ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

ARTIFICAL INTELLIGENCE IN FEDERAL AGENGIES
Bias and Government Artificial Intelligence

July 29, 2020

TRANSCRIPT
(Not Reviewed for Errors)

Panelists

David Super, Carmack Waterhouse Professor of Law and Economics, Georgetown University Law Center

Kristin Johnson, McGlinchey Stafford Professor of Law, Tulane University Law School

Alex Givens, President and Chief Executive Officer, Center for Democracy & Technology

Moderator

Chai Feldblum, Partner, Morgan Lewis & Bockius; former Commissioner, U.S. Equal Employment Opportunity Commission; Public Member, Administrative Conference of the United States

****

AUDIO TRANSCRIPTION OF

Administrative Conference of the United States

Artificial Intelligence in Federal Agencies:

Bias and Government Artificial Intelligence

****

MR. WIENER: Well, good afternoon. I'm Matt Wiener, the Vice Chair and Executive Director of Administrative Conference of the United States. And I'd like to welcome you to this panel. This is the third panel of our symposium on Artificial Intelligence in Federal Agencies, sponsored by the Administrative Conference, or ACUS for short, and the Institute for Technology Law and Policy at Georgetown Law School.

Today's panel is on bias and artificial intelligence. As far as I'm concerned, there's no more important topic in the artificial intelligence area than bias. And we have an outstanding panel to address the topic this afternoon. I'm especially pleased that our panel is being moderated by Chai Feldblum. There's no one more qualified to moderate the panel. And you'll note on her -- in our program materials that it lists her many affiliations, which includes now as a partner at Morgan Lewis. And before that she was a Commissioner at the EEOC for nine years. And before that, a very distinguished law professor at Georgetown Law School. And she also happens to be a member of the Administrative Conference of the United States. And we are very, very happy to have her as a member and she's a very good friend to the Conference. And having said that, let me turn it over to you, Chai, for what I think will be an outstanding and interesting discussion.

MS. FELDBLUM: Great, thank you, Matt. And I thought you were going to start with; among the various things she's a public member of ACUS, which I would have started with, because really, I just as you know, I think ACUS just plays an incredibly important role in thinking through tough issues. And I think this panel is one example in the whole series on "Artificial Intelligence" shows that role that ACUS is playing.

So I'm very excited to be moderating. I'm very excited that you-all are going to hear from really three incredible folks. And instead of me reading three sentences from their bios, what I'm going to do is just ask each of them to tell you, obviously, their name, where they are and just a few sentences about how they got into this area of AI and bias.

So Kristin Johnson, we'll start with you and then Alex, go to you, and then David. So Kristin?

MS. JOHNSON: Great. Thanks so much, Chai. I am Kristin Johnson, the McGlinchey Stafford Professor of Law and Associate Dean of Faculty Research at Tulane University Law School. I am delighted to join you and I have to join Chai in thanking Matt Wiener, Todd Phillips and Todd Rubin as well as Jeff Gary and ACUS as well as the Georgetown Institute for Technology in the Law for organizing this summer symposium in general and this panel in particular.

I'm tremendously grateful for the report that was distributed last February, "Artificial Intelligence in Federal Agencies" that explores the role that artificial intelligence has or plays -- machine learning plays in the context of federal agency adjudication, rulemaking and other regulatory activities. I was delighted that the organizers committed time in their report and this symposium for a discussion of what Matt correctly describes as one of the most critical and concerning areas in the adoption and employment of artificial intelligence technologies. Specifically today, I'll be focused on algorithms --

MS. FELDBLUM: Wait, wait, wait, wait. Kristin, I'm going to stop you, because I was like, wow, that's more than an introduction. I'm just going to stop you before you head into your substance, okay. I just want Alex and David to introduce themselves and then we'll get into substance. So Alex?

MS. GIVENS: Sure. Thank you so much. So I'm Alex Givens. I'm the CEO of the Center for Democracy and Technology which is a civil society organization based in Washington, D.C., that for 25 years has worked to fight employees' individual rights in the center of the digital revolution. We focus on a very wide range of issues from consumer privacy to preserving an open and accessible Internet to security and surveillance issues to platform governance, which is to say that everybody I usually work with are staring intently at their screens as we speak, and they are not listening to me, they're listening to the House hearing that is happening right now.

So I'm very grateful to those of you who have dialed in and to join this conversation. I will say that most relevant for our conversation today, we do a lot of work focusing on how data is collected and used to make decisions that impact people's lives. And that's the crux of the conversation we're having today. So I'm thrilled to join you-all.

MS. FELDBLUM: Great. And David Super? Oops, David, make sure your video is on. I mean, make sure your audio is on. That's what we didn't -- okay. You know what I'm going to do, David, in terms of time, I'm going to tell people how awesome you are, and you're going to figure out the audio, because David is a law professor at Georgetown Law School who does lots of incredible stuff including, as you will hear, some issues around AI and government. And the way we're going to do this is Kristin, Alex and David are all going to tell you about some specific issues they've been working on; it's sort of like case studies.

And our goal for this panel for what we hope you walk away with this panel from, one, hopefully having learned something new and interesting. Two, if you're a in-government agency or a lawyer working with these issues, that you come out with some concrete ideas of what you might want to do. And third, especially if you're in academia that you have some new and interesting ideas that you might want to research.

