ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

ARTIFICAL INTELLIGENCE IN THE TRENCHES
A View From The Agencies

August 13, 2020

TRANSCRIPT
(Not Reviewed for Errors)

Panelists

Marco Enriquez, Senior Applied Mathematician, Office of Data Science, Securities and Exchange Commission

Kurt Glaze, Program Analyst, Office of Analytics, Review, and Oversight, Social Security Administration

Sean Khozin, Global Head of Data Strategy, Janssen R&D; Formerly Associate Director, Oncology Center of Excellence, Food and Drug Administration

Krista Kinnard, Director, AI Center of Excellence, General Services Administration

Moderator

Stephen Sanford, Director, Center for Strategic Foresight, Government Accountability Office

Transcription of Audio File:

Artificial Intelligence In The Trenches_ A View From Inside The Agencies

Audio Runtime: 59:04

(Beginning of audio recording.)

MR. WIENER: Good afternoon. I'm Matt Weiner, the vice chair and executive director of the Administrative Conference of the United States, ACUS for short. Welcome to the fourth and final panel of our symposium on AI in federal agencies, which ACUS has been pleased to cosponsor with Georgetown's Institute for Technology, Law, and Policy.

Today's panel, like the prior three panels, today's panel will be both recorded and transcribed. The video and transcription will soon be available on the ACUS website.

In the first three panels, we largely heard the perspectives of experts outside federal agency, eight agencies today. We'll hear from several innovators in and practitioners of AI within federal -- inside federal agencies, and we're very thankful that Stephen Stanford, director of the Center for Strategic Foresight at the Government Accountability Office, GAO, has joined us to moderate our discussion today.

I'll turn things over to the Institute's Jeff Gary in a minute. I'd first like to thank, though, him and the Institute's director, Hillary Brill (phonetic) for cosponsoring this symposium with ACUS, ACUS attorneys Todd Phillips and Todd Rubin for

organizing this symposium so expertly, all of our panelists, of course, and especially our many attendees, 200 today I'm told, many themselves experts in AI.

With that, let me turn it over to you, Jeff, with my thanks once again.

MR. GARY: Absolutely. Thank you, Matt, and I'd also just like to extend a great thanks to ACUS and especially to Todd Phillips and Todd Rubin, who have been doing just phenomenal work on this, planning even before we switched to a webinar format.

Everyone, I'm Jeff Gary. I'm a project manager at the Institute for Tech, Law, and Policy at Georgetown. The Institute is a think tank. It's run through the law center, and we focus deeply on some of the cutting-edge questions raised by new and emerging technologies. And we've been so pleased that we can be a part of bringing this program to life with ACUS.

You know, we believe strongly that as AI and new technologies develop, they bring new opportunities, and we've been able to explore those. But at the same time, we can't dismiss the challenges that these technologies pose to existing social issues such as systemic racism and discrimination and the lack of social mobility.

And I really feel that this series has shed a strong long on really the nuance and the realities of utilizing and engaging with these technologies, and we've had a really strong and open discussion about where things are working and where things are not and the benefits and challenges faced by agencies.

And with that in mind, I'm so glad that our panel today focuses on the realities of the agency's use of these technologies and how they're really grappling with them.

We'd love to stay in touch. So please do follow us on Twitter. We're at Georgetown Tech Law -- sorry, Gtown Tech Law. And visit our website, which is GeorgetownTech.org, where you can sign up for a newsletter if you want to attend other events like this.

On that note, I'm going to turn things over to Steve Sanford, who's going to introduce our panelists and get things moving into the discussion. Thanks again, everyone.

MR. SANFORD: Thank you very much, Matt. Thank you very much, Jeff, for those intro remarks. Very pleased to be here today. My name is Steve Sanford. I am the director of the Center for Strategic Foresight at the U.S. Government Accountability

Office. Much like ACUS, we are interested in improving the efficiency and the effectiveness of the federal government. So that's certainly a mission that, as an independent, nonpartisan agency of the legislative branch, we feel very akin to.

Very pleased to be joined by the panel today. In my work at GAO, I led GAO's first technology assessment on AI in 2018. I've done some work on deep fakes, and I'm currently working on a program of work related to AI governance, and the Comptroller General is going to convene an expert forum this September and will be subsequently issuing some findings from that work. So we're really excited to be a part of this panel today.

And without further ado, I would like to introduce our panelists. So I'm going to ask our panelists just to say a few words about themselves. Let's start with Krista, please.

MS. KINNARD: Yes. Hello, everyone. Thank you so much for having me. It's really a pleasure to be here. I'm the director at the Artificial Intelligence Center of Excellence, housed within the Technology Transformation Services at GSA. And so because we sit in GSA, we really -- we have a government-wide focus. And our main mission is to partner with the rest of

the federal government and across federal agencies to help accelerate the adoption of artificial intelligence in government agencies. And we do that in a number of ways, both through the sharing of best practices and lessons learned in our federal AI community of practice and also through direct partnership through our centers of excellence program. So thank you very much.

