



Solving for Tech Ethics™

Episode 2: Multi-lens technology ethics with NVIDIA's Nikki Pope

Moderator: Beena Ammanath, US Leader, Technology Trust Ethics, Deloitte Consulting LLP

Guest speaker: Nikki Pope, AI and legal ethics leader at NVIDIA

Beena: Hi, my name is Beena Ammanath. I lead Trustworthy & Ethical Technology at Deloitte. Today, on Solving for Tech Ethics, we have Nikki Pope, AI and legal ethics leader at NVIDIA. It's great to have you on the show, Nikki. I am so happy we're having this conversation. Welcome to the show.

Nikki: Thanks, Beena. It's great to be here.

Beena: Can you share with our audience a bit about your background and how you got to your current role?

Nikki: I guess you could say I got interested in AI because I've been involved in criminal justice. I'm a lawyer, although I'm a corporate lawyer, not a litigator. I was on the board of the Northern California Innocence Project for a number of years. I started paying closer attention to issues involving criminal justice, and I read about predictive algorithms. That's what got me down the path of looking into AI because of the way predictive algorithms are used in the criminal justice system for bail, whether someone gets a longer or a lighter sentence, all sorts of ways

in which those algorithms are used. There are all kinds of studies showing the bias that's inherent in those algorithms. That really drove me towards AI.

Beena: Nikki, you've made this interesting transition from academia to corporate, anchoring on this ethics lens. How has that transition been for you? I also want to touch on your unique experience with nonprofits, but first, let's talk about how has it been moving from academia to the corporate world?

Nikki: Well, it's an interesting and challenging transition. In academia, you have an idea or there's something that you're interested in doing research on, and you do it. You do the research, you write a paper or write an article. There's really not an awful lot of constraints on where your mind goes and what you want to work on. You can work on anything. I was the managing director of the High Tech Law Institute at Santa Clara Law School at Santa Clara University. I spent my time focusing on algorithms and criminal justice and the intersection between technology and criminal justice and social issues. No one stopped me because it's what you do. There's not the same degree of freedom when you come into a company because they have businesses that they are running. What you do has to advance that business objective in some way. I'm fortunate that I came into NVIDIA, which is a tech company working in AI. What I'm interested in is very important to the company, but there's not the same degree of freewheeling, do what you want, inside a company as there is in academia.

Beena: What about between nonprofits and for-profit companies? Is there a similar set of differences?

Nikki: I think there can be. I think it depends on what the nonprofit is. I run a nonprofit in my spare time that works with exonerees, people who were wrongfully convicted, and helping them transition back into society. That's a different initiative or it has different objectives and goals. I think it's different when you're working in a company because, as I said before, a company has stakeholders that are owners of the company, stockholders that you have to respond to, there may be regulatory agencies that you have to respond to. It could be just the SEC, the Securities and Exchange Commission, but there could also be the Food and Drug Administration, and if you're dealing as we do with autonomous vehicles, and you have all the transportation regulatory agencies. In a nonprofit, there

aren't a lot of regulatory agencies that you're having to deal with, at least in my nonprofit. There are a lot more stakeholders and there's a lot more interest, I guess, from outsiders in what you're doing.

Beena: Having straddled these two worlds as well, I think the reason, in addition to the stakeholders, there's also a more focused effort. To your point, you just cannot pick up any point and start any topic and start looking at it. It needs to be aligned with the business goals, and there's more well-defined focus as opposed to being completely greenfield.

Nikki: I think so, and I think that goes right into one of the questions that we discussed previously that had to do with how you or how I or anyone in this position does their job inside a company, and it's really essential to collaborate with others, especially if you're introducing ideas that are novel to the company or if you are going to affect someone's business in a way that they don't particularly appreciate at the time. I think it's essential to move slowly and deliberately, understanding what the business issues are, and getting buy-in from the people whose businesses you're going to suggest changes to.

Beena: A role like yours is new to the industry. In general, to companies these roles didn't exist. Even 5 years ago, 10 years ago, these roles just didn't exist. How would you describe your role? What does a day in your job look like?