So that's our overall goal. We're going to do sort of the case studies, I'm going to then ask the panelists to ask each other questions or make observations. And then we're going to open it up for questions. So please go ahead and write your questions in the questions box. I'm sure a lot of you have been doing lots of these Webinars so you would know that.

Okay. So Kristin, I'm going to move it over to you to talk about a very interesting piece in terms of credit and government agencies. So on to you.

MS. JOHNSON: Great. Thanks so much. So the use case I'll present initially, and I'm happy to follow up with a little bit of a discussion regarding how regulatory agencies might directly integrate AI into their platforms, but the use case I'll open with in terms of our panel is a credit scoring use case. In other words, my discussion for the next few minutes will focus on the integration of algorithms, or more specifically, machine learning algorithms into consumer credit scoring platforms.

This will pair quite neatly with what I believe Alex will discuss in the employment context. And in fact, one of the examples I will offer for the purpose of setting the stage for those of you who haven't joined for the entire series, and maybe new to AI, involves an employment related algorithm.

So let me set the stage by describing the background for this research. In a co-authored paper with Frank Pasquale and Jennifer Chapman and during my testimony this time last summer before the House Financial Services Committee Artificial Intelligence Task Force, I outline a number of concerns arising from the integration of nonconventional types of data in the consumer credit scoring -- or credit underwriting process.

More specifically, my research is focused on the integration of what I will describe as alternative data in the consumer finance phase. I'm focused on legal and ethical implications of the commodification of this type of data and the outcomes for our society.

Many of you will be familiar with what I'm going to describe, because personally, you may know your credit score, right? And that's why I like this use case, it's immediately accessible. So what I'll do in my remaining time is describe the integration of alternative data into innovative financial services platforms. I'll explore some concerns regarding the potential for technological advances to deliver on the promise of including many who historically have been excluded from the financial services space as well as raising some concerns and offering up some questions or suggestions for research that we may explore as academics or that consumer advocacy agencies have -- or advocacy organizations have undertaken more careful research and thoughtful analysis.

And finally, a few reflections that may be generally applicable across government agencies that interface with the public. As you may well know, individuals and families increasingly rely on credit to finance household purchases and overcome significant unanticipated expenses. Without access to credit on fair and reasonable terms, it can be extraordinarily expensive to be poor. For families with fragile financial circumstances, credit may serve as a lifeline enabling consumers to make short-term debt obligations and to pay for education, transportation, housing, medicine, child care and even food. And in today's pandemic, these concerns are especially poignant.

Two critical developments create promise for the 26 million Americans who are credit invisible, meaning they do not have credit histories. And the 19 million Americans with thin, impaired or stale credit histories. We would describe these as unscorable citizens. The birth of big data, the collection, storage and analysis of vast volumes of consumer data fuels artificial intelligence are automated decisionmaking platforms.

Similar to the proliferation of AI platforms and healthcare, education, employment and criminal law enforcement, the rise of AI in finance monetizes consumer data. Together consumers Web browsing data, click stream data, social media data and other bits of information aggregated through the consumers interaction with the world, and in the Internet more broadly, creates consumers' digital interface.

This consumer digital interface reveals intimate details about consumers' financial transactions including their checking and savings cash flows, their credit and debit card transactions, data that has tremendous value. Data mining engenders a new set of behavioral scoring criteria for evaluating credit worthiness. We describe this criteria as alternative data. Early studies reveal three significant challenges that arise as we integrate alternative data and in our endeavor to regulate the integration of this information.

First, alternative data may advantage or disadvantage. It is not immediately and implicitly clear that one outcome is more likely than the other, but both are probable. Particularly for those who are legally -- who are part of legally-protected classes or who are marginalized or vulnerable as a result of, for example, their immigration status or other personal financial circumstances or attributes.

Under the behavior scoring model, your friends on Facebook, the people in the pictures you post on Instagram, all of this, whether or not you participate in protests, may influence the interest rate that you receive on your next mortgage or car loan or whether or not you're eligible for an education loan.

Second, learning algorithms evaluate facially neutral alternative data. Facially neutral being the descriptor but maybe, in fact, not completely accurate, yet the results of automated decisionmaking processes may unintentionally use variables that function as proxies for protected traits or characteristics. As a result, the use of machine learning algorithms or these highly sophisticated algorithms may lead to disparate impacts for members of legally- protected classes.

Consider, for example, Amazon's experiment with the learning algorithm tasked with reviewing resumes for a software programmer position. Armed with thousands of resumes from previous hires and general instructions regarding qualification, the algorithm went rogue. Because previous hires had predominantly been men, the algorithm began to discount any reference to women or women's colleges. So in other words, those who had preferences to serving as a women's chess club president or having participated in women's tennis clubs or women's teams, were unfortunately discounted in the algorithm's calculation regarding which resumes might be preferred among the class of resumes. Unknowingly, the algorithm replicated historic discrimination and hiring biases.

