Past Event: 2024 NCSBN APRN Roundtable - Regulatory Decision Making for Nurses: Is the Regulatory Environment Ready for AI? Video Transcript

Event
2024 NCSBN APRN Roundtable

More info: https://www.ncsbn.org/past-event/2024-ncsbn-aprn-roundtable

Presenters
Anna van der Gaag, CBE, Visiting Professor, Ethics and Regulation, University of Surrey;
Robert Jago, MPhil (Cantab), Professor of Law, Royal Holloway, University of London

>> Anna: Good afternoon, everyone, and I'm really delighted to be here presenting this research to you today, and I am joined by my colleague, Professor Robert Jago from the University of London. He's going to be appearing shortly on your screens, and he's going to talk about the qualitative aspects of our work together. But the first thing I want to say to you is that I am not a data scientist, I'm not a technical person, but I have spent many, many months in the company of people with impressive technical skills, and without them this work would not have been possible. So I want to first of all say thank you to my colleagues at Royal Holloway, and I also want to thank you NCSBN for taking a leap of faith with us back in 2018 and funding this work at a time when AI was hardly even in the public discourse, and certainly wasn't being talked about in regulatory circles at all. So we were extremely fortunate to work with the Texas Board of nursing, who were instrumental in making this happen. I again want to particularly thank -- [listing names] Who were fantastically supportive of what we are trying to do, and work very closely with us. So, is regulation ready for AI? Every generation brings innovation and change, and I don't know how many of you recognize this little box, but it is a room thermostat designed in the U.S. in the 1980s. I can remember my grandparents looking at this little box with great skepticism. They certainly didn't think that a machine was capable of calibrating the temperature in a room without the intervention of humans. But these boxes were probably an early form of AI, and they did things that humans had previously done. In this session is about exploring a modern day equivalent the room thermostat, and inviting your reflections on whether you feel, as regulators, program certifiers, leaders of professional associations, and clinicians, are you comfortable with this development? Let me start with some definitions. What do we mean when we talk about AI? There are many, many definitions, but at its most basic, AI is technology that imitates aspects of human decision-making, using mathematical calculations, probabilities, to predict and interpret data. We see it everywhere in our daily lives, on our smartphones. We know that it makes decisions all the time about what we wear, what we eat, what we read, and companies are using these technologies to help predict our habits and to sell us what they think we are most likely to want on the basis of those calculations. And ChatGPT, which is probably one that you are almost familiar with, which stands for generative pretrained transformer, is one example. It is a

form of generative AI, which can be used to create completely new content, text images, videos, using these technologies, these algorithms. But, actually, most of the artificial intelligence used in health research to date has been nongenerative AI, and it relies far more on direct human intervention and control. Another term you will see, machine learning, natural language processing techniques, these are all forms of AI, and they learned from data patterns, either in supervised ways or in self-supervised ways. That the distinction we make between generative and more traditional forms of AI. For example, the AI used to improve the speed and accuracy of diagnosis in cancer has undergone years of rigorous supervised learning using tens of thousands of images and millions of pieces of anonymized data to arrive at those accurate predictions of risk. And research has shown that these kinds of tools are actually better, more consistent, at arriving at the correct diagnosis in humans. So this technology is advancing with incredible speed. This graph, which was published last December by the U.S. trade and technology Council, showing how AI is being used across an industry. To improve the speed and accuracy of processing information. That underpins so much of the sorts of initiatives. This data actually comes from 2019. You will see a more recent trajectory. We can probably all agree that AI is here and here to stay. We certainly need to debate the dangers it brings, its uses and abuses, in many jurisdictions where they are now putting in place regulatory frameworks to protect us from the misinformation that is being generated by these tools. Our question for today is, is our regulatory environment, our Nursing Regulatory Body and in, ready for this? Robert and I certainly felt this was worth exploring. We had been researching, complaints-handling, in health professional regulation for many years, and we wanted to see whether these tools could be used particularly in disciplinary work, so complaints that are made about nurses, there is a very formal process I'm sure you're familiar with. A complaint comes in and it is assessed as to whether there is a high risk of harm whether the case needs to be taken forward through illegal process or not. But what we do know from the evidence is that there's a small number of what I would describe