ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

# ARTIFICAL INTELLIGENCE AND ADMINISTRATIVE LAW DOCTRINES

## Challenges and Opportunities

July 9, 2020

TRANSCRIPT
(Not Reviewed for Errors)

**Panelists**


Cary Coglianese, Edward B. Shils Professor of Law, University of Pennsylvania Law School; Public Member, Administrative Conference of the United States

Deirdre Mulligan, Associate Professor, University of California Berkeley School of Information

Arti Rai, Elvin R. Latty Professor of Law, Duke University School of Law


**Moderator**


David Vladeck, A.B. Chettle Chair in Civil Procedure, Georgetown University Law Center; Senior Fellow, Administrative Conference of the United States

**************************************

AUDIO TRANSCRIPTION OF

Administrative Conference of the United States

Artificial Intelligence and

Administrative Law Doctrines:

Challenges and Opportunities

**************************************

MR. WIENER: Well, welcome everyone. I'm Matthew Wiener, the Vice Chair and Executive Director of the Administrative Conference of the United States. Welcome to the second panel of our Symposium on Artificial Intelligence in Federal Agency Programs, which the Administrative Conference, or ACUS for short, is cosponsoring with the Georgetown University Law Center of Technology -- excuse me, Institute For Technology Law and Policy.

Today's panel is on the relationship between artificial intelligence and administrative law doctrines. We have an outstanding panel lined up this afternoon, which our moderator will introduce in a moment. Today's panel will be -- is being recorded and it will be transcribed. The recording will appear on our Web site before not too long, along with the transcript, and I hope everyone makes good use of the transcript.

With that, let me turn it over to David Vladeck, who's not only a distinguished member of the Georgetown Law School faculty but also a Senior Fellow of ACUS and a longstanding and very good friend of our agency. David?

MR. VLADECK: Well, thank you, Matt. On behalf of the Administrative Conference and Georgetown

Law School's Institute For Technology Law and Policy, I want to welcome everyone to today's discussion of the impact of the government's growing use of artificial intelligence on our administrative state. The Administrative Conference's commissioned report government by algorithm, artificial intelligence in federal administrative agency drives home that the use of AI tools raise fundamental questions about how agencies will perform their vital functions.

The report gives a clear-eyed understanding about both the benefits and challenges caused by the government's increasing use of AI. The report finds that in many areas, AI tools are already substantially increasing the efficiency and quality of government decisionmaking and the delivery of services, but there are still questions to be answered. And they include the following: Algorithms are only as good as the data they're trained on. They may fail to detect new sophisticated frauds, perhaps significant error rates. So one question is, what are the impacts of these limitations?

Next, decisions made by algorithms are opaque. There's an answer, but there's never really an explanation, but reason giving is at the heart of the traditional notions of due process and fairness. Should agencies refrain from using AI in making decisions involving individual rights until decisions can be adequately explained? Algorithms need data. Much of the data these days that's being collected by government is personal data. Some of this is biometric data, can be uniquely associated with an individual.

We tried China by becoming a surveillance state, what impact will AI have on the United States? Artificial intelligence tools such as border controlled by AI machines and not the border patrol agents suggests that AI may displace thousands of government employees. Is that an acceptable tradeoff? AI technology is often hard to use and there's often a mismatch between the expertise of existing agency staff and the need for a technology-proficient staff. Is this a temporary problem or is this endemic?

And last but hardly least, the AI technology government employs is often not government generated or government owned. Increasingly, government contractors are going to carry out core governmental functions. What are the implications of government by contractor? So these issues and more are the subject of today's discussion.

We are fortunate to have the foremost thought leaders on this issue as our panelists. You will first hear from Cary Coglianese who is the Edward B. Shils, Professor of Law and Political Science at the University of Pennsylvania. He is also the founding director of the Penn Program on Regulation and he's also authoring yet another report on AI for the Administrative Conference.

Next up is Professor Deirdre Mulligan who is a professor at the School of Information at the University of California at Berkeley. She's also affiliated with Berkeley Centers on Long-Term Cybersecurity and its Center For Technology, Society and Policy.

Last, but hardly least is Arti Rai, who is the Elvin R. Latty Professor and Founding Director of the Center for Innovation Policy At Duke Law. She also holds an appointment at the Duke-Margolis Center for Health Policy.

So each of the panelists is going to take about 5 minutes to introduce topics that they are particularly interested in working on, and then we'll open this up to some questions. So Cary, you're first.

MR. COGLIANESE: Okay. Thank you very much, David, and thank you to the Administrative Conference for hosting this session. I am a public member of the Administrative Conference, and as David said, I'm working on a report right now commissioned by the Chairman on Artificial Intelligence and its use by government agencies. I'll talk -- some of my remarks certainly will make its way into the report and certainly my report will be greatly informed by the comments that Deirdre and Arti share as well.

Let me try to make four main points here at the outset, just really points that I think are fundamental to keep in mind and provide a good foundation for discussion about the use of artificial intelligence tools. So the first point is just definitional. What is it that we're talking about here?

Algorithms, some people say, is what we're talking about, but we're really talking about a very small category of algorithms. And algorithms have been with us since the dawn of humanity. Two plus two equals four is an algorithm. My peach cobbler recipe is an algorithm. Section 553 of the Administrative Procedure Act is an algorithm for how an agency creates a rule. Statistical analysis, regression models and all of that, those are algorithms, too.

What at least what I'm talking about when I'm talking about artificial intelligence is the use of machine learning algorithms. This is a particular type of algorithm that is very good at forecasting, making predictions. And I think, like with other kinds of statistical tools that make predictions, the point is to try to identify something that's likely to happen in the future or to give some kind of probabilistic assessment of that event happening.

