

[MUSIC PLAYING]

**KIMBERLY
NEVALA:**

Welcome to *Pondering AI*. My name is Kimberly Nevala. I'm a strategic advisor at SAS and your host this season as we contemplate the imperative for responsible AI. Each week we're joined by an expert to discuss a different aspect of the quest to ensure that artificial intelligence (AI) is deployed safely, fairly, and justly for all.

Today I'm so excited to be joined by Lama Nachman. Lama is an Intel Fellow , and director of Intel's Human & AI Systems Research Lab. She also led Intel's Responsible AI Program. And as you're going to hear, I think she's involved in some of the most fascinating and impactful work and applied artificial intelligence today. So welcome, Lama.

**LAMA
NACHMAN:**

Thank you. Glad to be here.

**KIMBERLY
NEVALA:**

So let's start. Tell us a little bit about what inspired and led you to your current work at Intel.

**LAMA
NACHMAN:**

So I would say I've been really constantly frustrated by how technology didn't really understand people. It's very active, it feels like it just has zero memory of anything that we interact with. So that really started a lot of the context that we are computing - where we're trying to bring much more context to our interaction with technology. But then another aspect of this that became quite obvious as we started to work through some of these capabilities is that you start to see a lot of the inequity that exists in areas like education or elderly support or people with disabilities, for example. So one of the areas that we've been looking at is: how can technology essentially help level the playing field? How can it be more about amplifying human potential and bringing the best of human potential? And humans and AI can actually collaborate rather than this notion of human AI competition where one of us is going to win and one of us is going to lose? So a lot of the focus of the research within my lab is really focused on that: which is how do we bring AI to help amplify human potential. It's really a multidisciplinary team that spans ethnography, design and AI. So we usually would start from some of those unmet needs to start to understand how do we dig deeper into these types of usages and look at the capabilities that are needed to enable that more broadly.

**KIMBERLY
NEVALA:**

So it sounds like the naming of the lab to be human AND AI systems was very, very deliberate. Can you dive a little bit more into that objective, that collaborative intent and some of the types of projects you folks are working on currently.

LAMA

NACHMAN:

Absolutely. So I talked about this notion of equity and improving human potential. So maybe I'll talk about a couple of areas that we're looking at. Education, and specifically, early childhood education as one of the areas that we've been very focused on for quite some time. And it started with this notion of we know, from the pedagogy research, that if you look at engagement and learning outcomes, they're very much correlated. So if you're able to actually improve engagement of students, you can actually get much better learning outcomes. Clearly in many cases, you'll do that with lower ratio of students to teachers. There are different ways of actually improving that. But we started to ask the question, how can technology really come in and help understand the learner engagement? And then use that to better empower the teacher to personalize the learning experience?

And if you think about it, with early childhood learning, you don't want to necessarily have kids behind screens. Because part of what we've learned from a lot of the ethnographic research that we've done is that parents are very, very worried about screen time. And at the same time they're worried about their kids being left behind because they're not really utilizing the power of what technology brings. So a lot of the work that we've done there marries the context where computing in human and AI, is to actually start to think about these learning environments where kids can learn the way they would expect them to learn. They're on the rug, they're playing with manipulators, tangibles, things like that. But then the AI system is able to actually watch over them, understand what's happening. I don't know if you've ever looked at a child communicating, it doesn't sound anything like how we communicate with Alexa. They point, they gesture, they go into silent modes because they're trying to understand what's going on.

So essentially, you need technology to be able to understand that level of interaction and essentially be situated in the same context that the kids are in. And then help facilitate that learning experience and then bring a lot of these analytics back to that teacher. Because if you think about it as they're trying to do something, you will understand where they're struggling, where they're confused, where they might be excited and I'm really more engaged. So, to bring that deeper insight to the teacher to help further personalize that experience.

