Federal AI Law and Policy Panel

BCLT Symposium on AI Law and Governance

February 28, 2025

Berkeley
Center for Law
& Technology

From Principles to Practice

Research and Policy Leading to the OMB AI Guidance

Sorelle Friedler

Shibulal Family Professor of Computer Science

HAVERFORD
COLLEGE

DEPARTMENT OF COMPUTER SCIENCE

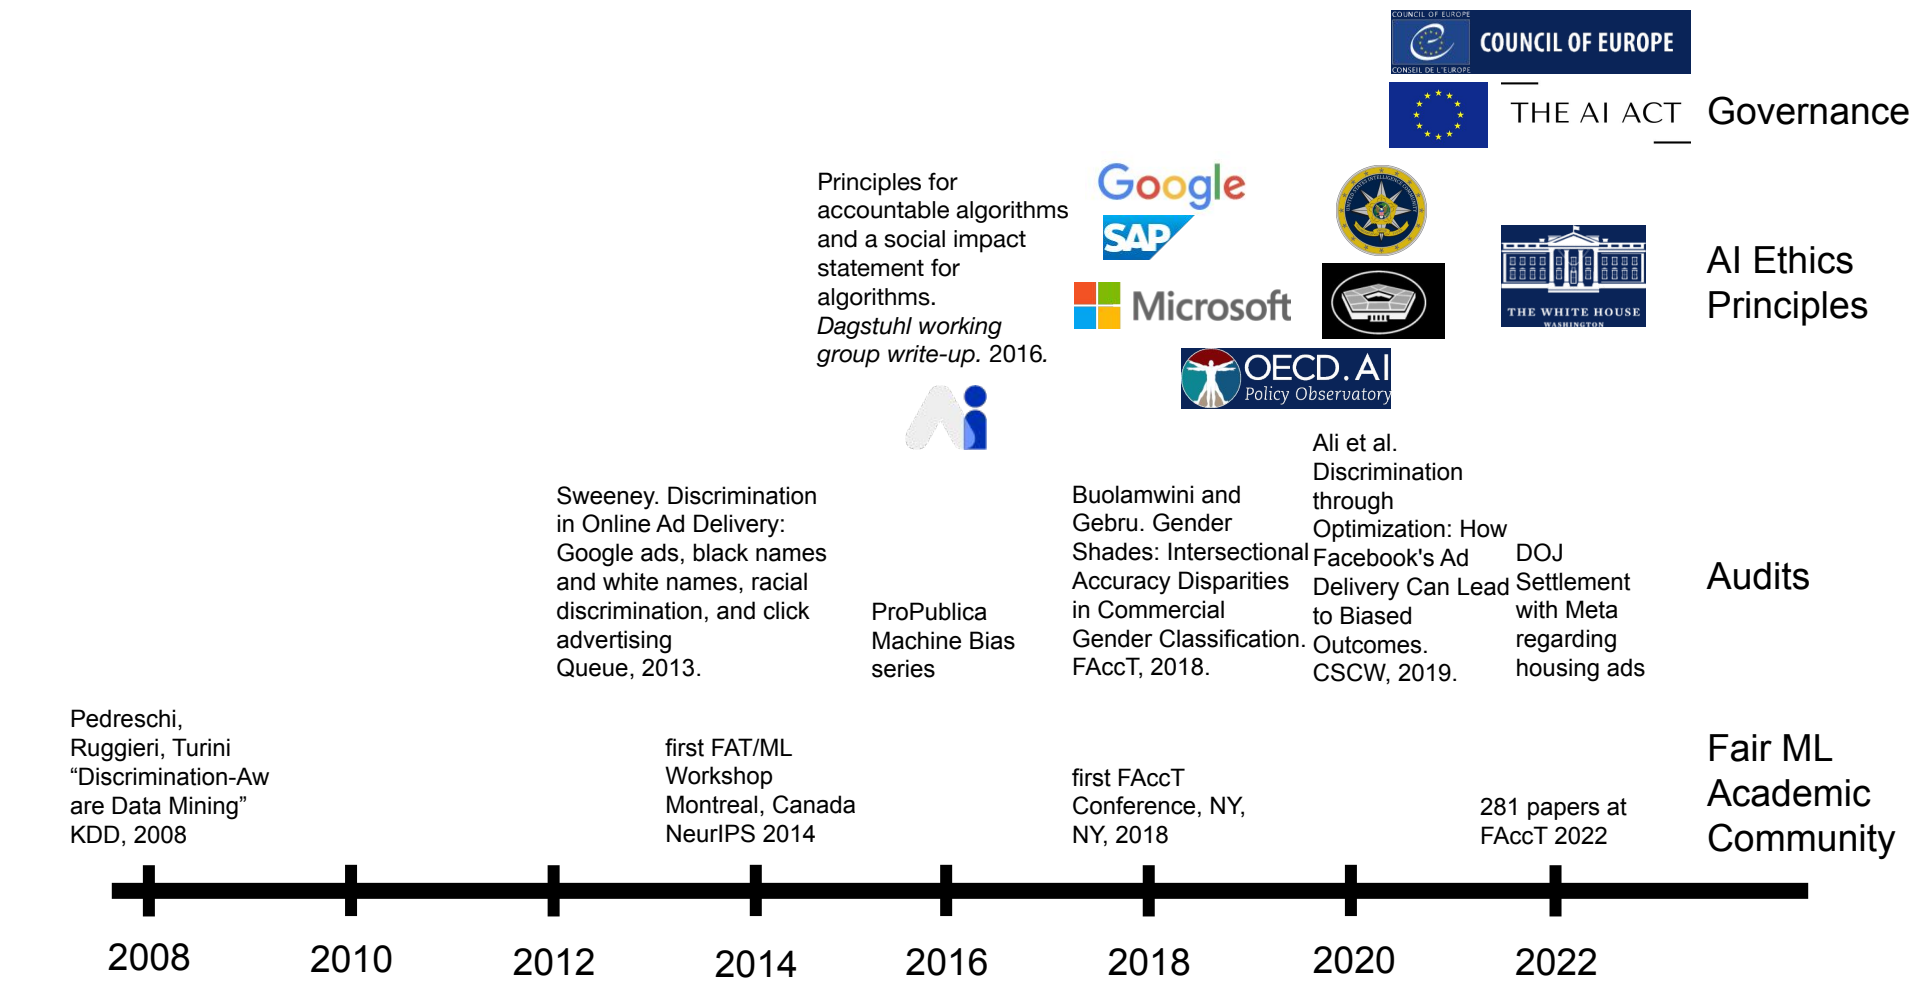

An Incomplete History of Responsible AI

BLUEPRINT FOR AN

AI BILL OF RIGHTS

MAKING AUTOMATED SYSTEMS WORK FOR THE AMERICAN PEOPLE

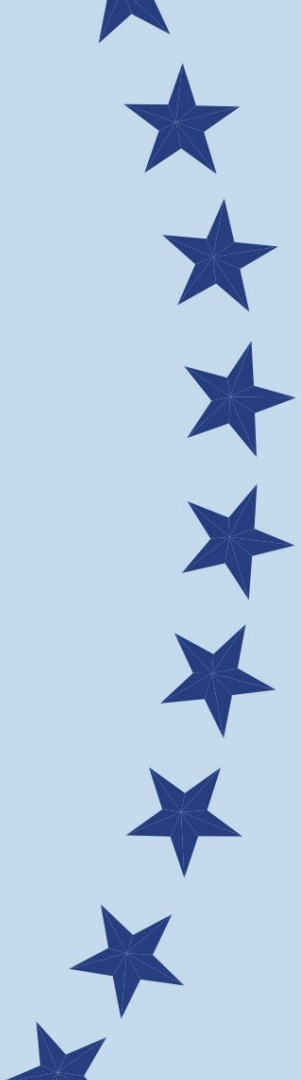

OCTOBER 2022

THE WHITE HOUSE
WASHINGTON

Blueprint for an AI Bill of Rights

THE WHITE HOUSE

Safe and Effective Systems
You should be protected from unsafe or ineffective systems.

Algorithmic Discrimination Protections
You should not face discrimination by algorithms and systems should be used and designed in an equitable way.

Data Privacy
You should be protected from abusive data practices via built-in protections and you should have agency over how data about you is used.

Notice and Explanation
You should know when an automated system is being used and understand how and why it contributes to outcomes that impact you.

Human Alternatives, Consideration, and Fallback
You should be able to opt out, where appropriate, and have access to a person who can quickly consider and remedy problems you encounter.

