



ON THE PLATFORM

BUILDING DATA AND AI ETHICS COMMITTEES

Host: Steven Tiell, Senior Principal Accenture Labs

Guest: Ron Sandler, Professor of Philosophy,
Chair of the Department of Philosophy and Religion,
Director of the Ethics Institute, Northeastern University

Steven: Hello, and welcome to **On the Platform**, where we're talking to the most influential and innovative thinkers in platform technology on the hottest topics and trends. My name is Steven Tiell. I'm a senior principal in Accenture Labs, which strives to be the tip of the spear for innovation at Accenture. My team focuses on data ethics, what we currently call responsible innovation.

On the Platform is a podcast series produced by Accenture's communications, media, and technology group, Software and Platforms Practice. Today we are going to explore committee-based governance for ethics and AI. I'm joined by my friend and collaborator Ron Sandler, professor of philosophy, chair of the department of philosophy and religion, and

director of the ethics institute at Northeastern University. Ron, welcome to the program.

Ron: Thanks Steven, it's good to be here.

Steven: Over the past year, we've been collaborating on some research, and we recently published a paper that talks about how to establish committee-based governance for ethics and AI. We're going to talk about that now, and about the issues that are confronting companies and governments as it relates to ethics as technology roles out artificial intelligence. Ron, what do you see through your lens of what the need is for data and AI ethics?

Ron: From our perspective, as you well know, it's almost not a day goes

by without some news about some ethical issue or ethical problem associated with machine learning or machine bias or privacy or security, or transparency, or explainability or misuse of data, or something. And when you think of these as one off issues, they're kind of interesting problems that need to be addressed. But when you look at them all together, when you see this many issues come up with this regularity, it's an indicator to us that there's a broader capacity issue here. There's a broader societal problem. And that broader problem is that we lack the requisite resources, whether it's institutional structures or regulatory structures or expertise or background knowledge to be able to identify and prevent these things from coming up in the first place. And so as you know, quite a lot of our work is to do with thinking about how can we build societal capacity to anticipate and therefore prevent ethical issues coming up with big data, machine learning, AI and such things?

Steven: And we're seeing these problems not just pervasive in businesses, but pervasive in governments as well. It's the case that governments are trying to get out in front of these technologies, but what it results in is mostly that they're reactive. And maybe reactive not in a quick manner either. And so what we're seeing is that a lot of companies are struggling with this

notion of here's how far and fast the innovation has gone, but there's a number of questions that start to arise about how does this technology interface with society and they look to governments for answers but a lot of times governments don't have those answers.

Ron: Yeah I think that's right. I think because the technologies move so fast, because the products and the services are being introduced using machine learning or using facial recognition or developing autonomous decision making systems and these sorts of things; because it's moving so much quicker than regulatory or oversight processes, it's going to be the case that organizations and institutions are going to have to build capacity to address the ethical issues that go significantly beyond what the regulations are going to provide.

Steven: One of the things I really liked about what you and John Basl brought to the program was this notion of looking at, standing on the shoulders of giants. Let's look at other fields and how they've addressed this. And that's something that all of your research was really built on top of right?

Ron: That's right. One of the other things that we like to emphasize is that we can actually learn a lot about how to build capacity in AI and data ethics from previous technologies

that have raised issues in the past. And in this case we thought that bioethics was a pretty good analog because it was a case where organizations were facing a lot of issues and questions and there weren't great models for how to deal with them and so models were developed. In that case, things like IRB boards or health ethics committees or hospital committees to help people make decisions or support physicians and researchers. So we thought that could be a good model for AI and data ethics committees within organizations. So we looked at that and tried to adapt them to what we see as – it's a different technology and it's a different environment, but there's still a lot to be learned about what does a committee need in order to work well. And that was the primary question that we were asking.

Steven: One of the things I'm hearing from clients is that they want something on a spectrum of governance. In some situations, they say oh we're interested in essential governance and what I parse that to mean is we don't really want to delay cycles very much. We don't want to go in front of a board or a committee that might turn us down after we've spent some time working on something. We just really want to have these things infect our practices naturally in an organic way and something that our people just

do by default. Is that something that you think is realistic?

