



GW Law Faculty Publications & Other Works

Faculty Scholarship

2022

Artificial Intelligence Accountability of Public Administration

Francesca Bignami

George Washington University Law School, fbignami@law.gwu.edu

Follow this and additional works at: https://scholarship.law.gwu.edu/faculty_publications



Part of the [Law Commons](#)

Recommended Citation

The American Journal of Comparative Law, <https://doi.org/10.1093/ajcl/avac012> (June 1, 2022).

This Article is brought to you for free and open access by the Faculty Scholarship at Scholarly Commons. It has been accepted for inclusion in GW Law Faculty Publications & Other Works by an authorized administrator of Scholarly Commons. For more information, please contact spagel@law.gwu.edu.

FRANCESCA BIGNAMI*

Artificial Intelligence Accountability of Public Administration[†]

In law and policy debates, there are many terms that get thrown around—big data, algorithms, artificial intelligence. One way to understand the technology is to think of the big data as the fuel, algorithms as the rockets; and artificial intelligence as the planet that computer scientists seek to reach.¹ That is, to go in reverse order, the goal of artificial intelligence is to empower computers to replicate all the things that humans can do—see, hear, speak, even think. And how is that to be accomplished? By algorithms that use big data. Or, more precisely, by machine learning algorithms that use big data.

Machine learning algorithms represent a newer generation of computer science. Old style algorithms are based on complete models with relatively few explanatory variables and contain a comprehensive set of if-then statements that give instructions to a computer. Machine learning algorithms are very different: based on an initial algorithm and the data inputs and the desired output, they do the work of generating what can be an extraordinarily complex, operating algorithm. To quote from the computer scientist Pedro Domingos:

Every algorithm has an input and an output: the data goes into the computer, the algorithm does what it will with it, and out comes the result. Machine learning turns this around: in goes the data and the desired result and out comes the algorithm that turns one into the other. Learning algorithms—also known as learners—are algorithms that make other algorithms. With machine learning, computers write their own programs, so we don't have to.²

* Leroy Sorenson Merrifield Research Professor of Law, The George Washington University Law School. Many thanks to Abhi Shelat, John Basl, Katie Anne Creel, and Christo Wilson for their help with navigating the AI literature and to Devin Sullivan for his excellent research assistance.

[†] <https://doi.org/10.1093/ajcl/avac012>

1. Pedro Domingos, *CSEP 546: Data Mining; Machine Learning, Lecture 1*, YOUTUBE (Mar. 29, 2016), <https://www.youtube.com/watch?v=LnlW9gdjWfc&t=7528s>.

2. PEDRO DOMINGOS, *THE MASTER ALGORITHM* 6 (2015).

© The Author(s) [2022]. Published by Oxford University Press on behalf of the American Society of Comparative Law. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.

The difficulty with machine learning, at least for the law, is that the actual content of that algorithm is often not fully known or knowable to the humans operating the code, not even to the scientific expert:

State-of-the-art machine learning deploys far more complex models to learn about the relationship across hundreds or even thousands of variables. Model complexity can make it difficult to isolate the contribution of any particular variable to the result. . . [R]elatedly, the machine learning outputs are often *nonintuitive*—that is, they operated according to rules that are so complex, multi-faceted, and interrelated that they defy practical inspection, do not comport with any practical human belief about how the world works, or simply lie beyond human-scale reasoning. Even if data scientist can spell out the embedded rule, such rules may not tell a coherent story about the world as humans understand it, defeating conventional modes of explanation.³

Only recently has the law sought to address the new scientific and human reality. The novelty that the emerging legal frameworks seek to capture is the ability of machines to do what only a few decades ago most people thought only humans could do. So far, in the United States, no one legal definition of the technology and the policy problem has emerged as dominant. As the Administrative Conference of the United States puts it:

There is no universally accepted definition of “artificial intelligence,” and the rapid state of evolution in the field, as well as the proliferation of use cases, makes coalescing around any such definition difficult. . . Generally speaking, AI systems tend to have characteristics such as the ability to learn to solve complex problems, make predictions, or undertake tasks that heretofore have relied on human decision making or intervention. There are many illustrative examples of AI that can help frame the issue for the purpose of this Statement. They include, but are not limited to, AI assistants, computer vision systems, biomedical research, unmanned vehicle systems, advanced game-playing software, and facial recognition systems as well as application of AI in both information technology and operational technology.⁴

There are a couple of definitions that have gained currency in the law. The first is contained in the John S. McCain National Defense Authorization Act for Fiscal Year 2019,⁵ and is used in the

3. DAVID FREEMAN ENGSTROM ET AL., GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES: REPORT SUBMITTED TO THE ADMINISTRATIVE CONFERENCE OF THE UNITED STATES 11 (2020) [hereinafter ACUS REPORT].

4. Admin. Conf. of the U.S. Statement #20, Agency Use of Artificial Intelligence, 86 Fed. Reg. 6615, at 6615 n. 1 (Jan. 22, 2021).

5. John S. McCain National Defense Authorization Act of Fiscal Year 2019, Pub. L. No. 115-232, § 238(g), 132 Stat. 1636, 1697–98 (2018).

AI in Government Act of 2020,⁶ Executive Order 13960,⁷ and Office of Management and Budget, M-21-06.⁸ It reads as follows:

- (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
- (2) An artificial system developed in computer software, physical hardware, or another context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
- (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
- (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.
- (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

The second definition is contained in the National Artificial Intelligence Act of 2020. It reads:

The term artificial intelligence means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to—(A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.⁹

Lastly, there are definitions that have been formulated as part of sector-specific proposed legislation, for instance bills on facial recognition, driverless cars, and social media. One example is the proposed Algorithmic Justice and Online Platform Transparency Act,¹⁰ which is targeted at commercial uses of algorithms on online platforms. Section 3 defines “Algorithmic process” as:

a computational process, including one derived from machine learning or other artificial intelligence techniques, that

6. AI in Government Act of 2020, H.R. 2575, 116th Cong. (Sept. 15, 2020).

7. Exec. Order No. 13,960, 85 Fed. Reg. 78939, at 78942 (Dec. 3, 2020).

8. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-21-06, at 1 n. 2 (Nov. 17, 2020).

9. Division E of the National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5002, 134 Stat. 3388 (2021).

10. Algorithmic Justice and Online Platform Transparency Act, H.R. 3611, 117th Cong. (May 31, 2021).

processes personal information or other data for the purpose of determining the order or manner that a set of information is provided, recommended to, or withheld from a user of an online platform, including the provision of commercial content, the display of social media posts, or any other method of automated decision making, content selection, or content amplification.

In keeping with the state of terminological flux, this Report uses the terms AI and algorithms interchangeably. Where, as is often the case, the law finds its origins in the 1970s and earlier computer practices of public administration, algorithm can refer to either old style computer programming or machine learning. There are a couple of other preliminaries to keep in mind before turning to the questionnaire. Unless otherwise stated, the discussion below refers to federal public administration, not to the state law governing the operation of the administrative agencies of the fifty states. Last, following conventional practice, the discussion of administration excludes national security agencies and defense agencies because of the different concerns and legal frameworks that apply in those domains.

I. CONTEXT OF ALGORITHMIC ACCOUNTABILITY IN PUBLIC ADMINISTRATION

A. *The Extent of Adoption of Algorithmic Decision Making in Public Administration*

- 1) How has the rate of adoption of public algorithms evolved? How does it compare to that of private algorithms? Of note, a public algorithm is an algorithm which is used by public administration whether it was internally or externally designed.

Computerized data and algorithms have been widespread in public administration since at least the early 1970s. At that time the U.S. Privacy Act (1974) was adopted to combat the potential misuses of personal data contained in the computer systems of federal agencies.¹¹ However, the complex algorithms and big data associated with AI are generally traced to the 1990s.¹² In public administration, some of earliest uses were post-9/11 data mining for purposes of crime prevention and counterterrorism. In the present day, machine learning in the federal administration is widespread. A study conducted for the period from January to August 2019 found that of the 142 government agencies surveyed, nearly half (sixty-four agencies, or forty-five percent) had experimented with machine learning technology to

11. The Privacy Act of 1974, 5 U.S.C. § 552a. (2018).

12. See NAT'L SCI. & TECH. COUNCIL, THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC DEVELOPMENT PLAN 5 (Oct. 2016).

carry out a core function.¹³ Many of these agencies had used machine learning for more than one purpose, resulting in 157 use cases across the sixty-four agencies.¹⁴ Of these 157 use cases, one-third (fifty-three or thirty-three percent) were fully deployed, while the others were still in the “planning” phase or being “piloted or partially deployed.”¹⁵

- 2) Among all the public services, which ones use algorithms the most in their decision making?

At the federal level, law enforcement stands out as the heaviest user of machine learning. Other areas that figure prominently are health, financial regulation, social welfare, and commerce.¹⁶

- 3) Are algorithms evenly used across central, subfederal or member-state, regional, or local governments?

