A PRIMER

AI and Procurement

Mona Sloane
Rumman Chowdhury
John C. Havens
Tomo Lazovich
Luis C. Rincon Alba

Summer 2021

AI and Procurement Primer
New York University
Summer 2021
https://doi.org/10.17609/bxzf-df18

CONTENTS

INTRODUCTION

Artificial intelligence (AI) systems are increasingly deployed in the public sector. As these technologies can harm citizens and pose risk to society, existing public procurement processes and standards are in urgent need of revision and innovation. This issue is particularly pressing in the context of recession-induced budget constraints and increasing regulatory pressures.

The AI Procurement Roundtables Project brought together leading experts in the public sector, data science, civil society, policy, social science, and the law to generate a structured understanding of existing public procurement processes and identify how they can best mitigate risk and support community needs. These experts had three separate conversations focused on mapping data science solutions used by public institutions, algorithmic justice and responsible AI, and governance innovation and procurement in the context of AI. This primer is the authors' reflection on these conversations, as well as their reflections on the research the project team has undertaken on the topic of AI procurement in the United States. It does not speak for or represent the opinions of the roundtable participants.

Rather than prescribing abstract pathways to procurement innovation in order to account for AI risk, this primer sets out to equip individuals, teams, and organizations with the knowledge and tools they need to kick off procurement innovation as it is relevant to their field and circumstances. Therefore, the primer is structured around five elements that build on each other. First, it sets the scene by examining the histories and current issues related to procurement and AI. Second, it outlines six tension points that emerge in the context of procurement and AI, each of which are paired with a set of questions that can help address these tension points.

INTRODUCTION

Third, the primer identifies five narrative traps that can hinder equitable innovation in AI procurement, together with ways and strategies to avoid these narrative traps. Fourth, the primer proposes four calls for action as concrete steps that can be taken to create environments in which AI procurement innovation can happen. And fifth, the primer closes with a list of resources that individuals, teams, and organizations may find helpful for delving into the topic further, or putting their AI procurement innovation strategy into action.

We extend our sincere gratitude to the roundtable participants for bringing their expertise and critical voices to the table. We are also grateful for the funding that we have received for this project from the Northeast Big Data Innovation Hub.

Additional support was given by the Institute of Electrical and Electronics Engineers (IEEE) and the NYU Tandon School of Engineering, and the NYU Alliance for Public Interest Technology.

Defining Public Use Technology

In this primer, the notion of “public use technology” refers to technological systems used by public agencies. Public use technology includes AI systems. Public use technology is slightly different from “public interest technology” (PIT). PIT is an emerging field in which technologists, designers, strategists, researchers, and policymakers leverage digital technologies to create more sustainable and inclusive economic and governance systems. This does not necessarily need to happen within public agencies, but can be driven by other organizations and actors.

Defining Artificial Intelligence (AI)

There is much uncertainty around what constitutes an AI system or even what should be considered an algorithm in the first place. For example, a recent "algorithm" that was used at Stanford Medical School to apportion COVID vaccine doses was essentially a human-designed set of decision criteria, rather than a complex system trained on available data.[1] For the purposes of discussing issues in procurement, it is useful to think of AI more generally as "automated decision systems" (ADS).[2] These decision systems include "any systems, software, or process that use computation to aid or replace government decisions, judgments, and/or policy implementation that impact opportunities, access, liberties, rights, and/or safety."[3] The fact that these systems are increasingly embedded into or even replace government functions is crucial in understanding why they present significant issues for current procurement processes. AI systems are typically optimized for a particular goal, and how well that goal aligns with the ultimate function for which they are deployed can change significantly. They have potentials for improving the processes of public agencies and their interactions with citizens - for example, by utilizing chatbots to facilitate communication and information retrieval. They also bear risks - for example, with police facial recognition incorrectly identifying an individual as a criminal, leading to their improper arrest.

---

1 Kristian Lum, Rumman Chowdhury, "What is an "algorithm"? It depends on whom you ask," *MIT Technology Review*, February 26, 2021, https://www.technologyreview.com/2021/02/26/1020007/what-is-an-algorithm/

2 ADS include systems where humans make the final decision. For example, systems in which the ADS completes the data analysis, but a human then makes a decision based on that data analysis.

3 Rashida Richardson, "Defining and Demystifying Automated Decision Systems," *Maryland Law Review* 81 (March 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3811708

SETTING THE STAGE
The Challenges with AI Procurement

The COVID-19 pandemic has underlined how biases can manifest in many different aspects of public use technology. For example, federal COVID-19 funding allocation algorithms have favored high-income communities over low-income communities due to historical biases prevalent in the training data.[4]

AI solutions that can be implemented fast are typically provided by private companies. As more and more aspects of public service are infused with AI systems and other technologies provided by private companies, we see a growing network of privately owned infrastructure. As government entities outsource critical technological infrastructure (such as data storage, and cloud-based systems for data sharing and analysis) to private companies under the guise of modernizing public services, we see a trend towards losing control over critical infrastructure and decreasing accountability to the public that relies on it. Unlike private companies, which have a responsibility to serve their shareholders and to generate profit, public entities must consider their entire population when delivering a solution and are tasked with mitigating harms introduced by AI to the communities they serve. There are two problems that persist:

AI public use is different. The public use of an algorithmic decision-making system has different requirements than a private use product. While there is an expectation that a technology built for public use is able to address the needs of all citizens, this is not necessarily the expectation of privately built products; in fact, many products are targeted specifically for a particular audience and later generalized to a broader user group. This gap is rarely considered.

---

[4] Pragya Kakani et al., “Allocation of COVID-19 Relief Funding to Disproportionately Black Counties,” JAMA 320, no. 10 (August 2020): 1000-1003, https://doi.org/10.1001/jama.2020.14978

Public use has higher compliance standards. Public use technologies are subject to different compliance and legal criteria than most private use technologies. In New Orleans, the undisclosed use of Palantir's Gotham platform was argued to be in violation of Brady vs Maryland, a Supreme Court case that places a responsibility on prosecutors to disclose any information that could cast doubt on any evidence presented against the accused.[5] For Kentrell Hickerson, who received a prison sentence related to gang involvement, the use of Gotham was not disclosed to his lawyers, allowing him to petition for a new trial (although it was not granted). In other cases, the use of algorithmic decision-making systems has been in violation of due process law, in particular the use of allocative algorithms to determine Medicare and Medicaid benefits.