Unlike traditional statistical tools, machine learning has, I think, two qualities that make this conversation important to have. One quality of a machine learning algorithm is that it is autonomous in the type of variables that are selected and how -- and then even sometimes the mathematical relationships between them. It's not the variables that are selected, the weights to be given to them, the mathematical relationships between them, are autonomously generated by the algorithm itself, rather than specified by humans in advance as with a traditional regression model.

That automaticity, if you will, that learning nature is one reason why it's -- the machine learning algorithm has a second quality in that it's more opaque. It's not as easily or intuitively obvious how to explain the results that the algorithm achieves. Unlike with a regression model, where you could look at statistical weights on each coefficient and describe how much of the variation is explained by a particular variable, that's not as easily accomplished with a machine learning model.

Machine learning typically also, I should add, is really about correlation, not causation. And that's another reason why it's hard to explain. It's not possible, usually, to be able to say; because of A, B occurs with a machine learning model. But these are very highly accurate in many context, they're being used a lot in the private sector, in the medical fields and in many contexts have been shown to be much more effective and accurate in making these predictions. So that's the first point, just sort of what we're talking about.

The second point that I think is fairly -- should be fairly obvious is that the kind of concerns that David and other -- you know, David outlined for us at the introduction and that we'll talk about here today, are not distinctive or unique to machine learning. So again, other kinds of statistical or analytic tools have errors in them. Can they have biases in them that come in through data? Yes, they can. Can other statistical tools be used to support a surveillance state, you know, raise privacy concerns? Absolutely. Will there be other things that cause the loss of jobs for people in society? Sure. I mean, modern computers have done that. There's a lot fewer typists employed by the government today than there was 50 years ago.

So that's really the second point is just to say that we're not moving to a completely new world with artificial intelligence. The problems that we're going to talk about here today are problems that can exist and have existed with other types of statistical tools or methods of decision making.

Third point is that there is really, I think, nothing at its core about machine learning that should pose any insurmountable legal obstacles to its use by government officials, assuming they're acting in a responsible manner and careful about their development of these tools in the same way that they should be careful about their deployment and use of any other tools. This is an argument I'm happy to get into further, but I've made this in a couple of articles recently; transparency and algorithmic governance in the Administrative Law Review regulating by robot in the Georgetown Law Journal work that I did collaboratively with David Lehr.

So I'm happy to go into that, but I think the basic idea here is that if you take any of the legal concerns that people have raised about artificial intelligence, it is, I think, possible to justify its use with a careful work and analysis.

For example, let me just briefly say that we're probably going to talk a lot about explainability and the obscurity of algorithms. The standards that we have for transparency right now under due process or arbitrary and capricious review are pragmatic and not absolute. And you can give an adequate reasons for why a decision is made under a machine learning model, okay? And I'm happy to get into that further.

Fourth and final point that I want to make here at the outset is that when we are thinking about the legal issues and the policy issues about artificial intelligence, we need to always keep in mind; compared to what? Artificial intelligence compared to what? Well, if I were to tell you, you know, that we're going to have governmental decisions made in a process that's prone to error, fatigue, racial bias, delay, inconsistency, you would probably say, if that's the world that you want to advocate, we want to be very cautious, we want to resist that kind of world. But that actually is the world we have today with human beings and their decision making.

So it is not as if the status quo is perfect nor necessarily acceptable, especially today as the nation realizes, even more acutely than perhaps ever before in a more widespread fashion, the problems of systemic racism, the implicit bias that human beings can have. If we can -- if we can develop technological tools that can overcome some of the inherent biases and errors and inconsistencies that exist in human decision making, I think we ought to be open to that, and again, should not necessarily think about artificial intelligence tools or machine learning algorithms as some kind of brave new world, but actually as an opportunity when used well to improve on the status quo. Thank you.

MR. VLADECK: Thank you. Deirdre?

MS. MULLIGAN: Hi, thank you so much for having me here today. I want to start just a little bit more broadly. So I've been interested in the questions of how technology embodies policy and what that means for the mechanisms through which technology is brought into government processes in a lot of different domains. And to give a like non-machine learning example, I was part of a team that had a National Science Foundation grant for many years looking at the security, accuracy, privacy and other sorts of implications of the move to electronic voting systems.

And one of the things that came up there is that counties were procuring machines that really didn't meet public policy objectives. And to give a really crystal-clear example, the early direct record electronic voting systems basically collapsed casting and counting. And so normally, right, we think about marking a ballot and then that ballot is logically independent of the system of counting, right, whether it's individuals counting the ballots or it's a optical scan system counting the ballot, we know we can audit the counting process because we have the fixed ballots to go back to.

Because there was no concrete description of what it meant to be a ballot, when the developers designed these direct record electronic voting system what would happen is it would render a ballot on a screen, you would select your inputs, and it would incrementally add them to the counter, and then discard the ballot, right. So if the purpose of having an audit was to audit the counting process, the design of these systems made it completely impossible. So if you said could we have a recount, they'd say, sure, let's put the button and we'll get the same answer, right.

And you could look at that as a real failure of the technologist, or you could look at that as a kind of failure at an institutional level for us to kind of thought through, what were the assumptions in a paper-based world and what were the new -- the way in which we needed to kind of think through the design implications in this new electronic world so that we would still have the same level of kind of logical independence between casting and counting that's required for a meaningful audit, right.

And I just offer that as an example that, while the questions around machine learning models and the way in which they embed policy, I think, are really driving this conversation about how do we think about administrative law and technology. There are a lot of other areas where we adopt technology where they embed really profound policy choices that, I think, bypass many of the mechanisms that we normally rely on for scrutinizing them to make sure that they comply with or embody the policy choices that we need them to reflect.

