ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

ARTIFICAL INTELLIGENCE IN FEDERAL AGENCIES
Government by Algorithm

June 25, 2020

TRANSCRIPT
(Not Reviewed for Errors)

Panelists

Justice Mariano-Florentino Cuéllar, Supreme Court of California and Stanford Law School

David Freeman Engstrom, Professor of Law and Associate Dean for Strategic Initiatives, Stanford Law School; Public Member, Administrative Conference of the United States

Daniel E. Ho, William Benjamin Scott and Luna M. Scott Professor of Law, Stanford Law School

Catherine M. Sharkey, Crystal Eastman Professor of Law, New York University School of Law; Senior Fellow, Administrative Conference of the United States

Moderator

Hillary Brill, Interim Executive Director, Institute for Technology Law and Policy

Symposium on Artificial Intelligence in Federal Agencies

Audio Runtime: 1:12:04

(Beginning of audio recording.)

MR. WEINER: Good afternoon. I'm Matt Weiner, the vice chair and executive director of the Administrative Conference of the United States, ACUS, for short. We'll begin now -- we are waiting for two of our participants. We're not waiting so much as trying to connect them, and in particularly, we're waiting for -- or we're trying to connect Justice Cuellar and Professor Engstrom at Stanford Law School.

I'll begin with just a few introductory remarks before I turn it over to our moderator. This is ACUS's and the Georgetown Law Center's Institute for Technology, Law, and Policies summer symposium on artificial intelligence in federal agencies.

Thank you for all -- thank you to all of you, to all of our attendees. We have a lot of people on the phone, many of you with real expertise on our subject. And so we're really happy to have you.

For those of you on the phone who are unfamiliar with ACUS, let me just say that ACUS is an independent federal agency within the executive branch that studies and makes recommendations to improve rule-making, adjudication, and other administrative processes.

Our symposium this summer will consist of four virtual panels on four separate days. Each panel will be recorded and transcribed, and the recordings and transcriptions will be available at some point on ACUS's website.

For today's panel, we have the distinguished authors of a report, soon to be introduced, prepared for and commissioned by ACUS titled "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies".

And before turning it over to our panelists for our moderator, I just want to thank a few people. First and foremost, Hillary Brill -- first -- well, our panelists, but also Hillary Brill, the interim director of the institute at Georgetown. Two ACUS staff members who have done an extraordinary job in organizing this symposium, namely Todd Rogan (phonetic) and Todd Phillips (phonetic). And then the institute itself and in particular Hillary and Jeff Gary (phonetic), its project manager.

And with those brief introductory comments, let me turn it over to Hillary, our moderator. Hillary?

MS. BRILL: Hello, and thank you for that introduction. As you said, I'm Hillary Brill, and I lead Georgetown's Institute for Technology, Law, and Policy, and it is my privilege to be moderating today's panel with our esteemed panelists and to be part of this symposium cohosted with ACUS.

And we are the pilot program of our four-part summer series. So I hope you enjoy it. There will be different series along the way, and as you just heard, you can binge watch the entire series at the end, as well, if you love it so much. So I hope you enjoy it.

We had planned at the institute -- had planned to host you on Georgetown Law's campus. As you can see because of current events, we are unfortunately unable to do that. But we are so glad that so many of you have joined us virtually, and as discussed, in so many areas that are important to this issue, you are experts. And we are pleased to have you as part of this conversation.

Thank you, ACUS. Thank you for working with the institute on this project to make this symposium possible. And thank you to everyone at the institute who worked on this, including my predecessor, who should be named, Alex Givens (phonetic), who is speaking, I hope, on one of these panels along the way. And she is the one who worked quite a bit on the organization with ACUS. So I want to thank Alex Givens and Jeff Gary for all the work that they put in to making this happen.

The institute, it's a think tank at Georgetown Law, where we do original policy work, and we also convene with collaborators, technologists, our faculty, students, and experts like yourselves in cross-disciplinary fields, bringing together technologists and policy makers, and government agencies like today.

At the institute, we spend a lot of time thinking about how to train lawyers and policy makers to better understand the way technology impacts our society. Today's collaboration with ACUS is part of that longstanding commitment to studying the impact of AI on society.

The institute has hosted in February a symposium on AI and disabilities, previous legislative workshops, and many panel series. But today -- today is about the report that these esteemed panelists have worked so hard in putting together, and it was no easy task. This report is going to highlight the many benefits that there are to using AI systems in government. And it will also explore how agencies truly use AI, a preliminary groundwork discussion that is necessary.

These benefits are benefits we want in our personal lives and benefits we want by our government

-- decreased costs, greater efficiency, improved quality, and the ability to harness vast amounts of data. These benefits are the reasons why, as the report notes, almost 50 percent of all agencies are using some sort of artificial intelligence.

However, today's report also addresses that increases AI adoption comes with increased concerns, especially if the systems are being used by government agencies in decision-making processes and used by our law enforcement.

Many AI systems have well-documented racial biases. The software itself is often not necessarily transparent or understandable, sometimes even to its own creators. And remedies, what happens if there is a bias in AI? There often (inaudible) after the fact, and they are erroneous determinations caused by AI, and what do you do after the fact? That may not be sufficient.

So this issue of bias inherently is tricky, as bias in AI itself isn't accidental. It's part of its function. The systems are there to make discriminatory decisions. But if we can't determine whether the program is discriminating in a permissible way rather than in an impermissible or frankly illegal way, well, then we need to deeply explore how our government agencies should be using AI.

AI in government, as you will hear soon in more detail today, can provide tremendous benefits. But in some contexts, AI raises legal and moral concerns and may implicate due process rights or other civil liberties.

So thank you, panelists. We're going to explore these topics in context of their authorship and work on the report, Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies.

We are first -- we are going to hear from Justice Cuellar, Supreme Court justice of California and professor at Standard Law, who is an expert in administrative law and legislation in cyber law, and has served in the Clinton and Obama administrations and has taught as a professor at Stanford since 2001. He received his B.A. from Harvard, J.D. from Yale, and Ph.D. from Stanford.

Then we will hear from Professional Engstrom from Stanford Law, who is an expert in administrative law, Constitutional law, and legal history.

I have to say there's so much more. You need to read their bios. It's just a short summary of the incredible panelists that we have.