Ron: Absolutely but you have to build capacity within organizations to be able to do that. One dis-analogy between the kind of committee-based structures that we were looking at for AI and data ethics and something like IRB boards is that they don't need to be gatekeepers. They can be consultatory. They can be voluntary. They can be resources that people can go to when issues arise for help. But what's crucial about them and what we really think is important about them is that they can help to set policies within organizations and they can help also to train people so that there will be more people in the organization that will be able to deal with these issues in the way you've described.

Steven: That's a great point. What I hear is that there's this spectrum of governance that we can look at. And it's everything from maybe nuclear essential governance, all the way up to something that's much more bureaucratic, stakeholder rich, counsel based. And we can do some kind of spectrum therein. We can take flavors of each that we like and move forward. And is that something that you think is a viable way to go?

Ron: I don't know how else to go. There's ethical issues and questions that come up everywhere from the design and development of

technologies and services, up through implementation, up through roll out, and there's a lot of questions about the context in which they're rolled out. They're really significantly complex. They can be significantly complex problems. So there's no way to deal with these problems without having people and structures in place to be able to address them. Now some things are going to be not that big an issue. That's just a consult, just trying to figure out, is there even an issue here? But other things are going to be more sensitive. And if you're talking about deploying facial recognition technology and certain kinds of services for example, that might be something you want to hold up to stricter scrutiny. You need some kind of mechanism to figure out within an organization, what are the things that need stricter scrutiny, and what are the things that are not so much? And the way we do that is by building up cases. And the way you build up cases is by having people bring forward questions and seeing what the best ways to respond to them are.

Steven: And so as we move forward with building this institutional capability, one of the things that we've talked about in the past, and I'm talking with other universities about as well is this pipeline of talent. And how do we have people inside of organizations who all of a sudden have this whole new world of cares,

this new field of responsible innovation, where we need to have this cross pollination between disciplines. It's not enough to be a PhD in computer science and be very good at machine learning or neuromorphic computing or whatever the latest and greatest technology is. It's really a need that we have this humanities muscle as well. And how do we build those pipelines for candidates so that when we get into these situations we have these multi-disciplinary groups that can participate in that. One of the things I guess we look at is there's not an abundance of good examples of governance right now in ethics and AI. Do you think that that will change over time? Does it require those pipeline of students in order to really affect that change at scale?

Ron: Yeah, I think those are both really good points. On the first point, it's – when we think about capacity and societal capacity or organizational capacity, it's important to have good structures and it's important to have good regulations so far as is possible, but it also is going to require people. And the sort of people that if you imagine what an AI and data ethics field or profession would be like, it would be enormously interdisciplinary. Not just computer scientists, not just philosophers. It'll have social scientists, it will have people in law, it will have people in business, it will have people who are health

professionals, it will have data management practitioners. It's going to be a really diverse field because these are complex problems. And you're going to need lots of different skills and perspectives to deal with them. Which is another reason actually we like the committee-based approach, because no one person is going to have all the right perspectives. You have to bring together people of diverse perspectives and diverse skills to handle some of these issues.

And I think the other point that you made about sharing best practices is really crucial. So it's very hard if we're going to develop a sort of AI data ethics capacity and norms and institutions and structures, if organizations aren't sharing information on what works. If everybody is trying to do it themselves, it's going to be a lot harder than if everyone is learning from what other people are doing in a process that's going to develop some standards, some best practices, some shared knowledge base.

Steven: And as you know we've started a data ethics salon series to really address that need for practitioners in the space to get together and learn from each other. I think one of the telling things that came out of that experience is that we're actually in the process of spinning that out to a nonprofit

because we really do want to be able to publish best practices and have that set a bar by which others can run their programs and design their programs. One of the things that we started hearing about a year ago is that all of these people all of a sudden at different organizations have been beating on this drum for a number of years, and then they were given the opportunity to start a program. And then they said oh my goodness, where do I start? And so part of that building of that road map is really important and how we share those lessons learned is very important. Ultimately I think we can get to a place where we have an established set of best practices for this emerging field. And I think this is something where the committee-based governance paper is an absolute bedrock as part of that foundation.