“ This blueprint is for the older Americans denied critical health benefits because of an algorithm change.
The student erroneously accused of cheating by AI-enabled video surveillance.
The fathers wrongfully arrested because of facial recognition technology.
The Black Americans blocked from a kidney transplant after an AI assumed they were at lesser risk for kidney disease.
It is for everyone who interacts daily with these technologies—and every person whose life has been altered by an unaccountable algorithm. ”

https://bidenwhitehouse.archives.gov/ostp/news-updates/2022/10/04/blueprint-for-an-ai-bill-of-rightsa-vision-for-protecting-our-civil-rights-in-the-algorithmic-age/

Applying the Blueprint for an AI Bill of Rights

This framework describes protections that should be applied with respect to all automated systems that have the potential to meaningfully impact individuals' or communities' exercise of:

Rights, Opportunities, or Access

Civil rights, civil liberties, and privacy, including freedom of speech, voting, and protections from discrimination, excessive punishment, unlawful surveillance, and violations of privacy and other freedoms in both public and private sector contexts;

Equal opportunities, including equitable access to education, housing, credit, employment, and other programs; or,

Access to critical resources or services, such as healthcare, financial services, safety, social services, non-deceptive information about goods and services, and government benefits.

FROM PRINCIPLES TO PRACTICE

A Techincal Companion to the Blueprint for an AI Bill of Rights

1 Why this principle is important:

This section provides a brief summary of the problems that the principle seeks to address and protect against, including illustrative examples.

2 What should be expected of automated systems:

- **The expectations for automated systems are meant to serve as a blueprint for the development of additional technical standards and practices that should be tailored for particular sectors and contexts.**
- This section outlines practical steps that can be implemented to realize the vision of the Blueprint for an AI Bill of Rights. The expectations laid out often mirror existing practices for technology development, including pre-deployment testing, ongoing monitoring, and governance structures for automated systems, but also go further to address unmet needs for change and offer concrete directions for how those changes can be made.

3 How these principles can move into practice:

This section provides real-life examples of how these guiding principles can become reality, through laws, policies, and practices. It describes practical technical and sociotechnical approaches to protecting rights, opportunities, and access.

http://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights

What should be expected of automated systems

Protect the public from algorithmic discrimination in a proactive and ongoing manner

- ❑ Proactive assessment of equity in design.
- ❑ Representative and robust data.
- ❑ Guarding against proxies.
- ❑ Accessibility ensured during design, development, and deployment.
- ❑ Disparity assessment.
- ❑ Disparity mitigation.
- ❑ Ongoing monitoring and mitigation

Demonstrate that the system protects against algorithmic discrimination

- ❑ Independent evaluation.
- ❑ Reporting.

Legal Disclaimer

The Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People is a white paper published by the White House Office of Science and Technology Policy. It is intended to support the development of policies and practices that protect civil rights and promote democratic values in the building, deployment, and governance of automated systems.

The *Blueprint for an AI Bill of Rights* is non-binding and does not constitute U.S. government policy. It does not supersede, modify, or direct an interpretation of any existing statute, regulation, policy, or international instrument. It does not constitute binding guidance for the public or Federal agencies and therefore does not require compliance with the principles described herein. It also is not determinative of what the U.S. government's position will be in any international negotiation. Adoption of these principles may not meet the requirements of existing statutes, regulations, policies, or international instruments, or the requirements of the Federal agencies that enforce them. These principles are not intended to, and do not, prohibit or limit any lawful activity of a government agency, including law enforcement, national security, or intelligence activities.

The appropriate application of the principles set forth in this white paper depends significantly on the context in which automated systems are being utilized. In some circumstances, application of these principles in whole or in part may not be appropriate given the intended use of automated systems to achieve government agency missions. Future sector-specific guidance will likely be necessary and important for guiding the use of automated systems in certain settings such as AI systems used as part of school building security or automated health diagnostic systems.

The *Blueprint for an AI Bill of Rights* recognizes that law enforcement activities require a balancing of equities, for example, between the protection of sensitive law enforcement information and the principle of notice; as such, notice may not be appropriate, or may need to be adjusted to protect sources, methods, and other law enforcement equities. Even in contexts where these principles may not apply in whole or in part, federal departments and agencies remain subject to judicial, privacy, and civil liberties oversight as well as existing policies and safeguards that govern automated systems, including, for example, Executive Order 13960, Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government (December 2020).

This white paper recognizes that national security (which includes certain law enforcement and homeland security activities) and defense activities are of increased sensitivity and interest to our nation's adversaries and are often subject to special requirements, such as those governing classified information and other protected data. Such activities require alternative, compatible safeguards through existing policies that govern automated systems and AI, such as the Department of Defense (DOD) AI Ethical Principles and Responsible AI Implementation Pathway and the Intelligence Community (IC) AI Ethics Principles and Framework. The implementation of these policies to national security and defense activities can be informed by the *Blueprint for an AI Bill of Rights* where feasible.

