GOVERNING AI THROUGH ACQUISITION AND PROCUREMENT

Testimony by
Rayid Ghani,
Distinguished Career Professor
Machine Learning Department and the Heinz College of Information Systems and Public Policy
Carnegie Mellon University

Before the
United States Senate Committee on Homeland Security and Governmental Affairs Hearing on "Governing AI Through Acquisition and Procurement"

Thursday, September 14, 2023

Chairman Peters, Ranking Member Paul, Members of the Committee, thank you for hosting this important hearing today, and for giving me the opportunity to submit this testimony.

My name is Rayid Ghani and I am a Distinguished Career Professor in the Machine Learning Department and the Heinz College of Information Systems and Public Policy at Carnegie Mellon University. I've worked in the private sector, in academia, and extensively with government agencies and non-profits in the US and globally on developing and using Machine Learning and AI systems to tackle social and public policy problems across health, criminal justice, education, public safety, human services, and workforce development in a fair and equitable manner.

The promise of AI in helping build a better society

Artificial Intelligence has enormous potential in helping tackle critical problems we face in society today, ranging from improving the health of our children by reducing their risk of lead

poisoning[1], to reducing recidivism rates for people in need of mental health services[2], to improving educational outcomes for students at risk of not graduating from school on time[3][4], to improving police-community relations by identifying officers at risk of adverse incidents[5], to supporting proactive inspections to improve health and safety conditions in workplaces[6] and in rental housing[7], to improving healthcare practices[8], to designing more effective organ exchange systems[9]. There is tremendous potential for every federal agency to use AI - in helping them design, implement, and evaluate their programs to help improve outcomes for everyone and result in a better and more equitable society.

However, any AI system affecting people's lives has to be explicitly designed to focus on increasing equity and promoting our societal values, and not just narrowly optimizing for efficiency. AI can have a massive, positive social impact but we need to make sure that we put guidelines in place to maximize the chances of that positive impact. If not designed, deployed, and used appropriately, it can risk harm to people who have been traditionally marginalized in society. An AI system, designed to explicitly optimize for efficiency, has the potential to result in leaving "more difficult or costly to help" people behind, resulting in an increase in inequities. It is critical for government agencies and policymakers to ensure that AI systems are designed, developed, and used in a responsible manner to ensure that they result in supporting equitable outcomes for everyone. In a policy brief published by the Responsible AI Initiative at the BlockCenter at Carnegie Mellon University, we highlighted some of the unique challenges of AI Accountability and lay out a set of policy recommendations[10].

Scoping and Procurement of AI systems needs to be a focus area for policymakers

While the entire lifecycle of AI systems - scoping, procurement, designing, testing, deploying, and using needs to have guidelines and best practices in place that maximize the societal benefit and minimize potential harms (such as the efforts around the AI Risk Management Framework[11]

---

[1] Predictive Modeling for Public Health: Preventing Childhood Lead Poisoning. Potash et al. Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2015)
[2] Reducing Incarceration through Prioritized Interventions. Bauman et al.. ACM SIGCAS Conference on Computing and Sustainable Societies, 2018.
[3] A Machine Learning Framework to Identify Students at Risk of Adverse Academic Outcomes. Lakkaraju et al. Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining
[4]http://www.dssgfellowship.org/project/identifying-factors-driving-school-dropout-and-improving-the-impact-of-social-programs-in-el-salvador/
[5] Early Intervention Systems – Predicting Adverse Interactions Between Police and the Public. Helsby et al. Criminal Justice Policy Review, 2017.
[6]http://www.dssgfellowship.org/project/improving-workplace-safety-through-proactive-inspections
[7] http://www.datasciencepublicpolicy.org/projects/public-safety/san-jose-housing/
[8] Kilic A, Dochtermann D, Padman R, Miller JK, Dubrawski A (2021). Using machine learning to improve risk prediction in durable left ventricular assist devices. PLOS ONE 16(3).
[9]https://aaai.org/tuomas-sandholm-wins-2023-aaai-award-for-artificial-intelligence-for-the-benefit-of-humanity/#:~:text=This%20year%2C%20the%20AAAI%20Awards,on%20both%20practice%20and%20policy.
[10] https://www.cmu.edu/block-center/responsible-ai/index.html
[11] https://www.nist.gov/itl/ai-risk-management-framework

being developed by NIST), there has been a lack of attention to the earlier phases of this process, specifically scoping and procurement. Many of the AI systems being used in federal, state, and local agencies are not built in-house but procured through vendors, consultants, and researchers. This makes getting the procurement phase correct critical - many costly problems and harms discovered downstream can be avoided by a more effective and robust procurement process.

We need to ensure that government procurement of AI follows a "responsible" process, and in turn requires AI vendors to follow a "responsible" process in designing such systems, and results in the selection, deployment, and use of a system that promotes accountability and transparency, and leads towards equitable outcomes for those impacted.

