

AI Business School Podcast

Episode3

Presenter: David Carmona; Guillermo Carmona, Nick McQuire; Natasha Crampton; Sarah Bird

AI BUSINESS SCHOOL PODCAST – EPISODE 3

[MUSIC]

David Carmona: In a previous episode of this podcast, I told a story of my father, how he was a carpenter at heart and that made him better at his business. What I didn't mention though, is that he also wanted me to be a carpenter and follow his path, and we all know how that ended, I'm now presenting this podcast about AI. I'm now making the same mistake with my son, Guillermo. I really want him to learn computer programming, and that's why I was so happy when he was ten and he told me that he wanted to learn Python. However, he told me something that made me think. Guillermo, do you remember what you told me?

Guillermo Carmona: Yeah, I think I told you that you had to be with me the entire time when I was learning.

David: Yep, that's exactly what you told me. And why did you tell me that?

Guillermo: I was scared that I would create an AI that could turn against me or something.

David: Yeah, I remember that you told me that. I know that it was partly a joke, but behind it there was something that we should all think about. Guillermo, I appreciate that you joined me for this for this podcast. How do you feel about AI now?

Guillermo: Well, now I feel much better about the situation. Back when I was a kid, my imagination was at an all-time high and since there were all these movies about AI that I watched, I could imagine a world with Skynet taking over. But now that I know more about programming, I know that it's impossible that I create something like that by myself. So, please tell everybody that they should be responsible while developing AI, you know, just in case.

David: Thank you, Guillermo. I'll do that. That's exactly what I'm planning to do in this podcast. Thank you so much, Guilli, I'll see you at dinner.

Guillermo. Okay, you're welcome. Goodbye. Thank you.

David: Bye!

[MUSIC]

David: So, welcome to the AI Business School Podcast from Microsoft. I'm David Carmona and I have three kids, one of them you just met, who are scared of their own father's job. This may seem funny, but actually it reveals something deeper that is happening in society. AI is advancing very quickly, and it's imperative that we can do it in a way that can earn society's trust. Without trust, AI won't have any meaningful impact. As business leaders, it is critical that we develop AI responsibility, so our customers, our employees, or our 10-year-old kids can trust it. We're going to have that conversation today. We'll start with Natasha Crampton, Microsoft's Chief Responsible AI Officer who leads the company's recently formed Office of Responsible AI, which specifically addresses these issues.

Natasha Crampton: I think it's critically important that we treat responsible AI like we have treated privacy and security and we see responsible AI as a core element of trust. We know that people don't use technology that they don't trust, and so making sure that we are baking in responsible AI considerations when we're building the technology, also when we're deploying the technology, is really just an essential part of unlocking the value of these promising new AI technologies.

David: The responsible AI journey starts by acknowledging the challenges of AI. Like any other technology breakthrough in the past, AI comes with associated challenges and risks. And in the case of AI, it comes with very unique challenges that make even more important that we understand them.

My colleague, Sarah Bird, who was Microsoft's leader of Responsible AI for Azure Machine Learning, can take us on a deeper dive.

Sarah Bird: Both the power of the technology and the way in which we've seen it fail has led to a push around this responsible AI space and really adopting techniques to mitigate these failures and detect these issues up front. But it also started a much bigger conversation about ethics and society and technology.

David: Our attitude toward this conversation should be one of curiosity and growth mindset. We have to ask the difficult questions, even if we don't have the solutions yet. Sometimes it's not about what technology can do, but what it should do, according to Natasha.

Natasha: It's really important to have very open lines of communication in the beginning. I think to approach the engagement with humility and to be candid about the fact that we increasingly know what the questions are, but we don't always know what the solutions are. And also just to recognize that different people will have different motivations for being involved in this work.