[5] Laura Hooper, et al., "Treatment of Brady v. Maryland Material in United States District and State Courts' Rules, Orders, and Policies." *Federal Judicial Center (*October 2004), https://www.fjc.gov/content/treatment-brady-v-maryland-material-united-states-district-and-state-courts-rules-orders-a-0

Defining Public Procurement

"Public procurement" is one of the most legislated and regulated fields of government, and yet it is not clearly defined through an agreed-upon framework, with government agencies often using the terms "purchasing," "contracting," or "acquisition," interchangeably.[6] Typically, public procurement *activities* are outlined in the procurement guidelines of the public agencies and organizations. The absence of a clear framework of public procurement can complicate the development accountability mechanisms that cut across organizations. For the purpose of this primer, we adopt the following definition of public procurement:

Public procurement is the designated legal authority to advise, plan, obtain, deliver, and evaluate a government's expenditures on goods and services that are used to fulfill stated objectives, obligations, and activities in pursuant of desired policy outcomes.[7]

---

6 Robert E. Lloyd, et al., "What is Public Procurement? Definitional Problems and Implications," *International Public Procurement Conference Proceedings*, vol.3, (2004), http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.512.4209&rep=rep1&type=pdf

7 Eric Prier, et al., "The Implications of a Muddled Definition of Public Procurement," *Journal of Public Procurement* 9, no. 3 and 4 (2009): 326-370, http://www.ippa.org/jopp/download/vol9/issue-3-4/Art2_Public_Procurement.pdf

The Procurement Stages

Broadly, there are six steps of the public procurement process:[8]

1. Defining a need: first, the government must identify exactly what needs they would like the particular product or service to satisfy.
2. Writing a Request for Proposals (RFP): at this stage, the government has solidified its needs and translates them into a document that "describes the problem, the ideal solution for that problem, and invites companies to submit proposals for how they could fill the need."
3. Bid solicitation: next, the "competitive bidding process," a process by which companies submit their proposals and budgets, occurs; these "bids" are sometimes viewable by the public, depending on the specific RFP.
4. Vendor selection: once the RFP's deadline has passed and bids are submitted, a selection is made; this is often decided by a "procurement officer" or other official within an agency who has acquisition power.
5. Contracting: after the vendor is selected, the government engages in contract negotiations and ultimately agrees on terms.
6. Execution: once the deal is finalized, the vendor executes the contract as agreed.

While this is the formal definition of public procurement, it needs to be added that this process is continually evolving and changing, often with a view for making procurement more iterative and agile. Furthermore, the initiation of an RFP does not always come from a public agency. Agencies may not clearly articulate their need and vendors or private consultants then enter the conversation and shape the framing of the need. Or vendors may even create the perception of a need in the first place, which then kicks off the procurement process. There are inherent risks of this approach, such as the lack of a systematic analysis of whether or not an AI system is the optimal solution to meet a need, but many public agencies are faced with a shortage of technical talent that could steer this process and therefore they can be prone to lean on outside help.

---

8 Kenneth Cunnan, "Procurement 101," Code for America, accessed April 23, 2021, https://web.archive.org/web/20210423224200/https://www.codeforamerica.org/how-tos/procurement-101

SETTING THE STAGE
Procurement in the Global Context

The acquisition of external tools or services from third-party vendors is a core component of government operation. The process of acquiring these products, also known as procurement, allows governments to outsource resource-intensive and expertise-specific development to businesses with which they then contract.

Government procurement procedures can vary widely by country, and in fact occur at many different scales. Procurement rules and policies on the federal level in the United States shape procurement on the international level via trade agreements.[9] Beginning with the 1979 Government Procurement Agreement (GPA), the U.S. has entered into a series of agreements with foreign governments through which U.S. providers can compete for foreign procurement contracts. The 1996 revision of the GPA led to the creation of the World Trade Organization (WTO), and the most recent version of GPA went into effect in 2014.[10]

In addition to the GPA, the U.S. rolls certain procurement requirements into its free trade agreements. As such, the President (and via Executive Order 12260,[11] the U.S. Trade Representative) has the authority to waive any domestic purchasing requirements that would be in conflict with the terms of the international agreement. Such agreements often include reciprocal opportunities between the countries who are parties to the treaty. There are a number of institutions, on both international and national levels, that are central to procurement, particularly with regards to digital services:

---

9 "Government Procurement," Office of the United States Trade Representative, accessed May 12, 2021, https://ustr.gov/issue-areas/government-procurement

10 Agreement on Government Procurement (as amended on 30 March 2012)," World Trade Organization, accessed May 12, 2021. https://www.wto.org/english/docs_e/legal_e/rev-gpr-94_01_e.htm

11 "Executive Order 12260 of December 31, 1980--Agreement on Government Procurement," National Archives Federal Register, accessed May 12, 2021, https://www.archives.gov/federal-register/codification/executive-order/12260.html

World Trade Organization Government Procurement Agreement (WTO GPA): The WTO GPA is an international agreement regarding government procurement to which some, but not all, members of the WTO have agreed.[12] There is no requirement for WTO members to join the GPA, and the membership currently consists of twenty countries.[13]

Office of Federal Procurement Policy (OFPP): Residing under the Office of Management and Budget (OMB), the OFPP is the U.S. government's primary source of directives and policies for procurement that are used across federal agencies. In particular, they aim to establish procedures that "promote economy, efficiency, and effectiveness" in government procurement (themes which are often used to measure procurement manager success).[14]

The United States Digital Services (USDS): The USDS[15] was established in 2014 as a way to incorporate expertise from the private technology sector into federal government digital services. In the wake of the failure of the rollout of healthcare.gov in 2013,[16] the U.S. identified a need to boost its internal digital expertise, in particular for services that would be used by the public. As part of this, one of USDS's explicit objectives is to "rethink how the government buys digital services," therefore bringing such expertise to bear on the procurement of technology as well. These efforts include modernizing procurement to better account for the needs of technological systems, as well as establishing the use of common tools across different federal offices.

---

12 "WTO Government Procurement Agreement," Office of the United States trade Representative, accessed May 12, 2021. https://ustr.gov/issue-areas/government-procurement/wto-government-procurement-agreement

13 There are many countries currently negotiating to join the GPA, including China, Russia, and the Kyrgyz Republic.The countries that are currently part of the agreement are Armenia, Australia, Canada, the European Union, Hong Kong China, Iceland, Israel, Japan, the Republic of Korea, Liechtenstein, the Republic of Moldova, Montenegro, the Netherlands with respect to Aruba, New Zealand, Norway, Singapore, Switzerland, Taiwan (Chinese Taipei), Ukraine, and the United States.

14 A list of the main procurement institutions and regulators at the U.S. Federal and State level can be accessed via the U.S. Office of the United States Trade Representative. "Additional Information on U.S. Procurement." Office of the United States Trade Representative, accessed May 12, 2021. https://ustr.gov/issue-areas/government-procurement/additional-information-on-US-Procurement

15 "Our Mission." United States Digital Service, May 20, 2021. https://www.usds.gov/mission.

16 Bleiberg, Joshua, and Darrell M. West. "A Look Back at Technical Issues with Healthcare.gov." Brookings. Brookings, July 29, 2016. https://www.brookings.edu/blog/techtank/2015/04/09/a-look-back-at-technical-issues-with-healthcare-gov/

The Technology Transformation Services (TTS): The TTS[17] houses a series of programs aimed at boosting the digital expertise of the federal government. The 18F program acts as a central digital consultancy service, offering ad hoc technology consultations to various government agencies as needed. The Presidential Innovation Fellows (PIF) program pairs skilled technologists with agency employees for a one year period to tackle a specific public interest technology problem. Additionally, TTS operates IT Modernization Centers of Excellence (CoE) as a team of experts aimed at modernizing agency technologies. The overall aim of all of these programs is both talent recruitment and expertise-building on digital technologies within the federal government.

17 "Technology Transformation Services," U.S. General Services Administration, last modified February 23, 2021, https://www.gsa.gov/about-us/organization/federal-acquisition-service/technology-transformation-services

SETTING THE STAGE
AI and Procurement in the United States Today

Today, U.S. government agencies at the federal, state, municipal, and local levels can all follow their own procurement procedures subject to any applicable laws or regulations.[18] This often leads to complex, heterogeneous, and cumbersome sets of guidelines for vendors to navigate.[19] Nonetheless, a typical series of steps occurs: the definition of a need, writing a request for proposals which translates these needs into an invitation to companies to submit proposals for filling that need, the bid solicitation, vendor selection, contracting, and the execution. These procurement procedures are centered around criteria that are important to the procuring managers, such as getting the most out of their budgets and ultimate usefulness of the products to the agency in question.

