

The Social Science of AI: Intel's Elizabeth Anne Watkins

SAM RANSBOTHAM: We know that artificial intelligence tools are augmenting human performance, but how do people really feel about that? On today's episode, find out how one company develops AI tools with end users in mind.

ELIZABETH ANNE WATKINS: I'm Elizabeth Anne Watkins from Intel, and you're listening to *Me, Myself, and AI*.

SAM RANSBOTHAM: Welcome to *Me, Myself, and AI*, a podcast on artificial intelligence in business. Each episode, we introduce you to someone innovating with AI. I'm Sam Ransbotham, professor of analytics at Boston College. I'm also the AI and business strategy guest editor at *MIT Sloan Management Review*.

SHERVIN KHODABANDEH: And I'm Shervin Khodabandeh, senior partner with BCG and one of the leaders of our AI business. Together, *MIT SMR* and BCG have been researching and publishing on AI since 2017, interviewing hundreds of practitioners and surveying thousands of companies on what it takes to build and to deploy and scale AI capabilities and really transform the way organizations operate.

SAM RANSBOTHAM: Welcome. Today we've got a great guest: Elizabeth Anne Watkins, a research scientist at Intel. Elizabeth, thanks for joining us today. Let's get started.

ELIZABETH ANNE WATKINS: Thank you so much for having me today.

SAM RANSBOTHAM: As the world's largest semiconductor manufacturer, Intel is probably a company that most people already know, but can you tell us about Intel Labs in general and maybe your role specifically?

ELIZABETH ANNE WATKINS: Just like you said, we're not always top of mind in the big discussions around AI and the AI industry and the AI field right now, but there is so much fascinating work happening inside Intel, and we have such a unique perspective on the field and a unique way of entering that field that I'm really excited to bring some of that to light today in our conversation.

I just joined Intel in August of last year, and already it's been a really incredible experience meeting so many different teams. I joined Intel as a research scientist in the social science of artificial intelligence and work under Lama Nachman in Intel Labs. The group is called Intelligent Systems Research.

SHERVIN KHODABANDEH: Elizabeth, you mentioned Intel Labs is doing some unique things with AI. Do you mind sharing with us some of the things you're working on?

ELIZABETH ANNE WATKINS: A project that I'm particularly excited about is called

MARIE, an acronym which stands for Multimodal Activity Recognition in Industrial Environments. So basically, I'm going to start with a metaphor: Imagine that your computer could watch you put together, say, a piece of furniture that you ordered on the internet. When you got to a tough part of the manual, or you're holding a screwdriver, you're holding a piece of plywood, and you can't get back to the manual, imagine that your computer could see what you were doing, knew what the manual was going to tell you to [do], and then help you to connect those two. Imagine that your computer could actually tell you, "Hey, I think you are about to screw shelf A into bracket B," or something of that nature. And I know that every time I have received that flat pack that they say has an armchair in it, it's a really tough time for me to get from A to B.

And it's processes like these that our people are doing inside of our facilities, where they're actually building and manufacturing the semiconductor chips. So the folks who work inside of our factories, our technicians, are doing very involved and very delicate work, handling parts and tools for all kinds of manual operations happening on the factory floor. And so all of the work that they're doing is just as complicated — sometimes even more complicated — than getting that flat pack into an armchair. There's a lot of tools involved. There are all kinds of different processes, different pieces of equipment, different sizes of equipment. And so we are building computer systems a little bit like the one that I described that said, "Hey, did you mean to put screw A into bracket B?" We're building systems to help the people

inside of our factories do this kind of really careful and really complicated work.

SAM RANSBOTHAM: Actually, that's a fun analogy. I mean, I think we all find flat packs challenging, perhaps, though I have to admit I kind of enjoy them. But I'm sure it's much more complicated within Intel. And what I liked about that example is, I feel like so often we're talking about automation: Can we get machines to learn how to do something that humans do? That's machine learning [ML] at its core. And then we talk about augmentation and, well, how can machines help humans make a decision? This strikes me as a little bit different. This is going the next step.

So in this case, we're not trying to get the machine to assemble the flat pack. We're not trying to get the machine to assemble the semiconductor. It's still the humans doing it. And so it's the humans who are needing to learn here. That seems like pushing that a little bit further than we've [about] talked before.

ELIZABETH ANNE WATKINS: I'm so glad you picked up on that, Sam. That's exactly what we talk about inside Intel Labs, is that the human is the expert. The human is the one who knows how to interact with these systems, and they know how to screw the bits together. And they know all of the physical and dynamic intricacies of what it takes to put these tiny and very delicate components together.

