

MAY 2023

Considerations *for* Public Algorithmic Registries

A REPORT TO THE PITTSBURGH TASK FORCE
ON PUBLIC ALGORITHMS

Carl DiSalvo, Ph.D

Georgia Institute of Technology
College of Computing

Rosta Farzan, Ph.D

University of Pittsburgh
School of Computing and Information

Robert Gradeck


University of Pittsburgh
Center for Social and Urban Research



University of
Pittsburgh

Institute for
Cyber Law, Policy, and Security

Contents



Introduction 3

Developing a Reporting Framework 3

General information about algorithms and the process of reporting them 4

Registry of individual algorithms 5

Implementing agency context 5

Developer of the algorithm 6

Technical details 6

The model 6

Data used by the model 7

Assessing impacts of an algorithm 8

Managing reporting burdens 9

Limitations 10

Sources 10

Introduction

The goal of this work was to identify the kinds of information that could be captured about algorithms in use by public-sector agencies. In this project, the University of Pittsburgh Institute for Cyber Law, Policy, and Security (Pitt Cyber) wanted to identify information that can contribute to the design of a website supporting the Task Force on Public Algorithms. Pitt Cyber asked the authors to identify the kinds of information that could be shared about public algorithms with the community, and develop some initial design concepts/principles that can be included in a reporting website.



Although a growing number of communities are reporting information on public-sector algorithms that are in-use, a standard for what should be included in a reporting process has not yet emerged. This document can inform conversations within the Task Force and between the task force and public-sector agencies about the kinds of information that is needed to support accountability efforts. This document supports the task force recommendation #5 that a website be created to share information about public sector algorithms with the community.

Developing a Reporting Framework

Evaluating what information about public algorithms can be captured was a two-part process. The first step involved looking at existing efforts to share information about algorithms in use. We did find that a few communities, including New York City, Helsinki, Amsterdam, and Seattle had already started to share details about algorithms and surveillance technologies deployed within their communities. We captured the different elements of these websites in our assessment.



The second stage of this analysis included a review of the rapidly growing body of literature that recommends specific information about algorithms important to supporting a community accountability framework. These suggestions were shared in reports published by various research groups and in articles published in academic journals. Sources for all websites and reports reviewed appear at the end of this document.

The list below includes suggestions on the kinds of information about algorithms that can be considered by local governments for inclusion in any reporting system. We grouped the information that local governments should consider into several different categories, including information about the governance and reporting process, information about individual algorithms, background on the developer and agencies using the tools, and information related to impact assessments.

General information about algorithms and the process of reporting them

This first category of information that might appear on a website includes general information about algorithms, including what they are, and how they're used. Several reports highlighted the need to break down the opacity around algorithms in order to engage the public in a conversation about algorithms. The website can also establish standardized definitions of algorithms and key terms. Community members can also learn more about public participation processes related to this initiative.



- Educational materials to build algorithmic literacy for both public and public-sector audiences.
- Description of the accountability and governance process.
- “Official” definition of what is an algorithm for the standpoint of this accountability process, and definitions of key terms.
- Copy of relevant legislation, executive order, or policies related to public algorithms.
- Description of the accountability/reporting/assessment process, including thresholds for reporting more-detailed information about algorithms.
- Describe opportunities for public participation, including public meetings, public comments, public events, and processes to request additional information.
- Processes that allow people to challenge decisions, fix mistakes, and correct incorrect information.
- Archive of meeting notes/recordings, presentations, and other artifacts from prior engagement events and processes.
- Provisions for research access.

Registry of individual algorithms

A registry capturing details about individual algorithms can appear through the website. The information provided can differ if this effort adopts different reporting thresholds depending on the potential harms or impacts of the algorithm.

- A listing of the algorithms that have been included in or excluded from the registry.
- Tags along with other descriptions and classifications to aid in discoverability.
- A human readable narrative describing the algorithm, including goals, motivation for use, purpose, what triggers its use, anticipated benefits and anticipated outcomes.
- Description of how decisions would be made if the algorithm did not exist, or description of decision-making processes prior to adoption or deployment of the algorithm.
- Links to any publicly-accessible uses of an algorithmic tool.
- Date for when the algorithm was placed into service, revision history, and length of use.
- Description or listing of other communities that have implemented this algorithm.
- Ownership of the intellectual property related to each algorithm, and licensing information.
- Information about impact assessments or audits of the algorithm, including if they have been conducted, who conducted them, when they took place, which standards or frameworks may have been used, and any plans to revisit or update them



Implementing agency context

Details on the agency's implementation of algorithms can help provide important context about the agency, the people that use the algorithm, and the implementation environment. Materials in this section can also serve as important resources for people responsible for public reporting.


