

Caitlin McDonald:

On today's podcast, I'm joined by Dr. Rumman Chowdhury, global lead for responsible AI at Accenture Applied Intelligence. Rumman has been named by Forbes as one of five who are shaping AI, one of BBC's 100 women for 2017, and she is a member of the Royal Society of the Arts. She co-chairs the RSA's Citizen AI Journey Jury, and she is on the board of AI4ALL. Welcome Rumman.

Rumman Chowdhury:

Thank you. Thank you for having me.

Caitlin McDonald:

We're delighted to have you. I wanted to start the conversation today around the fairness tool at Accenture. So one of your projects there has been an algorithmic tool to identify and mitigate bias in AI systems. Tell us a little bit about what fairness means in an algorithmic context and how the tool works.

Rumman Chowdhury:

Yeah, so the fairness tool was what put at least my practice on the map, and it was the first time anyone had really thought about creating an enterprise tool that works with your tech stack to detect and mitigate bias and ensure fairness. But looking back, it was almost three years ago at this point, it seems almost naive because of the purely algorithmic approach. And it's not to say that... I'm a quantitative social [inaudible 00:01:35], so it's not to say that we didn't think about the socio-technical implications, it was just like, "Oh wow, there's this tool and we can use it to to add to the model development process." But since then, I think all of our thinking has gotten a lot more sophisticated and we're really trying to incorporate at least at the enterprise level more of what it means to really truly understand the context in which something is used.

So what does fairness mean in an algorithmic context? The answer is many things, and that's kind of problem. This is why creating a simple quantitative only tool or checklist doesn't really solve the entire problem because it nudges people towards purely optimizing for some metric or checking off all the boxes on the list, when really what is asked for is critical thinking. There are so many ways to measure fairness, to mitigate bias, and really having the expertise to choose the correct one for your situation is where the experts come in.

Caitlin McDonald:

Yeah, thank you. I think that's such a great point, that the contextualization is really critical. Also increasingly at the LEF, one of the things that we're thinking about as researchers is there's a drive towards metricizing everything or using data to really be the arbiter of all outcomes. Every organization wants to be data driven. When in reality, what we think they really want is to be data informed and to accept that there are some contextual flavors outside of the purely quantitative that needs to be taken into account. So I'd be really interested to hear what evolved in your thinking as you move towards this more contextual approach.

Rumman Chowdhury:

Yeah. I really love how you're you mentioned data informed versus data driven. I agree, there's this data primacy that comes as notion of being data driven. It's like you're not making the decision, the data is, and there's so many problems to unpack. Number one, even just the concept of this quote unquote objective truth in data, which does not exist, but also the human agency component of things and the

techno-centrism or the techno-chauvinism as Meredith Broussard would say, in the development of these tools. So driving context's actually quite complicated, and there's not really an easy answer to it. It doesn't mean host a design led thinking session, which is probably the way people get around it. And by the way, design led thinking sessions are great ways to brainstorm, but then the hard part becomes driving it into practice.

What I will say on the positive side is I'm seeing more and more companies interested in figuring out how to do this. And again, it's not a solved problem, but I will say that, for example, in Coursera, there's now an a certified ethics assessor training program that's actually meant for qualitative people, and there's also quantitative assessments out there. Part of creating this is about good governance, and we do have organizational structures within companies that already do these kinds of things. So risk and compliance in any sort of regulated industry is a great place to get started.

Every social media company has a trust and safety division, and even if you're none of the above, you have a legal arm and maybe it's not developed in that particular way, but it can be. So part of it's also looking to existing infrastructure and how it can be leveraged. And again, to your distinction between data informed versus data driven, data driven even structurally at an organization puts at the top the people who are able to code, to build technology, versus saying it's an ecosystem of people all involved in creating technology. Why would you give legal and compliance quote unquote respect to influence your product if you are saying we are data driven and data is at the top of everything.

Caitlin McDonald:

Yeah, and I think some of that comes back to, I'm also speaking as a social scientist on the qualitative side, recognizing that there are different kinds of data. So quantitative data is one kind and it is very valuable, but you also have other kinds of data that need to be taken into account as well when you're creating these kinds of structures and systems. To take us into the governance piece for a second, you mentioned some governance models that are specific around regulated industries, and then a couple of other options for groups that are not involved in regulated industries. And I'm wondering, there's also the idea that if you're a small place, you might not have a huge amount of governance going on, but also I think governance is changing over time. There's a shift in terms of how governance is being done. And I wonder what you might have been observing in that space.