Nikki: I don't know that a day in my job looks like something specific to the next day. One day does not look like the next. I could spend one day, I'll give you yesterday for an example, I had about eight meetings. They started out with talking with engineers about risk management and risk assessment and assignment of risk. Then, the next one was dealing with explainability and how do we do that. We have our GTC

conference coming up in November and I had a meeting on that talking about bias and how you assess bias or attempt to mitigate it. Then, there was a meeting with a potential toolmaker on assessment tools and mitigation tools with various people from engineering. Then I had a legal meeting where we were talking about licensing of our models, codes, technology. Then, at the end of the day, I had a really short meeting on privacy policy.

Beena: It's so varied. I completely agree. There is so much to do in this space that it's almost hard to figure out all the possible things that you could do within a day because not only the company that you are in, but also the breadth of topics that you need to get into. There are so many lenses to look at when we start talking about technology ethics, AI ethics.

Nikki: Absolutely.

Beena: NVIDIA works with so many industries. It is not just about figuring out ethics for a specific industry, your supercomputers are powering the entire world and the ethical implications have to be considered from all these different lenses. How do you ensure that the outcomes that you drive are positive? What are some of the scenarios that you try to avoid? How does your team prioritize these things?

Nikki: Prioritization is a challenge because, as you mentioned, not only are we across a lot of industries, but we're global. We have to consider when we talk about ethics, and I really talk about AI being trustworthy or responsible, not ethical, because I think ethics is a charged word and it means different things to different people, but you understand what being trustworthy is and you can understand what being responsible is. When we build responsible technology or trustworthy technology, you have to consider not only what would engender trust in a particular industry, but a particular

region. For example, if we have an AI model, we don't do these, but if an AI model is a recommendation model and say it's going to recommend that you buy this red dress, there's not a lot of harm that might come from that model for recommending clothes or recommending shoes or what have you, but if you have a recommendation model that is recommending a drug you should take, that's a completely different situation with a lot of ethical questions and a lot of regulatory oversight. There may be a difference between how that question is addressed in the US versus the UK versus Singapore or South Africa or Brazil. You have to really consider not just the industry, but the society and the societal norms and values.

Beena: So true. The challenge is also we don't have all the regulations figured out. It is still very nascent, early days on, forget about regulation, but what are the best practices? There's no playbook talking on how do you solve for technology ethics across all these industries. What are your thoughts on when will regulation catch up or is it going to be taking a while before we see policies start emerging on this?

Nikki: I think it's going to be a while, and I think it's going to be a while before the different regions have some consistency across their regulations and policies. But for me, the legal regulations or legal obligations are like the baseline, that's like the bare minimum that you should be doing and you should be doing it because the law requires you to do it. But I think it's possible for companies to have these conversations now where you are talking internally about what, for your particular business, and it's important that they know what business they're in and what their mission is, what the ethos is of their company. One car company may have a different mission and a different ethos than another car company, even though they're competitors. So understand what your company's ethos is and deliver on it, be true to it. You don't need the law

to tell you how to do that. You need, my grandmother used to say, common sense. Common sense can help you get to a lot of these solutions, doing what's right. I don't think it is inconsistent for a company to do what's right and still be profitable. Lots of companies do it.

Beena: Yes. Do you see any company, without naming names, do you have a favorite example of an organization that has applied tech ethics in the real world successfully?

Nikki: I have an example, but it's multiple companies. My example is there was a study done at MIT by a data scientist on facial recognition technology and the biases, skin tone racial bias and gender bias in facial recognition technology. The report was published. The data scientist called a number of companies that build facial recognition software and convinced them to stop providing their software to various organizations like law enforcement when there was such a bias built into the software. I think that's a huge win for AI ethics or responsible AI.

Beena: So true. You and I have talked about this in the past. There is a lot of noise around tech ethics in general. You hear a lot of headlines. You hear there's a lot of fear that is being pushed out. What do you think are some of the misconceptions? What are some of the things that are misunderstood by people who are not really working on solving for ethics? What are the big misconceptions that you've seen in your experience that you would like to share with this audience?

