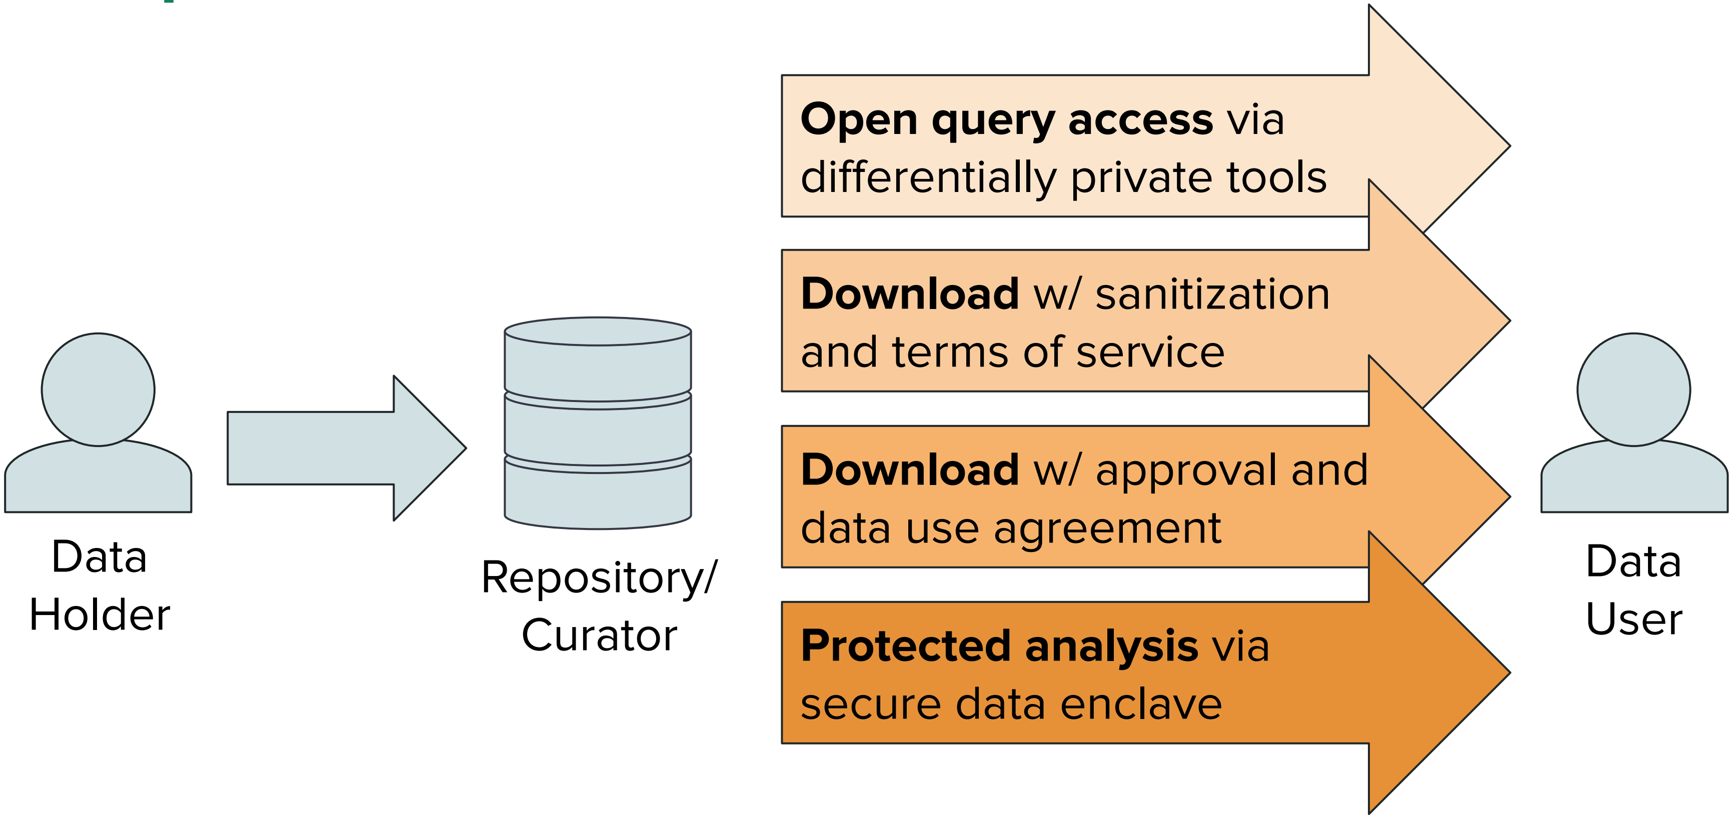

Data Holder

Repository/ Curator

Open query access via differentially private tools

Download w/ sanitization and terms of service

Download w/ approval and data use agreement

Protected analysis via secure data enclave

Data User

Inspired by Privacy Tools Project Model

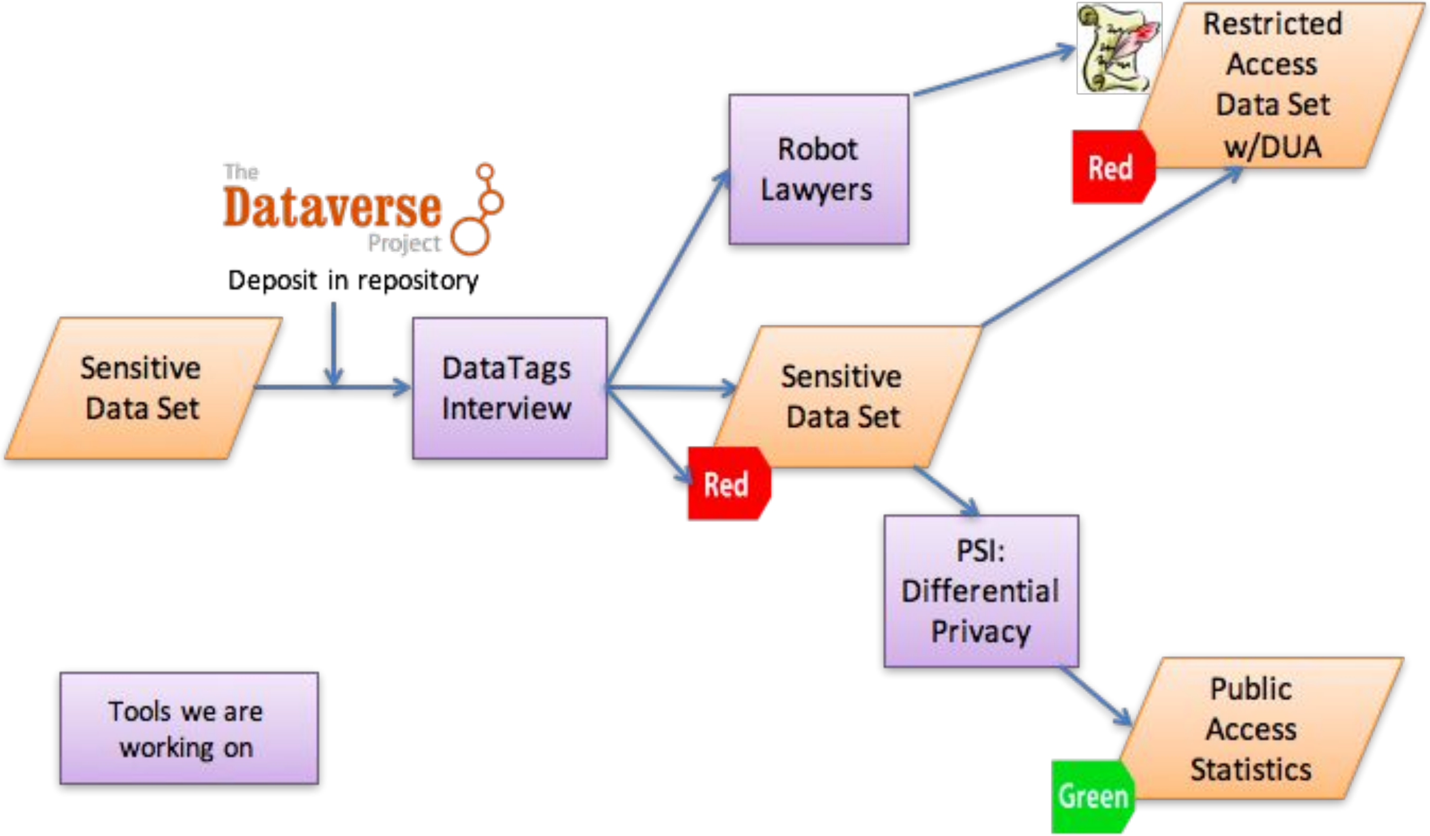

Opportunities for Bridging Privacy Definitions

Approach #3: Interpreting Formal Privacy Guarantees

Interpreting the Differential Privacy Guarantee

- Legal requirements relevant to issues of privacy in computation rely on an understanding of a range of different privacy concepts.
- None of the privacy concepts that appear in the law refer directly to differential privacy
    - However, the differential privacy guarantee can be interpreted in reference to these concepts—while accommodating differences in how these concepts are defined across contexts.

Common Privacy Concepts in the Law

- Personally identifiable information
- De-identification
- Linkage
- Inference
- Risk
- Consent and opting out
- Purpose and access restrictions

These concepts are interpreted differently across laws. They also appear in the technical literature, often with different definitions and interpretations.

Example: Personally Identifiable Information

Personally identifiable information is a central concept appearing in privacy law.

- Legal protections typically extend only to PII, and information not considered personally identifiable is not protected (e.g., FERPA, HIPAA Privacy Rule).
- Although definitions vary significantly, they are generally understood to refer to the presence of pieces of information that are linkable to the identity of an individual or to an individual's personal attributes.