Rumman Chowdhury:

Yeah, I actually just had a really interesting conversation about that at this conference called House of Beautiful Business, and we did a whole panel on governance. But my training is in political science, so I think it's very interesting to use this word governance and not actually tap into a lot of the research around what makes good governance. And by the way, I will first add that governance means different things to different people. If you talk to somebody who builds technology, they think of governance at the model level, so what is a system governance for a model, [inaudible 00:06:50] checks and balances at each stage of development, deployment, and then post deployment. But if you talk to most other people in the organization, they think of governance as organizational governance.

And then of course, there's this other layer of extra-governmental... I'm sorry, extra-organizational governance. So these might be influential bodies or groups, for example, LEF would be one. Any research organizations, civil society groups, which actually have had an impact on how organizations are talking and thinking about ethics. And then we have literal government. So if we think about these processes that we're building there, any existing institutions or new ones being formed, but the one thing I'm seeing that's pretty common across all of it, which is kind of troublesome again, as a scholar of politics, is that none of them are actually built democratically. I jokingly [inaudible 00:07:42]

citizens panel as well, that we would actually never want to be a citizen of a country that is governed the way that we're creating AI governance at international level.

Because what happens is we invite the quote unquote smartest minds in a room, and I am sometimes one of those minds, and we all come up with these thoughts and we're like, "Don't worry, trust us. We're the experts." And then we go and we tell everybody what good AI is. And this is not to say... Certainly there are people who are experts, who do this stuff for a living, but I worry if our voices are the only voices being heard and the only voices in the room, again, I don't think most of us want to live in a country that is constructed, the governance model is constructed the way we've created AI quote unquote governance. And then of course it's unsurprising when you see these meta reviews of fairness principles and policies and they say, "Oh, look, there are these common themes," and I kind of want to [inaudible 00:08:37] because the same 20 of us were all [crosstalk 00:08:42].

Caitlin McDonald:

Yeah. That's a really great point, and it actually takes me back to some work I did even before I joined the LEF where I was looking at the idea of AI citizens and data scientists, and thinking about if you work in... If you're normal citizen in your everyday life, you have these representatives who create the law for you, whereas in your algorithmic mediated life, which you can't see, because I'm quoting my hands, podcast listeners, which of course we all now are subject to. Our lives are increasingly digitally involved and automated and there are lots of digital decisions that happen behind the scenes. So we're citizens in that world of this digitally mediated life and essentially our representatives in that case are data scientists and policymakers and those 20 experts who are in the room who were making these kinds of governance processes.

And my argument was if you are a politician, you have practices, codes of ethics, and clear expectations in terms of what kinds of things you have to do to make sure that you're behaving the right way towards the citizens who elected you. Whereas if you are a hired data scientist or an appointed AI ethics board member, those responsibilities are often a lot less clear. And certainly to your point, the influence and the accountability that essentially the citizens, the consumers of these systems have is a lot less clear in terms of the mechanisms that we have for the ins and outs of that. Of course I'm not saying that we should move to a world where you vote your data scientists in and out of teams, but it would be really interesting to re-imagine the kinds of mechanisms that we do have, or could have in terms of governance in that sense.

Rumman Chowdhury:

Yeah, absolutely. And one thing I love about what you're saying is, extending the analogy to politicians, it's also quite clear who is a politician's constituency. So if I live in a very liberal part of New York versus in a very conservative part of Texas, my representative represents me, or at least it's supposed to represent me. And agreed, it is often unclear who people are representing when they go to create these governance systems. So are they representing the company they work for, or the body they work for? Is this some sort of a personal agenda that they're driving? And if so, does one person's individual agenda around AI carry more weight than any other citizens? So it can be a problematic structure.

Now I really want to write a paper on lessons learned from... Basically as a political scientist, you have to take intro to democracy, and it's actually about what makes democracies long-lasting and representative et cetera. And just today, again, on this panel, we were talking about constitutions and there's an entire... One can be literally a constitutional scholar and just do research into what makes a good constitution. When we think about these fairness and ethics principles, and there are so many, we're drowning in them at this point, it's hundreds, [inaudible 00:11:32], are there any, or are we

building them in a way that's actually in alignment with this literature on what has made good, long lasting stable constitutions? I think that'd be frankly fascinating.