Another example I would touch on is the work with assistive computing. So this started with some of the early work that we've done with Professor Stephen Hawking to try to really answer the question: can technology enable people with very limited ability to speak or move? Because if you think about it, if you can speak, you can today control your machine with speech. And if you can type, you could do anything. But if you have limited mobility and inability to speak, then you're left out.

And it's really not just only about communication. It's really about participation because today, if you have access to the digital world, you can actually pretty much do anything. So then we started to really ask the question: how can technology help capture whatever intent that the user might have? Anywhere from very limited muscle movement to brain computer interface, to be able to control their machine, communicate with the world, and do all the things that we today can do with our machines? And from there, really grow into how do we then bring in much more capability so that we reduce that type of interaction that's needed to convey that thought?

We also look at things like smart manufacturing and other areas where we can see that teaming between humans and eye for the eye to assess the human and for the human to be able to assess they AI to become more robust and can deploy into different environments.

KIMBERLY

NEVALA:

Yeah. It's interesting, I think that ability just to make ourselves heard and understood or even seen is so central to the human experience. So the ability to use these technologies and those abilities that we take for granted - those of us who don't have limitations in those areas. So the ability to use the technologies to improve that human experience. And I think there you're talking about in one level just creating parity, unlocking and augmenting our core capabilities. And in the other side, is amplifying the potential of humans, so whether it's the teacher or the student.

And it seems as you talk about those things that there would be no reason why we would say, hold up, wait, why would this be a bad thing? This couldn't possibly be a bad thing, or there's no negative implications of this, why wouldn't we want to do that? But I do wonder, what do we need to be aware of in the use of these technologies in this ways, do we run the risk of undermining unintentionally human connections or human empathy?

LAMA

NACHMAN:

Yeah. Actually, I would say there are tons of things to worry about, and it's something that, frankly keeps me up at night all the time. There are the obvious things around privacy and awareness of a lot of intimate things, especially when you talk about marginalized populations. In a sense, you're enabling somebody to communicate, but at the same time, you're enabling the machine to mimic the person in some way. And there are a lot of ethical considerations and worries that you would have. Including somebody's voice, like, what is authentic and what is not? That becomes really problematic.

Maybe a point that might not be as obvious, and this is something that we've been thinking a lot about in the design of the system. So I mentioned this notion of being able to control the machine with very limited input that somebody can control. But if you think about it, to really enable people to easily communicate, you have to elevate that level of communication. So one of the areas we've been working on is using technology is that response generation to help reduce that silence gap. So imagine somebody is having a conversation with someone who has that disability and the system can actually listen to the conversation and highlight some things that the user can quickly choose from, or, of course, they can type out the whole thing, but then if you choose something that's already there, you're enabling that response to happen quickly, which is a pain point that we hear often from people with earless and specifically in these conditions, because you have somebody speak and then somebody is trying to articulate the thought, and it takes a long time, it messes with this whole social interaction.

So we started to think about, can actually the system interact at that higher level? The problem is that, now, imagine this, every time somebody starts to choose something that's not really optimal or what they want to say, just for the sake of expediency, you're all of a sudden reaffirming or reinforcing the choice that was not optimal for the user. And over time, the system will start to actually pigeonhole you into something that's not really your intention.

So one of the areas in the design of that system that we're researching is to say, well, can I just give a hint to the system as I'm doing this to make it know that that's really not optimal. And then, when I'm not under the constraint of having to have a conversation, I can go back and really give feedback to the system on how it can better express my own thoughts. But at a time where I'm not really just pressed for that interaction expediency. So thinking through every aspect of that design to enable more agency, because you're giving up control. And a lot of times we give up control because it's like it's close enough, fine. But the pain point for people with disabilities is so much higher that as you start to take away that agency it becomes really much more problematic. So that's, for example, an ethical consideration that we struggle with quite a bit as we think about the design of that system. And this is why it's really key to have multidisciplinary teams that look at these things, because everybody comes at it from a different perspective and then you can bring a solution that's really much more responsible and cognizant of these needs.

