

COVID-19 Brings Data Equity Challenges to the Fore

H. V. JAGADISH, University of Michigan

JULIA STOYANOVICH, New York University

BILL HOWE, University of Washington

The COVID-19 pandemic is compelling us to make crucial data-driven decisions quickly, bringing together diverse and unreliable sources of information without the usual quality control mechanisms we may employ. These decisions are consequential at multiple levels: They can inform local, state, and national government policy, be used to schedule access to physical resources such as elevators and workspaces within an organization, and inform contact tracing and quarantine actions for individuals. In all these cases, significant inequities are likely to arise and to be propagated and reinforced by data-driven decision systems. In this article, we propose a framework, called FIDES, for surfacing and reasoning about data equity in these systems.

CCS Concepts: • **Social and professional topics** → **Computing/technology policy**; • **Information systems** → **Data management systems**;

Additional Key Words and Phrases: Data equity, data ethics, responsible data science

ACM Reference format:

H. V. Jagadish, Julia Stoyanovich, and Bill Howe. 2021. COVID-19 Brings Data Equity Challenges to the Fore. *Digit. Gov. Res. Pract.* 2, 2, Article 24 (March 2021), 7 pages.

<https://doi.org/10.1145/3440889>

1 INTRODUCTION

We are awash in data today. Most of the data available to us was collected by some means we did not design, for some purpose that was not ours. Shareable data could be a secondary byproduct of a transaction; for example, getting tested for COVID-19. Data could be created by a person with the intent to share, such as a social media post or a response to a poll. It could be the intentional creation of an organization for administrative or operational reasons, for example, of a hospital to be able to bill insurance. Data created in all these different ways is subsequently aggregated and repurposed: used, for example, to study the prevalence of COVID-19 in a population or the attitudes about the government's handling of the pandemic; to inform a state's testing strategy or

This work was supported in part by National Science Foundation (NSF) Grants No. 1934464, 1934565, 1934405 1926250, 1740996, 1741022.

Authors' addresses: H. V. Jagadish, University of Michigan, 2260 Hayward Avenue, Ann Arbor, MI 48109-2121; email: jag@umich.edu; J. Stoyanovich, New York University, Tandon School of Engineering, 370 Jay Street, 11th Floor, Brooklyn, NY 11201; email: stoyanovich@nyu.edu; B. Howe, University of Washington, 5222 36th Ave NE Seattle, WA 98105; email: billhowe@uw.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

2639-0175/2021/03-ART24

<https://doi.org/10.1145/3440889>

the staging of reopening activities; to schedule access to a physical resource such as an elevator or a workspace within an organization; and to inform contact tracing and quarantine actions for individuals.

All datasets acquired opportunistically will reflect the system from which they were collected: They will often be noisy and systematically biased. Yet, a specialized, task-specific data collection effort is a luxury we can rarely afford, and so we must find ways to draw insight and make decisions based on repurposed datasets while safeguarding against social inequities. Put another way: Social inequities inherent in the data pose a risk, but so does deliberate inaction! In this article, we articulate a common vocabulary and provide a taxonomy for data equity issues and the corresponding risks that may arise in data-driven decision systems (Section 2). We then consider these issues as motivating requirements for a Framework for Integrative Data Equity Systems called FIDES that is currently the focus of our collaborative research and that is designed to surface, negotiate, and mitigate the risks due to inequity and to support accountability (Section 3). We conclude with a reflection on the long road ahead towards rebuilding the public's trust in the government's ability to develop and manage data-driven decision systems (Section 4).

Our discussion draws upon the insights developed during a two-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, 2020.¹ Planning for the workshop was scheduled to take place in New York City in March 2020 but was moved online in response to COVID-19. Our collaborative work on these topics started much before the onset of the pandemic. Yet, as is abundantly clear, the prolonged crisis is bringing data equity into ever-sharper focus.

2 DIMENSIONS OF DATA EQUITY

There is concern about fairness today whenever data-driven systems are used. It is no longer believed that data are impartial and neutral. Nevertheless, the scope of fairness considered is often narrow. Our goal is to address data equity broadly defined. Therefore, we consider the following four distinct classes of data equity issues:

- (1) *Representation equity*: increasing the visibility of underrepresented groups that have been historically suppressed 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, then it becomes hard to discover systematic biases that may exist, let alone correct for them.
- (3) *Access equity*: providing for equitable and participatory 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.

