Federal Register Notice 87 FR 35250, https://www.federalregister.gov/documents/2022/06/09/2022-12432/request-for-information-on-advancing-privacy-enhancing-technologies, June 9, 2022

Request for Information (RFI) on Advancing Privacy Enhancing Technologies

Altman, Micah; Cohen, Aloni; and Vadhan, Salil

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.

July 28, 2022
Via electronic filing

Jeri Hessman
Technical Coordinator
Fast-Track Action Committee (FTAC) on Advancing Privacy-Preserving Data Sharing and Analytics.
NITRD/NCO
Email: <*PETS-RFI@nitrd.gov*>

Re: *RFI Response: Privacy-Enhancing Technologies*

Dear Members of the Committee,

This comment is informed by research with collaborators through the *Privacy Tools Project* at Harvard University.[1] In this broad, multidisciplinary project, we are exploring the privacy issues that arise when collecting, analyzing, and disseminating research datasets containing personal information. Our efforts are focused on translating the theoretical promise of new privacy protection and data utility measures into practical tools and approaches. In particular, our work aims to help realize the tremendous potential from social science research data by making it easier for researchers to share their data using privacy-protective tools.

Academic research in theoretical computer science, statistics and information science has demonstrated many challenges related to managing information privacy in the modern world.

In previous scholarly publications, we have offered several recommendations that we believe would help enable the wider sharing of research data while providing privacy protection for individuals.[2]

[1] See Privacy Tools for Sharing Research Data, http://privacytools.seas.harvard.edu. Also see prior Privacy Tools Group responses to related policy RFI's here: https://privacytools.seas.harvard.edu/policy-engagement .

[2] See Altman M, Wood A, O'Brien D, Vadhan S, Gasser U. Towards a Modern Approach to Privacy-Aware Government Data Releases. Berkeley Technology Law Journal 2015; 30(3):1967-2072;
Vayena E, Gasser U, Wood A, O'Brien D, Altman M. Elements of a New Ethical and Regulatory Framework for Big Data Research. Washington and Lee Law Review. 2016;72(3):420-442.
Altman, Micah, Alexandra B. Wood, David O'Brien, and Urs Gasser. "Practical approaches to big data privacy over time." International Data Privacy Law 8(1):29-51. (2018).
Fluitt, J. Aaron and Cohen, Aloni and Altman, Micah and Nissim, Kobbi and Viljoen, Salome and Wood, Alexandra, Data Protection's Composition Problem (September 9, 2019). European Data Protection Law Review (EDPL), Vol. 5, Iss. 3 (2019)
Altman, Micah and Cohen, Aloni and Falzon, Francesca and Markatou, Evangelia Anna (Lilika) and Nissim, Kobbi and Reymond, Michel Jose and Saraogi, Sidhant and Wood, Alexandra, A Principled Approach to Defining Anonymization As Applied to EU Data Protection Law (May 9, 2022). Available at SSRN: https://ssrn.com/abstract=4104748 or http://dx.doi.org/10.2139/ssrn.4104748

`

Although our previous writings do not comment directly on all of the privacy-enhancing technologies under consideration by the committee, the writers judge that the risks discussed in these works apply to protected health information and that the broad findings and recommendations are readily applicable here. For these reasons, we recommend that the committee read and incorporate these recommendations, which are summarized below.

As a general framework, we have recommended the development of rules and guidance based on the following principles of a modern approach to privacy:

- Calibrating privacy and security controls to the intended uses and privacy risks associated with the data;
- When conceptualizing informational risks, considering not just reidentification risks but also inference risks, or the potential for others to learn about individuals from the inclusion of their information in the data;
- Addressing informational risks using a combination of privacy and security controls rather than relying on a single control such as consent or deidentification;
- Anticipating, regulating, monitoring, and reviewing interactions with data across all stages of the lifecycle (including the post-access stages), as risks and methods will evolve over time; and
- In efforts to harmonize approaches across regulations and institutional policies, emphasizing the need to provide similar levels of protection to research activities that pose similar risks.
- Recognizing that a single set of privacy and security controls is frequently not appropriate for all intended uses of the information. And designing for access using a tiered model is one in which data are made available to different categories of data users through different mechanisms.
- Recognizing that privacy risks constitute a specific form of informational harm. Privacy is not sufficient, even in theory, to ensure that decisions or algorithms based on personal data will be secure, non-discriminatory, explainable, reasonable, or immune to dangerous misuse.
- Recognizing that as the volume and complexity of data uses and publications grow rapidly across a broad range of contexts, it has become impossible to monitor all past data releases and anticipate all future attacks. Instead, PETs must address and control the cumulative information risks to participants.
- Rejecting security by obscurity, and instead recognizing that security requires protections to be based on public algorithms and protocols built and vetted by the greater security community.
- Recognizing that the risk of harm from the use and disclosure of information is not limited to data represented in any specific format. To be effective, definitions of use and disclosure in new or proposed legislation should be generally applicable for any type of data release—whether in the form of microdata, a summary table, an information visualization, statistical model coefficients, a trained model output by a machine-learning algorithm, a textual summary, or any other form.