We don't want to train a computer to do that. A computer cannot do that. The way that this both assembly and cleaning process — because of course all of these

processes take place within these super-clean rooms, where everyone has to wear hazmat suits and gloves — we want to support the humans who are doing it. There is no computer who could do this. It takes human judgment; it takes human expertise to do these processes in a way that is comprehensive and truly dynamic and can respond to new pressures.

If a larger piece of a machine bends in a particular way, or if a screw falls into a hole in a particular way — it's going to be many, many years before computers are able to assess and diagnose and fix these kinds of constant dynamic problems. But you know who's really good at doing that is humans. And so we're trying to center humans and center human expertise. And that points to this dual reason that I'm really excited about MARIE, is that there's both a tech side and there's a human side. And on the tech side, I think MARIE is really exciting because it's multimodal. It combines a lot of different kinds of AI systems through ambient sensing. It combines computer vision, which uses activity recognition, with audio sensing, combined with natural language processing [NLP] in a way that can help build an ambient environment around the technician, ultimately to help them what they do.

But before the system can help them [with] what they do, the system has to actually learn, and it's the human experts that are teaching the system how to learn and teaching the system what it needs to know. And so this is, as far as I know, a brand-new way to deploy AI systems into new domains and apply them to new problems. ...We are confronting and tackling one of the big

challenges of AI development, which is data collection.

You need a lot of data to get into a new domain. You need a lot of data that's labeled the appropriate way, labeled according to the kinds of problems that you're trying to solve. And so we've kind of flipped the script: Instead of trying to get all the data we can before a system is deployed, we go ahead and deploy a system that then works in partnership with the experts on the ground, and they teach the system — through speaking aloud and through dynamic data labeling — they teach the system what it needs to know. And so hopefully the data that is going to be used to help people to do these tasks down the line is going to be produced by the very first batches of people who are using the tool.

SAM RANSBOTHAM: That seems like a great analogy because as you're saying that, it reminded me of what you would do if it was a new human partner joining the project. You'd go through those same steps.

Shervin and I have some research where we find that 60% of the people out there are thinking of AI as a coworker, and that's exactly the sort of relationship that you're describing. But it made me think of a little danger here. Let's say that when you put that system out there, it doesn't know much, and you've got the human training them. All right; I'm a human. I feel like I'm a little bit annoyed. I'm like, "Why am I having to work so hard to train this coworker?" How do you keep that dynamic? You mentioned there being a technical side of this project, but the social side as well. That

strikes me as maybe harder on the social side than on the technical side.

ELIZABETH ANNE WATKINS: That's precisely what I want to talk about next because that's precisely what I get to do. I get to talk to the people about what they actually think about these systems — and you are right on the money. Sometimes they do perceive it to be annoying, where they're asking, "Wait — so I'm putting together this really complex instrument, and the system asks me what I'm doing, and it's wrong? So how am I supposed to spend time not only doing my job but also teaching this machine how to do my job?"

And so I am really gratified that I get to be the social scientist who is there right on the ground. We're still deep in development, and it's common for companies to do as much testing as they can, of course, before something is deployed and then, when the actual tenets of deployment come — things such as user experience and user interface designs — those typically come closer to the end of a development process. But here, the social science and the understanding of what people need and what people might find annoying and what they need to trust a machine that is a coworker with them is right in the DNA of the development. And so I get to consistently speak with the technicians that we are building the system for and ask them questions just like the one that you described, asking them, "Do you think this is annoying? How might this be more helpful? What are some other affordances that we can build into the system to make it more helpful for you?"

And also asking deeper questions, not just about their dyadic relationship with the system but also about their larger sociotechnical context of work. Asking about, like, "What does your schedule look like?" and, "What have your bosses said about the system?" and, "Is there a way that we could try to facilitate more productive relationships, not just with your bosses but also with your other coworkers, through this system?" And so our goal is to build a very deep understanding, not just of one task but an entire space of work, and ask how our system can help in the best way it can to amplify human potential by being a coworker for humans within this space, but also [asking], "How can it be built with an understanding of their work context in order to facilitate a long-term relationship?"

Because one of these things that's demanded by the fact that we're leaning on people to provide data and leaning on people to help us actually do the data labeling is that we need them to be engaged, and we need them to like interacting with the system, and we need them to find it trustworthy and reliable and useful. And understanding how all those goals can be achieved requires a deep understanding of social context and social habits and work habits and task flows. And so that's precisely what we're working on now.