The *Blueprint for an AI Bill of Rights* is not intended to, and does not, create any legal right, benefit, or defense, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person, nor does it constitute a waiver of sovereign immunity.

The Blueprint for an AI Bill of Rights is non-binding and does not constitute U.S. government policy. It does not supersede, modify, or direct an interpretation of any existing statute, regulation, policy, or international instrument. **It does not constitute binding guidance for the public or Federal agencies and therefore does not require compliance with the principles described herein.** It also is not determinative of what the U.S. government's position will be in any international negotiation.

THE WHITE HOUSE

OCTOBER 30, 2023

Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence

BRIEFING ROOM › PRESIDENTIAL ACTIONS

By the authority vested in me as President by the Constitution and the laws of the United States of America, it is hereby ordered as follows:

Section 1. Purpose. Artificial intelligence (AI) holds extraordinary potential for both promise and peril. Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure. At the same time, irresponsible use could exacerbate societal harms such as fraud, discrimination, bias, and disinformation; displace and disempower workers; stifle competition; and pose risks to national security. Harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks. This endeavor

https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

Sec. 10. Advancing Federal Government Use of AI.

10.1. Providing Guidance for AI Management. (a) To coordinate the use of AI across the Federal Government, within 60 days of the date of this order and on an ongoing basis as necessary, the Director of OMB shall convene and chair an interagency council to coordinate the development and use of AI in agencies' programs and operations, other than the use of AI in national security systems. The Director of
interagency council. The interage
minimum, the heads of the agenci
of National Intelligence, and other
agencies designate their permaner
guidance described in subsection 1
represented on the interagency co
Assistant Secretary level or equiva
agency.

(b) To provide guidance on Fed
of the date of this order and updat
OMB, in coordination with the Director of OSTP, and in consultation with the interagency council established in subsection 10.1(a) of this section, shall issue guidance to agencies to strengthen the effective and appropriate use of AI, advance AI innovation, and manage risks from AI in the Federal Government. The Director of OMB's guidance shall specify, to the extent appropriate and consistent with applicable law:

(iv) required minimum risk-management practices for Government uses of AI that impact people's rights or safety, including, where appropriate, the following practices derived from OSTP's Blueprint for an AI Bill of Rights and the NIST AI Risk Management Framework: conducting public consultation; assessing data quality; assessing and mitigating disparate impacts and algorithmic discrimination; providing notice of the use of AI; continuously monitoring and evaluating deployed AI; and granting human consideration and remedies for adverse decisions made using AI;

THE WHITE HOUSE

OCTOBER 30, 2023

Executive Order on the Safe, Secure, Development and Intelligence

PRESIDENTIAL ACTIONS

ent by the Constitution and the laws by ordered as follows:

nce (AI) holds extraordinary potential for both promise and peril. Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure. At the same time, irresponsible use could exacerbate societal harms such as fraud, discrimination, bias, and disinformation; displace and disempower workers; stifle competition; and pose risks to national security. Harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks. This endeavor

https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF MANAGEMENT AND BUDGET
WASHINGTON, D.C. 20503

THE DIRECTOR

March 28, 2024

M-24-10

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: Shalanda D. Young

SUBJECT: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence

Artificial intelligence (AI) is one of the most powerful technologies of our time, and the President has been clear that we must seize the opportunities AI presents while managing its risks. Consistent with the AI in Government Act of 2020,[1] the Advancing American AI Act,[2] and Executive Order 14110 on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, this memorandum directs agencies to advance AI governance and innovation while managing risks from the use of AI in the Federal Government, particularly those affecting the rights and safety of the public.[3]

https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf

M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of AI

Leading by example in the government's own use of AI

Olivia Zhu
February 2025

M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of AI

- Legal grounding: AI in Government Act of 2020 directs the Office of Management and Budget (OMB) to "issue a memorandum [...] that shall...
  - (3) identify best practices for identifying, assessing, and mitigating any discriminatory impact or bias on the basis of any classification protected under Federal nondiscrimination laws..."
- The AI in Government Act also requires OMB to update the policy every 2 years.