In the past decade, public procurement of information technology initially struggled to keep up with new developments while also satisfying these goals. For example, in December 2014, the U.S. Congress passed the Federal Information Technology Acquisition Reform Act (FITARA)[20] as an attempt to streamline government acquisition of information technology systems. This law directed government agencies to coordinate the purchase of commercial-off-the-shelf (COTS) software in order to make better use of the overall federal budget.[21] While this law led to OMB recommendations regarding centralization of software licensing across the federal government where possible, it did not address the tensions in the procurement process specific to AI systems. Additionally, in August 2014, the federal government created the United States Digital Service (USDS) as an attempt to recruit technologists to improve government systems that were being used by the public. While this effort has led to more technological expertise within different agencies, it did not directly address the gaps in understanding of AI systems.

---

18 For a comprehensive overview of public procurement in the U.S., please see https://content.next.westlaw.com/Document/I2ef1290d1ed511e38578f7ccc38dcbee/View/FullText.html?transitionType=Default&contextData=(sc.Default)&firstPage=true.

19 Kenneth Cunnan, “Procurement 101,” Code for America, accessed April 23, 2021, https://web.archive.org/web/20210423224200/https://www.codeforamerica.org/how-tos/procurement-101

20 Management and Oversight of Federal Information Technology,” Federal CIO Council, accessed April 21, 2021. https://management.cio.gov/

21 Anne E. Rung and Tony Scott, “Category Management Policy 16-1: Improving the Acquisition and Management of Common Information Technology: Software Licensing,” Office of Management and Budget, June 2, 2016, https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2016/m-16-12_1.pdf

More recently, there has been a significant push within the federal government to modernize its understanding of AI. In 2019 and 2020, a series of executive orders were issued, aimed at growing the government's talent pipeline for AI, establishing a shared ethics framework, and initiating an inventory of current AI use within governmental agencies.[22] Additionally, in January 2021, the Office of Science and Technology Policy created an AI Initiative to centralize the government's national AI strategy.[23] In March 2021, the COVID stimulus American Rescue Plan Act included one billion dollars of funding toward government IT modernization.[24] In May 2021, the White House launched AI.gov, a site specifically dedicated to keeping the general public apprised of the federal government's efforts in AI innovation.[25] The increasing pace of these efforts is indicative of the government's ever-increasing interest in AI technologies.

As these initiatives related to talent development and AI innovation have grown, so has government acquisition of AI systems. In the U.S., these systems have been deployed for everything from determination of Social Security benefits to immigration regulation.[26] Each of the procurement steps described above can often run into problems when applied to the acquisition of an AI system. Fundamentally, this is because the deployment of AI systems within government encodes judgments made by the vendor about issues that extend beyond the scope of the product itself. Design choices that vendors make about what datasets to use, how (and on which populations) to test their systems, what objective to optimize (usually a poor proxy for the actual task), and even the framing of the decision to be made by the algorithm all encode policy choices.

---

22 Jory Heckman, "Trump Executive Order Looks to Expand AI Talent Pipeline into Agencies," *Federal News Network*, December 3, 2020, https://federalnewsnetwork.com/artificial-intelligence/2020/12/trump-executive-order-looks-to-expand-ai-talent-pipeline-into-agencies/

23 GCN Staff, "White House Opens AI Initiative Office," GCN, January 14, 2021, https://gcn.com/articles/2021/01/14/ostp-ai-initiative-office.aspx.

24 Aaron Boyd, "Central IT Modernization Fund Gets $1B in COVID Stimulus," Nextgov, March 10, 2021, https://www.nextgov.com/cio-briefing/2021/03/central-it-modernization-fund-gets-1b-covid-stimulus/172587/.

25 "The Biden Administration Launches AI.gov Aimed at Broadening Access to Federal Artificial Intelligence Innovation Efforts, Encouraging Innovators of Tomorrow," The White House. May 5, 2021, https://www.whitehouse.gov/ostp/news-updates/2021/05/05/the-biden-administration-launches-ai-gov-aimed-at-broadening-access-to-federal-artificial-intelligence-innovation-efforts-encouraging-innovators-of-tomorrow/.

26 David Rubenstein, *Federal Procurement of Artificial Intelligence: Perils and Possibilities*. New York: The Great Democracy Initiative, 2020. Accessed May 12, 2021. https://greatdemocracyinitiative.org/wp-content/uploads/2020/12/Artificial-Intelligence-Report_121320-FINAL.pdf

The procurement of AI is treated with the same off-the-shelf purchasing philosophy as other IT systems, but this mindset leads to policy making by third party design, ultimately resulting in a lack of government accountability for these decisions.[27] On top of this, current processes by which AI systems are trained are known to be rife with bias issues, such that these technological biases can encode biases from the real world.[28]

Here, it is instructive to consider the fraud detection system used by the U.S. Department of Agriculture to detect fraud within their food stamp program against the backdrop of the steps in the procurement process. Starting from step one, the defined agency need was to detect fraud. The vendor made an inherent assumption in building their product, namely that automating the flagging of "anomalous" transaction patterns would identify fraudulent transactions, thus satisfying the government need. However, it ultimately became clear that in some cases, anomalous transactions actually corresponded to unique cultural or social behavior patterns rather than actual fraud. This is a potentially subtle point, but when procurement officials rely on an off-the-shelf solution for AI procurement (as is currently suggested in the U.S. government's own training materials for procurement officers), such intricacies are unlikely to be probed.[29]

---

27 Deirdre Mulligan, et al., "Procurement as Policy: Administrative Process for Machine Learning" *Berkeley Technology Law Journal* 34, no. 3 (February 2020): 773-852, https://doi.org/10.15779/Z38RN30793.

28 David Rubenstein, "Acquiring Ethical AI," *Florida Law Review* 73, (November 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3731106.

29 Artificial Intelligence in Federal Procurement," U.S. General Services Administration, accessed May 12, 2021, https://www.youtube.com/watch?v=XJsgbGk8BIw

AI in Public Use Gone Awry: Automated Fraud Detection

In order to illustrate the incompatibility of AI systems and the typical procurement process, let's consider a specific example.[30] In 2002, the U.S. Department of Agriculture sought to use an algorithmic system to detect fraud within their food stamp program through analysis of electronic transactions (this is the "defined need" as described above). With the system in use, the agency ultimately disqualified a set of grocery stores in Muslim East African communities from accepting the benefit, based on behaviors that were algorithmically flagged as anomalous. Specifically, "the suspicious transactions at issue included large purchases made minutes apart, transactions for even-dollar amounts...and instances in which a few households made several unusually large purchases within a single month." Ultimately, these supposedly fraudulent transactions were explainable by specific cultural practices of the customers in question, namely the fact that they tended to shop in groups and make large purchases of meat that was then frozen and eaten over the course of weeks. In short, a set of minority communities suffered significant harms due to decisions made by an algorithmic system that was fundamentally incorrect. More recent examples of fraud detection gone awry, both in the U.S. (Michigan's unemployment fraud detection system) and abroad (welfare fraud detection in the Netherlands), additionally illustrate the pitfalls that government deployment of AI can bring.

[30] Deirdre Mulligan, et al., "Procurement as Policy: Administrative Process for Machine Learning" *Berkeley Technology Law Journal* 34, no. 3 (February 2020): 773-852, https://doi.org/10.15779/Z38RN30793.