SHERVIN KHODABANDEH: I think it's really interesting how you seem to be thinking beyond just this project. There's the element of "OK, how do I get this particular project in place," but underlying a lot of things you're talking about are things like,

“How do we figure out the best way to introduce, essentially, a new system? How do we figure out the best way for how to transform and get to our overall vision?” And that transcends just a single project. How are you trying to pick up on those lessons?

ELIZABETH ANNE WATKINS: That is a really great question. I am really proud of our labs and some of the big bets that we’re making on human-AI collaboration and how much the input and contributions and insights drawn from social science are being picked up, and have been picked up historically at Intel, throughout their research projects. As we’re doing research projects and building products across verticals like education and manufacturing and accessibility, there’s deep investment across the teams that facilitates and encourages conversations with folks like me — with the social scientists.

We often get these really fantastic teams in a room where we have engineers and data scientists and policy makers and social scientists, including not just me but a wonderful team headed up by Dawn Nafus of anthropologists, as well as another team of social scientists that has psychologists. And we all work together really closely to ensure that the products we build are not just designs to answer one problem but that they’re properly engineered around what the best solution is, that we take advantage of all the different kinds of multimodal affordances that our tools can provide, and that we’re deeply understanding of the social and organizational context that all of our products are going to be deployed into.

SAM RANSBOTHAM: I think so often there’s a transition that seems to be happening between when we were first talking about these models and tools and using data: Things were very heavily data- and tool-oriented and very heavily science-oriented. And that made a lot of sense because we didn’t know how to do some of these things. And now, as a society, we’re learning how to do many of these tools and model buildings. But a lot of the team you mentioned there — anthropology, psychology — these are not traditional things that we might have thought of being integral to producing an AI application, but those are the ones you happened to mention. Why did you pick those? I guess that’s some of your background in social science, right?

ELIZABETH ANNE WATKINS: Yeah, throughout my graduate studies and postdoc work, I was always really invested in the users and in people. And I was always really fascinated by just how weird people can be and how creative and how innovative, and how folks often use technologies in ways that their developers did not anticipate and did not foresee.

And while there are some dangerous elements to that, as we’ve seen in various misuse and dual-use applications, there are also a lot of really fantastic and wonderful ways that people have figured out for technology to be more suited to them in their context or to work a little bit more smoothly for them. And so seeing how Intel was different because of their history of being a semiconductor manufacturer and being deeply invested in hardware and having an ecosystem approach to the way

that AI tools are deployed really showed me that they were a company that cared, just like I did, about the people who were going to be using these tools and the social structures in which these people were embedded, and building tools that could match not just the people who bought or used the tools but also the lives and the communities into which these technologies were going to be interjected.

SHERVIN KHODABANDEH: Elizabeth, clearly a lot of your background has been coming through as we've been talking, but could we step back for a second and ask you about how you ended up in this role at Intel Labs?

ELIZABETH ANNE WATKINS: I guess my road to this role started way back, probably at the beginning of my graduate career. After I graduated MIT and started my Ph.D. at Columbia, there was an opening for someone to contribute to a project on security and privacy behaviors among journalists. And my adviser said, "Hey, I got you an interview for this project." And I said, "I don't know anything about security and privacy." And she said, "Just go to the interview. I got you the interview. Just go." And so the night before the interview, I thought, "Oh, I'd better have something to talk about. PGP — Pretty Good Privacy — everyone's talking about PGP. I better download PGP so I can act like I know what PGP is." And I tried to figure it out, and it was hard. I couldn't figure out ... I went to the website and I was like, "The website doesn't really explain what's happening here. Where's the key? I download the key, but the key's also kept in a database, but I write the key on my emails? Oh, I'm so ... I

can't do this. I'm so bad at this. This is not for me."

And I went into the interview the next day with the wonderful project lead, Susan McGregor, and I said, "You know what? This job's not for me. I tried to do PGP; it's weirdly hard. There's something about it that I just ... I can't quite grok the language." And she said, "I think that makes you perfect for this job because we are trying to figure out why are security and privacy are so hard for people at work, especially journalists who are high-value targets for attack from a lot of different actors, and we are trying to figure out how we can make security and privacy protocols better for them, so it sounds like you're frustrated with how security is designed." And I said, "Yeah, that was frustrating." And she said, "OK, let's do some work."

So we ended up working together for several years studying journalists. We had one particularly harrowing project studying the journalists who published the Panama Papers. And for that, they had terabytes of data. There were journalists working all over the world; none of them were colocated, and they didn't have a single breach. Never once did they have a breach and in all of that data. And so we saw that as a success story. And we said, "OK. We're always hearing the bad stories about breaches and attacks; we never hear the success stories." And so we got to interview the journalists who contributed to the Panama Papers and talked to them about their organizational culture and how making protocols uniform across all journalists helped to instill a sense of teamwork and a sense that they were protecting each other.