We consider data equity in the context of automated decision systems while recognizing a broader literature around the role of administrative systems in creating and reinforcing discrimination. Spade argues that administrative systems facilitate state violence encoded in laws, policies, and schemes that arrange and define people

¹<https://midas.umich.edu/fides-workshop-program/>.

by categories of indigeneity, race, gender, ability, and national origin [Spade 2015]. Hoffman considered how these effects are amplified through data technologies and their purveyors [Hoffmann 2020]. Decision systems, regardless of consideration of equity, mechanize existing structures, such that any effort to define and address data equity issues are at risk of manifesting as technological “happy talk” [Benjamin 2019], which involves a willingness to acknowledge and even revel in cultural difference without seriously challenging ongoing structural inequality. To combat these outcomes, we emphasize a primacy of equity in the design: The goal is not to automate and correct for equity, but design systems that exist to further anti-oppression. For example, a machine learning system to help submit insurance claims to maximize payment is designed to counteract the discrimination effected by corporate models to minimize payments.

Even in the best of circumstances, underlying structural inequities in access to health care, employment, and housing exhibit themselves in the data record and are propagated through decision systems, automated or otherwise, to become reinforced by policy. Under the current regime of haste, we not only reinforce existing structural inequities, but also risk introducing and normalizing new inequities, such as those arising if contact tracing apps and certificates of immunity become the norm.

That social inequity is reinforced and amplified by data-intensive systems is not new. We know from other domains that advances in data science and AI can be undermined by similar problems: Automated decisions based on biased data can operationalize, entrench, and legitimize new forms of discrimination. For example, a defendant’s immediate social network may reveal many convictions, but that information must be interpreted through the lens of socioeconomic conditions and prior structural discrimination in the criminal justice system before concluding that an individual is at a higher risk of recidivism or bail violation. Similarly, standardized test scores are sufficiently impacted by preparation courses that the score itself says more about socioeconomic conditions than an individual’s academic potential. Over the past 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.

3 A FRAMEWORK FOR INTEGRATIVE DATA EQUITY SYSTEMS

To address data equity concerns discussed in Section 2, we consider design principles for FIDES, a socio-legal-technical data sharing and management system that is the subject of our ongoing joint work. The system would be deployed and used in contexts similar to an open data repository: participating institutions and individuals would publish data through FIDES to facilitate broader use. Unlike typical open data repositories, FIDES aims to accept heterogeneous, sensitive, and potentially biased data from both public and private sources as input; facilitates manual and computational procedures for data transformation, cleaning, linking, and publishing; produces integrated (and bias-adjusted) data products (e.g., visualizations, integrated datasets, trained models) as output; and manages access to the data, data products, and provenance to protect privacy, facilitate accountability, and generally enforce compliance with relevant laws. The goal of FIDES is 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.

FIDES provides a computational infrastructure for data manipulation, analysis, and model training, but emphasizes a socio-technical data governance infrastructure, with a focus on engaging a broad range of stakeholders in system design, evaluation, and oversight. This kind of system—coupling 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.

3.1 Representation Equity

To support *representation equity*, FIDES implements methods for adjusting bias in source data, producing new semi-synthetic datasets that offer better statistical properties, while preserving other relevant properties [Asudeh

et al. 2019b; Lin et al. 2020; Rodriguez et al. 2018; Salimi et al. 2019; Stoyanovich et al. 2020]. The first step is to find groups, possibly intersectional sub-groups, that are inadequately represented in the data. For example, if many essential workers are on hourly wages without health insurance, then they may be under-represented in insurance data. To the extent that the exposure patterns for such workers are different than for others in the population, this under-representation can lead to bad decisions in downstream analysis. For specific questions (e.g., “How many food service workers are potentially exposed to COVID-19?”), the appropriate corrections can be applied manually by epidemiologists, statisticians, and economists working on the study. But the same bias can affect many data products across many studies; registering and versioning bias-corrected datasets in multiple contexts warrants computational support. FIDES provides methods to identify this under-representation, to address it by statistical means where possible, and manage the use and reuse of the results.

