

AI Agents in Federal Agencies

Author : Christopher Walker

Date : September 4, 2020

David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*, **Report for the Administrative Conference of the United States** (2020), available at [SSRN](#).

The use of artificial intelligence is on the rise at federal agencies, and administrative law scholars aren't paying enough attention. In a [forthcoming article](#), [Ryan Calo](#) and [Danielle Citron](#) question whether this increasingly "automated administrative state" presents a legitimacy crisis. After all, legislatures delegate broad law-implementation authority to administrative agencies because of regulators' "ability to accrue expertise and the prospect of flexible and nimble responses to complex problems." Yet these agencies are increasingly subdelegating such implementation authority to "systems in which they hold no expertise, and which foreclose discretion, individuation, and reason-giving almost entirely." Despite these concerns, Calo and Citron ultimately do not demand a deconstruction of the automated administrative state. Instead, they argue that "agencies should consciously select technology to the extent its new affordances enhance, rather than undermine, the [expertise] rationale that underpins the administrative state."

The Calo-Citron article deserves its own Jot. But it also raises a more fundamental question: How automated is the federal administrative state today? Although some work has documented the use of artificial intelligence and machine learning (AI/ML) at a handful of federal agencies, we had lacked a system-wide study. Recognizing this deficiency, the [Administrative Conference of the United States](#) commissioned an all-star group of scholars to comprehensively examine this regulatory landscape. In February, those scholars—[David Freeman Engstrom](#), [Daniel E. Ho](#), [Catherine M. Sharkey](#), and [Mariano-Florentino Cuéllar](#)—issued their 122-page report, entitled [Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies](#).

There is so much to like *lots* in this report—too much to cover in this short Jot. But I'll flag a few highlights.

Part I of the report takes inventory of AI/ML across the federal bureaucracy. To do so, the researchers—including the report authors and a large team of law students, political scientists, and computer scientists—focus on the 142 most significant federal agencies. They find that 45% (64 agencies) have embraced AI/ML to some degree. As depicted in Figure 2 of the report, the most common use of AI/ML is for "regulatory resource, analysis, and monitoring," followed by (in order) enforcement, public services and engagement, internal management, and adjudication. Roughly half of AI/ML use cases (84 of 157) were developed in house. Part I also breaks down the data by agency, subject matter, and implementation stage, among other things. A 43-page [online appendix](#) provides even more granularity, including details on all 157 use cases identified at these agencies. In reviewing the findings, I was struck by the fact that most agencies (though not all) experimenting with AI/ML have substantial resources. Many under-resourced agencies, such as agencies focused on immigration and veterans affairs, don't seem to be doing much AI/ML innovation. I'll return to this observation later.

After presenting the big picture in Part I, Part II details eight fascinating case studies into how federal agencies have utilized AI/ML. The case studies were not chosen at random, but as illustrations of how

agencies utilize AI/ML in various regulatory settings. To illustrate agency enforcement practices, the report has case studies on the Securities and Exchange Commission and Customs and Border Protection. On regulatory analysis, it presents case studies on two pilot programs at the Food and Drug Administration. The report includes three case studies on citizen engagement—at the Consumer Financial Protection Bureau, the Federal Communications Commission, and the U.S. Postal Service.

On administrative adjudication, the report presents case studies on the Social Security Administration and the U.S. Patent and Trademark Office. As Cary Coglianese and Lavi Ben Dor have [observed](#), federal agencies—compared to Article III federal courts—seem much more willing to embrace and experiment with AI/ML when adjudicating. Such experimentation in adjudicating individual claims may seem particularly concerning to many observers. The report explores how the Social Security Administrative has tried to improve the quality of administrative adjudication through the use of AI/ML—by clustering appeals by issue to be decided by specialized appellate adjudicators, by accelerating appeals based on predicted likelihood of success, and by leveraging natural language processing for quality assurance. Similarly, in its case study of informal adjudication at the U.S. Patent and Trademark Office, the report explores how the agency has utilized AI/ML with respect to patent classification, patent prior art search, trademark classification, and prior trademark search. In the patent prior art search context, for instance, the agency developed an in-house tool called Sigma. Yet it ultimately did not implement Sigma agency-wide because the pilot program revealed that the tool “improve[d] efficiency only for examiners with a computer science background.”