as high risk nurses. Nurses who are harming patients. And it comes to regulators around the world, requiring no regulatory action. There is no removal of licenses. Up to 70% of cases are in that category. Rob and I were particularly interested in whether we could use a tool like the ones we are talking about to distinguish between high-risk and low-risk cases. So the decisions could be made more quickly, particularly about the cases where there was very unlikely to be any regulatory action taken. Could we separate the lower order risks from the high order risks using technology? Through that, improve the speed and accuracy. What we were seeking to test here is whether we could develop an AI tool to provide multiple types of information to aid decision-making. Starting with risk level, calculate the risk level from existing complaints data. And then, to link those cases to regulatory rules and standards. Finally, to provide a link to previous similar cases so that the case manager who is making the assessment of that complaint, in that case, has access to a far greater range of data than they would otherwise. And this means that, if the tool is successful, it could potentially allow teams to access data on tens of thousands of cases, and use the risk categorization and outcome for those other cases to assist them in their decision about the current case. So that is the aim in a nutshell. But, really important, as with all AI developments, getting the governance and design principles right is essential. And we used government guidelines to ensure that we adhered to these throughout the project, and we referred to David Leslie's work as the guidance that we use, and you will find in our published work. Really important that those design principles around fairness, accountability, sustainability, and transparency were adhered to throughout the project. The methodology -- I'm not going to spend too much time on this. If you are interested in the technical side of this, please do have a look at some of the published papers that had been written by my colleagues. But for the quantitative analysis we used 5,700 disciplinary cases, complaints data from the U.S., and that was the Texas Board of nursing. From the U.K., and the Australian nurse regulated. We used that data to build the tool, built the prototype,

incorporating text from complaints from the status of the case, the source of the complaint, the human judgment on the risk level, and the degree of harm that was reported in the text. We tested the accuracy of these risk predictions I've talked about using five different AI classifiers. Those of you who are in the audience and you know about these classifiers, we use gradient boosting, adaptive boosting, CNN, and an unchaste classifier, and a combination of three of those to test the reliability and accuracy of those risk predictions. We also were very keen to recognize bias in the data. We started by using gender techniques because we had sufficient information on the gender of those cases. Our plan is to go on to use, when I have a bigger data set, to be able to test out and check for race, ethnicity, and so on. So debiasing techniques are pretty standard within the suite of tools we are talking about, and it was very important that we were able to give assurance that the risk predictions were not based on gender. Just as eventually we want to show that these risk predictions are not based on race. Finally, we did a lot of qualitative testing with regulatory staff, which Rob is going to talk about shortly. But I'm just going to show you now some screenshots of the prototype we've developed. Here you have the landing page, see have a secure login for any case manager who is using the tool, and what the case manager sees first is this page, which shows some of the text for each of those complaints, and then the independent calculation of the risk category. That is under their predicted risk column. Then you have predicted probability and confidence score, which is essentially calculating the probability that the risk category is correct, and the confidence that the case is showing similarities to other cases. As you see, the final column is left blank, because that final column is for the human, for the case manager to make the final decision, using the information from the tool but not basing the final decision on the tool. That is entirely for the human to make that judgment. The next page gives the case manager a little more information about how the decision has been arrived at. So the complaint text used by the algorithm, the keywords that are used to make that risk classification. It gives transparency to the process and an assurance of how the tool has arrived at that decision. And then, as I said earlier, the second task is to relate the case to relevant parts of the regulatory code, and finally to find previous similar cases allowing the case manager to make comparisons on the classification of the case. So this is not about replacing human judgment, but allowing case managers to obtain this independent risk calculation or classification to triangulate that with other cases and with the relevant parts of the regulatory rules, all in a matter of minutes. But the decision is a human one, and humans have the last word. Finally from me, before hand over to Rob to talk about the qualitative findings, I'm just going to give you one slide on the reliability results. This