The HCLTech Trends and Insights podcast

Nick Ismail
Hi Heather, thank you for joining me today on the HCLTech Trends and Insights podcast. You've just joined HCLTech as the head of responsible AI, so let's start there. And what sparked your interest in this subject and you know what excites you about the future of AI.

Heather Domin
Yeah, Nick, it's great to be here and talking about this topic and it is an exciting topic. I certainly have had an interest in the space for a long time and have been fortunate to be able to work in it. You know, my career, I started off really looking at, you know, AI, we didn't necessarily think about it in the same way we do today, but, you know, I was always fascinated, even from my early education and career days and the intersection between, you know, the possibilities, the things that we can do with technology, but also the potential, you know, for risk and how we can mitigate that, right, so that we can get the full, full potential, because there is so much value in technology when done well.

Nick Ismail
Sure and you mentioned, it's changed a lot and there's no doubt that AI is evolving very fast and the conversation around responsibility is growing just as quickly. So from your experience, what's the biggest misconception people have about responsible AI?

Heather Domin
So I think that people generally are aware that there are risks where they either hearing the headlines, they're seeing things in the news and trying to understand, you know, if they should be concerned. I think most people are concerned about some of the bigger, you know, safety issues and longer term risks and they tend to also be thinking about impact to the way we work and jobs. Those are very common, common things that that people think about. And I think, you know, because of that, people maybe are not aware that there's a lot that we can do right. They may overlook some of the shorter term, more focused things that we can do at this time, including setting up the right processes within organizations, committees and groups to help manage the risks offices such as such as the one I lead. And you know, they also are not aware that you can do a lot of testing and techniques within the life cycle of building an AI system and deploying it that help protect people from potential risks. And you know, so I think sometimes that misconception is just that there's only these bigger picture, longer term risks that we need to be concerned about, when really there are quite a few shorter term, smaller issues that we can actively address within organizations.

Nick Ismail
That's great and we're going to be talking about that today in the light of the release of HCLTech's whitepaper in partnership with MIT on implementing responsible AI in the generative age. The findings from that report found that while 87% of executives believe responsible AI is important, only 15% feel ready to implement it. So why is responsible AI such a high priority now in the short term, as you mentioned and why do you think businesses are struggling with implementation?

Heather Domin
Yeah, so, you know, I think it's, it's interesting. I've had the opportunity to work in this space for a number of years, and what I've seen is that leaders, I think even five years ago, were quite aware that responsible AI, was important. But what you see in the study, and what I'm hearing and seeing with my peers and in many organizations, is that they are actually now taking active steps. And I think sometimes it takes time for people to understand the issues. It takes time. Time to build awareness on the practical steps, and we also now have many tools, techniques, certifications, pathways that can support the implementation. And so I think we're really now starting to see more investment, and it's reflected in the study, right, that people are thinking about responsible AI as a major area of investment in the next 12 to 18 months, right? This is a short term focus and a strong focus for many organizations and I think that this is people are understanding that, and we've seen it happen where, you know, AI deployments actually don't become successful unless you have responsible AI there and it can cause you a lot of reputational damage. It can cause a lot of rework and expense, potentially fines. Now, with with legal, you know, requirements in place that people are understanding, you know, I need to invest, I need to make sure that I'm taking the appropriate steps to make AI successful. And why do they want AI to be successful? There's huge productivity gains. People often find it's actually more enjoyable, right? A lot of the generative AI tools are actually fun to use. So it can prove the way we work. It helps with upskilling and and creativity and many things that organizations want to take advantage of because they also understand that can lead to a competitive advantage. And as you are able to leverage AI and really take your business to the next level, take advantage of the productivity gains and other aspects, then the responsible AI is an enabler that allows you to do it with quality, you know and have a successful long term implementation of your AI.

Nick Ismail
Sure and if It's not implemented responsibly, what could go wrong?