`

We further note the importance of using protections that provide formal privacy guarantees where feasible. Many data-sharing models are compatible with a formal privacy guarantee called differential privacy. Differential privacy is a strong, quantitative notion of privacy that addresses both known and unforeseeable attacks, and is provably resilient to a very large class of potential misuse. In recently published work, we provide guidance on designing systems that integrate differential privacy protections throughout the information lifecycle and apply it in conjunction with a range of other complementary informational controls.[3]

In addition, the research cited above finds that addressing privacy risks requires a sophisticated approach, and the privacy protections currently employed in government releases of data do not take into account recent advances in data privacy research. We note that there is a wide range of technical, procedural, legal, educational, and economic controls available for managing privacy risks. However, most government data releases rely almost exclusively on a narrow set of interventions, namely redaction of identifiers and binary access control. This focus on a small set of controls likely fails to address the nuances of data privacy and utility, as well as the differences between data releases, which vary widely in terms of the intended uses of the data and the privacy risks involved.

This research also notes, as paraphrased, that advances in science and technology enable the increasingly sophisticated characterization of privacy risks and harms and offer new interventions for protecting data subjects. In our work,[4] we describe a lifecycle approach that supports a systematic decomposition of the factors relevant to data management at each information stage, including the collection, transformation, retention, access or release, and post-access stages. Additionally, we propose a framework for developing guidance on selecting appropriate privacy and security measures that are calibrated to the context, intended uses, threats, harms, and vulnerabilities associated with specified research activities.

Figure 1 provides a partial conceptualization of this framework.[5] In this diagram, the x-axis provides a scale for the level of expected harm from the uncontrolled use of the data, meaning the maximum harm the release could cause to some individual in the data based on the sensitivity of the information. This scale ranges from low to high levels of expected harm, with harm defined to capture the magnitude and duration of the impact a misuse of the data would have on an affected individual's life, and we have placed examples as reference points along this axis. The y-axis provides a scale for the

[3] Wood, Alexandra, Micah Altman, Kobbi Nissim, and Salil Vadhan. "Designing Access with Differential Privacy." Handbook on Using Administrative Data for Research and Evidence-based Policy,, Shawn Cole, Iqbal Dhaliwal, Anja Sautmann, and Lars Vilhuber (Eds.). Abdul Latif Jameel Poverty Action Lab, Cambridge, MA (2020).

[4] Altman M, Wood A, O'Brien D, Vadhan S, Gasser U. Towards a Modern Approach to Privacy-Aware Government Data Releases. Berkeley Technology Law Journal. 2015; 30(3): 1967-2072.

[5] This diagram originally appeared in Altman M, Wood A, O'Brien D, Vadhan S, Gasser U. Towards a Modern Approach to Privacy-Aware Government Data Releases. Berkeley Technology Law Journal. 2015; 30(3): 1967-2072.

`

post-transformation identifiability or the potential for others to learn about individuals based on the inclusion of their information in the data. Several examples are provided as anchor points, ranging from data sets containing direct or indirect identifiers, to data shared using expertly applied rigorous disclosure limitation techniques backed by a formal mathematical proof of privacy.

The level of expected harm from uncontrolled use and the post-transformation identifiability of the data, taken together, point to minimum privacy and security controls that are appropriate in a given case, as shown by the shaded regions in the diagram. Regions divided by a diagonal line correspond to categories of information for which an actor could reach different conclusions based on the intended uses of the data or privacy standards that vary based on the applicability of regulation, contract, institutional policy, or best practice. The sets of controls within the shaded regions focus on a subset of controls from the more comprehensive set of procedural, economic, educational, legal, and technical controls we catalog in the work cited above. In practice, the design of a data management plan should draw from the wide range of available interventions and incorporate controls at each stage of the lifecycle, including the post-access stage. Also, note there are regions of this diagram that deviate from current practice in some domains. For example, we argue that data that have been de-identified using simple redaction or other heuristic techniques should in many cases be protected using additional controls, even though some existing standards do not expressly call for the use of additional controls when using such techniques.

