



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

## 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.<sup>1</sup> 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~~previously  
6 exercised by human officials.

7 Consistent with its statutory mission to promote efficiency, participation, and fairness in  
8 administrative processes,<sup>2</sup> the Administrative Conference offers this Statement to identify issues  
9 of which agencies should be mindful when adopting or modifying AI systems and developing  
10 practices and procedures for their use. The Statement draws on a pair of reports commissioned

---

<sup>1</sup> 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.

<sup>2</sup> See 5 U.S.C. § 591.

**DRAFT November 30, 2020**



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

11 by the Conference,<sup>3</sup> as well as the input of AI experts from government, academia, and the  
12 private sector.

13 The issues highlighted in this Statement are relevant to an array of agency personnel. To  
14 minimize the risk of unforeseen problems involving an AI system, ~~the agency~~ agencies should,  
15 throughout ~~the an AI~~ system's lifespan, solicit input about the system from ~~an array~~ various of  
16 offices—including, at a minimum, their legal, policy, financial, human resources, and technology  
17 offices. Agencies should also keep in mind the need for public trust in their practices and  
18 procedures for use of AI technologies.

### 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 consider publicly identifying the processes' goals and the rationales behind them.  
23 For example, an agency might prioritize transparency in the service of legitimizing its AI  
24 systems, facilitating internal or external review of its AI-based decisionmaking, or coordinating  
25 its AI-based activities. Different AI systems are likely to satisfy some transparency goals more  
26 than others. Where possible, agencies should use metrics to measure the performance of their AI-  
27 transparency 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 higher degree of transparency and interpretability from the agency

---

<sup>3</sup> DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY, & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (2020), <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%20Learning.pdf> (draft report for Administrative Conference of the United States).



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

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  
36 militate against providing too much information (or specific types of information) to the public  
37 about the AI system's processes. In ~~each~~every context, agencies should consider whether  
38 particular 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 specific  
44 context, including the 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 about~~ the AI technique. Agencies should strive to anticipate  
49 such 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.

### 51 *2. Harmful Bias*

52 At their best, AI systems can help agencies identify and reduce the impact of unwanted  
53 human biases.<sup>4</sup> 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 harmful biases that might result from the use of the AI  
56 system as well as the biases that might result from alternative systems that rely on human actors  
57 (such as an incumbent system that the AI system would augment or replace). Because different

---

<sup>4</sup> 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.



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

58 types of bias pose different types of harms, the outcome of the evaluation will depend on the  
59 agency's unique circumstances and priorities and the consequences posed by those harms in that  
60 context.

61 AI systems can be biased because of their reliance on data reflecting historical human  
62 biases or because of their designs. Biases in AI systems can increase over time through  
63 feedback, ~~which~~ Such feedback can occur, for example, if the use of a biased AI system leads to  
64 systematic errors in categorizations, which are then reflected in the data set or data environment  
65 the system uses to make future predictions. Agencies should be mindful of the interdependence  
66 of the models, metrics, and data that underpin AI systems.

**Commented [A1]:** Sentence as drafted has two "which" clauses in a row.

67 Identifying harmful biases in AI systems can pose challenges, as when the bias affects a  
68 particular population but information about which individuals are in that population is not  
69 directly available. To identify and mitigate such biases, agencies should, to the extent practical,  
70 consider whether other data or methods are available.

71 Data science techniques for identifying and mitigating harmful biases in AI systems are  
72 developing. Agencies should stay up to date on developments in the field of AI, particularly on  
73 algorithmic fairness; establish processes to ensure that ~~people with diverse~~  
74 ~~perspectives~~ adequately trained personnel are able to inspect AI systems and their decisions for  
75 indications of harmful bias; test AI systems in environments resembling the ones in which they  
76 will be used; and make use of internal and external processes for evaluating the risks of harmful  
77 bias in AI systems and for identifying such bias.

### 78 *3. Technical Capacity*

79 AI systems can help agencies conserve resources, but they can also require substantial  
80 investments of human and financial capital. Agencies should carefully evaluate the short- and  
81 long-term costs and benefits of an AI system before committing significant resources to it. ~~Each~~  
82 ~~Agencies agency~~ should also ensure ~~it~~ they ~~have~~ has access to the technical expertise required to  
83 make informed decisions about the type of AI systems ~~it~~ they ~~requires~~, how to integrate those  
84 systems into ~~its~~ their operations, and how to oversee, maintain, and update those systems.



