# Agency Use of Artificial Intelligence

## Ad Hoc Committee on Agency Use of Artificial Intelligence

## Draft Statement for Ad Hoc Committee | November 30, 2020

1      Artificial intelligence (AI) techniques are changing how government agencies do their

2   work.[1] Advances in AI hold out the promise of lowering the cost of completing government tasks

3   and improving the quality, consistency, and predictability of agency decisions. But enhanced

4   agency use of AI also raises concerns about the discretion being vested in AI systems and the

5   extent to which those systems are exercising authority that ought to be handled by human

6   officials.

7      Consistent with its statutory mission to promote efficiency, participation, and fairness in

8   administrative processes,[2] the Administrative Conference offers this Statement to identify issues

9   ~~of which~~ agencies should consider ~~be mindful~~ when adopting or modifying AI systems. The Statement

10   draws on a pair of reports commissioned by the Conference,[3] as well as the input of AI experts

11   from government, academia, and the private sector.

---

[1] The National Institute of Standards and Technology has offered the following basic definition of AI:

> AI technologies and systems are considered to comprise software [or] hardware that can learn to solve complex problems, make predictions or undertake tasks that require human-like sensing (such as vision, speech, and touch), perception, cognition, planning, learning, communications, or physical action. Examples are wide-ranging and expanding rapidly. 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 (IT) and Operational Technology (OT).

NAT'L INST. OF STANDARDS & TECH., U.S. LEADERSHIP IN AI: A PLAN FOR FEDERAL ENGAGEMENT IN DEVELOPING TECHNICAL STANDARDS AND RELATED TOOLS 7–8 (Aug. 9, 2019). The Administrative Conference adopts that definition for purposes of this statement.

[2] *See* 5 U.S.C. § 591.

[3] DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY, & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (2020),

12  The With this Sstatement , we are recommending that agencies focus on the following issues when

13  implementing AI: 1) Transparency; 2) Harmful Bias; 3) Technical Capacity; 4) Obtaining AI

14  systems; 5) Data; 6) Privacy; 7) Security; 8) Decisional authority; and 9) Oversight. These issues

15  highlighted in this Statement are relevant to an array of agency personnel. To minimize the risk

16  of unforeseen problems involving an AI system, the agency should, throughout the system's

17  lifespan, solicit input about the system from an array of officesall of these stakeholders—including, at a minimum, the

18  legal, policy, financial, human resources, and technology offices.

*19*                                    *1. Transparency*

20          Agencies' efforts to ensure transparency in connection with their AI systems can serve

21  many valuable goals. When agencies set up processes to ensure transparency in their AI systems,

22  they should publicly identify the processes' goals and the rationales behind them. For example,

23  an agency might prioritize transparency in the service of legitimizing its AI systems, facilitating

24  internal or external review of its AI-based decision making, or coordinating its activities.

25  Different AI systems are likely to satisfy some transparency goals more than others. Where

26  possible, agencies should use metrics to measure the performance of their AI-transparency

27  processes.

28          In setting transparency goals, agencies should consider to whom they should be

29  transparent. For instance, depending on the nature of its operations, an agency might prioritize

30  transparency to the public, courts, Congress, or its own officials.

31          The appropriate level or nature of transparency and interpretability in an agency's AI

32  systems will also depend on context. In some contexts, such as adjudication, reason-giving

33  requirements may call for a high degree of transparency and interpretability from the agency

34  regarding how an AI system functions. In other contexts, such as enforcement, an agency's

35  legitimate interests in preventing gaming or adversarial learning by regulated parties could

https://www.acus.gov/sites/default/files/documents/Government%20by%20Algorithm.pdf; Cary Coglianese, *A Framework for Governmental Use of Machine Learning* (Oct. 2020), https://www.acus.gov/sites/default/files/documents/Coglianese%20Report%20-%20A%20Framework%20for%20Governmental%20Use%20of%20Machine

36  militate against providing too much information (or specific types of information) to the public

37  about the AI system's processes. In each context, agencies should consider whether particular

38  laws or policies governing disclosure of information apply.

39      In selecting and using AI techniques, agencies should be cognizant of the degree to which

40  a particular AI system can be made transparent to appropriate people and entities, including the

41  general public. There may exist tradeoffs between explainability and accuracy in AI systems, so

42  that transparency and interpretability might sometimes weigh in favor of choosing simpler AI

43  models. The appropriate balance between explainability and accuracy will depend on the

44  agency's circumstances and priorities.

