

Podcast: Navigating Ethical AI

00:00:01:00 - 00:00:04:05

The HCLTech Trends and Insights podcast.

00:00:05:02 - 00:00:38:02

Hello everyone and thank you for tuning into the HCLTech Trends and Insights podcast where we discuss the latest technology trends and stories that are impacting business and society. I'm Nick Ismael, the global head of brand journalism at HCLTech. And today I'm delighted to be joined by Kristen Ladd, senior responsible AI program Manager for Cognitive Services at Microsoft and Andy Packham, chief architect and senior Vice President for the Microsoft Business Unit at HCLTech.

00:00:38:03 - 00:01:12:03

Kristen and Andy, how are you both doing? Well, good, thank you. Excellent. And thank you for joining us for this episode. So as we know, artificial intelligence has rapidly advanced in recent years and it's reshaped industries and even our daily lives. But with greater power, of course, comes greater responsibility. And as with all innovations, we must approach A.I. development ethically and responsibly for it to impact society positively.

00:01:12:04 - 00:01:48:01

Concerns about ethical conduct are widespread today and are driving the adoption of related legislation worldwide, including in the EU and the US. When implementing A.I. technology, businesses are ultimately responsible for weaving in principles like transparency, safety, unbiased principles and so few things lead to discrimination. And today, we're going to discuss the importance of ethical AI and how organizations can scale it effectively.

00:01:48:02 - 00:02:21:15

So let's get started. First, talking about scaling ethical AI, what are the key challenges while identifying and mitigating bias in AI systems? And how can these challenges be addressed to develop fair and unbiased AI models? Kristen, Would you like to start us off? Absolutely. So

while there's a large set of different responsible AI harms that we could consider in this domain, let's zero in on bias and fairness related harms like you mentioned.

00:02:22:01 - 00:02:46:12

So I'd say there are a few key challenges to identifying and mitigating bias and it really does start at that identification step. So first it's considering what type of AI system are you working with. If you are training the model yourself. The ways that you'll identify are going to be quite different than if you're picking up and using a generative AI model developed by another company.

00:02:46:14 - 00:03:12:15

So identifying what type of AI are you interacting with? And then there's also an important step of figuring out what type of fairness related harm is most important to you. Is it quality of service where it's going to be more differences in quality for how it performs for different people? Is it harms of allocation? So is it making decisions that impact people differently or particularly in the context of generative AI?

00:03:13:03 - 00:03:37:09

That's where you start to get into representation harms. So is it a racing certain groups? Is it outputting stereotypes or demeaning them? And then considering how your deployment, environment, geography influences all of that. So really understanding what is the context in which your system is operating and then identifying what type of fairness is going to be most relevant to your system.

00:03:37:10 - 00:04:08:11

Because answering those questions is going to directly influence how you need to mitigate across the entire stack. For some systems, you'll be able to measure and retrain the model yourself. For others, you're going to need to use different strategies to mitigate where you may need to just understand what are the limitations of the model you're working with. On top of that, what safety systems like content filters are available to you that can influence how it performs?

00:04:08:12 - 00:04:49:01

And how do you add additional mitigations on like communicating limitations to your end users and having good feedback mechanisms in place? So it varies by system and across the full stack of the AI system, how you mitigate. Absolutely. And I'm sure that variety kind of adds to the challenge as well. Absolutely. The variety really just underpins the need for deeply understanding and including responsible AI in part of the conversation so that it just becomes a regular part of quality and development.

00:04:49:04 - 00:05:18:15

And as you get more used to interacting and leveraging AI systems in your own solutions that you're building, you get a better sense of how you might need to adjust it in a new application. Okay, great. Thank you. And Andy. From your perspective, what are the kind of key challenges while identifying and mitigating bias in AI systems and how can they be addressed at this and really go into some of that core detail?

00:05:18:15 - 00:05:42:13

And this is just, you know, it's such a critical subject, but actually it's still relatively new. Looking at number of Google searches on responsible AI and, you know, for years and years it was almost flat. And then just about this time last year, funnily enough, it just, you know, the whole thing's exploded. So there's no definitive answer.