Caitlin McDonald:

Please write it. We'd love to read it.

Rumman Chowdhury:

I need another political scientists to write this one. I feel like I can't call myself one anymore. It's just been a while.

Caitlin McDonald:

To continue extending the analogy, you also then have this interesting situation where yes, you have extremely specialized constitutional scholars, but you also have the ordinary citizen on the street who's like, "How do I understand how this is impacting my life, and I have 15 mcgillion things to worry about, and here's this referendum on the ballot and I have no idea what this means." Not to name any specific referendum that I might be referring to, but there are many that we could be thinking about at this present time. So even just down to the local school board election or whatever it might be, as an ordinary citizen, trying to be an informed one or understand decisions that are being made about your life or contribute to those, it's very difficult, and the same is true, I think, with any automated decision making system as well.

Rumman Chowdhury:

Yes. I agree. Especially with automated decision making systems. So I do have research I'm working on, something I call the retrofit human, and it's building on the techno-centrism in Silicon Valley or just in tech in general. And again, it's interesting because we use this term Silicon Valley and really we just mean the culture of tech, which is what originates out of Silicon Valley, but really is embodied just no matter where you go in the world. It's sort of the tech culture mantra. And the mantra being that humanity is flawed and technology will save us. So as a result, it is actually incongruent to say, we're going to create human centric AI or human in the loop AI if fundamentally what you believe is we're the problem and technology's the solution. Why would you put the problem at the center?

And this is the part that doesn't match when people use those words. It's not part of the intrinsic design of the system. So retrofit human happens when we want to reshape society or human beings to fit the limitations of technology. And as an example, I use the Uber self-driving car litigation, but also self-driving cars in general. And how they were pitched to us as this magical pod that would whisk you away to wherever destination, but has turned out to actually be literally the worst of driving and all of the liability. So taking into account Madeleine Elish's a moral crumple zone concept. So in Uber example, the self-driving car hit a woman in the street and the liability has actually fallen on the driver, not the company making the quote unquote self-driving car, I'm doing air quotes as well, but the person who was sitting behind the wheel.

And realistically we think about human beings and how they drive, but also what is the value we would extract out of a self-driving car, other than, I guess, bragging rights? Theoretically it should be to free up time, to reduce stress from travel. The stress from a commute is a very, very real and measured thing. But it doesn't actually. [inaudible 00:14:48] we'd have to be sitting at the wheel at attention the entire time, so what is the benefit to humanity in building this thing?

Caitlin McDonald:

Yeah. And to your point, I think it's the worst of both worlds in that you have to be paying attention, but for 90% of the time, you're not actually doing anything, and it's only in the worst possible situation that you as the human driver are asked to intervene, but because you're out of practice, it's much more difficult.

Rumman Chowdhury:

Spot on. And now that now the talk is about redesigning entire cities to be built around self-driving cars. The state of Michigan has said that they want to start investing in Detroit. Again, the incongruence of what we know to be healthy for human beings in urban design and planning, especially in a COVID era. So Paris wants to do this 15 minutes city where everything you need's in 15 minutes. That is not actually conducive to a quote self-driving car city. So what is it that we're optimizing for? Human health and wellbeing, or having this technology widely available because capitalism.

Caitlin McDonald:

Yeah, because of a vision of the future that perhaps not everyone shares equally in.

Rumman Chowdhury:

And that's where I talk about democratic processes and who's in the room and who's deciding what a good vision of the future even is.

Caitlin McDonald:

Those are really great questions. I want to take us as well to do a little bit of a deeper dive in terms of when we think about the way that these systems of governance are evolving in organizations. If you're a company that has accepted the idea that digital ethics is important and you're concerned about these things and you don't want to be the next Uber with the headlines, staying out of the news, where is it that you should start as an organization, and what are some of the way stations as your maturity model evolves? Because you mentioned already, for example, that the fairness tool you think is earlier, and now you're thinking about it as evolved and you're offering essentially different advice, and that's just in the space of three years. So if I'm trying to build a program in this, how do I get started? Where do I go?