So I come to this conversation, at least in part, by the Loomis decision which I think many of you are probably familiar with, it involved the use of a risk recidivism tool. I'm not going to go through all the details, but the case ends up being appealed to the U.S. Supreme Court. And the Solicitor General weighs in on this case and basically says that the challenge to the use of this compass tool is not a suitable vehicle for looking at the question of whether or not risk recidivism tools can basically be used -- be normed differently for male and female offenders, because it's completely unclear from the record below about how the compass tool accounts for gender.

And for me, you know, that means this case has gotten all the way up to the Supreme Court and there's a lack of clarity in the record about how gender is being used, right. And the debate was about was it actually an attribute that was being used in the model, or were there different scales being used for men and women, right, both of which raise interesting questions, but the problem is like there was no clarity in the record, right. And this is a tool that had been put in use in the system of justice, right. And to be that raises such enormous questions about how we are unboarding technology into various processes, whether they're in the courts or in other administrative agencies.

And that might not have been so troubling if you viewed it as an outlier, but then there was some additional work by Brauneis and Goodman in 2018 where they looked at the way in which states and counties were coming to use different sorts of algorithmic assessment tools. And they found that government simply did not have many records concerning the creation and implementation of algorithms, either because those records were never generated, or because they were generated by contractors and never provided to the governmental clients.

And this meant there were no records that modeled design choices, data selection, factor weighting, validation designs, and at a really -- at the most basic level the governments didn't even have a record of what problems the models were supposed to address and what the metrics of success were.

And I think when we read the ACUS report most recently, we see a similar strain here where there was a -- there's a finding that for most government applications, 61 percent, there was insufficient publicly-available technical documentation to determine with precision what methods were employed. In some cases the agency description appears more like marketing language or concerns of tools still under development. In other cases agencies describe use of neural networks, natural language processing or facial recognition technologies, but do not provide enough technical details to discern whether a use case is a simpler or a more sophisticated version thereof, right.

And so we see like a lack of information about what tools are being used. Now, the ACUS report finds, I think, different variations of expertise in different agencies, and it was really interesting to me in the Brauneis and Goodman work at the state and county level they find that most tools are being outsourced. And I think one of the really interesting findings in the ACUS report is that many of these tools aren't being developed in house. But similarly, there's some real questions, I think, about the extent to which both agencies have the level of expertise available to make good decisions about base models.

And secondly, the extent to which things that are really core policy issues, and I would say things, questions about how we operationalize target values, the choice of model to use, thresholds, the training data, where it is sourced from, how it is cleaned, when it is updated, all of these issues require both a profound level of expertise, but also because they are essential policy choices, we also need techniques to make sure that they are visible both to the public in ways that allow for public participation about such choices, as well as for agency staff, because one of the things that we know if we want to maintain some discretion in engagement is figuring out how agency staff understand the outputs of models, understand the reasoning of models, even if they can't understand the inner workings of them, are essential for them to kind of align well with agency mandates.

So I think an important part of this question is not like what are the problems, but really where can we look for solutions. And what I want to suggest is that I really, at this point, think that we need to figure out a way to create more centralized communities of expertise within the government. The rise during the Obama Administration of the U.S. Digital Services and 18F, which provides kind of a skunkworks effort that can be both developing resources and guidance documents and methods for sharing expertise across agencies is very important, coupled with the use of things like impact assessments, but I also want to suggest prototypes and other things that allow the public to understand some of the policy choices that are so knitted into the technical designs of these systems are really essential for us to think about developing. And I know my time is up. So I'm going to sit down.

MR. VLADECK: Thank you so much. Arti, you're next.

MS. RAI: Well, thanks so much to my fellow panelists and to the Administrative Conference and to Georgetown Law for inviting me to this event and to David for moderating. So I think that both Cary and

Deirdre have spoken eloquently and at a very acute way about the ways in which database machine learning may or may not raise challenges that are different from other technologies, because as they've both pointed out, we've had algorithms forever, including secret algorithms, including somewhat opaque algorithms.

So Loomis, which involved an algorithm was not a machine learning algorithm as far as I'm aware. Now, we don't know for sure, because it was secret, but most commentators believe it was not, repeat, not a machine learning algorithm. So these issues, in lots of ways, are not new.

I want to spend just a few minutes highlighting some issues that do arise that are perhaps a little bit newer, even in relatively utilitarian contexts, that don't implicate individual rights and bias to the same extent that some of the context to which Deirdre was speaking, do implicate such concerns of rights and bias. And I think these issues are not entirely new, but they have to do with the fact that database machine learning does involve such extremely high levels of expertise. And in the context that I'm going to focus on on the part of the private sector that is going to agencies to get various rents; the economic term for the types of things that the private sector seeks from agencies.

So I'm going to focus on the patent system of which I am extremely well versed in, but I think that the context that I'm talking about, the patent context, is not dissimilar from other contexts where very highly learned, highly expert entities can seek a rents including by using machine learning from agencies.

So let me, again, as I said, mention the patent office as my specific use case. And then -- but first show the general principles of the patent use case illustrates and then get into, dive a little bit into the specific case which was, by the way, in the ACUS report. And so if you want to read it in greater detail, the ACUS report has a chapter that draws from my work on the patent office.

So I think the case study of the patent office highlights two points that are generalizable in commercial context that involve sophisticated players. First, and this is in keeping with what's already been said, effective use of machine learning by the public sector will require real expertise, not necessarily expertise to protect individual rights in all cases, but actually expertise to deal with the expertise on the other side that the agencies are faced with, because the commercial context has involved very sophisticated players on the other side. So that's one thing.