- Organization name and steward/contact person.
- Description of who uses or has access to the algorithm, along with general information about access controls.
- Information about training requirements, frequency of training, certification, curricula, and training materials or user manuals.
- Procurement process details including copies of development/maintenance contracts, solicitations, responses, and source of funding.
- Resources describing how public sector employees can get help and assistance with the reporting process.

- Checklists or other reporting guidance for public sector employees as to what algorithms are exempt from reporting requirements, and situations where a higher degree of information about algorithms is required.
- Information about protections for whistleblowers.
- Copies of reporting forms, guidance, checklists, or instructions.

Developer of the algorithm

Information about who developed the algorithm can be captured whether the algorithm was developed in-house, or by an outside party.

- Name, organization, and capabilities of the developer, including contact information.
- Qualifications of key developers.
- History of their work with public-sector algorithms.
- History of claims/lawsuits/actions that have been taken against the developer.
- Payments made to the developer for development and maintenance.




Technical details

The model

Technical details about the algorithm provide additional insight about how the algorithm performs. The reporting website can provide or adapt to standardized frameworks as they emerge. Details about models can be captured in “model cards,” which provide structure for capturing details about how algorithms function for both technical and non-technical audiences.

- Data science and statistical techniques used by the algorithm.
- In scope and out of scope use cases used to develop the model.
- How the algorithm is optimized to emphasize particular outcomes.
- Actions and Processes that have been implemented to mitigate harms.
- Details on human roles in the decision-making process.
- Copies of technical manuals or documentation.
- Costs for development and maintenance.
- Version history.
- Software license.
- Link to a working test version of the algorithm.



- Link to code (use caution if code could be re-engineered to game the system or violate data privacy protections).
- Applicability of laws, standards, and governance processes to the design and use of the algorithm (IRB, HIPAA, Privacy policies, etc).

Data used by the model

Details about the individual datasets used to develop and operate an algorithm will also be important for assessing performance. Frameworks for capturing contextual details about datasets used by algorithms are emerging, and one notable example cited in several of the reports we reviewed is the “Datasheets for Datasets” framework.



- **Motivation** for the dataset, including the purpose, owner, creator, steward, and funding organization.
- **Composition** of the dataset, which can include what the data represents, what is included in the data, what is missing, relationships with other data, errors, harms, biases, sensitive information including PII, classification systems, and any laws, regulations, or processes that govern use of the data.
- **Collection processes** used to assemble the data, including workflows and procedures, frequency, use of technology, methodologies, and ethical considerations (including informed consent, governance, protections for subjects of human research, and whether compensation was provided, along with working conditions of the data labor force).
- **Transformations** are often used to prepare data for use with models, and the framework can capture who performed the labor, where the work was conducted, procedures, techniques, and technologies, and processes that were used to prepare data for use in the model, label, tag, and categorize data, and ensure quality output.
- **Uses of the data** can also be captured by the reporting website, and can inform the community as to how specific datasets are used to train models, make predictions, and inform decisions using other techniques. The reporting website can also capture limitations, warn others of harmful or inappropriate uses of the data, and offer suggestions as to how else the data can be used.
- The Datasheets for Datasets framework also covers **data sharing**, and can document if there were any licenses, permissions or restrictions applied to data used by the algorithm. This section can also include links to entries in relevant data catalogs and data repositories.
- Other policies and plans related to **data management and maintenance** that impact the data used in algorithmic systems may be important to capture and include on a reporting website.

Assessing impacts of an algorithm

Algorithmic risk assessments can identify the risks that people may face by use of automated decision systems. Literature we reviewed suggests that these assessments focus on the harms that may be experienced by people who historically lack power in our society, and that these assessments adopt an intersectional framework.

Intersectional analyses look at factors that describe power imbalances across multiple dimensions (such as a person's race, age, gender, ethnicity, sexual orientation, disability, religion, or place of birth) at the same time rather than in isolation.