Nikki: Well, I think the biggest misconception is one that I had before I started researching in this space, and that is that the technology is a lot further along than we think, that the technology that Hollywood is the reality and it's not. There's no Cyberdyne, there's no Skynet. Machines

aren't going to come and take over the world. I had this false sense of where the technology is and what AI is and what it does. What I think would help people to understand is that artificial intelligence, it's weird that it's named artificial intelligence because it's really a machine that learns what we teach it to learn, and it learns it over and over and over and over. If we don't present the information, it doesn't learn whatever the thing is. It reflects back what we put in and to me that's not intelligence. It's not really bringing in more information and thinking, it is looking for patterns of behavior in the data that we feed it. I think when I understood that, it made all the difference like, "Okay, now I know what I'm dealing with here. This is not a thinking machine, it is a pattern finding or pattern observation tool." I think it would help for people to understand that and for it to be explained in a way that regular people who are not scientists and data scientists can understand.

Beena: That's so true and I know you and I have talked about this. The nonprofit that I started called Humans For AI, which is really about just driving more AI literacy so that everybody can have that basic fluency understanding about even what AI is, and think about it, like even other technologies as they come to the fore and it creates value, there are going to be negative implications. There are going to be risks associated. When we talk about AI ethics, it does tend to get very quickly focused on bias and you've mentioned a little bit about explainability. What are some of the factors that you think fall under the umbrella of building trustworthy AI?

Nikki: If you take a step back and think about what trustworthiness means and how you build trust, if you're building trust one party to another party or human to human or machine to human, however you're doing it, the person who is receiving that benefit needs to understand how it works and you

need to build the trust. I think getting to the question that you asked just before about what people don't know or what people need to know is I think we all, not all, but generally as a society, accepted this idea that machines can do better than humans. Machines can solve problems better. They can solve them faster. Machines are better. Therefore, a machine solution is better than a human solution. That's not necessarily the case. It's important to understand that AI has the ability to solve a lot of problems and to help us solve problems, but it's not the problem solver. Once we started to realize that there were problems in the way these algorithms or these models or systems were designed, we lost trust in them. We had, I don't know, a childlike trust originally. "Oh, wow, these machines are going to be great." Then, they didn't perform great. They were biased against men or women or black people or people who speak with an accent or whatever it is and we started to lose trust in the ability of those machines or those systems to be fair. I think in order to build trust, one of the essential elements of trust is explainability. I think we need to be able to explain to a person who, for instance, was denied a mortgage loan, why they were denied and how that happened and for the loan officer who is making the decision, that person needs to understand how that algorithm reached the outcome that it did, how it got to that suggestion, and everything that went into it. I don't think it's impossible to explain AI. You just have to do it in bites and you have to figure out who it is that is asking the question and what it is they need to know. An example that I give is a self-driving car. If I'm buying a self-driving car, I want to know how it works, but I don't want to know exactly how it works. I just want to know how it decides to take the freeway instead of taking a surface street. I want to understand how it identifies a person versus a truck or a bicycle, but that's really about the extent of what I need to know, but if I'm a regulator or I'm with the National Highway

Transportation and Safety Commission, I want to know details of how the engineering works, how the technology works. I want a completely different level of explanation. When you talk about explainability, it's really important to know who it is you're explaining to and what it is they want to know, and your explanation is going to vary. You'll have different explanations for different audiences.

Beena: So true. Explainability is crucial. You make a very good point, Nikki, as to it has to be targeted for that audience. You need to be able to build your explainability in a way that is understandable to the group that you're targeting to. It's not a one size fits all. It's almost like you need to use AI to understand who the target audience is and then personalize that.

Nikki: It also comes into play when you're developing your model. In privacy law, there's this concept of privacy by design, where you build in the privacy into the product. You don't tack it on at the end. I think we are moving towards, companies are moving towards this idea of ethics by design or trustworthiness or responsibility by design. You have these conversations, development teams have these conversations when they're thinking about, "Oh, I got this great idea for an AI model that will do X" and then have the conversation then as opposed to at the end going, "Oh, I need to check off that ethics box." Then it might be too late. You've built something that you cannot fix after the fact. Then you have to go back to square one to begin with. I like what I'm seeing with more ethics by design being brought into the development process.