Heather Domin
Well, so there are many things that we have to be concerned about. We can think about things like fairness and bias. We've seen issues in, you know, if you're if you're using AI and hiring or, you know, with public benefits, and many things are, there are cases where, unfortunately, you know, without the right training data, without the right testing techniques, unfortunately, you know, people's lives can be impacted, and we need To really be thinking about mitigating for potential issues with with bias and ensuring fairness. And you know, obviously, security and privacy remain a very big concern. We need to be having the appropriate guardrails so that you know if, if people's personal information is not protected right or leaked. In the case of generative AI, that can be an issue. And we also need to think about things like transparency and if we don't have the appropriate communications and enabling people understand that they're interacting with AI information about the models and the systems that are being deployed that can lead to breakdown in communication and a lack of understanding that often is problematic. So a lot of this is, you know, there's some that also involve, like, having a human involved and checking, right? So we think about that lot in the context of generative AI and making sure that humans are actually, you know, when they're using the output or checking appropriately.

We've had cases like in the legal domain, where people, unfortunately, you know, leveraged, you know, court cases that were actually not not real, right? They were generated by the system and then that came out later, right? So if you have, as the humans, consume the output, right, we we need to also be aware that we also have a responsibility to to check right and use that data responsibly.

Nick Ismail
Sure I was, I was actually going to ask you, what are some of the biggest AI missteps? That you've seen and you reference the legal example. Are there any others?

Heather Domin
Yeah, I mean, you know, I think there's some funny ones, like with, you know, the recommendation, you know, to keep it was to keep cheese from falling off pizza, right? To use glue, right? I think we've all heard of that one, AI summarization and some of it, so some of it's funny, right? And, but when we think about though more practically, right? So some of these and this, I think, is a large part of what people are concerned about nowadays is like the misinformation aspects, and we've seen, unfortunately, you know, bad actors can actually, you know, misuse that and have misuse that in camp, in political campaigns and sometimes It's actually for profit motive. We've seen, you know, people use that in the case. So there's one, one that's particularly concerning, I think, for a lot of people also, is the impact on children, right? So if you see the use of generative AI to generate like science videos that are not true, right? So we've seen this and you know, I think that organizations, as they develop content, really need to be checking and making sure that the content is used in a way that is not misleading. So I think some of that is to people over trusting the content and not taking appropriate steps to check.

Nick Ismail
Okay. And I'd just like to ask about regulation now, because you did touch upon it earlier with AI regulations like the EU AI act coming into play, it's a good thing and a positive thing, but businesses often see compliance as a burden rather than an opportunity. Do you think ultimately regulation will help drive responsible AI adoption, or will it hinder AI innovation?

Heather Domin
I'm a big believer that it will help. I've seen this in practice. I know as I work with developers and data scientists and others throughout the organization, what happens when you don't have an appropriate amount of regulation and standards in the industry is that people essentially don't know what to do, they don't know how to align, there has to be a line drawn somewhere and people struggle with what that line should be. So you end up spending a lot of time, especially organizations that want to do things right and but when you don't have clarity on what that is or what the expectation should be, it can slow you down. Basically, you don't have clarity, so you can't move forward quickly and so you spin on decisions and even small ones. For example, when we think about measuring bias, it actually requires many times looking at the protected attributes, the characteristics that we want to avoid being biased against so that could be gender, age, you know, that sort of thing, right? So we have that, of course, but the specifics of that are actually harder to define.

So at what age should we cut off? You know, what groups should we put people in? And how do we test for that? That actually, in many cases, is now defined by law and in some cases, like there's a New York City automated employment decision to law, which actually specifies the calculation for the expectation, expected way in which they want you to calculate and measure bias and fairness. And you know, this, this, this couldn't have its own challenges in terms of being, you know, so specific, but it can also help people move forward, because there's real clarity. And so I think that, in my view and in my experience, regulations when done well and sensibly and I do think many policymakers these days are seeking feedback from industry, I think that's the right way to do it. And, have, you know, the appropriate input and the appropriate regulations. It actually helps speed that up. And I think that's why you see a lot of major tech companies, executives from from tech companies supportive of, you know, the appropriate regulations.

Nick Ismail
Okay? The report found that only 23% of executives feel ready to manage user adoption and change when introducing AI. How do you build trust in AI so that people actually want to use it? Is regulation part of that, or are there other elements that need to be incorporated?