MR. SANFORD: Thank you, Krista. Welcome. Sean, please.

MR. KHOZIN: Hi, everyone. I'm Sean Khozin. I must highlight the fact I've recently left the federal government. So I believe I'm the only one here who is not currently in federal government.

So currently, I'm the global head of beta strategic, data science innovation at Johnson & Johnson Pharmaceuticals, and a few months ago before joining J&J, I was at the FDA as associate director of the (inaudible) Center of Excellence and also as a founding director of Information Exchange and Data Transformation that we launched in 2015 with special authority from the Department of Health and Human Services as an incubator at the FDA for developing new capabilities in artificial intelligence and other data science-related areas.

I'm a thoracic oncologist, and I've been doing data science academically and also in the private sector for a number of years going back to my (inaudible). Proud to be part of this conversation.

MR. SANFORD: Thank you, Sean. Kurt, please.

MR. GLAZE: Yes, thank you, Stephen and ACUS and the Institute for having this panel today, and I'm honored to be here. My name is Kurt Glaze. I formerly was an appellate attorney at the Social Security Administration, and now I am a program analyst there with the Analytic Center of Excellence within SSA. I'm also the creator and lead of a decision support application project known as Insight, which is a software product designed to help disability program adjudicators with their adjudication activities at the hearings and appeals levels of adjudication. So thanks again for having me.

MR. SANFORD: And finally Marco, please.

MR. ENRIQUEZ: Hey. My name is Marco Enriquez. I'm a senior applied mathematician in the U.S. Securities and Exchange Commission. I run the largest analytics program at the SEC that's (inaudible). We utilize and evaluate AI technologies in support of the SEC's mission.

Also, I'd like to add thank you to ACUS and also the views that I'll express today are my own and don't reflect the commission's. Thank you.

MR. SANFORD: Thank you, Marco. Appreciate that. I'll add my own disclaimer that any remarks that I make today are also my own and don't necessarily reflect the institutional views of GAO. So thank you, Marco, for that reminder.

Let's hear next from all of our panelists. I'd like to hear from them each in series. And if you could tell us a little bit about your organization's journey and AI and machine learning -- where you started and where you are now, what specific ways your agency is using AI and machine learning to fulfill its mission. And as you've been going through that, what have been some of the key enablers of success as that journey has been underway?

So again, we'll start with Krista, please, if you could do that. And I'd also invite our audience, who's listening live, please feel free to post questions during these initial comments from the panelists. We'll do a quick follow-up round after these remarks to answer some of those questions, and then we'll do more questions at the end. But we do want this to be interactive. So if you have any

questions as you're hearing from our panelists, please put those forward. So we'll be hearing from each panelist for about five minutes. So starting with Krista, thanks.

MS. KINNARD: Great. Thanks so much, Stephen. So the Artificial Intelligence Center of Excellence and the broader (inaudible) AI portfolio is a very new organization. It was set up last October. So we're almost a year old in light of the executive order on maintaining American leadership on artificial intelligence.

And so we sit in a part of GSA called Technology Transformation Services, whose entire mission is to engage with the rest of the federal government, as I kind of alluded to earlier.

So we're not really doing development within GSA's internal programming. We are a program that exists to partner with the rest of the federal government. And so we do that in a number of different ways.

And so really the whole genesis of this and why we have an Artificial Intelligence Center of Excellence was that agencies are excited -- and I think we're about to hear from a group of folks who are going to show you that agencies are already

starting to embrace artificial intelligence. And I think that there's a lot of really interesting and engaging use cases around how this technology can really be applied to help agencies meet their mission.

But there's a lot of questions about it, right? I think everyone here would agree that AI is not a solved technology, right? You can't just buy one AI and get it out of the box. It's complicated, it's nuanced, and as agencies start thinking about, you know, what are my principal business and mission challenges that I'm trying to meet, how can AI as a technology be used, and if I really want to use artificial intelligence to support my mission, what does that look like? What do I do? Right? And so that's why we exist.

And part of what we do, right, so part of what we established when we set up this group is a community practice. Because first and foremost, what we know is that we're not in this alone, right? And going off and trying to build individual artificial intelligence systems in a black box is not going to serve anyone.

So we wanted to create a space for the federal government to come together. And we do through a series of webinars. So far, we've led a series on acquisition in AI, right? I mean, there are some

agencies who are going to develop a lot of this in-house. They've got deep technical staff to be able to do this.

But if we're speaking realistically, acquisition and procurement is going to be at least part of an AI solution. So really trying to figure out how can the federal government be smart about what we're buying, right, is really something that our community wanted to learn about.

The other thing that we've done, we've done some tech talks, right? We had a Ph.D. computer scientist come in and start talking about how you can start using some open sourced tools to start playing around with data in your organization.

