Privacy & Ethical Considerations for AI and Big Data Research

Alexandra Wood
Berkman Klein Center for Internet & Society at Harvard University

Discussion Forum: Facing Artificial Intelligence
Montréal, Québec · December 3, 2018

These opinions are my own. They are not the opinions of the Berkman Klein Center, any of our funders, nor (with the exception of co-authorship on previously published work) my collaborators.

Motivation

Vast quantities of rich, fine-grained data related to human biology, behaviors, and relationships are being created by new services and methods

- e.g., social networks, mobile apps, DNA sequencing

AI and big data analysis hold tremendous potential for scientific inquiry, public health, and clinical practice

- e.g., computational social science, biomedical big data research

Motivation

New risks and gaps in oversight for big data and AI research

- Facebook emotional contagion study: **gaps in scope of coverage**
- Genomic DNA databases: **new privacy risks**
- Cambridge Analytica: **other emerging risks (e.g., manipulation)**

Commonly used measures for protecting research subjects fail to address oversight challenges in big data and AI research

Oversight Challenges

- Narrow scope of regulatory definitions (e.g., *human subjects research*, *identifiable information*, and *private information*)
- Emphasis on the study design and collection lifecycle stages—and, to a much lesser extent, dissemination and re-use of data
- Lack of oversight for privately-funded research
- Inadequacy of common approaches (e.g., informed consent and de-identification)
- Reduced utility resulting from heavily redacting or withholding data

Failures of De-identification

Advances in the scientific understanding of privacy have demonstrated that privacy approaches relying exclusively on de-identification fail to provide reasonable protection.

While they may reduce some risks, traditional de-identification approaches

- do not prevent all disclosures or protect information in the manner that most individual subjects would expect,
- address only a subset of privacy attacks and attackers,
- are not readily scalable for use by non-experts, and
- often result in the redaction or withholding of useful information.

Recommendations for a New Ethical Framework

- Robust human subjects protection is essential to
  - safeguarding the interests of research subjects;
  - ensuring trust, transparency, and accountability in the research community; and
  - maintaining continued support for, funding of, and participation in research studies.
- Recognition of the human right to participation in the production of scientific knowledge.
- Gaps in the current framework demonstrate that changes are needed to continue to advance these values in light of AI and big data research.

Lifecycle Approach to Data Management

- Review of risks and intended uses at each lifecycle stage, as information is used over time.
- Selection of appropriate controls at each lifecycle stage.