`

Figure 1. Calibrating privacy and security controls.

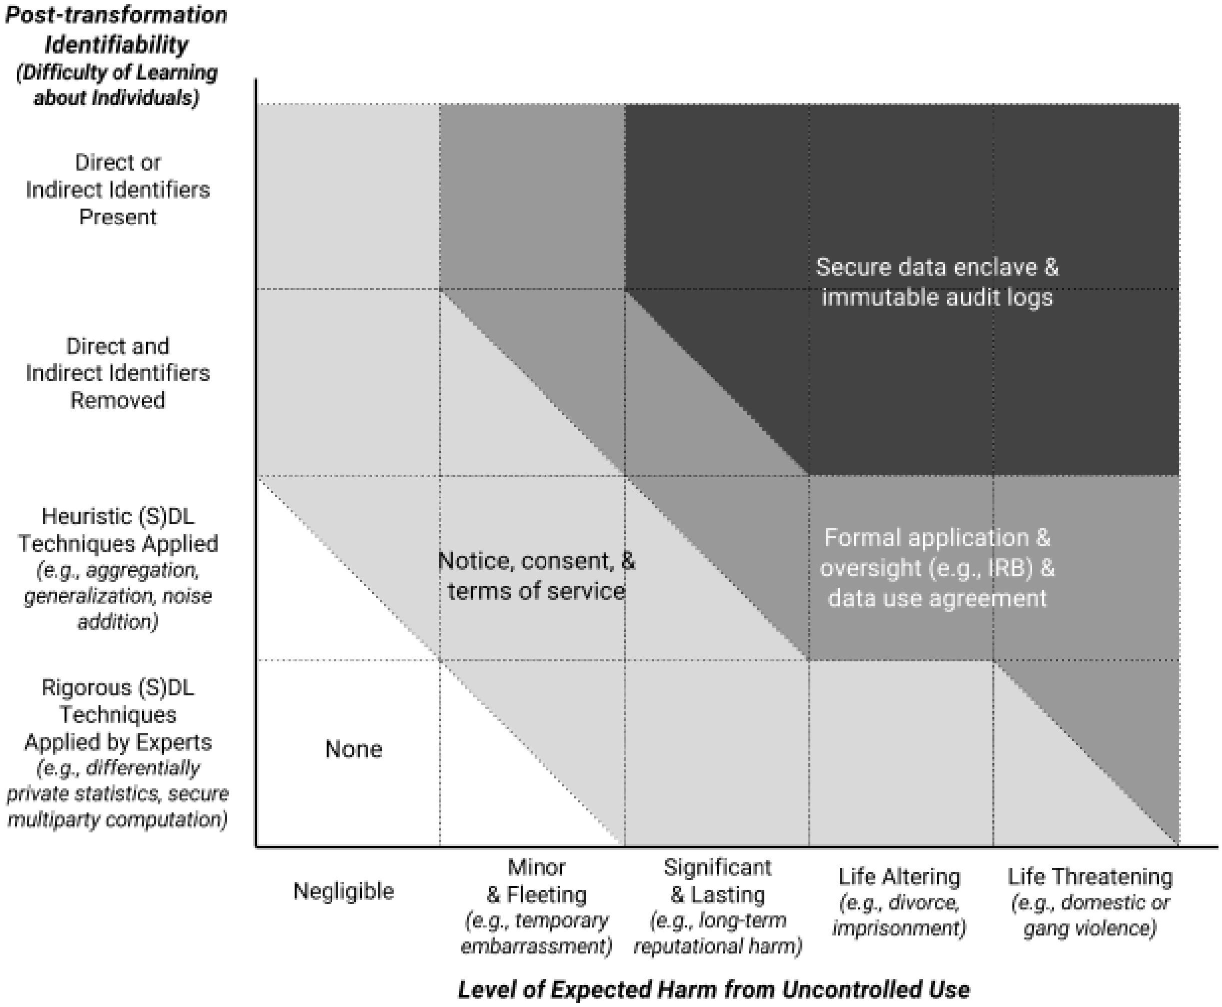
Post-transformation Identifiability (Difficulty of Learning about Individuals)
Direct or Indirect Identifiers Present
Direct and Indirect Identifiers Removed
Heuristic (S)DL Techniques Applied (e.g., aggregation, generalization, noise addition)
Rigorous (S)DL Techniques Applied by Experts (e.g., differentially 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

For many activities, implementing a single set of privacy and security controls may not be appropriate for all intended uses of the information. For this reason, we generally recommend that regulators and data controllers implement a tiered access model. A tiered access model is one in which data are made available to different categories of data users through different mechanisms.