And really the goal here is to bring people together. And I have just been overwhelmed by the number of people who are just really excited, and they'll say something like, oh, I heard such and such agency can speak to your community. I'm trying to do something similar. Can you connect me, right, so that we can -- we can share those lessons learned and share -- there are other things, right? Share resources, share tools, share frameworks, right?

And our community is incredibly strong. We have almost 1,000 members, and we are less than a year old.

And they're really engaged, which I think is -- is incredibly powerful but also speaks to the demand and the interest in bringing this technology to the table.

And to be honest with you, I've just been incredibly impressed by partner groups who have spoken up, and they just -- they ask such good questions about things like how do you do this responsibly, how do we mitigate bias, right, how do we start thinking about privacy when we are building these solutions, right?

And I don't know the answer to all of those, right. And really the community isn't here to solve your problems. The community is here to bring those to light so we can solve them and start addressing them together.

The other side of what we do through the centers of excellence is we do actually partner, right, one-on-one with agencies, and more broadly our centers of excellence program focuses in six different areas. So we have our Artificial Intelligence Center of Excellence, which is our newest, but we also have a longstanding Data and Analytics Center of Excellence infrastructure, cloud, customer experience, and call center modernization, right?

And we have kind of this demonstrated history of

being able to partner one-on-one, bring in our technical expertise to drive projects. And we engage on an executive level because we know that for these types of projects to really gain traction, to really have the impact that I think that they really should and can have in a way that is meaningful, in a way that is bringing value, in a way that is responsible, you need to have that championship, right? You have to have those leaders stepping up and saying this is something that is important to our organization, and this is something we want to do, and it's something we want to do well.

And that's really where our centers of excellence thrive, begin able to get in there, roll up our sleeves, and with an organization, speak to their executive leadership and go across the organization to say we need these partners to come together to actually do an implementation.

So I lead one of our engagements with the department, and we're doing some automation and bringing in intelligence to some of the processes that they're looking at kind of in their procurement cycle.

So I'll stop there. I know I've taken up a lot of time, but I -- I'm happy to field questions about what we're seeing in the federal landscape, what are

some of the common challenges not just in one specific agency but across agencies, and where are agencies both finding opportunities and challenges to come together and really make an impact in artificial intelligence adoption.

MR. SANFORD: Great, thank you, Krista. It's remarkable to hear about some of the momentum right now behind AI and the awareness of the issues in the federal space, maybe even just compared to two or three years ago. It sounds like there's been quite a lot of movement there, and I'm sure we'll come back to some of those issues you've raised.

Next, let's hear from Sean, who's seen both the experience of how machine learning is deployed on the federal side with FDA and then also in the private sector. So Sean.

MR. KHOZIN: Sure. So the FDA, when we look at the application of AI and machine learning at the agency, I think one can divide it into three different categories. The efforts that are aimed at automating workflows and business process. So there are a few pilots that, when I was at the FDA, I was involved in and that are being entertained right now.

And those are some of the same themes that any organization and any business can take advantage of.

And in fact, at J&J, we're also applying machine learning and piloting some of these methods in streamlining and potentially automating certain business processes. So that's one category.

And the second category is the way that one applies machine learning and AI to analyzing biomedical data assets. For the FDA, that's about, you know, approving drugs (inaudible) devices. And for us at J&J, it's about extracting insights from the data we're generating as part of our trials and developing programs.

What's interesting is that, you know, the FDA in a lot of cases have been applying modeling and simulation and more recently methods that can categorize as AI and machine learning in very specific contexts.

A great example is pharmacometrics. Pharmacometric review of drugs involves a lot of modeling and simulation, and the FDA does have great expertise in that area, and they've been increasingly applying machine learning methods.

And again, we are also applying machine learning to a lot of the pre-clinical discovery work that's being done. And (inaudible) component, and there's been several publications. One was by the FDA, and I

led several AI efforts on the academic front that produced published academic papers that we felt could advance the priorities of the FDA, the mission of the FDA, and also provide actually insights to the industry in terms of, you know, the art of the possible, in some cases.

Because FDA is in a very unique position where it has access to essentially all the drugs and biologics that have been ever approved in the United States, and that data can be leveraged in very unique ways and a perfect substrate for AI.

And then the third component of that is how the FDA approaches the approval of AI algorithms, and there are several pilots there. For example, (inaudible) has (inaudible) that address that, and there are a number of different programs that are underway. And as many of you may have heard, (inaudible) technology modernization plan that is aimed at providing a framework for scaling data science and also AI machine learning across the agency.

For Jansen and J&J as healthcare and pharmaceutical company, we are incorporating (inaudible) analytics, including machine learning methods, across essentially the entire drug

development continuum and bringing in new talent and building infrastructure, and that was one of the many reasons that I joined the company and was really providing valuable insights to (inaudible) some of our programs and also providing an opportunity to just ask fundamentally different questions that one typically doesn't ask as part of, you know, traditional drug development paradigm.

