



Agency Use of Artificial Intelligence

Ad Hoc Committee on Agency Use of Artificial Intelligence

Draft Statement for Ad Hoc Committee | November 9, 2020

1 Artificial intelligence (AI) is transforming how government agencies do their work. In
2 many cases, advances in AI are lowering the cost of major governance functions, improving the
3 quality of agencies' decisions, and making agencies more consistent and predictable. But
4 enhanced agency use of AI also raises important questions about how agencies might best design
5 algorithms and user interfaces; the risk that agencies' uses of AI will exacerbate biases in the
6 agencies' operations; the appropriate spheres of human and machine decisionmaking in
7 administrative processes; the compatibility of machine decisionmaking with foundational
8 administrative law concepts like due process, transparency, and delegation; and the need for
9 different forms of oversight of machine decisionmaking.

10 This Statement highlights some of the most important issues confronting agencies in their
11 uses of AI, whether the technology is developed internally or procured from an external provider.
12 It identifies principles agencies should apprehend and apply in deciding when to use AI, what
13 type of AI to use, and how to develop, implement, and maintain an AI system. Because agencies'
14 use AI in myriad and diverse ways, this Statement does not delve into specific applications of AI.
15 Instead, it focuses on overarching considerations that ought to inform the wide range of decisions
16 agencies confront in using AI. Those considerations include bias, transparency, procedural due
17 process, capacity building, and delegation.

18 *Transparency*

19 In some cases, the mechanisms that give rise to recommendations made by AI are
20 undiscoverable and unknown — a phenomenon known as “black box AI.” Because AI can so

Commented [MT1]: Committee chair John Cooney recommends adding a couple of sentences at the very start of the statement. One sentence would identify to whom the statement is directed and identify the purpose of the statement. The second sentence would inform the reader about the studies that underpin the statement and also explain that the statement reflects the consensus of committee members and also reflects the input of a variety of AI experts from government, academia, and the private sector, whose shared experiences suggest some pervasive issues in agency use of AI.

Commented [MT2]: Committee may wish to consider defining AI and limiting certain parts of this draft statement to particular types of AI (e.g., machine learning). On the other hand, the difficulty of defining terms related to AI may mean it's ultimately not worth the trouble.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

21 readily become an inscrutable black box, it is important that agencies take transparency concerns
22 seriously in developing and deploying AI. Among other things, agencies should carefully
23 consider to whom they should be transparent and for what purposes. For instance, depending on
24 the nature of its operations, an agency might prioritize transparency to the public, to courts, or to
25 its own officials. And the agency might prioritize transparency in the service of various goals,
26 such as legitimizing its AI decisionmaking, facilitating internal and external review of the
27 agency’s AI decisionmaking, or ensuring the smooth coordination of agency activities. Different
28 types of AI are likely to serve some of these concerns better than others.

29 Among other reasons, transparency in AI is valuable because it results in more reliable,
30 checkable, and trustworthy agency decisionmaking processes. To that end, in choosing and using
31 AI tools, agencies should be cognizant of the degree to which a particular AI is explainable. At
32 the same time, however, agencies should be aware that there is often a negative correlation
33 between explainability and accuracy in AI. The appropriate balance between explainability and
34 accuracy will depend on the agency’s unique circumstances, including its priorities when it
35 comes to transparency.

36 The appropriate level of transparency in an agency’s AI will inevitably depend on
37 context. For example, when an agency deploys AI in adjudication, the need to give regulated
38 parties a full explanation of the decisionmaking process can require a high degree of
39 transparency from the agency regarding how the AI functions. By contrast, when an agency uses
40 AI to make enforcement decisions, the agency’s legitimate interest in preventing gaming or
41 adversarial learning by regulated parties could militate against the agency’s publicly providing
42 too much information (or specific types of information) about the AI’s processes. Agencies
43 should be sensitive to such context-based distinctions in evaluating transparency in an AI.

44 It is also important for agencies to think about how they will explain decisions made by
45 their AI. Technologies for explaining AI decisionmaking processes — commonly called
46 “explanatory AI” or “xAI” — are rapidly evolving. Different types of xAI can offer different
47 ways of explaining AI processes, some of which may be more useful than others in different



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

48 contexts. In deciding what type of AI to employ, therefore, agencies should consider the extent to
49 which different types of AI can be explained.

50 Connecting all of these transparency issues are questions about intellectual property,
51 which agencies should consider as early as possible when developing or procuring AI. When an
52 agency's AI relies on protected or proprietary technologies or algorithms the agency does not
53 own, intellectual property protections can limit the agency's access to information that would
54 allow it to understand or explain the AI. For that reason, agencies should make it a point to
55 consider — at the beginning of the AI development or procurement process — the extent to
56 which intellectual property limitations might affect their ability to use AI in a transparent and
57 explainable fashion.