His current scholarship focuses on the

intersection of law and artificial intelligence. He is a facility affiliate at the Stanford Institute for Human-Centered Artificial Intelligence and the Stanford Center for Legal Informatics, and the Regulations and Evaluation and Governance Lab. He received a J.D. from Stanford, a Master's of Science from Oxford, and a Ph.D. from Yale, and has clerked for Chief Judge Wood on the Seventh Circuit.

Professor Ho from Stanford Law is an expert in administrative law, regulatory policy, and antidiscrimination law. He's an associate director for the Stanford's Institute for Human-Centered Artificial Intelligence, and directs the Regulations, Evaluation, and Governance Lab at Stanford. He received his J.D. from Yale, Ph.D. from Harvard, and clerked for Judge Williams on the Appeals Court.

And Professor Sharkey. Professor Sharkey from NYU is one of the nation's top authorities on many different issues, including economic laws rule, punitive damages, and federal preemption. She is an appointed public member of our very own ACUS, Administrative Conference of the United States, and an elected member of ALI. She received her Master's of Science from Oxford, J.D. from Yale. She clerked for Judge Calabresi of the U.S. Court of Appeals for the

Second Circuit and Justice Souter.

So we are very fortunate to have you here today. With that, we'll proceed with the presentation by each author of the report and the discussions of some of your process and findings, of course, within the short, limited time that we have. We will go into more detail as the panels -- as the series progresses. But today is a true table setting of these issues.

So if we can begin, I'd like to start with Justice Cuellar.

JUSTICE CUELLAR: Hello, can you hear me?

MS. BRILL: Yes.

JUSTICE CUELLAR: Oh terrific. Thank you. It's been a little bit more challenging to get my kids to stop using their cellphones. Thank you, Hillary, for that terrific introduction, and thank you to ACUS for supporting this project. You are going to hear more from my colleagues about what we learned. I want to give a little bit of context for why, in some ways, this report was 64 years in the making.

So yesterday Wayne County prosecutor Kim Worthy, probably many of you know, admitted that a faulty facial recognition identification was responsible for a suspect's erroneous 30-hour detention and interrogation.

And stories like this make it pretty easy to see why the public is getting interested in how government uses this mix of analytical techniques in computing systems capable of learning that go under the heading of AI.

But I want to just take four minutes to start earlier in 1956. It's a humid New Hampshire summer, and several scholars are organizing a workshop on a topic they just decided to call artificial intelligence. This motley crew is led by the quirky mathematician John McCarthy but also includes logicians, electrical engineers, cognitive scientists, shockingly enough, no lawyers. And they waste little time in sketching out an agenda that summer that is just striking to look at, the topics they were discussing, because some of the very words they used to describe their scope of discussion could be taken out of the report that we have just been working on and released 64 years later.

Building knowledge bases for digital computers, natural language processing, computer vision, and even neural networks. They're all men and confident enough to expect very rapid progress in the ensuing years.

Four years later, Senator John F. Kennedy loses New Hampshire but wins the presidency, and the arrestable and brilliant James Landis, scholar of administrative law, probably known to many of you, writes a report emphasizing to the President Elect the crucial role of administrative agencies. So he was talking about the tricky balance between political responsiveness and agency insulation, the value of government-wide efforts to make them work better, and that effort eventually culminates in the establishment of ACUS, which began operations in 1968.

It's fair to say that in the ensuing decades, at least some of the projects that Landis sketched out in his transition report to John F. Kennedy got more traction and moved more quickly than the agenda that John McCarthy and his colleagues sketched out at Dartmouth, which tended to be much more technically daunting than they expected.

But things began to change in stages. On the national security front, research never abated on AI and produced important changes in areas like avionics and even RAND Corporation advised geopolitical strategy. And a few years later, of course, the internet plus cheaper computing power brought massive disruption, and the rest is history.

So this rising interest in AI in the private sector in its current incarnation naturally triggered among a lot of us pretty intense questions about essentially the intersection of the legacy of this Dartmouth workshop and the concerns that Landis spent most of his life on like what can we delegate to AI. How can we comply with law in an AI-spiked world? How do we stress test AI technology to detect its hidden biases so we can avoid what just happened in Wayne County? How can society change its civic institutions to use algorithms in a more efficient way to write rules, to adjudicate? And how, given that change, can we define more stable goals against which to measure change?

And let's be clear. These questions are obviously relevant not only to the federal government, but speaking of Wayne County, to states and localities that spend more than 80 percent of all government dollars, leaving aside entitlements, debt service, and defense.

All this is heady stuff, but the four of us almost simultaneously ran into a problem that ACUS helped us turn into an opportunity. It was hard to engage with these questions thoughtfully when we didn't even have a basic working knowledge of how much AI was being used by agencies around the country.

So with the help of ACUS and with colleagues that

I just have delighted in working with at every turn, we set out to pursue a project focused on getting a baseline picture of how AI use was playing out in government agencies, beginning with the federal government. We recruited some superb students from Stanford and NYU to work with us. We did our best to survey available testimony, press coverage, agency disclosures. We put to one side national security agencies for others to work on in the future, and we delved more deeply into particular agencies and issues benefitting enormously from the wisdom of federal officials at a vast range of agencies, maybe some of you are on this webinar.

And since the goal wasn't just to chronicle what agencies told us but to analyze the composite picture that emerged, we have the beginnings of this -- in this report of a taxonomy of concepts and ideas, the structure and agenda of reform, and a research that will last for some time, maybe for another 64 years.

I think it's fair to say that AI use is already extensive and varied in federal agencies and will become more so, and as you're going to hear from David and then Dan and then Cathy, even the current picture offers its share of striking surprises.

But the bottom line that I want you to remember is that this report was in some ways 64 years in the making. Thank you, Hillary.

MS. BRILL: You're welcome. You're welcome. Thank you. And that was a great history lesson and true table setting. And I appreciate it. It was a great story and narrative to set -- set the stage for the rest of you.

Professor Engstrom.

PROFESSOR ENGSTROM: Sure, thank you. So I'll start by echoing Tio's (phonetic) thanks to ACUS. ACUS was absolutely tremendous in supporting this project from the start, and that ran from Matt Weiner at the top all the way -- all the way down. So thank you. We couldn't have produced a report that we're as proud of without your support throughout.

(Inaudible) quite a few agency officials on staff, perhaps, in the audience today. So this may be my best chance to thank them and to say that, you know, many of you are unsung in the report. We don't cite you by name. Agencies didn't want us to. But we couldn't have produced a report that was quite as rich as it was without your help. So thank you.