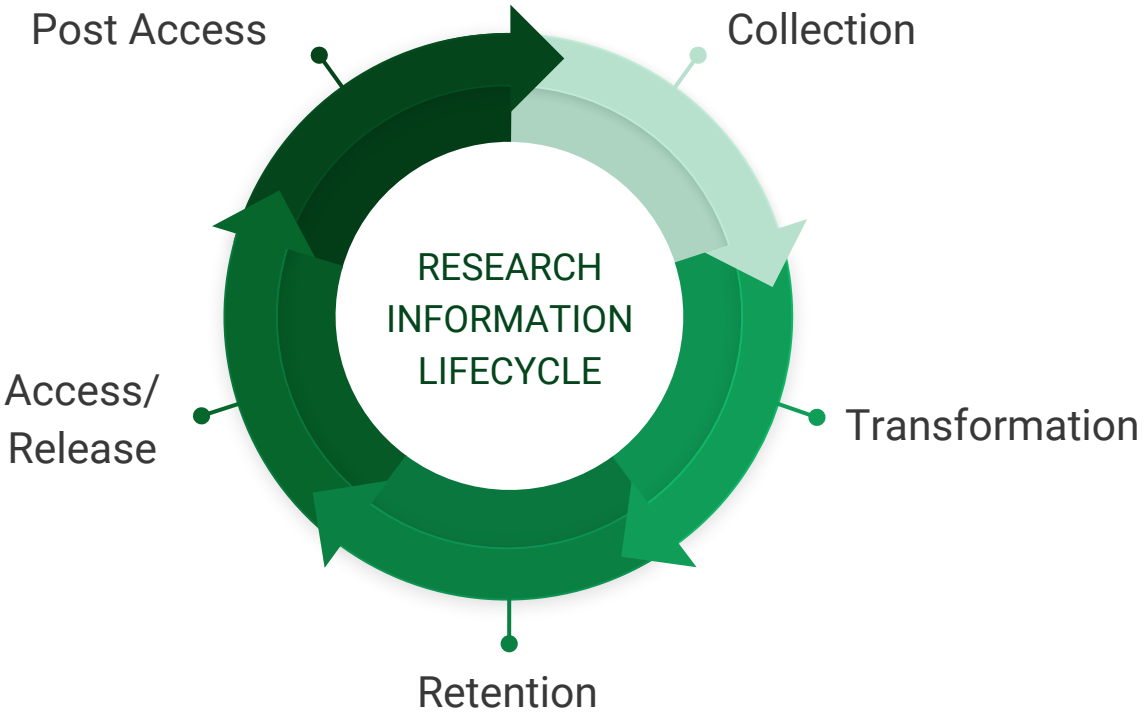

Recommendation #1: Universal Coverage

Oversight should cover the full scope of human subjects research.

- Addressing gaps in coverage for research involving:
  - Many categories of information deemed to be public or non-identifiable,
  - Privately funded research, and
  - Research activities across all stages of the lifecycle including the storage, processing, analysis, release, and post-release stages.
- Sharing responsibilities with consumer review boards, participant-led review boards, and personal data cooperatives.
- Exploring technological solutions for automating decisions.

Recommendation #2: Conceptual Clarity

Revised definitions and guidance should be developed by an expert body.

- Reliance on a sharp binary determination based on the presence of *identifiable private information* leads to inconsistency and uncertainty.
- Revised definitions should be developed based on a modern understanding of privacy, moving away from a binary view of *identifiability* or *public availability*.
- Regulators and review boards should consider consulting with ethics and privacy experts—or establishing a regularly-convening advisory committee—to provide concrete recommendations on clarifying definitions, practices, methodologies, and guidelines for implementation.

Recommendation #3: Risk-Benefit Assessments

Researchers and review boards should be encouraged to perform comprehensive, systematic risk-benefit assessments.

- Risk-benefit assessments should evaluate
  - the benefits that would accrue to society as a result of a research activity,
  - the intended uses of the information,
  - the privacy threats and vulnerabilities associated with the research activity, and
  - the potential harms to human subjects as a result of the inclusion of their information in the data.

Recommendation #4: Adopting New Solutions

Incentives to explore and adopt emerging technical and procedural interventions are needed.

- Researchers should draw from the wide range of available privacy measures, rather than rely on consent and de-identification alone.
- Requirements for consent and de-identification could be amended to provide that in many cases they should be combined with additional controls, including information security controls.
- Regulators, in consultation with an expert body of privacy researchers, IRB administrators, and researchers, could compile a list of approved techniques that provide a strong guarantee of privacy protection.

Wide Range of Available Interventions

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

Recommendation #5: Tailored Oversight

Oversight should be tailored to the specific risks and uses involved in a given research activity.

- Different research activities can be reviewed by different review boards.
- Data releases can be tailored through tiered access, for example:
  - **Public access**: contingency tables generated with formal privacy guarantees
  - **Intermediate access**: protected server access for approved researchers
  - **Restricted access**: raw data shared through a monitored environment
- Similar mechanisms can be implemented at other stages.