Then second, even in these cases that don't involve individual rights in the same way as some of the cases upon which Deirdre was focussed, the use of machine learning will raise opacity issues, multiple opacity issues, that can be flags for what, at least some would say, are due process problems. Now, we may not necessarily, all of us may not necessarily think that these are particularly salient due process problems, but certainly lawyers will raise them as due process problems even when the entities that are being affected are, you know, large, sophisticated corporations.

All right. So let me talk about that opacity piece in particular, because I think that is a little bit new. And it's not that secrecy is new, it's not that complexity is new, but I do think that database machine learning combined with secrecy raises the stakes to an even greater level of opacity. So here I'll bring in my intellectual property background a little bit and note that reverse engineering database machine learning algorithms is, for the most part, a little bit harder than reverse engineering ordinary algorithms.

And so as a consequence, if you can't reverse engineer very readily and you don't have access to any of the training data and the decision making is very complex, that's going to create perhaps a slightly greater level of opacity than we've seen in the past.

And that's where I think it's sort of the combination of all the different types of opacity. It's the complexity of opacity plus the secrecy of opacity, neither of which individually was unique, but in combination, I think, that level of opacity becomes a little bit different from what we've seen in the past. So the synergy is what I'm really interested in. And this synergy really does come out in the patent case that I have focussed my initial research, at least, on. So let's turn to this patent examination case study.

So the patent office faces a really difficult challenge. It gets hundreds of thousands of patent applications a year and it has a very small labor force to deal with those applications. And because it's completely funded by user fees, not very much money either, because users don't want to pay very much money. All right. Unless you think, by the way, that bad patents are an arcane issue that doesn't -- don't have real-world social consequences, let me just throw out one statistic that relates to drug pricing and bad patents.

There's a recent study from Health Affairs just came out in June 2020 that examined the extension of patent term caused by questionable patents on just a few drugs. And it found that patent term extension on just five key drugs cost Medicaid programs about half a billion dollars from 2010 to 2016. So questionable patents on just five drugs; half a billion dollars. And that's just Medicaid which is a very small part of total health care spend.

So you know, these are real issues with real social welfare consequences where the commercial states are extremely high and we have extremely sophisticated players on all sides, well, really on the private sector side and then the public sector side has to sort of strive to keep up, in other words, the public examiner core.

Okay. So how does this work in terms of trying to integrate machine learning into what the Patent Office does. So there's pretty good evidence. And one of my colleagues at Duke, Michael Frakes is responsible for generating most of it, quite frankly. But poor quality patents are granted, because examiners just don't have the time to search for the prior invention to determine whether the application covers territory that is either not new or is obvious. So that's where some of these bad drug patents come from, for example, and he's got a recent paper on that as well.

So machine learning could automatically help find this prior invention to help with the time problem. But it's going to be opaque, perhaps, what the prior -- and how it finds the prior invention. And as a consequence there are good patent attorneys who are already very much in the mode of telling the patent office you can't have anything that's opaque, because that's a due process violation, and we are going to be very suspicious if you try to implement any of that.

And as a consequence, the Patent Office, which has been working on trying to implement machine learning immediately put out this statement saying that, oh, no, no black box for us, we're going to be completely transparent. Well, of course being completely transparent also means that you create opportunities for gaming by the very sophisticated players on the other side.

So what's happened? Well, the Patent Office was, to its credit, very much on top of the idea that machine learning -- could you believe these things called concept semantic tools could be used. And so they came up with this homegrown, so that Deirdre's point is well taken that this was homegrown, it was actually pretty good. It was a latent semantic analysis tool called Sigma. Multiple problems.

First, it wasn't as transparent as the Patent Office, the lawyers who prosecute patents before the Patent Office, wanted. And that was probably for a good reason, but nonetheless, that was a problem. Second, the -- and this goes to expertise on the patent examination core side. The algorithm was difficult to use even by the geeky people at the Patent Office, because those who were geeky but not in computer science didn't -- couldn't use it well. So the biochemists of the world weren't able to use it well.

So multiple problems, and as a consequence, their homegrown algorithm, good as it was, didn't really get off the ground. So then they have to try to contract with the private sector and then the private sector then adds the layer of trade secrecy on top of complexity and opacity caused just by a complexity. And they couldn't end up contracting with the private sector contractor that was willing to give away all its trade secrets essentially. And as we've just discussed, the Patent Office had announced that it wasn't going to take anything that had trade secrecy protection, because that's not what the patent lawyers -- or they weren't going to allow the Patent Office to do that.

So the Patent Office found itself in a box, and I think this is a real problem for these sophisticated commercial contexts where the public sector is just racing to keep up with what the privacy sector is leaps and bounds ahead in doing, including for purposes of gaming the system. And so that is my intervention for my 5 minutes. Thank you very much.

MR. WIENER: Okay. So this is a nightmare for the moderator, because there's a consensus about too many things. There's a consensus that the status quo isn't great and that machine learning algorithms can help. I think there's a consensus that there's an expertise issue in government; that existing staffing may not have the expertise to actually use these tools to their best advantage. And third, there's an opacity issue. You know, you can't interrogate a machine learning algorithm, you just can't. They won't talk, and you can't make them talk.

So I think what I'd like to get comment on and we'll do 3-minute rounds this time, and we'll start with Arti because she was so patient the last time around. What impact does that have on how government audit will follow? Do we need to have as some, including Deirdre, have suggested, sort of a common core of people who are expert who can then sort of help out?

Do we need to have embedded expertise in every agency? Sort of how do we deal with that and how do we deal with the opacity problem? Because again, machine learning algorithms may be very expert in forecasting or making predictions, but they also, you know, they may be wrong or they may not be accurate at times. So Arti, you go first.

MS. RAI: Great. So I'm glad you brought up Deirdre's point about expertise, concentrating expertise or at least having core expertise resources. And I think that idea is a very good one. It seems to me that having both centralized expertise and more agency-specific expertise would be useful. Of course, that's expensive and that's part of the problem --

MR. VLADECK: Right.