Ron: As you know we put this paper out not as the final word on how to go about building effective AI and data ethics committees within organizations, but as a place to start. And our hope is that people will take it up, try it, add new things, adapt it to whatever the context is that they are involved in, and then importantly share out what works and what doesn't work. We're trying to develop similar communities through Northeastern's network of campuses. And we have campuses all across the United States and Canada. And one of our initiatives is called Cascadia AI

which is being run out of our Vancouver campus in collaboration with our Seattle campus. And it's aiming to not only encourage organizations to adopt ethics principles around AI and data ethics, and not only implement them – that is come up with an implementation strategy – but also to evaluate them and then share back to the other people who have signed on to what we are calling the AI Cascadia Commitment, AI Ethics Cascadia Commitment – what works and what doesn't work. So it's an intentional building of community to share knowledge and develop best practices.

Steven: I know there's something we went back and forth on quite a bit when we were writing this paper, and it's this notion that we really need that feedback loop. I look to one instance in particular where we were talking with a client about this process and about that feedback loop. And they had some pressure – they said no we want to keep this stuff from our executive team. It's better that they don't know about some of these things. And our legal counsel is telling us we shouldn't document these things in email either. And my response was really that, no this is the kind of things that they need to understand. Having that institutional knowledge building is profoundly important. And so how do you expect to evolve if you're not looking at that feedback loop of what

works and what doesn't? And I think that that's really where this is headed, is that one of the things that committee-based governance can do is that self-reflection. Organizational level self-reflection. We know that that's a telling trait of strong leaders, is their ability to self-reflect. But the same is true of strong organizations. And I think that this plays a critical role in that process.

Ron: Yeah and you and I have talked about this many times, but one nice feature of the development of a robust AI and data ethics field and societal capacities is that it's the right thing to do. It's good for people, it's good for democracy. It's good for justice. It's good for autonomy. But it's also good for businesses and institutions to do as well because it can do things like give competitive advantages or help manage risk. Or create opportunities.

Steven: If you had told me a number of years ago when I was starting down this path that I would be talking with clients about culture change, I would have thought you were nuts. But very much many of these conversations end with or lead to this notion that organizations really need to double down on what their values look like because those values are going to help architect what their employees care about and what they do and how can you articulate how your values show up in your products and services that you deliver and

how do you use your values to design principles that act as the guard rails for what you ought not do? I think that I just keep coming back to those things time and time again because ultimately we're talking about how are we granting people improved agency in their lives? And if we're not reflecting on oh, are we really giving them more agency or are we making decisions on their behalf? And how does this fit with our values as a company? I think we're in a real quagmire if we don't have a handle on those things.

Ron: I think that's a really good point. And it's a really good point that it's one thing to make the commitments, so to adopt the principles or adopt the value statement. But it's a whole other thing to implement it. To operationalize it within organizations. And that takes a non-trivial amount of commitment and resources, but again it's not only the right thing to do but it's also good for recruiting talent and keeping people and bringing in the best people who care about these things.

Steven: That's another great point. We've heard from a number of clients that the reason that they're reaching out is because the folks that they've hired in the last couple years are demanding this. They don't want to work for an organization that doesn't attend to the ethics of the products and services they put out into the world.

Ron: So our approach, and I think this is an important point – isn't to say, this is exactly what an AI and ethics data committee should look like. That's impossible because what it should look like depends on the organization. It depends on the mission and the value statement. It depends on whether it's a private or a public organization. Depends on the kinds of – it depends on so much. So what we do in the report is we try to identify the key questions that an organization would need to ask and answer to build the committee that would function right for them. So that's the goal of this report is to help people think through the questions that they would need to think through in order to build a committee that will serv their purposes, and it won't be the same for everyone. As we talked about before, some committees might have greater power to prevent certain things from happening. Some might just be consultancies. Some might be involved at the research level. Some might only be involved in product. There's so many different ways you can constitute and empower the committees. And knowing what questions to ask is crucial to developing a committee that will actually work to accomplish its goals.

Steven: Great points. And unfortunately that's all the time that we have for today. Ron, thank you so much for taking part in today's conversation.

Ron: Thanks for having me.

Steven: We hope you have all enjoyed this episode. Please help us get the word out and be sure to subscribe, share, rate, and review our series. We would love to hear from you and hope that you tune in again for the next episode of ***On the Platform.***