All right, I'm going to talk about enforcement. That was the part of the report where I ran point, and I don't have to explain to you that enforcement is a critical part of governance. If you have too little of it, then there's probably costly lawbreaking going on out in the world. But going after the wrong people is also costly, and it's unfair.

And so several agencies within the federal administered state have begun using machine learning to support enforcement decisions. And the report -- part 2 of the report is where the really rich case studies are, and in that part of the report, we profiled two tools, in particular, at the SEC that the SEC has fully implemented and is using.

One of those tools examines transaction data. So how do we structure data, numbers, to catch insider trading? Another tool used at the SEC parses the narrative disclosures of investment advisors. So these are registrants. People have to register with the SEC in order to do what they do, and this is very unstructured data. These are just paragraphs of text, and the SEC is using a machine learning tool to predict which among those investment advisors might be the bad apples, might be violating the federal securities laws.

There are plenty of other agencies that are developing or deploying machine learning in the enforcement space. The SEC, by no means, exhausts the set. The Centers for Medicare and Medicaid Services is using some machine learning to catch healthcare fraud. The EPA is developing some tools that will predict Clean Water Act violations. The IRS is applying some ML in the tax fraud context.

There's also interesting stuff at the state level, and we assume that lots more of this is going to come online as machine learning continues to proliferate throughout the federal government and is joined up to the mountains of data on which many agencies sit.

So going forward, I think there are three kinds of work to be done. I'm expanding on what we've already done in the report. One is to continue to surface use cases, to slice and dice them, to understand their different dimensions. And I've started in on some of this in some of the follow-on work that I've done.

Obviously, all of these enforcement tools are united by this common focus on shrinking the haystack of a pool of violators. So you can think of these systems as recommender systems. They're not fully automated. They don't fully displace agency discretion. Rather, they help agencies decide where to allocate their scarce enforcement resources.

But beyond this, if you look across the tools, they're very much a varied lot. They differ in their sophistication. They range from logical, rule-based AI to some fairly sophisticated forms of machine learning. They differ in the types of predictions that they make. They differ in their sourcing, whether they were developed in-house by agency technologists or whether they were acquired through the procurement process.

And we think that one of the great contributions of the report is to bring to light some of these technical and operational details because as we think about how we might want to regulate this, how we might want to try to build an accountability structure around these tools, those details are really going to matter.

I have a skinny five minutes. I don't even know how much I have left, but let me just say two more things by way of placeholder that might inject some ideas into the conversation that's going to follow these short little presentations.

So I think two fairly good things to think about as we think about enforcement tools and algorithmic enforcement tools, in particular, one is how these tools are going to reshape the internal agency structure and operation. And then another important question is how these tools are going to press on doctrine and force us to think about agency accountability in new ways.

So on the first of those, I'm doing some writing here, and the way I like to think about it is that these tools are going to shift the citing of discretion within agencies. And one nice way of thinking about is that these tools, as they become more and more pervasive within administrative agencies, they're going to push discretion up, over, and out.

So up, they're going to increase the managerial control of the managers over the more dispersed line level enforcement staff. Over, they're going to shift discretion to technologists. One way of thinking about administrative laws, it's an effort to allocate power within different types of agency stakeholders, among different types of agency stakeholders. So think lawyers, scientists, the political appointees of the top of an agency. And I think these tools are going to add technologists to that mix and some discretion is, therefore, going to be lodged in the technologists who will have control over the coding of the algorithms.

Second really important thing to think about is how we can build an accountability structure around these. The lawyers in the audience know that for a long time, administered law has hived off enforcement decision making from judicial review. Part of that is that we don't trust generalist judges to second guess agencies, especially around budgetary matters. But part of it, too, is that we don't think we can really reconstruct individual enforcement decisions well enough to permit judicial review.

And so here's an example where the advent or the increasing uptake of these tools could really press on doctrine in significant ways. And we might want to rethink, for instance, that hiving off of enforcement decision making from judicial review.

So that's all I'll say. Those are mostly just placeholders. I'll assume we'll come to the question of judicial review and accountability later on. I'll assume we'll come back to thinking about how this is going to alter the internal operation of agencies across different types of governance tasks. But hopefully that's a helpful injection of at least a couple of ideas into our conversation.

MS. BRILL: Thank you, Professor Engstrom. Now, Professor Ho, would you please discuss some of your issues from your part of the report?

PROFESSOR HO: Great. Thanks, Hillary. And I want to thank Georgetown and academic, Matt and Todd, in particular, for facilitating all this. David already thanked the many agency officials who participated in the research for this report.

The other really important element of all of this was the way in which we brought 30 law and computer scientists, students from Stanford and NYU together, to really wrap their heads around these issues and really peek underneath the hood of what kinds of ML techniques were being deployed. So a big thanks goes out to our students, as well.

I've been asked to just say a couple of opening remarks around how AI is being used in agency adjudication. Before I turn to two of those examples, I just want to highlight one of the first parts of the report, which is that with these students, we looked at the top 140 agencies by FTEs, really to get a rigorous sense of the extent to which agencies were deploying AI. And two basic findings from that canvas were that out of these 140 agencies, nearly half had really given serious consideration of the use of AI and machine learning.

That said, when the computer scientists started to look underneath the hood to really ascertain the level of sophistication, I think it was quite varied, and there was only about 12 percent of the use cases that were rated particularly high. Many were simply providing insufficient detail to really come up with a rigorous understanding of the level of sophistication. So there is yet a fair amount of work to be done on that front.

Let me speak briefly about two examples of innovation for mass adjudication. The first is in the Social Security Administration. As ACUS knows better than most agencies, ensuring the accuracy and consistency of mass adjudication is a major challenge for the administrative state.

So we've known for decades and decades that there can be a disturbing amount of arbitrariness in the grant rates when judges within an office are randomly assigned to cases where grant rates, for instance, for Social Security Disability can vary from as low as 8 percent for one judge to 98 percent to another judge, leading some to decry this as a form of disability roulette.

And due process, that is the kind of constitutional underpinning for mass adjudication, is resource-intensive. It can take years at the Board of

Veterans Appeals, upwards of seven years for an appeal to actually be resolved from the time that it is filed. So there is tremendous -- there are tremendous gains here, potentially, for using AI, and the story of innovation in this space really comes from the Social Security Administration's appeals council, where the head of the appeal council, Gerald Ray, was really creative in prototyping potential solutions to overcome IT hiring rules. Judge Ray started to identify lawyers who also could code and bring the kind of structured information in to develop tools like predictive tool of the kinds of cases that were likely to be easy grants, therefore allowing the agency to skip hearings and make early grant determinations.