Third, machine learning algorithm engage in machine learning. And by "machine learning", because I'm the first speaker up to bat, I'll just describe quickly, applies inductive techniques to large data sets to enable the algorithm to learn rules that are appropriate to a particular task. In other words, the intelligence of machine learning is oriented to outcomes, not process. A smart algorithm is designed to consistently reach accurate results based on a chosen task and designated parameters. Like a calculator multiplying 15 digit numbers faster than any human brain could, in a narrow, well-specified area it can reach conclusions faster than humans might be able to. However, the reality is that the dimensions are difficult and the issues here are nuanced.

As I described earlier in the context of credit scoring algorithms, the data shared with the algorithm can make all the difference. The process of planning, selecting, storing and partitioning data among other concerns might lead to data sets that deeply influence how the algorithm learns. Unintentionally, leading developers have released algorithms trained to be neutral only to discover that the algorithms performed in blatantly discriminatory ways.

In the narrow context of consumer finance, federal agencies exercising oversight of fair consumer lending practices such as the CFPB, the OCC, the FDIC or the FTC and others who have historically focused on the enforcement of the Equal Credit Opportunity Act, and other laws ensuring fair lending, would want to be exceptionally thoughtful about the firms they regulate and their integration of algorithms. In particular, in the context of adverse notices where an algorithm has been deployed by a third-party credit rating agency, or acquired by a credit agency from a third party, it will be imperative for the credit rating agency to describe why and how the credit rating agency has taken a particular perspective on an applicant's credit rating or an applicant's credit score.

MS. FELDBLUM: You need to -- one minute warning.

MS. JOHNSON: Down to the last three sentences. In the modern context of federal agencies' adoption of machine learning algorithms, we might note that integrating historic law enforcement or prosecutorial data may lead to efforts -- may lead us to direct efforts and resources to areas that have been historically heavily policed, right, sort of replicating the pattern in our society.

Similarly, law enforcement agencies that use facial recognition technology in a variety of context may discover as did the authors of the darker shades or gender shades study that those with darker skin tones may not be as easily identified using facial recognition technology.

These kinds of implications of bias by the algorithm that's presumed to execute based on neutral data are issues that we must continue to explore, carefully evaluate and effectively regulate which may mean submitting impact statements requiring regulated entities, impact statements regarding the algorithms' use and integration or other interventions to ensure fairness.

MS. FELDBLUM: Thank you. Clearly, we're enough on the way (unintelligible) in terms of where we are on time. So David Super, let's see how your audio is doing. And it's not. And I sent an e-mail. So would anyone on ACUS please see about e-mailing or calling David and seeing whatever you can do in terms of helping? And sometimes, as people know who are listening, I'm sure who have been on these, I don't know if sometimes it helps to sign off and sign back on, but as a panelist, you may not want to do that. So sending out the request to ACUS to help David.

Okay. So Alex, I'm going to go to you instead and take it away, and actually, we'll follow very nicely on what Kristin has been talking about.

MS. GIVENS: Sure. And David, feel free to say "test, test" while I'm speaking if you need to keep testing your audio. I won't be distracted. So I'm going to turn us -- I'm actually covering two different topics in my remarks today. One is going to focus on benefits determinations and then the second is going to pick up on employment, which draws on the example that Kristin just talked about.

In the benefits section, the piece that I will focus on is the increasing use of algorithms to help inform eligibility determinations for benefits programs. This is happening at an increasing number of state programs both in the United States and around the world. And in many instances we're seeing devastating effects from errors and miscalculations in how these new tools are developed and deployed.

For those of you who are looking for further reading or kind of a user- friendly guide to this, one of the most detailed accounts of these types of issues can be found in Virginia Eubanks' book "Automating Inequality". And I'm going to touch on one of the case studies that she raises there, because it's a very useful illustration.

She writes very powerfully about how the massive systemic problems that arose in Indiana in the mid-2000s when the state moved to automate all of their welfare eligibility processes. There were record numbers of errors in the transition, for example, people being required to resubmit all of their documentation establishing eligibility for services. And if there were errors in the patients's record as a result of that, it created a finding of failure to cooperate which led to automatic termination of benefits.

From 2006 to 2008 the State of Indiana denied more than one million applications for food stamps, Medicaid and cash benefits, which is a 54-percent increase in rejections compared to the three years prior to the switch to automation. It's a really important number when you actually think about what the human impact is in, not just outright rejection, but even delays in issues like food stamps or Medicaid or other cash benefits. These are vital services that people depend on. And what we see is a really serious human impact that comes from these changes.

The Indiana story is one of experiment. We're seeing an increasing number of these issues reported in various programs in the United States and the United Kingdom, Australia and around the world. I think it's hopeful as we talk about things to think about a brief taxonomy of how problems can arise. Some problems arise from shared flaws in data entry or database linkage. So errors in systems transition from one system to another.

These are perhaps inevitable, but when they scale into massive problems when the systems lead to automatic suspension or termination of benefits. And when we think about who is most impacted by those, they are the most vulnerable members of our communities. But there were other problems beyond just those technical questions of translating data from one set to another. Other problems arose from design flaws that come when you formalize benefit policies into the code that these programs need to operate in.