David: In Microsoft, we started having those conversations very early. And in 2018 in the book "The Future Computed", we published our set of principles for responsible AI. In those principles, we identify the challenges that we believe were associated with AI, and we also established our approach to those challenges as a company. That's the second step that you will need to accomplish in your organization. It's not only about acknowledging the challenges, is about taking a stand on how your company will address them. To do this, we launched a committee called Aether, which stands for AI Ethics and Effects on Engineering and Research, I know, that's a mouthful. We consider the Aether Committee like a think tank with many different backgrounds and perspectives on it, and we discuss there the issues surrounding responsibility. Sarah can explain it further.

Sarah: So, Aether is an organization inside of Microsoft of many different types of experts coming together to work on responsible AI issues. And so, the different working groups do everything from recommend, for example, a new interpretability algorithm, help people understand when they should use that algorithm, to sort of giving overall guidance for the company.

David: This work resulted in six principles about our shared responsibility in AI in Microsoft—fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. Aether is still functioning because responsible AI is a journey in

continuous evolution. As the technology evolves, we have to make sure that we adapt to new challenges. For example, since we established those principles in Microsoft, we've had to address new obstacles around facial recognition or deep fakes, just to name a few.

Even if Aether's work is still in motion, we hope that sharing our learnings can be useful for other companies that are starting similar efforts. And we've seen many companies in every industry experiencing the same growing pains.

[MUSIC]

Nick McQuire: This year we've seen the importance of interpretability, security into the system, and privacy, in particular, becoming more significant in terms of their priorities for investment. So, how do I ensure that I have systems that I'm building that are interpretable? How do I ensure that they're secure and robust for my business. That shift is fairly new, particularly last 12 months, and it indicates that companies are starting to pay attention much more to some of these areas of responsible AI than we've seen before.

David: This is Nick McQuire, Lead Analyst for CCS Insight. What he saw as main blockers just a couple of years ago, like availability of data or staff skilling, are now very different. Organizations are acknowledging the risks associated with AI, and they're shifting their efforts into understanding how an AI model can have unintended consequences.

Nick: It's also, I think, becoming clearer for companies is that if they want to move into that kind of production phase and they also want to really look at the technology evolving to becoming very mission critical for them, and as they kind of embed it into their operational processes, they're starting to realize much more quickly around kind of what are the kind of real world risks and responsibilities that they're going to have when they actually start depending on machine learning processes and technologies as part of their business.

[MUSIC]

David: As we establish our definition of responsible AI here at Microsoft, we also realized that AI needs to be both fair and inclusive, and they both go hand in hand. An AI model that excludes a group of people due to their skin color, gender, or postal code is of course unfair, and it could even run afoul of the law. Those are very obvious cases, but in many others,

the bias is more subtle and even more dangerous because of that. What if the bias already exists in the real world and AI learns that bias through the training data? The algorithm may seem correct, but in fact it can perpetuate an unfair bias in society. That is not acceptable. Identifying that bias and addressing it is critical. The people who develop the models need to follow their organization's principles and put checks in place for accountability. We should not point to an algorithm as the ultimate arbiter of responsibility on any process. This is a human's job. That means people need to be included in the loop and they must understand the issues of inclusion and fairness. A team with diverse backgrounds always helps.

Sarah: So in a lot of cases, hopefully you can just go to the scenario-specific industry best practice. The other option is of course to do user testing with those groups, reach out, do user research specifically, kind of targeting those groups. There's also organizations that of course represent different communities, nonprofits that represent different committees that are growing AI experts to help give guidance on these issues. So, there really isn't one answer, but there's a lot of options.

David: Sarah's insights are very important, but in many ways what she's just talked about are external considerations. Building diverse teams from the start to tackle these issues goes a long way toward developing principles, governance, and tools.

Sarah: We do have quite a diverse team where everyone is coming for their own reason and with their own perspective. And honestly, I think that's been amazing in helping us develop these tools because there's more communities, more mindsets represented.