shows you the published results of the first 1241 cases. What it shows you is that all five of those risk classification systems that I talked about significantly outperform the majority baseline in terms of accuracy. The ensemble model, which is a combination of C1, C2, and C3, gives the highest level of accuracy. And we anticipate the accuracy will increase with the increasing number of cases that are added to the prototype. So, if that hasn't sent you to sleep, really looking forward to Rob talking to you about our qualitative findings, and then I'm going to come back and wrap up at the end. Rob, over to you. >> Robert: Thanks very much, Anna. I wanted to confirm that I am Robert Jago, and I was one of the P.I.s on the research project. It was my role to coordinate the quality data, which I'm going to talk about later. But I also wanted to confirm that in our research team we did have a variety of perspectives that ensured a range of voices informed the direction of the project and the lens within which the research was actually conducted. You will see that we have, for example, computer scientists, a lawyer, which was me, and it was my role, for example, to resolve the data sharing concerns that emerged. We have a clinical and regulatory perspective, which understood how the complaints process routinely impacts the lives of practitioners, and we also have a social scientist he was able to ensure that our research methods were effectively utilized. Now, we are also mindful as a research team that, when considering the ethical concerns in the AI space, we were drawn to the seminal work of Michael

sandaled. In 2018, he identified three key ethical concerns. Firstly, he talked about the concern as to privacy and surveillance, and here he was particularly concerned about the need to collect vast amounts of data, which in turn can be used by a range of operators for a range of different purposes. Secondly, there are concerns around bias and discrimination, and a key concern here is, just as human judgment can be biased, whether that is unconscious or conscious, so the data we collect can also generate biases. In the final ethical concern which is raised by Sandel is related to the role of human judgment. Here there are concerns that AI will replace humans, whether that be through replacing them in their work or the willingness to focus on the potential gains of AI rather than being reminded of some of the potential losses. In our qualitative work, we ran three online focus groups with nurse regulators. One with ten nurse regulator participants in Australia, a second with 11 nurse regulators in the U.K., and a third with seven nurse regulators in Texas in the U.S. So we had 28 participants overall. The online focus groups were actually facilitated by three members of the research team. Qualitative data from the focus groups was then analyzed in accordance with the framework, and our method of data analysis actually utilized the six phases of their thematic analysis. We did this as follows. We familiarized ourselves first of all with the data. This meant transforming, reading, and noting initial ideas. We then generated initial codes, so the coated the data across the three focus groups using transcriptions, we then search the themes and we were able to collate codes to potential fames, make a thematic map, and as we develop to those we defined and named the themes, which meant we had to refine the specifics of each themes and the overall story that the analysis revealed. Finally, we produced our report, which finalized the themes and supported those themes with data extracts which related findings to the research question, and the literature, and of course formulating a discussion. So, in terms of our themes, we can see these on the slide. Negotiating trust and trustworthiness. Our participants focused much attention on trust, mistrust, and trustworthiness in relation to the inclusion of AI and an AI tool in decision-making related to complaints and nurse regulation. The focus group participants were acutely aware of the consequences of error if it is to practice processes, and the fact that these were acute. Because poor decision-making can result in registrants either being able to continue to harm patients, or being incorrectly judged to lack fitness to practice. Either way, the consequences are serious, and they remind regulators of their responsibilities to ensure decision-making processes are trustworthy. Within this theme, our participants talked about prioritizing honesty and transparency, ensuring a balanced appraisal of human/machine capability, and the need to think about impacting language and whether AI and its varying connotations could impact on trust of any subsequent developments. The second theme highlights participants' awareness of some of the most challenging ethical considerations which relate to AI/machine learning in professional regulation fitness to practice processes. When looking at affirming fairness and nondiscrimination, there was a concern in this theme that outputs from AI tools could actually incorporate bias and result in unfair and discriminatory decision-making. The three sub themes relate to minimizing bias, the use of an algorithm which results in bias or unfair decisions, and relation to race, age, or indeed gender, for example. There's also a feeling of the need to avoid fabrication. Where it needs to be understood which values or