45      The proprietary nature of some AI systems may also affect the extent to which they can

46  be made transparent. When an agency's AI system relies on proprietary technologies or

47  algorithms the agency does not own, the agency and the public may have only limited access to

48  the information needed to understand the AI technique. Agencies should strive to anticipate such

49  circumstances and address them appropriately, such as by working with outside providers to

50  ensure they will be able to share sufficient information about such a system.  [RICHARD

    PIERCE]:  An agency should not enter into a contract to use a proprietary AI system unless it is

    confident that the people inside the agency and outside the agency who need access to information

    about the system have adequate access to that information.

51                          *2. Harmful Bias*

52      At their best, AI systems can help agencies identify and reduce the impact of unwanted

53  human biases.[4] Yet they can also unintentionally create or exacerbate those biases by encoding

54  and deploying them at scale. In deciding whether and how to deploy an AI system, ~~therefore,~~

55  agencies should carefully evaluate the biases that might result from the use of the AI system as

56  well as the biases that might result from alternative systems that rely on human actors (such as an

57  incumbent system that the AI system would augment or replace). Because different types of bias

**DRAFT November 30, 2020**

[4] The term *bias* has a technical meaning in the machine learning literature related to model characteristics. Under some circumstances, increasing bias (roughly the error of the average prediction) can improve system performance, if it reduces the risk of overfitting. Here, the Administrative Conference uses the term more generally to refer to common or systematic errors in decision making, especially those implicating normative concerns related to fairness and equal treatment.

58  pose different types of harms, the outcome of the evaluation will depend on the agency's unique

59  circumstances and priorities and the consequences posed by those harms in that context.

60      AI systems can be biased because of their reliance on data reflecting historical human

61  biases or because of their designs. Biases in AI systems can increase over time through

62  feedback~~., which~~ This can occur if the use of a biased AI system leads to systematic errors in

63  categorizations, which are then reflected in the data set or data environment the system uses to

64  make future predictions. Agencies should be mindful of the interdependence of the models,

65  metrics, and data that underpin AI systems.

66      Identifying biases in AI systems can pose challenges, as when the bias affects a particular

67  population but information about ~~which individuals are~~ those in that population is not directly

68  available. To identify and mitigate such biases, agencies should, to the extent practical, consider

69  whether other data or methods are available. Agencies should periodically examine and refresh AI algorithms and other protocols to ensure that they remain sufficiently current and reflect new external data and  circumstances relevant to the functions they perform.

70      Data science techniques for identifying and mitigating biases in AI systems are

71  developing. Agencies should stay up to date on developments in the field of AI, particularly on

72  algorithmic fairness; establish processes to ensure that people with diverse perspectives are able

73  to inspect AI systems and their decisions for indications of harmful bias; test AI systems in

74  environments resembling the ones in which they will be used; and make use of internal and

75  external processes for evaluating the risks of bias in AI systems.

76                          *3. Technical Capacity*

77      AI systems can help agencies conserve resources, but they can also require substantial

78    investments of human and financial capital. Agencies should carefully evaluate the short- and

79    long-term costs and benefits of an AI system before committing significant resources to it. Each

80    agency should also ensure it has access to the technical expertise required to make informed

81    decisions about the type of AI systems it requires, how to integrate those systems into its

82    operations, and how to oversee, maintain, and update those systems.

83          Given the data science field's ongoing and rapid development, agencies should consider

84    cultivating an AI-ready workforce, including through recruitment and training efforts that

85    emphasize AI skills. When agency personnel lack the skills to develop, procure, or maintain an

86    AI system that meets the agency's needs, the agency should consider other means of expanding

87    its technical expertise, including by relying on tools such as the Intergovernmental Personnel

88    Act,[5] prize competitions, or cooperative research and development agreements with private

89    institutions or universities.

90                              *4. Obtaining AI Systems*

91          Decisions about whether ~~or how~~ to obtain an AI system can involve important trade-offs.

92    Buying an AI system from an external source might allow the agency to acquire a more

93    sophisticated tool than it could design on its own, access that tool sooner, and save some of the

94    up-front costs associated with developing the technical capacity needed to design an AI system.