More generally, the decision to optimize for this task in an attempt to meet the stated need encodes assumptions about what is “typical” behavior, thus embodying policy decisions in the way that other government acquisitions do not. It is unclear whether the current competitive bidding process and vendor evaluation are structured in a way that makes it possible to detect this misalignment between agency and vendor goals. Additionally, current government contracting may not adequately capture the fact that ongoing evaluation of an AI system’s performance is often necessary. In practice, AI errors may be harder to detect due to the lack of organizational or technical transparency or to AI literacy within the agencies, and they may have cascading effects as AI systems interact with other systems within and beyond the agencies that deploy them. The U.S. government is attempting to address this by improving its own regulations and evaluation criteria,[31, 32] and modifications to procurement will be a key component in this process.

Addressing the limitations of public procurement for AI systems is no easy task, but certain changes in policy and mindset can be helpful in moving forward. In the remainder of this primer, we identify specific points of tension that procurement officers, vendors, and others involved in the process will likely encounter. Additionally, we lay out certain narrative traps that often arise when discussing AI procurement.

---

31 Asa Fitch, “White House Nears New Rules on Artificial Intelligence,” The Wall Street Journal, October 21, 2020. https://www.wsj.com/articles/white-house-nears-new-rules-on-artificial-intelligence-11603300366.

32 OMB. “Guidance for Regulation of Artificial Intelligence Applications,” November 17, 2020. https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf.

SIX TENSION POINTS

During the roundtable discussions, six distinct points of tension emerged. These points of tension serve as a way to focus on issues that can come up in the context of a particular technology. They also serve as a way to identify the most urgent areas of improvement within existing procurement processes and practices in order to address the societal impact of AI technologies.

SIX TENSION POINTS
DEFINITIONS

The AI space in general is challenged by a panoply of terms that remain undefined. The first set of definitional challenges pertains to technologies and procedures: there is no agreed-upon notion of AI, or even algorithm.[33] This can hinder, for example, the cataloging of existing socio-technical systems within government, but also the development of procurement innovation that is specific to these technologies. Agencies and local governments sometimes define these technologies for themselves, for example in registries [34] or compliance reporting,[35] or as part of new regulation,[36] but there is no cross-agency coordination. Similarly, there are no agreed upon definitions and procedures of AI impact and risk assessments or audits.

The second set of definitional challenges pertains to legal frameworks and principles, particularly justice. There are vastly different notions of what constitutes justice in the context of AI.[37] In order to arrive at a workable definition of justice that, indeed, *is* just, it is necessary to include those who are affected by AI *in*justice.[38]

The third set of definitional challenges pertains to metrics, and particularly metrics of success, both for the AI system, as well as the process through which it was procured. There can be an absence of appropriate metrics of success, or the metrics of success can be conflicting. For example, financial fraud detection models measure success based on identifying anomalous behavior (e.g., out of the norm for an individual), which may mean (but is not limited to) fraudulent behavior. Anyone who has had their credit card frozen while on vacation likely appreciates this distinction. These systems are not only used by banks, public agencies may also deploy them to detect benefit fraud. While the technical definition of success may evolve around detecting as much fraud as possible, the practical definition of success may be to only detect those cases that are the most likely to constitute fraud.

---

33 Kristian Lum, Rumman Chowdhury, “What is an “algorithm”? It depends on whom you ask,” *MIT Technology Review*, February 26, 2021, https://www.technologyreview.com/2021/02/26/1020007/what-is-an-algorithm/

34 Khari Johnson, “Amsterdam and Helsinki Launch Algorithm Registries to Bring Transparency to Public Deployments of AI,” VentureBeat, September 28, 2020, https://venturebeat.com/2020/09/28/amsterdam-and-helsinki-launch-algorithm-registries-to-bring-transparency-to-public-deployments-of-ai/.

35 Jeff Thamkittikasem, “Implementing Executive Order 50 (2019): Summary of Agency Compliance Reporting,” NYC Algorithms Management and Policy Officer, 2020. https://www1.nyc.gov/assets/ampo/downloads/pdf/AMPO-CY-2020-Agency-Compliance-Reporting.pdf.

36 “Ordinance No. 7,592-N.S.,” Council of the City of Berkeley, accessed May 12, 2021, https://www.berkeleyside.org/wp-content/uploads/2018/03/2018-03-27-Item-02-Ordinance-7592.pdf

37 Alan Lundgard, 2020, “Measuring Justice in Machine Learning.” Paper presented at *Conference on Fairness, Accountability, and Transparency*, Barcelona, January 2020.https://doi.org/10.1145/3351095.3372838

38 Ginia Bellafante, “The Landlord Wants Facial Recognition in Its Rent-Stabilized Buildings. Why?” *The New York Times*, March 28, 2019, https://www.nytimes.com/2019/03/28/nyregion/rent-stabilized-buildings-facial-recognition.html

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE DEFINITIONS TENSION POINT:

How can procurement innovation serve as an incentive and platform to arrive at clear definitions of AI technologies and related procedures, such as risk assessment and audits, that work across agencies?

What are processes through which “justice” can be defined in an equitable way and then serve as a leading principle for AI procurement innovation?

How can metrics of success be defined adequately through AI procurement innovation?

SIX TENSION POINTS
PROCESS

The procurement process has been designed to prevent abuse, but as a result has solidified sets of procedures that tend to only allow large vendors to meet standards and compete for government contracts. Long-term contracts that result from these sets of procedures lock in government agencies over years and can create path dependencies.[39] These path dependencies can also spill over into the newly emerging space of algorithmic auditing, where large vendors which already have contracts with government agencies add on algorithmic auditing to their service portfolio, incentivising government agencies to contract these services through their existing vendors. Tension in the procurement process also occurs with regards to bottlenecks that can exist in the procurement of various services, for example when there are delays or distortions in one of the procurement stages (particularly bid solicitation, vendor selection, contracting, and execution). As the disproportionate impact of AI systems on individual citizens and communities becomes more evident, there also is a pressing need to define at what point risk and impact assessments should be required in the procurement process, which - importantly - should also include the execution stage.[40] As the procurement process is redefined and redesigned, points at which public participation expertise and reasoned deliberation can happen and be integrated need to be defined.

---

39 Ryan Johnston, “Illinois officials say tech contract holding up availability of nonbinary IDs,” *StateScoop*, June 8, 2021, https://statescoop.com/illinois-drivers-license-nonbinary-idemia/

40 It is possible that issues with AI systems only emerge at the execution stage, when a model is deployed in its intended context. AI acquisition processes that build on the idea that technical and design requirements can be fixed and locked at the development stage cannot account for this dynamic.

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE PROCESS TENSION POINT:

How can procurement be redesigned to abolish the systemic preference for established vendors (which can be large or small) and avoid path dependencies?

What role can algorithmic auditing play in the procurement process?[41]

How can *meaningful* public participation, such as participation that happens early and in a non-extractive way, be integrated into the procurement process?

---

[41] As mentioned earlier, there is no clearly defined definition of what precisely constitutes AI auditing, but four types of problematic machine behavior have been identified that are typically addressed in audits: discrimination, distortion, exploitation, and misjudgement. However, it has been suggested that algorithmic audits need a more holistic approach in order to understand and mitigate potential harms. Inioluwa Deborah Raji et al., 2020. "Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing," Paper presented at the *Conference on Fairness, Accountability, and Transparency*. Barcelona, January 2020. https://doi.org/10.1145/3351095.3372873; Jack Bandy, "Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits," *The Proceedings of the ACM (PACM) Human-Computer Interaction*, CSCW '21 (2021), https://arxiv.org/abs/2102.04256.; Mona Sloane, "The Algorithmic Auditing Trap," *OneZero*, March 17, 2021, https://onezero.medium.com/the-algorithmic-auditing-trap-9a6f2d4d461d.