Beena: Love it. That's a great point too. You need to look at the existing processes and make sure that you're thinking of ethics early on in the process, but how do you empower your organization, the people, the talent? Is it just the data scientist? Is it just the technologist? Is it every employee? What are your thoughts on how do you empower

everybody in the organization to think about ethics and understand what the company's stand is on it?

Nikki: It's a great question. I think companies, or people in companies, employees are, well, they're people, they're consumers. There already exists that concern about ethics and about responsibility. You can see it in things like climate concern. That perspective is already there. It's just a matter of management in a company, making sure that employees understand this is important and this is important for the company and you are all a part of this process. If you see something that isn't working, say something about it, or if you have an idea of a way to make something better, say something about it. No matter where you are. I mean, the idea can come from someone on the development team, but it also could come from someone in marketing and communications or someone in HR. You don't know where the idea or the record with this suggestion is going to come from, and it's important because AI, if it doesn't touch everything now, it's going to. It's going to be in reviewing resumes and deciding what brand of coffee to buy in the cafe. I mean, it'll be in everything, in all sorts of solutions. It's already in Netflix and Amazon and every place else that we use online. We're already interacting with it, and it's just going to be the same inside companies and outside. We should start thinking about that and building the vocabulary for it and understanding what it is we can do to improve it because we are AI. I mean, we are it. It's a reflection of us.

Beena: It's everywhere. Going back to your point on ethics by design, that's great for companies who are very early in their journey of looking at AI, but there are companies that are more advanced in their AI journey and they hadn't really thought about ethics early on. Now they have to do some level of catch-up. What's your advice

to companies that are more advanced in their AI/tech journey? How should they be thinking about ethics? How do they play catch-up? What's the best way to approach it?

Nikki: That's a tough question. It's going to be a challenge for companies that have a lot invested already in AI, but there are two things. There's one really good example and that's the GDPR, the EU Privacy Regulation. That came well after lots of companies had already collected data on people and have databases and personal information. In companies like financial institutions and in the health care industry, they already had regulations dealing with handling personal information and sensitive personal information. But companies like retailers and social media platforms, they didn't have that sort of regulation and all of a sudden now they have to have it. My advice to companies that are further along in the process may sound a little cruel and heartless, but it is regulation is coming and so start now. You can look at the EU's Artificial Intelligence Act as a road map for what potentially could come, but there's likely to be much more stringent, much more restrictive regulation coming down as well. Even if you're in the US, you can look to the Federal Trade Commission. Even the FTC put out guidance on artificial intelligence with respect to anticompetitive behavior. Earlier this year, they said, I'm going to paraphrase this, but it's basically, "If you are producing, if you have an AI system that is biased or that does not treat people fairly, that is anticompetitive and there's a way that people can complain, file a complaint against the company that does that." It's already here, and there are markers here for companies that are more advanced to know what they should be doing. To me, there's no excuse for not starting the conversation.

Beena: Solving for ethics is hard. You need to understand the technology, you need to understand the regulation policies, you need to understand human behavior, business motivation. There are so many

aspects to it. What are some of your tips on how do you get your arms around solving for ethics in a large organization? No matter which stage you are in, what are some of the dimensions that companies should think about?

Nikki: You definitely need to have cross-functional or interdisciplinary teams. I don't have a computer science background, I don't have a tech background, an engineering background, but I do have a legal background and I've been around for a while and I've worked for a few companies, but it took me talking with our engineers to understand how a particular AI system works. It's different from industry to industry. When you're talking about financial services algorithms, I talk to the people in our financial services industry to find out what's important for that industry and how do these algorithms work, how do these systems work? That informs what it is that you need to address because you have to prioritize. There's so many things, there's so many businesses, so many models, and so many opportunities, you have to figure out where you're going to focus first. I do that with the help of people in cybersecurity because they've done this a lot already. Engineering. I turn quite a bit to the privacy law folks because they had to deal with this in a very dramatic way with GDPR. Then the other thing that I did when I first joined NVIDIA is I reached out to my counterparts such as they were at other companies, mostly tech, but some not tech. Some of them were further along in this process than we were, and some of them were not as far along as we were. It was really helpful to talk to people who have gone through a lot of what we're going through to find out the best ways to move forward and also things not to do. How to navigate in this space.