David: This is so important. People within an organization should be building responsible AI because they want to. Responsible AI requires a culture transformation. We want everybody to be part of it. When you undergo that transformation, you may encounter some good surprises. For example, as Sarah can tell you, what our engineers might have thought of as a tax, like dragging the development process, was actually received very positively.

Sarah: One of the things that's been really great to experience in responsible AI over the years is originally people thought that this was all tax, that you would just build your model and then you would have to do these responsible AI kind of compliancy things and it would slow you down and maybe the model would be a little bit worse and you put it out there.

And actually, what we're seeing in practice is a lot of these techniques help people find errors in their models sooner. They help the model work better for more people.

[MUSIC]

David: So, how do we foster this culture in our organization? How do we make sure that people developing AI are following our company's principles? While Aether falls into more of the theoretical, Natasha's home base, the Office of Responsible AI or ORA takes on much of the heavy lifting of putting the thought leadership into practice within Microsoft. One of the functions of ORA is to establish a governance model, so we all know the principles and how to stay within those parameters.

Natasha: Governance frameworks for responsible AI need to acknowledge that there's no steady state in terms of the rules that need to be followed right now. It's still a very nascent area. It's evolving quickly. And I think it's really important to develop a governance system that acknowledges that. So, it is critically important to be clear about roles and responsibilities. I think putting in place really strong infrastructure for support so that people know where they can go to get helpful advice when they need it is really critical.

David: One of Nick's clients is taking this very seriously and is an important example of how even legacy companies can realize that a new approach will open them up to new markets.

Nick: I was speaking to one large, global bank fairly recently. They've just recently put in place a process that means that no machine learning model that they build can move into production without some inherent interpretability or explainability into the model, and some checks around bias in the data to ensure that they are having a representative approach and diverse approach in terms of the data that's being fed into their machine learning strategies.

[MUSIC]

David: One feature of our governance framework is that we don't put one single person or one organization in charge of responsible AI.

Natasha: It's actually a shared responsibility across a number of different teams. And so, making sure that people understand the specific role that they play and how that accrues to the overall effort is a really important way of setting up that basic framework for governance.

David: An important part of governance, of course, is to build in metrics or KPI's, so you can measure how you're reaching your goals. Collaboration, as Sarah will tell you, is key.

Sarah: If you're picking the fairness metric for an application, you of course need the data scientists in there to understand how to actually implement that metric, if that's even feasible, but you also need the business decision makers, you might need the legal experts in this. You might also need people representing different communities to help understand what does fairness mean in this context to these different communities?

David: How do you scale this governance across your organization? The Office of Responsible AI addressed this issue with another important concept that they implemented, the Responsible AI Champs Program. These are, in essence, the people who take the guidelines and practices and scale them across the company to help create not just awareness of responsible AI, but actually implement it.

Natasha: We have our champs acting in a fairly hands-on role, but we want them to be able to engage across the organization with the people doing the work on the day-to-day basis but also what their leadership. And a huge part of what we're asking them to do at the moment is to generate awareness about the practice of responsible AI, how to get help, where the resources are. But we also want to partner with them to advance our cultural approach here, which is one of responsible innovation and driving that sort of culture change and reinforcing that really does require leadership.

David: Part of the responsibilities of these AI champs is to identify sensitive use cases for the Office of Responsible AI. To make that happen, they follow a process called impact assessment. Natasha can tell us more about this process.

Natasha: One of the specific things that we've been working on quite closely with our pilot teams is the process of undertaking an impact assessment. So, understanding how the technology that you're building might impact people and society. There's no tool that helps

you with this part of the process. This is really just thought work, it's really about reframing the way that you think about the technology that you're building. And for me, it's just really heartening to see that more holistic perspective coming out.