There are no comprehensive data on the use of AI at the state and local level. However, one area that figures prominently in the media is policing and criminal justice, which is predominantly a state and local function in the United States. For instance, PredPol is a program used by local police departments to predict crime hot spots and to devote police resources to those neighborhoods.¹⁷ The criminal justice system uses algorithms that predict recidivism. Based on a convicted criminal’s answers to a questionnaire as well as demographic and other types of data, the model assigns a risk score to the individual. Depending on the jurisdiction, that score can be used to inform a judge’s decision at the sentencing phase or, later, after the person has begun serving their sentence, a correction department’s decision on whether to grant parole.¹⁸

- 4) Are there public policy areas which particularly serve to drive standards in public algorithms?

Because of their powers, the police’s use of algorithms has generated considerable backlash and has driven calls for greater accountability and oversight. At all levels of government, law enforcement is a big consumer of the AI technologies that have been developed by the tech industry. Law enforcement agencies have subscriptions to the various big data consolidators and to commercial vendors of risk prediction models. One example that came to light in 2021 was the use by Immigration and Customs Enforcement (ICE) of private utilities data to detect individuals with irregular immigration status.¹⁹ Located in

13. ACUS REPORT, *supra* note 3, at 15.

14. *Id.* at 16.

15. *Id.* at 18.

16. *Id.* at 16.

17. CATHY O’NEIL, WEAPONS OF MATH DESTRUCTION 85 (2016).

18. *Id.* at 25; State v. Loomis, 371 Wis. 2d 235, 245 (Wis. 2016).

19. Drew Harwell, *Utility Giants Agree to No Longer Allow Sensitive Records to Be Shared With ICE*, WASH. POST. (Dec. 8, 2021), <https://wapo.st/37loQ9P>.

the Department of Homeland Security, ICE has civil and criminal enforcement powers over illegal immigration and transnational crimes. Utilities companies sell so-called “header data” on their consumers to Equifax, one of the major credit rating agencies and a big data consolidator. This header data includes the names, home addresses, Social Security numbers, and other details of roughly 170 million people. Equifax then turns around and sells that information to various types of tech companies, including Thomas Reuter’s CLEAR, which is a subscription service available to private investigators, government agencies, and the police. ICE has used the utilities data, easily searchable through its CLEAR subscription, to track potential immigration offenders. When this data and algorithm-fueled enforcement practice came to light, the private utilities companies and Equifax came under pressure from Senator Wyden in the U.S. Congress, and they have since stopped selling the data.

Facial recognition technology is another example of how law enforcement has driven public debate on algorithms—even though there has not yet been a legislative response, at least at the federal level. The big data of images associated with names, available through social media sites like Facebook, together with machine learning algorithms, have stimulated the production of facial recognition technology. This technology can be used for multiple purposes, including the identification of individuals of interest to the police. Many police forces have subscriptions that enable them to use an app to run photographs through a private vendor’s facial recognition service.²⁰ The technology, however, has been shown to be less accurate at identifying the faces of people of color as compared to whites. Moreover, regardless of its accuracy, facial recognition has obvious implications for basic human rights in a liberal democratic society.

At the federal level, the response to the police’s use of facial recognition has been numerous congressional hearings and a bill proposing a moratorium but so far no concrete action has been taken.²¹ At the state level there are also a number of bills pending, and in Washington State, a law backed by Microsoft has been adopted, which allows public agencies to use facial recognition but subject to certain limitations and regulatory requirements.²² Moreover, at the local level, San Francisco, Cambridge (Massachusetts), and a couple of other cities have banned the use of facial recognition by the police and other public authorities.

20. Kashmir Hill, *The Secretive Company That Might End Privacy as We Know It*, N.Y. TIMES, Jan. 18, 2020, <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>.

21. See Facial Recognition and Biometric Technology Moratorium Act of 2021, S. 2052, 117th Cong. (2021–2022), June 15, 2021.

22. 2020 Wash. Sess. Laws ch. 257 § 3(2) (“Facial Recognition—State and Local Government”).

B. Designing a National or a Multi-level Strategy for the Use of Algorithms in Administrative Decision Making

- 1) When was a national strategy on algorithmic decision making in public administration adopted?

The United States has been gradually adopting a national strategy on AI. The official sources setting down the principles are, in chronological order:

- “Maintaining American Leadership in AI,” Executive Order 13859 of February 11, 2019, 84 Fed. Reg. 3967;
- Office of Management and Budget, M-21-06, “Guidance for Regulation of Artificial Intelligence Applications,” Nov. 17, 2020;
- “Promoting Use of Trustworthy AI in Federal Government,” Executive Order 13960 of December 3, 2020, 85 Fed. Reg. 78939;
- AI in Government Act of 2020 (Division U, Title I), H.R. 133, January 3, 2020;
- National Artificial Intelligence Initiative Act of 2020 (Division E, Sec. 5001), January 1, 2021.

To date, Executive Orders 13859 and 13960 are the only two legal sources that aim to set down comprehensive standards for AI. Roughly speaking, the two Executive Orders divide the universe between AI in the industry and research sectors and AI in the public administration sector. The first, Executive Order 13859, is directed at government funding and regulation designed to promote AI throughout the economy and in the national security and defense establishment. The second, Executive Order 13960, is squarely focused on the theme of this Report. It is directed at internal government uses of AI that can improve public administration’s execution of their existing functions. It covers all federal government agencies, with the exception of national security and defense.

OMB Memorandum M-21-06 provides guidance on implementation of Executive Order 13859. It specifically says that the subject matter of Executive 13960 is outside its scope. In principle, therefore, OMB’s guidance on AI law is outside the scope of this Report. However, the distinction between the two types of AI use and development is not always clear, and where relevant, this Report will also discuss the standards set down in Memorandum M-21-06.

The two congressional laws are structural and institutional in nature. The first, the AI in Government Act of 2020, establishes a program in the General Services Administration called the AI Center of Excellence, which is tasked with improving adoption of AI in the federal government. The second, the National Artificial Intelligence Initiative Act of 2020, creates a coordinated federal infrastructure for AI in all its various dimensions. The centerpiece of this infrastructure

is the National Artificial Intelligence Initiative Office, part of the White House's Office of Science and Technology Policy.

- 2) Is the national strategy on algorithmic decision making closely related to an open data policy?

Yes. Executive Order 13859 sets out, among its objectives: "Enhance access to high-quality and fully traceable Federal data, models, and computing resources to increase the value of such resources for AI R&D, while maintaining safety, security, privacy, and confidentiality protections consistent with applicable laws and policies."²³

Among the important documents listed on the website of the National Artificial Intelligence Initiative Office are the 2020 and 2021 Federal Data Strategy Action Plans. These are designed to implement OMB's Federal Data Strategy, which promotes open data practices while at the same time respecting privacy and other accountability principles.²⁴

- 3) Under the national strategy, what types of algorithmic systems are promoted: basic or advanced (machine learning and deep learning algorithms); decision making aid or decision-making algorithms; or a combination of these?

Any form of AI that will promote competitiveness, innovation, and efficiency at the same time as respecting various good governance principles is promoted.

- 4) What reasons are mainly given for the adoption of algorithms in government?

The most comprehensive, official statement of the reasons for adopting AI in public administration is given in Executive Order 13960. It says the following:

Section. 1 Purpose.

Artificial intelligence (AI) promises to drive the growth of the United States economy and improve the quality of life of all Americans. In alignment with Executive Order 13859 of February 11, 2019 (Maintaining American Leadership in Artificial Intelligence), executive departments and agencies (agencies) have recognized the power of AI to improve their operations, processes, and procedures; meet strategic goals; reduce costs; enhance oversight of the use of taxpayer funds; increase efficiency and mission effectiveness; improve quality of services; improve safety; train workforces; and support decision making by the Federal workforce, among other positive

23. Exec. Order No. 13,859, 84 Fed. Reg. 3967, at 3968 (Feb. 11, 2019).

24. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-19-18, FEDERAL DATA STRATEGY – A FRAMEWORK FOR CONSISTENCE (June 4, 2019).

developments. Given the broad applicability of AI, nearly every agency and those served by those agencies can benefit from the appropriate use of AI.²⁵

5) Are nonpersonal or anonymized data likened to a common good, an essential facility?

The various documents that have been produced as part of the Federal Data Strategy repeatedly emphasize the importance of public use of government data that is respectful of privacy. However, the data is not specifically identified as nonpersonal or anonymized and it is not analogized to an essential facility.

6) Does the law provide for mandatory sharing of nonpersonal or anonymized data by public administration?

The OPEN Government Data Act, passed as part of the Foundations for Evidence-Based Policymaking Act of 2018, requires that agencies make data open by default.²⁶ As explained in a report on the implementation of the Act by the General Accountability Office:

Recognizing the need to make federal government data accessible and usable for the public, Congress passed and the President signed the Foundations for Evidence-Based Policymaking Act of 2018 (Evidence Act). Title II of the Evidence Act—the Open, Public, Electronic and Necessary Government Data Act of 2018 (OPEN Government Data Act)—requires federal agencies to publish their information as open data using standardized, nonproprietary formats, making data available to the public open by default, unless otherwise exempt. The act codifies and expands on existing federal open data policy including the Office of Management and Budget’s (OMB) memorandum M-13-13 (M-13-13), Open Data Policy—Managing Information as an Asset.