Several papers explore what frameworks for algorithmic accountability assessments can include. Government agencies, ideally in consultation with the public, will ultimately be the actors to define the information that an impact assessment will capture, and the techniques and frameworks used to assess and report on the impacts of automated decision systems as part of a larger system of governance and accountability. The types of questions impact assessments and case studies may address can include:

- Which people is the algorithm most-likely to impact, and what are impacts to society from using the algorithm?
- What intersectional outcomes are reported about algorithm performance?
- How likely are specific harms to occur to both individuals and society? What are the impacts of these harms?
- What are some of the unintended consequences or harms that could come from use of the algorithm?
- Does the use of an algorithm codify existing biases in our society, and/or create new ones?
- Who defined the problem being solved by the algorithm as worth solving?
- Why is an algorithm the appropriate method for solving this problem?
- Does the decision to deploy an algorithm reflect community values or needs?
- What trade-offs have been made in the design and deployment of the algorithm?
- Does the community view use of the algorithm as legitimate?



Managing reporting burdens

Capturing information about every algorithm in use by public sector agencies may pose a significant reporting burden. Being aware of these burdens and designing compliance processes can lead to a greater likelihood that a reporting website will fulfill its role of sharing information about algorithms with the community. To minimize this burden, New York City developed a reporting framework as part of its compliance process related to [Executive Order 50](#) requiring the creation of a [public report of algorithmic tools](#). The process is managed by the Algorithms Management and Policy Officer (AMPO) housed at the Mayor's Office. New York City administrators designed the compliance process to minimize reporting burdens. Key lessons from the initial New York City experience include:

- Designing a two tier reporting process, where a somewhat-limited amount of information is collected about individual algorithms can minimize staff burdens. As outlined in the Report of the Pittsburgh Task Force on Public Algorithms, transparency initiatives can capture additional information about algorithms that are identified as higher risk. They encouraged other communities to not be too-ambitious in an initial data collection process. In the future, they plan to build a higher degree of structure into the process.
- Providing training and education to staff about algorithms and the reporting standard was instrumental to their success. Some staff needed context on the executive order and benefitted from background on what algorithms are and their reporting responsibilities. The culture in departments they worked with ranged from “data savvy,” to those with lower levels of data maturity. Training and Q&A sessions were offered separately to each department, and the team managing the process made themselves available for office hours and 1-on-1 conversations over the four-month compliance process. This engagement was handled in a virtual environment due to the pandemic.
- The compliance process for algorithms was designed to closely resemble the compliance process for the City's open data program managed by the Mayor's Office of Data Analytics (MODA). The process for reporting algorithms also relies on a hub and spoke model where coordinators/liasons in each department report to AMPO. Other compliance processes they incorporated into the design of this effort included those related to data privacy and performance management. They encouraged other communities to build on existing compliance frameworks, which could include data catalogs, and data governance processes.
- New York City also will explore opportunities to use the procurement process (Task Force Recommendation 4) as an opportunity to collect data about algorithms.





Limitations

Building public trust in the use of algorithms requires both transparency and accountability. Holding agencies that violate government policies and the public trust accountable can be enabled by sharing information about algorithms, but accountability will not come through simply providing more detail about algorithmic decisionmaking systems. Accountability will require members of the community to use information describing proposed and deployed algorithms as they engage in public discourse and in the political process. It will also require public-sector agencies to share power with community members in making decisions about the use of algorithms. A 2021 report developed by Data & Society emphasized the importance of appropriate governance structures in order for accountability to be meaningful, and stressed the importance of defining a source of legitimacy for the accountability framework in the form of legislation or strong norms.

This report does not explore the kinds of website design features that can help move from information to accountability, though the authors developed a series of design sketches that are available as a [separate document](#). Case studies of algorithmic systems can be one way that governments can encode accountability processes and community values into the design of the reporting website and the structure of the governance process. ■

Sources

Most of the research to review reports and websites was conducted in 2021.

ACLU Washington. “Algorithmic Equity Toolkit” (website). The Critical Platform Studies Group, the Tech Fairness Coalition, and the ACLU of Washington, 2021. aclu-wa.org/AEKit.

Algorithmic Accountability Act of 2019, H.R. 2231, 116th Congr. (2019). congress.gov/bill/116th-congress/house-bill/2231.

Anderson, David, Joy Bonaguro, Miriam McKinney, Andrew Nicklin, and Jane Wiseman. “Ethics and Algorithms Toolkit (beta)” (website). Johns Hopkins University, Center for Government Excellence, 2018. ethicstoolkit.ai.