00:05:42:15 - 00:06:12:00

So in another study recently, the vast majority of, you know, senior leaders across every industry see ethical AI and responsible AI as being super important because there's only about 6% so far have actually managed to successfully implement. So it's certainly it's certainly complex. I think, you know, AI is based on data. So any problems with the quality of the data, the bias in the data are seriously going to be addressed.

00:06:12:01 - 00:07:05:07

But I think it's also important to look broader across the organization responsible. I must be framed within a broader ethical and responsible framework for the whole organization and getting the right culture right. So it's not just that the system itself is not biased, but it's been

used responsibly and there's a lot of focus on maintaining privacy and so it's more about focusing that across the whole of the organization, which means really establishing a very strong governance function with executive leadership that really has that, you know, authority to control what the, you know, which each of the teams that's building these solutions is working on, but should see their role very much in terms of

00:07:05:07 - 00:07:44:00

a guide and advisor. You know, like I said, this is all new stuff. So a lot of advice is needed and making sure that you're building that out with very multi-disciplinary, diverse, inclusive teams. There are the solutions are not simply technical and they've got to, you know, respect a whole variety of inputs and views that and then, you know, making sure that everybody has been trained and really understands not just the technology but the impacts and how to recognize how to respond to and avoid bias in the systems.

00:07:44:01 - 00:08:24:03

Thanks Andy. And you mentioned that it's very much the beginning of this journey. Are there any kind of outliers? And if so, can you highlight any real world examples of organizations that have successfully scaled ethically AI and the benefits that they've achieved? So I think there are organizations that have already moved quite considerably and I think there are ones that are already kind of, you know, potentially in, you know, a regulated industry, but already been giving a lot of free, you know, thinking to how that they were viewed and, you know, in a broader scale.

00:08:24:05 - 00:08:54:12

I think for me, I was reading a slide management review a few days back about Unilever and how they've implemented their responsible AI network. The two things that really jumped out to me was that they started by framing this within a much broader framework. There was across the organization that were very clear about accountability. And there was always a there was always a Unilever person in the process.

00:08:54:12 - 00:09:27:10

And the blame was never attributed to the system. It was always a it was a person. So I think that's important that there's a level of trust built by this, this person of accountability. And I think also it was really interesting for me, it stood out is that in in that process of looking and understanding the data, they started to look beyond not just whether or not that diet, that the data was bias, but whether or not actually it was it was it was actually useful.

00:09:27:11 - 00:09:54:00

So there were areas where they had improved their business processes by thinking about data in a different way. So there are organizations that are moving very quickly in this space, you know and I think the publicly available frameworks, you know, like from Microsoft are really helping organizations kind of pick the path through this complex journey.

00:09:54:01 - 00:10:22:12

Thanks Andy. And Kristen, how about yourself? Can you provide any kind of real world examples of organizations that have successfully scaled ethical A.I.? Is that anything that you can tell us about what's going on at Microsoft? Yeah, definitely. I'd say it's been really exciting seeing responsibly AI become a much larger part of the external conversation. And at Microsoft, we've been our on our responsibility journey for several years at this point.

00:10:22:13 - 00:10:58:15

So the first version of our Responsible AI standard came out internally a few years ago and we were just able to release it externally for the first time once it reached a different level of maturity in 2022. But it's been several years of really internalizing these lessons, figuring out how do we go from principles to practical implementation. And I think part of what be touched on is the importance of sharing this across the broader ecosystem, making sure that there's a good governance structure in place and really having some solid education in place as well.

00:10:59:02 - 00:11:41:06

So while we've been on our responsibility journey for several years at this point in terms of how do we implement it, a lot of that focus really has been on how do we embedded in product development across the company. So we've got a governance organization called the Office of Responsibly AI programs like Sensitive Uses that interact with some of the higher risk A.I.

systems and then really try to spread it out through all of our business groups, through our different responsibly AI champs and partnering directly with customers is a huge part of figuring out where is this working and where do we need to continue to adjust?

00:11:41:08 - 00:12:05:11

That's great. Thank you. I now want to kind of pivot topic slightly and talk about effective communication of ethical benefits. And I suppose that comes in into some of the things you were talking about in terms of it needs to be a strong function in the leadership. Everyone needs to be trained, everyone needs to be accountable. So to start and Andy, let's let's start with you.