Heather Domin
Yeah, well, I think regulations and standards, so a lot of it, you can verify. You can do, you know, check nowadays they're, you know, standards like ISO 42,001 and and others actually provide, I think you know a good amount of confidence to people about that there are appropriate controls in place, and so that can certainly help with trust, and of course, that you're checking and complying with the laws, yes, but beyond that, training and helping people understand what has been done to ensure protections is important. So I think introducing the AI system and its functionality, you need to do that, of course, so that people feel comfortable with the technology itself, but then also helping communicate the protections and guardrails that have been put into place that can really go a long way in helping people trust.

Nick Ismail
Okay and looking at things like shadow AI, which is when teams start using AI tools without oversight. We've heard of shadow it before. Of course, it can lead to security and compliance risks. Is this a big problem? And what can companies do about it? If it is?

Heather Domin
Yeah, so certainly it can be an issue that's that's absolutely right and I think it goes and this is why people within the industry have often talked about things like creating inventories of AI systems and making sure that we have awareness of, you know, what's happening in our AI environment. There's lots of tools nowadays that can help actually detect shadow AI and help organizations manage that. It is, though, of course, something that has to be actively managed. And you know, but if you don't have awareness of that, it's absolutely true, you could have something happening environment, just like with security, right? That's why we do active security monitoring and controls. We need to think about that also in a similar way for AI and what's happening with our AI systems, where are they? What data are they accessing,

what models are being used, API calls, things like that. And then you know from there, once you're monitoring, taking appropriate action when, when there's risky instances or or data or models being used.

Nick Ismail
Okay. I think we've spoken in the context of the report about the importance of responsible AI, the challenges of implementing it, how to manage user adoption, the importance of regulation. Are there any other findings or insights from the report that you found particularly interesting or surprising? Perhaps.

Heather Domin
So. I think that the really interesting thing to me is that a significant percentage of the course of the respondents is about 58% are really confident about their data privacy and security practices right that they understand how to do data security and privacy to a greater extent than some of the others, like fairness and bias. And what I think is interesting about that is I worked with many privacy professionals and what they tell me is that AI is sort of like 510, years behind where privacy and, you know, security really have been right. So there was initially when the privacy laws and requirements started to. Come out a similar situation where, you know, people didn't understand. What we see in many organizations is a really robust way of managing that, and in their certifications and professional societies, and, you know, lots of ways which we support privacy professionals and AI is now slightly behind, but it's getting there, and that gives me confidence. When I see that privacy and security measures are people are more confident about that, that in the future, we're going to see the same thing for AI and so that, I think was a positive thing that I really stood out to me. I think, you know, it was a interesting also, just to see that organizations are really prioritizing investments in the space. I think that that is slightly surprising to me. I've actually done quite a bit of research on you know that the investments that organizations are making over time and in the last few years, it's been challenging for people, right? They've, they've, they've organizations, even the leaders spoken to many, many leaders in organizations all over the world and various types and when they've gone to make cases for investment in responsible AI, they're often met with challenges around prioritization and justifying the return on investment and that sort of thing. And what is surprising and a good way is to see in this study that, you know, investments are a priority and they should be in the in the short term, actively increasing. So that is a that's a great thing,

Nick Ismail
Amazing. I'd now like to talk about some practical steps that organizations can follow to implement responsible AI and explainable AI and embed a culture of responsible AI. So let's start with the first one. Can you provide in your experience, practical steps that organizations can follow to implement responsible AI, effectively and sustainably.