Judicial Review of AI Decisions

58
59 Certain agency uses of AI may lead people to file lawsuits challenging the agency's
60 actions as violating procedural due process. Such a challenge would be analyzed under the
61 framework set forth in *Mathews v. Eldridge*. Many AI systems will fare well under that
62 framework, particularly when the systems deliver more accurate decisionmaking at lower costs
63 to agencies. There are, however, many considerations relating to agencies' uses of AI that would
64 not be cognizable on judicial review but may nevertheless implicate principles underlying
65 procedural due process. Those considerations include norms such as reasoned decisionmaking,
66 fairness, and providing an opportunity to be heard. Agencies should bear those considerations in
67 mind when assessing their uses of AI.

68 Agency uses of AI may also result in lawsuits alleging violations of the Administrative
69 Procedure Act's prohibition on agency actions that are "arbitrary, capricious, an abuse of
70 discretion, or otherwise not in accordance with law."²¹ In adjudicating such lawsuits, courts will
71 consider whether the AI decision "was based on a consideration of the relevant factors and
72 whether there has been a clear error of judgment."²² Complying with that standard will require

¹ See 5 U.S.C. § 706(2)(A).

² See *Citizens to Preserve Overton Park, Inc. v. Volpe*, 401 U.S. 402, 416 (1971)



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

73 that agencies be able to understand and provide at least some explanation of any decisions made
74 by AI that are the basis for agency action.

75 Error correction is closely related to both procedural due process and the APA's
76 prohibition on arbitrary or capricious agency actions. Because AI is not foolproof, agencies that
77 use it should have plans in place for diagnosing and correcting false positives and other
78 systematic errors that threaten regulated parties' legal rights. Appeals in individual cases are not
79 an especially effective means of correcting the sorts of systematic errors that are common with
80 AI. Consequently, when it comes to correcting errors by AI, agencies should focus their efforts
81 on systematic reviews and oversight mechanisms, rather than relying exclusively on individual
82 challenges.

83 *Auditing AI*

84 Regular audits of agency AI are particularly important in this regard. Agencies should
85 account for the need for such audits early on, even during the procurement process, so that they
86 can negotiate agreements that allow them to avoid the intellectual property issues that sometimes
87 make it difficult for agencies and others to understand and audit AI systems. Once an agency has
88 procured an AI, it should conduct regular audits of the AI, particularly if the AI or the
89 circumstances in which it is deployed are liable to change over time. Use of an AI over an
90 extended period can reveal defects that were not apparent earlier in the AI's existence. Regular
91 auditing can help agencies become aware of deficiencies in the AI before those deficiencies
92 create widespread procedural due process problems.

93 Internal AI audits should not be the only mechanisms by which agencies check their AI
94 for potential procedural due process problems. Because auditing itself is prone to biases and
95 other sources of error, agencies should consider how they can most effectively solicit external
96 input about their AI. In addition to providing agencies with a more diverse set of perspectives for
97 identifying potential problems with the AI, external auditing of agency AI can also help
98 legitimize the AI in the eyes of both those who are affected and those who will conduct oversight
99 reviews of the agency's work.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

100 Agencies should also beware that AI can affect how agency staff do their jobs,
101 particularly as staff grow to trust and rely on the AI. In addition to monitoring and auditing AI,
102 therefore, agencies should pay attention to how their staffs use the agencies' AI and whether the
103 AI might be impairing agency staff in performing other job functions.

104 *Bias*

105 As with human decisionmaking, AI decisionmaking can be biased. While AI can help
106 agencies identify and reduce the impact of human biases, it can also exacerbate those biases by
107 encoding and deploying them at scale. In deciding whether and how to deploy an AI, therefore,
108 agencies should carefully evaluate whether the AI is likely to be more biased than the human
109 decisionmaker it might replace. In conducting that evaluation, agencies should also be cognizant
110 that biases have different consequences in different contexts, meaning the consequences of
111 certain types of AI biases will be different depending on the nature of the decisions the AI is
112 tasked with making.

113 Biases can creep into algorithms in several ways. AI systems learn to make decisions
114 based on training data, which can include biased human decisions or reflect historical or social
115 inequities. Bias in AI can also be caused by flawed data sampling, in which groups are over- or
116 underrepresented in training data. Or bias in AI can result from applying an AI in one setting
117 even though it was trained on data from a materially different setting.

118 Agencies should work diligently to identify and combat biases in the AI they use. Doing
119 that requires agencies and their personnel to stay up-to-date on developments in the field of AI;
120 establish processes to ensure that people with diverse perspectives are able to inspect the AI and
121 its decisions for hints of bias; test algorithms in regimes resembling the ones in which they will
122 be used; and make use of technical and external tools, like red teams and third-party audits, to
123 supplement internal agency processes for evaluating bias risks in AI systems.