And perhaps the most innovative tool here is one that uses natural language processing to catch errors in draft decisions. So for instance, it will parse the draft language by an administrative law judge and then look at the functional impairment that's identified in the set of facts and compare that against a kind of table of job classifications to flag potential internal inconsistencies in the decision so that judges can go and review those draft decisions.

That's an extraordinary story of innovation within government. But there are also two kind of ways in which it connects to core tenets of administrative law, much in the way that Professor Engstrom sort of alluded to.

One is that for the better part of the modern due process jurisprudence, we focused on accuracy, at least since Matthews versus Eldridge. And one question there is if you're (inaudible) hearings whether that might lead us to at least reconsider the kind of dignity prong of due process. At least we have anecdotal evidence of litigants who come in and report really knowing they're going to lose the case but finding real value in simply being heard.

And it's possible that by easing the burden of processing these kinds of cases that AI could actually recover that kind of lost constitutional value.

And the other one alluded to also by Professor Engstrom is about the internal allocation of authority. ACUS has thought a lot about the decisional independence of administrative law judges, and the adoption of something like the inside tool tends to be higher amongst staff attorneys. And so there's a question there about the internal allocation of decisional authority within the agency.

Second use case I'll just highlight briefly is an example of informal adjudication in the U.S. Patent and Trademark Office, which is also no stranger to backlogs with 9,000 patent examiners. And the PTO has been prototyping methods to improve the classification and search of trademarks and patents. So the idea is the most time-consuming part for any patent examiner is identifying relevant prior art. And if you can build better search methods to reduce that search cost, that could help cut down the backlog of the agency.

On the trademark side, one of the most innovative tools is actually a computer vision model that allows trademark examiners to take a trademark sort of application and see whether there are visually similar prior registered marks based on a kind of computer vision algorithm.

Two last points just on that example is that I do think there are really important governance questions as the prior speakers have alluded to. One point here is that we've learned from the computer vision literature in the past few years there are lots of opportunities for adversarial learning, meaning gaming of brittle computer algorithms. And so if trademark examiners no longer actually themselves inspect visually similar marks, it's possible for sophisticated parties potentially to game the trademark registration process if they know what kind of computer vision algorithm is built out. And that raises some serious questions about accountability and fairness with sophisticated parties are better positioned to fool trademark examiners.

And then the last thing that our -- the last point I'll make here is that this PTO case study highlights one of the tricky dimensions in terms of the role of contractors in building out AI solutions. About a third of the use cases we uncovered were developed by outside contractors. And oftentimes those use cases can be locked behind proprietary source code. But in the PTO case, there was an even more sort of apparent potential conflict of interest where the very contractor that was -- had built out a natural language processing-based engine to classify patents for assignment to different art units was also advertising selling the ability for patent applicants to be able to write their patents in a way to gain particular arguments.

So I think there are real kind of governance issues that need to be tackled in this space to make sure that AI is not abused in particular ways.

MS. BRILL: Thank you. And Professor Sharkey.

PROFESSOR SHARKEY: Thank you. So I want to thank Georgetown and Hillary, as well, and also give a shout out to Alexandra Givens. A decade and a half ago, she was one of my very first students, law students, and teaching assistants, and was phenomenal in that role. It was wonderful to reconnect through this work.

The second is, again, to thank ACUS. Just to be clear, ACUS is sort of like a hallmark of ACUS that they engage with academics and have us serve as consultants but then also serve as kind of a calling card to get input from a variety of different federal agencies and officials. And it's enormously helpful and productive, these kind of partnerships that they enable.

My -- I have a history with ACUS. As Hillary, as you mentioned, I'm an elected member but also back in, I think, 2010 when Paul Verkuil resurrected ACUS as its first head, I started as a consultant on a different project in which I also enormously benefitted from interviewing federal official, agency folks, many of whom might be on the call, and then under Matt's leadership was really honored to participate in this endeavor.

Final prefatory remark is just in some ways, I think our project was unique, drawing together not only, you know, academics from different institutions, Stanford and NYU, also drawing in expertise from Justice Cuellar from the California State Supreme Court, but the students that we gathered were both lawyers and technical, computer science folks. And in a way, I think our project is kind of like a microcosm of what's needed and what distinguishes this from the earlier Dartmouth project, you know, that you mentioned because it brings together both legal expertise, policy input, along with technical savvy kind of at the outset of thinking about some of these problems.

And just a smaller footnote, it includes women as well as men. Female students, we have some female collaboration. I've been in contact (inaudible) reporters who say, oh, are there women interested in machine learning and artificial intelligence, and the future is bright if we look at (inaudible) students in this project.

But a few remarks from my perspective, I want to think about some of the findings that surfaced in the report with respect to the Food and Drug Administration as a kind of window onto the future of AI in regulatory analysis. And by regulatory analysis, we include standard setting, guidance documents, and ultimately, rule-making.

And interestingly, the FDA is a great exemplar or lens into some of these issues because most people know that the future of healthcare, in particular, is going to be increasingly mediated by machines, by machine learning, by AI technologies.

And the other thing is that AI tools are extremely data hungry. This is sort of a theme that emerges throughout our report. And it's important to note that the FDA, which is the world's leading drug regulator, sits on an extremely large repository of data from clinical trials. And so the potential of being able to harness enormous datasets using these kinds of AI tools is really mindboggling.

There's also not only this existing data reserve but lots of emerging sources of data with respect to electronic health records, with respect to wearable technologies, and the like.

And so there are two main points that I want to bring out into our discussion. The first relates to something that Professor Engstrom foreshadowed, which is the way in which AI and machine learning might be actually quite transformative with respect to an agency's mission.

So the FDA is a great example of this. And from my perspective, at least, what this does, the machine learning and AI kind of fuels a transformative shift in the regulatory paradigm from being primarily a pre-market and a clearance for drugs and medical devices, shifting much more into the post-market surveillance. And that's going to harness using machine learning and AI with respect to dramatically improved ways to collect real world data and analyzing it kind of an ongoing basis.

But our report kind of uncovered with respect to the FDA there is that the FDA is kind of at a crossroads. On the one hand, they can go down an avenue of further refining existing AI tools, including primarily natural language processing. They've been using some pilots that look remarkably similar to what David was mentioning with respect to the SEC, namely using natural language (inaudible) processing to kind of sift through adverse event reports and try to figure out which ones deserve the agency's priorities and the like.