- PII is also related to the concept of *de-identification*, which refers to a collection of techniques, that if performed successfully, used as to remove PII, or transform PII into non-PII.

PII: Interpretation of DP Guarantee

PII does not have a precise technical meaning.

In practice it can be difficult to determine whether information is personal, identifying, or likely to be considered identifying in the future.

Further, the meaning of PII in releases that are not in a microdata or tabular format, such as statistical models or outputs of a machine learning system, is unclear.

However, when differential privacy is used, it can be understood as ensuring that using an individual's data will not reveal essentially any personally identifiable information specific to her, regardless of the definition of PII that is used.

- Here, *specific* is used to refer to information that is unique to the individual and cannot be inferred unless the individual's information is used in the analysis.

Opportunities for Future Regulation

- **Regulations should articulate clear goals for privacy protection.**
  - These goals should be line with the scientific understanding of privacy.
  - Regulations should move away from implicitly or explicitly endorsing ad hoc de-identification techniques.
- **Example: Guidance on European data protection law outlines goals that go beyond the traditional notion of de-identification.**
  - Protection from singling out, linking, or inferring an individual’s personal data from a dataset.
  - These concepts have not yet been defined precisely and formally from a mathematical perspective, but they aim to describe a clearer goal.

Conclusion

How can practitioners and regulators grapple with tensions between current regulatory standards and advances in both privacy attacks and technologies?

- Challenging due to gaps between technical and normative definitions.
- There are at least four promising approaches to this problem:
   1. Extracting a formal mathematical model of a regulation to make a claim that a privacy technology satisfies a legal requirement.
   2. Integrating modern privacy approaches into a data management plan.
   3. Interpreting formal privacy guarantees in relation to terms used in law and policy.

References

Kobbi Nissim, Aaron Bembenek, Alexandra Wood, Mark Bun, Marco Gaboardi, Urs Gasser, David R. O’Brien, and Salil Vadhan, **Bridging the Gap between Computer Science and Legal Approaches to Privacy**, 31 *Harvard Journal of Law & Technology* 687 (2018), https://dash.harvard.edu/handle/1/37355739

Alexandra Wood, Micah Altman, Aaron Bembenek, Mark Bun, Marco Gaboardi, James Honaker, Kobbi Nissim, David O'Brien, Thomas Steinke, and Salil Vadhan, **Differential Privacy: A Primer for a Non-technical Audience**, 21 *Vanderbilt Journal of Entertainment & Technology Law* 209 *(2018),* https://dash.harvard.edu/handle/1/38323292

Kobbi Nissim and Alexandra Wood, **Is Privacy *Privacy*?**, *Philosophical Transactions of the Royal Society A* 376:2128 (2018), https://dash.harvard.edu/handle/1/38021438

Micah Altman, Alexandra Wood, David R. O’Brien, Salil Vadhan, and Urs Gasser, **Towards a Modern Approach to Privacy-Aware Government Data Releases**, 30 *Berkeley Technology Law Journal* 1967 (2015), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2779266