Figure 1 illustrates the relationship between transformation and release controls, and suggests how controls could be selected for different access tiers. For example, an investigator could provide public access to some data without restriction after robust disclosure limitation techniques have transformed the data into differentially private statistics. Data users who intend to perform analyses that require the full dataset, including direct and indirect identifiers, could be instructed to submit an application to an oversight body such as an institutional review board, and their use of the data would be restricted by the terms of a data use agreement. We argue that this framework, implemented through a data management plan and tiered access model, would help

`

data providers, data users, and oversight bodies calibrate the use of privacy and security controls to the contexts, threats, harms, and vulnerabilities associated with each specified research activity, as well as the purposes desired by different categories of data users.

In our prior work (cited above) we also call special attention to advanced data-sharing models and emerging formal approaches to privacy. We note that there are many privacy methods, PETs, and data-sharing models that can provide stronger privacy protection than traditional de-identification techniques that are in wide use today.

Although PETs have advanced rapidly, none offer a plug-and-play, high-performance, and high-quality solution over all common data protection use-cases: There remain many open questions that require technology-specific research and development: e.g. adapting different PET methods to specific computations and data types; performance tuning; and ease of use for developers, analysts, and the end-users of analytics -- including the public and policy-making communities.)

A broader and perhaps more important question is how to match and align PETs as appropriate to enable different intended uses, and protect against different harms and threats. The broad classes of harms that different PETs aim to mitigate, and the specific threats that they are effective in protecting against, vary widely across technologies. For example, secure multiparty computation aims to prevent harm to institutional data holders from other adversarial data holders while enabling computation on their joint data, by requiring collective approval for any computation on the joint data. In contrast to SMC, differential privacy aims to protect against privacy harms to individual data subjects by limiting the inferences that can be made from the informational outputs -- regardless of the specific type of computation; while personal data stores and data coops aim to reduce individual privacy harm and individual and group economic harm by limiting the domain of human activity in which computations and inferences are used.[6] Further, most PETs aim to mitigate threats to individuals or specified institutional actors, and are not designed to provide wholesale protections to larger groups, such as marginalized communities.

Finally, we note that there is a paucity of research that examines the consequences of PET adoption or PET regulation on large interconnected social systems or society as a whole. It is well established that even privacy interventions that are provably effective and a Pareto-improvement in the short term can have a longer-term perverse effect on risk and global social welfare under many conditions (e.g. bounded rationality, externalities, adaptive system dynamics).[7] And more and more frequently, the

[6] For more details see Altman, Micah, Alexandra B. Wood, David O'Brien, and Urs Gasser. "Practical approaches to big data privacy over time." International Data Privacy Law 8(1):29-51. (2018).

[7] Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. "The economics of privacy." Journal of economic Literature 54, no. 2 (2016): 442-92.

`

introduction of PETs has broad consequences for data use and policy.[8] Thus in the policy arena, it is important to consider not only the technical properties of the PET but to also consider the potential consequences of rules and policies requiring PETs and the way that both these rules and the widespread use of a PET may affect actor information and incentives.[9]

Data releases should incorporate more advanced data sharing models, including formal privacy models, where possible, as such techniques can enable wider access and use of data while providing robust privacy protection.

Thank you for your consideration of these comments.

Respectfully,

Micah Altman, Research Scientist, Center for Research in Equitable and Open Scholarship, MIT Libraries

Aloni Cohen, Assistant Professor of Computer Science and Data Science, The University of Chicago

Salil Vadhan, Vicky Joseph Professor of Computer Science and Applied Mathematics, School of Engineering and Applied Sciences, Harvard University

[8] For a recent and particularly important case, see Boyd, Danah and Sarathy, Jayshree, Differential Perspectives: Epistemic Disconnects Surrounding the US Census Bureau’s Use of Differential Privacy (March 15, 2022). Harvard Data Science Review (Forthcoming) , Available at SSRN: https://ssrn.com/abstract=4077426.

[9] Altman, Micah, Alexandra Wood, and Effy Vayena. "A harm-reduction framework for algorithmic fairness." *IEEE Security & Privacy* 16, no. 3 (2018): 34-45.

`