Benjamin, Ruha. *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity, 2019.

Caplan, Robyn, Joan Donovan, Lauren Hanson, and Jeanna Matthews. *Algorithmic Accountability: A Primer*. Prepared for the Data & Society Caucus, Tech Algorithm Briefing: How Algorithms Perpetuate Racial Bias and Inequality, April 18, 2018. datasociety.net/wp-content/uploads/2018/04/Data_Society_Algorithmic_Accountability_Primer_FINAL-4.pdf.

City of Amsterdam Algorithm Register (website). City of Amsterdam, the Netherlands. algoritmeregister.amsterdam.nl/en/ai-register.

City of Helsinki AI Register (website). City of Helsinki, Finland. ai.hel.fi/en/ai-register.

City of Seattle. “About Surveillance” (website). City of Seattle Information Technology. seattle.gov/tech/initiatives/privacy/surveillance-technologies/about-surveillance-

City of Seattle. *City of Seattle Master List of Surveillance Technologies*. Revised December 2019. seattle.gov/Documents/Departments/Tech/Privacy/12-2019%20Revised%20Master%20List%20of%20Surveillance%20Technologies.pdf.

“Consultation des Algorithmes publics” (website). Nantes, France. data.nantesmetropole.fr/pages/algorithmes_nantes_metropole.

D'Ignazio, Catherine, and Rahul Bhargava. “Approaches to Building Big Data Literacy.” Bloomberg Data for Good Exchange, New York, NY, 2015. kanarinka.com/wp-content/uploads/2015/07/Big_Data_Literacy.pdf.

Diakopoulos, Nicholas, Sorelle Friedler, Marcelo Arenas, Solon Barocas, Michael Hay, Bill Howe, H. V. Jagadish, Kris Unsworth, Arnaud Sahuguet, and Suresh Venkatasubramanian. *Explaining Decisions Made with AI*. Information Commissioner's Office (UK) and the Alan Turing Institute, 2022. ico.org.uk/media/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-artificial-intelligence-1-0.pdf.

European Commission. *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence*. Brussels, Belgium, April 2021. digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.

Finch, Kelsey and Omer Tene. “Smart Cities: Privacy, Transparency, and Community.” In *Cambridge Handbook of Consumer Privacy*, eds. Evan Selinger, Jules Polonetsky, and Omer Tene, 125–48. Cambridge, MA: Cambridge University Press, 2018. ssrn.com/abstract=3156014.

Foard, Alex. “Interview with Alex Foard, NYC Algorithms Management and Policy Officer and Annie Gurvis from the NYC Mayor's Office of Operations.” By Robert Gradeck. April 27, 2021.

Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé, and Kate Crawford. “Datasheets for Datasets.” arXiv, 2018. arxiv.org/abs/1803.09010.

“Introducing #NoMoreDataWeapons” (website). Data for Black Lives. February 2021. blog.d4bl.org/introducing-nomoredataweapons.

Kaminski, Margot E. and Gianclaudio Malgieri. “Algorithmic Impact Assessments under the GDPR: Producing Multi-layered Explanations.” University of Colorado Law Legal Studies Research Paper, no. 19–28 (September 2019). Also in *International Data Privacy Law* (forthcoming). ssrn.com/abstract=3456224 or dx.doi.org/10.2139/ssrn.3456224.

Kingsley, Christopher and Stefania Di Mauro-Nava. *First, Do No Harm: Ethical Guidelines for Applying Predictive Tools Within Human Services*. Metrolab Network: Washington, DC, 2017. metrolabnetwork.org/wp-content/uploads/2017/09/Ethical-Guidelines-for-Applying-Predictive-Tools-within-Human-Services_Sept-2017.pdf.

Koene, Ansgar, Chris Clifton, Yohko Hatada, Helena Webb, Menisha Patel, Caio Machado, Jack LaViolette, Rashida Richardson, and Dillon Reisman. *A Governance Framework for Algorithmic Accountability and Transparency*. European Parliamentary Research Service, April 2019. [europarl.europa.eu/RegData/etudes/STUD/2019/624262/EPRS_STU\(2019\)624262_EN.pdf](https://europarl.europa.eu/RegData/etudes/STUD/2019/624262/EPRS_STU(2019)624262_EN.pdf).