Procuring solutions to specific problems rather than procuring "AI"

Too often, organizations go on the market to buy "AI" without completely understanding, defining, and scoping the concrete problem they want to tackle, without assessing whether AI should even be part of the solution, and without including individuals and communities that will be affected. AI systems are neither applicable for all problems facing government agencies, nor are they one-size-fits-all. By starting with the concrete problem at hand, and understanding how it's being tackled today, an effective, collaborative, and inclusive scoping process can help determine the requirements that the AI system needs to fulfill. For example, consider procuring a system to support matching unemployed individuals with training or skilling programs that they can be enrolled in to get them back into employment most effectively. The requirements for such a system are not primarily AI requirements, but rather come from the intended goals of the program being administered - whether this system will result in increased employment rates for the individuals impacted, whether it will propagate existing disparities in training and employment outcomes, whether it will enable and empower employment agencies, unemployment counselors, and unemployed individuals, among other requirements that will be surfaced through the scoping process, such as the one we've previously developed[12].

AI systems optimize for what their developers tell them to optimize for (and the procurement process needs to tell the vendors what to optimize for)

AI algorithms are neither inherently biased nor unbiased (in the societal sense) or have inherent, fixed, "values". The AI developers designing and building the AI system (implicitly or explicitly) make hundreds of design choices that result in the eventual system and its behavior. If the developers make design choices that explicitly focus on societal outcomes we care about and evaluate their systems against those intended outcomes, the AI system can help achieve what we want it to achieve. The procurement process needs to define these goals and values, and ensure that the vendors address those appropriately in the system being procured.

[12] http://www.datasciencepublicpolicy.org/our-work/tools-guides/data-science-project-scoping-guide/

AI is forcing us to make societal (and public policy) values explicit (and the procurement process needs to define what those values should be)

Because an AI system requires us to define exactly 1) what we want to optimize it for, 2) which mistakes are costlier (financially or socially) than others, and 3) by how much, it forces us to make these ethical and societal values explicit. These values are implied in any decision-making process, including all the human decision-making processes that exist today, but are not necessarily made explicit. These implicit values (coded through human decisions) when biased and unfair, result in inequitable outcomes.

For an AI system to be built, these values need to be provided as a critical input[13]. For example, for a system that is recommending lending decisions, we may have to 1) specify the differential costs of highlighting someone as unlikely to pay back a loan and being wrong about it versus predicting that someone will pay back a loan and being wrong about it, and 2) specify those costs explicitly in the case of people who may be from different gender, race, income, or education level groups. While that may have happened implicitly in the past and with high levels of variation across different human decision makers (loan officers in this case), with AI-assisted decision-making processes, we are forced to define them explicitly.

One key question the procurement process has to answer is who and how we should come up with these sets of values for a given problem setting, what information to ask for around these values, and how to evaluate the correctness of the values, and the fidelity of the procured and designed system to these values. Unfortunately, today, these decisions are too often left essentially by default to the AI system developer or an arbitrary set of individuals who define those values in an AI algorithm (explicitly or implicitly). The recommendations at the end of this testimony go into more detail on what I recommend should be done but it certainly should not be left to the AI system developer making those choices alone; the team and process should include all stakeholders including policymakers and the community being impacted by this system.

What does it take to create responsible AI systems for society?

The following steps need to be taken to create Responsible AI systems for society and the procurement process needs to set expectations and accountability for vendors in each of these steps:
1. **Defining** the goals and policy and societal outcomes the system needs to help achieve(which includes the societal values and a collaborative, multi-stakeholder process).

---

[13] From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research M. Feffer, M. Skirpan, Z. Lipton*, and H. Heidari. The AAAI /ACM Conference on Artificial Intelligence, Ethics, and Society (AIES), 2023.

2. **Translating/Mapping** those desired outcomes and values into analytical and technical requirements that the vendors should design the AI system to achieve.
3. **Building** an AI system that fulfills those analytical requirements and releasing documentation and additional artifacts enumerating all the design choices (including around the choice and use of data, the AI algorithms, and the downstream use and impact), demonstrating and providing evidence of how it was built to achieve those goals.
4. **Validating** through a trial (and providing evidence) that the AI system did, in fact, fulfill those requirements and achieve the initial outcomes defined in step 1 before deploying the system.
5. **Continuous Monitoring & Evaluation** of the entire system (the AI system followed by human decisions) during its lifetime to ensure that it continues to achieve equitable outcomes from step 1.