And seeing the power of culture and shared mission on behaviors around something as difficult as PGP and security protocols ... it was really inspirational.

And around that time, I started paying attention to facial recognition, and from there it was a short bridge over to looking at AI and asking similar questions around “Why is this hard to understand? How can we communicate it to people in a more strategic way that is more understanding of the work that they’re doing and the work that they need to do? And how can we bring some transparency to these systems in a way that is meaningful and understandable for real people?”

SAM RANSBOTHAM: One of the things interesting about that is, security’s always like that. If it’s a little bit hard to do, then we all tend to not do it, because it’s tomorrow’s problem versus today’s problem. ... That’s something we as humans are terrible about. And you know, I think what we’ve failed [at] in many ways here is not building that into the infrastructure so it’s a default. And there’s a lot of analogies for artificial intelligence: We’re building things, and if we don’t make the defaults easy to use or easy to do the right thing, then people will do the wrong thing — or people will do the easy thing or the short-term thing. And you mentioned facial recognition: What are the kinds of things that we should be building into our infrastructure so we don’t propagate AI-based mistakes?

ELIZABETH ANNE WATKINS: That’s a great question. Something that drives my research, both throughout my graduate career and here at Intel, is a recognition

that people are the experts in their own lives and that we really need to engage with the people and with the users throughout the development pipeline in order to build not just systems but solutions, and solutions for the real problems that are happening on the ground. And as close as a company can get to including the expertise of social science and what social scientists can bring about rigorous and robust tools to study how people live and the kinds of languages they use and the kinds of values they have and what’s truly important to them and what they want to protect and what they want to keep safe, as well as building in different kinds of options and different kinds of pathways into the same technology that might be used differently by people of different levels of accessibility, are all things that I would love to see the entire field take up.

SAM RANSBOTHAM: That sounds good. I mean ... much like buying low and selling high sounds great. How do companies actually do this? What steps do they need to take to make progress on these?

ELIZABETH ANNE WATKINS: That’s a great question. The way that we’re doing it at Intel is that we are facilitating conversations between our technical teams — people who are building really amazing systems around robotics and accessibility tools and educational tools — and embedding social scientists across all these teams so that we can ask some questions around “Hey, what are the presumptions of this system? Have you talked to teachers? Have you talked to the folks who are going to be using this robot on the ground? What are your

conversations like with the people who are going to be served by these solutions?”

And I’m really lucky that I get to be in a place where we’re embedding this expertise way back into the problem-formulation stage and into the project-formulation stage, all throughout the development pipeline. I’ve also been really gratified that Intel established the ethical impact assessment process along with the Responsible AI [Advisory] Council. And so this is a process for robustly and rigorously building into the development pipeline that teams all across Intel and across all business units [use]. This is a way to inject the expertise, not just of social scientists but everyone who sits on the Responsible AI Council, including engineers and policy makers and folks from legal and folks from standards, and facilitate conversations with this multidisciplinary Responsible AI Council with development teams through their submission of the ethical impact assessment.

And the ethical impact assessment is a way for us to put values into practice and ensure that values around human oversights and human rights and privacy and safety and security and diversity and inclusion — these are ways for us to make sure that those values are built into the tools that Intel BUs [business units] are putting together across the organization.

SHERVIN KHODABANDEH: OK, we have a segment where we ask our guests a series of five rapid-fire questions. Just answer the first thing that comes to your mind. First: What is your proudest AI moment?

ELIZABETH ANNE WATKINS: Oh, goodness. That’s a great question. I am really proud of being able to represent the voices of the technicians who we have on the floor to these teams of engineers and data scientists who are building these systems. And because of the process that Intel has built, where I get to sit with the engineering teams who are building the computer vision and the action recognition and the NLP and choosing the phrases in the semantic frames around tasks and about how tasks are built and how they can be recognized, being able to do the interviews that I do with technicians to ask them about their concerns when this technology is introduced, what we can do to make sure that their concerns are addressed, asking what they need to know around transparency, what they need in order to trust the system, all in the service of enhancing the work that people are doing.

The fact that I get to grab that information and deliver it consistently to the engineers and to see how quickly they respond. Like, “Oh, well, we can build this and we can build that. And what if we built this into the UI? And what if we built this into the dialog?” It’s such an incredibly compelling process, especially after coming from academia, where I would write a paper and then it would take a year to get published. And then ...

