

Remarks Delivered at the Administrative Conference
of the United States (ACUS) Plenary Meeting, June 13, 2019

Legal Issues in the Use of Artificial Intelligence by Administrative Agencies

Cary Coglianese[†]

The world of artificial intelligence has arrived. At the highest level, literally, commercial airplanes rely on machine-learning algorithms for auto-piloting systems. At ground level, again literally, self-driving cars now appear on public streets—and small robots automatically vacuum floors in some of our homes. More profoundly, algorithmic software reads medical scans to find cancerous tumors. These and many other advances in the private sector are delivering the benefits of forecasting accuracy made possible by the use of machine-learning algorithms.

What about the use of these algorithms in the *public* sector? Machine-learning algorithms—sometimes referred to as predictive analytics or artificial intelligence—can also help governmental organizations make more accurate decisions. Just as these algorithms have facilitated dramatic innovations in the private sector, they can also enable governments to achieve better, fairer, and more efficient performance of key functions.

I was asked to open today’s lunch session with some brief framing comments about possible legal issues presented by federal agencies’ reliance on machine-learning algorithms. What might be the administrative law concerns about governmental use of machine-learning algorithms?

I can only provide here a short answer to this question. Longer answers can be found in two recent articles of mine: “Regulating by Robot: Administrative Decision Making in the Machine-Learning Era”¹ and “Transparency and Algorithmic Governance”²—both co-authored with my former student, David Lehr.

For now, I simply want to do two things. First, I want to explain *why* machine learning tools could raise legal issues at all. What makes these tools distinctive? Second, I want to sketch four doctrinal issues implicated by machine learning: delegation, due process, equal protection, and reasoning. Overall, my conclusion is that administrative lawyers should be able to work with agency officials to design highly consequential algorithmic tools that fully comply with prevailing legal standards.

These are real issues—not sci fi. As others on today’s panel will show, machine-learning technologies are already being put into use by federal agencies in the service of domestic policy

[†] Edward B. Shils Professor of Law and Professor of Political Science, and Director of the Penn Program on Regulation, University of Pennsylvania Law School. Professor Coglianese is a Public Member of the Administrative Conference of the United States (ACUS) and the Chair of its Rulemaking Committee. These remarks reflect the author’s personal views and not necessarily those of ACUS.

¹ Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147 (2017).

² Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1 (2019).

implementation. Of course, the vast majority of these uses actually raise few interesting legal questions. No one seriously thinks there are legal problems with the Postal Service using learning algorithms to read handwriting on envelopes in sorting mail, or with the National Weather Service using them to help forecast the weather. And, with *Heckler v. Chaney*³ in mind, relatively few legal questions should arise when agencies use algorithms to help with enforcement, such as to identify tax filings for further auditing.

But we are rapidly moving to a world where more consequential decision-making, in areas not committed to agency discretion, could be aided by, and perhaps even replaced by, automated tools that run on machine-learning algorithms. For example, in the not-so-distant future, certain government benefits or licensing determinations could be made using artificial intelligence.

Such uses will raise interesting legal questions because of *two* key properties of artificial intelligence systems: automation and opacity.

The first property—*automation*—should be pretty obvious. Machine learning algorithms make it possible to cut humans out of decision-making in qualitatively important ways. When this happens, what will become of a government that, in Lincoln’s words, is supposed to be a “of the people, for the people, and by the people”—not *by the robots*?

By itself, automation should not create any legal bar to the use of machine-learning algorithms. After all, government officials can already legally and appropriately rely on *physical* machines—thermometers, emissions monitoring devices, and so forth.

It is the second key property of machine learning algorithms—their *opacity*—that, when combined with the first, will appear to raise distinctive legal concerns. Machine learning algorithms are sometimes called “black-box” algorithms because they “learn” on their own.

Unlike traditional statistical forecasting tools, machine learning does not rely on human analysts to identify variables to put into a model. Machine learning algorithms effectively *do the choosing* as they work their way through vast quantities of data and find patterns on their own. The results of a learning algorithm’s forecasts are not causal statements. It becomes harder to say exactly *why* an algorithm made specific determination or prediction.

This is why some observers will see automated, opaque governmental systems as raising basic constitutional and administrative law questions.

Yet, my second objective today is to suggest that, with proper planning, governmental reliance on algorithms should not face insuperable or even significant legal barriers. Let me take each of four key legal doctrines briefly in turn.

First, let us look at the nondelegation doctrine. If Congress cannot delegate lawmaking authority to private entities, then maybe government cannot legally delegate decision-making authority to machines. Yet, algorithms do not suffer the same dangers of self-interestedness that make delegations to private human individuals so “obnoxious,” as the Supreme Court put it in *Carter v.*

³ *Heckler v. Chaney*, 470 U.S. 821 (1985).

Carter Coal.⁴ Moreover, the math underlying machine learning necessitates that officials program their algorithms with clear objectives, which will surely satisfy anyone’s understanding of the intelligible principle test.

Second, with respect to due process, the test in *Mathews v. Eldridge*⁵ requires balancing a decision method’s accuracy with the private interests at stake and the demands on government resources. The private interests at stake will always be exogenous to machine learning. But machine learning’s main advantage lies in accuracy, and these systems can economize government resources. Algorithms should thus fare well under due process standards.

Third, consider equal protection. Artificial intelligence does raise important considerations about algorithmic bias, especially when learning algorithms work with data that have biases built into them. But machine-learning analysis can be constructed to reduce these biases—something which is sometimes harder to achieve with human decision-making. Moreover, due to the unique ways in which machine learning operates, federal agencies would likely find that even explicit inclusion of variables related to protected classes will still likely be permissible under the Fifth Amendment. The “black box” nature of machine learning will, for example, preclude inferences of discriminatory intent.

Finally, what about reason-giving? Despite machine learning’s black-box character, it should still be possible to satisfy reason-giving requirements. It will always be possible, for example, to provide reasons in terms of what algorithms are designed to forecast, how they are constructed, and how they have been tested validated. Just as agencies now show that physical devices have been tested and validated to perform accurately, they should be able to make the same kind of showing with respect to digital machines.

In the end, although adjudicating by algorithm or rulemaking by robot will sound novel, the use of machine learning—even to automate key governmental decisions—can be legally accommodated into administrative practice.

When used responsibly, machine learning algorithms have the potential to yield improvements in governmental decision-making by increasing accuracy, decreasing human biases, and enhancing overall administrative efficiencies.⁶ The public sector can lawfully find ways to benefit from the same kinds of advantages that machine learning algorithms are delivering in the private sector.

⁴ *Carter v. Carter Coal* 298 U.S. 238, 311 (1936).

⁵ *Mathews v. Eldridge*, 424 U.S. 319 (1976).

⁶ See Cary Coglianese, *Optimizing Regulation for an Optimizing Economy*, 4 U. PA. J. L. & PUB. AFF. 1, 11 (2019).