Each case study could be the subject of a standalone law review article. Indeed, for scholars looking for research agenda ideas in this area (or law students for student note topics), check out the implications section of each case study, which flags some of the most pressing issues related to that use of AI/ML.

The final part of the report (Part III) zooms back out to examine several cross-cutting implications of the report’s findings. The authors focus on six categories of implications and recommendations:

- (1) building internal capacity;
- (2) transparency and accountability;
- (3) bias, disparate treatment, and disparate impact;
- (4) hearing rights and algorithmic governance;
- (5) gaming and adversarial learning; and
- (6) the external sourcing challenge.

Similar to the case studies in Part II, these half-dozen implications categories could each be the subject of a full-length article. Consider, for instance, the first category of building internal capacity. The report recognizes that federal agencies generally face a “make-or-buy” decision when it comes to incorporating AI/ML in their regulatory activities. The procurement and development of AI/ML will continue to be a challenge for federal agencies. Outsourcing development has important advantages, as those outside of the agency likely have much deeper expertise in developing these AI/ML tools. Yet, as the report underscores, subject-matter expertise can be critical in developing the AI/ML, such that outsourcing the entire process to outside groups would be problematic.

Outsourcing in some circumstances may also be in tension with core values of administrative governance. As Calo and Citron [argue](#), AI/ML developments at federal agencies should be

oriented toward the furthering of substantive commitments and values, such as access, quality, and self-assessment. They are not designed simply to save costs (and in the process undermine procedural commitments without garnering more efficiency) but rather to enhance the capabilities of the administrative state itself—both agencies and officials—to engage in more effective and fair governance. In general, they would not outsource agency functions requiring expertise and discretion to third parties whose software and hardware deliver neither. These efforts have potential to enhance the justification of the bureaucratic state by, ideally, generating knowledge, enhancing expertise, tailoring outcomes, and increasing responsiveness—the purported reasons Congress created agencies to carry out its will in the first place.

Not surprisingly, the report authors strongly recommend internal development even when an agency procures AI/ML from outside the agency. This internal development includes: (1) increasing technical infrastructure and data capacity; (2) developing staff capacity and competency; (3) investing in AI/ML strategy and “regulatory sandboxes”; and (4) incorporating “accountability and transparency by design.” The remaining implications categories in Part III similarly explore how AI/ML should be developed across the administrative state to address some of the dangers of automating regulatory activities.

A critical challenge to developing AI/ML internally, as I hinted above, is that federal agencies face significant budget constraints. Buying off-the-shelf solutions may thus be more tempting for agencies with limited resources, yet such outsourcing often carries greater risks from a Calo-Citron legitimacy perspective. And some agencies that may benefit the most from AI/ML innovations face greater budgetary challenges than others, such that buying AI/ML is not even possible. In the final implications category, the report wisely observes that make or buy are not the only options. Federal agencies can also borrow. These approaches include collaborating with outside non-commercial entities to develop AI/ML, sponsoring AI/ML innovation competitions, and collaborating with other administrative agencies. This borrowing category deserves much more attention in the literature and in the real world. Indeed, we may find similar parallels and helpful lessons from the history of military software procurement, as my colleague Bryan Choi explores in [Software as a Profession](#), where the government pursued such a hybrid development and procurement strategy.

Government by Algorithm is a massive contribution to the administrative law literature, and it lays the foundation for an ambitious research agenda for this emerging subfield of administrative law. The report could not have come at a better time, as federal agencies turn more and more to AI/ML to regulate.

Cite as: Christopher Walker, *AI Agents in Federal Agencies*, JOTWELL (September 4, 2020) (reviewing David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*, **Report for the Administrative Conference of the United States** (2020), available at SSRN), <https://adlaw.jotwell.com/ai-agents-in-federal-agencies/>.