And these methods do allow researchers and drug developers to really ask fundamentally different questions that without applying such methods wouldn't -- we wouldn't be able to answer.

MALE VOICE: But once you justify handing them out to law students, you can't justify not handing them out to anyone basically who wants one.

MR. KHOZIN: Yes, was that a question? I believe I heard a question.

MR. SANFORD: I'm not sure if that was -- I'm not sure who asked the question. Well, so Sean, was that the end of your remarks, Sean?

MR. KHOZIN: Yes.

MR. SANFORD: Great, thank you. Thank you very much. So Kurt, next from you, you already mentioned a little bit that you're actively deploying some tools in the fulfillment of the mission. So I would love to

hear about that.

MR. GLAZE: Sure. So I suppose before I begin, I'll give a little bit of back story. I started my career as an appellate attorney at SSA for its disability program. And then in 2015, after years of engaging with that program from, you know, a purely legal side, I pitched software -- a concept for software, and that's now called Insight, and that is my current major project, and I've been working on it in more or less fulltime since 2015.

So some background on Insight is essentially in a nutshell, it is decision support software, again, designed to provide adjudicators with a series of things. First of all, it reads the text of disability decisions written at the hearing, written or reviewed at the hearing or appeals levels of the disability program. We have multiple levels of adjudicative review.

It extracts information from the (inaudible) texts of those decisions, combines it with other information about the case and claim in our current systems, and ultimately offers feedback to adjudicators about potential quality issues that are present on the face of the decision.

In addition to that, we offer them contextual

reference information. So you know, kind of targeted reference information that is specific and helpful for that claim only. And we also offer them a series of tools that are basically enabled through the information that we extracted upstream.

All of this is designed to help adjudicators improve the quality, efficiency, and consistency of our disability decisions in service to the public.

So I think to understand the why as to why Insight, you know, we even rolled this out or why SSA funded the development of Insight, it helps to kind of understand how SSA adjudicated claims prior to Insight, and I think the SSA story pre-Insight is very common among mass scale adjudicative agencies in the federal government.

So prior to Insight, the workflows essentially consisted of individual attorneys at either the hearings or appeals levels working individually on a case, you know, preparing a work product, and kind of proactively seeking out, based on their training and experience, the resources they need, looking up the right regulatory or subregulatory guidance for the claim, and then passing their completed work product to another lawyer or adjudicator for independent review by them.

And if a quality issue was found or additional work needed to occur, that other individual would send it back to the first individual. But -- so you guys can glean that this was a very manual process, independent action by independent action.

So how Insight breaks that paradigm is it intervenes at the individual level to try and proactively bring relevant information to the individual rather than them constantly having to seek out that information.

It also, as I said, its probably most notable feature is its capacity to potentially flag quality issues that merit further attention by that individual before they push their work product down the line. And as you all can imagine, it is much more inefficient to deal with these quality issues if they're embedded in work products downstream such as on appeal, or after the work product has been handed off to another staff person than it is to address them as soon as possible when the worker is right there looking at, you know, potentially thousands of pages of evidence in an individual case and trying to digest all of this.

So that is the aim of Insight. As you mentioned, Insight is fairly far along. It's not conceptual.

We've actually deployed Insight software to all adjudicators and adjudicative staff, by which I mean essentially kind of like law clerks, attorneys who write decisions, or at the appeals level, appellate attorneys who conduct an initial review of an appealed case.

We've rolled that software out for use by all attorneys, essentially, at the hearings and appeals level, and that's been the case since essentially -- at both those levels since 2018.

I'm happy to discuss this later, but we've also conducted several analyses of Insight's impact in business value for Insight's disability program and have seen some positive results both in improvements to the quality of the decisions as far as we can tell but also to the efficiency of adjudication, which at SSA I can tell you as such a large body is extremely important that we do whatever we can to serve the public efficiently.

I think some of the most important success enablers that led to Insight -- Insight wasn't developed in a vacuum. We benefitted at SSA from having an electronic system, which some adjudicative bodies are, you know, more advanced or less advanced stages.

We also had a fairly robust case -- electronic case processing system, and both of these systems upstream were kind of the foundational building blocks upon which Insight could possibly build some of the technologies that it has.

SSA is also -- in this recent decade underwent many efforts to globalize data and enterprise data warehouses and otherwise make experimentation more feasible for teams across the agencies.

So some of those have been major enablers for us that, you know, we stood on the shoulders of those to build Insight. So once again, thank you for having me.

MR. SANFORD: Thank you, Kurt. I appreciate those remarks. Sounds like you've come a long way in that journey from 2015 and continuing to move forward.

Marco, please let's hear from you.

MR. ENRIQUEZ: Thank you. So like Hurt, allow me to just maybe give a little bit of background. So my interest with machine learning really just started in graduate school. I have a doctorate in applied mathematics, and the dissertation topic that I chose was in the field of optimal control theory, which is really the precursor to what people call reinforcement learning in the AI community.