95    Creating an AI tool within the agency, by contrast, might yield a tool that is better tailored to the

96    agency's particular tasks and policy goals. Creating an AI system within the agency can also

97    facilitate development of internal technical capability, which can yield benefits over the lifetime

98    of the AI system and in other technological tasks the agency may confront.

99          Certain government offices are available to help agencies with decisions and actions

100    related to technology.[6] Agencies should make appropriate use of these resources when obtaining

101    an AI system. Agencies should also consider the cost and availability of the technical support

necessary to ensure that an AI system can be maintained and updated in a manner consistent with

its expected life cycle and service mission.

102

## 5. Data

AI systems require data, often in vast quantities. An agency should consider whether it
has, or can obtain, data that appropriately reflects conditions similar to the ones the agency's AI
systems will address in practice; whether the agency has the resources to render the data into a
format that can be used by the agency's AI systems; and how the agency will maintain the data
and link it to the agency's AI systems without compromising security or privacy.

## 6. Privacy

Agencies have a responsibility to protect privacy with respect to personally identifiable
information in AI systems no less than in other aspects of agency operation. In a narrow sense,
this responsibility demands that agencies comply with requirements related to transparency, due
process, accountability, and information quality and integrity established by the Privacy Act of
1974, Section 208 of the E-Government Act of 2002, and other laws and policies.[7] More broadly,
agencies should recognize and appropriately manage privacy risks posed by an AI system.
Agencies should consider privacy risks throughout the entire development life cycle of an AI
system and assess those risks, as well as associated controls, on an ongoing basis. The Office of
Management and Budget and the National Institute of Standards and Technology have developed
risk management frameworks that agencies may find useful in implementing AI systems.[8]

## 7. Security

Agencies should consider the possibility that AI systems might be manipulated, fooled,
evaded, and misled, including through manipulation of training data and exploitation of model
sensitivities. An agency must ensure not only that its data is secure, but also that its AI systems
are trained on that data in a secure manner, make forecasts based on that data in a secure manner,
and otherwise operate in a secure manner. Agencies should continuously consider and evaluate
the safety and security of AI systems, including resilience to vulnerabilities, manipulation, and
other malicious exploitation.

## *8. Decisional Authority*

Agencies should be mindful that most AI systems will involve human beings in a range of capacities—as operators, customers, overseers, policymakers, or interested members of the public. Accordingly, any decision to deploy an AI system should account for the human tendencies and preferences of humans in those roles.

Human factors may sometimes undercut the value of using AI systems to make certain determinations. There is a risk, for example, that human operators will delegate ~~devolve~~ too much responsibility to AI systems and fail to detect cases where the AI systems yield inaccurate or unreliable determinations. That risk may be tolerable in some settings—such as when the AI system has recently been shown to perform significantly better than alternatives—but intolerable in others.

Similarly, if agency personnel come to rely reflexively on algorithmic results in exercising discretionary powers, use of an AI system could have the practical effect of curbing the exercise of agency discretion or shifting it from the person who is supposed to be exercising it to the system's designer. Agencies should beware of such potential shifts of practical authority and take steps to ensure that appropriate officials have the knowledge and power to be accountable for decisions made or aided by AI techniques.

Finally, there may be some circumstances where, for reasons wholly apart from decisional accuracy, an agency may wish to have a decision be made by a human being, even if the law does not require it. In some contexts, accuracy and fairness are not the only relevant values at stake, and an AI system may be difficult to sustain if human beings perceive it as unfair, inhumane, or otherwise unsatisfactory.[9]

## *9. Oversight*

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. An agency's mechanisms of internal oversight will be shaped by the demands of external oversight. And the more effective an agency's internal oversight mechanisms, the better it is

158 ~~likely~~ to fare with external oversight. An agency should be cognizant of both forms of oversight

159 in making decisions about its AI systems.

160 External oversight of agency use of AI systems can come from a variety of government

161 sources, including inspectors general, the Government Accountability Office, and Congress.