Rumman Chowdhury:

So what we realized, the big gap when we introduced the fairness tool was actually that we didn't have the right organizational structures and processes. So all of this sounds very [inaudible 00:16:59] who's a technologist, but there isn't a fancy platform you can buy or single GitHub repo you can use or a cool tool you can buy from somebody to solve this for you. And actually I have a paper coming up. Well, it's an article in Sloan Management Review coming out. Oh, actually I think in two days. Two or three days. It's actually about organizational change dynamics. So it's myself, my data scientist at Accenture, Bogdana Rakova, Henriette Cramer from Spotify Labs, and Jingying Yang from Partnership on AI. And we got together and we asked ourselves a question. The people who are tasked with applying this in companies, what are the problems they're seeing and how do they think organizations can really leverage the change they need to integrate this concept of responsible use of technology into their DNA?

And we don't look at the history of fairness and technology. We actually look at concepts of organizational change management. And specifically, it's not in this paper, but I did a talk last year in Data & Society, where I looked at theories of change in organizations and where we've taken moral and social imperatives and tried to combine them with capitalistic incentives. So how does it work for

climate change or organics or sustainability, and how successful or not successful were they relatively? And it's really fascinating to follow. But in this specific paper that is coming out in SMR, we actually came out with a few things that companies can do. It's going to be in the article.

But the two that I love highlighting one is this concept of diffuse accountability. So you may want to have a centralized lead, and I frankly think it's a good idea, and obviously I have a bias towards that. I do think it also just helps have a coherent message across an organization, to have a central distribution point. But that's the thing, is [inaudible 00:18:49], there's some level of accountability that lies with the different people. And this is not to say all lawyers need to learn how to program. You actually have to have the critical literacy to understand what's happening in the system and how it relates to your expertise, and how you can communicate your expertise to others. And the second one I want to mention, which is I think the most important, but again, my bias as a data scientist, it's how do you create success metrics for both models and people at your company to incentivize them to do the right thing?

And that's actually much harder to do than it seems, because, this is what the sustainability world has been struggling with with measuring impact metrics for the SDGs. So this is not an easy problem to solve, and I think organizations have struggled for a long time, with, if we take a step back, responsible technology is an innovation, this concept of innovation metrics versus performance metrics comes up in Eric Ries's Lean Startup. So truly building something that is innovative and different is by definition, in the current status quo, not going to optimize for status quo. So we have to respect that and figure out how to create ways of measuring and rewarding that innovative behavior.

Caitlin McDonald:

Yeah. I think those are some really great points. So just to unpack a few things. One, this podcast magically is coming out in the future, so we will have a link to the paper which will already be out. It will exist by that time. So we'll put that in the links that go with the podcast. And the Sustainable Development Goals by the UN are what Rumman is referring to there, and that point about incentives and competing incentives I think is a really good one. And one of the things I think that's really coming out of this podcast series is that interdisciplinary teams are a really good idea. And that sounds simple and it sounds it should be obvious to everyone, but it is organizationally and structurally, we are simply not optimized to work in that way typically, because often what you have is a concentration of expertise in one area, rather than in some sense, organizing around product teams or groups of people who bring their own unique expertise to the table. So are there some examples that you've seen of that working really successfully?

Rumman Chowdhury:

I think interestingly, where I might draw from for an analogy for that would be in cybersecurity and their use of things red teaming. And we're seeing some of that being used to understand the ethical implications of technology. So internally at Accenture, when I've built things akin to a red team, ours was comprised of data scientists, a few lawyers, an industrial organizational psychologist. And we had a project manager leading us all along. But even then it was a pretty tedious process. It was really difficult to do. And going back to this idea of reward metrics, we all had to do it side of desk, even myself. It wasn't actually... I did it because we needed to get it done and we wanted to do this work, but it wasn't actually built into our day to day, so it got pretty difficult to do.

The other thing about these things is it's a little difficult to scale. So you want to create an interdisciplinary team. Does that mean you just grab a lawyer, any lawyer? And what is the kind of literacy somebody needs to have to be able to contribute to this space? And I mentioned the phrase

critical literacy earlier, and this is where we get into critical literacy. There was a book that was recommended to me by Chris Wiggins who's the chief data scientist of the New York Times and professor at Columbia. And he loves to recommend, and I ordered it and I think it's great. It's called *Multiliteracies for a Digital Age*, written in 2005 by an English professor. And it actually outlines what sort of literacy do people need to have to understand technological systems? And he brings up different kinds of literacy and what you might need. But it's just a very, very helpful rubric when we even use this term, use this concept of literacy, but also in the sense of meaningfully contributing to a quote interdisciplinary group in assessing AI systems.