The OPEN Government Data Act requires that agencies develop and maintain comprehensive data inventories and collaborate with non-governmental entities, such as stakeholders, researchers, and the public, to understand how users value and use government data. The act also includes requirements related to agency strategic planning efforts, including a requirement that each agency have an open data plan that prioritizes data sets for disclosure on the Federal Data Catalogue.

25. Exec. Order No. 13,960, 85 Fed. Reg. 78,939, at 78,939 (Dec. 3, 2020).

26. The OPEN Government Data Act, Pub. L. No. 115-435, 132 Stat. 5529 (Jan. 14, 2019).

In addition to these agency requirements, the OPEN Government Data Act directs OMB to facilitate agency implementation by issuing guidance and reporting on agencies' progress in meeting their open data statutory requirements. OMB must also collaborate with the General Services Administration (GSA) to maintain a federal data catalogue of agency inventories and collaborate with GSA and the National Archives and Records Administration (NARA) to develop a repository of tools and resources to support agencies' open data efforts.²⁷

- 7) Has government-industry partnership or collaboration been implemented to accelerate artificial intelligence adoption in government agencies?

There are a number of avenues for introducing AI technologies in government agencies. In the examples of law enforcement given earlier, the technology is produced and maintained entirely by the tech industry, and it is purchased by agencies through the government procurement process. In other cases, it has been developed in-house. In yet other cases, it is the product of collaboration between government agencies and the university sector. To give a rough idea of the breakdown among these different modes of producing AI, the ACUS Report is instructive. Of the 157 federal use cases that were identified, fifty-three percent of the AI applications were developed by in-house agency technologists; thirty-three percent were purchased from the private sector through the government procurement process; fourteen percent were the product of noncommercial collaborations, including partnerships with research universities and agency-hosted competitions.²⁸

- 8) Were institutions created to put algorithms into administrative practice?

There are many offices in the federal government dedicated to developing AI, both within public administration and in public policy aimed at the private and university sectors.

- The National Artificial Intelligence Initiative Office (NAIIO) is located in the White House Office of Science and Technology Policy. Its role is horizontal, spanning the entire federal administration. NAIIO provides interagency coordination on AI and serves as a central point of contact for agencies, technical experts, and the public.

27. U.S. GOV'T ACCOUNTABILITY OFF., GAO-21-29, OPEN DATA: AGENCIES NEED GUIDANCE TO ESTABLISH COMPREHENSIVE DATA INVENTORIES; INFORMATION ON THEIR PROGRESS IS LIMITED 1-3 (Oct. 2020).

28. ACUS REPORT, *supra* note 3, at 88.

- The General Services Administration has an AI Center of Excellence which promotes the adoption of AI technologies across the federal government.
 - The National Institute of Standards and Technology has research, standards, and policy stream devoted to AI.
 - There are also offices within specific government agencies. The following is an illustrative list: Joint AI Center (Department of Defense); Artificial Intelligence and Technology Office (Department Energy); Office of Data Science Strategy (National Institutes of Health); National Artificial Intelligence Institute (Veterans Administration).
- 9) Are there calls for the creation of an algorithm regulatory agency with authority over both public and private algorithms?

No. There are a couple of proposals in Congress to create a U.S. privacy agency, which would have jurisdiction over the big (personal) data necessary for algorithms, but there are currently no proposals for a freestanding algorithm regulatory agency. In many of the algorithmic accountability bills, the Federal Trade Commission is given regulatory authority.²⁹

II. GENERAL CHARACTERISTICS OF THE LAW OF ALGORITHMIC ACCOUNTABILITY IN PUBLIC ADMINISTRATION

A. *Definition of Algorithmic Accountability in Public Administration*

- 1) What does algorithmic accountability in public decision making mean under the law?

The federal legal sources identified at the outset of this Report do not contain a definition of algorithmic accountability. It is worthwhile, therefore, looking to state law for insights into potential regulatory trajectories. Although there are no state laws applicable to algorithms generally speaking, Washington State has recently passed a law on facial recognition.³⁰ It contains three important accountability features: whenever a state or local government agency intends to use facial recognition technology it must file an accountability report with the legislative authority;³¹ the agency must solicit comments and hold consultation meetings on the draft accountability report;³² and decisions based on facial recognition that produce legal effects must be subject to meaningful human review.³³

29. See, e.g., Algorithmic Accountability Act of 2019, S. 1108, 116th Cong. (2019); Online Privacy Act of 2019, H.R. 4978, 116th Cong. (2019–2020).

30. 2020 Wash. Sess. Laws ch. 257 (State and Local Government).

31. *Id.* §3(2).

32. *Id.* §3(3).

33. *Id.* §4.

A closer look at the required accountability report sheds further light on the concept of algorithmic accountability. Overall, the accountability report is designed to guarantee transparency of the algorithmic system. Many of the specific elements are targeted at the data protection concerns that are implicated by algorithms: purpose, use, and sharing limitations; data security; and data breach reporting. However, these requirements are applied not only to the data inputs, but also to the algorithm's outputs. Moreover, data protection's standard of accuracy is applied and tailored to algorithms more broadly speaking. Thus, the government agency must report the results of testing of "the facial recognition service in operational conditions."³⁴ It must also report on facial recognition error rates. Last, data protection's civil liberty guarantees for special types of data is applied to the algorithm's potential "disparate impact on marginalized communities."³⁵

B. Scope of Algorithmic Accountability in Public Administration

- 1) Does the law feature rules governing both the use of algorithms and the design of algorithms implemented in public administration?

Executive Order 13960 sets down the "Principles for Use of AI in Government."³⁶ It shall be: lawful and respectful of our nation's values; purposeful and performance-driven; safe, secure, and resilient; understandable; responsible and traceable; regularly monitored; transparent; and accountable. These Principles apply to the entire life cycle of algorithms—"[w]hen designing, developing, acquiring, and using AI in the Federal Government." The principles for governmental AI are set down in section 3 of Executive Order 13960.

- 2) Does the law impose specific requirements on public algorithms and private algorithms?

As explained earlier, there are separate Executive Orders for the use of algorithms by government agencies (Executive Order 13960) and for the regulation of industry and research-sector algorithms (Executive Order 13859). The principles contained in the two Executive Orders overlap in large measure, especially when the guidance set out in OMB Memorandum M-21-06 (which applies specifically to the regulation of industry and research-sector AI) is taken into account.

- 3) Does the law distinguish between significant public algorithms and minor public algorithms?

No.

34. *Id.* §5.

35. *Id.* §3(2)(g).

36. Exec. Order No. 13,960, 85 Fed. Reg. 78,939, at 78,490 (Dec. 3, 2020).

4) Is the applicability of algorithmic treatment in a given decision-making process determined in light of the degree of complexity of the decisions involved? That is, are only low-complexity decisions subject to algorithmic treatment while medium and high-complexity decisions are left to human discretion?

No.

5) Does the law cover both individualized decision making (adjudication) and rulemaking?

Yes.

C. Striking Development Features of the Law of Algorithmic Accountability in Public Administration

1) Is the law of algorithmic accountability in public decision making couched in a general statute or regulation; sector-specific statutes or regulations; or both? Does it consist of a few rules?

Executive Order 13960 is general in application.

2) Is this law rather hard law or soft law?

Like all executive orders, Executive Order 13960 is hard law that is binding on the federal administration. Executive Orders remain in force in successive presidential administrations, but they can be repealed at will by the President.

3) What are the varying degrees of importance of international law, subfederal law, regional law, local law, and caselaw in this law?

At the international level, the United States adheres to OECD Recommendation on AI (2019). Executive Order 13960 is not enforceable in federal court and therefore case law cannot be expected to develop under Executive Order 13960.

4) Does this law feature a specific legal regime for personal data and pseudonymized data? Does it distinguish between general personal data and specific, sensitive personal data?

No. However, there are general privacy laws that are potentially applicable. The Privacy Act of 1974 applies when federal agencies establish a “system of records” involving personal information.³⁷ Under the Privacy Act, there is special treatment for the sensitive category of data on how individuals exercise their First Amendment

37. The Privacy Act of 1974, 5 U.S.C. § 552a(a)(3) (2018).

rights.³⁸ Federal agencies also have a duty to publish privacy impact assessments when they employ systems that use “personally identifiable information” (PII).³⁹

5) Do personal data include juridical persons’ data?

No. Neither the E-Government Act of 2002, section 208, nor the Privacy Act of 1974 give rights to corporations or other types of juridical persons.

6) Does the law provide for the use of autonomous drones and automated biometric identification in police and emergency/first responders operations?

Aerial surveillance by drones is common in federal and state law enforcement. At the federal level, the Federal Aviation Authority exercises regulatory authority.⁴⁰ To date, the Supreme Court has not considered the issue of whether drone surveillance is covered by the Fourth Amendment. In a series of cases decided in the 1980s, the Court found that there was no reasonable expectation of privacy triggered by conventional aerial surveillance and therefore the Fourth Amendment did not apply.⁴¹ However, the pervasive and constant character of drone surveillance is significantly different from earlier forms of aerial surveillance, and it is unclear how the Court will rule when presented with the issue.