**KIMBERLY
NEVALA:**

So there's a few interesting things in there. One, we've talked a lot about the need to create these engagement systems where you can give the system constant feedback. This is such a great example of such good intent where you may unintentionally pigeonhole a user. Again, with the intent of expediting something that may lead them down a path and they may not feel comfortable being able to correct that. But you're also talking about bringing not only that user experience and the user themselves but different functional perspectives into the discussion. So how do you design the process and set that up?

So for instance, in the circumstance you were talking about, what does the mechanism look like for that individual to say: here is how I want to express myself, here are the idioms I want to use. And what are the other disciplines you bring within your teams during this design process to make sure that all of these aspects are really being thought through critically and thoughtfully?

**LAMA
NACHMAN:**

Great question. Basically just to follow up that thought. So essentially, if you are able to flag that in the moment in a way that's very, very low weight: imagine you just gaze at this thing in the screen that tells it, just mark this for later. And then when you have time later, what the system can do is go and play back. Think about it as, I'm going to train my AI session, and yes this will take more time, but I have time now, and I can actually dedicate it to make the system work better for me. So then you can actually ask your system to say, can you bring up these interactions, and the system records enough of that interaction so that you can replay.

And once you do that, then you have a different type of interaction, where you are you really spelling out what it is that you would have said in this specific conversation. Because the interaction can be at this high level where you select one of a few options and it plays it like that. Or, you have a way to actually spell out letters and words and just rely on less automation and ability to actually express, but that will take much longer. But if you do that in your off time, what the system will do is take these as really strong evidence to retrain and fine-tune the system. So that next time around, what you will get is something that's more consistent with what you really wanted to express. So that's on the first question.

In terms of the disciplines, we have a lot of different expertise, in people with ethnographic research background, psychology background, interaction design, HCI background, AI. And typically, the way we think through this is we try to understand the users' needs. When we work in education, we actually have people who specifically have a background in education. And so we look at these different needs, and we start to hypothesize and take a lot of user input.

Like frankly, when we worked with Stephen, I always joke that Stephen Hawking was a designer on this project, he was a debugger of this project, he was someone who really brought a lot of interesting aspects for us to consider. But clearly, the voice of the users is very important. So when we work in education, we work with teachers, we work with students, we get input from parents. Because you have to make sure that it's really working for all the stakeholders in here. With people with disability, how does that change? Caregivers that are taking care of the individual, how do they change their workflow, what does that mean and what do they need to do to be able to support the user? So there are all of these aspects that you really need to bring to even start ideating about a system that would actually solve these issues.

Then clearly, there's a very iterative approach. So with Stephen as an example, we will get him in front of the system over and over again, as we add more features and will sometimes go back to scratch. I'll tell you, multiple times, I was like: is this a solvable problem? There's a reason why no one was able solve this problem for him at some point in time. So you start to go back to the drawing board and say, what are we missing? Where is this not really working?

**KIMBERLY
NEVALA:**

I'm a bit of a perfectionist, I'll say a recovering perfectionist and I've realized that over the years that a lot of things that I like to do are things that I'm not necessarily good at. And waiting to be good at them means I'm never going to do them. I wonder if we do have a bit of that tendency in systems design - even before artificial intelligence - of not wanting to bring people in when a system isn't perfect. But we know with AI they are never going to be perfect, they will continue to make mistakes. And so this idea that if we bring people in earlier and ask them openly and honestly to give us that critique and to be open to the suggestion that maybe this isn't a solution that works and we know that it's going to be wrong. Are they more prone to or more accepting and willing to give you that honest feedback, to accept mistakes and help you design around those mistakes? Because it's not if, it's when they happen, then if we are trying to program a perfect reality and system before we get them into the fray?