For example, in the mid-2000s there was an instance in which a California program cancelled Medicaid for over 5,000 qualified beneficiaries because they failed to obtain annual redeterminations of their eligibility. In that instance, neither Federal Law nor State Law required annual redeterminations for some individuals, but it had been coded into the system. So what you had was people being penalized even though they were in compliance with the law.

We can call this a design error in some instances, but a more accurate statement is to really reflect that new policy decisions can be embedded into code sometimes unintentionally. And that can have really devastating results, and in this area in particular that agencies and people that care about administrative law really need to focus on and care about.

If you indulge me, I'm going to go deep on one more case study to illustrate the use and the benefits system. The area that we have a particular project focused on at CDT is the use of algorithms and benefits determinations and how that impacts disabled people.

This example focuses on the use of home and community-based services, credits under the Medicaid system. And in this instance, there are actually a surprising number of cases that are already being brought around the country of people having significant reductions in benefits, sometimes having their eligibility revoked altogether after the adoption of new systems.

In the Armstrong case out of Idaho, a series of decisions that came out in 2016 and 2017, Plaintiffs were a class of adults who had developmental disabilities who were eligible for home and community-based services that were funded through Idaho's Medicaid Program. In 2011 Idaho switched their system to a new algorithm-driven program that worked like this. So a human would visit the individual and complete an assessment form with track boxes representing the individual's needs.

So for example, the form would asked about feeding and ask the assessor to rate the person's need for assistance in feeding on a scale of 1 to 4. The person would then manually enter that data into a digital budget tool which automatically calculates what Medicaid would pay to cover the need. The budget tool would calculate a total assigned budget amount and generate an automatic notice that would tell the beneficiary recipient how much money they are permitted to use under their care plan.

The person could appeal that budget amount to a human, but the exception to the decision was granted only if they show an immediate threat to health and safety, which is a very high standard and a term that was undefined. And more importantly, that appeal was very lengthy, took months, people couldn't see the assessment forms where their needs were actually assessed, because the company asserted a copyright interest in it and trade secret interests as well. And then most importantly, when we think about who the recipients are in this instance, people with developmental disabilities, there was no support provided for or financially covered for pursuing those appeals.

You can see the superficial appeal of adopting a program like this. There are arguments about how it helps ground determinations in data, about how this creates a more objective measure of translating needs into a budgetary amount that is allocated. But the switch had a significant impact on participants in the home services program.

On appeal, 62 percent of the decisions were increased following reconsideration. Again, what I say what is interesting about these programs is that these have actually gone to court, so there are judicial opinions analyzing what went wrong after extensive discovery. And in this instance, the discovery revealed really telling an interesting facts about the program that need to be on people's radar if you look in its face.

One is that there were very significant design errors, so the budget tool was developed based on 3,500 participant records from earlier years, but of that sample, one-third were discarded for sample or for facial errors. So the data in there that was driving the algorithm was really egregiously flawed.

What the Court then found was that there were also instances of very significant input errors. So the person that was in the home doing the assessment had to manually transfer their scores from a number of different pages to three separate worksheets that then went in to form the budget tool. There were huge -- there were significant findings of human error in doing this, but remember, because the proprietary concerns that were raised by the company that developed this tool, Plaintiffs weren't allowed to review their sheets if it looked like there may have been an error in the system.

Finally, there was no updating or auditing of this system, so the Court observed that, although Idaho knows that the tools needs to be recalculated annually, basically to appropriately match needs to what the budgetary allowance should be, Idaho wasn't doing that. And really importantly, no auditing. So Idaho had never checked to determine how many participants were actually assigned insufficient budgets.

Because these cases are going to court, we're starting to see a body of case law develop around this. It's grounded in the Goldberg v. Kelly precedent establishing that welfare recipients have a right to adequate hearing before their benefits can be terminated and that that process includes timely and adequate notice and fair hearing.

So this has given Medicaid recipients a hook to challenge the programs that are being adopted when these types of flaws are being found. Sadly, a remedy has been much harder to come by, so thinking through how do we actually get these tools to work well still seems to be alluding many people and their smart minds that are really working on hard on this. But again, for those of you who work in relevant government agencies and also for administrative law professors, this is a fascinating area of the law, a fascinating series of cases that are unfolding that I highly commend to your attention.

And I should say that CDT will be publishing a report on this exact issue analyzing this range of cases in the coming months. And so that will be out there as a resource for folks to reference.

I'm going to pivot, Chai, if you'll allow me -- or should I pivot to employment now or should I wait for David?

MS. FELDBLUM: Let's see about getting David on and then we'll see our time. And by the way, folks, I see one question on the question check box. Please go ahead and put in your questions so that we can make sure to answer them. And David, I see we've got the old-fashioned phone here calling in. So let's see about hearing you.

MR. SUPER: Can you hear me?

MS. FELDBLUM: Yes, we can. Go right ahead.

MR. SUPER: Great. Well, I apologize for all these technical difficulties. My computer says I'm on, the application says I'm on, but I wasn't on. In any event, I want to talk about a particular problem in SNAP, the Supplemental Nutrition Assistance Program, that replaced food stamps. And it concerns something called SNAP trafficking. This is the idea of trading SNAP benefits for something other than eligible food at an eligible store, because people have severe food needs but also severe needs for other things like toilet paper and soap and whatever that SNAP doesn't cover.