Caitlin McDonald:

Yeah, that's definitely going on my reading list immediately. And I think that idea of technical literacy is really important. And it's something that we're seeing the way that that plays out in a number of different fields. And aside from reading the book, are there any strategies you would recommend for people who are looking to develop that kind of literacy?

Rumman Chowdhury:

Yeah, so I keep alluding to the organizational structures that need to happen. We actually have an entire governance guidebook that we make available through Accenture. I'll outline the few steps that we recommend people to do. Number one is creating your North star, your ethics guides and fairness principles. And again, to our conversation around what makes a good constitution, the rule of thumb would be, it needs to be generalizable but actionable. So saying something like AI should do no harm is not really actionable. Harm to whom, et cetera. Second, I always like to talk about this concept of constructive dissent. You have to build channels of constructive dissent. Often when you follow whistleblower cases, you're alluding to the companies want to not want to be in the media, et cetera, and that completely makes sense.

And some of these whistle blowers are friends of mine, and they all have the same story. They say, "I really loved the place I worked at. I felt incredibly passionately about what we were building. I really thought we were doing a good thing. And then I found this thing and it all fell apart for me," not because they found a bad thing, but because nobody listened to them. When they brought it up to their superiors, they were either ignored or punished. They tried to bring it up the ladder. So having formal channels of constructive dissent, where in these people are actually protected if they want to raise ethical issues, I think is important. I think a lot of companies have introduced things like open forum, town halls.

I personally don't like things like that because it just ends up being the loudest voices being heard, not necessarily the most constructive, and people talk very, very differently when they're talking to a large audience than they would if they were actually trying to express a concern and trying to fix something. And also, I will say you'll find that the people really trying to fix something and raising problems are probably not the ones speaking up in a public forum environment. So those would be a few I would definitely recommend organizations do, besides what I mentioned earlier about this concept of diffuse accountability and creating metrics.

Caitlin McDonald:

And especially for those voices that genuinely want to create change in an organization. If you're an organization that wants to create those forums for creative dissents, you don't want to do a town hall because you recognize all the problems you raise. What are some alternative mechanisms that people can use or organizations can use to promote those?

Rumman Chowdhury:

Some places introduced this concept of a digital ombudsperson. So this neutral body in the organization that is there to answer questions, because sometimes, by the way, these ethical issues are sometimes just questions people have about how technology is going to be used and the intent, and that individual can answer that question. Similar in other parts of the organizations where you have whistleblower protection. So interestingly, which is why I draw a lot from legal, risk, and compliance folks. Sarbanes–Oxley introduced some really, really interesting concepts of ethics in organizations that was unprecedented. And some of that discussion is about this idea of whistleblower protection. So what are the rights you are afforded as somebody who is, in good faith, trying to help an organization.

And another good way of channels of the dissent, I mean, frankly are just about good community building, and I wish there were, like I said, an office I could make, or an easy metric I could say that would make that happen, and the culture just comes from the top. So one of the things that we look at in this paper are how close people are to the channels of leadership, how many levels between them. And it's quite interesting to map the relative success of ethics movements and organizations to how close they are to leadership and what is the message that leadership is giving throughout the company. So a lot of that really falls on leadership to strike the right tone.

Caitlin McDonald:

Yeah. That makes a lot of sense. And also, the other thing that occurred to me as you were speaking is that no fault idea around the whistleblower also really relates to some other successful examples of ethical mechanisms or structures being introduced in particular around retrospectives in teams that are looking at what went wrong without blame. Looking at medical ethics, you have a whole ethical review board there. And in particular flight crashes and review boards around that are always independently operated, because the idea is if we want to create less problems and fewer problems in the future, what we need to do is not figure out who's the scapegoat for this, but figure out what were the mechanisms that went wrong, what were the channels that went wrong, and then address those. And the only way that we are going to be able to do that is if everyone feels safe enough to speak out about where the problems were. And if people are trying to hide in order to deflect blame, then we'll never get to the bottom of what happened.