00:12:05:12 - 00:12:37:08

What strategies can organizations employ to effectively communicate ethical benefits to customers? I think, you know, every customer has got this reasonable. I mean, it's absolutely valid assumption that they're going to be treated fairly and that their personal information is huge, responsible. So it's important to build trust with the customers. And, you know and like I said about framing this, then it needs to be more than just trust about AOA.

00:12:37:09 - 00:12:58:05

It's a whole brand trust. So I think with AI, one of the key things is, you know, make sure that your authentic, you know, if you're implementing, you know, for example, if you're implementing an AI driven chat bot, make sure you know, the customers are clearly aware that they're not speaking to a human. Don't try and pretend to be something you're not.

00:12:58:06 - 00:13:31:07

And because, you know, if you if the customer sees through that, then trust will be lost immediately. So it's very much about building trust. And I think a lot of people really worry about the loss of autonomy and agency and see AI itself as being, you know, unemotional and uncaring. And actually that's true. It is. So having, you know, having the human or the expert in the loop is really important.

00:13:31:07 - 00:13:53:14

The still needs to be a a channel of getting through to someone who can help rather than just be routed through the same systems time and time again. So I think it's an it's an extension of, you know, how do you effectively communicate your brand and your brand values and the trust that that brings to the customer without trying to hide?

00:13:54:00 - 00:14:27:10

I as AI as something that it's not in the background. Absolutely. And Kristen, is there any kind of strategies that or ways you've seen the benefits of ethical AI being communicated that you can share with us? I think I like to divide this into two big buckets. One is responsible development and the other is empowering, responsible use. So the first responsible development bucket is how do you as the organization in building the AI tech infuse responsibly AI into it?

00:14:27:10 - 00:14:51:08

So it's just part of the experience when customers are interacting it with it. And this includes some of what India was touching on, including disclosing when content is generated, including some clear feedback mechanisms pointing to help pages to learn more and really doing some of that behind the scenes testing and development to minimize some of the responsible AI issues.

00:14:51:09 - 00:15:23:06

So one is just creating that out standing product experience. But the other part of the story is empowering, responsible use and that's where a lot more of the transparency and communication strategies come into play. So either putting together transparency notes or documentation to share more information about the system or even in some of the sales or the marketing materials, making sure that responsibility is part of the story because this is something that people are asking about and care deeply about it.

00:15:23:08 - 00:16:06:03

And it can be really powerful to share your approach as a differentiator. Thanks very much. And in terms of those strategies, I, assume that key and kind of overcoming potential skepticism or resistance from customers, but are there any other kind of tactics you can share to overcome that traditional skepticism? Yeah, I'd say the first step there is understanding the

reason for the skepticism, because that's really going to inform what's most important to continue in that conversation.

00:16:06:04 - 00:16:53:10

Sometimes this skepticism comes because there's fear of a lot of implementation difficulty because this is new order to people. So if that's where the hesitation is coming from, that this is going to be difficult or costly to implement or it won't be achievable, then it becomes valuable to talk about. Here are the resources and tools available to empower this, whether that's pointing to tools like stair, learn or interpret email or content filters that can help make it more achievable and implementation patterns that can help overcome some of the resistance if it's implementation difficulty or if some of the resistance is just this is really new or they need to see some examples, then that's where

00:16:53:10 - 00:17:08:11

more of the focus can be, sharing some first principles and ways of thinking about this and engaging in some of those earlier educational conversations around what this means and how to consider it as part of your AI strategy.

00:17:08:11 - 00:17:13:06

Thank you. And Andy, did you have any thoughts on that before we moved on to the next topic?

00:17:13:06 - 00:17:29:09

So, Nick, I think, you know, Kirsten covered it absolutely perfectly. It is all about making sure that you address, you know and there's you address the customer concerns and it's a feedback mechanism.

00:17:29:10 - 00:17:53:11

So the trust continues. We still got we've still got to go some way with trust. But I think the more that we talk about, you know, I more that we talk about responsibility, you know, in

multiple forums and think we'll have more kind of awareness that it's that it is something that we you know, certainly we have an industry have got a huge opportunity to control.

00:17:53:12 - 00:18:21:04

It is an exciting time. But I think that also does I think, as you said at the beginning, you know, we you know, the companies are a lot of responsibility on that to make sure that we are guiding and leading the industry on this. Thank you. And now just moving on to the final topic of our ethical discussion, which is the role of ethical AI in driving innovation.