At the state level, there are numerous states that regulate the police’s use of drones. For instance, Florida, Minnesota, and Virginia require that the police obtain a warrant before deploying drones.⁴²

Among biometric identifiers, facial recognition is one of the most used by police and immigration agencies.⁴³ In the case of the entry and exit data of noncitizens, collection is authorized by the Intelligence Reform and Terrorism Prevention Act of 2004.⁴⁴

7) What protocols have been imposed, notably in terms of transparency and participation, to combat bias and discrimination in the design process of algorithms?

38. *See id.* at 5 U.S.C. § 552a(3)(7); FRANCESCA BIGNAMI, COMMITTEE ON CIVIL LIBERTIES JUSTICE AND HOME AFFAIRS (LIBE) EUROPEAN PARLIAMENT, *THE US LEGAL SYSTEM ON DATA PROTECTION IN THE FIELD OF LAW ENFORCEMENT SAFEGUARDS, RIGHTS, AND REMEDIES FOR EU CITIZENS* 10–14 (2015).

39. *See* The E-Government Act of 2002, Pub. L. No. 107-347, § 208, 116 Stat. 2899 (2003).

40. FEDERAL AVIATION ADMIN., *DRONE RESPONSE PLAYBOOK FOR PUBLIC SAFETY*, www.faa.gov/uas/public_safety_gov (last visited Sept. 2020).

41. *California v. Ciraolo*, 476 U.S. 207 (1986); *Florida v. Riley*, 488 U.S. 445 (1989); *Dow Chemical v. EPA*, 476 U.S. 227 (1986).

42. FLA. CRIMINAL CODE § 934.50; MINN. STAT. CH. 82—S.F.No. 3072; VA. CODE § 19.2-60.1.

43. ACUS REPORT, *supra* note 3, at 30.

44. The Intelligence Reform and Terrorism Prevention Act of 2004, Pub. L. No. 108-458, 118 Stat. 3638, 3817 (2004).

OMB Memorandum M-21-06 provides guidance on bias and discrimination. As explained earlier, Memorandum M-21-06 applies specifically to government regulation of industry and research-sector AI. However, it may very well influence also the approach to internal, government AI. It says:

Fairness and Non-Discrimination

Agencies should consider in a transparent manner the impacts that AI applications may have on discrimination. AI applications have the potential of reducing present-day discrimination caused by human subjectivity. At the same time, applications can, in some instances, introduce real-world bias that produces discriminatory outcomes or decisions that undermine public trust and confidence in AI or be used in other ways that violate antidiscrimination statutes. When considering regulations or non-regulatory approaches related to AI applications, agencies should consider, in accordance with law, issues of fairness and nondiscrimination with respect to outcomes and decisions produced by the AI application at issue, as well as whether the AI application at issue may reduce levels of unlawful, unfair, or otherwise unintended discrimination as compared to existing processes.⁴⁵

D. Algorithmic Automation and Human Discretionary Power

- 1) Is personal data processing the only triggering factor for the regulation of automated administrative decision making?

Under federal law, there are no triggering factors specific to AI. Rather, AI is analyzed as an element of the ordinary forms of administrative action that can trigger legal requirements under general administrative law and constitutional law. An agency's use of AI may be akin to the adoption of a federal rule, in which case the rulemaking requirements of the federal Administrative Procedure Act (APA) may apply. In the case of individualized determinations, the use of AI may be subject to constitutional due process requirements or APA adjudication requirements. Moreover, any type of determination resulting from an AI system could be challenged as a violation of the APA's arbitrary and capricious standard.⁴⁶

- 2) Does the law recognize a right to a human in the administrative decision-making process? What requirements must be met to avail oneself of the right? Does this right apply from the initial decision or only at the agency appeal level?

45. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-21-06, at 6 (Nov. 17, 2020).

46. See Admin. Conf. of the U.S. Statement #20, Agency Use of Artificial Intelligence, 86 Fed. Reg. 6615, at 6622–23, (Jan. 22, 2021).

Under federal law, there is no such right. Some commentators suggest that the constitutional right to due process includes the right to a human decisionmaker in federal administration, but the future of “robo-judges” is too remote for this claim to have made its way into the courts yet.⁴⁷

- 3) What internal or external oversight frameworks are provided to ensure a balance between human and machine? External oversight frameworks such as the one performed by the executive branch, the legislative branch, independent regulators, the public, and public interest or civil society organizations.

There are no oversight frameworks designed specifically to ensure a balance between human and machine. With respect to AI broadly speaking, as said earlier, there is external oversight by the National Artificial Intelligence Initiative Office, located in the White House’s Office of Science and Technology Policy, and the General Services Administration’s AI Center of Excellence. There are also all the usual administrative accountability bodies. The Administrative Conference of the United States explains how these various accountability bodies, including internal oversight, might be called upon to monitor public AI:

It is essential that agencies’ AI systems be subject to appropriate and regular oversight throughout their lifespans. There are two general categories of oversight: external and internal. Agencies’ mechanisms of internal oversight will be shaped by the demands of external oversight. Agencies should be cognizant of both forms of oversight in making decisions about their AI systems.

External oversight of agencies’ uses of AI systems can come from a variety of government sources, including inspectors general, externally facing ombuds, the Government Accountability Office, and Congress. In addition, because agencies’ uses of AI systems might lead to litigation in a number of circumstances, courts can also play an important role in external oversight. . . .

Agencies should establish a protocol for regularly evaluating AI systems throughout the systems’ lifespans. That is particularly true if a system or the circumstances in which it is deployed are liable to change over time. In these instances, review and explanation of the system’s functioning at one stage of development or use may become outdated due to changes in the system’s underlying models. To enable that type of oversight, agencies should monitor and keep track of the data being used by their AI systems, as well as how the

47. *ACUS Report*, *supra* note 3, at 84.

systems use those data. Agencies may also wish to secure input from members of the public or private evaluators to improve the likelihood that they will identify defects in their AI systems.⁴⁸

III. ALGORITHMIC TRANSPARENCY AGENCY IN PUBLIC ADMINISTRATION

A. *Definition of Algorithmic Transparency*

- 1) What is the definition of transparency in the context of algorithmic accountability in public decision making?

Executive Order 13960 contains several principles for the use of AI in government, at least three of which are critical to transparency: AI should be understandable; the various elements of AI systems should be clearly identified and traceable; and information on AI use should be made available to stakeholders.⁴⁹ For purposes of clarity, these three principles (sections 3(e), 3(f), 3(h)) are reproduced in full below:

(e) Understandable. Agencies shall ensure that the operations and outcomes of their AI applications are sufficiently understandable by subject matter experts, users, and others, as appropriate.

(f) Responsible and traceable. Agencies shall ensure that human roles and responsibilities are clearly defined, understood, and appropriately as-signed-for the design, development, acquisition, and use of AI. Agencies shall ensure that AI is used in a manner consistent with these Principles and the purposes for which each use of AI is intended. The design, development, acquisition, and use of AI, as well as relevant inputs and outputs of particular AI applications, should be well documented and traceable, as appropriate and to the extent practicable.

(h) Transparent. Agencies shall be transparent in disclosing relevant information regarding their use of AI to appropriate stakeholders, including the Congress and the public, to the extent practicable and in accordance with applicable laws and policies, including with respect to the protection of privacy and of sensitive law enforcement, national security, and other protected information.

48. Admin. Conf. of the U.S. Statement #20, Agency Use of Artificial Intelligence, 86 Fed. Reg. 6615, at 6618 (Jan. 22, 2021).

49. Exec. Order No. 13,960, 85 Fed. Reg. 78,939, at 78,940 (Dec. 3, 2020).

In addition, OMB Memorandum M-21-06 provides guidance on transparency. As explained earlier, the Memorandum applies specifically to government regulation of industry and research-sector AI. However, there can be expected to be significant crossovers in the federal government's approach to AI externally and internally. On transparency, OMB emphasizes the importance of making AI applications understandable to both technical experts and the public that can be expected to be impacted by the AI. At the same time, it underscores competing considerations in favor of opacity, including harms from exploitation of the information by bad actors, the technical limits on explainability, and the benefits of the AI application. The relevant passage is reproduced in full below:

Disclosure and Transparency

In addition to improving the rulemaking process, transparency and disclosure can increase public trust and confidence in AI applications by allowing (a) non-experts to understand how an AI application works and (b) technical experts to understand the process by which AI made a given decision. Such disclosures, when required, should be written in a format that is easy for the public to understand and may include identifying when AI is in use, for instance, if appropriate for addressing questions about how the application impacts human end users. Disclosures may be required to preserve the ability of human end users and other members of the public to make informed decisions, although agencies should be aware that some applications of AI could improve or assist human decision making. Agencies should carefully consider the sufficiency of existing or evolving legal, policy, and regulatory environments before contemplating additional measures for disclosure and transparency. What constitutes appropriate disclosure and transparency is context-specific, depending on assessments of potential harms (including those resulting from the exploitation of disclosed information), the magnitude of those harms, the technical state of the art, and the potential benefits of the AI application.⁵⁰

- 2) Who is targeted to benefit from transparency: the public, courts, legislatures, or agency officials?