**LAMA
NACHMAN:**

Yeah. And that's actually a fantastic question. And if you think about it, there are a lot of different methods to be able to get at what you want even despite a technology not being ready. So part of it is really understanding the limitations of the technology and whether it works or not. But even earlier, before you have designed the specific AI capability, there's a lot of wizard of Oz type approaches that you could take to even understand if this system was perfect, would it support that needs? So we usually iterate at different levels of fidelity of prototyping within that journey.

The other thing that I would say, specifically about AI, if we wait for AI to be perfect, we will never deploy AI. That's just, I think, a statement that I can stand by even if I live to be 100. I think the issue is a lot of times is that we're not thinking about how do we use the uncertainty in AI to help design systems that take that into account. Because if you bring uncertainty - in that decision or in the inference - as a first class citizen in this design, you can actually do things very differently. You can bring in the user, you can bring in the here, and this is why I feel like we really need to be thinking about human AI systems as these mixed systems that can actually bring the ability of both of these very, very diverse actors to actually bring the best and the most out of this.

And the nice thing about that is that what humans tend to really be good at is very, very different what AI tends to be good at. And thinking about not they are replacing the human, but thinking about the AI really kind of bringing in where the human is maybe struggling, and with things that are so easy for humans to do, it can augment the AI capability. So that's where I see really a lot of the potential to even have many of those technologies deployable without this notion of, oh, eventually it will be perfect. Because I honestly don't believe it would.

KIMBERLY NEVALA: Yeah. I think this really brings home the point that distrust (or people not trusting the system) isn't necessarily because they're uncertain, it's because they're uncertain about their uncertainty, if that makes sense.

LAMA NACHMAN: That's absolutely true.

KIMBERLY NEVALA: So you've spoken a lot about this need to really change the dialogue and our mindset from AI as a competitor to AI as a collaborator, and I think you're starting to reference that. Can you talk a little bit more about why does that mindshift matter, and how does that change in understanding impact our approach to developing these technologies at all?

LAMA NACHMAN: Yeah. So I do talk a lot about this. And I think it's an area as you can see that I'm very passionate about. Maybe it's easier probably to talk about an example.

KIMBERLY NEVALA: Perfect.

LAMA NACHMAN: When we think about, let's take manufacturing, it's clearly where Intel-- it's an area that we do a lot of work in. But if you think about the tendency of trying to deploy AI, very typically, we'll go for, what is the most obvious thing that we can replace the human with an AI system with? So what subtask is repetitive, predictable, has tons of data we could easily train, and then have just the human not have to do this?

However, those are limited, where you can actually bring the AI capability. Because when you think about bringing an AI capability, what that really means is it needs to be something that will be resilient, does not have to change every five minutes, and then you have to retrain everything from scratch. So there are a lot of other qualities that need to be true for that to actually make sense.

Now, if you step back and say, look at a complete workflow from beginning to end. There are outcomes that I'm interested then in manufacturing and there are humans who are performing tasks within this total workflow. And it makes sense to try to actually help support these people as they make mistakes, as they're trying to improve their productivity, et cetera. AI doesn't have to understand every single aspect. It has to just augment the gaps that people have.

So if you look at it from that perspective, you can start to see how the AI can alert people to and bring information when it's needed, learn from experts as they perform task,s and then support people who are maybe not experts to do these tasks. You could start to see a lot of different places where AI can come in even at that level of fidelity and really bring a lot of value to that complete workflow. But it's depending on the fact that in the cases where things don't look like what I've seen before in training and where I'm uncertain - to your earlier point - there's a human that I can rely on to help really bring much more clarity and help support me learn what I don't understand. Or where I'm uncertain, I can ask.

So there's a lot of work that's happening today in trying to really do better at uncertainty estimation and where it doesn't know. Because when it knows that it doesn't know and you're actually in a mixed decision system with a human, there's so much more that you could do. So that the total performance of that whole system is so much better. Despite the fact that this is not perfect, the human is not perfect, the combination of those two things is so much better. And that's really where I mentioned manufacturing as an example, but we see this in pretty much almost every field that we can think of.

