



AI's Promise and Peril for the U.S. Government

David Freeman Engstrom, Daniel E. Ho,
Catherine M. Sharkey, and
Mariano-Florentino Cuéllar

While the use of artificial intelligence (AI) spans the breadth of the U.S. federal government, government AI remains uneven at best, and problematic and perhaps dangerous at worst. Our research team of lawyers and computer scientists examined AI uses among federal administrative agencies – from facial recognition to insider trading and health care fraud, for example. Our report, commissioned by the Administrative Conference of the United States and generously supported by Stanford Law School, NYU Law School, and Stanford's Institute for Human-Centered AI, is the most comprehensive study of the subject ever conducted in the United States. The report's findings reveal deep concerns about growing government use of these tools, and so we suggest how AI could be unleashed to make the federal government work better, more fairly, and at lower cost.

In March 2019, the [Stanford Institute for Human-Centered Artificial Intelligence](#) funded research exploring the topic of AI's growing role in federal agencies. The project culminated in the 122-page [report](#), "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies," which was commissioned by the Administrative Conference of the United States, an agency that provides advice across federal agencies.

In the big picture, AI promises to transform how government agencies do their work by reducing the cost of core governance functions, improving decision-making, and using the power of big data for greater efficiency. Many benefits exist. In the enforcement context,

KEY TAKEAWAYS

- Few federal agencies are using AI in ways that rival the private sector's sophistication and prowess, yet AI use is widespread and poses numerous governance questions.
- AI tools used by the federal government need to reflect transparency and society's longstanding legal, political and ethical foundations.
- At federal agencies, many of the most compelling AI tools were created from within by innovative, public-spirited technologists – not profit-driven private contractors.



the Securities and Exchange Commission can use AI to “shrink the haystack” of potential violations of insider trading, and the Centers for Medicare and Medicaid Services use AI to identify fraud, for example. AI tools can help administrative judges spot error in draft decisions adjudicating disability benefits and help examiners at the Patent and Trademark Office process patent and trademark applications more efficiently and accurately. The Food and Drug Administration, the Consumer Financial Protection Bureau, and Housing and Urban Development currently task AI to engage the public, by sifting through millions of citizen complaints. Others have experimented with chatbots to field questions from welfare beneficiaries, asylum seekers, and taxpayers.

While the benefits are real and tangible, key issues and problems remain. Questions arise, for example, about the proper design of algorithms and user interfaces, the respective scope of human and machine decision-making, the boundaries between public actions and private contracting, the capacity to learn over time using AI, and whether the use of AI is even permitted in certain contexts.

Research Outcomes

To more fully understand this dilemma, our research team studied uses of AI at the top 142 federal agencies, conducted in-depth case studies of adoption, development, and deployment, and analyzed the results from technical, legal, and policy angles.

Our five main findings include:

- **The government’s AI toolkit is diverse and spans the federal administrative state.** This includes conventional machine learning to more advanced “deep learning” with natural language and image data. In fact, nearly half of the federal agencies studied (45%) have experimented with AI and related machine learning tools. These activities reflect enforcing regulatory mandates, adjudicating government benefits and privileges, evaluating public health and safety, extracting

TOP TEN AGENCIES AND SUBAGENCIES BY NUMBER OF USE CASES

Agency Name	Number of use cases
Office of Justice Programs	12
Securities and Exchange Commission	10
National Aeronautics and Space Administration	9
Food and Drug Administration	8
United States Geological Survey	8
United States Postal Service	8
Social Security Administration	7
United States Patent and Trademark Office	6
Bureau of Labor Statistics	5
Customs and Border Protection	4

The above list excludes overarching department-level agencies.



information from the government's massive data streams, and communicating with citizens.

• **Despite the federal embrace of AI, the technological sophistication lags in the public sector.** Our computer scientists found that AI used in federal agencies either lacked detail and only 12% could be rated as high in sophistication. This raises a red flag because federal agencies will find it harder to realize gains in accuracy and efficiency with less sophisticated tools. Without significant public investment, the public-private technology gap could widen.