elements are focused on in a decision-making tool or algorithm, which could present something of a quandary. As some elements, there is a need to ensure the data is reliable. Related to this is the final sub team, which is to ensure accountability. Participants explored the objective versus the subjective nature of decision-making in this area, and the importance and challenges of clarity regarding accountability. Nuanced decision-making is challenging for humans, and it is therefore also challenging for a learning support tool, but for different reasons. Whilst it was recognized that there would probably be technological developments which would assist this process, it was felt by the participants that we should always remain alert to the potential for discrimination. In the final theme here was managing burdens and benefits. It was a strong awareness in our three focus groups that

regulatory decision-making is complex, and the need to take into account context, uncertainty, and ambiguity. Hence this third theme of managing burdens and benefits has sub themes of shades of gray with a need for consistency and to negotiate complexity. There are also fears of falling through the cracks, whether decision-making has significant consequences for all, and the need for professional judgment. And also, regarding what any AI decision support tool could deliver in the context of professional regulation. Here it was felt there was probably a need for humility as to her expectations of AI. And then, finally, there were discussions around effectiveness and burden reduction, and the potential benefit of any tool vis-a-vis the emotional content of fitness to practice complaints. What, then, do you think the future holds in this area, based on our research? The first point to make is, if we are to utilize AI tools for this purpose, we need to ensure that we engage in the deliberative process with a range of stakeholders. We also need to balance the enthusiasm and caution and recognize that humans have bias, whether conscious or unconscious, and this informs the position that they will take in the deliberation. Now, we identified and argued that participants tended to fall into 1 of 4 groups. The first of these were the ideologues, and they are absolutely committed to the idea that AI is the future, and they can have a tendency to minimize some of the concerns raised about the extension of its use. The second are the romantics, and they are perennially optimistic about the future of AI. They believe that it has the potential to remove human inefficiencies and biases, and this can only be a force for good. The third of the pragmatists. They are pragmatic enough to believe that AI is here to stay. We cannot outrun it, so we need to embrace it, warts and all, and work out how we as humans can maintain and regulate what we see as its shortfalls. Finally, the fourth group are the objectionists, and they all cost think we should reject AI as a source for good. According to them, we should be both critical and suspicious of the contribution that AI can make here. Now, I leave you to decide which of these groups you think you may fall into, and we can discuss this in the Q&A afterwards. But it's also clearly important to prepare the workforce, which is both regulators and professionals, so they can understand the utility of AI and any barriers or benefits that may exist in its use. It's also necessary to prepare policy, which will in turn prepare the public so that they understand what any AI is being used for. Much of what regulators do is to protect the public, and the public need to understand that any use of AI will continue to protect them whilst also adding value to the regulatory framework that they are currently being protected by. Finally, we refer here to what is dubbed the politics of accommodation. In other spheres, this is often referred to as the establishment of fruitful cooperation with groups who tend to sit in rival blocks. We should celebrate the different perspectives that exist, not minimize those differences, but understand how they can inform our journey to our final destination, wherever that may be. On that note, I look forward to discussing some of these issues with you in the Q&A. But for now I would like to hand back to Anna for some concluding thoughts. Thank you. >> Anna: Thank you, Rob. Finally, from me, to return to our question, is regulation ready for AI? Our conclusion is that we are just at the start, and we must continue, we must invest in this. I hope you have demonstrated that this is feasible and has the potential to bring benefit. Ultimately, transformation comes from people, not from tools, and we have to engage everyone in the debate about how we move this forward in regulation. I hope that we have shown through this work that there is a potential they are to improve regulatory decision-making, and to shift that precious human resource from the low risk cases to the high-risk cases, where there will always be a need for proper hearings and weighing up the evidence for and against a case. But there's also room for embracing new technology to make regulation better, faster, without losing or compromising that essential human ingredients which is human judgment. So that's all for me and from Rob. We are really looking forward to hearing your questions and getting into a discussion with you about your thoughts and reflections on the question of whether regulatory bodies are ready for AI. Thank you so much. >> Michelle: Thank you so much, Anna and Rod, for that terrific