This is something that happens in small numbers, people's benefits are so small that they're generally used up on food. USDA studies show that the trading of benefits for something else is rare, but it does happen and the program tries to stamp it out. And the story I want to tell, I think, has two basic lessons to it. One is how algorithms can be unfair, and the second is how an algorithm that is unfair but not very important can become a lot more important and do a lot more damage with its unfairness.

The starting point is that with the conversion of food stamps into SNAP and everyone having electronic benefits is an enormous volume of data about transactions made with SNAP and stores and recipients. USDA obviously can't investigate everything and everyone so they have used algorithms to narrow in on what are thought to be suspicious transactions. Unfortunately, what this ends up doing is identifying transactions that are abnormal, that are outside the usual patterns, but not necessarily in suspicious or dishonest ways.

For example, someone drives past a large supermarket, a Kroger or Safeway or whatever, to go to a smaller store and buys a lot of their food there. That is seen as potentially trafficking, that they know someone at the smaller store who will buy their SNAP benefits for cash. Possible, but it also may mean that that store is the one that stocks the food that their ethnic group enjoys and values and that the supermarket doesn't.

So it identifies people, some people who probably are trafficking, but it also identifies immigrants, it identifies people from racial and ethnic minorities, it identifies people who are relatively informal and value shopping from someone who maybe speaks their language if their first language isn't English or who understands other cultural preferences that they have.

Many of the other items in these trafficking-prone profiles that USDA's developed have similar dual purposes. They're not irrelevant to trafficking, but they also identify informality or people with ties to subcommunities rather than to the broad mainstream that shops at Kroger's. And this was, I think, has always been a problem, it has -- it is not helpful to have any algorithms in government that target people who are doing nothing wrong other than being members of a subgroup, such as immigrants or an ethnic minority, but it's gotten a great deal more significant.

Originally, this was used at the very beginning of the process to identify stores that would be investigated. A store that had a large number of transactions that were seen as suspicious under these algorithms would have an undercover investigator sent in, they'd try to sell food stamps or SNAP benefits for cash. If they were successful, they'd criminally charge the store, and part of the plea bargaining they'd get the store to name the people who had sold food stamps there. Now, I mean, not a flawless process, but a sane one, a sensible one, one that one can start with.

However, after the 911 attacks U.S. attorneys around the country lost interest in SNAP trafficking and were no longer willing to bring criminal charges in these cases, and without criminal charges there is no plea bargain. So the program ended upstanding on its head.

And now, someone being flagged, an individual recipient being flagged for these transactions and shopping at a store that is believed to be trafficking is used as a basis to disqualify people from SNAP. And often, because trafficking is criminal, what happens is not even a Goldberg hearing, as was mentioned before, but rather a fraud investigator often armed, often wearing a badge, tells the recipient on the basis purely of the algorithm having flagged them as doing things that are suspicious, that the government has been (inaudible) that they will have to sign a confession unless they want to be prosecuted criminally.

In almost all of the country, criminal defense lawyers do not understand SNAP, do not understand trafficking, do not understand the algorithms. And if people are prosecuted criminally, they will almost certainly have to plead to something and get a criminal record. So people say, oh, you're only going to throw me off the program I use to feed my family for a year? Well, that's pretty bad, but it's better than a criminal record. Where do I sign?

So the bulk of these things don't go to any sort of hearing. If they did go to a hearing, people are not represented, and even if they are represented, the government refuses to share information underlying these algorithms that could be used to impeach them. In essence, the fraud investigator says I am an expert in fraud, I trust the algorithm and the algorithm says that Mr. Super here is guilty, and that's the end of it.

MS. FELDBLUM: Not -- not a good story, right, in terms of any of these stories. So actually we're going to end with -- I mean, I want to leave 15 minutes for the question and discussion. So Alex is going to have a few things to say about employment issues and then Kristin, a few things to say about commenting on regs which is, of course, our favorite activity for many of us on this call. So Alex.

MS. GIVENS: Sure. Well, we're in luck because Kristin teed this up really beautifully talking about the Amazon example. So yes, to pivot from my earlier remarks about benefits, another area that we work on is the use of AI in hiring. There's an increasing amount of reliance on AI in various aspects of the employment life cycle. So this can range from determining who sees certain job ads on social media services to screening resumes and assessing candidates to reviewing employee performance on the job and far more as well, assigning shifts, et cetera.

I'm going to focus on the use in hiring because hiring, of course, is a gateway to economic opportunity -- is the gateway to economic opportunity. And there is an increasing use of these tools. To give it a little bit more of a flavor, Kristin used the example of resume screening tools that can be used to help reduce the stack from a thousand applicants for a job to a more manageable level for humans to review, and some of the problems that can arise there.

There are also examples like assessments that are based on interactive games so where people will go through a series of exercises on a computer screen and their performance on that exercise will be compared to a sample set, a pool of ideal candidates as to how the test thinks a person should respond. That also comes up in the realm of video interviews, so there's at least one company that's marketing video interviews and then purporting to run AI analysis on your vocal modalities and your facial expressions in the course of that recorded interview, and other examples besides this as well.