MS. RAI: -- that, you know, all of this would be very expensive to generate, because as would be, I think, as obvious to many people, individuals who have this expertise are highly sought after in the private sector. They can earn maybe 10X what they could earn in the public sector in the private sector. And so this is a very -- and I've seen personally and engaging with various agencies situations where the person with whom I was engaging left for a much higher-paying job in the private sector. And so it's, you know, dispiriting.

I do wonder whether there's a possibility of doing something that's a little bit just outside the public sector, some sort of third-party certification model, at least for certain machine learning areas. So this comes to me because of some work I've been doing in the healthcare space where the physicians and other healthcare providers are thinking about how to set up third-party certification organizations at least for some types of machine learning systems so that every hospital doesn't have to have that level of expertise. And I wonder if that could happen in a way that would be useful for government agencies as well.

As for the opacity issue, so yes, I think opacity a really unique issue in terms for all the reasons of secrecy plus complexity that I noted. I think it's a trickier issue when you're talking about sophisticated commercial players though than when you're talking about individual rights, because a full transparency with sophisticated commercial players gives them opportunities to game the system, and that, I do worry about. And that is where I wonder if, you know, the fact that you have to contract out for trade secret protected stuff is a good cover for -- at least in certain contexts -- basically being able to hide that algorithm from those who might use it to game the system.

MR. VLADECK: Thanks. Deirdre?

MS. MULLIGAN: So I wanted to jump in, I guess, first on -- there's another piece in the ACUS report and they say that no agency examined in the report had established systematic protocols for assessing the potential for an AI tool to encode bias. And to me, one thing that centralized expertise could provide is some background knowledge about different ways to think about bias and different ways to think about its relationship to fairness, right.

So we know fairness can mean lots of things, right, there's the dignitary interest, there's you know, different ways of thinking about fairness at a legal level, right. When we get into how to translate that into the design of a technical system, we get into all these tricky issues about how fairness is measured. Is it by group level, demographic parody, equal positive predictive values, equal negative predictive values, yada, yada, yada, yada, right. It goes on and on and on. And right now, like I'm not particularly surprised that no agency had established systematic protocols for assessing.

My guess is that no agency has yet established systematic ways to think through how to consider kind of biases that are in the data, in the model selection, in the selection of target variables. And so it's not just how to measure it and audit it over time, but it's even like what to aim for and how to build towards that goal. And so to me, I can't imagine a world in which every single agency is going to be expected to do all of that detailed work on their own without some scaffolding at a central level.

Now, the one other thing I want to say about expertise though is a lot of the conversation tends to be, oh, we need data scientists, we need machine learning experts, we need people who do neural networks. And I would suggest that I think a lot of the relevant expertise is actually bridge players. So people who understand statistics and understand enough about machine learning and understand enough about the law that they often tend to be these hybrid players. And for them to facilitate good reason decision making about the use of machine learning, they need what are called, and what we like to call, boundary objects to bring other people along in the conversation.

So if we want to have a conversation about bias and its relationship to fairness, we need ways to tease out, are we talking about the bias embedded in the system or embedded in the overall system of justice and how it relates to that? Are we looking to root out particular forms of bias in the data? How can we do that? What sorts of techniques do we have? So I just want to like slightly problematize the kind of expertise in that I think it's not just about machine learning experts. It's about kind of which teams that can do a lot of boundary spanning and provide meaningful advice to other government agencies.

And then the second question that you asked about, opacity. I think Arti's absolutely right that I think like we might have different tolerances for different levels of opacity in different domains and in different areas, depending upon whether or not we're talking about the deployment of enforcement or the allocation of benefits and the extent to which the use of the tool constrains agency discretion.

And I think when we're talking about something that's used to surface patterns that can inform decisions about enforcement priorities that's used in a decision assistance manner, right, that's very different than, for example, if you have a tool that is kind of being used in a more formulated way and where there are real costs to agency personnel for deviating from the output of a technical outcome.

MR. VLADECK: Thank you. Cary?

MR. COGLIANESE: Well, Arti and Deirdre have laid out a lot of great points as I will try not to repeat those. I would just want to make -- I guess say three things. First, on the expertise issue. It's much bigger than just the need for expertise to use machine learning and we have, I think, a denigration of expertise in our larger culture today and certainly in certain levels and parts of the government as well. And so we need to think about expertise more broadly.

We need to think about it more broadly, I think, also not just because government might want to use artificial intelligence tools, but because the private sector is using them. And in an article I wrote called optimizing government for an optimizing world, or "Optimizing Regulation for an Optimizing Economy", we have increasingly private sector actors who are using these tools, and government is called upon to think about how to regulate and oversee those private sector uses. So the government needs to get that expertise in house whether it's using the tools or just overseeing private sector harms that might come about from these tools.

Second point is that the type of oversight, the type of public engagement, all of the best practices that I think Deirdre has outlined very nicely, I think those are going to vary, though, depending upon the type of use. And this may be just restating it in a little bit different way something that Deirdre was just saying, but we might care a little bit more about having robust transparency and public participation for systems that really make a tangible difference and actually override human decision making in areas of great consequence to people.

But maybe we don't need all of that if the National Weather Service, as it does, is using machine learning tools for weather forecasting, or the Postal Service which was actually one of the first parts of the Federal Government to use machine learning, is using it to read handwriting on envelopes and deliver mail. You know, these more banal uses maybe even today with the FDA, perhaps relying on machine learning tools to scour the 30,000 papers that have been written so far just since January 1 in the scientific literature on COVID-19, for example. These are important uses, but they're tools for finding things and helping humans make decisions and that may be different than actually substituting for human decisions.