Moving Forward to Governments Procuring and Using AI Systems that Result in a More Equitable Society: Our Recommendations

It is critical and urgent for policymakers to act and provide guidelines and regulations for both the public and private sector organizations procuring, developing, and using AI in order to ensure that these systems are built in a transparent and accountable manner and result in fair and equitable outcomes for society. As initial steps, we recommend:

1. Focused Procurement for Specific Use-Cases

We need to ensure that AI systems are procured for specific use-cases, and to support intended outcomes around that use case, rather than as generic AI systems. This is intended to both promote better outcomes as well as to prevent harm through misuse. An AI system that yields beneficial and equitable outcomes in one context might yield just the opposite in another. While AI algorithms across different areas have a lot in common, developing a generic and complete framework for AI that works well across all possible uses is likely to be an unrealistic proposal. Rather, the need for application-grounded procurement processes is important to achieving policy and societal goals across different government agencies.

2. Development of common procurement requirements and templates

While this may seem contradictory to the previous recommendation, the documentation and artifacts needed to assess the appropriateness and effectiveness of an AI system are common across many use cases. The specifics of the use case define the concrete values, the goals, and the evaluation criteria, and the common procurement requirements and templates are used to assess how well the system is able to achieve them. These common procurement and RFP templates should include the set of artifacts that should be provided during the evaluation of the AI software. This includes information on:

1. How the system was built and what it was designed to optimize for
2. What tests were run to check if it did do what it was intended to do?
3. What types of people was it effective for? Who does it fail for?
4. How long was it in trials for, when, and how did the effectiveness change over time?
5. What risks need to be considered and what are the mitigation plans for each of these risks?
6. Any extended data collection process and infrastructure that may need to be set up to collect additional data attributes (such as race, gender, or income) that may not already be collected but are necessary to measure equity outcomes
7. How to set up evaluation standards to compare the performance of these systems to the human decision-making processes (if any) currently being used.
8. How the vendor supports explainability and interpretability of the AI systems in order to provide recourse to individuals who may be adversely impacted by the decisions made using the system.
9. A continuous improvement plan to ensure that the system continues to not only be evaluated but also improved upon to achieve the desired outcomes.

RFPs for AI systems should include an explicit initial project phase to gather requirements for the values and goals of the system. This process should include a diverse team and work with stakeholders including: developers who build and deploy AI systems, decision-makers who implement the systems in their workflows, and the community being impacted by these systems.

Ideally this should be put in place for any process involving decision making of any kind, whether human decisions or AI-assisted decisions but becomes critical in cases where the scale of deployed AI systems increases the risk. This is not an exhaustive set of questions and will need to vary based on the problem being addressed and the impact this system can have on people's lives.

3. Community Participation

Create guidelines that ensure **meaningful involvement of the communities that will be impacted** by the AI systems right from the inception stage. Engage in continuous dialogue and feedback to understand their concerns, values, and suggestions which should guide the design of RFPs, and of the design, deployment, and use of the AI systems.

4. Create Trainings, Processes, and Tools to Support Procurement Teams

As the procurement teams expand their role and start procuring AI-augmented systems, they will need to be supported by increasing their capacity to fulfill this role. We recommend creating trainings, processes, collaboration mechanisms, and tools to help them:

1. Understand where existing processes may and may not be well-adapted to systems using AI.
2. Understand and define what process and outcomes standards to set.
3. Understand how to evaluate whether the requirements created for an AI system were in fact aligned with the identified societal equitable outcomes.
4. Understand how to evaluate whether the AI system did in fact do what it was designed to do.
5. Develop a continuous monitoring and audit process and tools.
6. Create standards for when a system should "expire" and a corresponding renewal process.
7. Create technical software tools to support the end-to-end procurement process - help with scoping the need, to write RFPs, to identify issues with RFPs, to analyze responses, to conducting technical evaluations of AI software from vendors.
8. Avoid common pitfalls that result in costly downstream impact such as:
    a. Locking into unnecessary long-term contracts
    b. Inability for the underlying data systems to ingest new or external data,
    c. Inability to export data into other systems for further linkage and policy analysis
    d. Lack of interoperability with commonly used systems within and across federal, state, and local government agencies
    e. Lack of customization and configuration based on changing needs of the use case
    f. Hidden financial costs that may be involved in various processes such as for scaling the system, or in customizing it.
    g. Over-focus on trivial, software metrics such as up-time, at the expense of use-case focused metrics around effectiveness or equity.

These steps need to take a phased approach and be iterative:

- Near term: Getting started by partnering with external organizations such as universities to help create these guidelines and to provide training to government agencies
- Medium term: Collaboratively developing more rigorous procurement processes and tools to support the agencies in the medium term
- Longer Term: Providing the agencies with the resources to expand their internal capacity.

The overall goal behind these recommendations is to set some standards around procurement of AI by government agencies and to support and enable the agencies to implement those standards effectively and procure AI systems that help us achieve their policy and societal goals.