162 Courts can also play an important role in external oversight of agency uses of AI systems.

163 Agency uses of AI systems might lead to litigation in a number of circumstances. Those affected

164 by an agency's use of an AI system might, for example, allege that use of the system violates

165 their right to procedural due process.[10] Or they might allege that the AI system's determination

166 violated the Administrative Procedure Act (APA) because it was arbitrary and capricious.[11]

167 When an AI system narrows the discretion of agency personnel, or fixes or alters the legal rights

168 and obligations of people subject to the agency's action, affected people or entities might also

169 sue on the ground that the AI system is a legislative rule adopted in violation of the APA's

170 requirement that legislative rules go through the notice-and-comment process.[12] Agencies should

171 consider these different forms of potential external oversight as they are making and

172 documenting decisions and the underlying processes for these ~~about~~ AI systems ~~and as they are~~

173 ~~developing processes for making those decisions~~.

174 Agencies should also develop their own, internal evaluation and oversight mechanisms

175 for their uses of AI systems. Successful internal oversight requires advance and ongoing

---

[10] Courts would analyze such challenges under the three-part balancing framework from *Mathews v. Eldridge*, 424 U.S. 319, 335 (1976).

[11] *See* 5 U.S.C. § 706(2)(A). Courts would review such challenges under the standard set forth in *Motor Vehicle Manufacturers Ass'n v. State Farm Mutual Automobile Insurance Co.*, 463 U.S. 29, 43 (1983).

[12] *See* 5 U.S.C. § 553(b)–(c).

176 planning and consultation with the various offices in an agency that will be affected by the

177 agency's use of an AI system, including its legal, policy, financial, human resources, and

178 technology offices. An agency's oversight plan should address how the agency will pay for its

179 oversight mechanisms and how it will respond to what it learns from its oversight.

180      Agencies should establish a protocol for regularly evaluating AI systems throughout the

181 systems' lifespans. That is particularly true if a system or the circumstances in which it is

182 deployed are liable to change over time. In these instances, ~~, since, in that case,~~ review and

183 explanation of the system's functioning at one stage of development or use may become outdated

184 due to changes in the system's underlying models. To enable that type of oversight, agencies

185 should monitor and keep track of the data being used by their AI systems, as well as how the

186 systems use that data. Agencies may also wish to secure input from members of the public or

187 private evaluators to improve the likelihood that they will identify defects in their AI systems.

188      To make their oversight systems more effective, agencies should clearly define goals for

189 their AI systems. The relevant question for oversight purposes will often be whether the AI

190 system outperforms alternatives, which may require the agency to benchmark the system against

191 the status quo or some hypothetical state of affairs.

192      Finally, AI systems can affect how agency staff do their jobs, particularly as agency

193 personnel grow to trust and rely on the systems. In addition to evaluating and overseeing their AI

194 systems, agencies should pay close attention to how agency personnel interact with those

195 systems.

196

197      In summary, we offer this statement to highlight the various issues agencies should

198 evaluate when using, investing in, developing, or obtaining AI.

199

---

[5] 5 U.S.C. §§ 3371–76.

[6] Within the General Services Administration, for example, the office called 18F routinely partners with government agencies to help them build and buy technologies. Similarly, the United States Digital Service has a staff of technologists whose job is to help agencies build better technological tools. While the two entities have different approaches—18F acts more like an information intermediary and the Digital Service serves as an alternative source

**DRAFT November 30, 2020**

for information technology contracts—both could aid agencies with obtaining, developing, and using different AI techniques.

**DRAFT November 30, 2020**

200

---

[7] *See, e.g.* 5 U.S.C. § 552a(e), (g), & (p); 44 U.S.C. § 3501 note.

[8] *See, e.g.*, Nat'l Inst. of Standards & Tech., *NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management, Version 1.0* (Jan. 16, 2020); Nat'l Inst. of Standards & Tech. Special Publication SP-800-37 revision 2, *Risk Management Framework for Information Systems and Organizations: A System Lifecycle Approach for Security and Privacy* (Dec. 2018); Office of Mgmt. & Budget, Circular A-130, *Managing Information as a Strategic Resource* (July 28, 2016).

**DRAFT November 30, 2020**

201

---

[9] *Cf.* Admin. Conf. of the U.S., Recommendation 2018-3, *Electronic Case Management in Federal Administrative Adjudication*, 83 Fed. Reg. 30,686 (June 29, 2018) (suggesting, in the context of case management systems, that agencies consider implementing electronic systems only when they conclude that doing so would lead to benefits without impairing either the objective "fairness" of the proceedings or the subjective "satisfaction" of those participating in those proceedings).

202

**DRAFT November 30, 2020**