SIX TENSION POINTS
INCENTIVES

The incentives that underpin both the procurement process in general, as well as the different organizations that come together through procurement, can be detrimental to creating AI impact assessment structures. Vendors are driven by capitalist incentives and are focused on profit-generation in the offering of their services to public entities. The primary responsibility most vendors have is toward their shareholders, not necessarily their clients, or the people who use their services and systems. This means the current procedural and cultural set-ups of procurement are propped up by incentive structures that are not conducive to measures that would protect the public or mitigate algorithmic harm, not least because these could potentially slow down the procurement process. At this point, government and vendor incentives actually are aligned: both are incentivised to solve a problem fast and contract it out, once it has been identified and budgets have been approved. Both actors are also incentivized to frame a technological solution as the most efficient solution to any given problem. Government employees who are tasked with managing the procurement process may also not be incentivized to change the procurement process in order to account for the potential harms of AI technologies. Their task is to procure as efficiently as possible, there is no organizational or career reward for changing existing procedures.

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE INCENTIVES TENSION POINT:

How can renewed procurement procedures create incentives for upstream impact assessments?

What are the changes needed in the procurement process to prioritize the public interest, not stakeholder interest?

How can government agencies create incentives for procurement professionals to change the procurement process in order to address the potential harms of AI technologies?

SIX TENSION POINTS
INSTITUTIONAL STRUCTURES

The institutional structure of government can challenge procurement innovation to account for the specifics of AI systems, particularly when it comes to the timeframe for solving problems and achieving such innovation. According to former U.S. Chief Data Scientist DJ Patil, government policies are generally constructed with a 10-year time frame. Policymakers seek to create impact within their (much shorter) tenure, often with a view for reelection. Industry is moving at an even faster pace. Policymakers and industry speeds converge when decision-makers are confronted with having to solve issues within their tenure, which does not necessarily align with the organizational priorities and timeframes of government. Another issue is that "government" is often treated as monolithic, excluding other forms of government (such as Indigenous governments) from the conversation, and therefore from efforts to innovate.[42] Furthermore, the ways in which procurement is currently organized can negate and foreclose infrastructural sovereignties for First Nations.[43]

---

42 Many important opportunities for learning are foreclosed by that. For example, during the COVID-19 pandemic, tribal nations have built on existing social and community infrastructures and networks, such as gathering places, to distribute shots, as well as on many years of experience of bringing medical care to remote communities, resulting in very high vaccination rates. Alex Brown, "In Hard-Hit Indian Country, Tribes Rapidly Roll Out Vaccines," *Pew Charitable Trust*, February 9, 2021, https://www.pewtrusts.org/en/research-and-analysis/blogs/stateline/2021/02/09/in-hard-hit-indian-country-tribes-rapidly-roll-out-vaccines

43 "Procurement in Indigenous Communities," First Nations National Building Officers Association, accessed May 12, 2021, https://www.fnnboa.ca/procurement-in-indigenous-communities

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE INSTITUTIONAL STRUCTURES TENSION POINT:

How can procurement innovation for AI serve as a gateway for synching up timeframes and speeds of policy-makers, government agencies, and industry?

How can the notion of “government” be expanded in order to amplify and support Indigenous sovereignty through procurement?

How can procurement be re-organized so that government and citizen and community more closely align?

SIX TENSION POINTS
TECHNOLOGY INFRASTRUCTURE

Procurement is the gateway for technology infrastructure implementation, and therefore has long term effects on cities, communities, and on agencies themselves. Due to the lack of capacity and/or resources to build technology in-house, procurement invites the building of large-scale technology infrastructure, including AI, through private vendors. This dynamic does not foster, but prevents transparency: often protected by tradelaw, private vendors are not obligated to open up the “black box” and share insight into their training data or their models. Relatedly, a promised outcome or a narrative often is treated as more important than the technological backend. These promised outcomes and narratives often are a point of convergence for agencies and the private sector, and they are treated with higher priority than the establishment of accountability in the procurement process, for example in the context of the climate emergency and “clean tech,” or technologies that are deployed for public health management in the context of the COVID-19 pandemic.

AI systems that are implemented in this way, and that can cause harm to communities, such as through over-surveillance and -policing, become infrastructural and therefore are unlikely to be taken down, even when harm is proven, for example in the context of smart sensors installed in street lights. More transparency and oversight in the procurement process (vs. attention to the promise of a technology) can prevent the implementation of AI infrastructures that can prove to be harmful by design. Agencies are in need of resources to build up literacies and capacities around the impacts of procuring AI systems, which includes knowledge sharing across agencies.

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE TECHNOLOGY INFRASTRUCTURE TENSION POINT:

What are the mechanisms through which procurement processes can mandate the opening up of the “black box?”

What steps need to be taken to ensure that innovation in the procurement process (vs. the outcome) puts emphasis on assessing the actual impacts of AI infrastructure?

How can AI procurement innovation foster the provision of resources for agencies to build up AI literacies and capacities for knowledge sharing across agencies?

SIX TENSION POINTS
LIABILITIES

Procurement innovation is impossible without considering legal implications, particularly pertaining to protecting the procuring institutions from liabilities. Public services are often held to a higher standard for their services and outcomes than private companies, and liabilities are often treated as net risk, rather than as an adaptive framework for risk management. Additionally, the impact of AI systems creates new complexities that challenge existing liability practices and regimes.

Currently there are few, if any, meaningful legislative safeguards to protect against the evolving discriminatory impacts of AI systems, such as the violation of human rights or antidiscrimination laws. Even when public agencies are subject to legal obligations covering AI or other technologies, these obligations are not necessarily considered by private companies in the design and testing of a technology product for use in the public sector. Rather, companies tend to use spaces with little citizen and public sector protection as a liability-free testground. The technologies that are tested in these liability-free spaces are then deployed by local agencies.

In addition, the inherent uncertainty about the capabilities and functionalities (ex ante) and impacts (ex post) of AI systems may require a re-evaluation of the distribution of liabilities across AI supply chains. Direct contractors of public agencies may have many different suppliers themselves, calling into question who holds the liability for the performance of the AI system that gets deployed. Important measures that can help develop more clarity on the distribution of liabilities across this AI supply chain include improved clarity in applicable vendor contracts (e.g., allocation of liability, warranties, clarity on trade secret protections, or insurance for AI incidents); rigorous vendor diligence; post-deployment monitoring; and quality standards.

QUESTIONS THAT CAN BE ASKED IN ORDER TO ADDRESS THE LIABILITIES TENSION POINT:

In what ways can procurement systems be encouraged to incorporate liability as part of design considerations around AI systems and their deployment?

How can procurement processes be changed in order to integrate an adaptive framework of risk assessment?

What updates are needed to procurement systems or laws in order to offer a gateway for assessing the distribution of liabilities across AI supply chains?

FIVE NARRATIVE TRAPS

The roundtable conversations brought about five narrative traps that appear to hinder innovation in procurement, and AI procurement specifically. Paying attention to these narrative traps is important, because narratives - ensembles of texts, images, spectacles, events, and cultural artifacts - are powerful tools that shape the trajectory of policies and practices, not least because they mobilize our collective imagination into certain directions over others. In order to improve the procurement process so that it can mitigate the potential harms of AI and foster innovation in government services, it is important to stay mindful of these narrative traps and find ways to avoid them.