Guiding the selection of appropriate privacy controls

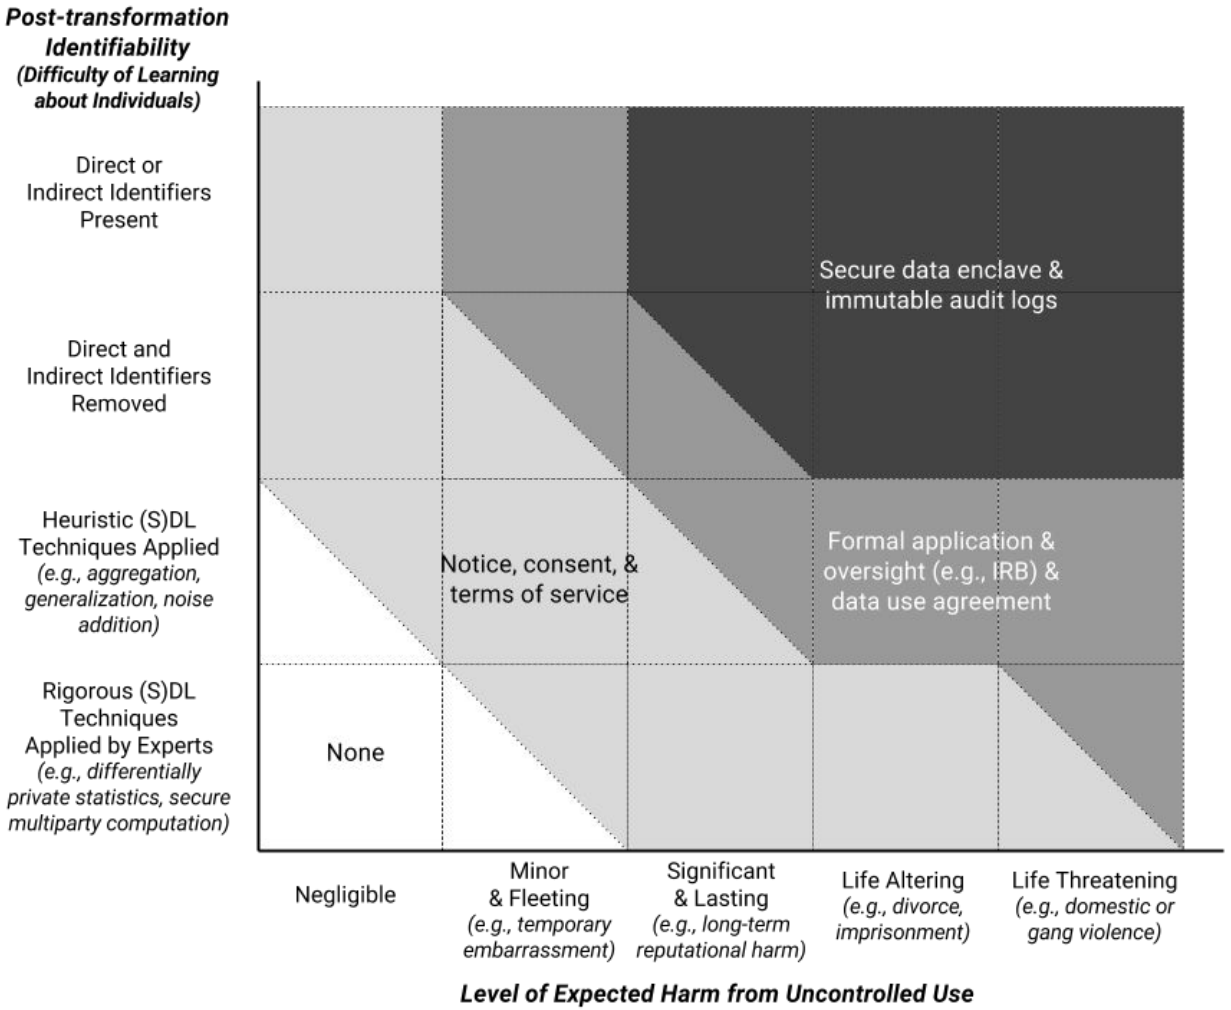

Example Tiered Access Model

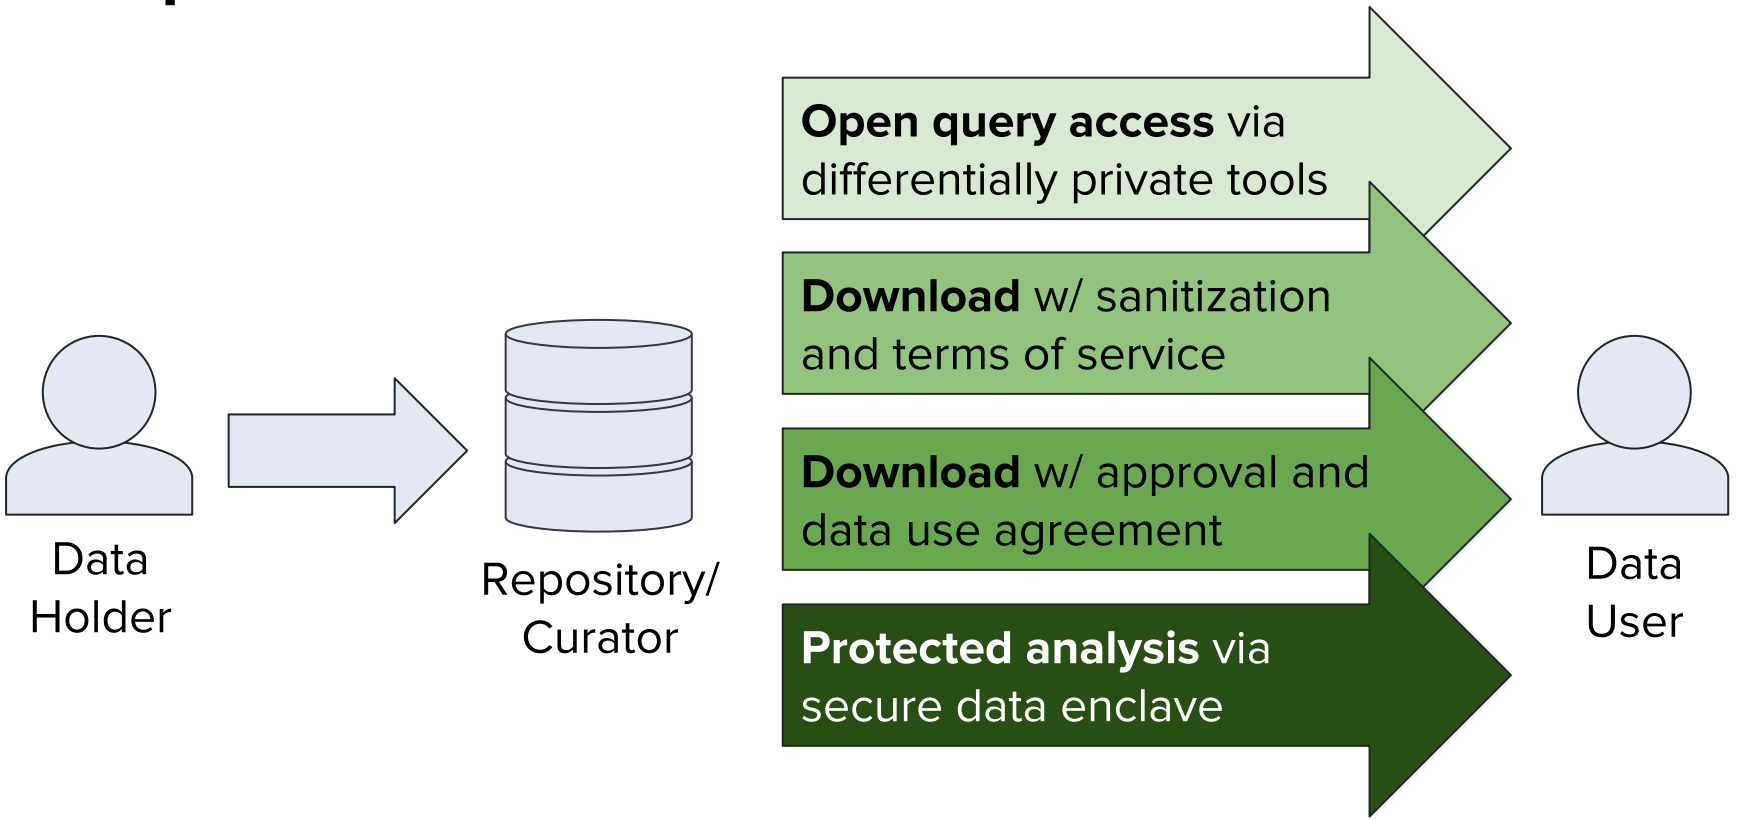

Harvard University Privacy Tools Project

The Privacy Tools Project is a broad effort to advance a multidisciplinary understanding of data privacy issues and build computational, statistical, legal, and policy tools to help address these issues in a variety of contexts. It is a collaborative effort between Harvard's Center for Research on Computation and Society, Institute for Quantitative Social Science, Berkman Klein Center for Internet & Society, and Data

LATEST NEWS & BLOG POSTS

Alexandra Wood and Stephen Chong to speak on "Robot Lawyers: Automating Legal Compliance for Transferring Private Data"

Latanya Sweeney Holds Public Debate with French Secretary of State for Digital Commerce; Consults with President Macron on Tricolor Strategy on Artificial Intelligence