The second avenue and maybe a very, very promising one that the FDA is really thinking about is collecting more structured or focused data and the ways in which they can go directly to sources, some of which I mentioned at the outset with respect to this real world data.

Second, very briefly, a point that I want to bring out is the relationship to this idea of building internal capacity. I think the FDA, to me at least, provided some really interesting, surprising examples -- I'll mention one -- about the kinds of internal embedded expertise that's being developed.

And internal, what they call, incubator of machine learning, AI technology is called INFORMED at the FDA. It stands for Information Exchange and Data Transformation Initiative.

It was described to us, and we interviewed various federal officials who are involved with this as a regulatory sandbox. It's basically an internal incubator within the FDA of some of these machine learning AI tools.

And so to me what stood out is this is a way that an agency like the FDA can "fail cheaply". Right? The FDA is otherwise kind of an agency that would have a low-risk tolerance. Their decisions are life and death decisions. So you -- the margin for error there is pretty small. But having this internal incubator, they can try to kind of have exploration of some of these tools.

And the second and final point I'll make is that the FDA is an interesting agency because they not only are going to be using these technologies internally, they regulate AI out in the real world. So they have been approving medical devices, for example, that use machine learning/AI technologies. And so there's a way in which they have this internal incubator, and they are publishing their findings from the deployment of some of these technologies. They can search "e-risk" (phonetic) certain machine learning/AI tools that the private sector then can have more confidence as they go about using them. Thank you.

MS. BRILL: Thank you. Thank you so much to all of our panelists. And many of you mentioned all of the different people who participated in the report and the important collaboration between technologists and lawyers. And that has always been at the core, also, of what the institute is trying to do to bring technologists and lawyers together, and we are really pleased to see a report like this come from such collaboration.

You also mentioned a variety of things about accountability. You mentioned relationship building and internal capacity, Professor Sharkey. Surprising examples of AI in different government, a lot of benefits in AI because we focus on concerns of AI, but there really are benefits to be recognized here, and discussion of gamesmanship and how do we handle third-party vendors.

We're going to discuss some of these in the time that we have and then have some questions from the audience. I want to first address bias, bias in AI, we know that that's a concern, and the report notes it as well. And racial bias, especially, is a concern with the problem of artificial intelligence and machine learning systems. This is true across the board, no matter what the systems are designed to do.

You note in your report that bias can come from a variety of different factors. It can come from whether the human coders -- themselves unintentionally, it could come from unrepresented data, linking datasets that might not otherwise be connected, or other sources.

But what I want to ask you on a preliminary basis is what are some of the use cases where the potential for bias that you guys learned about in your research that the potential for bias concerns you the most. And who wants to start?

JUSTICE CUELLAR: I'm happy to start. I will just tell you that obviously, bias is a huge problem.

But I will -- and then I'll give you one context where the team that I was working with, they certainly raised concerns about it, and then I'll just highlight one of the difficulties in talking about bias in the AI context.

So we looked at some published reports and testimony highlighting how agencies that do border-related enforcement use facial -- or beginning to use facial recognition. And this is an area, of course, where not surprisingly, the full extent of what is currently happening is sort of probably well beyond what we were able to wrap our minds around.

But there was enough that we could get a sense of that highlighted some of the very -- Professor Engstrom, for example, was talking about where the shift of discretion out includes the extent to which an agency purchases a set of software that reflects an architecture for thinking about visual data and working through visual data that may have probabilities of failure modes, even in the absence of any adversarial effort to make that happen.

So when the Wayne County stuff was reported, I was not surprised. I would say there's no question that there are probably some contexts where processing of visual data can be useful to an agency and probably advance social welfare. But we have to be pretty careful.

Now, let me highlight one way in which the whole discussion of bias gets really tricky. There are going to be a lot of situations, and facial recognition will be one example where you can look for bias around race and around gender, where the definition of bias is pretty clear.

But then there are other contexts where the trade-offs are really about values, and one person's bias is another person's legitimate decision to prioritize one outcome over another.

So if you think about the choices that state regulators and NHTSA and the private sector will have around self-driving vehicles and the trolley problem-like choices that have to be made about where you put the risk, I think that we have to acknowledge that there are some blurry areas where questions that are partly technical, partly policy responsiveness questions are also, in some sense, at risk of triggering concerns about bias, given questions of who's in the room when the decision making happens, so to speak.

PROFESSOR HO: Yeah, I'm happy to follow on that. I think Justice Cuellar is right to kind of point to the use case around facial recognition technology. We have a kind of in-depth case study of the Customs and Border Patrol Agency, where there were significant errors, and it was really hard for the agency even to ascertain what the source of errors were because the system that had been built up was proprietary.

And we have a pretty substantial evidence base that documents the potential for bias in facial recognition technology. And the National Institute of Standards and Technology itself has done really terrific work on performance benchmarking of facial recognition technology that also corroborates the fact that across some 47 vendors, there are really significant performance differences in terms of the accuracy of FRT when applied to minority groups. So that is an area of real concern.

Let me say two other things. One is that the scope of the report focused more on the civil side of things, specifically carved out sort of national security and military applications, which are also some of the areas of greatest concern. And so, Hillary, when your question, what are the things things that concern you the most, in a sense, we have to be careful here about most -- you know, given that it was covered within the report, given we excluded sort of military and national security things, just because we realized it was going to be very challenging to gain any sort of transparent insight into a number of those applications.

But the kind of thing generally that you worry about that I think also connects to the civil enforcement side is what we now have seen in terms of the performance of predictive policing algorithms. So there's a great paper by a set of machine learners that basically shows that if you target and allocate police officers based upon arrests and then feed arrest data back into the model and refine it in that particular way, you could send police over to the exact same zip code over and over and over again in a kind of runaway feedback loop, even if the underlying crime rates were random.

So that is a really important thing for agencies to get right. And there is actually some important work to be done in terms of how to properly build in information as it comes in so that enforcement algorithms don't result in that kind of a runaway feedback loop.

The second thing I'll just say is that as Professor Engstrom had alluded to, there are really important doctrinal implications here. On the one hand, a lot of anti-discrimination law has shifted towards sort of anti-classification as an undergirding principle. And what we know from the past decade of work in fairness, accountability, and transparency in machine learning is that blinding yourself to features like race and gender are really imperfect ways to account for the potential disparate impact of algorithms.