The transparency principles of Executive Order 13960 contemplate the public, Congress, and agency officials as the primary beneficiaries.

50. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-21-06, at 8 (Nov. 17, 2020).

B. *Transparency of the Design of Public Algorithm*

- 1) Is there a task force in charge of making recommendations for, or spearheading the deployment of, explicable public algorithms? What is the role of nongovernmental organizations such as think tanks or associations in developing transparency?

There are two main government bodies that conduct and fund research on AI generally speaking, including the issue of explainable algorithms. First, the National Science Foundation has established National AI Research Institutes across the country in partnership with research universities. Many of their projects involve the issue of explainability. Second, the National Institute for Standards and Technology (NIST) is responsible for fundamental research, development, and standards for AI technologies. AI Fundamental Research covers bias, explainability, and security, and therefore it can be expected to contribute to the explainability of public sector algorithms.

With respect to civil society, the numerous government offices with AI responsibilities regularly conduct requests for information and public consultations. For instance, the Department of Homeland Security has solicited public comments on facial recognition technologies, including concerns about bias, security, and privacy.⁵¹ There are also advisory committees comprised of members from the scientific community, the tech industry, federal laboratories, and the non-profit sector. The primary one is the National AI Advisory Committee (NAIAC).

- 2) Is the use of privately designed algorithmic systems forbidden in some areas to preserve explainability?

No.

- 3) What safeguards are provided to minimize the opacity claims which the private designers of algorithmic systems may derive from trade secret?

Intellectual property (IP) law applies to privately developed data and algorithms and it is acknowledged that IP can hinder algorithmic transparency.⁵² To the extent that the limitations on disclosure come into conflict with general administrative law duties of notice and reason-giving, public administration is either precluded from using the technology or it must negotiate waivers to IP rights in the procurement process.⁵³ In recent litigation, discussed below, a federal magistrate judge has ordered the private vendor to produce the system's

51. DEPT. HOMELAND SECURITY, PUBLIC PERCEPTIONS OF EMERGING TECHNOLOGIES, 86 Fed. Reg. 61285 (Nov. 5, 2021).

52. ACUS Report, *supra* note 3, at 77.

53. See generally Cary Coglianese, *Contracting for Algorithmic Accountability*, 6 Admin. L. Rev. 175 (2021).

training data and algorithms subject to a protective order in light of the trade secret concern.⁵⁴

4) During the design process of algorithmic systems, are private companies required to gather input from the public administration, which will put the algorithm to use?

No.

5) Is the launch of the design process of a public algorithmic system subject to notice to the public?

No.

C. *Transparency During and After the Algorithmic Decision Making*

1) Is the decision to deploy an algorithmic (whether individualized decision making or rulemaking) decision making subject to public notice?

There are no procedural rules specifically applicable to public administration's use of algorithms. However, if the algorithm is analogous to a rule, then the APA's requirements of informal rulemaking would apply—i.e., notice and comment procedure.⁵⁵ The APA defines a rule as “the whole or a part of an agency statement of general or particular applicability and future effect designed to implement, interpret, or prescribe law or policy or describing the organization, procedure, or practice requirements of any agency. . . .”⁵⁶ An algorithm that requires agency officials to take certain actions with respect to industry actors or other regulated parties would likely be considered a rule.

2) How does the decision-making context (individualized decision making or rulemaking) impact the degree of required transparency?

Administrative authorities are vested with different types of functions—among others, guaranteeing compliance with the regulatory duties set down in congressional law, distributing social security entitlements and other types of benefits created by law, and giving further guidance on the requirements set down in congressional law through binding rules or other types of broadly applicable administrative acts. These functions are generally referred to as enforcement, adjudication (individualized decision making), and rulemaking. The type of administrative function maps on to different types of procedural

54. *Flores v. Stanford*, 2021 WL 4441614 (S.D.N.Y. Sept. 28, 2021).

55. 5 U.S.C. § 553.

56. 5 U.S.C. § 551(4).

requirements, which also apply to AI and whether it is used in the context of enforcement, adjudication, or rulemaking.

With respect to agency enforcement, there are competing sets of considerations. On the one hand, enforcement decisions are typically treated with extreme deference by courts, meaning that agencies are often not required to give an explanation or if they are, the explanation required is minimal.⁵⁷ Moreover, full disclosure of the considerations used by agencies, including the algorithmic ones, can be used by industry to “game” the system—i.e., to avoid detection and therefore violate the law with impunity.⁵⁸ On the other hand, once the criteria for agency enforcement are set down in an algorithm, the algorithm could potentially be considered a rule, in which case it would be subject to APA informal rulemaking unless one of the exceptions apply.⁵⁹

In U.S. administrative law, the duty to disclose the facts and evidence in favor a proposed determination as well as the duty of reasoned explanation in support of the final determination varies as between adjudication and rules. The variation, however, turns more on the types of facts and arguments that must be disclosed than it does on the extent of the disclosure required.⁶⁰ While rules are associated with policy-based facts and arguments, adjudications are thought to turn on the specific circumstances of the person’s situation. AI carries the prospect of blurring the distinction between these two types of agency action by converting the fact-specific and human element of adjudications (individualized determinations) into the rule-like operation of the algorithm.

Although the future of “robo-judges” is still far off, there is the very real prospect that aspects of individualized determinations in social welfare programs and other areas of government intervention will be taken away from civil servants and will be made by algorithms. This is seen as a potentially beneficial development because it carries the prospect of greater efficiency, accuracy, and evenhanded treatment across the millions of individuals that come within the scope of government programs. As a matter of law, there are no general constitutional or administrative law objections since agencies are afforded great discretion in whether they choose to operate by rule or adjudication in undertaking their statutory mandates.⁶¹ At the same time, if algorithms come to occupy more space in adjudications, such as a person’s eligibility for benefits, then they will be considered rules and they will have to satisfy the notice-and-comment and reasoned

57. *Dunlop v. Bachowski*, 421 U.S. 560 (1975); *Heckler v. Chaney*, 470 U.S. 821 (1985).

58. ACUS REPORT, *supra* note 3, at 86.

59. *See generally* *Chamber of Commerce of the U.S. v. U.S. Dep’t of Labor*, 174 F.3d 206 (D.C. Cir. 1999).

60. *See generally* RONALD A. CASS ET AL., *ADMINISTRATIVE LAW: CASES AND MATERIALS* 417–23 (8th ed. 2020).

61. *Heckler v. Campbell*, 461 U.S. 458 (1983).

explanation (“concise general statement of basis and purpose”) requirements of APA informal rulemaking.⁶²

- 3) Is the use of algorithmic systems limited to simpler algorithmic systems to guarantee explainability and accuracy?

As of yet, no such requirement has developed in the law. With respect to agency practice, ACUS has acknowledged that there is a trade-off between explainability and accuracy:

Consideration of actual use cases reveals hard trade-offs between accountability and efficacy. Imposing constraints on model choices—by, for example, limiting the number of data features or prohibiting more sophisticated modeling approaches—trades off interpretability against a tool’s analytic power and, thus its usefulness. As just one example, requiring the SEC to deploy a less sophisticated but more interpretable algorithmic tool in making enforcement decisions may make it easier for regulated parties or agency overseers to evaluate the tool’s workings but may also bring substantial costs, subjecting regulated parties to undue prosecutions and wasting scarce agency resources in the process.⁶³

- 4) Does the law mandate algorithmic opacity in certain areas?

Executive Order 13960 makes exceptions to the transparency principle “with respect to the protection of privacy and of sensitive law enforcement, national security, and other protected information.”⁶⁴ In these areas, the subject-specific legislation prohibits disclosure.

- 5) Do public administrations have the duty to inform the person concerned by the individualized decision of the use and the functioning of the algorithm? What does the duty to inform entail? Is it a proactive duty, reactive duty, or both?

If the determination is subject to constitutional due process (as would be, for instance, a benefits determination or government employment), the individual has a right to notice detailing the reasons for the proposed determination. At least one federal court has held that this includes notice of the data and algorithm used to produce the adverse decision so that it is possible for the adversely affected party to correct an erroneous determination.⁶⁵ In addition, the agency’s final determination must be accompanied by an explanation of any adverse decision, which would include similar information.

62. ACUS REPORT, *supra* note 3, at 84.

63. *Id.* at 76.

64. Exec. Order No. 13,960, 85 Fed. Reg. 78,939, at 78,490 (Dec. 3, 2020).

65. See *Hous. Fed’n of Teachers, Local 2415 v. Hous. Indep. Sch. Dist.*, 251 F. Supp. 3d 1168 (S.D. Tex. 2017) (scoring algorithm resulting in termination of public-school teacher). See also *KW v. Armstrong*, 789 F.3d 962 (D. Idaho 2015) (statistical model resulting in reduction of benefits for developmentally disabled adults).