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

85 Given the data science field’s ongoing and rapid development, agencies should consider  
86 cultivating an AI-ready workforce, including through recruitment and training efforts that  
87 emphasize AI skills. When agency personnel lack the skills to develop, procure, or maintain an  
88 AI system that meets the agency’s needs, the agency should consider other means of expanding  
89 its technical expertise, including by relying on tools such as the Intergovernmental Personnel  
90 Act,<sup>5</sup> prize competitions, ~~or~~ cooperative research and development agreements with private  
91 institutions or universities, and consultation with external technical advisors and subject-matter  
92 experts.

### 93 *4. Obtaining AI Systems*

94 Decisions about whether or how to obtain an AI system can involve important trade-offs.  
95 ~~Buying~~ Obtaining an AI system from an external source might allow the agency to acquire a  
96 more sophisticated tool than it could design on its own, access that tool sooner, and save some of  
97 the up-front costs associated with developing the technical capacity needed to design an AI  
98 system.<sup>6</sup> Creating an AI tool within the agency, by contrast, might yield a tool that is better  
99 tailored to the agency’s particular tasks and policy goals. Creating an AI system within the  
100 agency can also facilitate development of internal technical capability, which can yield benefits  
101 over the lifetime of the AI system and in other technological tasks the agency may confront.

102 Certain government offices are available to help agencies with decisions and actions  
103 related to technology.<sup>7</sup> Agencies should make appropriate use of these resources when obtaining  
104 an AI system.

---

<sup>5</sup> 5 U.S.C. §§ 3371–76.

<sup>6</sup> Agencies may also obtain AI systems that are embedded in commercial products. The considerations applicable to such embedded AI systems should reflect the fact that agencies may have less control over their design and development.

<sup>7</sup> 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 for information technology contracts—both could aid agencies with obtaining, developing, and using different AI techniques.



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

105

### 5. Data

106 AI systems require data, often in vast quantities. An agency should consider whether it  
107 has, or can obtain, data that appropriately reflects conditions similar to the ones the agency's AI  
108 systems will address in practice; whether the agency has the resources to render the data into a  
109 format that can be used by the agency's AI systems; and how the agency will maintain the data  
110 and link it to the agency's AI systems without compromising security or privacy. An agency  
111 should also review and consider existing statutes and regulations that impact its use of AI as a  
112 potential consumer of data.<sup>8</sup>

113

### 6. Privacy

114 Agencies have a responsibility to protect privacy with respect to personally identifiable  
115 information in AI systems ~~no less than in other aspects of agency operation~~. In a narrow sense,  
116 this responsibility demands that agencies comply with requirements related to transparency, due  
117 process, accountability, and information quality and integrity established by the Privacy Act of  
118 1974, Section 208 of the E-Government Act of 2002, and other applicable laws and policies.<sup>9</sup>  
119 More broadly, agencies should recognize and appropriately manage privacy risks posed by an AI  
120 system. Agencies should consider privacy risks throughout the entire ~~development~~ life cycle of  
121 an AI system from development to retirement and assess those risks, as well as associated  
122 controls, on an ongoing basis. The Office of Management and Budget and the National Institute  
123 of Standards and Technology have developed risk management frameworks that agencies may  
124 find useful in implementing AI systems.<sup>10</sup>

<sup>8</sup> For example, among the requirements established by Title II of the Foundations for Evidence-Based Policymaking Act (i.e., the Open, Public, Electronic and Necessary Government Data Act, or OPEN Government Data Act), Pub. L. No. 115-435, 132 Stat. 5529 (Jan. 14, 2019), agencies must submit annual plans that include the data, methods, and analytical approaches used to develop evidence to support policymaking.

<sup>9</sup> See, e.g., 5 U.S.C. § 552a(e), (g), & (p); 44 U.S.C. § 3501 note.

<sup>10</sup> 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).