Comments Cited in Future of Privacy Forum's Final Report on City of Seattle's Open Data Program

Alexandra Wood Presents "Differential Privacy: A Primer for a Non-technical Audience" at Differential Privacy Webinar

Kimia Mavon wins Venky Award for Community Building

Aloni Cohen and Kobbi Nissim Inducted into

Overview of Privacy Tools Project

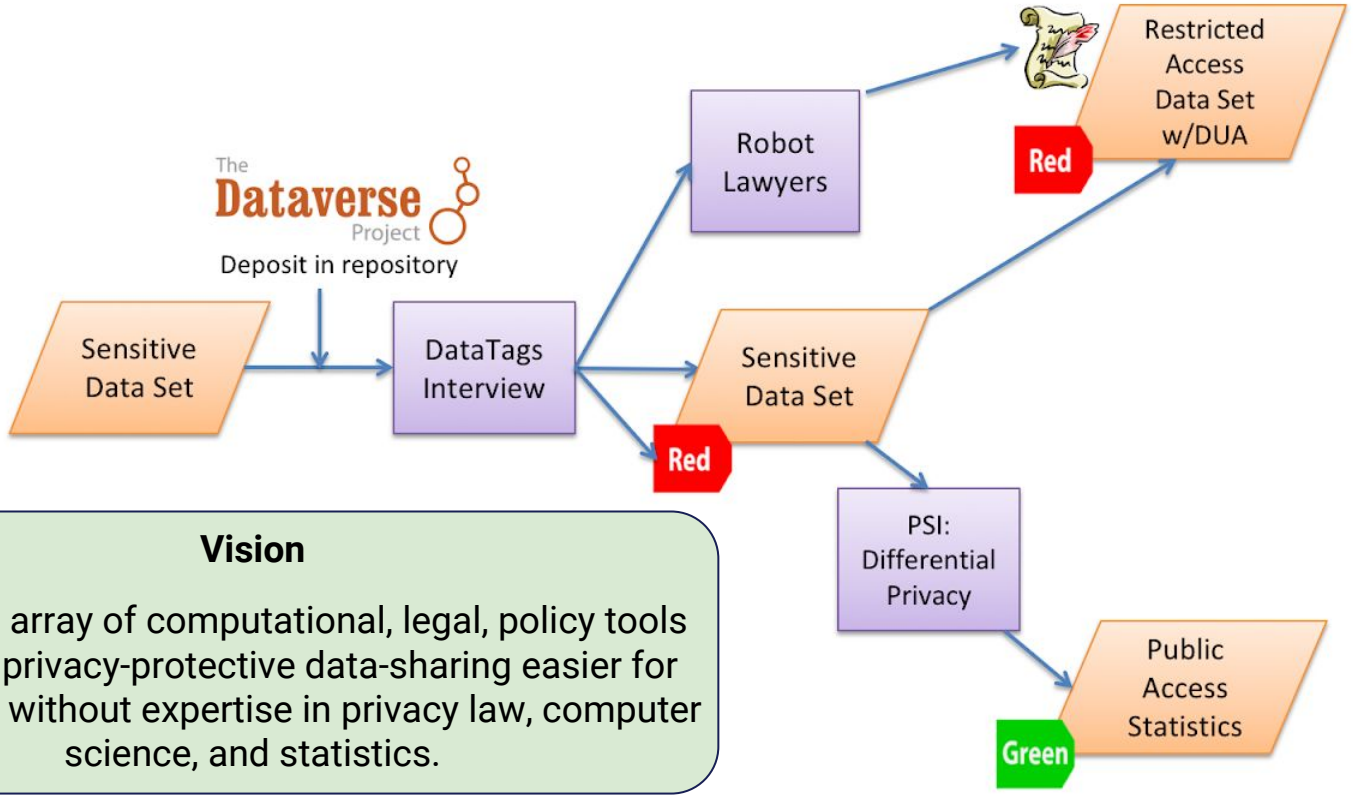

Vision

To build an array of computational, legal, policy tools that make privacy-protective data-sharing easier for researchers without expertise in privacy law, computer science, and statistics.