3.2 Feature Equity

To support *feature equity*, FIDES implements dataset search and integration services. If representation equity is about adding missing rows, feature equity is about adding missing attributes. For example, Louisiana was the first, and for a time, the only state to release demographic information in their COVID-19 case data. To understand racial disparities, the choices were to study Louisiana only (limiting relevance), extrapolate to other states and locales (which would be unsound in general), or to make inferences about demographics based on other datasets. The latter case is challenging in practice, but is amenable to computational support [Grechkin et al. 2018]. Given a dataset X for which demographic information is absent, the process is to (1) find another dataset Y sampled from the same population (e.g., patients) that possess demographic information, and (2) infer the demographic information for X using the distribution of Y . This inference may be simple or complex, but both the search step and the inference step can be system-assisted.

In addition to the computational capabilities, a data governance infrastructure helps navigate the continuum between open data (available for any use) vs. closed data (available for no use whatsoever) with the help of semi-synthetic data generation techniques to preserve privacy, protect competitive advantage, and correct for bias [Young et al. 2019]. These techniques were first developed in the context of urban transportation and mobility—a domain that has many similarities with location-based management of public health data, because of the spatio-temporal nature of the data and because data linkage is both necessary and difficult.

3.3 Access Equity

To support *access equity*, FIDES includes automatically constructed *nutritional labels* that provide succinct information about the fitness for use of a dataset for a particular purpose or task [Stoyanovich and Howe 2019; Sun et al. 2019; Yang et al. 2018]. These labels help triage fitness for use to ensure data is released when it is necessary and beneficial, breaking a dichotomy between restricting access to only privileged institutions or individuals vs. opening data for broad public use where risks of misinterpretation and privacy violations can harm specific groups. For example, Chowkwanyun et al. describe how releasing racial patterns in COVID-19 cases can give rise to incorrect biological, behavioral, or geographical explanations for racial health disparities [Chowkwanyun and Reed Jr 2020].

3.4 Outcome Equity

To support *outcome equity*, FIDES develops methods to increase fairness and diversity, both in the data, but also in result presentation [Asudeh et al. 2019a; Schelter et al. 2020; Stoyanovich et al. 2018; Yang et al. 2019, 2020; Yang and Stoyanovich 2017]. For example, Marion County in Ohio had the highest per capita COVID infection rate in the country at one point in Summer 2020. But the numbers were inflated by an outbreak at a prison, a severe problem nationwide that has more to do with systemic policies of mass incarceration than any particular mitigation strategy applied in Ohio specifically. Outcome equity issues imply a need to ensure accountability for decisions

related to release and use of data, especially in emergency situations where rules around data access and use may be relaxed to meet urgent needs. FIDES provides a framework to reason about risks, and to instrument the decision-making process via provenance and auditability, supporting a post hoc review of where the data was used and for what purpose, with accountability associated with the decision to break the steady-state rules. We discuss these scenarios in Section 3.5.

3.5 Accountability in Emergency Situations

The COVID-19 pandemic is also making acute the imperative to define and enact effective policy in the face of uncertain risks and of difficult risk-risk tradeoffs. In the context of data-driven decision making, this 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.²

The FIDES framework, introduced above, can help ensure data equity, but also requires extensive consultation with stakeholders and careful deliberation. These processes take time, which is not available in an emergency. But short-cutting these processes can introduce major risks to equity and could even expose decision-makers to liability. In an emergency situation, we need approved procedures for taking shortcuts and managing risk and liability. We must be able to break the glass and pull the fire alarm, but we accept some responsibility for the risk of a false alarm. In the current COVID-19 emergency, there is pressure to relax rules to make fast progress (funding approval, vendor selection, data sharing, public announcements). A process to move into an emergency condition with relaxed rules should be designed so as not to remove all accountability. The damage from making mistakes (e.g., wasted resources or privacy violations) must still be mitigated. For example, if private data is leaked in the context of an emergency, the responsible party must be required to show that emergency conditions were in effect and that the benefit of violation outweighed the harm. 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).

In our context, data sharing is the motivating example: The risk of releasing premature and therefore misleading information, and the risk of violating privacy are weighed against the potential benefit of answering questions in days rather than months. 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.