Heather Domin
Yeah, so you know, I do think that starting with the people and looking at the organizational aspects is a great place to start. Many organizations, even small organizations, can do that right kind of get there, whether that's an AI committee of some sort that helps to assess risks and make decisions when needed, that's something that I do know many organizations actively take on, and it's a great first step and that means you're starting to really think about the operational risks and some of the challenges that you may face in AI. It can also be part of strategic decisions around what AI investments to make and where to deploy and what to do with AI you may set up then, you know, in larger organizations and maybe some smaller ones too, an office like what we have, you know, an office of responsible AI and governance, you can certainly do that, and that'll help to institutionalize and, you know, provide a way to roll out the responsible AI aspects more broadly. And then you know, from there, what you'll find with with dedicated focus, right, that people can then build on that and that it it will grow over time into other things you need to, of course, be thinking about. You'll have technical teams that will need to focus on the the AI life cycle, doing the appropriate testing throughout the life cycle the AI system, the data, as well as pre deployment and post deployment. Ongoing monitoring need to have testing that you're doing. And then, you know, obviously there are quite a lot of things ongoing as you monitor the you know users needs and you're gonna have to do upskilling and training within your organization, there are a lot of stakeholders involved in an AI life cycle, and that's one of the things that actually makes it. Different than other types of IT implementations. It's a socio technical system that has a lot of different stakeholders involved. So to really get that right level of user adaption, you will have to consider and work actively with many different stakeholders. It's also important, as we've seen many times during the AI life cycle development itself, that you have diverse teams involved. And then over time, you know you can, you can, essentially, you know, work to improve. So you have all of your you know foundational policies, your organizational management processes in place, your testing. And then over time, you're going to improve and there will be opportunities more and more so over time, as you get feedback from users and from employees and as the standards and requirements evolve to continue to improve and grow your program and that's that's often a trajectory that we see happen.

Nick Ismail
Yeah and you mentioned it's something that's adopted by many stakeholders. This technology isn't just something for the IT department. It's a cultural and leadership challenge to implement responsible AI more so than it is a technology challenge. Would that be fair to say?

Heather Domin
Yeah, I think it's fair to say there's certainly a big difference between the way we developed and deployed it systems in the past and how we are to doing it today. So we have AI systems that are often deployed in a matter of weeks or months right now that people are able to leverage generative AI and foundation models in the past, you had to go through a long process of gathering user requirements and did user testing and it was, you know, a much more. It process focused deployment and now you have users that you know, are actively involved and and from an early stage and actually the the deployed setting right? And so I think it's, you have a, also a broader user base, potentially right as many organizations look to deploy across broad areas of their business for productivity reasons.

It could be chat bots or other generative AI systems that actually help to generate responses to customer inquiries and that sort of thing, right? So there are many times where IT systems, I think, were more focused in the past, you have a broader user base and that does require quite a lot of user engagement, feedback and at an earlier stage and faster. And I think that is part of the challenge with the adoption aspect, right? Is just how quickly we can get these systems out. It's a good thing in many cases and it's also a challenge to make sure that people are comfortable and able to adopt in that short time frame.

Nick Ismail
Thank you. And finally, I just want to talk about HCLTech and our approach to responsible AI. We've just announced a partnership with the responsible AI Institute. What is that and how will it benefit us and our efforts and contributions to responsible AI as well as our clients.

Heather Domin
Yeah, so the responsible AI Institute membership is very exciting. We just kicked that off and we're thrilled to be, you know, a member of the responsible AI Institute and that gives us quite a lot as an organization and that we can then also help our clients with in terms of a community that is actively pursuing the next level of responsible AI and how do we continue to evolve that as a field with best practice sharing and information and opportunities to collaborate and learn from our peers, all of that really takes us to a level where we can grow our own practices and help our clients do the same. Um, and it's exciting, you know, we need that, just like in any area, we all need to be part of a broader community, part of, you know, understanding what's happening in the space as it continues to evolve. And I see that responsible AI Institute membership as a key piece of that picture.

Nick Ismail
That's great. And finally, can you just outline HCLTech's approach to responsible AI, both in the programs that we undertake for our clients and also in our own internal developments?

Heather Domin
Absolutely. Yeah. So we do take a very comprehensive approach to responsible AI, it's rooted in our values that we have as an organization and as we take that and we actually look to implement what we think of as our cornerstones around accountability and fairness and security, privacy, transparency, those cornerstones help form what we can do, both internally and with clients, in terms of things like consulting with clients, helping them contain and manage issues and continue on their path of growing their AI deployments and actually growing their business and monitoring for opportunities. So that is, you know, we built quite an extensive framework, risk management framework and program at HCLTech, which aligns our internal work, and we take that, we can, we can actually leverage that to work with clients who are often facing very similar challenges. And so it's one of the benefits as a is a large organization, technology organization works across many different industries and verticals, where we can take that and share those learnings with organizations of all sizes. And I think it's a wonderful thing.

Nick Ismail
Heather, thank you so much for your time and insights.

Heather Domin
Absolutely, it's great to great to be with you. Nick.