Third point. Third point here is just with respect to the obscurity and particularly the issue about private contractors who are doing a lot of this work. I just think we need to be more thoughtful about government procurement when engaging private contractors and making sure that legitimate trade secrets can be protected. I don't think that the source code needs to be disclosed to withstand procedural due process or arbitrary and capricious review.

But do you need to be able -- does the government need to be able to disclose the outcome variables that are being used? Does it need to disclose the objective function that the algorithm is designed to optimize for? Sure. Those things -- but neither of those things should be -- we should worry about trade secrets about. Even with validation runs and the like, I think, can be insisted on in the procurement process when government is relying on private consultants to make sure that that's disclosable down the road to be able to withstand any concerns about adequate transparency.

MR. VLADECK: Well, thank you. Because that's the bridge to the next set of questions that I want to ask. So Arti recounted an effort by the Patent Office to develop its own machine learning algorithmic tool and ultimately had to resort to the private sector.

Do we worry about government increasingly relying on tools that are not generated by government?

Just take one example which is drawn from the ACUS report. You know, the CBP, the border folks are trying to use face recognition to substitute for boarding passes. And they've contracted with a number of outside organizations and, you know, not surprisingly, because face recognition technology is not perfect, there have been all sorts of problems.

And but one of the ironies here, and this is sort of driven home in the ACUS report, was CBP wasn't able to explain any of the problems that it was facing, because most of the sort of intellectual property there was bought and used to some extent by third-party contractors. And so the question is, what do we do about that? Is that okay? Is normatively, is it okay for government to essentially outsource core governmental functions? We've done it in other spheres, I mean, we have private prisons that now, you know, hold lots of federal inmates. But here, when we're making these kinds of choices, is it okay to rely on outsourced government contracting. And Cary, you go first on this one and we'll go back around.

MR. COGLIANESE: Again, I don't see that there's anything necessarily different here than in any other context except for the fact that, at the end of the day, you have to keep in mind that the output of a machine learning model will not be a set of coefficients that can allow the government, very intuitively and easily, to explain why a forecast was made or what proportion of a variance is explained by particular variables.  That's just not the way machine learning works.

But can the government, even if they're relying on a private contractor to disclose what I think are the essential elements to withstand due process, describe the system, how is it structured, what is its goals, what are the data that are being used, what were the validation tests that were done, what were the results of those validation tests, can we show that there's some, you know, increase in accuracy?

Under Mathews v. Eldridge that's one of the three factors in the balancing test that the Court uses for procedural due process in particular.  And machine learning, generally speaking, when designed well has the potential to really improve accuracy and reduce errors. That doesn't mean it will eliminate errors, and the errors that machine learning and algorithms make will often be ones that humans would not make, and maybe we can't even understand why they were made, but if we can get fewer errors, you know, you just need to be able to make sure at the end of the day that you contract in such a way that the government will be able to have access to those essential elements to demonstrate what it did, why it was designed the way it did.

The ineffable sort of black box nature of machine learning, you know, is not something that I think current law requires government to somehow be able to address. I mean, I think of machine learning tools as just tools. And just as, you know, it's sufficient for government to be able to explain if it's using a thermometer if this thermometer has been validated to read certain temperatures accurately and not necessarily provide some kind of, you know, phenomenal logical explanation of why mercury does what it does or why the physics underlying the thermometer. It's a tool, it's been designed for a certain purpose, it's been validated, it works well.

And if you could think about that as the way of explaining what a machine learning algorithm does, I think we're basically going to be fine. Just make sure that when you go through the procurement process with third parties, firms, that you are going to be able to have enough information at the end of the day to demonstrate that. And I don't think that requires them giving up the precise, you know, innovations that they have for the particular kind of machine learning algorithm, because there's lots of different types of machine learning algorithms. They don't have to disclose that, but just give us enough about how this thing is working.

MR. VLADECK: So Deirdre, Cary used the magic word "procurement". Do you want to talk a little about that as you respond more generally?

MS. MULLIGAN: Sure. First, I want to respond, like I think the use of the example of a thermometer is just like really misleading, right. Thermometers, we have a whole set of standards, we understand what we're trying to measure, we understand not just that there's testing and validation and we do a whole bunch of different things to make sure that they are doing the task that we want and actually doing it in the right way.

And these machine learning tools, we need to be concerned, not just that they're giving the right answer, but that they're giving the right answer for the right reasons, right, if we're thinking about using them for making important decisions. And as we've been discussing, you know, what machine learning does typically, it's used not to learn from detailed decision trees that experts put out and say, here's how we reason. What it is used to do is to try to look at past decisions by agency or other experts and to, from that, develop its own logic that it uses, right.

So it is displacing logic in some way, it's coming up with its own reasoning. And as we know, because machines are good as detecting patterns that we don't see, it's often not intuitive to us and it often, you know, like all the classic examples of identifying the snow rather than the wolf, right, like we know it's actually reasoning wrongly. And that, you know, the known unknowns I think here, we know that it is reasoning based on things that we wouldn't reason on. And those can be super dangerous. And so I think it's really important to keep that in mind here.

So I think that the same way we think about like security, right, we understand we're not going to be able to like root out security, but what we want to do is locate it in places where we can best manage it. And I think when we're thinking about something like bias and a lot of the issues that come up in relationship to facial recognition systems are concerned about different performance on different segments of the populations, different distributions of false positives and false negatives and what that might mean in practice for the population.

I think there's a broader set of questions that I actually think are more important about whether or not using surveillance technologies given the racism and other issues in our criminal justice and policing system generally poses particular kinds of systemic risks that suggest that we shouldn't be using those kinds of tools at all in the current environment, but kind of cabining those off and just talking about kind of a performance of these algorithms.