And when I graduated, I joined industry around the time where big data was the latest buzzword, and it was really an exciting time because everybody was trying to really analyze and parse like hundreds of gigabytes and terabytes of data and also considering how to apply machine learning algorithms to scale.

After a few years in industry, I decided to kind of like steer my career towards civil service, which is how I ended up at the SEC. And at the SEC, my assessment really is that I would say that we're kind of at the tail end of the kind of proof of concept and kind of prototyping phase for some of our AI programs.

It's really exciting because some of the more successful programs we're seeing a move to systemize and rapidly deploy these systems at the enterprise scale. Furthermore, our infrastructure both in terms of compute and also data storage, it's rapidly maturing.

At the SEC, we use artificial intelligence technologies for a lot of different applications, and hence, it was really important for us to kind of carefully think about these applications and categorize them.

So we've come up with two course buckets, really -- and it ties to really the algorithm's potential to

harm the targets or sort the consequence levels, if you will. So I'm going to discuss those two levels, so going to talk about the high consequence bucket and the low consequence bucket.

With regards to high consequence type models, think of surveillance and risk assessment workflows, so how we find insider traders or us applying risk assessment algorithms to our registrants.

Obviously if we were careless in how we deployed AI models, so if the AI had some implicit bias, we could cause a lot of reputational harm to someone. And furthermore, we would cause harm for ourselves because we would lose the public's trust in our ability to really carefully and responsibly use these technologies.

With high consequence level algorithms, you know, and I think my colleagues will agree, it's really important to have humans in the loop. I think it's really irresponsible to deploy these systems end to end. So you will not be seeing any time soon, right, an algorithm that is the judge, jury, and executioner in a lot of these cases.

To my knowledge, no AI system is sophisticated enough to understand securities laws. So it will not be supplanting our human subject matter experts any

time soon.

So for us, especially in this high consequence level kind of bucket, AI is just an extra piece of evidence that a human or humans might use. Furthermore, it's conceivable that there are actual multiple models kind of helping kind of, you know, derive insights to support an outcome.

Okay. And finally, one thing I kind of wanted to discuss is that with this particular type of task, it's really important when feasible to use explainable AI technologies, right? For us, it's not enough to say, hey, this firm -- the model says it's high risk. The lawyers at the SEC will just laugh in my face.

What really they need is, right, an explanation, right, where should I look next to find corroborating evidence. As Krista had said, however, I want to note that these technologies are far from solved. In fact, in the field of natural language processing, interpretability is still very much an open problem. Right? So we have to be careful about our use of these technologies in that space. Nonetheless, we should try and also evaluate to see if it's appropriate.

Finally, I wanted to discuss, right, this low consequence bucket. And so you hear about these types

of tasks a lot less simply because it's not really as exciting. Right, we kind of are interested in stories about like, you know, potential future Sky Net in the making, and these tasks are not that. But I believe that they're a huge opportunity area.

And my research staff and I are actually investing a lot of time here and really -- the crux of this is that rote tasks can be automated in part or fully using artificial intelligence. It also is nice because you don't necessarily need to worry about explainable AI or even things being a black box fully because these tasks, again, are fairly low level and inconsequential that, you know, if you get a few wrong, it's not really that big of a deal.

But these tasks, if solved correctly, could save a lot of time. And if you roll up the time saved that really amounts to millions of dollars, right, across enterprise. And you know, we're not even going to talk about the intangible benefits here too.

And so what's nice about this is that at the SEC, really, we consider AI technologies as an enabler, right? So we want to enable staff to do more with less. And really -- so these capabilities exactly do just that. And again, I wanted to push forward this idea of the low consequence use cases because I

believe that there will be plenty in every organization, and if you solve them, the resource savings can be big.

With regards to enablers, I think having mature and democratizing competition in platforms is important. When I first joined the SEC, you were tied -- your power was tied to your (inaudible), and that is no longer the case. So right now, we have the capability to self-service fairly high-powered machines, even some of them have GPUs enabled, so that we could write train models.

Also I think just really bringing the IT division to bear with collaborating with them instead of having an adversarial relationship, and I could discuss more on that later. I think that really transformed our capabilities, and we were able to progress really rapidly as a result. Thank you.

MR. SANFORD: Thanks very much, Marco. Really insightful. I really like the distinction you drew between high consequence and low consequence AI. I think we all have some intuitive understanding that an algorithm designed to make movie recommendations is very different from one that's designed to assist with sentencing in a criminal justice environment, and I think that kind of paradigm is very useful to keep in

mind as we all think about this.

I wanted to come back to something Krista mentioned in her opening about the level of partnerships and collaboration going on and maybe drill down a little bit there. Ask where, when you look at the federal space and you see this community coming together that's focused on trying to advance machine learning and AI into mission support, where do you see the greatest traction occurring for partnerships or collaboration? Is it with data sharing? Is it developing principles? Is it actually, you know, joint application development? Where is the most traction happening, and what do you think are some of the prospects for future such engagement?