In each of these instances the algorithmic piece of this, and where the AI fits in, is in analyzing these tools against an idealized set of traits or a profile that's been associated by the designers with good fit for the job. The appeal for the employers is clear, right, so vendors market these tools as evidence-based hiring assessments. They claim that they make the process more objective than human review and may even help reduce bias. For example, the gamified assessments really are very aggressively marketed as being alternatives to traditional in-person interviews where human bias can skew the outcome.

In our work we look at this and say that that may be while in theory, but is still very problematic in terms of the execution for many of the reasons that Kristin alluded to in the Amazon example. When you think about what the training data set is for these tools, in very many instances what companies are doing is drawing on a sample from their existing employees. That means that existing patterns of inequality discrimination are perpetuated into the future in ongoing hiring decisions.

In addition, not only do you have the risk of an individual HR interviewer having bias, but we are now thinking about testing at scale. So value judgments about what skills and what abilities are required of somebody are now being applied at a large grand scale, sometimes not just at a company level, but when it's a vendor that is selling very similar products across companies, across entire industries or fields. That is a very dangerous recipe that I find we need to focus on very significantly.

The last piece I'm actually going to cut my remarks short just because I know you want to have time for discussion. The last piece that I will flag on this, and I think it's a good one for discussion, is that there is a lot of conversation around testing for bias in these tools. The hiring area is one where we're hearing probably the most about it, because the vendors are very eager to reassure employers that they have heard the Amazon example, they've heard other instances and they're taking corrective measures. I'll just put a flag in there that it's far more problematic than it sounds. There isn't an easy way to test who your algorithm is screening out and what the consequences are. And so we need to be really careful when we think about that dynamic as well.

MS. FELDBLUM: Thank you so much. Last sort of case study example, Kristin, talk to us about reg tech.

Nope, unmute yourself. Very good to mute when you're not talking. I try to do that.

MS. JOHNSON: Apologies. Finding it when you have a thousand apps open complicates it, right. So now I've revealed what you would see in the closet of my computer desktop if I shared my screen. But Alex's comments perfectly dovetail what the reflections I will share here at the end of the panel.

So my reflections will focus on highlighting the tensions, if you will, that we've talked about today, that algorithms, in particular machine learning algorithms, offer many and efficient and arguably accurate and effective mechanisms of executing a rote task (inaudible) in the context of algorithms, it is imperative that we are thoughtful about the effects of applying or deploying the algorithms.

And the final example that I will describe today relates to several points, that Alex just mentioned as well, that sort of grow out of some of the concerns that Alex described. So the final case study I'll reference is administrative agencies' integration of algorithms or artificial intelligence in the machine learning tools in the context of customer service or, more broadly, we might describe it as the notice and comment period of the rule-making process. We could also imagine similar platforms being deployed in the context of complaints.

Earlier in my comments, I referenced the CFPB, the Consumer Finance Protection Bureau, and the platforms that we can imagine they will deploy that would solicit comments from consumers raising flags regarding sort of predatory practices, for example, in the lending -- in the consumer lending space, right. So how would artificial intelligence be integrated into platforms for these two purposes?

Well, in the context of the notice and comment period in the rule-making process, we could imagine agencies soliciting directly from consumers in a way historically that was delayed or at least slowed by the need to receive the comments from the consumers or from various regulated entities directly, whether they be by letter or whether they be directly by telephone commentary. In either case, we can facilitate the development of thoughts that would essentially efficiently review thousands of comment letters or review thousands of consumer complaints instantaneously almost, and sort of attempt to categorize those complaints or comments in a way that is machine motivated, right.

So the machine learning algorithm would carefully review each of the comments or letters and attempt -- or complaints -- and attempt to classify them in the first instance based on the substantive remarks that are made in those comments or complaints and attempt to produce a report that would enable the agency to more efficiently review what the concerns of the citizens might be or regulated entity might be.

Well, a number of challenges arise immediately with this type of automated customer service interface. So much like Alex was describing earlier in the context of screening interviews, first, the training data set that was used to train the algorithm may not effectively capture the comments, concerns or complaints of certain groups in our citizenry.

So if someone's first language is not English, if the person's lexicon isn't sort of consistent with what we would expect them to use in terms of commonly-adopted language to describe a concern, those concerns might be under-included, right. So we can imagine immediately that deploying the bots that might facilitate customer service or receive comments or complaints as being unlikely to include or possibly -- sorry, likely to exclude and, therefore, under-inclusive of some of the concerns from certain groups within our community.

They may also misread sort of volumes of information. So were grassroots groups mobilized to disburse a significant number of similar comment letters, the bots that facilitate classifying complaints or comments might misread that as data that has been generated by a bot, in fact, faking a comment, letter or faking a complaint or faking concern. So carefully distinguishing between what is a community grassroots-based movement and what might otherwise be some type of campaign deployed by a third party that isn't generally reflecting what citizen concerns is one of the kinds of issues we can imagine arising.

And therefore, as a result, we are worried or concerned that there may be over-inclusion, right, of certain groups because the bot is unable to effectively distinguish between the kinds of concerns that are the types that we want to include and integrate into our thoughtful analysis in these context, and those that we would exclude because they are the results of automating a process and permitting others in our society to deploy technology in responding to that process.