4 REBUILDING TRUST

In his testimony before the U.S. 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.” He added that there need to be “boots on the ground,” especially near minority communities that, he said, “have not always been treated fairly by the government.” 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. It is as damaging to our collective ability to contain the pandemic. And it reaches far beyond vaccines and exhibits, for example, in individuals declining to share information with contact tracers.³

²<https://www.nytimes.com/2020/03/10/us/coronavirus-testing-delays.html>.

³<https://www.nytimes.com/2020/06/21/nyregion/nyc-contact-tracing.html>.

Fair treatment of minorities goes well beyond equity in data. Furthermore, ensuring data equity today will not immediately build the trust that is missing; instead, it will help avoid a lack of trust down the road. Nonetheless, data equity is an important dimension of the responsible design, development, and use of automated decision systems, and of building trust, and so it is increasingly important as our society becomes ever more data-driven.

With Big Data and AI, where there once was unalloyed enthusiasm for what could be accomplished, there is today a trepidation about the harms that could result. This fear leads to a lack of trust, which manifests itself particularly in times of emergency. In this article, we discussed a framework through which to ensure data equity and to help build trust through accountability.

Thinking through the issues of equity and accountability is not easy. As part of our work, we have devoted considerable effort to educating current and future data scientists [Khan and Stoyanovich 2020; Stoyanovich and Lewis 2019] so they have a framework with which to assess the societal impacts of their actions. There is also an urgent need to educate policymakers, auditors, regulators, and members of the public about the opportunities afforded to us by the responsible use of data and technology—opportunities that we simply cannot afford to miss during this COVID-19 crisis.

REFERENCES

- Abolfazl Asudeh, H. V. Jagadish, Julia Stoyanovich, and Gautam Das. 2019a. Designing fair ranking schemes. In *Proceedings of the International Conference on Management of Data (SIGMOD'19)*, Peter A. Boncz, Stefan Manegold, Anastasia Ailamaki, Amol Deshpande, and Tim Kraska (Eds.). ACM, 1259–1276. DOI: <https://doi.org/10.1145/3299869.3300079>
- Abolfazl Asudeh, Zhongjun Jin, and H. V. Jagadish. 2019b. Assessing and remedying coverage for a given dataset. In *Proceedings of the 35th IEEE International Conference on Data Engineering (ICDE'19)*. IEEE, 554–565. DOI: <https://doi.org/10.1109/ICDE.2019.00056>
- Ruha Benjamin. 2019. Race after technology: Abolitionist tools for the New Jim Code. *Soc. Forces* 98, 4 (12 2019), 1–3. DOI: <https://doi.org/10.1093/sf/soz162>
- Merlin Chowkwanyun and Adolph L. Reed Jr. 2020. Racial health disparities and Covid-19—Caution and context. *New Eng. J. Med.* 383 (16 July 2020), 201–203.
- Maxim Grechkin, Hoifung Poon, and Bill Howe. 2018. EZLearn: Exploiting organic supervision in automated data annotation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*. International Joint Conferences on Artificial Intelligence Organization, 4085–4091. DOI: <https://doi.org/10.24963/ijcai.2018/568>
- Anna Lauren Hoffmann. 2020. Terms of inclusion: Data, discourse, violence. *New Media Soc.* DOI: <https://doi.org/10.1177/1461444820958725>
- Falaah Arif Khan and Julia Stoyanovich. 2020. Mirror, mirror. *Data, Respons. Comic Series* 1 (2020). Retrieved from <https://dataresponsibly.github.io/comics/>.
- Yin Lin, Yifan Guan, Abolfazl Asudeh, and H. V. Jagadish. 2020. Identifying insufficient data coverage in databases with multiple relations. *Proc. VLDB Endow.* 13, 11 (2020), 2229–2242. Retrieved from <http://www.vldb.org/pvldb/vol13/p2229-lin.pdf>.
- Luke Rodriguez, Babak Salimi, Haoyue Ping, Julia Stoyanovich, and Bill Howe. 2018. MobilityMirror: Bias-adjusted transportation datasets. In *Proceedings of the 1st Workshop on Big Social Data and Urban Computing (BiDU@VLDB'18) (Communications in Computer and Information Science)*, Jonice Oliveira, Claudio M. de Farias, Esther Pacitti, and Giancarlo Fortino (Eds.), Vol. 926. Springer, 18–39. DOI: https://doi.org/10.1007/978-3-030-11238-7_2
- Babak Salimi, Luke Rodriguez, Bill Howe, and Dan Suciu. 2019. Interventional fairness: Causal database repair for algorithmic fairness. In *Proceedings of the International Conference on Management of Data (SIGMOD'19)*, Peter A. Boncz, Stefan Manegold, Anastasia Ailamaki, Amol Deshpande, and Tim Kraska (Eds.). ACM, 793–810. DOI: <https://doi.org/10.1145/3299869.3319901>
- Sebastian Schelter, Yuxuan He, Jatin Khilnani, and Julia Stoyanovich. 2020. FairPrep: Promoting data to a first-class citizen in studies on fairness-enhancing interventions. In *Proceedings of the 23rd International Conference on Extending Database Technology (EDBT'20)*, Angela Bonifati, Yongluan Zhou, Marcos Antonio Vaz Salles, Alexander Böhm, Dan Olteanu, George H. L. Fletcher, Arijit Khan, and Bin Yang (Eds.). OpenProceedings.org, 395–398. DOI: <https://doi.org/10.5441/002/edbt.2020.41>
- Dean Spade. 2015. *Normal Life: Administrative Violence, Critical Trans Politics, and the Limits of Law*. Duke University Press. Retrieved from <http://www.jstor.org/stable/j.ctv123x7qx>.
- Julia Stoyanovich and Bill Howe. 2019. Nutritional labels for data and models. *IEEE Data Eng. Bull.* 42, 3 (2019), 13–23. Retrieved from <http://sites.computer.org/debull/A19sept/p13.pdf>.
- Julia Stoyanovich, Bill Howe, and H. V. Jagadish. 2020. Responsible data management. *PVLDB* 13, 12 (2020), 3474–3489. DOI: <https://doi.org/10.14778/3415478.3415570>
- Julia Stoyanovich and Armanda Lewis. 2019. Teaching responsible data science: Charting new pedagogical territory. *CoRR* abs/1912.10564 (2019).