If the determination is subject to the requirements of APA formal adjudication, then a logic like that of the constitutional due process cases would apply. The agency might have a duty to give notice of the data and the algorithm in the hearing before the administrative law judge.⁶⁶ The administrative law judge's recommended decision might also need to afford an explanation of the algorithm as part of the "findings and conclusions, and the reasons or basis therefor."⁶⁷

If the determination is not considered APA formal adjudication but rather APA informal adjudication, then the agency's duties of notice-giving and explanation are minimal and would in all likelihood not include explanation of how an algorithm operates.⁶⁸ An example of informal adjudication that involves AI is the grant or denial of patents and trademarks by the U.S. Patent and Trademark Office (USPTO).⁶⁹ The USPTO uses machine learning to classify patent and trademark applications and to search for prior art or prior mark. These are critical initial steps in the USPTO's determination of whether a patent or trademark application is meritorious. The demands of informal adjudication are so minimal that the legally required statement of grounds for any denial of patent or trademark most likely does not include an explanation of these algorithms.⁷⁰

Finally, it should be noted that the right of notice under constitutional due process and the APA turn on the specific circumstances of the case. The algorithm may be considered a rule that must satisfy the requirements of APA informal rulemaking, in which case the agency's duty to give notice of the reasons for the proposed determination and the duty to explain the final determination may have been satisfied in large part by the prior rulemaking proceeding.⁷¹ Moreover, algorithms may be used to facilitate administrative adjudication in various ways but they still may allocate ample discretion and decisional authority to agency adjudicators.⁷² In such instances, the case for disclosure is relatively weak since the algorithm is several steps removed from the final decision and the basis for the agency determination continues to rest principally with the human adjudicator and the facts and reasoning articulated by the human adjudicator.

All of the above duties (under the Constitution and the APA) are pro-active.

- 6) Are private juridical persons subject to a similar duty to inform the person concerned by an individualized decision?

66. 5 U.S.C. § 556(d).

67. 5 U.S.C. § 557 (c)(3)(A).

68. 5 U.S.C. § 555(e).

69. ACUS REPORT, *supra* note 3, at 46–52.

70. See Arti Rai, *Machine Learning at the Patent Office: Lessons for Patents and Administrative Law*, 104 IOWA L. REV. 2617 (2019).

71. See generally *Ark. Dep't of Human Servs. v. Ledgerwood*, 530 S.W.3d 336, 344–45 (Ark. 2017).

72. See ACUS REPORT, *supra* note 3 at 44–45.

There are no general legal duties applicable to private firms or other types of juridical persons. There are, however, sector-specific requirements. For instance, the credit-reporting industry, which is probably the most heavily regulated industry from this perspective, is required to make certain disclosures to consumers about their credit scores. Upon the request of a consumer, the consumer reporting firm must communicate their credit score, the range of possible credit scores under the model used, all of the key factors that adversely affected the credit score of the consumer in the model used (which shall not exceed four), and other information.⁷³

7) In the case of individualized decisions, is the duty to inform conditioned on a personal data processing requirement?

None of the notice or explanation requirements discussed above are linked to the Privacy Act or the E-Government Act and the privacy duties established under those statutes.

8) Are there public registers for public sector algorithmic systems?

Section 5 of Executive Order 13960 requires that each covered federal agency compile an inventory of its AI use cases. When compiling this inventory, agencies are to ensure that their AI use cases are consistent with the Principles set down earlier in the Order. They are also directed to make their inventories available to the public.

9) Is the source code of an algorithm communicable under the right to access administrative records?

There has been considerable litigation under the Freedom of Information Act on access to the algorithms used by federal agencies. In principle, the government duty of disclosure applies equally to computer code as it does to other types of information.⁷⁴ One example of how FOIA was successfully used to gain access to an agency algorithm is the Department of Homeland Security's Risk Classification Assessment (RCA), which uses algorithms to determine a migrant's flight risk and risk to public safety in order to determine whether the person should be subject to civil detention by Immigration and Customs Enforcement. A group of academics filed a number of FOIA request with DHS and then spent almost three years in the agency appeals process and in litigation in the federal courts.⁷⁵ Among other things, they ultimately obtained

a summary of 1.4 million RCA entries with limited data fields, and a total of 2500 RCA entries with expanded data

73. 15 USCA § 1681g.

74. See Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1 (2019).

75. Kate Evans & Robert Koulish, *Manipulating Risk: Immigration Detention Through Automation*, 24 LEWIS & CLARK L. REV. 789 (2020).

fields. This production occurred over the course of 2017 and 2018, resulting in over 1700 pages, printouts of the RCA online training course for ICE officers, matrices of the RCA's algorithm and scoring rubric for evaluating risk and recommending detention including every change during the first five years.⁷⁶

However, it does not appear that the researchers obtained access to the RCA computer code. Further, there are many exceptions to FOIA's right to information, and they often used by agencies to deny access to AI algorithms.⁷⁷

D. *Non-algorithmic Transparency v. Algorithmic Transparency*

- 1) How does algorithmic transparency compare to and blend with nonalgorithmic transparency? Be it in individualized decision making or in rulemaking, consider the components of non-algorithmic transparency: the right to information, publication, the right to reasons.

As explained above, the lion's share of the law on algorithmic transparency derives from the general duties and rights of constitutional and administrative law.

IV. PARTICIPATION IN ALGORITHMIC DECISION MAKING IN PUBLIC ADMINISTRATION

A. *Definition of Algorithmic Participation*

- 1) What is the definition of participation in the context of algorithmic accountability in public decision making?

There does not exist a special legal definition of public participation in the context of algorithms. OMB Memorandum M-21-06, which as explained above applies to the regulation of industry and research-sector AI but can be expected to influence thinking on government uses of AI, underscores the importance of public participation in AI regulation to improve public trust and confidence. The relevant passage says in full:

Public Participation

In accordance with Executive Order 13563, "Improving Regulation and Regulatory Review," regulations "shall be adopted through a process that involves public participation." Public participation, especially in those instances where AI uses information about individuals, will improve agency

76. *Id.* at 854.

77. See Coglianese & Lehr, *supra* note 74, at 26.

accountability and regulatory outcomes, as well as increase public trust and confidence. Agencies must provide ample opportunities for the public to provide information and participate in all stages of the rulemaking process, to the extent feasible and consistent with legal requirements (including legal constraints on participation to, for example, protect national security and address imminent threats or respond to emergencies). Agencies are also encouraged, to the extent practicable, to inform the public and promote awareness and widespread availability of voluntary frameworks or standards and the creation of other informative documents.

B. *Participation in the Algorithm Design*

- 1) Is the design process of an algorithmic system subject to an impact assessment requirement in general or in some cases?

There are two types of impact assessments that potentially apply to AI design—Privacy Impact Assessment (PIA) and Regulatory Impact Assessment (RIA).

Federal agencies are required to conduct privacy impact assessments (PIA) before “developing or procuring information technology that collects, maintains, or disseminates information that is in an identifiable form,” information that has come to be referred to as personally identifiable information (PII).⁷⁸ They must also do so before initiating a new collection of information that “will be collected, maintained, or disseminated using information technology” and that includes PII that can be used for the “contacting of a specific individual. . . .”⁷⁹ Among other things, the PIA must include information on purposes and uses, notice and consent, and security. The big data that is used to develop algorithms might or might not trigger a PIA depending on what form it takes and whether it is considered PII.

Federal agencies are also required to conduct a regulatory impact analysis of new regulatory initiatives, which may, depending on the circumstances, include their use of AI. Under Executive Order 12,866, “significant regulatory actions” must be notified to the White House’s Office of Information and Regulatory Affairs (OIRA) in the Office of Management and Budget and must be accompanied by a regulatory impact analysis. What counts as “significant regulatory actions” is defined broadly based on a variety of characteristics, including for instance whether the actions “[m]aterially alter the budgetary impact of entitlements, grants, user fees, or loan programs or the rights and obligations of recipients”; or “[r]aise novel legal or policy issues arising

78. E-Government Act § 208 (b)(1)(A)(i).

79. *Id.* § 208 (b)(1)(A)(ii).

out of legal mandates.”⁸⁰ As one textbook puts it: “practically speaking, the only limit on the kinds of regulatory actions reviewable by OIRA is whether the Director finds it sufficiently important to merit the attention of the White House—though depending on the budget and staff available for review, that could be a very significant constraint.”⁸¹ There are two types of RIA. The first is a relatively simple assessment of costs and benefits; the second, reserved for economically significant regulations, is a comprehensive assessment of costs and benefits, including the underlying data, sources, and methods, quantification of costs and benefits, and information on regulatory alternatives. Many government uses of AI are likely to trigger regulatory review.