I think questions that we need to ask is, who do we think is going to be more attentive to the issues of over and under inclusion in training data, who's going to be in the best position to understand and to do the testing and validation to understand different performance on different populations?  And what level of transparency do we want around both the standards being set both for the data and for the technology and for the audits of that technology?

And so I think as in, you know, Jody Freeman's famous, "The Private Role in Public Governance", yes, right, there's -- there are different ways to structure private and public relationships around important governmental functions, but I think here we really need to be attentive because of the level of opacity, as Arti described, that's not just at the trade secrecy level but at the technical level itself.

And I think that like the last issue is kind of the brain drain from government. And we see a data drain from government, and I think that's very problematic. I was on the Oakland Privacy Advisory Commission for a while here in Oakland and we have a surveillance ordinance that we were applying to different kinds of technologies. And one of the things that seems to become the norm is that private sector comes in, says we'll do this function for you, we're going to suck up all the public data and then we're just going to give you reports, right. And that means that governments aren't even able to assess whether or not something is performing well, because they don't end up with the raw data to do their own validation, right.

So I think we also need to be concerned about kind of how expertise and how raw assets end up being redistributed in these public/private arrangements.

MR. VLADECK: Thank you. Arti, you want to weigh in?

MS. RAI: Yes. So I'm going to weigh in with my intellectual property hat on, because I've been doing a lot of research recently into what sorts of things private sector firms consider their core trade secrets and what can be disclosed either fully publicly or to trusted parties, quote/unquote, which might include the government agency but not necessarily the whole public.

So you know, one thing that is relevant, I think, to this discussion is that summaries of training data, demographic characteristics, you know, how it was collected, all of those procedural issues surrounding training data, I have not heard anyone say from any of my private sector interviews that that is something they consider a trade secret. So summaries of training data, they get into some relative detail about demographics and the like.

And also -- and this gets a little more tricky with respect to labeling, you have to label the data in order for it to be good training data. Labeling is a little more tricky, but you can -- you know, that can be publicly disclosed to some extent. So I think that mitigates some of the concerns that one might have in these contexts where individual rights and bias are really very salient, which I do think is a really important context even though, for the most part, you know, I had not been talking about those context.

So I think that's one way to do it with respect to training data, because the training data does seem to consider, you know, the secret sauce or the real gold for these private sector firms. One can talk about training data without reviewing the training data even perhaps to the government agency, although ideally, I think a government agency should be trusted enough with the training data that they should be able to get the training data.

They may not have the expertise to know what to do with it exactly, but at least given that their FOIA exemptions and all the rest of it, I don't see why a government agency can't get the training data, or even the source codes so long as it's exempt from FOIA. But fully public information can be made available regarding summaries.

And then everything that Cary said with respect to the reasoning of the model, it seems to me, nobody with whom I have spoken in the private sector has disagreed with the idea that the key factors, confidence intervals with respect to predictions, that sort of thing, that's all really good best practices.

And then finally, last but not least, in terms of validation, some of the bias and other concerns about which we are rightly focused -- on which we are rightly focused, can be addressed by making sure that your model is tested -- well, is created on really diverse populations and all the rest of it, but then also tested on a totally separate data set from the data set from which the training data was devised, so you get performance measures before it's even put out into the wild on a totally independent data sets.  Again, that's all expensive, but it's not trade secret, it shouldn't -- none of that should be trade secret.

MR. VLADECK:  Okay.  So we have a question from the audience that I'd like to read to each of you. And Arti, you get the first crack at this.

The question is this: IT systems in use in government agencies tend to be entrenched -- become entrenched and obsolete.  Given limited AI experience and the high cost of labor, how is this tendency not to become exacerbated for AI systems?

And when I was in government we were using the computer system that, you know, I think there were still Kaypro computers around.  So you know, any thoughts about that question, which I think is a serious question?

MS. RAI:  Yeah, I completely agree.  When I was at the Patent Office we were still, believe it or not, there were faxes being used which --

MR. VLADECK:  Oh, yeah.

MS. RAI:  -- faxes still existed.  So yeah, so the -- I think this is a really serious issue which is why I do think that the private sector has to be involved. I don't see how this can be a completely homegrown effort at the end of the day. I mean, at the PTO, as I've said, they did come up with a pretty good homegrown effort but it wasn't user friendly. And so it does seem to me that the private sector has to be involved which involves money, which involves these careful, you know, safeguards to make sure that what the private sector is giving you has been properly validated and so forth. So yeah, I think it's a real problem and a real concern.

MR. COGLIANESE: If I could jump in?

MR. VLADECK: Sure.

MR. COGLIANESE: The Government Accountability Office, you know, a few years ago did a study and I guess something on the order of 75 percent of all IT spending in the Federal Government goes to legacy systems. So you know, we are pouring a lot of money into really antiquated technology. So it's not just the expertise, the human capital that we need in government but also the actual hardware systems, and I'll add to that, the data. We have a lot of data, but it's often not organized, there aren't ways of linking up to separate data sets.

So there's some work to be done if

government is going to take a lead in developing these systems. And I agree with Arti that, at least in the world we're in today, I mean, if we don't make those improvements in the infrastructure, the IT infrastructure and the expertise, government will have to be relying a fair amount on the private sector for help with this.

MR. VLADECK: Deirdre?

MS. MULLIGAN: I don't know that I have much to add. I do think that it is interesting to look at the way USDS and 18F, for example, their healthcare.gov, 18F, USDS come in and provide some different kinds of expertise and really change outcomes as far as usability and system design. And I think that the positioning of expertise and what we need to bring in house and what we can outsource and the positioning of the technology itself, right, what we can rely on, in the private sector and what agencies need to bring in house is really going to be domain and problem specific.