- Julia Stoyanovich, Ke Yang, and H. V. Jagadish. 2018. Online set selection with fairness and diversity constraints. In *Proceedings of the 21st International Conference on Extending Database Technology (EDBT'18)*, Michael H. Böhlen, Reinhard Pichler, Norman May, Erhard Rahm, Shan-Hung Wu, and Katja Hose (Eds.). OpenProceedings.org, 241–252. DOI : <https://doi.org/10.5441/002/edbt.2018.22>
- Chenkai Sun, Abolfazl Asudeh, H. V. Jagadish, Bill Howe, and Julia Stoyanovich. 2019. MithraLabel: Flexible dataset nutritional labels for responsible data science. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM'19)*, Wenwu Zhu, Dacheng Tao, Xueqi Cheng, Peng Cui, Elke A. Rundensteiner, David Carmel, Qi He, and Jeffrey Xu Yu (Eds.). ACM, 2893–2896. DOI : <https://doi.org/10.1145/3357384.3357853>
- Ke Yang, Vasilis Gkatzelis, and Julia Stoyanovich. 2019. Balanced ranking with diversity constraints. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*, Sarit Kraus (Ed.). ijcai.org, 6035–6042. DOI : <https://doi.org/10.24963/ijcai.2019/836>
- Ke Yang, Joshua R. Loftus, and Julia Stoyanovich. 2020. Causal intersectionality for fair ranking. *CoRR* abs/2006.08688 (2020).
- Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. ACM, 22:1–22:6. DOI : <https://doi.org/10.1145/3085504.3085526>
- Ke Yang, Julia Stoyanovich, Abolfazl Asudeh, Bill Howe, H. V. Jagadish, and Gerome Miklau. 2018. A nutritional label for rankings. In *Proceedings of the International Conference on Management of Data (SIGMOD'18)*, Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein (Eds.). ACM, 1773–1776. DOI : <https://doi.org/10.1145/3183713.3193568>
- Meg Young, Luke Rodriguez, Emily Keller, Feiyang Sun, Boyang Sa, Jan Whittington, and Bill Howe. 2019. Beyond open vs. closed: Balancing individual privacy and public accountability in data sharing. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT'19)*. ACM, 191–200. DOI : <https://doi.org/10.1145/3287560.3287577>

Received September 2020; revised November 2020; accepted December 2020