00:18:21:05 - 00:18:54:05

So Andy, we've talked about how you scale ethical AI, how you communicate the benefits, but how can it actually help drive innovation within an organization? Well, I think innovation innovation actually takes a considerable investment. It's, you know, it's expensive and it's hard, hard work. So making sure that the decisions that your the outcomes and decisions you're taking based on the data is without bias.

00:18:54:07 - 00:19:26:07

And so that you're going to be able to develop something that at least gets to market the people responsible with positive outcomes. I think we're going to see a lot of regulation. This area, you know, especially in some of the regulated industries with governments and regulatory bodies taking a lot more interest in this. So any company that's already thinking through ensuring that they're able to meet the expectations of the guidelines is going to be is going to be ahead of the curve.

00:19:26:08 - 00:20:00:13

You know and I think if you know, if you see algorithms, the continued existing qualities, then, you know, it's probable that you're going to miss entire markets or you're going to be offering wrong products or solutions to part of that market, which is mainly that's going to be just a lost opportunity. And whether or not that's in a that's opportunity which to measure financially or in terms of reputational risk and impact, that will have a negative impact.

00:20:00:14 - 00:20:28:02

So I think ensuring that you've got the right frameworks in place, you've got the right controls in place will be something that will ensure that innovation is much more successful. I read recently about it was in a UNESCO report where there was a, there was some sort of research done against various job sites that showed, depending on the gender selected, different roles were presented.

00:20:28:02 - 00:20:54:15

If you, you presented the gender as male, there were more leadership roles. If you presented it as female, there were there were less higher paid roles. Now, that platform must have taken a lot of investment and time to build that because they hadn't really thought this through. You know, it's likely that some absolutely brilliant candidates would not have that role presented to them and they would never apply.

00:20:55:00 - 00:21:19:12

So, you know, with that bias, without, you know, being able to understand that that return on investment, you know, whether or not that will eventually come through is hurting reputation that you're losing an opportunity there in the market. Yeah. I mean it's a fundamental consideration now moving forward how to innovate and that will kind of impact the success of innovations.

00:21:19:12 - 00:21:46:00

I think I understand. And Kristen, is there anything you want to add on how ethical A.I. is actually driving innovation? I love this question because there's so many ways in which responsible AI leads to better products. A lot of the responsibly AI practices that we see organizations implement fundamentally improve the product quality. So we take some of the generative AI examples.

00:21:46:01 - 00:22:08:15

A lot of the responsible AI benefits come from appropriately figuring out ways to constrain the system to really only do what you want it to do for your application. Well, which does take work because they're very, very broad systems and that both mitigates a large number of responsibly AI concerns but leads to products that work better for your scenario.

00:22:09:01 - 00:22:54:11

And then I'd say the other way in which it drives innovation within organizations is responsibly. AI requires deep, deep stakeholder empathy because this is a very socio technical problem. And with that stakeholder empathy in terms of considering who might be impacted by the system, it can lead to some really significant product changes and improvements. So we've been talking a lot about fairness in this episode and it really requires understanding who are all of the different users, what are their different demographic backgrounds, how might they interact with the system and addressing those issues often leads to better overall quality.

00:22:54:12 - 00:23:43:07

Yeah, that makes sense. And can you provide any examples of these innovative solutions that been enabled by ethical homes? Supply Yeah. So if we take speech to text or speech recognition systems focusing on fairness, really required diving deep into people of different linguistic backgrounds. So people from black communities, Latin communities all across the world, different communities in different ways of speaking, to make sure that these types of speech recognition technologies, that the quality was really looking at the intersection between those different demographic groups and ensuring that we're getting the best quality that we can across those.

00:23:43:08 - 00:24:11:06

Another example is if we take a look at a bunch of our different copilots, so I'll pick on some of the M365 copilot where it uses AI generative AI to help you search over a lot of your different Microsoft Office apps and interact with it responsibly is infused throughout the user experience, throughout the incident response process and in the actual instructions that are given to the model.

00:24:11:08 - 00:24:42:12

Andy as we worked with teams to really improve those responsible experiences and metrics, we saw a lot of improvements in quality overall in terms of the system just working better, doing what people wanted it to be doing. And in cases where it was off track or didn't get it quite right, incident response procedures were improved because there is a better understanding of the types of issues that might be encountered and how to respond to them.