And so I don't think we're going to have one size fits all here, but I think there's a constant kind of denigration of government capacity. And I think making government an interesting place where one gets to solve the most important problems facing the world can be a really compelling thing for people who, right now, are trying to figure out how to get you to click on (inaudible). And I think that bringing experts often brings new technology, like the experts are the people who design it.

And so I think that there is a real desire roaming in the relevant technical communities to solve important problems. And if we can bring more of those people into government in different ways, whether it's on tours of duty or if it's in specialized skunkworky kind of ways or on agency staff or on expert advisory committees, I think we can bring some technology with it. And so you know, I think there won't be a one-size solution and I'm optimistic about the government.

MR. VLADECK: So we have about 8 minutes left. What I'd like you to do in the last quick round here is sort of talk about how government is going to be able to attract and retain people of this kind of sophisticated knowledge.

When I was at the Federal Trade Commission when I started in 2009, there was not a single technologist on staff, retention is very difficult. Arti pointed this out earlier, what are we -- you know, do we concentrate our expertise in OTA or something like that, or do we diffuse it among the agencies? If we're going to take advantage of AI, how can government best do this? And you know, in answering it, is outsourcing really the only answer?

Your patent example, I think, is a cautionary tale. The agency tried -- the office tried to develop its own algorithm and ended up having to abandon it. So Cary, let's start with you. Each of you have about 3 minutes because we have to wind up at 3:15.

MR. COGLIANESE: Well, I think there's not really a, you know, going to be something that these tools inherently by themselves are plug and play in different context. So I do think that agencies need to develop their own inhouse expertise about how to use these tools for the type of problems that they confront. And some agencies are doing that. The Securities and Exchange Commission, for example, has developed an inhouse staff focusing on machine learning tools to identify -- help identify fraud on the market. So that's an example.

How do we get there? I mean, I think that we have a need for people in government who do a lot of types of analysis and not just use machine learning tools. And we always have to think about inspiring our younger people, you know, to the ideals of public service and to the real value that people in government can provide and the challenges, quite frankly, too. I mean, I think in some respects if I had to leave the audience with anything is to say that I think the technical issues here are in some sense almost the easiest ones.

When it comes to government and government using machine learning tools, the value choices and the policy issues are some of the really tough nuts that we have to crack. You know, inevitably there will be tradeoffs between things like accuracy and fairness, or even within particular context tradeoffs between achieving one goal and not harming people in another way. And how we actually make those choices are policy normative ones where we need -- we need at the end of the day, not just technology, but we need people. We need good people in government who can interface with, interact with people in society overall and help us make those policy choices in a way that will make these systems be viewed as legitimate and not just as somehow technically accurate.

MR. VLADECK: Okay, thank you. Deirdre?

MS. MULLIGAN: So David, you know this and Cary and Arti and others may know. You know there is a movement right now to try to develop career paths and professional identity for students with a certain level of technical expertise in a stem field, but also with an interdisciplinary orientation. So some understanding of a particular social context, particular discipline along with kind of ethical and legal competence to create a public interest technology field, similar to the way we developed a public interest law field.

And I don't think there's a reasonable expectation that the financial rewards of working in government are ever going to be the same as those working in the public sector, and I think all of us on the call right now realize this. And I don't think that should be the goal. I think that there are lots of reasons to want to work in the government on these sorts of problems in helping technical people see themselves as problem solvers of large social problems in teams or as part of, you know, a social justice movement, part of a good government movement that we need them and we need to partner with them and that we want to partner with them and we're making space at the table for them, are really important signals.

And so I think it's hopeful that there's a set of universities around the country that are trying to help technical people view themselves in public service, view themselves in social justice. And I think that that can have a real -- just setting the right tone and developing the career pathways can help address some of the lax that we see in government today.

MR. VLADECK: Arti, you get the last word.

MS. RAI: Oh, great. I do want to end on a more optimistic note, because I think I've been somewhat pessimistic through a lot of my talk here. So I do agree with Cary and Deirdre that there are, I think, there's a significant cadre of people coming up through the ranks who are motivated by the desire to improve the functioning of government. And it seems to me that one way to harness that energy perhaps, perhaps, would be to have something similar to what we've established in some agencies through offices like the Office of the Chief Economist.

So those are offices where you can go in and head up the office for a couple years and then go back to wherever you came from, so you don't have to be a permanent, you know, resident of Washington, D.C. for the rest of your life. You can come in and out. And that has been a model that's been used with technologists as well, although perhaps not as robustly as they could. And it seems to me that that sort of model would work very well for technologists as well. And it was -- it has been used, as Deirdre has pointed out, to some extent, but perhaps even more so within each agency would be something worth considering moving forward.

MR. VLADECK: Well, the FTC has a chief technology officer. We brought in Ed Felten early on ---

MS. RAI: Right.

MR. VLADECK: -- and it really transformed the agency, so maybe that's one tool that we might try to use more uniformly. Always have a chief technology officer. Well, thank you all. You've been terrific. I would applaud you, but no one I think can hear. But great job and we've ended on time and so thanks so much.

MS. MULLIGAN: Thank you, David.

MR. COGLIANESE: Thank you, David.

MS. RAI: Thank you.

MR. VLADECK: Thank you, guys.

MR. COGLIANESE: Thank you Georgetown and ACUS.

MS. RAI: Yes, indeed.

MR. VLADECK: And Matt.

MS. RAI: And Matt.

(End of audio file)

CERTIFICATION

I, Carmel Martinez, TX CSR No. 8128, FPR No. 1065, do certify that I was authorized to and did listen to and transcribe the foregoing recorded proceedings and that the transcript is a true record to the best of my ability.

Dated this 7th day of August, 2020.

Carmel Martinez,

TX CSR No. 8128

FL FPR No. 1065