Rumman Chowdhury:

You're spot on. Red teaming is not intended to be a witch hunt. You're not trying to find who the bad actors are and weed them out if the people disagree with you. I actually see these things as a form of quality assurance. And the way I explained this to programmers and engineers and designers is you have the QA that you do around your code. By the way, 20 years ago, people didn't do a lot of this QA, because all of the stuff was not standardized, so this has been introduced over time. The simplest would be testing for scaling. That makes sense. And [inaudible 00:28:33] very, very quickly.

But newer ones are, for example, testing for security and testing for privacy, which we did not do for a very, very long time and then had disastrous results. And now companies are very concerned with notions of privacy and cybersecurity. But again, this is put under good quality assurance of your product or your model. The point is not pushing forth a moral agenda, it's to say you are just not building a good product if you are not looking at the ethical considerations and not just these optimization metrics around accuracy, et cetera.

Caitlin McDonald:

Yeah. And those are all really great points. And I think that idea of these mechanisms evolving over time is really related to some of the mechanisms that we talk about in the research that we've done around... Essentially when things are in genesis phase, those mechanisms don't exist, but then as products mature and projects mature and companies mature, those more structured approaches begins to take hold. So actually that's really useful for teams that are thinking about, "Okay, so how is this going to be governed in the future? You can actually learn to predict some of the mechanisms that might be coming down the line so that you can really build those things in advance, rather than waiting someone to tell you that you're not compliant.

Rumman Chowdhury:

Yeah, absolutely. But one way this is interestingly almost at odds, and one of the culture clashes that I see is... Well, maybe not. It's about the culture of data science and data scientists. Why has data science and AI completely dominated quote unquote analytics and all of this stuff in such a short period of time? It's because of the subculture of open source and this culture of openly sharing everything we do, and almost painful transparency. I say painful because I think to a lot of traditional organizations, this idea of building something and giving it away is just bizarre. Why [inaudible 00:30:22], right? But that's just not how our culture works.

Sometimes I think people have this knee-jerk reaction to any guidelines or regulations or standards, cause it's like, "No, no, no, no. This is a free for all environment. Everyone does it." And I think all of these things should be part of our culture. I mean, data science is in adolescence. We're still very, very young. I came to Silicon Valley to be a data scientist in 2013, so I'm probably part of one of the earliest waves of professionalized data scientists. When I came, there were really no programs to teach you this kind of thing. The other day I was told that they have lots and lots of people coming in with master's degrees in data scientists. They didn't exist. In fact, I taught at a bootcamp because there was insufficient education for data scientists.

And that's how the culture has been. And it's great to keep that spirit of sharing and openness, but also appreciate the responsibility that comes with the cool tools and toys that we have. So I think part of an interesting thing to watch is how the culture of data science and data scientists evolve. And I guess as an example, I think NERPS is interesting, where they've introduced an impact statement and the controversy and discussion around this impact statement is a really good example of what I'm talking about.

Caitlin McDonald:

Yeah. That makes sense. I could talk about this obviously for hours and hours, but unfortunately we have limited time. So are there any final thoughts or key takeaway messages that you have for people listening to this podcast?

Rumman Chowdhury:

Oh gosh, so many. I would say the biggest one I think about a lot lately really is governance. It does come up a lot. And this concept of the equality in the systems, and/or the systems of power that exist in creating models and the governance of these models. So it's easy to say that, "Oh yeah, we should unsilo teams and we should create ABC or whatever." But I think fundamentally it's actually about the systems of power whereby people who can code and build the technology, people like myself, are often centered in this conversation as the arbiters of good and bad and right and wrong and the gods who bestow upon everybody else information. And if there's one thing I try to do in all of the talks I give is to tell people that they too have agency over these things. They have their expertise from which data

scientists can learn a lot. So in these conversations and in these rooms where you're... even if it's an interdisciplinary conversation, observe the power dynamic and make sure you're asserting for the individuals who are not necessarily the traditional people in power.

Caitlin McDonald:

Such a great final thought for people to take away, especially as they're considering their future development of their organizational structures. Rumman, thank you so much for joining us today.

Rumman Chowdhury:

Thank you for having me, Caitlin.

Caitlin McDonald:

Thanks for listening to the Growing Digital Ethics and Practice podcast. You can find out more about Rumman's work at www.rummanchowdhury.com, and you can find out more about the Leading Edge Forum perspective on digital ethics just by Googling the phrase stemming sinister tides, where the first link should be our 2019 position paper [Stemming Sinister Tides: Sustainable Digital Ethics Through Evolution](#).