presentation. We would like to jump right into the Q&A and get going and talk about this really interesting topic you have presented to us. I would just like to start off -- Rob, I was very interested to hear you mention that the focus groups had concerns and identified the value of humility around the potential for AI. Could you expand on that a little bit further? >> Robert: Yes, she would. Thank you very much. I think the key thing is, with humility -- inevitably, with the groups and the people, the voices that we listen to, I think what sometimes happens is people overestimate what is possible. In this particular focus group, somebody just explained to us that, actually -- and he was a computer scientist. He said, "actually, don't get too excited yet." In some respects he felt we were running before we could walk and we needed to recognize there were actual limits. He said we need a bit of humility about what can happen here and what we can do, and what we need to be thinking about in the context of what is possible, what we should be doing in response to what's possible, and, whilst we should be ambitious, we shouldn't overestimate what we think or might think is going to happen in the future. I think he felt particularly that the conversation often being had about the negative side of it was often considering what might happen. I think he felt there was a sense we could over exaggerate it. He just said we need to have more humility about what is possible. Not just from that perspective, but the other side, as well. >> Michelle: Thank you, Rob. We do have some questions in the Q&A, so I'm going to jump into those. The first question, "in the U.S., many of APRN's practice is tied to physicians, via collaboration or supervision. How can AI account for those expectations and regulations?" >> Anna: Shall I start off? >> Michelle: That would be great. >> Anna: Is the question about accountability, then? The routes to accountability? I mean, I suppose it really -- I suppose it leads to another question, which is about how professions who work jointly together become jointly accountable. My sense is, as we move forward with these tools, and they will come, I think, to medical regulation as to nurse regulation, physical therapist regulation -- that they will come. Actually, we need to move at the same pace and in collaboration with each other. I don't know whether that quite answered the question that's being asked, but in terms of progress >> Michelle: Please, go ahead, and I will follow up. >> Anna: I was just going to finish and say, in terms of accountability, if a nurse practitioner is working to a physician and AI is being used, then clearly there has got to be a recognition of the contribution of those tools to practice as to regulation. >> Michelle: I think the question also had to do with differentiating the two providers when it came to discipline. I'm not certain if that is what the person was trying to get at, but that is certainly something that we hear boards of nursing talk about. In that case, how would AI differentiate the nurse practitioner, for example, from the physician? In the case of a disciplinary complaint. >> Anna: Of the complaint would have come to the board of nursing? So all the documentation on that case would be pertinent to the nurse practitioner. It may make reference to the medical practitioner, but I think in terms of captivating risk it would center on the role of the nurse practitioner in that instance. But I think it's a really good question. It's a question of whether we need more research or differentiation, because there is a level of risk that is carried by the team very often, not just trying to differentiate between two, between a nurse and a medical practitioner, but there are the people in the room, other practitioners who may also be contributing to that complaint or that error. >> Michelle: Thank you. Rob, I don't know if you have any comments there. >> Michelle: Yes, thanks. I think it feeds into each side, this idea that it's really important to work together. If we can recognize that there are these opportunities and think about, how would this apply to a nurse, and how might this apply to a physician? Then obviously the idea would be that you may well end up with some shared values and shared ideas, of course. But there may be some differentiations that can be built into the system is being used without any form of compromise. But recognizing that, if it is to do with accountability, I think you need to be sure that, obviously, the wreck person is attributed to her particular behavior. As Anna said, it is often quite complex and nuanced, and there will be different points in time,