MS. KINNARD: Would you like me to start?

MR. SANFORD: Yes, please, Krista. Yeah.

MS. KINNARD: That's a really great question, and I think that there's a lot of collaboration that we're seeing. So I also want to echo your point, Marco, your calling-out of high consequence versus low consequence is so important. And you're right, it's not glamorous, but honestly, that low consequence AI is actually where I think there is the most opportunity for partnership and collaboration.

So you can look at this from the perspective of using AI to directly meet your mission. There's a whole lot of functions that happens across many organizations that support mission delivery. And so one of the areas that I'm seeing a lot of collaboration across is those mission support areas.

So procurement, for example, every agency, I don't care who you are, you got a procurement office. Maybe they do things a little bit different, but there's regulations that dictates how procurements happen.

And so if one agency has used artificial intelligence or automation to help with a part of that, that's something that can be shared.

And then I think another big area where we're seeing a lot of opportunity for partnership is in policy and governance models. So it's one thing to build an AI solution. But I mean, it can't exist in a vacuum, right? There needs to be oversight. There needs to be people looking in on it. There needs to be people monitoring it saying this is continuing to do what we expect it to do and meet the need of the organization in a way that is providing value, right?

And so I know on projects that I have worked on, I've heard of maybe an organization not working on

exactly the same project, but that project that they did that was different prompted them to, I don't know, make the security checklist that is really valuable and continue to share the checklist, right, it's -- it's something that can, at a very minimum, inform what we're going to do at this agency, right? Or policy or governance around how are you sharing data, how are you assigning ownership of different datasets, of different data systems, of different models, right?

And so that framework for how people are setting this up in their organization is absolutely something that I think can -- can be shared. So I'll stop there. I'm sure the rest of the panel has comments and ideas around that as well.

MR. SANFORD: Yeah. So if I could maybe pivot to Sean and ask Sean to weigh on this concept of partnerships, and I think you've seen partnerships, again, through two different lenses in your previous role and your current. But maybe talk a little bit about the nature of the relationship between government entity and the private sector when it comes to come of these things.

MR. KHOZIN: Sure. I think when it comes to data science and AI, (inaudible) partnerships are a critical component of really starting to understand

how to manage high-risk use cases but also how to scale some of the existing successes.

So when I was at the FDA working on the (inaudible) I mentioned before, that was really -- the essence of that was to develop portfolio (inaudible) collaboration. So we started to work with a number of start-ups in this domain, and in fact, published a number of papers on foundational AI research as it pertains to amassing the interest of patients and developing new therapies. And so partnerships are a critical component of that.

I wish there were more partnerships, public and private, around (inaudible) because I think in a lot of cases, we do have expertise. We do have the technical infrastructure, but the data isn't there. And data sharing among federal agencies, for example, NIH and the FDA, and also in the context of a public/private collaboration enterprise can be highly, highly beneficial.

FDA is engaging in several public/private partnerships, and obviously, we all know -- we've all heard things are happening in the context of COVID-19, developing a vaccine, and (inaudible) NIH has, for example, collaboration (inaudible) called active consortium, and data sharing is a component of that.

But there's another effort to (inaudible) collaborating with a nonprofit organization called Project Data Sphere on developing data sharing frameworks, and there's a workshop coming up, I believe, in October that talks about (inaudible) registries, sharing data to develop new insights in rare cancers.

But we have to think about new ways of scaling responsible sharing of data. What's quite interesting is that there are creative ways that in the past couple of years federal agencies are thinking about in terms of enabling (inaudible) medical data in the evolving (inaudible) economy.

One of them is the efforts that is being led by the Office of the National Coordinator, basically the body within NHHS that's responsible for certifying electronic health records.

In 2016, as part of 21st Century (inaudible), there was a mandate directed to (inaudible) to eliminate what's called information bulking electronic health records, and it was traditionally -- did not share the data, even in some cases for patients. It's been very hard to get their own date. And typically still things are faxed around and in best-case scenario, it's put on a CD-ROM.

Now, NC currently as part of certification will be mandated in January of 2021 will mandate open APIs, where anyone, any individual can gain access to their own health data, and they can share it with anyone they wish.

So putting the patients at the center and in charge of having their own data sort of shared with researchers, companies, the federal agencies, as they choose.

So that is likely going to be transformative. In fact, those APIs have already been incorporated into the Apple health kit, which has now a new electronic health record future, where you are able to extract a lot of the data, your own data, in a structured fashion.

So that -- these policy decisions in combination with the organizational public/private partnerships are already evolving and I believe are a critical component of really extracting maximum use from existing machine learning methods and to be able to develop better and more nuanced and more predictive and precise algorithms.