OMB Memorandum M-21-06 provides guidance on RIA with respect to regulation of nongovernmental AI and it may be illustrative of RIA of governmental AI. It articulates the standard elements of an RIA and includes the privacy, equal protection, and other concerns distinctive to AI:

After identifying a set of potential regulatory approaches, the agency should conduct a benefit-cost analysis that estimates the benefits and costs associated with each alternative approach. The benefits and costs should be quantified and monetized to the extent possible and appropriate, and presented in both physical units (*e.g.*, number of accidents avoided) and monetary terms. When quantification of a particular benefit or cost is not possible, it should be described qualitatively. The analysis of these alternatives should also evaluate, where relevant and appropriate and consistent with Executive Order 13859, impacts to equity, human dignity, fairness, potential distributive impacts, privacy and civil liberties, personal freedom, and other American values. The agency’s analysis should be based on the best available scientific, technical, and economic information.⁸²

Memorandum M-21-06 also articulates the various risks associated with AI, which are to be taken into account in AI adoption and regulation and also, presumably, in RIA. While some point in favor of limiting AI applications, guaranteeing greater transparency and participation, or both, others point in the opposite direction:

Assessing Risk

When humans delegate decision-making and other functions to AI applications, there is a risk that AI’s pursuit of its defined goals may diverge from the underlying or original

80. Exec. Order No. 12,866, 58 Fed. Reg. 51,734, at 51,738 (Oct. 4, 1993).

81. RONALD A. CASS ET AL., ADMINISTRATIVE LAW: CASES AND MATERIALS 603 (2020).

82. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-21-06, app. A at 13 (Nov. 17, 2020).

human intent and cause unintended consequences—including those that negatively impact privacy, civil rights, civil liberties, confidentiality, security, and safety. Because traditional forms of delegated decision-making are accompanied by risks that present some—although not all—of the dynamics present in the case of delegation to AI, existing approaches to risk continue to be relevant. However, the kind of AI adopted and the way it works in decision-making may present new demands on existing risk frameworks. In addition, because components of AI applications, such as algorithms or the data they are trained on and use, may be sensitive or subject to legal protections (e.g., privacy or intellectual property), agencies should consider the risks of inadequate protections to algorithms and data throughout the design, development, deployment, and operation of an AI system, given the level of sensitivity of the algorithms and data. Agencies should also consider that an AI application could be deployed in a manner that yields anticompetitive effects that favors incumbents at the expense of new market entrants, competitors, or up-stream or down-stream business partners.

Managing Risk

The management of risks created by AI applications should be appropriate to, and commensurate with, the degree of risk that an agency determines in its assessment. In general, as emphasized above, the agencies should also be comparing risks unique to the AI application to other similar risks associated with not using such applications within a regulatory framework or risks mitigated by the adoption of AL. For AI applications, agencies should adopt a tiered approach in which the degree of risk and consequences of both success and failure of the technology determines the regulatory approach, including the option of not regulating. Agencies should be aware that there is always likely to be at least some risk, including that associated with not knowing what is currently unknown. For AI applications that pose lower risks, agencies can rely on less stringent and burdensome regulatory approaches—or non-regulatory approaches—such as requiring information disclosures or consumer education. For higher risk AI applications, agencies should consider, for example, the effect on individuals, the environments in which the applications will be deployed, the necessity or availability of redundant or back-up systems, the system architecture or capability control methods available when an AI application makes an

error or fails, and how those errors and failures can be detected and remediated.⁸³

An example of how impact analysis, both Privacy Impact Analysis and Regulatory Impact Analysis, might in the future be applied to AI is the facial recognition technology used by the Customs and Border (an agency within Department of Homeland Security). Since 2004, CBP has been collecting and using biometric identifiers for purpose of border and immigration enforcement.⁸⁴ The most important type of identifier that has emerged over the decades is facial recognition technology. At the time that CBP first established the program and then subsequently, with each major expansion, it has been required to conduct both a Privacy Impact Assessment and, in conjunction with the rulemaking procedure, a Regulatory Impact Assessment.⁸⁵ At no point, however, does it appear that the AI technologies have been included in these assessments, perhaps because they are generally not produced in-house but rather are acquired from tech companies.⁸⁶ Expansion of the scope of PIA, RIA, or both is one possibility for improving AI accountability in the federal government.

- 2) Does the impact assessment include a public comment period or other type of input by independent internal experts, external experts, or consultation with a neocorporatist advisory body?

With respect to Privacy Impact Assessment, there is no public comment period. They are generally approved by the agency's Chief Information Officer, submitted to the White House's Office of Management and Budget, and made publicly available.⁸⁷

With respect to the regulatory review process and Regulatory Impact Assessment, as originally designed it was conceived as a tool of presidential oversight of federal administrative agencies. Public comment occurs through the parallel, and historically older, process of APA informal rulemaking. RIA and regulatory review by OIRA is a bureaucratic process without legally guaranteed opportunities for public involvement.

Even though it still is true that the rulemaking process is the main vehicle for public participation, OIRA has become more transparent over the past decades. As one author reports:

83. *Id.* app. A at 13–14.

84. See ACUS REPORT, *supra* note 3 at 31–32.

85. DEPT. OF HOMELAND SEC., PRIVACY IMPACT ASSESSMENT FOR THE AUTOMATED BIOMETRIC IDENTIFICATION SYSTEM (IDENT) (July 31, 2006), www.dhs.gov/sites/default/files/publications/privacy_pia_usvisit_ident_final.pdf; DHS/OBIM/PIA-001 AUTOMATED BIOMETRIC IDENTIFICATION SYSTEM (IDENT) (Dec. 7, 2012), www.dhs.gov/sites/default/files/publications/privacy-pia-nppd-ident-december2012.pdf. RIAs: 69 Fed. Reg. 468 (2004); 69 Fed. Reg. 533,318; 73 Fed. Reg. 77,473.

86. See ACUS REPORT, *supra* note 3, at 32.

87. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-03-22, OMB GUIDANCE FOR IMPLEMENTING THE PRIVACY PROVISIONS OF THE E-GOVERNMENT ACT OF 2002 (Sept. 26, 2003).

Once criticized as the regulatory black hole,” OIRA undertook a number of improvements during the Clinton and Obama Administrations and, in many ways, had become one of the more publicly engaged federal entities. OIRA discloses on its website when it has begun deliberating on a rule, it logs all meetings with nongovernmental entities, and it makes certain interagency communications with top-level officials available to the public.⁸⁸

Moreover, in some cases RIAs are released early in the rulemaking process, and therefore they can provide important technical background and regulatory information that can be used by stakeholders in the public comment period required in APA informal rulemaking.⁸⁹

3) What are the other characteristics of the impact assessment requirement?

The are no other significant characteristics.

C. *Artificial Intelligence as a Public Participation Tool*

1) Are algorithms used to analyze and respond to mass comments in rulemaking proceedings that are intrinsically nonalgorithmic or only slightly so?

A number of government agencies are experimenting with AI and machine learning algorithms to sort through the increasing volume of public comments being submitted in agency rulemaking proceedings. The aim is to identify the many duplicative or inauthentic, bot-generated comments and to classify comments by the types of concerns raised.⁹⁰

D. *Participation in the Algorithmic Decision Making*

1) Is the decision to deploy an algorithmic decision making subject to public input or to consultation or authorization in addition to the impact assessment or consultation required for the adoption of the algorithm?

As explained in response to the questions on impact assessment, the general rule of American administrative law is that APA rulemaking is the main vehicle for obtaining public input. This is specifically contemplated for algorithmic regulation in OMB Memorandum M-21-06:

88. JASON A. SCHWARTZ, INST. FOR POLICY INTEGRITY, ENHANCING THE SOCIAL BENEFITS OF REGULATORY REVIEW 25 (Oct. 2020).

89. *Id.* at 27.

90. See ACUS REPORT, *supra* note 3, at 60–63.

Public Consultation

The informal rulemaking process under the Administrative Procedure Act provides predictable and meaningful opportunities for interested stakeholders to provide input on draft regulations and scrutinize the evidence and analytic bases of regulatory proposals. In soliciting public input on Notices of Proposed Rulemaking (NPRMs) that relate to AI applications, agencies will benefit from the perspectives and expertise of stakeholders engaged in the design, development, deployment, operation, and impact of AI applications, and facilitate a decision making process that is more transparent and accountable.

To the extent feasible, agencies should also provide opportunities for stakeholder consultation before the NPRM stage, including through the issuance, when appropriate, of RFIs [requests for information] and Advance Notices of Proposed Rulemaking (ANPRMs) to inform decisions about the need to regulate. Agencies should also consider holding stakeholder and public meetings both prior to issuing an NPRM and during the public comment period.⁹¹

In the context of public administration, if the government's use of AI meets the definition of a "legislative rule," informal APA rulemaking is required. Although the definition of legislative rules in contrast to nonlegislative rules is notoriously slippery, the standard textbook analysis generally looks to the extent to which the rule constrains agency discretion—the more binding and the more narrowly formulated the agency's rule, the more likely it is to count as "legislative" and therefore require informal rulemaking.⁹²

Algorithms that result in benefits determinations or employment decisions clearly count as legislative rules.⁹³ Algorithms that influence agency enforcement priorities or that target certain types of regulatory violations might also very well be considered a legislative rule. Consider *Chamber of Commerce of the United States v. U.S. Department of Labor*.⁹⁴ There the Occupational Safety and Health Administration (OSHA) established a program to incentivize employer self-regulation: an employer that adopted a comprehensive safety and health program would face a lower probability of OSHA workplace

91. OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, OMB MEMORANDUM M-21-06, app. A at 12–13 (Nov. 17, 2020).

92. See generally JOHN F. MANNING & MATTHEW C. STEPHENSON, LEGISLATION AND REGULATION: CASES AND MATERIALS 935–84 (4th ed. 2021).