00:24:42:14 - 00:25:07:13

Yeah, it creates a more personalized, receptive experience for the user. And Andy, is there anything you'd like to add on specific examples of innovative solutions using at the on? I think what really excites me is the opportunity here to turn this around 380 degrees. You know, we know in areas such as Gopnik you know, example about a job site.

00:25:07:14 - 00:25:31:06

So we know, you know, there's areas and use cases where, you know, the impact of bias has been well documented to turn this around and use responsible AI to ensure that you are not being biased. I think it's a massive opportunity. So you know, you can do that in hiring and you could do that in maybe in financial product design.

00:25:31:07 - 00:25:51:15

You know, we know that there has been bias in a lot of the models in in the past being able to use responsible AI not just to ensure that the model is, you know, is not biased as it should be, but also to point out, you know, where biases and unconscious biases may be coming in so that we can tailor those solutions far more appropriately.

00:25:51:15 - 00:26:22:15

And then, you know, very sensitive. But in healthcare, you know, a lot of medical models have been trained on on research data that's predominantly, you know, I'm, you know, predominately based on males white from Western Europe. So not only ensuring that that bias is removed, but also helping adopters understand and recognize where the, you know, they might be making assumptions that are now, you know, proven to be no longer valid.

00:26:22:15 - 00:26:53:06

I think it's hugely exciting how we can turn this round and make it rather than a kind of a role to avoid bias, making it a role to actually enhance the way that we engage with the customers. That's great. And just before we close Andy and Kristen, is there anything you'd like to add on the future of A.I. and, you know, enabling it responsibly and ethically?

00:26:53:07 - 00:27:26:03

I think the biggest thing is going to be just staying up to date with the technology as it's changing. But really considering how do we do this across the full application stack, if you think about stacked layers of Swiss cheese or a defense in depth strategy, no single mitigation is going to be perfect. And that's where we really need to think across the model itself, the safety systems we build on top of it, the user experience, transparency and incident response and ongoing improvement procedures.

00:27:26:03 - 00:27:48:07

So I very much encourage folks to consider responsibly AI across the full product development lifecycle and every layer of it. Thank you. And Andy. Yeah, I think for me two things. One, you know, we've got to work in a way, would it work to remove the label? You know, it should be AI and it shouldn't be responsible.

00:27:48:07 - 00:28:19:06

I am. But I think what I'd be really looking for and really excited about is when we start seeing, you know, ethics and the responsible use of technology beginning to be embedded in maybe university training programs in computer science. So the it becomes core and fundamental to culture and our behavior. I think that's a nice way to end and I'll just do a little some up of the discussion we spoke about a lot of things.

00:28:19:06 - 00:28:51:04

So I can't include everything. But I think we started with talking about how ultimately ethical A.I. is absolutely fundamental. But many organizations and are at the very beginning of that journey, a strong governance function is needed within executive leadership. There's lots of advice needed as well and everyone needs to be trained and accountable. Responsible A.I. has become a large part of the external conversation and now it's a case of moving from principles to practical implementations.

00:28:51:05 - 00:29:18:11

There are questions that are rising about how to embed responsible A.I. and product development and a partner ecosystem for this is crucial for sharing the benefits or communicating the benefits of ethical AI requires responsible development and also empowering response for use. We need to make sure that data that is used to make decisions, to innovate is to innovate, is without bias and ethical and responsible.

00:29:18:14 - 00:29:47:00

AI ultimately will lead to better practices and improve products quality. And it also requires greater stakeholder empathy, which can lead to products, improved methods. And just finally, what you ended on Andy, it's important that we move to a place where we remove the label and make responsible ethical AI the norm and bring it into the mainstream. So with that, I would like to thank you both very much for your time on today's podcast.

00:29:47:01 - 00:30:00:02

Your insights have been really valuable and interesting. Thank you very much. Thank you. And thank you to the audience for tuning in. We'll see you at the next one. Goodbye.

00:30:00:02 - 00:30:01:03

Goodbye.

00:30:01:03 - 00:30:08:12

Goodbye.

To learn more about what HCLTech and Microsoft are doing together visit- [Microsoft Ecosystem Unit](#) page