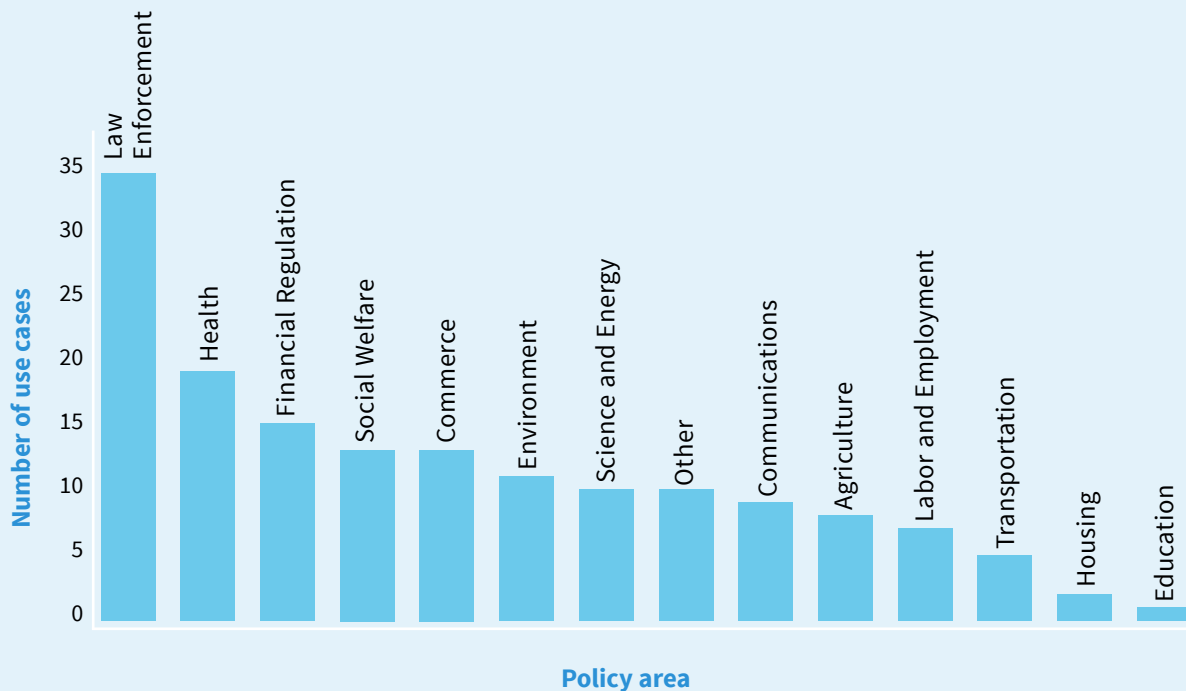
• **AI technical capacity must come from within federal agencies.** While many agencies rely on private contractors to build out AI capacity, a majority of profiled use cases (53%) are the product of in-house efforts

In-house expertise yields AI tools that are better tailored to complex governance tasks and more likely to be implemented in a lawful, policy-compliant, and accountable fashion.

by agency technologists. This underscores the critical importance of internal agency capacity building as AI continues to proliferate. In-house expertise promotes AI tools that are better tailored to complex governance

AI USES BY POLICY AREA

Each bar represents the count of the number of use cases by policy area.



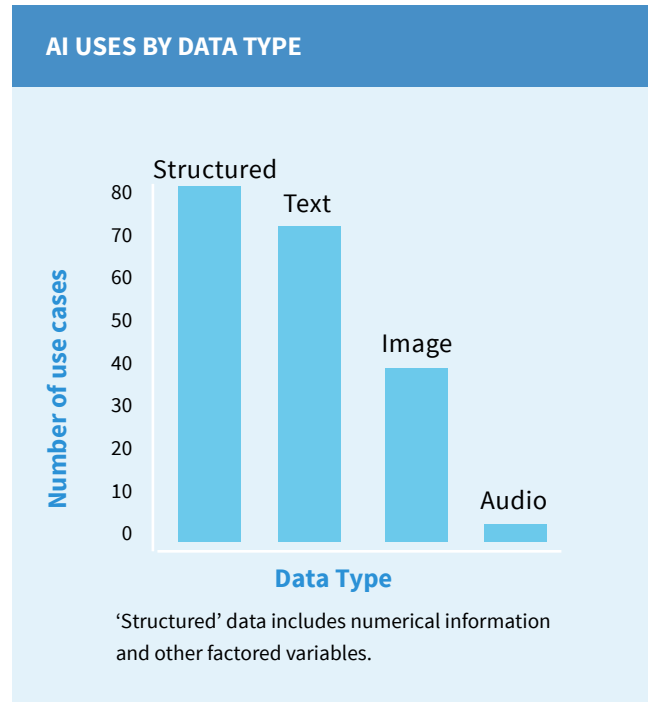


tasks and more likely to be designed and implemented in lawful, policy-compliant, and accountable ways.