MR. SANFORD: Thank you, Sean. Thanks very much. I wanted to turn to Marco and ask a question on the partnership factor as well. Again, with how you're

interacting with the private sector on some of these things. And also to come back to a point you made, the types of internal partnerships that were required within SEC, you mentioned working with your IT support function on some of this. So if you could address those two, I think that would help folks understand some of the issues too.

MR. ENRIQUEZ: Sure. So at the SEC, obviously, we have a lot of registrants, and a lot of them are technologically advanced, and they're looking for some guidance with regards to how they utilize AI. And we often have to walk a really fine line. We are enforcing securities laws. We don't enforce good practices and best practices for AI and ML use.

And frankly, again, going back to what Krista said, a lot of these kind of things, concepts in AI, again, still an active area of research. So they're not solved.

So it goes back to this notion of, frankly, really (inaudible) for us to try to issue guidance on some things that are not solved, right? But we do try to engage our registrants, and we actually have financial -- it's called Ben Hub (phonetic). It's our financial innovation hub. So registrants can request to meet with us to try to present ideas and maybe

solicit some feedback.

And so that's been actually exceptionally really successful. To echo what Sean said, though, we do have some barriers. So data sharing is still a huge kind of undertaking. In fact, it's virtually nonexistent, to my knowledge.

Furthermore, that's also true, you know, like Krista, we tried to work with other kind of financial -- or other agencies just in the financial regulatory domain across the globe. And even in that space, sharing data and sharing code is really difficult, and so we have to kind of get the blessing of a lot of lawyers in succession to even share like slides, for example.

So I think that there's a lot of work to be done and maybe common MLUs and frameworks to basically facilitate just sharing of information because I think it really would be to everyone's advantage -- and there are a lot of common tasks that we all need to solve. And that's certainly true for the financial services industry.

With regards to our IT staff, I will say when I first joined SEC, it was like guerilla data science. You just had to install whatever you needed to do to install it on your computer to just get the job done.

And because asking IT often meant waiting usually a week or two, which is unacceptable, especially now for those of you guys who've actually coded, you literally type in a command in your computer, and it installs instantaneously.

And so that whole two weeks of wait time is just ridiculous. But you know, recently, I would say the last two years or so, we started really talking to IT folks and saying, look, we need to come together and come up with a solution because you guys don't like what's happening, which is basically people circumventing your policies behind your back. We don't like the long times it takes to really install, quite frankly, like common components that data scientists use.

So de facto now, we made Python, for example, the primary language for scientific computing at the SEC, which led to a lot of like kind of IT sanctioned support. So now, I could actually type -- you know, install this module on my computer and have it installed instantaneously.

At the same time, we know it's been sanctioned by IT, so they -- you know, it's safe and, you know, won't violate any sort of security protocols. And I think that's just a really kind of good example of us

just coming to the table, discussing kind of like pain points, and then just trying to see, like, work in IT really help alleviate those pain points.

And so by engaging the business more, I think we have now a really great IT division that I think is making a lot of big strides, again, to facilitate a lot of these data science and AI and (inaudible) workflows.

MR. SANFORD: Thanks, Marco. You know, there might be some federal agencies that are actually -- would be envious of a two-week turnaround for software. Everything's a little different depending on where you are.

I want to come back to something Kurt mentioned earlier, but first, given the time, I want to invite the audience members to -- attendees, please submit your questions. If you have questions, we'll try to get to some of those in the last minutes of the webinar.

And as you're doing that, as you're composing your questions and sending them in, I wanted to ask Kurt, just to follow up, you had talked about beginning to assess the effectiveness of some of these systems. Could you tell us a little bit about how you're doing, what you learned from the process of

evaluating these types of systems?

MR. GLAZE: Yeah, sure. I think I can give -- to start, I think it'd be helpful to have a broad overview of how we've attempted to assess and ensure the value and quality of what we're putting out in this systems.

So first during the development process, we engaged heavily with subject matter expert attorneys in the disability program to design some of the discrete decision support features that we offer through the program and to validate them prior to release. That's normal processes.

But once they're released, that then transitions into the evaluation stage, as you mentioned. So some of the studies -- I have not personally executed these, but we have statistical staff at SSA who partnered with us to study outcomes of use of our decision support software in disability cases during our phased rollout.

So initially, we were a voluntary use product, so kind of a natural experiment as we trained and advertised our product to staff, and they could voluntarily choose to engage with Insight for a period. So we had a kind of natural use and nonuse group at scale, and our statisticians, for instance,

to measure Insight's impact on case processing efficiency executed a regression analysis to evaluate based on our case processing system the timeliness of case closure when Insight was used -- essentially, when Insight was used versus not used as part of the workflow.

And I wish they were here because they could explain these much more elegantly than I could, but effectively, the result was an observed reduction in case processing time both at the hearings and appeals level.