So these are just some high-level thoughts and concerns about what could happen if we integrate artificial intelligence technology into the notice and comment or complaint processes as they currently exist. But I think they illustrate a number of the concerns in the context of the agencies' actual deployment of artificial intelligence as we've talked about over the course of our panel today.

MS. FELDBLUM: Thank you. And we got these case studies by 2:45. So we're going to be able to do 15 minutes of conversation. And I'm going to pull out one of the questions that was asked. And then before you answer, I'm also going to make a few observations from when I was a Commissioner at EEOC and we were dealing specifically with employment issues.

So the question is; what are the thoughts about setting up a federal agency for oversight of these AI/ML, you know, machine learning artificial intelligence algorithmic black boxes? Okay, so that's the question.

And I want to make these three observations and get your comments on these observations. One is that companies are going to use AI, right. So in terms of the issues that you raise as concerns, in terms of possible responses, well, I just made that as an assertion but I guess I want to ask it also as a question. Do you think that one of the facts on the ground, people who care about these issues have to take into account is that you're probably not going to be able to deal with it by not having these AI/ML approaches being developed?

The second piece is so much of this is from bias that's been there before -- you know what and David (echo/unintelligible) -- your phone is muted.

Okay. So the second is they have to use what is already there. That's what they're going to tell you. So what are the approaches that you can do to undo that, right. And then finally, my observation again, from EEOC is that there's really not a lot of conversation between the people who are substantively trying to use these AI tools and the tech people who are building it. I mean, it's just they're different worlds and different conversations. So what are the possibilities for overcoming that?

So I guess that's my three questions (echo) when we think about this thought of federal agencies. So Alex, I'm going to maybe start with you. Kristin and then David.

MS. GIVENS: Sure. I want to answer all your questions but I'm not going to. I'm going to restrict myself. So first, on the AI commission. In my mind we need an all-of-the-above approach. We need more expertise and the benefit of centralized expertise thinking about the challenges raised by AI. But we also need to think about AI's impact in specific sectors. So to me, the EEOC has to be thinking about the use of AI in hiring and it can't outsource that to a central AI commission that's going to do the thinking for it. You need the agencies, whether it's HUD, whether it's CFPB, you name it, you need the individual agencies that police these silos on the ground.

I can think of far greater understanding of how algorithmic systems are impacting the issues that they are here to serve the American people on. And that is kind of Core Mission Number 1. I do think there are benefits to coordination across agencies in some manner. Historically, the Office of Science and Technology Policy in the White House would help play that type of role. There are other commission structures that have been suggested in Congress that could help play a coordinating function so that agencies can learn from one another.

That conversation is happening. ACUS itself is doing a wonderful job helping agencies think about it. But I do think we need a more robust infrastructure, and very importantly, one that is thinking not only about affirmative uses in AI, right, so a lot of the narrative right now coming out of the executive branch is around just how do we win the race for AI and how are we kind of encouraging strengthening these tools. We also need the agencies that are the cops on the ground, policing for bias, policing for discrimination to have a far better grounding in how AI is affecting their work.

On your three questions, I may just engage with Question 2, because I want to hear what the other panelists have to say, which was about how do we get over the fact that systems, of course, just replicate the bias in the systems that they learned from, right, that's kind of an inherent flaw built in here is that if you're reliant on teaching assistance based on training data, garbage in leads to garbage out.

So there are a couple different ways to think about that. I think, first of all, in the reg tech example, there needs to be really thoughtful efforts on how you approve the range of training data so what natural language processing looks like. If you are learning from Twitter in general or from Facebook streams, you're going to get access to a far broader range of dialogs and conversational techniques than you will if you are just studying, you know, the Oxford English dictionary, right.

And so I think there's creative work to be done there and there are movements in the computer science field that are working about these issues. But the other piece of that, I think, is a really thoughtful conversation about when we need to rely on these tools and when there should be more thoughtful intervention. And employment is a perfect example of this. I really question whether the right way to hire employees is to see how your current employers are doing on a game and then hire people that play the game like they do, right.

Where is the analysis that is actually looking at what are the essential functions of the job? Like a real job analysis of what are the skills that are needed to perform this role, how do we measure whether somebody has these skills, and how do we do that in a focused and applied way. So when we work on these issues, one of the things that we really caution people against is just the shiny object of somebody, you know, they're descending from heaven telling you they're going to fix all of your hiring problems, because here's just a tool that will magically sort through people.

We don't get that, like you don't have that luxury as an employer. The responsibility's on you to really think through what are you testing people for, what are the skills that you need and how do you measure those. And I would argue that AI may be a small piece of that sometimes, but it cannot be the full answer and employers and vendors need to be far more thoughtful than they currently are.

MS. FELDBLUM: Great. So we have about 8 minutes left. So Kristin, if you could give us some thoughts and then David on these issues.

MS. JOHNSON: Hi. So I am so grateful for those questions, Chai. I want to start with your first question which essentially asks whether we could anticipate that the businesses and other institutions in our society have adopted various forms of algorithms and that the movement toward adopting those types of platforms is one that is here to stay; not likely to go away in the near future. And I think that that's a fair observation.