SAM RANSBOTHAM: Oh, a year sounds fast.

ELIZABETH ANNE WATKINS: Right. If I were lucky, it would take a year. I think my longest record was maybe three and a half years it took to get published. I don’t know. What’s your record?

SAM RANSBOTHAM: Uh, 10. So ...

ELIZABETH ANNE WATKINS: Ten!

SAM RANSBOTHAM: Yeah. Let's don't talk about it though, because I'll get all sad.

ELIZABETH ANNE WATKINS: Oh, no.

SAM RANSBOTHAM: So you mentioned concerns that people have. What worries you about artificial intelligence?

ELIZABETH ANNE WATKINS: Well, one of their biggest concerns, when I started talking to the technicians, was that they were going to get replaced.

I said, "OK, let's have a conversation. What do you think this is for?" And they said, "I've seen the news. I know you're trying to build a robot just to do exactly what I do, so I figure I'm training my replacement." And I thought, "Oh, no! That's terrible." And I got to have a lot of conversations with our technicians where I got to say, "That's not what we're doing. We are not trying to replace you. You are the expert. We need you. In fact, we need you to teach the system so that the system can go and help other people like you. So by you training the system, you're actually helping a lot of other folks in Intel fabs [semiconductor manufacturing plants] across the planet, and you are ensuring that they get help in the way that you would like to get help."

SAM RANSBOTHAM: Some research that Shervin and I were involved in a couple years ago focused entirely on this organizational learning aspect that you were

alluding to — that this is a way for everyone to learn more quickly and to spread knowledge. What's your favorite activity that does not involve technology? What's not AI going on in your world?

ELIZABETH ANNE WATKINS: Well, there's so much happening recently, it seems tough sometimes to figure out what's AI and what's not.

[In] my personal life, I do a lot of baking. It's been a while since I made a loaf of bread. I do a lot of cooking. I am very lucky to live with my husband in the city of New York, and so we do a lot of exploring — in fact, just getting outside. That's probably the biggest non-AI activity, is walking outside on our own two feet and looking around with our own eyes. And that always feels very refreshing.

SHERVIN KHODABANDEH: What's the first career you wanted when you were young? What did you want to be when you grew up?

ELIZABETH ANNE WATKINS: I wanted to be an artist, and it has been a topsy-turvy winding road. I did my undergrad at UC Irvine, and I studied video art.

SHERVIN KHODABANDEH: What's your greatest wish for AI in the future?

ELIZABETH ANNE WATKINS: There's so much that AI can do that humans cannot, but there's also so much that humans can do that AI cannot. And I think a big challenge, at least for me, and I hope the people with whom I work going forward for the next few years, is going to be figuring

out exactly what that balance is and how can we systematize “What do people really need help with that AI systems can do?” but with a really thorough understanding of what it is that people do and how they do it.

Now that the AI systems are becoming so advanced, but oftentimes in a lab or in a vacuum, as they mature into the real world, I think it’s really exciting what they can do, but we’re facing a lot of work in getting them there. And I think it’s going to be other social scientists and multidisciplinary teams like the ones that we have at Intel all working together to make sure that these systems can really be deployed as solutions.

SAM RANSBOTHAM: That’s a good cap for the episode. Thanks for talking to us. I think one of the things that has come through [is], you’re talking about projects and things you’re working on, but you’re also giving us a hint about what’s going happen in the future as we move off of the, let’s say, tool-focused “Can we get the ML right? Can we get the model right? Can we get the data right?” You’re talking a lot about what happens next, what happens once we check off some of those check marks — what are

those checks? And I think a lot of people can learn from that. Thanks for taking the time to talk with us today.

ELIZABETH ANNE WATKINS: Oh my gosh, it’s been such a pleasure.

SAM RANSBOTHAM: Thanks for listening. Next time, Shervin and I speak with David Hardoon from Aboitiz Data Innovation. We’re once again talking about chemical engineering, so you won’t want to miss this one.

ALLISON RYDER: Thanks for listening to *Me, Myself, and AI*. We believe, like you, that the conversation about AI implementation doesn’t start and stop with this podcast. That’s why we’ve created a group on LinkedIn specifically for listeners like you. It’s called AI for Leaders, and if you join us, you can chat with show creators and hosts, ask your own questions, share your insights, and gain access to valuable resources about AI implementation from *MIT SMR* and BCG. You can access it by visiting mitsmr.com/AIforLeaders. We’ll put that link in the show notes, and we hope to see you there.