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

125

### *7. Security*

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

133

### *8. Decisional Authority*

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

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

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

150 Finally, there may be some circumstances where, for reasons wholly apart from  
151 decisional accuracy, an agency may wish to have a decision be made by a human being, even if



## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

152 the law does not require it. In some contexts, accuracy and fairness are not the only relevant  
153 values at stake, and an AI system may be difficult to sustain if human beings perceive it as  
154 unfair, inhumane, or otherwise unsatisfactory.<sup>11</sup>

### 155 9. Oversight

156 It is essential that agencies' AI systems be subject to appropriate and regular oversight  
157 throughout their lifespans. There are two general categories of oversight: external and internal.  
158 An agency's mechanisms of internal oversight will be shaped by the demands of external  
159 oversight. And the more effective an agency's internal oversight mechanisms, the better it is  
160 likely to fare with external oversight. An agency should be cognizant of both forms of oversight  
161 in making decisions about its AI systems.

162 External oversight of agency use of AI systems can come from a variety of government  
163 sources, including inspectors general, the Government Accountability Office, and Congress. In  
164 addition, because agency uses of AI systems might lead to litigation in a number of  
165 circumstances. ~~C~~ourts can also play an important role in external oversight ~~of agency uses of AI~~  
166 ~~systems.~~ ~~Agency uses of AI systems might lead to litigation in a number of circumstances.~~ Those  
167 affected by an agency's use of an AI system might, for example, allege that use of the system  
168 violates their right to procedural due process.<sup>12</sup> Or they might allege that the AI system's  
169 determination violated the Administrative Procedure Act (APA) because it was arbitrary and  
170 capricious.<sup>13</sup> When an AI system narrows the discretion of agency personnel, or fixes or alters  
171 the legal rights and obligations of people subject to the agency's action, affected people or  
172 entities might also sue on the ground that the AI system is a legislative rule adopted in violation

<sup>11</sup> 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).

<sup>12</sup> Courts would analyze such challenges under the three-part balancing framework from *Mathews v. Eldridge*, 424 U.S. 319, 335 (1976).

<sup>13</sup> See 5 U.S.C. § 706(2)(A). Courts would likely 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).





## ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

173 of the APA’s requirement that legislative rules go through the notice-and-comment process.<sup>14</sup>  
174 Agencies should consider these different forms of potential external oversight as they are making  
175 and documenting decisions about AI systems and as they are developing processes for making  
176 those decisions.

177 Agencies should also develop their own~~r~~ internal evaluation and oversight mechanisms  
178 for their ~~uses of~~ AI systems, both for initial approval of an AI system and for ongoing oversight  
179 of the system. Successful internal oversight requires advance and ongoing planning and  
180 consultation with the various offices in an agency that will be affected by the agency’s use of an  
181 AI system, including its legal, policy, financial, human resources, and technology offices. An  
182 agency’s oversight plan should address how the agency will pay for its oversight mechanisms  
183 and how it will respond to what it learns from its oversight.

184 Agencies should establish a protocol for regularly evaluating AI systems throughout the  
185 systems’ lifespans for efficacy and ethicality. That is particularly true if a system or the  
186 circumstances in which it is deployed are liable to change over time, since, in that case, review  
187 and explanation of the system’s functioning at one stage of development or use may become  
188 outdated due to changes in the system’s underlying models. To enable that type of oversight,  
189 agencies should monitor and keep track of the data being used by their AI systems, as well as  
190 how the systems use that data. Agencies may also wish to secure input from members of the  
191 public or private evaluators to improve the likelihood that they will identify defects in their AI  
192 systems.

193 To make their oversight systems more effective, agencies should clearly define goals for  
194 their AI systems. The relevant question for oversight purposes will often be whether the AI  
195 system outperforms alternatives, which may require the agency to benchmark the system against  
196 the status quo or some hypothetical state of affairs.

197 Finally, AI systems can affect how agency staff do their jobs, particularly as agency  
198 personnel grow to trust and rely on the systems. In addition to evaluating and overseeing their AI

---

<sup>14</sup> See 5 U.S.C. § 553(b)–(c).



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

199 systems, agencies should pay close attention to how agency personnel interact with those  
200 systems.