which means the individuals might find themselves subject to different disciplinary processes at different points in time. But I think collaboration is probably the answer. That would be my view. >> Michelle: Thank you. Next question. "Were all levels of nursing licensure exam and in initial testing?" >> Anna: The short answer is yes. In the sampling we didn't distinguish between different levels of nurse practitioners. In the sample, we had the whole level represented. Yeah. >> Michelle: Okay, terrific. Next question, "this is a fascinating use of AI. Many uses for AI come with a warning to make sure the user is familiar with the appropriate limitations and uses of AI. As an educator, I'm curious what your ideal vision for what the education about AI would look like." Rob, do you want to lead us off on this one? >> Robert: I'm happy to start on that one. I think it's a great question, actually. As an educator, I would say that. But I think what is important is to celebrate these different voices, and that is what we are keen to do in our research. It's very easy when you have a research team for you to pull together and say, this is our common view, this is our common vision. And we wanted to do that, but we also wanted to make sure that all the voices were heard, and I explained the talk that it was a multidisciplinary team but we did have very different perspectives, as well. Actually, when we look back at some of our groups I mentioned, I could pick people in our research team who had those different views and those different positions. I think it is important to ensure those voices are heard and not ignored, and to the contribution is not minimized, either. I think that's really important. If we are going to move ahead in a really purposeful and positive way, I think it is important to make sure we do take on board those voices which are perhaps not as enthusiastic or as positive, or understandably concerned about where to go with this and what the potential is. And some of the potential pitfalls, as well. I think the key thing is to make sure that all accounts are suitably considered and diverse and ensure all of those voices are heard, and don't feel they are being minimized at all. >> Michelle: Anna, would you like to comment? >> Anna: I suppose I come at it from a slightly different perspective, in terms of vision. I think, for me, I would want to reassure those who are new to this that, just as other innovations in technology, there are the enthusiasts who start off trying to persuade others to come along with them. And they may come in their enthusiasm, as Rob alluded to, omit the downsides. I think understanding, for example, the very important differences between generated AI, which is the form of AI that most people are familiar with, and probably the form of AI that most people are fearful of, and the nongenerated AI, where the kind of computations are based much more or controlled by human input. That's the kind of tool that we have designed, and that, I think, has to be the way forward in regulation. I suspect in lots of parts of health care, as well, particularly in diagnostics, it has to be that nongenerative model. So I think, with the vision for transformation and improvement, and efficiency and accountability, it has to come with a caveat around, of course they're going to be things that we don't have full control of, but let's design the system so we have a maximum amount of human input, and those important human judgments I talked about. >> Michelle: Thank you. Next question, "I've heard that AI can make up or falsify citations. Can we program into the algorithm safeguards as to the regulations it is holding vendors to?" Anna, why don't we start with you this time? >> Anna: We were particularly excited by the suggestion that came from regulatory staff about trying to match cases to regulatory rules. As you know, there are dozens and dozens of regulatory rules, hundreds of pages of rules that actually are used in making the assessment of an individual case. If the question is, how can I put guardrails and safeguards in place to ensure the right rules are applied? I think the human judgment that comes at the end of the process is the ultimate test of that. The idea of the tool is that it is able to process vast amounts of information and pages and pages of rules at a much faster pace than human can. That's about expertise and experience built over months and years, and that again isn't a foolproof system. The safeguards that you need in place for computerized decision-making, you actually need the safeguards in place for human judgment. And that is why we have people doing

internal audits and checking on decisions that are made by humans. What is interesting about the conversations we have had with regulatory staff, as this work developed, is that they really welcomed that additional layer of quality assurance to their own decision-making. I think it is checks and balances. The really important point is being made. There are absolute guarantees, but you design the system recognizing that there are risks associated with the automated judgment or computerized judgments, and the human judgments, as well. And then hopefully you end up with a well-informed decision that has pulled through the correct rule. >> Michelle: Thank you. Rob, anything additional? >> Robert: Not particularly. I was just going to add, and that is right. Arguably, the current system, we are trying to suggest is we know we need to minimize the risk, but nobody is ever going to eradicate it. But it is the efficiency and the scale involved that does reduce the time for the caseworkers to focus on some of those really high-risk cases. I think that is one of the benefits we can certainly see here with the tool and how it could work in practice. >> Anna: If I could just add, the feedback we had from regulatory staff on the ground, that sense of wanting to focus the human resource on those more complex, high-risk cases, rather than on the cases that they were pretty much certain to go into no further action. I think that is