**KIMBERLY
NEVALA:**

So how do you rationalize, or what's the process or the thinking people need to engage in to figure out: this is a problem that yes that has some uncertainty. It's got some fuzzy edges and we don't know what we don't know necessarily, but this is the level of uncertainty that we can tolerate. Here's the level of safety, the guardrails, et cetera, that are going to be OK so that we can deploy this in a way that is safe, it's resilient- even though it's not perfect. So how do we define how perfect it needs to be?

**LAMA
NACHMAN:**

Yeah. So I mean I would say it really just depends on the specific system that you're deploying and there are certain KPIs that you would look at and at the end of the day say, does this actually provide value? Are those outcomes that I'm measuring, whether it's errors in performing tasks or it's productivity, end to end time and performing the tasks? There are certain KPIs that you can measure and say if the AI performance isn't providing at the end an ability to improve that from for example where a human is, then what's the point?

But I think the more important point to think about - and the way we think about it - is that these systems continue to learn. So the idea is that you deploy them, and over time, they are learning from that deployment, because they're watching the human. The human is actually training the system. So one of the areas that we look a lot at is this notion of dynamic type deployments. Where essentially, think about it as an expert that's trying to teach the AI as they're performing their tasks.

And you could look at this and say, well, what's the point? That I just move a lot of the pain from people who develop to users to actually train these AI systems? But if you think about a lot of the techniques that we look at from ingesting a lot of information from manuals, from watching what's happening and using things like weekly supervision to be able to actually continue to ramp up the capability of the AI system, you're not going to necessarily give a decision making capability but it's not at that level. But over time, what's happening is that the AI system is actually learning in that deployment. And as it learns, it can take on more, and as it learns, it can help support and all of us, for example, performing that part of the task.

So if we think about the AI system as a static system, then we can say, well, at what level does it need to be? But if you think about it as something that continues to learn and improve over time, then really maybe the question is, when does it take over rather than when does it make sense to deploy it?

**KIMBERLY
NEVALA:**

So you talk about identifying what the objective outcome is and looking at that. Is it as - or even maybe more - important that folks are thinking and defining up front what it is that we don't want the machine or the AI to learn? Or things that we don't want to see. So that we're on the lookout for those as well? Because the upside and the downside of a learning system is it's going to learn from what it sees and that may or may not be what you thought or intended.

LAMA
NACHMAN: Absolutely. 100%. And if you think, there are a lot of different really connotations to this. So one aspect of this is that the minute you make the system dynamic and it learns, it could end up really degrading in performance. And the model performance will shift over time in a place and not necessarily upward, but downward, and that's a problem. So that's a problem just from a resilience perspective.

There are also things about that that have a lot of ethical implications. So we've seen, for example, in chat bots that learn, how do you make the chat bot not become racist? That's a problem. So there are a lot of interesting checks and balances that you really need to put in place. And typically when you are deploying systems out there at scale, you're actually essentially thinking about how can they system monitor itself or how can people monitor the system and make sure that it's not going to just run amok and we're not going to know about it.

And again, I would add the risk of sounding like a broken record, the beauty of a human-AI system is that these things can be corrected for and can be caught, and you could do something in real time about it. But that's absolutely important. You have to actually be thinking ahead of time what are all the things that could go wrong in such a deployment? How do I understand that the learning isn't moving in a direction that I don't want it to?

KIMBERLY
NEVALA: Yeah. So really acknowledging and accepting the fallibility both of the system itself and the humans that interact with that because these fallibilities can be reinforcing as well. So, as you look at the work that you folks are doing across all of these very interesting areas from accessibility to human augmentation and more mechanical applications - for instance, in manufacturing - are there areas that people want to apply these technologies where you think maybe we just shouldn't? Because they're just not ready yet or because there are ethical or other implications maybe to, again, human agency and human connectedness that require more thought?

