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

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

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,

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.