FIVE NARRATIVE TRAPS

"We must engage the public."

Why this narrative is a trap: First, it deliberately diffuses what "engagement" means - does it refer to a democratic process, citizen participation in administrative decision-making processes, public oversight, or anything else? A lack of defining engagement means that communities will find it difficult to demand it. Second, and relatedly, it leaves unclear what "the public" means and ignores those communities who have historically been excluded from that frame (see narrative trap 2). Those communities may even be excluded further in spaces that are centered on public engagement participation.[44] And third, this narrative unloads accountability onto "the public." When there is a call for the public to get engaged in order to mitigate the risk of AI systems, then this is an indicator that elected officials and administrators have failed, or are about to fail, in adequately representing the interests of the public.

How the trap can be avoided: This narrative trap can be avoided, at least partially, by using existing and new processes and protocols - such as in procurement - to clearly define what "engagement" means, and what "public" means, in particular contexts. These definitions can and must take their cue from local communities and advocacy groups, and they should be communicated publicly and widely. Similarly, there should be clear communication about the distribution of responsibility, and the different ways in which "the public interest" gets framed in different tenders for AI systems.

---

[44] Sloane, M., Moss, E., Awomolo, O., Forlano, L. "Participation Is Not A Design Fix For Machine Learning" Pre-Print, August 11, 2020, https://arxiv.org/abs/2007.02423. Accessed May 12, 2021.

FIVE NARRATIVE TRAPS
“We must find simple definitions of ‘X.’”

Why this narrative is a trap: Narratives that call for simple definitions of the complex contexts and issues around AI drive the deployment of monolithic framings of key terms such as “the public,” “government,” “bias,” and “algorithm.” These monolithic framings tend to exclude nuance and derive from dominance in the public discourse, therefore upholding historic power structures. For example, historically the term “the public” did not include communities of color, and “government” did not mean “indigenous government.” A continued use of those terms without nuance and historic context will perpetuate exclusion. Similarly, calls for simplification can function as incubators for new narrative traps, because they generally set the expectation that generalization - over nuance - is the key to successful technology design and policy processes. For example, a call for generalizing “bias” can actually cause terms like “harm” or “impact” to be generalized to a degree that they are removed from the actual lived experiences of individuals.

How the trap can be avoided: This narrative trap can be avoided through a number of steps. First, a reflexive use of monolithic terms, where they cannot be avoided, should be normalized. For example, acknowledging the historical exclusionary nature of “government” when describing government involvement in AI deployment should become part of the story. Second, a question that can, or should, always be added is “Who is left out of the narrative?” And third, nuance in description about AI functionalities and contexts (including socio-political and organization contexts) is needed in order to normalize the complexity in narratives around AI. This is important, because ultimately, complexity is needed to create and retain accountability structures that work.

FIVE NARRATIVE TRAPS

“The main threat is the government use of data.”

Why this narrative is a trap: The idea that government use of data is potentially a bigger threat than the private use of data is a fallacy. In many contexts, governments, or government agencies do not possess data archives of the same size as industry, nor do they have capacity to collect, clean, and/or analyze that data. Relatedly, they may have stricter rules as to the collection and processing of data, limiting their ability to use data effectively where it would be needed. Narratives that suggest the opposite are a trap, because they divert regulatory attention and scrutiny away from private companies. They also distract from the government capture of citizen data through private companies via vendor contracts or public-private partnerships.

How the trap can be avoided: This narrative trap can be avoided by cultivating public knowledge about the abilities and limits of government agencies related to data collection use, as well as the data exchange between private and public entities. Generating information on this relationship could be mandated in new procurement rules. This information should also focus on identifying exactly what data will be collected, owned by, and given to private companies, and to what end. The conversation about threats stemming from data collection and use should not be limited to data, but should also include data by-products, such as improved models.

FIVE NARRATIVE TRAPS

“One incentive shared across all actors can initiate change.”

Why this narrative is a trap: This narrative is a trap because it implies that a silver-bullet approach can serve as a solution for very complex and emerging problems that occur in the context of AI and procurement. It perpetuates the idea that goal-alignment and compliance can be achieved through introducing a shared incentive (e.g. avoid fines), which is a strategy that ignores that there are larger systems of (capitalist or bureaucratic) incentives at play for different actors that may override any single incentive. The narrative is also a trap because it can lead to oversimplified goals that are poorly defined and therefore hard to create accountability mechanisms for (such as “help citizens”). Lastly, it can distract from other change mechanisms that may be more effective for innovating AI procurement.

How the trap can be avoided: This narrative trap can be avoided by tracking the different incentive structures that actors bring to the table. Creating clarity about the ways in which government agencies, vendors, communities, regulators, and other stakeholders are incentivized at present can serve as a starting point for developing more nuanced narratives around what effective and targeted mechanisms can be developed for recalibrating procurement processes so that they can manage and mitigate AI risk.

FIVE NARRATIVE TRAPS

“We can create change in AI design and deployment through procurement alone.”

Why this narrative is a trap: This last narrative trap is one that we do not want to fall into with this project and with this primer. Narratives postulating that procurement can solve any and all issues related to government use of AI systems (ranging from potential harms to liabilities) are based on a silver-bullet approach that is likely to promise innovation and change in a way that it cannot deliver, especially across the many different government agencies. They also ignore the fundamental fact that it is exceptionally difficult to change and improve the procurement process iteratively in order to address the emergent nature of AI risk.

How the trap can be avoided: This narrative trap can be avoided by developing clear definitions and understandings of the procurement process (and its history), especially as it is organized in different agencies. As part of that, the abilities, and inabilities, of government agents (and procurement officers specifically) to intervene in AI design and deployment process should be clearly outlined. Relatedly, it is important that government agencies publicize the existing AI systems already in use, for example through public registries. The most important way to avoid this narrative trap is to build a network of procurement officers, AI researchers, representatives of advocacy groups, and more, to share knowledge, experience, resources, and build capacity.

FOUR CALLS FOR ACTION

In order to facilitate innovation in public procurement so that challenges, tension point, and narratives can be addressed, four concrete actions can be taken.

FOUR CALLS FOR ACTION

Redefine the process.

The growing need for developing and enforcing accountability structures in the context of the procurement and deployment of AI systems by public agencies suggests that said processes should be redefined so that public agencies can understand the tradeoffs and benefits of AI systems quicker and with more certainty. There needs to be ample time to define and document the problem that the AI system should address, and how it should address it. Affected communities must be heard. New AI legislation, on both the national and international levels (such as the new EU AI regulation), must be absorbed effectively and translated into change in AI design, procurement, and use. A redefinition of this process can create space for the notion of collective accountability to evolve through innovation in AI procurement, whereby government agencies, as buyers, can extend their power to demand accountability and transparency of vendors, and where agencies can go back and iterate when problems occur. It can also address the false dichotomy between fast approaches that are “wrong” and slower approaches that are “right.” This redefinition should recall that policy decisions are often encoded in the definition and construction of AI systems, and therefore procurement of these systems should be undertaken with the nuance and consideration given to other policy determinations.

FOUR CALLS FOR ACTION

Create meaningful transparency.