Levy, Karen, Kyla E. Chasalow, and Sarah Riley. “Algorithms and Decision-Making in the Public Sector.” *Annual Review of Law and Social Science* 17, no. 1 (2021): 309–34. annualreviews.org/doi/10.1146/annurev-lawsocsci-041221-023808.

Loi, Michele, in collaboration with Anna Mätzener, Angela Müller, and Matthias Spielkamp. *Automated Decision-Making Systems in the Public Sector: An Impact Assessment Tool for Public Authorities*. AlgorithmWatch, June 2021. algorithmwatch.org/en/wp-content/uploads/2021/09/2021_AW_Decision_Public_Sector_EN_v5.pdf.

Metcalf, Jacob, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. “Algorithmic Impact Assessments and Accountability: The Co-construction of Impacts.” FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (March 2021): 735–46. doi.org/10.1145/3442188.3445935.

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, Timnit Gebru. “Model Cards for Model Reporting.” Paper presented at the Conference on Fairness, Accountability, and Transparency, ACM, January 2019. arxiv.org/abs/1810.03993.

Model Cards. “About” (website). Google. modelcards.withgoogle.com/about.

Moss, Emanuel, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. *Assembling Accountability: Algorithmic Impact Assessment for the Public Interest*. Data & Society, June 2021. datasociety.net/wp-content/uploads/2021/06/Assembling-Accountability.pdf.

New York City Algorithms Management and Policy Officer. *Calendar Year 2020 Summary of Agency Compliance Reporting*. New York City Office of the Mayor, 2021. www1.nyc.gov/assets/oti/downloads/pdf/reports/ampo-agency-compliance-cy-2020.pdf.

Pittsburgh Task Force on Public Algorithms. *Report of the Pittsburgh Task Force on Public Algorithms*. University of Pittsburgh Institute for Cyber Law, Policy, and Security (Pitt Cyber), 2022. cyber.pitt.edu/sites/default/files/pittsburgh_task_force_on_public_algorithms_report.pdf.

Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." In the *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain*, 33–44. New York: Association for Computing Machinery, 2020. doi.org/10.1145/3351095.3372873.

Richardson, Rashida, ed. "Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force." AI Now Institute, December 4, 2019. ainowinstitute.org/ads-shadowreport-2019.html.

Richardson, Rashida. "Defining and Demystifying Automated Decision Systems." *Maryland Law Review* (forthcoming). ssrn.com/abstract=3811708.

Selbst, Andrew D. "Disparate Impact in Big Data Policing." *Georgia Law Review* 52, no. 109 (2017): 109–95. ssrn.com/abstract=2819182 or <http://dx.doi.org/10.2139/ssrn.2819182>.

Sieber, Renée and Ana Brandusescu. *Final Report–FAcCT CRAFT: Civic Empowerment in the Development and Deployment of AI Systems*. Critiquing and Rethinking Accountability, Fairness and Transparency Workshop at ACM's Fairness, Accountability and Transparency in Machine Learning, Montreal, Canada, June 30, 2021. aifortherestofus.ca/wp-content/uploads/2021/10/Civic-Empowerment-Workshop-Report-FAcCT_CRAFT.pdf.

Vincent, Nicholas, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. "Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies." In the Proceedings of the 2021 Association for Computing Machinery (ACM) Conference on Fairness, Accountability, and Transparency (FAcCT '21) (2021): 215–27. doi.org/10.1145/3442188.3445885.

Williams, Rebecca. *Whose Streets? Our Streets! (Tech Edition): 2020–21 "Smart City" Cautionary Trends & 10 Calls to Action to Protect and Promote Democracy*. Belfer Center for Science and International Affairs, Harvard Kennedy School, August 2021. belfercenter.org/sites/default/files/2021-08/WhoseStreets.pdf.

Wilson, Christo, Cong Yu, and Bendert Zevenbergen. "Principles for Accountable Algorithms and a Social Impact Statement for Algorithms" (website). FAT/ML. fatml.org/resources/principles-for-accountable-algorithms.

Young, Meg, Michael Katell, and P. M. Krafft. "Municipal Surveillance Regulation and Algorithmic Accountability." *Big Data & Society* 6, no. 2 (2019). doi.org/10.1177/2053951719868492.

We gratefully acknowledge the generous support of
The Heinz Endowments and the Hillman Foundation.



University of
Pittsburgh

Institute for
Cyber Law, Policy, and Security

4200 Fifth Avenue
2516 Cathedral of Learning
Pittsburgh, PA 15260

cyber.pitt.edu