M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of AI

- **Governance**: Chief AI Officers (CAIOs) at agencies, agency AI Governance Boards, AI use case inventories
- **Innovation**: Agency AI strategies, AI talent at agencies, open government data, releasing AI model weights and code, CAIO Council
- **Risk Management**

Section 5: Managing Risks

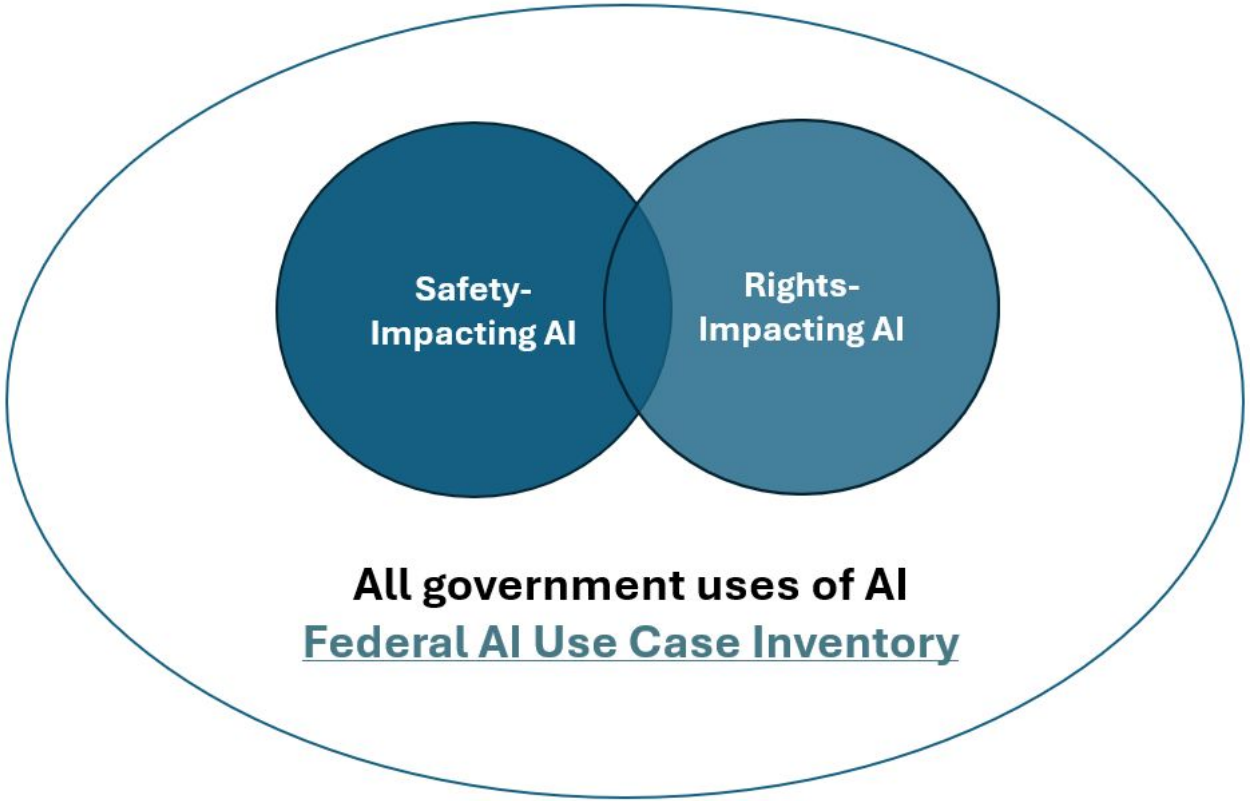

Section 5: Managing Risks

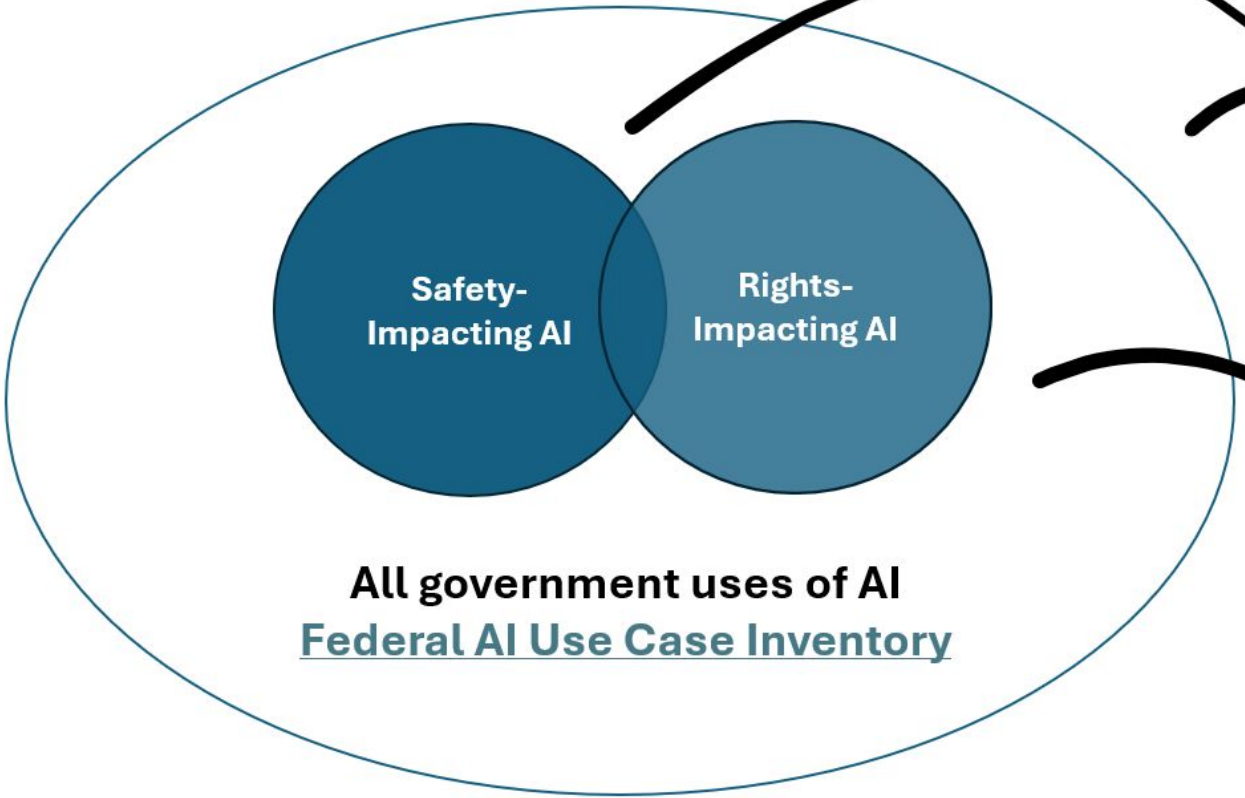

AI uses automatically presumed to be:

Safety-Impacting

1. Controlling electrical grids
2. Managing industrial waste
3. Maintaining election integrity
4. Etc.

Rights-Impacting

1. Law enforcement risk assessments
2. Detecting student cheating
3. Screening tenants
4. Etc.

Section 5: Managing Risks

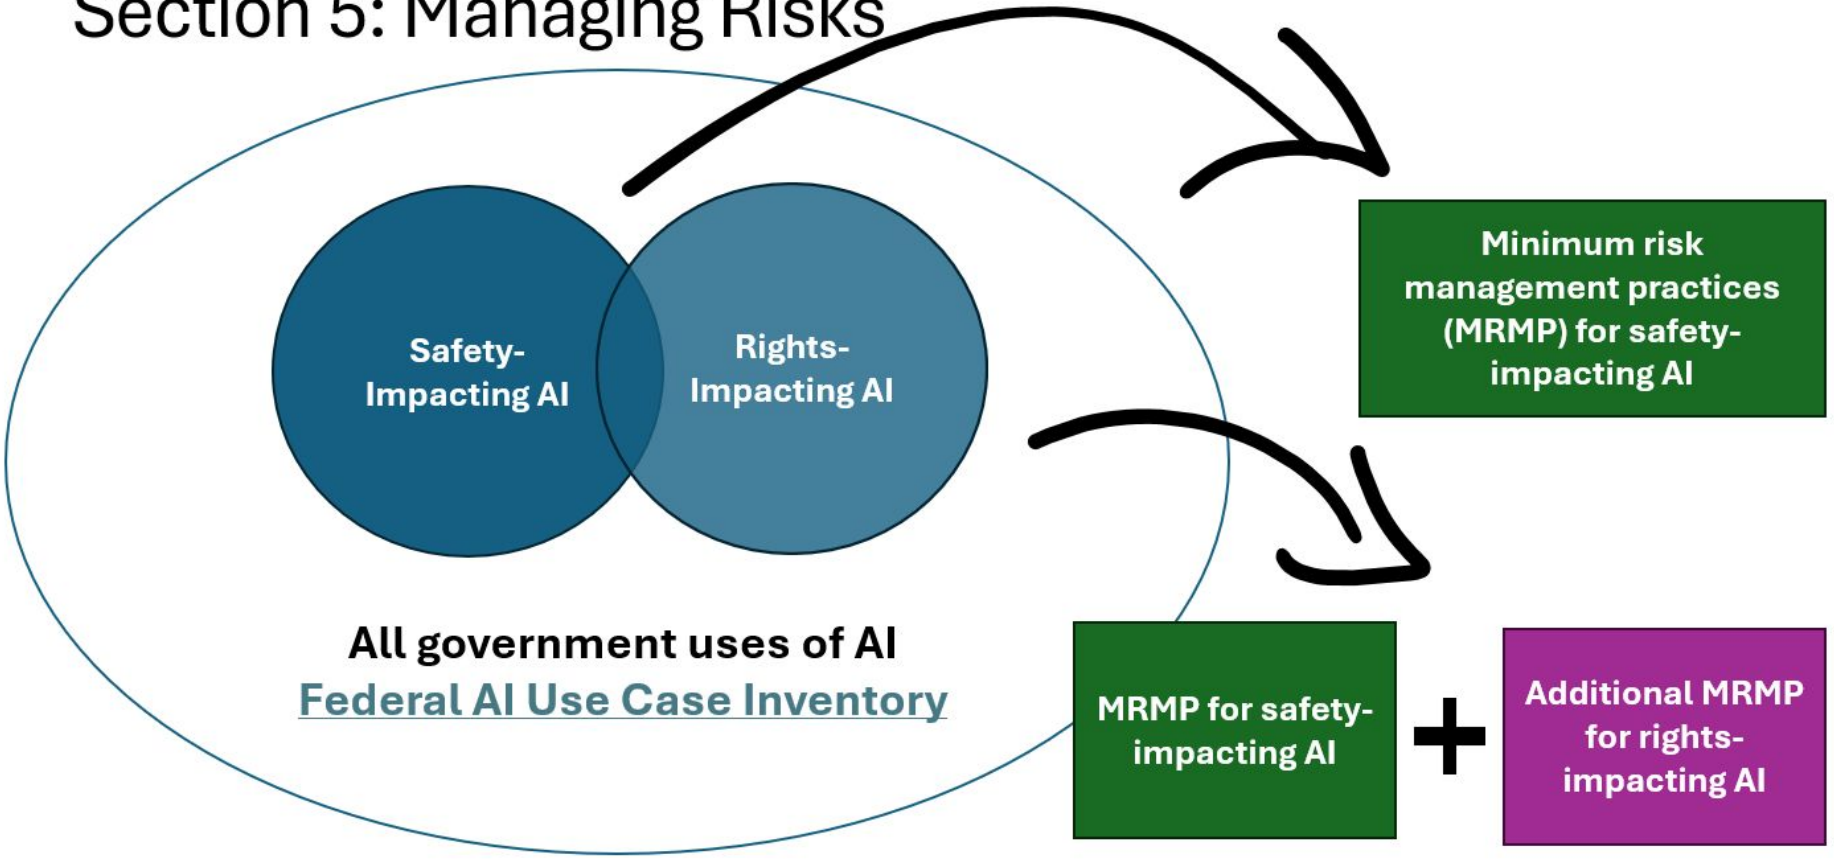

Minimum Risk Management Practices

Pre-Deployment

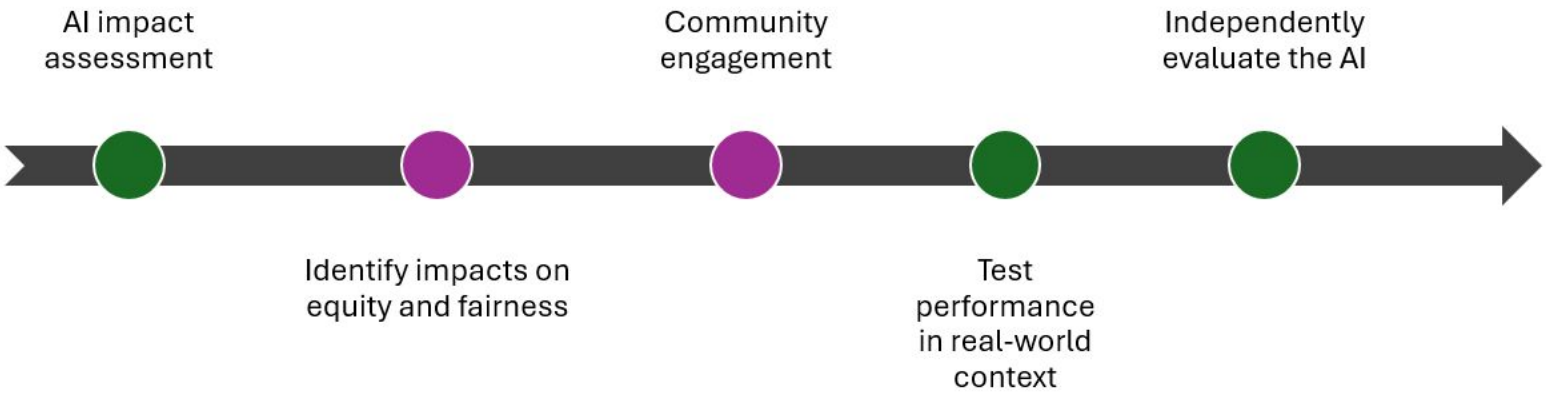

Key

- MRMP for safety-impacting AI
- Addl MRMP for rights-impacting AI

Minimum Risk Management Practices

Post-Deployment & Ongoing

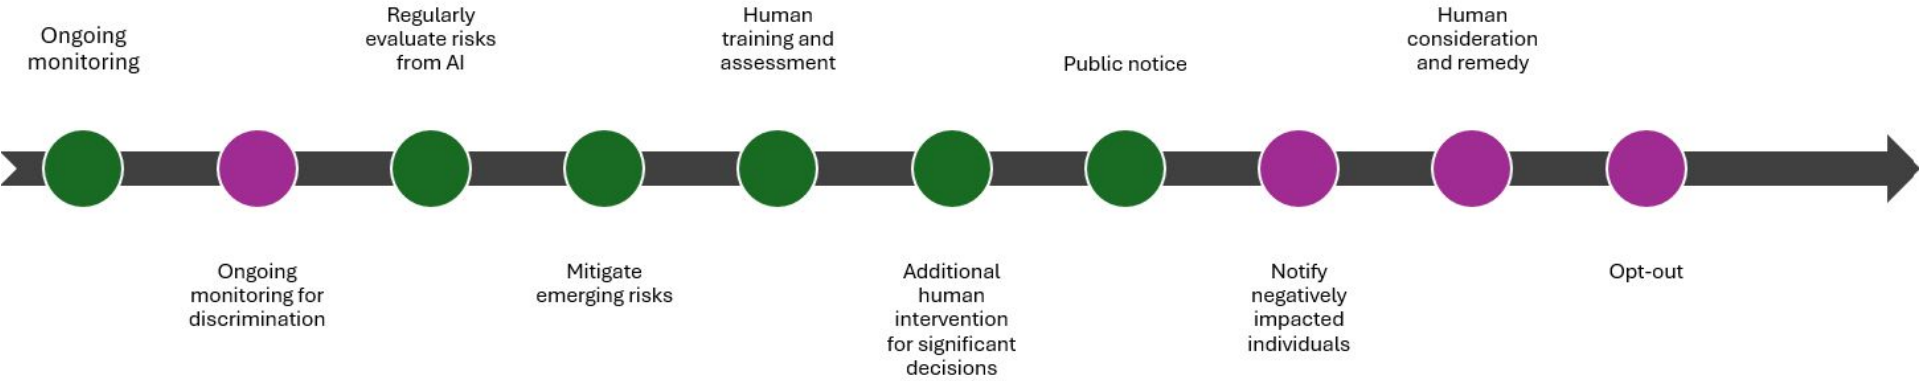

[Mona slides]

Government Tech Procurement

- **Legislative environment**
  - Competition in Contracting Act
  - FISMA
  - FITARA

- **Agency governance on IT**
  - Warranted contracting official + CIO
  - Post-award contract management by COR and CO
  - Authorities to Operate to deploy solutions
  - PIA/SORN

Government AI Procurement

- Commercial item acquisition
  - FAR Part 12
  - Commercial terms and conditions
- Different approaches at different part of the tech stack
- Different approaches for different agency requirements

For example
System integrators that bundle labor hours + solutions
OR
buying commercial AI + a separate vendor for software development
OR
COTS with no modifications at all

Government AI Procurement

- Buy with commercial procedures
- AI workforce includes the acquisition workforce
- Use-case specific governance
- What are the government's requirements?
- Resourcing for ***both*** AI governance and IT security governance
- Trump administration's focus on AI leadership
    - Executive Order on Removing Barriers to American Leadership in Artificial Intelligence
    - "AI Dominance"