Beena: We are all in this together and nobody has it all fully figured out. Coming together, and I think that's how you and I met, Nikki, is trying to connect. I'm a technologist by training, but I have long

admitted that this is one of those problems you need that diversity of thought from a professional background, educational background, geographic background, you need so many different experts coming together to be able to actually solve this extremely fuzzy, gray area that we are all trying to get our arms around. You're so right. Depending on the company, the industry you're in, it's going to be a different lens. It's going to be a different solution.

Nikki: Exactly. I will say that one of the things that was surprising to me and really welcome and was delightful is how much sharing of information there is between companies that otherwise are competitors. I mean, I have talked to people who are involved in tech ethics or responsible AI or just robotics in all sorts of companies and they've been really forthcoming in sharing what they've learned and likewise I have as well. I think it's because the seriousness of what we're doing and the opportunity, I mean, there's a tremendous opportunity to, I don't want to say save humanity, that's not really what it is, but solve some of the big problems that we have. It's possible, and all hands on deck.

Beena: That is so true. I mean, I'm definitely seeing more of that collaboration on this topic than anything else because we are all in this together. I agree whether it's save humanity or if we don't solve for this, I think we will hit a roadblock. At some point, we'll have to think about the progress we're making with technology and the risks associated with it. It is an absolute must for all businesses to figure this out. Nikki, any parting words for our audience as they think about this progress? I took a few notes— build your network, find your peers and reach out, figure out ways to collaborate, nobody's got this figured out, you need to collaborate internally with stakeholders from different teams to be able to solve for this. What are some of your parting words for our audience today?

Nikki: I would say particularly people who are going to come into a position like mine or who are interested in doing this work, educate yourself. I think it's the most important thing. There are lots and lots of books. I can't even tell you how many books I've read on AI ethics and machine ethics and business ethics and the future of humanity, all of those. There're some really good books out there that are...some are more to one side than the other, so I would say a good balance across the board of the books that are humanity is doomed and AI will save humanity and everything in between. I would say definitely educate yourself and part of that education also includes attending webinars or conferences on this topic because you learn an awful lot from just having conversations. You and I have talked a lot and I've learned an awful lot talking with you about just the everyday questions that pop up when you're doing

business. There aren't a lot of us in the corporate world. There are a lot in academia, but there aren't a lot of us in the corporate world. I think we should definitely reach out to each other and share the best practices and the knowledge and information because we'll only move forward faster if we do that.

Beena: Nikki, so well said. I'll just add on that there are not many people in the corporate world who are really focusing on the solution. There are people, to your point, in academia, but there's also a lot of marketing and clickbait headlines that create a lot of hype around it, but it is just headlines. I think we need more warriors like you who are focusing on solutioning for it, to actually figure out, "Okay, we understand there is bias. How do we solve for it?" Moving the conversation forward towards actual impact. We need more warriors like you and

that was actually one of the big reasons for starting this interview series is to feature people like you who are actually moving the needle on solving for it in the real business world that everybody can learn from, can connect with, and we can build our own tribe of warriors focused on solutioning. Nikki, I appreciate your time today. This was a great conversation and I know we must continue. Thank you so much for joining us today.

Nikki: You're welcome. It was fun.

Beena: Nikki, thanks again for being with us on the show, and I want to thank our audience for tuning into Solving for Tech Ethics. Be sure to stay tuned to [Trustworthy & Ethical Technology at Deloitte](#) for more of our latest thinking and market insights. Thank you and take care.

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to deloitte.com/about.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.