And so I think we're on a kind of collision course between anti-classification and what is known in the fairness in machine learning literature, which is that the way in which we understand and address bias is by developing formal algorithms that really build in these kind of fairness constraints that rely on having measures for protected attributes to really build in the appropriate safeguards.

PROFESSOR SHARKEY: Just a quick word, too, because, Hillary, you know, as you alluded to in your question we could -- bias, typically, I think, people put in the front of their minds these issues about disparate impact on various races, genders, et cetera. There's an optimistic story about the infusion of machine learning and AI with respect to bias, too, and the FDA story kind of captures part of that, which is that this is all about data, data, data, and how representative the data is.

And so for example, if we're worried that present clinical trial data is rather unrepresentative, so it doesn't include data on all groups in society, et cetera, to the extent that these machine learning/AI tools allows us to harness lots more real world evidence coming from all sorts of different groups, et cetera, and seeing how things play out, that could lead to, you know, a de-biasing in a way that people don't typically think about sometimes because they don't raise anti-discrimination type issues.

PROFESSOR ENGSTROM: So let me just -- let me say a couple of very brief things, and I'll loop back to some things that I talked about in my introductory remarks. I think here's one of the places where really taking apart these tools and understanding their technical and operational details can matter. So think about the two SEC tools that I told you about in my -- in my five minutes.

One of those looks at conduct that's already been engaged in. If you think of that as like a reactive tool, this is the insider trading tool. You're looking at transactions already completed.

The other tool, though, is more of a -- you can call it a preemptive tool. You're trying to build a profile of a likely violator of the law, and that's like a really important distinction if you think about it when we think about age discrimination and bias type concerns. A reactive approach where we have perfect transparency over the equities markets holds the promise of perfect enforcement and perfectly nondiscriminatory enforcement, if we can capture every instance of wrongdoing.

But the preemptive approach, where we're building a profile, is essentially a kind of profile. And so there's much potential there for discrimination. So that's the first plug I'll make as to the usefulness of our report in really trying to get under the hood of some of these tools.

The other thing -- and I'll go back to from my introductory remarks -- is this is another place where algorithms really press on doctrine when you think about these enforcement tools. They are mostly hived off from judicial review. Prosecutorial discretion is really important. The Armstrong case -- this is an equal protection case -- says that we don't permit selective prosecution claims unless there's a really strong evidentiary showing of both discriminatory intent and effect, and that's a really hard case to make out to even get review of these things.

So what you worry about in the enforcement context is that because (inaudible) in such substantial ways that there's a slow burn of discrimination that can go on.

Now, the use of algorithmic tools like criminal risk assessment at the state level, those have already been the subject of litigation. It's likely we're going to get a Supreme Court case on that. There will be guardrails built around those. You might not agree with what those guardrails are, but there will be guardrails built around those particular tools.

But I guess I worry more about that slow burn of bias, then, that can make its way into some of the data analytics that a lot of similar regulatory agencies are using, and they won't be reachable under current doctrine.

MS. BRILL: Thank you, all, for -- oh wait, did you have something else you wanted to say, Professor Ho? I didn't want to interrupt you.

There are a variety of different themes that came up with how to -- how should we potentially deal with the issue of bias. In addition to just finding bias, but how do we deal with the issue of bias. And one discussion that you mentioned was, you know, we have to first define what bias is, Justice Cuellar, and that's -- that also brings up a whole variety of issues of what are standard settings on trying to stop bias, and who makes those decisions, and is it the government, or is it the person creating the actual software, and is that a private entity?

I'm not asking you answer that question, although I would love it, but we don't have enough time. I want to move on to -- because they were getting ready to answer, and I wanted to move on, actually, to the question of transparency, frankly, as a possible mechanism for some oversight because oversight was mentioned by several of you, as well, as something that should be taken into account in artificial intelligence.

In fact, in the report, one recommendation was setting up an AI oversight board, for example. But many agencies, as you know, are using AI, frankly, built by the private sector. And I believe it was you, Professor Ho, who was talking about this concern of accountability and transparency with actual third-party vendors. Some people are calling it Black Box AI. Or whatever way you want to describe it, which can significantly limit the way that we can see how those systems are being used due to trade secrets, if it's a third party, or other types of IP protections.

So what -- it's so complex now. It's not easy to explain how these artificial intelligence systems work, particularly when there's someone outside of the government who is making it. So if we can't have transparency, what do we do? What do we do for accountability? Should we not have third-party vendors? Should we ban these "Black Box AI"? What paths do you see, in general, to increase transparency and the challenges with increasing transparency?

JUSTICE CUELLAR: I know you're trying to deliberately be proactive when you say if we can't have transparency do we ban all contractors. Like, I would say, no, we don't ban all contracting, and I wouldn't give up entirely on transparency. I know my colleagues have a lot to say about this.

But let me just frame it by saying one thing about why we ought not to expect too much for transparency. I think a fair read of our report is that you don't really get an insight into how AI is going to perform in government if you just look at the algorithm or the math behind the algorithm. It's at least a function of what is the algorithm, what data will the algorithm use, what is the reliability of the computing system and network that the algorithm and the data are being processed in, how does the user interface work to present data and recommendations to the user, and how does the organization perform.

So all that is to say that really this discussion has to be about benchmarking at least as much as it is about transparency. And there are going to be some situations where an agency can and will make a compelling argument, certainly one that could present a court, I would imagine under the appropriate statutes, that it has a good reason not to share every single thing about how it's working with AI with the public for enforcement-related reasons. But do we need more transparency? For sure. We just can't expect that will solve every problem.

PROFESSOR HO: Yeah, one thing that the report highlights that became really clear to us in talking to a range of the agency officials is how important what we refer to in the report as internal due process is in terms of at least the agency having enough of an understanding as to how a tool is really performing.

So that, to go back to the SSA example, was the brilliance of actually having a person like Kirk Lays (phonetic) who was both a lawyer and someone who can do forms of natural language processing really build the system out, and it was that internal capacity that really enabled him to scope out what is a problem worth solving. And I think a quote from him -- "I developed the flags I wanted to have available as an adjudicator". And that is one of the kind of challenges when delegating something out to a contractor.

And I think at least, as Justice Cuellar said, in the enforcement context, there may be reasons not to have the decision system be completely transparent. The IRS, for instance, in its audit selection system guards very carefully how it selects because you would really be worried about reverse engineering if everyone knew exactly how audits were selected by the IRS. But what's very important is that the agency itself, the domain experts, have a really clear understanding.