• **AI decision-making needs to take into account America’s unique legal norms, reasoning, and constraints.** Government agencies have an obligation to explain decisions, but such explanations become increasingly difficult as algorithmic decision making systems displace human judgment and its capacity for nuance. If AI tools are to successfully proliferate throughout federal agencies, they must be designed to ensure accountability and fidelity to legal norms of transparency, explanation, and non-discrimination. This means creating AI that truly functions reasonably within the complexities of human and institutional environments.

• **AI could raise fairness and equity issues, while also fueling social anxieties.** A concern is that the wealthy and privileged will “game” the federal government’s AI systems by maximizing their considerable resources and know-how. An enforcement agency’s algorithmic predictions could fall more heavily on smaller businesses that, unlike larger firms, lack computer scientists who can reverse-engineer the agency’s model and keep out of its cross-hairs. If citizens come to believe that federal AI systems are rigged, then broad societal support for an effective and tech-savvy government will evaporate quickly.

AI could raise fairness and equity issues, while also fueling social anxieties.





Moving forward on federal AI

As our report notes, the federal government's use of AI raises issues of accountability, technological quality, and societal conflict. Delving deeper, questions arise:

For example, how much transparency is necessary to judge an AI tool's fidelity to the law, accountability, and other societal norms and conventions?

Also, to what extent can existing legal oversight tools, particularly administrative law, achieve meaningful accountability, and to what extent will accountability require newly minted interventions? And, how do we weigh the merits of internally-built AI tools and growing government's own tech operations compared to using private contractors with AI-based solutions at the ready?

This conversation on federal AI is extremely important now. The White House recently announced plans to chart a national artificial intelligence policy for federal agencies – a welcome and long overdue step. Meanwhile, both funding for AI research and updating agency IT systems and guidance for agencies are in the works, and would be excellent next steps in terms of concrete support.

The question now is whether the government's unfolding of AI throughout federal agencies will be managed well or managed poorly

In closing, the federal government is at a crossroads with AI. Managed well, federal AI use can make the U.S. government more efficient, accurate, and fair. Managed poorly, federal AI can exacerbate the public-private technology gap, make agencies more vulnerable and less transparent, and heighten concerns about government arbitrariness and biases.

Much work remains to be done. Wherever the nation lands on AI policy issues at the individual level, federal AI use on a widespread basis is here to stay. The question now is whether the government's unfolding of AI throughout federal agencies will be managed well or managed poorly. How we answer that question will critically shape the future of governance.

The original report, ***Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*** can be found here:

<https://law.stanford.edu/ACUS-AI-Report>

Stanford University's [Institute on Human-Centered Artificial Intelligence \(HAI\)](#), applies rigorous analysis and research to pressing policy questions on artificial intelligence. A pillar of HAI is to inform policymakers, industry leaders, and civil society by disseminating scholarship to a wide audience. HAI is a nonpartisan research institute, representing a range of voices. The views expressed in this policy brief reflect the views of the authors. For further information, please contact HAI-Policy@stanford.edu.



[David Freeman Engstrom](#),

Professor of Law and Associate Dean for Strategic Initiatives, Stanford University; Affiliated Faculty, Stanford Institute Human-Centered Artificial Intelligence (HAI).



[Daniel E. Ho](#),

William Benjamin Scott and Luna M. Scott Professor of Law, Professor of Political Science, and Senior Fellow at Stanford Institute for Economic Policy Research, Stanford University; Associate Director, Stanford Institute Human-Centered Artificial Intelligence (HAI).



[Catherine M. Sharkey](#),

Crystal Eastman Professor of Law, NYU



[Mariano Florentino Cuéllar](#),

Justice on the Supreme Court of California; Herman Phleger Visiting Professor of Law, Stanford University; Affiliated Faculty, Stanford Institute Human-Centered Artificial Intelligence (HAI).



Stanford HAI: Cordura Hall, 210 Panama Street, Stanford, CA 94305-1234

T 650.725.4537 **F** 650.123.4567 **E** hai@stanford.edu hai.stanford.edu