LAMA
NACHMAN: Absolutely. It's interesting. So one area that we didn't touch on is we've been doing quite a bit of work in trying to understand collaboration and human to human communication. And it's funny enough that this started pre-pandemic, because we were looking at remote work but not in a pandemic context. So it was like just because of many how people are in the different geos, we have team meetings that span five different geos sometimes.

And so part of that, you understand that there is the people in the room and there the people who are calling remotely. And there's an inequity in this whole balance. So, how do you start to bring more of the in-person experience to that remote collaborator? And then, of course, the pandemic hit. And then there is equity, everybody has a worse experience across the board. But then you start to think about, well, there is this human to human thing that happens that we pick up as humans. Or, I would say, maybe a good portion of humans are very contextually aware. And they use a lot of the signals in very different ways to show empathy, for creating much better connection and interaction, and all of that, and then you take that away.

And then you start to think about your Zoom scenario where you have these tiny, tiny little boxes all over the screen and you're trying to actually connect, how do you do that? So part of what we've been looking at is, how do you use AI to help bring back some of that context into these type of conversations? But then there are a lot of other things to worry about. What does that mean? Is your computer system spying on you to try to actually use some of these things that you don't necessarily think that you're expressing, but you are expressing. Because if you were in front of a human, you would probably disguise more of these things than if you were not, as an example, and it's very easy to forget what's going on when the person is not right in front of you.

So a lot of these bits and pieces. So this is in terms of just the design and where does that apply and all of that.

There's another aspect of it, which is I think maybe more obvious, but I'll just mention it, is a lot of times we're building capabilities, because we have some specific usage or a really good usage in mind. But then you can start to think about how these things could be taken and applied in areas that it shouldn't be applied to.

And then the question becomes, well, what does that mean? Do you not build the capability, like there are a million good usages of one thing and some really hairy and very, very worrisome usages. And what does that mean and what type of checks and balances can you put in to make sure that technology isn't being misused when you're talking at the capability level? And it's really funny.

So one thing that we all are aware of defects and a lot of the issues would be fakes. And sometimes you think about it and say, is this really a good use of why we would make something sound like somebody? But then we just talked about accessibility. I can't tell you how many times I've had conversations with people who have loved ones with ALS who, or the person himself or herself who has ALS and have lost that ability to speak really thinking that they would give anything for being able to say things in their own voice and have other people listen to that voice.

And it's like a technology can be totally misused and you can think of more really inappropriate uses of that technology, inappropriate, in some sense, than the one that I would imagine is really very appropriate. It's such a struggle, it's so many ethical dilemmas in all of this work. And it's really very like the only way that we have taken both within my team and then across Intel is to really say we need to bring a very multidisciplinary approach. We have to bring very different voices to the table to help us think through all of these risks and harms, and what can be done to mitigate this.

Some of these might be in technology, some of these might be in what we do with licensing of our technology with policy, there are a lot of different things, there isn't one solution to this problem, but you have to be mindful of these risks and then doing something about these risks.

**KIMBERLY
NEVALA:**

Yeah. It's interesting, because you're talking about for instance using natural language generation to enable accessibility and to give people their voice back and then the downside of deepfakes and things like that. And those are different use cases. But you also get these issues within the context of a single use case. You talked at the top of the discussion about using artificial intelligence to understand when children are being engaged and not engaged. You can also then see it being used to identify kids who might just be distracted or distractible and potentially that being used to somewhat punish what is just natural child behavior. So these come up in a micro and macro sense.

**LAMA
NACHMAN:**

Absolutely. Spot on. In the same use case, what is this going to be really enabling? Are you actually enabling the teacher to say, well, these are things clearly that this kid has gotten, but then wait, what's going on here and how can I dig deeper to understand then help support that gap. It's funny. One of the use cases that we see always fly around is like, can we help understand who's raising their hands? And, well, then what does that become? Why are we doing this and what are we using this for?