Privacy Tools Project: Robot Lawyers

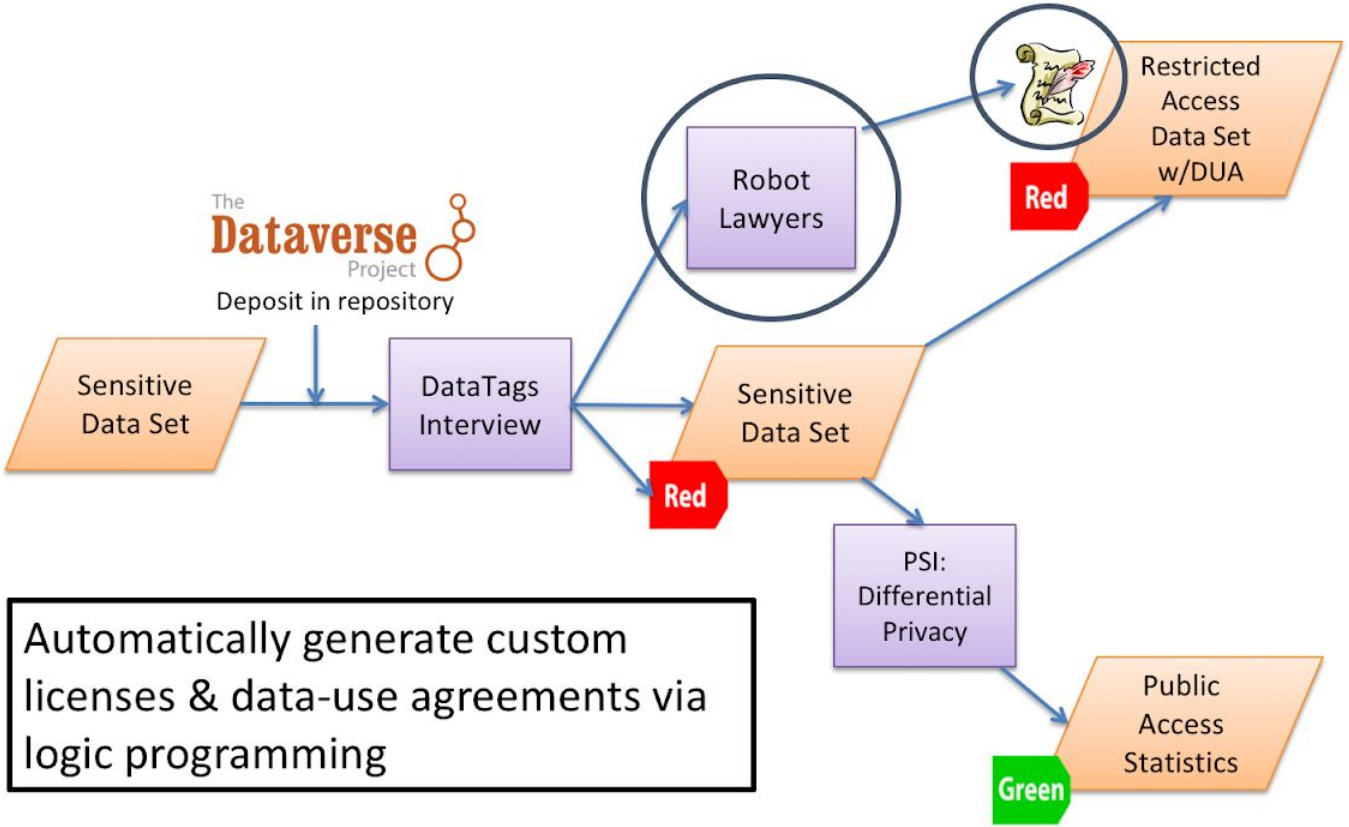

Privacy Tools Project: PSI Differential Privacy Tool

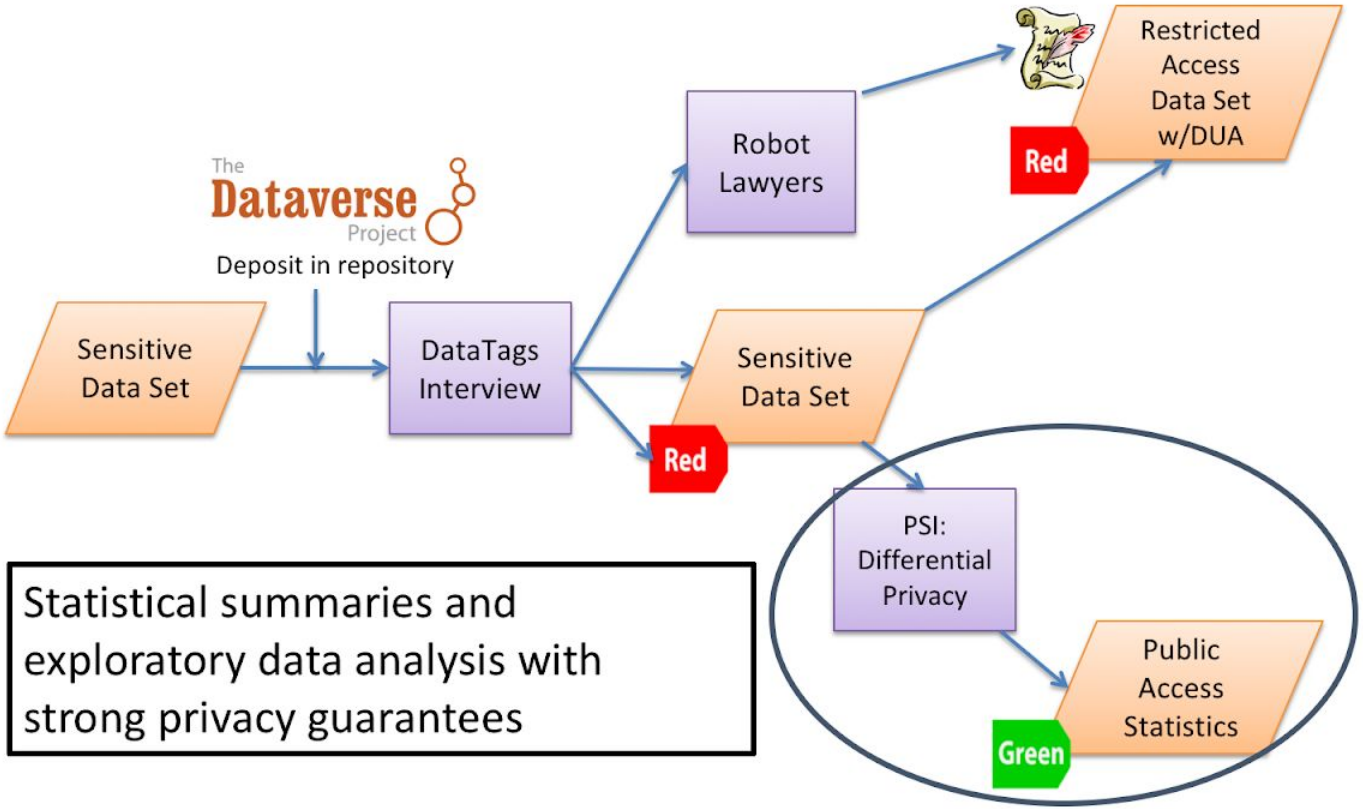

Privacy Tools Project: Bridging Privacy Definitions

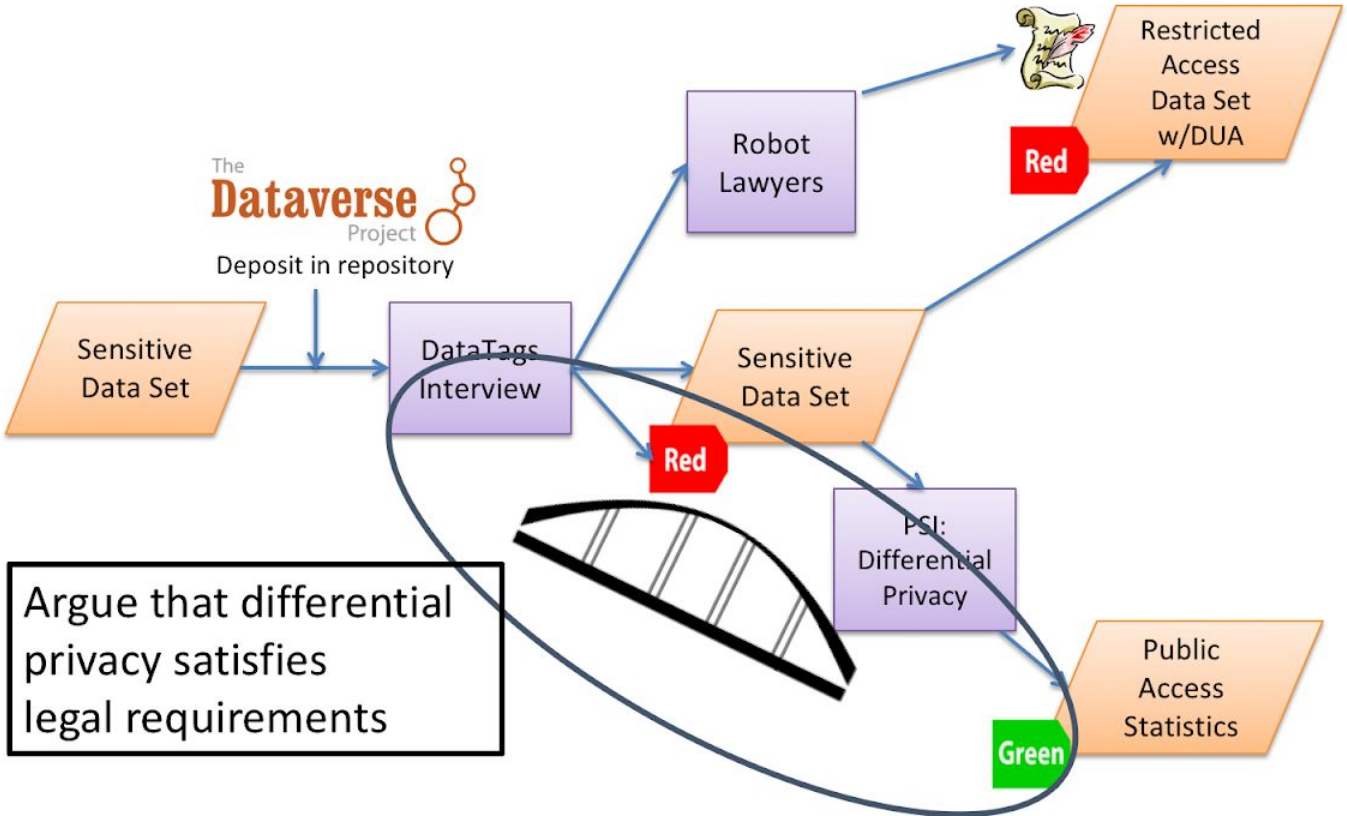

Conclusion

- New large-scale data sources and modes of analysis hold tremendous promise for research.
    - However, AI and big data research present new risks that the current regulatory framework is ill-suited to address.
    - How can the value of the data can be captured in a way that respects fundamental principles of ethics and privacy?
- Updating the oversight framework would help enable the collection, use, and sharing of big data in line with ethical principles, research community norms, and expectations of human subjects.