very much the space we were in with regulatory staff. That is where they saw the real benefit to this tool. >> Michelle: Kind of along those lines of what you were just saying, Anna, the next question, "are there concerns that they will see your legal action? They can make a recommendation that seems to oppose the results generated by the tool." Rob, do you want to start with this one since it's kind of a legal question? >> Robert: I think that's a really interesting point. It's one of those things about accountability that I think is really of interest in this whole area. I suppose the first thing is someone needs to be able to bring attention to the fact that there was an alternate decision-maker. So I guess as part of the appeals process you want to understand whether someone could actually ask what the tool suggested, but then I think it would be up to the regulator to make it clear that the tool is a first step. Ultimately, if there was other evidence that was considered and other issues that were identified, it's quite possible to go against that. Now, part of that will obviously result in some challenge, but I think the hope is of course you could argue the more accurate the tool and the predictability, et cetera, it gets more positive, it makes it more difficult for someone to go against that machine judgment. But one assumes that there would be evidence that people could point to to say that this is the reason we did go against it. Nobody is suggesting that the tool should replace human judgment. The whole point is to be able to say, in these circumstances, with this evidence, we were able to choose a different path. But I do think it raises important questions about the way in which that is presented to people, and the way in which everyone understands how the tool is being used. Otherwise I could see that would be a routine concern. He would have people saying, what is the machine view? And I guess that might make regulator slightly defensive in practice, because then they think, the tool said this should be followed. But I think with precedent and a range of activities and explanations, that is the sort of thing that can be reviewed and explained. >> Michelle: Thank you. Anna? >> Anna: I think this is a really important and interesting area, and we talk about it, for example, in relation to judgments that are made by different boards. So there are different outcomes from humans who are looking at cases as well as the difference between a machine learning tool and a human resource. I think that comes down to the amount of experience that a particular board or panel has of dealing with that case. I think that's really important that we recognize variation within human judgment, and we are always trying to work to create more consistency. Whether that is within states were between states, or indeed between jurisdictions, I think this is a big debate across regulation globally. How do you make informed decisions when actually you have relatively little experience? You are not making decisions about cases all the time. It is not your primary role. But you are doing it to the best of your ability. I think recognizing variation between humans is something that we probably need to talk about more

often, but I think, as Rob says, at the end of the day, this is very much designed with the human being having the last word and being the final decision-maker. Would that be subject to legal action? Possibly. But I think we have to push forward with this and recognize the risks as we do so. I don't think, given the degree of reliability that we are hoping to achieve through more and greater testing that we would see big differences between human judgment and the judgments that are made by the tool itself. >> Michelle: Thank you. We have time for one more question. This is kind of a big one. Anna, you had mentioned the concern about bias. This has been a popular topic around AI. So, how do we avoid bias when the machines are learning from human beings who have bias? Can you address that in our last couple of minutes? >> Anna: Well, sure. This is so important to us. We have spent a lot of time on this. I think, generally, people who are designing these tools are incredibly focused on the whole -- what is colloquially called "garbage in, garbage out." The danger of using data that is biased is absolutely a clear and present one, and we need to be developing ever more sophisticated tools. We used three different types of gender debiasing techniques, gender-removing, gender-neutralizing, and gender-swapping, effectively switching the pronouns in order to test the impact of gender on the outcome. But, I mean, the question is a really good one, and we absolutely need to be focused on ensuring that the data we put into these tools, we need to be checking it and we need to be recognizing particularly bias on race, and removing those biases in order that we build a tool that doesn't have those human judgments that are going to effectively give us the wrong outcome. So I think we have to be really, really focused on this. >> Michelle: Well, thank you both so very much for a terrific presentation, and your thoughtful and insightful responses to the questions from our audience. Everyone, please join me in thanking Anna and Robert for their time today. Thank you again very much. We greatly appreciate it. >> Robert: Thank you. >> Anna: Thank you so much for having us. >> Michelle: Take care.