That makes sense, to some degree, a modest reduction because we do offer decision support features that are designed to improve the speed with which they work these cases, but we also offer quality feedback, which of course necessarily entails taking another look, looking back at your work.

But net, we saw reductions. So that regression -- and it takes a lot of data, a lot of voluntary use data, a lot of nonuse data, and a lot of time to partner with folks who aren't necessarily familiar with your product and how it works and how it fits into the business process. So those are large studies.

In terms of the quality of decisions, some of the

studies we engaged in were evaluating the frequency of Insight-covered quality areas, the frequency of quality issues in those Insight-covered domains in issued decisions, again, where Insight was used or not used kind of as a natural experiment by our staff. And again, we saw consistent reductions in Insight-covered quality issues in final work products coming out of our disability adjudicators at both levels.

We also conducted some in-depth manual case studies of specific quality flags that Insight rolls out where we had a team of attorneys study actual cases where we raised a given quality flag and actually have eyeballs on all work products, you know, including manual kind of case analyses that were written by our attorneys, basically on an across the board look at everything they possibly touched to see how they reacted. Did they catch the issue more or less with Insight's help? How did it affect their ultimate work products in ways that are very difficult to measure through existing structured data sources about our case work?

And again, in both of those, we did see improvements in the recognition of quality issues with Insight as an intervention in cases where it raised that -- those flags.

So those are some of the studies, and this is, you know, in addition to, you know, many engagements with our users that also go to, as Marco said, kind of intangible value of these products, which are things like the dignitary interests of our claimants, the workflow experience of our staff working these cases. Like, what is their day to day experience with an assistive tool like Insight versus not?

And so we do conduct outreach. We've previously conducted surveys, all sorts of engagement mechanisms to try and get a sense for the reception of our staff and their perception of its value. And so far, we've been happy that the feedback has been positive, generally. So those are some of the ways.

MR. SANFORD: Thanks, Kurt. One quick follow-up from the audience to you, Kurt, and then I've got another audience question I think we have time for to the whole group. The question to you, Kurt, was some curiosity about what does the Insight tool actually return back to the user. Is it a reading list of things to look at. Is it actually marking up the decision? Is it making alternative text recommendations? What sort of actual product is it delivering back as its recommendations?

MR. GLAZE: That's a very good questions. I wish

I had visuals because this might be a little easier. But it's a web application that essentially when they kick off or access Insight, it pops up essentially a web browser page that displays quality feedback as line items for the user to review.

And I should underscore again that Insight, as a decision support product, is never the final arbiter of any element of a disability claim's adjudication. It is always an advisory service, much like if you had a personal assistant helping you with a case who looked over your draft before you moved it forward, and they offered some observations about what they saw, that this exactly what Insight is doing except at scale.

And so it pops up a web application. We provide any quality feedback we have to offer, and there may be none. In something like 45 to 50 percent of cases that are sent, we really don't have any substantive quality feedback to provide to the user. Their work product seems fine as far as Insight can tell.

We conduct about in the 30s about 30 discrete, specific analyses of elements of quality. So we are by no means comprehensive in our analysis.

But anyway, our feedback is provided to adjudicators, and it's really a jumping-off point for

further analysis by that adjudicator. They can engage with it. They can think -- you know, they can agree, disagree, think that they've covered the issue and rationale or that it's covered by the facts of the case. They can change the decisional language to adapt to the quality issues cited to hopefully remediate it before issuance.

But that's where we fit in to the workflow.

MR. SANFORD: That's great. The last question to the group and actually I'm wondering, Krista, maybe you can just answer this quickly, if you can, we got a question from the audience how are agencies addressing or planning to address ethics in AI? Is there -- are there efforts specific on the ethics question?

MS. KINNARD: Yes, but that's a great question, and the answer is yes. So again, ethics is something that's not solved, right? There is not any one single person who says I know all the things about AI ethics and responsible AI implementation.

There's many different groups thinking about this, and in fact, I think every agency who is thinking about AI is starting to think about this. So we actually have a working group as part of our community practice specifically focused on this.

We are not policy makers, right, we do not create

frameworks, we do not create policy. But we are creating a space where federal government can come together and share the resources they have, have those tough conversations, and learn from folks who may be a little bit further down the line, for example, the DoD released their ethics principles. You've got other organizations and groups of organizations starting to release their ethics principles. So you've got a lot of folks thinking about this.

Is there a go-to framework for how to solve AI ethics and responsible AI for your organization? No. But there are a lot of smart people thinking about it.

MR. SANFORD: Great. Thank you, Krista, and with that, I see we're at time. I want to thank all of the attendees at today's webinar, and in general, I want to thank ACUS and the Institute for Technology, Law, and Policy at Georgetown for this symposium series. I think we had some fantastic insights from our panelists today. I appreciate their time and wish everyone well in their machine learning and AI journey. Thank you very much.

(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 3rd day of September, 2020.

WENDY SAWYER, CDLT