[MUSIC]

David: When setting up your responsible AI strategy, consider an Aether like structure with different working groups. Some of them will be more technical to help, for example, develop algorithms. But others should be more interdisciplinary to define use cases, predict unintended consequences, and discuss situations that in real life don't always have a clear answer. Then create your own Office of Responsible AI to take those learnings back to the organization and develop practices and guidelines to help those development teams; infuse responsible AI into your existing processes in your organization and identify best practices that you can replicate broadly. Once you are there, I have very good news for you because technology can help. In Microsoft, we are developing tools specifically targeted at developing AI responsibly. They are focused on three main areas: understand, protect, and control. Understand includes tools that allow you to open the black box of AI models, which is critical for transparency. It also includes tools like FairLearn, which help you assess the bias in your algorithms and address that bias. Protect includes technologies to keep data private. There are many tools available here, including encryption technologies like homomorphic encryption that allows you to run AI on top of encrypted data or technologies to keep data anonymous like differential privacy. Control includes platforms like Azure Machine Learning which helps you implement governance and auditing processes in your organization. Sarah explained to me one of the tools in the understand category.

Sarah: We have a new toolkit called Interpret ML that enables people to use many of the different best interpretability techniques, either to explain their model behavior or to find errors, or to explain a particular decision to an end user. And so, what we do there is both design the toolkit to help people like actually understand what they're trying to achieve and then help them use that, the right technique based on their goal. Like if you're trying to find errors in your model, you may want a very different interpretability tool than if you're trying to explain the decision to an end user, for example.

David: The good news is that Azure has already these open source tools built in. Sarah has found that using them can be incredibly helpful. They speed up your timeline, since you don't need to reinvent the wheel all the time.

Sarah: In general, I recommend for the core algorithmic capabilities, the best technology available is in the open source. And in fact, we are building all of our responsible AI toolkits, such as Interpret ML and FairLearn in the open source. That way, the experts in the community can also contribute to them, people can verify them, you know that it is actually implementing the algorithm correctly and the best practice correctly. And it also just helps us move faster. This is a very new field and the state of the art is moving very quickly and open source is able to move faster than anything.

David: And Nick sees these in action as well with his clients.

Nick: What we're starting to see is kind of best practices emerge from some of the more mature organizations with respect to machine learning that are heavily pivoting to responsible approaches that fundamentally help them increase agility and bypass challenges that will slow them down, down the road.

[MUSIC]

David: As we finish this episode, I want to leave you with three recommendations for developing AI responsibly. And remember what my son, Guillermo, said. He's counting on you to follow these recommendations. So no, no pressure here. Let's go for it.

First, define your principles of responsible AI within your own organization. That will guide you to know the problems you are trying to solve, but also to establish a position in your company on how to address them.

Second, put these principles into practice through governance. Establish a process across your organization that infuses responsible AI into every step of your development cycle.

Third, understand the technology and leverage tools to help you in your journey of developing AI responsibly. AI is going to be full of shiny, new technologies to accomplish amazing things. Always assess those shiny technologies with a thoughtful analysis of the

unintended consequences. As we said before, just because technology can do it doesn't mean that it should.

And finally, remember that responsible AI is a continuous process. Be humble and acknowledge that you don't have all the answers. Learn from others, connect with external organizations, and bring diverse perspectives to the table. Responsible AI requires a growth mindset.

I'd like to thank our guests for this episode, Sarah Bird, Natasha Crampton, and Nick McQuire. You can visit Microsoft AI Business School at aka.ms/aibs. And you actually have a full track in there for responsible AI. You can also learn more about responsible AI at Microsoft and even to see our tools and guidance by visiting our Responsible AI Center in microsoft.com/AI.

I want to finish with one thought from Sarah Bird, which I love. Sarah defines responsible AI simply as a way to acknowledge who we build technology for. I leave you with that thought.

Sarah: And people are starting to actually more deeply remember and understand that like there are humans on the other side of the technology we're building. And so, I think we're just seeing technology that's designed for humans first, like really taking people into account from the beginning, happening just so much more from the beginning, which is like a huge change in how computer science used to go.

David: Thanks for listening.

[MUSIC]