And Professor Engstrom may be able to give us a little bit more insight into that dynamic at the SEC, where it was very much the sort of lawyers who were demanding greater transparency and intelligibility of the risks for selecting cases.

PROFESSOR ENGSTROM: Yeah, I can speak to that briefly, which is as we talked to the SEC officials and staff who put into place some of the algorithmic enforcement tools being used within the agency, they noted -- so take that second tool that I profiled, the one that tries to figure out which investment advisors might be the bad apples.

So these are predictions that are generated in some central part of the agency by technologists. And then those outputs are then handed off to line level enforcers. And what we heard from the officials and staff as we talked to them about the implementation of this tool is that those line level enforcers are not at all impressed by being told that, hey, this fancy machine learning system threw a flag as to this investment advisor but not this one, you know, but not that one.

They want to know why. They want to know, you know, why the flag was thrown. They want to know which part of the narrative disclosure threw the flag. And so this gave us at least some reason for optimism that the different -- that the splitting off of, say, the technologists and those line level enforcers in the agency actually creates its own internal form of due process and demands for explanation.

MS. BRILL: Would you like to add anything, Professor Sharkey? I also -- I also wanted to say to the audience that you can provide questions. We're open to a question-and-answer session, so please do provide your questions. But Professor Sharkey?

PROFESSOR SHARKEY: Yeah, just very briefly. So first, I'll defer, of course, to Justice Cuellar about what kind of arguments judges would find persuasive on judicial review with respect to transparency or not --

JUSTICE CUELLAR: I was talking about hypothetical judges.

PROFESSOR SHARKEY: Right, right.

JUSTICE CUELLAR: And hypothetical courts.

PROFESSOR SHARKEY: Mysterious point, though, comes back to what I said earlier about our project being in some sense this microcosm. I think that what it means in terms of having sufficient level of explainability, reasons given, is a question that is a merger of legal expertise, policy expertise, and scientific savvy.

So groups, for example, like NHSTA that were mentioned before who are coming up with standard settings in various areas, I think they are onto the idea that while they have these laboratories with scientists who are developing what's possible in terms of the scientific capability, it behooves them to reach out to legal policy analysts, not waiting -- sort of like the idea like we'll develop the technology first, and then the law will give us a thumbs-up/thumbs-down. Getting the legal policy input along with the generation of this emergent technology, I think, is really critical and key.

PROFESSOR HO: Yeah, I'll say one other thing, just in terms of where the case law is headed here. Many folks in the audience will know of the Wisconsin Supreme Court's case in the Loomis decision, where the criminal risk assessment score was challenged under due process.

And one of the claims by Eric Loomis in that case was that the risk assessment algorithm was closed source, and it was not possible to actually know how the risk score was calculated.

And the way the Wisconsin Supreme Court disposed of the case was to conclude that because the inputs into the algorithm in the presentencing report were all available and transparent, there was no due process violation in not being able to peek underneath the hood.

And I think if there's one thing a report really highlights is that it's going to be important for the future of algorithmic governance to actually understand how something really is engineered. It is not enough, given the complexity of models, to say the inputs are all transparent because there may be many things going on underneath the hood that are going to be important as a policy and as a legal matter.

MS. BRILL: So I want to open questions up to the audience, if we hopefully still have some audience members with us because this panel was scheduled to go from 1:00 to 2:15 and to allow you to have some opportunity after you heard the panelists to ask some questions.

And one question that was brought up by the audience was specific -- oh, to submit a question, please type it in the chat box. So please open up your chat box. All of you have one. It is there. There is a part that says questions. Please go into the questions, and if you have one, please go ahead and submit it.

One question from the audience was what part of the APA do you think AI challenges the most.

JUSTICE CUELLAR: Oh, I love that question. I'm actually curious to hear what my colleagues say about that.

MS. BRILL: Great.

JUSTICE CUELLAR: I'll just throw out what counts as arbitrary and capriciousness has been zipped up kind of at the heart of administrative law for a while, a portion of it, and I think that this is a really great moment where the question is being culled in a new way by the intersection of user interfaces, algorithms, data, and organizational performance. And I think we'll have to be a lot more specific about what that means.

PROFESSOR SHARKEY: Yeah, I would just echo -- I mean, a bunch of work that I've done outside of AI has focused on how courts increasingly try to scrutinize kind of the empirical basis for what makes something a reasonable decision on the part of the agency.

And so the focus, you know, this idea of what part of machine learning/AI that gets infused, particularly in the regulatory rule-making context, is going to, I think, require probably the most work.

MS. BRILL: There's a question -- oh I'm sorry, Professor Engstrom.

PROFESSOR ENGSTROM: I was going to say this is just a give a boy a hammer moment. Everything looks like a nail. But I think this (inaudible) enforcement has always been this tweener governance task. It's always existed in this kind of limbo. It's both a wholesale and a retail endeavor that has an adjudicative component. It has a rule-making component to it. And so I think that there is going to be some very interesting sort of near to midterm thinking that needs to be done about whether we might want to reshape ex-post review, allow relatively more challenges to enforcement decision-making that would require some kind of an amendment to the Heckler V. Chaney line of cases.

It's also conceivable that we would want to declare algorithms and algorithmic systems of various types rules that are -- that must go through notice and comment. That doesn't mean every algorithmic system would have to do so, but some would, and we could try to think about smart line drawings to determine which types of algorithmic systems do have to be pushed through that process. But -- and I think it could be done.

But I think that -- I do think there'll be some very interesting thinking there, and it does present all the usual trade-offs between ex-post review of enforcement decision-making or other types of decision-making or in a notice-and-comment context, we could kind of call that ex-ante review.

MS. BRILL: Thank you. There is a question specifically for Professor Ho. I'm going to put you on the spot. Professor Ho, can you speak more on the potential for reverse engineering of systems and overly transparent systems? Are there any policy remedies to mitigate them?

PROFESSOR HO: As a law professor, it's nice to have the tables turned on me to be cold-called (inaudible) like this. So just so I understand the question, what are the potential remedies for kind of reserve engineering, I take it?

Yeah, I think it's going to be one of those fairly domain-specific inquiries. So for instance, in the trademark context or the patent context, are there kind of good faith obligations. And so if there is someone who has used a reverse-engineered kind of mark to evade the computer vision algorithm, it's an open question whether that would potentially violate sort of the practice rules in front of the PTO.

