The Future of Privacy is Data Equity

Bill Howe, H. V. Jagadish, and Julia Stoyanovich

AI applications that affect people's lives create challenges for equity related to the collection and use of data. The advance of data science and AI in equity-sensitive domains such as criminal justice, hiring, and medicine has shown that automated decisions based on biased data can operationalize, entrench, and legitimize new forms of discrimination. Over the last few months, due to the COVID-19 pandemic, the processes that lead to data equity issues have played out at an accelerated pace, in full public view. COVID-19 motivates including data equity as a first-class design consideration for any infrastructure (legal, social, or technical) that supports data-driven decision-making.

Defining data equity We consider four classes of data equity issues:

1. *Representation equity*: increasing the visibility of groups that have been historically underrepresented in the data record. For example, confirmed COVID-19 cases require testing, and there can be racial disparities in both the availability of testing and in the desire of individuals to be tested, leading to systematic biases in collected data.

2. *Feature equity*: facilitating linkage across datasets to ensure access to features that help expose and quantify inequity. For example, if attributes such as race and income are not recorded along with other data, it becomes hard to discover systematic biases that may exist, let alone correct for them.

3. *Access equity*: providing for access to data and data products across domains and levels of expertise. For example, a state's COVID-19 case data should be shared broadly rather than held tightly by the state's own government, since combination with other sources can lead to better verification and greater insights.

4. *Outcome equity*: monitoring and mitigating unintended consequences for any groups affected by a system after deployment, directly or indirectly. For example, contact tracing apps may facilitate stigma, harassment, or retribution in the case of positive diagnosis.

Data equity issues are pervasive but subtle, requiring *holistic consideration* of the sociotechnical systems that induce them (as opposed to narrowly focusing on the technical components and tasks alone), and of the contexts in which such systems operate.

Integrative Data Equity Systems To address data equity concerns we consider design principles for integrative data equity systems, socio-legal-technical data sharing and management systems that accept heterogeneous, sensitive, and potentially biased data from both public and private sources as input; facilitate manual and computational procedures for data transformation, cleaning, linking, and publishing; produce integrated (and bias-adjusted) data products (e.g., visualizations, integrated datasets, trained models) as output; and manage access to the data, data products, and provenance to protect privacy, facilitate accountability, and generally enforce compliance with relevant laws. The goal of these systems are to support applications in critical domains that share common bottlenecks: the difficulty and risk of *sharing sensitive data* due to both privacy and equity, the difficulty and risk of productively *integrating data* from disparate and heterogeneous sources due to cost of effort in addition to privacy and equity, and *limiting the potential for misuse* of data and models due to the loss of context during sharing and integration.

Such systems provide a computational infrastructure for data manipulation, analysis, and model training, but emphasize a data governance infrastructure, with a focus on engaging a broad range of stakeholders in system design, evaluation, and oversight. This kind of coupled system — combining governance with computational services — aims to facilitate value-driven evidence-based decision-making, enabling us to take collective and coordinated action, even in emergencies, while supporting accountability and promoting trust.

Equity in emergency situations In his testimony before the US Senate on June 30, 2020 on the topic of COVID-19 vaccines, Dr. Fauci noted: “It is a reality: a lack of trust of authority, a lack of trust in government, and a concern about vaccines in general.” This lack of trust is a consequence of a long history of unfair treatment of minority and disadvantaged communities, and of lack of representation from these communities in decision-making. And it is damaging to our collective ability to contain the pandemic.

To help rebuild trust, we must adopt effective policy in the face of uncertain risks and of difficult risk-risk trade-offs. In the context of data-driven decision making, this requirement means providing access to information about tests, outcomes, and contacts of those infected to manage disease spread. As an example, in Seattle, there is evidence that regulatory hurdles around IRB and repurposing of capabilities prevented rapid response.[1]

In an emergency situation, we need approved procedures for taking shortcuts and managing risk and liability with respect to releasing data, making inferences, and defining policy without broad discussion with affected populations. If we are slow to act, we risk reinforcing a harmful status quo; if we are too fast, we risk releasing premature and therefore misleading information. These risks are weighed against the potential benefit of answering questions in days rather than months. We argue that a data equity infrastructure must include allowances for emergency situations that still keep accountability intact: The goal is to provide a balance between extreme risk aversion (resulting in lost opportunities to save lives or reduce harms) and extreme risk tolerance (resulting in lost accountability for harms incurred during emergency situations.)

We envision a formal transfer of liability: If I release data, I am asserting that I have evidence that the benefits outweigh the harms, and I am agreeing to investigation after the fact. This kind of “fire alarm” exception is not typically included in the agreements and laws that govern data sharing; they are designed for the common case only. A simulation exercise, as is common in emergency response agencies, could facilitate the design and testing of the fire alarm model, and could be scaled down from global situations to smaller scopes: If we have a window of opportunity to change policy around, say, housing, then the value of releasing data could outweigh the potential risks to privacy, as long as accountability is in place.

More information This piece draws upon the insights developed during a 2-day NSF-funded workshop on Frameworks for Integrative Data Equity Systems (FIDES) and Foundations of Responsible Data Science (FORDS) that took place on March 25-26, 2020[2].

[1] `https://www.nytimes.com/2020/03/10/us/coronavirus-testing-delays.html`

[2] `https://midas.umich.edu/fides-workshop-program/`