There is a need to improve communication and issue presentation of AI systems and the risk and harms they can pose. Procurement officers, policymakers, citizens, and vendors must gain a better understanding of how individual situations of AI harm and risk connect to bigger structural problems, and vice versa. In order to create meaningful transparency, standards for the communication of the goals and assumptions baked into an AI system, as well as the risks and harms the system can pose, should be established, alongside guidelines for the documentation and record keeping of such communication.

FOUR CALLS FOR ACTION
Build a network.

There is a need for interagency communication, as well as exchange and capacity building on issues related to the procurement of AI systems. There also is a need to more clearly define intra-agency responsibilities for the procurement of AI systems and the impact they can have on citizens as well as on the agencies themselves. Resources must be developed and shared for supporting individuals and communities within agencies who are working toward improving procurement processes in order to mitigate AI harm. Similarly, assistance for communities outside of agencies who are surfacing AI harms and issues should be made available. These resources and opportunities for capacity building should be pooled in a network of procurement officers, AI researchers, representatives of advocacy groups, and more.

FOUR CALLS FOR ACTION

Cultivate talent.

The field of public interest technology is growing significantly. Government agencies are increasingly using technology, including AI systems, across all aspects of their work. This means that there is a growing need for public interest technologists: professionals who are trained in both technical and social science fields and are able to adequately assess the social impacts of technology as they continually emerge. Luckily, there also is a growing desire among the next generation of technologists to engage in meaningful work that takes into account the social impact of technology. It is therefore paramount that this talent is cultivated early and equitably. Education must become more interdisciplinary and applied, while being grounded in theories of not just ethics and moral philosophy, but scholarship of inequality, racism, feminism, intersectionality, and more. This education must continue beyond the academe and be afforded to practitioners and communities alike. At the same time, government agencies and the private sector should focus on creating a whole range of public interest technology jobs, recruit and retain diverse talent early in their careers, and support public interest technology teams internally and externally.

TEAM

The team thanks the roundtable participants for their expertise and critical perspectives. The material presented in the primer is the authors' reflection on the conversations that took place during the roundtables, as well as the research the project team has undertaken on the topic of AI procurement in the United States. This primer does not represent the opinions of the roundtable participants.

TEAM

Mona Sloane, Principal Investigator

Mona Sloane, Ph.D., is a sociologist working on design and inequality, specifically in the context of AI design and policy. She frequently publishes and speaks about AI, ethics, equitability and policy in a global context. Mona is a Senior Research Scientist at the NYU Center for Responsible AI, and an Adjunct Professor at NYU's Tandon School of Engineering, as well as a Fellow with NYU's Institute for Public Knowledge (IPK), where she convenes the Co-Opting AI series and co-curates The Shift series. She also is the technology editor for Public Books, and a fellow with The GovLab.

From fall 2021, Mona will serve as the director of the *This Is Not A Drill* program, which will develop a public pedagogy on art, equity, technology and the climate emergency. Recent projects she has led as principal investigator include the Terra Incognita NYC project, an investigation of New York City's digital public spaces in the pandemic, as well as the Procurement Roundtables project, which she led together with Dr. Rumman Chowdhury (Director of Machine Learning Ethics, Transparency & Accountability at Twitter, Founder of Parity) in collaboration with John C. Havens (IEEE Standards Association). Currently, Mona works with Emmy Award-winning journalist and NYU journalism professor Hilke Schellmann on hiring algorithms, auditing, and new tools for investigative journalism and research on AI. With Dr. Matt Statler (NYU Stern), she also leads the Public Interest Technology Convention and Career Fair project which will bring together students and organizations in the public interest technology space across the United States and beyond.

Mona is also affiliated with the Tübingen AI Center in Germany where she leads a 3-year federally funded research project on the operationalization of ethics in German AI startups. From 2020-2021 she was part of the inaugural cohort of the Future Imagination Collaboratory (FIC) Fellows at NYU's Tisch School of the Arts.

Mona holds a PhD from the London School of Economics and Political Science and has completed fellowships at the University of California, Berkeley, and at the University of Cape Town. She has written for The Guardian, MIT Technology Review, Frankfurter Allgemeine Zeitung, OneZero Medium, and other outlets. You can follow her on Twitter @mona_sloane

Rumman Chowdhury, Co-Principal Investigator

Dr. Rumman Chowdhury's passion lies at the intersection of artificial intelligence and humanity. She is a pioneer in the field of applied algorithmic ethics, creating cutting-edge enterprise technical solutions for ethical, explainable, and transparent AI since 2017.

She is currently the Director of the META (ML Ethics, Transparency, and Accountability) team at Twitter as well as GP of a venture capital fund, Parity Responsible Innovation Fund, that invests in early-stage responsible technology startups. She was previously CEO and founder of Parity AI, an enterprise algorithmic audit platform company and formerly served as Global Lead for Responsible AI at Accenture Applied Intelligence.

Rumman has been featured in international media, including the Wall Street Journal, Financial Times, Harvard Business Review, NPR, MIT Sloan Magazine, MIT Technology Review, BBC, Axios, Cheddar TV, CRN, The Verge, Fast Company, Quartz, Corrierre Della Serra, Optio, Australian Broadcasting Channel, and Nikkei Business Times. She has been recognized as one of BBC's 100 Women, Bay Area's Top 40 Under 40, and honored to be inducted into the British Royal Society of the Arts (RSA). She has also been named by Forbes as one of the Five Who are Shaping AI.

As service to the field and the larger community, she serves on the board of Oxford University's Commission on AI and Governance, the University of Virginia's Data Science Program, and Patterns, a data science journal by the publishers of Cell.

Dr. Chowdhury holds two undergraduate degrees from MIT, a master's degree in Quantitative Methods of the Social Sciences from Columbia University, and a doctorate in political science from the University of California, San Diego. You can follow her on Twitter @ruchowdh.

John C. Havens, Collaborator

John C. Havens is Executive Director of The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems that has two primary outputs – the creation and iteration of a body of work known as Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems and the recommendation of ideas for Standards Projects focused on prioritizing ethical considerations in A/IS. Currently there are thirteen approved Standards Working Groups and one completed Standard in the IEEE P7000™ series. He is also Executive Director for The Council on Extended Intelligence (CXI) that was created to ensure society prioritizes people, planet and prosperity by promoting responsible participant design, control of data rights, and holistic metrics of prosperity. CXI is a program founded by The IEEE Standards Association and MIT, whose members include representatives from the EU Parliament, the UK House of Lords, and dozens of global policy, academic, and business leaders. Previously, John was an EVP of Social Media at the PR firm Porter Novelli and a professional actor for over 15 years. John has written for Mashable and The Guardian and is author of the books, Heartificial Intelligence: Embracing Our Humanity To Maximize Machines and Hacking Happiness: Why Your Personal Data Counts and How Tracking it Can Change the World. You can follow John on Twitter @johnchavens.

TEAM

Tomo Lazovich, Collaborator

Tomo Lazovich, Ph.D., is a machine learning scientist focused on issues at the intersection of technology and equity, in particular working to eliminate the disproportionate harms that AI systems have on marginalized communities. Tomo is currently a Senior Machine Learning Researcher on Twitter's ML Ethics, Transparency, and Accountability (META) team. They have previously done research at the Charles Stark Draper Laboratory and Lightmatter. Tomo received their PhD in Physics from Harvard University in 2016. You can follow them on Twitter @laughsovich.