But my other colleagues may also have other insights here as to the kinds of concerns that arise. There are, of course, other instances where transparency and like reverse engineering, if you've got the incentives set up right, may actually be the desirable thing.

So for instance, I think this is the interesting contrast between sort of benefits algorithms and enforcement algorithms. In the benefits context, if part of what happens is you're making much more transparent the conditions under which you're entitled to a disability benefit, that may not be reverse engineering. You may actually be more crisply communicating what the eligibility criteria are, and that may actually be something that is desirable, given the documented amount of discretion that we do see in making disability determinations.

PROFESSOR ENGSTROM: So can I tack on just a little bit there? Which is if you really look into the emerging, quite excellent literature on transparency around AI, both in the private sector and in the public sector, there's lots of interesting conversation about what transparency means, what types of explanations would satisfy it.

So a distinction you'll often see is between decision-level transparency and system-level transparency. Decision-level transparency is where a person might be entitled to some very thorough explanation as to the provenance of the particular decision. But it could be that transparency is hashed out quite nicely by more system-level explanation. Like what are the -- you know, what are the basics of the model, what are the basics of the data inputs, things like that.

And so I think here, it's important to note that the logics and the imperatives of different governance tasks are really different. The logics and imperatives of enforcement are very different from, say, the social welfare benefits context.

And so we might want decision-level transparency in the welfare context. But we can't provide that in the enforcement context because it kills the tool. If you open source the tool, you kill its usefulness.

MS. BRILL: So in a sense, you're saying it should be context-specific for potentially what level of transparency we have. And it's quite interesting -- we don't have enough time for everyone's questions, but many of the questions revolve around transparency and what can be done to -- to solve the issues with that.

And I can't say what's in the minds of the people that are writing it, but I think it comes from the not understanding and the concern and then everything around the fear. And maybe if we understood what these types of systems were doing, then we would feel a lot more comfortable.

And I think as -- for those of you who have questions that weren't answered, there are there more panels on this report, and there will be more discussions about these issues.

So I want to end on this question. As you're all lawyers, and you're all technologists, and you all have put together this great work, will Congress now need to enact new statutes to govern how agencies use new AI tools, or will existing statutes be adequate to the task? And let it be known, this is an audience question. I'm not cold calling you. Any professor can jump in whenever they feel comfortable. But what do you think about what Congress needs to do?

PROFESSOR SHARKEY: So I'll jump in with -- I'll take it on small bore in the following way. There was a dispute that surfaced that's actually not reported in our report because we didn't focus -- we didn't do a use case study for NHSTA, but we did review various officials for NHSTA, et cetera, and there was a disagreement as to whether or not their current regulatory status allowed them to mandate that car manufacturers give them direct data or not and what this meant for the future of their use of machine learning and AI.

And so I guess before -- what I would want to do is before answering that question about Congress does or doesn't have to do, it would be worth studying existing regulatory mandates, agency by agency, and engaging in these debates about what -- you know, how you can push the limits in terms of saying that the agency already has the authority.

And I would center many of these questions around gathering of data, at least for agencies like NHSTA and FDA that are going to be these regulators of health and safety out in the real world.

PROFESSOR HO: Yeah, I also am reluctant to speculate as to what exactly Congress should do. But I do want to kind of answer this broader theme, Hillary, that you're pointing out as to concerns about transparency because I really do think it goes back to the point about internal capacity. That is the most complicated machine learning models right now could be sledgehammers to kill a fly. That is, there are -- there's a kind of complexity/accuracy trade-off. But it may not always require a sledgehammer to solve particular problems. And the people who are going to be best situated to really understand that are going to be agency staff who have an insight into these tools and understand the domain. And that's really where we think some of the biggest gains are likely to be made.

There is also some real fruitful models here, for instance, in terms of academic agency collaborations to start to bring in some of that insight in-house into the agency. So I'm not sure about any specific recommendations, but I think it is important to think about the internal capacity within agencies to navigate that transparency, complexity, accuracy trade-off in choosing which tools are suitable for the problem at hand.

JUSTICE CUELLAR: I really like Professor Ho's answer, and I'll just build on it by making the following observation: there are two considerations that are normative and one that's descriptive, practical, that might inform any discussion of this topic.

The first normative consideration is how much we can realistically imagine a trans substantive approach to AI that is going to make sense at a high level of generality like APA style for NHSTA and for FDA and for the SEC and for SSA and so on. And I have my doubts.

I think part of what is helpful about this report is to highlight how there are definitely cross-cutting themes, but there's a lot of context-specific, subtle work that really is right at the intersection of some fairly bespoke technical issues and some very bespoke legal and factual and politically common issues that should live at the more specific level.

The other normative consideration is how much we want to preempt the kind of experimentation happening not only in the private sector but in the states. Laboratories of democracy with one or another crosscutting solution.

The descriptive, practical observation is Congress has had real trouble passing even basic cybersecurity legislation. It's not exactly in a particularly productive period of its history for any number of reasons. So query whether we can expect a lot of action.

MS. BRILL: We didn't ask if Congress would. Well, I guess -- if they should.

JUSTICE CUELLAR: Just throwing it in as an added bonus for what it's worth.

MS. BRILL: I don't want to cut you off, but we have just a moment in closing, and I wanted to thank all of you for the great insight and hard work that you put into this report and the table setting.

I mean, you brought up all the issues. You brought up all of the themes. You did most of my work. You just summarized the fact that a lot of this is context-specific, and a lot of the themes are across the board discussions that other people seem to want to hear about are about transparency, are about what is bias, what does that mean, and who decides it, and what are the actual agencies that are going to be determining that and setting those standards, and what is the public/private connection with that relationship.

And you said this is 64 years in the making. I don't think we have 64 more years to answer those questions, although -- you know, it could happen that someone else could be talking about remember that webinar back 64 years ago. But we do need good discussions, and the rest of this symposium will hopefully bring it to the experts that joined us. So thank you all today for -- plenty of thank you to the audience, and thank you again to ACUS and everyone who helped make this possible.

MULTIPLE VOICES: Thank you.

(End of audio recording.)

CERTIFICATE

I, Wendy Sawyer, do hereby certify that I was authorized to and transcribed the foregoing recorded proceedings and that the transcript is a true record, to the best of my ability.

DATED this 9th day of July, 2020.

WENDY SAWYER, CDLT