93. See generally Ark. Dep't of Hum. Servs. v. Ledgerwood, 530 S.W.3d 336, 344–45 (Ark. 2017) (computer algorithms allocating home-care hours to disabled low-income individuals).

94. Chamber of Commerce of the U.S. v. U.S. Dep't of Labor, 174 F.3d 206 (D.C. Cir. 1999).

inspections to detect regulatory violations (seventy to ninety percent lower). The program was considered binding and concrete enough to count as a legislative rule. The analogy between OSHA's enforcement regime and a more sophisticated algorithmic enforcement tool is not far-fetched.

- 2) How is the degree of algorithmic literacy factored in the participation and transparency regime?

Currently there is no attention in the federal legal sources to the challenges of making AI accessible to the lay public so as to guarantee real transparency and participation. The problem of long, dense, and technically forbidding explanations of proposed rules is a longstanding one in American administrative law.⁹⁵ So far there have not been any significant fixes to the problem, which is a general one but can be expected to be especially acute in the domain of AI.

E. Nonalgorithmic Participation v. Algorithmic Public Participation

- 1) How does algorithmic participation compare to and blend with nonalgorithmic participation in individualized decision making and in rulemaking? Consider the components of nonalgorithmic participation: the rights of the defense or the right to a fair trial and due process; the right to comment on draft decisions; indirect participation based on neo-corporatist consultation; and the right to petition.

As explained above, there does not exist a separate participation regime for algorithms. At present, the main legal opportunity afforded for participation is APA informal rulemaking.

V. JUDICIAL REVIEW OF ALGORITHMIC ACCOUNTABILITY IN PUBLIC ADMINISTRATION

- 1) What are the broad trends in, and the striking features of, the algorithmic accountability caselaw?

A survey of constitutional litigation involving machine learning algorithms has found that procedural due process challenges dominate as compared to equal protection (antidiscrimination and bias) challenges and privacy challenges.⁹⁶ Because of the focus on state action in U.S. constitutional law, all of the judicial opinions reviewed involved the public sector's use of algorithms.

95. See Wendy Wagner, *Participation in the U.S. Administrative Process*, in *COMPARATIVE LAW AND REGULATION: UNDERSTANDING THE GLOBAL REGULATORY PROCESS* 125 (Francesca Bignami & David Zaring eds., 2016).

96. Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 *CORNELL L. REV.* 1875, 1879, 1903 (2020).

The results of the Huq study are not surprising. The most prominent administrative due process cases have already been canvassed above, in the discussion of transparency. By contrast with procedural due process, neither equal protection nor privacy are promising grounds for mounting legal challenges to algorithmic governance. With respect to equal protection law, the Supreme Court has retreated from disparate impact theories under the Equal Protection Clause and requires a showing of discriminatory intent,⁹⁷ which does not translate well to the machine learning context. It is very difficult to impute intent to the process by which bias and unfair outputs can be generated by machine learning algorithms.

With respect to privacy, there is no robust right to data privacy in constitutional law.⁹⁸ Even the statutory law—the Privacy Act—does not afford obvious remedies for the potential privacy violations that lurk in the big data necessary for machine learning litigation.⁹⁹ Under the Privacy Act, individuals can sue the federal government for privacy violations involving their information only if that information is contained in a system of records, defined as a system from which the government agency retrieves information based on a personal identifier like a name or social security number. Much, if not most, of the data used in machine learning algorithms does not come from such systems of records. Moreover, even if the data did come, improperly, from such systems, individuals would have a very hard time showing Article III standing, since it would be virtually impossible to show the type of injury required under Article III.

- 2) In the absence of written law, do courts tend to apply the principle of algorithmic accountability and expand its scope in light of nonalgorithmic accountability?

At present, there is no such principle in U.S. law.

- 3) How do courts handle the explainability issue when confronted with a black box or machine learning algorithm? In such a case, do they tend to substitute reasonableness for explainability?

So far, there is not much experience in the courts with the black box issue. In *Houston Federation of Teachers*, which was already mentioned above, a Houston school district relied on a commercially developed and maintained algorithmic system to score teachers.¹⁰⁰ Teachers who performed poorly in the scoring system were fired. The local teacher's union and nine teachers challenged their terminations

97. Julie C. Suk, *Disparate Impact Abroad*, in *A NATION OF WIDENING OPPORTUNITIES: THE CIVIL RIGHTS ACT AT 50*, at 283 (Samuel R. Bagenstos & Ellen D. Katz eds., 2016).

98. Francesca Bignami & Giorgio Resta, *Transatlantic Privacy Regulation: Conflict and Cooperation*, 78 *LAW & CONTEMP. PROBS.* 231, 235 (2015).

99. BIGNAMI, *supra* note 38, at 10–14.

100. *Hous. Fed'n of Teachers, Local 2415 v. Hous. Indep. Sch. Dist.*, 251 F. Supp. 3d 1168 (S.D. Tex. 2017).

based on a number of grounds, including the procedural due process argument that they did not receive adequate pretermination notice of the grounds for termination, since they did not have access to the computer algorithm and data necessary to verify the accuracy of their scores. The district court agreed, saying that “without access to . . . proprietary information—the value-added equations, computer source codes, decision rules, and assumptions—[the teachers’] scores will remain a mysterious ‘black box,’ impervious to challenge.”¹⁰¹ However, the remedy was not disclosure, since the commercial algorithms and software were protected as trade secrets. Rather, the school district was precluded from relying on the algorithmic system for making termination decisions: “When a public agency adopts a policy of making high stakes employment decisions based on secret algorithms incompatible with minimum due process, the proper remedy is to overturn the policy while leaving the trade secrets intact.”¹⁰²

One area where there has been some doctrinal movement is with respect to recidivism risk assessments for purposes of criminal sentencing and parole determinations. Although the criminal justice system falls outside the ambit of this Report, this case law points to possible trends in the administrative domain. The Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) is a commercially developed and managed algorithmic system for assessing an individual’s risk of recidivism. It uses factors such as sex and age and data from the general population to develop an algorithmic system, which, based on the offender’s individual profile, generates a recidivism risk score for the offender. This risk score and report is used in a number of states to assist with criminal sentencing and parole determinations.

In *Wisconsin v. Loomis*, a defendant who was denied probation at sentencing in part because of his COMPAS risk score challenged the constitutionality of the system.¹⁰³ He argued that his right to due process was violated because he did not have access to the COMPAS system to verify the scientific validity of the algorithm and the accuracy of his score. Although the Court acknowledged that there were potential problems with COMPAS, it did not require disclosure of the algorithms, factors, and data. It accepted that this information was proprietary and a trade secret and was therefore protected from disclosure. Rather, the Court set limitations on how COMPAS is used in sentencing:

Risk scores may not be used: (1) to determine whether an offender is incarcerated; or (2) to determine the severity of the sentence. Additionally, risk scores may not be used as the

101. *Id.* at 1179.

102. *Id.*

103. *Wisconsin v. Loomis*, 371 Wis. 2d 235 (Wis. 2016).

determinative factor in deciding whether an offender can be supervised safely and effectively in the community.

Importantly, a circuit court must explain the factors in addition to a COMPAS risk assessment that independently support the sentence imposed. A COMPAS risk assessment is only one of many factors that may be considered and weighed at sentencing.¹⁰⁴

Recently, however, a federal district court in New York has taken a different tack. In New York, COMPAS risk assessments are used when deciding whether to grant offenders sentenced to life in prison discretionary parole. A number of prisoners who were convicted of committing homicides as juveniles and who were repeatedly denied parole sued the New York Parole Board.¹⁰⁵ One of their arguments was that the lack of information on the method used by COMPAS to calculate their risk scores violated constitutional due process. The district court allowed the due process claim as well as other constitutional claims to proceed.

In subsequent proceedings, a magistrate judge has ordered production of all the underlying data and analytics of the COMPAS system, subject to a protective order designed to protect the proprietary nature of the system.¹⁰⁶ The information to be produced is described as follows:

(1) the normative dataset used to create and normalize COMPAS (the “Norm Group Data”); and (2) the regression models for two COMPAS “scales”: (a) the General Recidivism Risk Scale, and (b) the Violent Recidivism Risk Scale (the “Regression Models” and collectively, the “Compelled Materials”). . . The Norm Group Data is a repository of offender information from several jurisdictions that Northpointe uses to generate the Regression Models, which are sets of inputs used to predict the likelihood of new offenses and new violent offenses after an offender’s COMPAS assessment date. . . Northpointe also uses the Norm Group Data to translate recidivism risk scores into data presented to individual Defendants before parole hearings.¹⁰⁷

The plaintiffs’ analysis of COMPAS has come to focus on the model’s use of age in calculating risk as well as the appropriateness of the raw data for generating scores for juvenile offenders. It remains to be seen how the due process and other constitutional claims will be resolved, but the case represents a promising effort to reconcile transparency with the proprietary information of private-sector vendors.

104. *Id.* at 275.

105. *Flores v. Stanford*, 2019 WL 4572703 (S.D.N.Y. Sept. 20, 2019).

106. *Flores v. Stanford*, 2021 WL 4441614 (S.D.N.Y. Sept. 28, 2021).

107. *Id.* at *2.