Luis C. Rincon Alba, Research Assistant

Luis Rincon Alba is a Colombian artist and scholar based in New York City since 2010. He has taught at the departments of Art and Public Policy and Performance Studies at New York University's Tisch School of the Arts. He is currently a doctoral candidate in the Performance Studies Department at New York University and a Public Humanities Fellow at Humanities NY and the Urban Democracy Lab. As an actor, performer, and oral narrator, he has collaborated with different artistic collectives in his home country and also in Brazil, Argentina, Mexico, the United States, and Italy. He is also the artistic director of the collective MUSA Presents.

We want to thank the Northeast Big Data Innovation Hub for funding this project through their seed grant program. We also want to thank David Rubenstein and Anna Gressel for their insightful comments on this primer, as well as Kayla Krieger, Melissa Lucas-Ludwig, Clara McMichael, and Janina Zakrzewski for their invaluable support.

ROUNDTABLE PARTICIPANTS

Chris Albon, Director of Machine Learning, Wikimedia Foundation

Stephanie Russo Carroll, (Ahtna-Native Village of Kluti-Kaah), Associate Director and Manager, Tribal Health Program; Assistant Research Professor, Udall Center for Studies in Public Policy; Assistant Professor, Public Health and American Indian Studies Graduate Program; Co-Director, Center for Indigenous Environmental Health Research, University of Arizona

Ashley Casovan, Executive Director, Responsible AI Institute

Christina J. Colclough, Founder, the Why Not Lab; Member of the Steering Committee, Global Partnership on AI; Board of Advisors, AI and Equality Initiative, Carnegie Council; Member, OECD One AI Expert Group, UN Secretary General Roadmap for Digital Cooperation

Jonas Eliasson, Director of Transport Accessibility, Swedish Transport Administration; Visiting professor of Transport Systems, Linköping University; Chair of the Division of Planning and Civil Engineering of the Royal Swedish Academy of Engineering Sciences (IVA); Vice Chair of the Board of Expertgruppen för Offentliga Studier (ESO)

Alex Engler, David M. Rubenstein Fellow, Brookings Institution; Adjunct Professor, McCourt School of Public Policy, Georgetown University; Affiliated Scholar, Massive Data Institute, Georgetown University

Linda van de Fliert, Project Manager Public Tech, Gemeente Amsterdam

Chris Gilliard, Professor of English, Macomb Community College; Visiting Research Fellow, Shorenstein Center, Harvard Kennedy School; Member, UCLA Center for Critical Internet Inquiry Scholars Council; Member, Surveillance Technology Oversight Project

Noel Hidalgo, Co-Founder, Executive Director, BetaNYC

Katie Kinsey, Staff Attorney, The Policing Project

ROUNDTABLE PARTICIPANTS

Mark Latonero, Senior Policy Advisor, Partnership on AI; Senior Associate, Human Rights Initiative, Center for Strategic and International Studies; Fellow, Carr Center for Human Rights Policy, Harvard Kennedy School; Affiliate, Data & Society Research Institute; Senior Research Fellow, Human Rights Center, UC Berkeley Law; Senior Research Fellow, Center for Communication Leadership & Policy, USC Annenberg

Graham MacDonald, Chief Data Officer, Associate Vice President of Technology and Data Science, Urban Institute

Sean McDonald, Senior Fellow, Centre for International Governance Innovation; Co-Founder, Digital Public; CEO, FrontlineSMS; Fellow, Duke Center on Law & Technology, Duke University; Advisor, Digital Democracy; Advisor, Ethics and AI Committee, IEEE

Brandeis Marshall, Professor of Computer Science, Spelman College; Faculty Associate, Berkman Klein Center, Harvard University; Co-Founder and CEO of DataedX

Emanuel Moss, Anthropologist, CUNY Graduate Center; Researcher, AI on the Ground Initiative, Data & Society Research Institute; Research Assistant, Pervasive Data Ethics for Computational Research

Deirdre Mulligan, Professor of Law, School of Information, UC Berkeley; Faculty Director, Berkeley Center for Law & Technology; Co-Organizer of the Algorithmic Fairness & Opacity Working Group; Affiliated Faculty on the Hewlett funded Berkeley Center for Long-Term Cybersecurity; Faculty Advisor to the Center for Technology, Society & Policy

Clara Neppel, Senior Director, IEEE European Business Operations

Safiya Noble, Associate Professor, Department of Information Studies, UCLA; Co-Founder and Co-Director, UCLA Center for Critical Internet Inquiry; Research Associate, Oxford Internet Institute, University of Oxford; Commissioner, Oxford Commission on Ai & Good Governance; Board Member, Cyber Civil Rights Initiative; Author of Algorithms of Oppression

ROUNDTABLE PARTICIPANTS

Kholoud Odeh, Vice President, Technology and Data Science, Chief Information Officer, Urban Institute

Eli Pariser, Co-Director, Civic Signals; Author of The Filter Bubble; Co-Founder of Upworthy; Co-Founder of Avaaz; Omidyar Fellow, New America Foundation; former Executive Director of MoveOn.org

DJ Patil, Senior Fellow, Belfer Center, Harvard Kennedy School; Board of Directors, Devoted Health; Former Chief Data Scientist of the United States Office of Science Technology and Policy

Desmond Patton, Associate Professor of Social Work; Associate Dean of Curriculum Innovation and Academic Affairs; Courtesy Appointment in Department of Sociology; Founding Director of SAFElab, Columbia University

Andrea Renda, Senior Research Fellow, Head of Global Governance, Regulation, Innovation & Digital Economy, Center for European Policy Studies; Non-resident Senior Fellow, Kenan Institute for Ethics, Duke University

Kathleen Riegelhaupt, Associate Director Digital Policy, New York Public Library

Fabian Rogers, Constituent Advocate, Office of NYS Senator Jabari Brisport

Hana Schank, Director of Strategy for the Public Interest Technology, New America; Author of Power to the Public

Renee Sieber, Associate Professor, Department of Geography, McGill University

Julia Stoyanovich, Assistant Professor of Computer Science and Engineering, Tandon School of Engineering; Assistant Professor of Data Science, Center for Data Science; Co-Founder and Co-Director, NYU Center for Responsible AI; New York University

ROUNDTABLE PARTICIPANTS

Andrew Strait, Associate Director of Research Partnerships, Ada Lovelace Institute

Vincent Southerland, Assistant Professor of Clinical Law, NYU Law; Director of the Criminal Defense and Re-entry Clinic, NYU Law; Executive Director, Center on Race, Inequality, and the Law, NYU; Board Member, the Bail Project; Board of Trustees, the Center for Constitutional Rights

Adrianna Tan, Product Director for San Francisco Digital Services

Jeff Thamkittikasem, Director of the Mayor's Office of Operations, New York City

Stefaan Verhulst, Co-Founder, Chief Research and Development Officer, The GovLab at NYU; Editor-in-Chief of Data & Policy Journal, Cambridge University Press; Research Director, MacArthur Research Network on Opening Governance; Chair of the Data for Children Collaborative, Unicef; Member of the High-Level Expert Group to the European Commission on Business-to-Government Data Sharing; Senior Fellow, Center for Democracy and Technology (CDT), Yale University

Bianca Wylie, Co-Founder, Digital Public; Co-Founder, Tech Reset Canada; Senior Fellow, Center for International Governance Innovation; Founder, Open Data Institute Toronto; Member, Toronto Public Library's Innovation Council

Meg Young, Postdoctoral Fellow, Digital Life Initiative, Cornell Tech

AI and Procurement Primer
New York University
Summer 2021
https://doi.org/10.17609/bxzf-df18