And then the funny thing is, also, there is the other side of it which is: kids will learn what type of behavior they should do to circumvent these type of systems. So then they become, by definition, totally useless. So it's all of these things we really need to be much more thoughtful about: what is this enabling, what are the risks associated with it, and how can we get the best out of this while mitigating the harms?

KIMBERLY NEVALA: So as you look forward then, what are the things that you're most concerned about as these technologies evolve and as we just rapidly deploy and adopt them?

LAMA NACHMAN: So many things that I'm really concerned about. Maybe we didn't touch much on privacy, but privacy is an area that I worry about heavily. Because in everything that we're doing, we're collecting so much data from so many different modalities, much of it is like very, very private data, if you will. I talked about agency and it's something that really worries me a lot, especially with people with disability, with kids, that worries me quite a bit.

We talked about misuse, and that's really almost a problem that has no end, because there are so many creative ways that you could misuse these technologies that I could be thinking about it all night long and somebody will come up with yet another one.

But one of the things that I also worry about, I keep thinking about how can technology and AI bring equity. And sometimes really what's maybe more obvious is that it can amplify inequity. Sometimes just by virtue of access, you essentially made that divide even larger. Because now you need to ask yourself the question of, if I make this really compelling and I'm not democratizing the access, then I'm actually doing the opposite of what I just said out to do, and that's another thing that really keeps me quite a bit worried. But also obviously bias in AI is another area that I'm very worried about, especially when you don't have the right data and the right attributes capture, which then means it only works for a subset of the population because those are the ones that we have data from, and if you want to go and get data from populations that might not be represented, then you're actually creating even more of a problem there.

So it's just like I mean there is nothing obvious in any of these things as well, and maybe that's why we're all interested and excited about doing that work, but it's also a lot of responsibility to actually do that.

KIMBERLY NEVALA: So as not to anyone to scare people out of the room... Because, obviously, I do think that sometimes the folks who are most targeted and focused and have the biggest concerns are also the folks who are the most excited about the technology and see the potential. So what most excites you? What would you like to see happen in the next three to five years?

LAMA NACHMAN: So what I love to see is really more and more deployments that cap that does amplify human potential. We're enabling people to learn quickly to improve their capabilities, especially in, like we talk a lot about shifting what the workforce looks like, and we're trying to upskill and reskill people. And I think many times when we are faced with automation, there are people that are always left behind. It's really interesting to say, can you turn that?

Unlike any automation, AI has an ability to help improve that learning capability of people. And can we actually then utilize AI itself to help in that up-skilling and that learning and bringing that potential back so that you're mitigating also some of these risks of automation and efficiency and things like that. So that, to me, and I'm a very optimistic person, although I talk more about the problems in there, but really optimism is what keeps me in this field. And I honestly see all of that potential for improving equity, accessibility, enable human potential. But I think that doesn't happen randomly. It happens by being very, very focused and intentional about it, and thinking about these things very early on in that design cycle.

KIMBERLY NEVALA: Thank you, Lama. I found this absolutely fascinating. And I think we are probably all more sensitive to context, maybe a little more comfortable with being uncertain and ready to embrace imperfection as we deploy these systems moving forward. So thank you again for joining us today with this great discussion.

LAMA NACHMAN: Thank you. I love the discussion.

KIMBERLY NEVALA: Awesome. So in our eighth and final episode this season, we are going to speak with Dr. Ansgar Keone, a senior research Fellow and Ernst Young's Global AI Ethics and Regulatory leader. He specializes in Social Media Data Ethics and AI Regulation, and he's going to help us all understand more about what we need to be more informed and intentional in our engagement with AI as individuals, citizens, and organizations. So make sure you don't miss it by subscribing now to *Pondering AI* in your favorite podcatcher.

[MUSIC PLAYING]