- Achieving this balance will be critical to ensuring the trust and support of the public and, ultimately, the viability of AI and big data research.

Related Work

- Micah Altman, Alexandra Wood, David O'Brien, and Urs Gasser, "Practical Approaches to Big Data Privacy Over Time," 8 *International Data Privacy Law* 29 (2018)
- Effy Vayena, Urs Gasser, Alexandra Wood, David R. O'Brien, and Micah Altman, "Elements of a New Ethical Framework for Big Data Research," 72 *Washington & Lee Law Review Online* 420 (2016)
- Micah Altman et al., "Towards a Modern Approach to Privacy-Aware Government Data Releases," 30 *Berkeley Technology Law Journal* 1967 (2015)
- Alexandra Wood et al., Comments to the Department of Health and Human Services Re: Federal Policy for the Protection of Human Subjects; Proposed Rules, Docket No. HHS–OPHS–2015–0008 (Jan. 6, 2016)
- Salil Vadhan et al., Comments to the Department of Health and Human Services Re: Advance Notice of Proposed Rulemaking: Human Subjects Research Protections: Enhancing Protections for Research Subjects and Reducing Burden, Delay, and Ambiguity for Investigators, Docket No. HHS-OPHS-2011-0005 (Oct. 26, 2011)
- David R. O'Brien et al., "Integrating Approaches to Privacy Across the Research Lifecycle: When Is Information Purely Public?," Berkman Center Research Publication 2015-7 (2015)