124 *Capacity building*

125 Implementing AI can help agencies conserve scarce resources. But it can also require a
126 major investment of human and financial capital — both in procuring or developing an AI and



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

127 also in maintaining it. Agencies should carefully evaluate the short- and long-term costs of AI
128 before committing to deploying it.

129 The costs of an AI will depend in part on whether the agency creates the AI internally or
130 procures it from an external source. There are different benefits to each approach. The decision
131 to create an AI within the agency may yield a tool that is better tailored to the agency's particular
132 tasks and more compliant with the agency's policy aims, whereas the decision to buy an AI from
133 an external source might allow the agency to acquire a more sophisticated tool than it could
134 design on its own.

135 The appropriate answer to the "make or buy" question regarding AI will ultimately
136 depend on the agency's unique circumstances, but agencies should be sure to study the
137 alternatives closely. Because it takes a high level of technological sophistication to make an
138 intelligent decision about whether to make or buy AI, agencies facing that choice should ensure
139 they have access to the relevant technological expertise. Given the relative dearth of experts in
140 the emerging field of AI, as well as the field's ongoing and rapid development, agencies should
141 be prepared to expend the financial and human resources to make sure their decision about
142 obtaining AI is well-informed by people with the requisite knowledge.

143 An agency that chooses to develop its own AI must ensure that it has access to adequate
144 technological expertise, otherwise it risks doing a bad job. In some cases, agency personnel will,
145 themselves, lack the skill to build an AI that meets the agency's needs. In those cases, agencies
146 should consider other means of quickly expanding their technological expertise, including by
147 relying on tools like the Intergovernmental Personnel Act, prize awards, or cooperative research
148 and development agreements with private institutions or universities.

149 An agency that elects to purchase an AI from an external source rather than building it in-
150 house should understand that the decision to purchase an AI often requires a high degree of
151 technological knowledge on the agency's behalf. Thus, agencies that outsource the development
152 of their AI should still invest in developing their internal capacity to make sound decisions about
153 the type of AI the agency requires. It is important, for example, that the agency personnel who



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

154 negotiate the purchase agreement know enough about AI so that they can reasonably foresee
155 what information the agency might need about the AI's functioning in the future. Armed with
156 such knowledge, agency personnel should make sure that any agreement they sign to procure AI
157 allows them access to the information they will need to provide about the AI, notwithstanding
158 the likelihood that developers may want to preserve some core trade secrets or innovations.

159 *Delegation and Accountability.*

160 Agencies that use AI may also confront issues involving different types of delegations.
161 One type of delegation involves the delegation of decisionmaking authority from the agency to
162 an algorithm that the agency itself has developed. This type of delegation raises questions about
163 regulated parties' rights to a human decisionmaker, rather than a computerized one. While the
164 law does not give clear answers those questions, values of transparency, accountability, and due
165 process will, in some cases, weigh in favor of providing human oversight to check and confirm
166 AI determinations.

167 Agencies that contract with private parties to develop AI, by contrast, should review the
168 Federal Acquisition Regulation (FAR) prohibition on using contracts "for the performance of
169 inherently governmental functions."³ Likewise, they should consult Office of Management and
170 Budget Circular A-76, which does not establish legally enforceable rights or obligations, but
171 does provide valuable insight into what constitute "inherently governmental."
172 Agencies that rely on AI to perform tasks that might be deemed "inherently governmental
173 functions" should consult with agency counsel to ensure they are complying with the FAR.

174 Finally, whenever an agency uses AI to make decisions, it raises questions of
175 subdelegation by the agency to an algorithm. Whether such subdelegations are permissible will
176 depend on the specifics of the action at issue, the relevant statutory scheme, and other factors
177 identified by the Supreme Court in cases such as *United States v. Morgan*.⁴

³ 48 C.F.R. 7.503(a).

⁴ 313 U.S. 409 (1941).



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

178 *Data collection, standardization, and security.*

179 To enable their AI, agencies must collect or obtain data. In collecting and organizing that
180 data, agencies should consider (1) whether they have the appropriate data set, meaning it reflects
181 conditions similar to the ones the AI will address in practice; (2) whether the data are in a form
182 that's usable without an excess, and (3) how the agency ensures that the data is collected is
183 maintained and linked to the AI.

184 Agencies should also remember that the collection and use of so much data can implicate
185 privacy and data security concerns. Notably, however, almost none of those concerns are unique
186 to agency use of AI. One important exception involves the possibility that an AI will reverse-
187 engineer or effectively discover, through machine learning, certain inferences based on private
188 information. That possibility is one reason why agencies should take care that, in compiling or
189 maintaining that information, they do not run afoul of the Privacy Act or other laws governing
190 agency use and storage of data.

191 *Potential Projects for the Conference Regarding AI*

Commented [MT3]: The committee may wish to identify specific issues with AI that could lend themselves to future ACUS studies and recommendations.