I think it's a fair observation for some of the reasons that we've indicated in our discussion regarding what some of the benefits of artificial intelligence might be. So we've signalled that there certainly is something more efficient about relying on artificial intelligence, and that would be highly attractive to the average business or institution or federal agency, right, being more efficient at executing a task, in particular a rote task, for which the consequences of relying on the algorithm might be less significant or severe is a highly-attractive pathway for many types of firms and institutions.

The challenge arises, I think, as Alex and David both point out, and as I hope to highlight as well, with the consequences of relying without appropriate checks and balances on science -- or on algorithms and machine learning algorithms as a form of science, right, wholly entrusting as the futurists might, the notion that the platform itself can perform as desired.

And some of the debunking that myth comes -- in order to debunk that myth we must recognize that artificial intelligence may not always be as intelligent as we would like for it to be. Therefore, the human intervention, the human in the loop really creates an opportunity for us to carefully evaluate what the probability in terms of outcomes might be for our society and the ethical implications in particular.

So I'll leave you with one last point of reflection that really kind of ties some of our examples together. In the last several weeks a number of employers and firms have begun to rely on contact tracing, digital contact tracing executed through privately-acquired platforms that Alex is describing that may have been historically deployed or at least were ramping up and focused on hiring or other elements of the employment cycle are now -- or technologists and developers focused on those, are now focused on using this technology in a different way that would impact the broader citizenry in the context of a public health pandemic.

We have so many questions about the uses of this technology that, as we begin to think about deploying it in this type of sphere, it is critical that we have thought through and began to create appropriate checks and balances, because the consequences are certainly likely to be deeply felt by some of the most vulnerable and those who are marginalized in our society and our economy. So as we think about integrating this kind of a technology in additional areas of our society, we have to appreciate it won't be going away, but there certainly must be guardrails that direct us in the appropriate uses of the technology.

MS. FELDBLUM: Great. And so David, in terms of reactions to the questions, as well as your thoughts about setting up this federal agency for oversight and probably technical assistance.

MR. SUPER: I agree that we need a all-of-the-above approach, that this is a huge problem. When I started working on food stamps and antipoverty program, the big obstacle is the lawyers. And as the lawyers constrained everything that happened, everything had to be cleared through the lawyer and you had a balky lawyer and nothing happened.

Now, the lawyers are pretty marginal; it's the programmers. And because what really happens is the algorithm, which doesn't get cleared through the lawyers, because how, the lawyers become pretty irrelevant. They write their rules and everyone thinks that's nice and no one pays any attention to them. It's between the operations people and the programmers.

Once upon a time, the Federal Government saw the importance of law and set up a super agency to deal with that called the Department of Justice. I think we're rapidly getting to a point where technology and automated decision making is as important and also needs a super agency. Not that all wall work is done in the DOJ, not that all tech work would be done in such an oversight agency, but it is necessary.

The other thing that I think we absolutely need to do is deal with transparency. There are law enforcement concerns about what I'm talking about, there are trademark and -- trade secret rather, and copyright issues about the in-home services assessment forms. And we need to make a decision that if this is going to be part of government, and it's going to be so dramatically outcome determinative that those concerns for secrecy are simply invalid and need to be overrun. And if people don't want to expose their products to public scrutiny, then they're welcome to not contract with the government.

MS. FELDBLUM: Great. So I have one remaining quick question for Alex and then some concluding comments, which is one of the questions that came in Alex is, do you have a date for when that report is going to be issued that you referenced on employment in people with disabilities?

MS. GIVENS: My team would kill me if I said yes, because it would put them to a public deadline, but it is coming soon. We had a wonderful workshop in January of all of the major litigators that have been bringing these cases, including the people that have served as the main plaintiffs which is a major undertaking when you're kind of taking on this extensive litigation against the State.

So we've done that work. It will be coming soon. A matter of -- next month, let's say that, that's safe enough. Right?

MS. FELDBLUM: That's safe enough and on behalf of your staff, thank you.

MS. GIVENS: I should just say it'll be available at cdt.org, not to do a plug, but hopefully that's useful or useless if people are looking.

MS. FELDBLUM: No, no, no. I've known it since it got started. So and also, I know that ACUS will be, I'm assuming will be making materials available and hopefully that among others.

So I -- in terms of concluding remarks, I would come back to where I started. I hope that folks have learned at least something new that they didn't know, that you have some ideas, certainly any lawyers and agencies to say, no, I do need to be relevant and I do need to be part of this conversation. And then future ideas for research and including research and work that ACUS can do.

Obviously, ACUS is in it already, but any other ideas that you have, I think, certainly I, the other members of ACUS and all the lawyers and academics who are working on this would appreciate, because this is a challenge. Just because it's a challenge doesn't mean that we don't face it and try to do something about it. So thank you Kristin Johnson, Alex -- and Alex goes by Alex, Alexandra Givens, and David Super for not only doing all this work, but sharing it with us in this panel. Everyone, have a good afternoon or morning wherever you are. Bye now.

MS. JOHNSON: Thanks for --

(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.

CMartinez

Carmel Martinez,
TX CSR No. 8128
FL FPR No. 1065