Published on *The National Law Review* https://natlawreview.com

AI 2030 Show With Nikhil Pradhan [Video]

Article By:

Nikhil T. Pradhan

Nikhil Pradhan (Senior Counsel, Boston) joined FinTech4Good's AI 2030 Show to explore the legal risks and rewards of generative AI as it becomes more integrated into the business environment.

Transcription

The below episode transcript has been edited for clarity.

Xiaochen Zhang

Good morning, good afternoon, good evening. Welcome to the AI 2030 show. My name is Xiaochen Zhang and I am the Executive Director of AI 2030 and also CEO of FinTech4Good. Today, we have Nikhil Pradhan on our show and before we dive deep into the topic can you introduce yourself?

Nikhil Pradhan

Yeah, thanks so much for having me. My name is Nikhil Pradhan, I'm an intellectual property lawyer with Foley & Lardner. The majority of my practice is helping companies protect their innovations and navigate what their competition is doing, particularly using patent protection. I also do a lot of work helping companies figure out how to understand AI technologies and solutions and how to implement them effectively in their businesses as well as to protect their innovations around AI. So, really looking forward to this conversation with you.

Xiaochen Zhang

Awesome. And when you talk about all the innovation happening around generative AI-based applications, there is definitely tons of interest, you know, from both employers and also employees –

a lot of them are now using generative AI-based applications to do their job. And then you know, if I'm an employer and all my employees are using different type of generative AI tools, what kind of risk is associated with that? And how can I as employer make sure that all those risks are managed?

Nikhil Pradhan

I think data is going to be a common theme through all these questions, both from the perspective of what data does for a model that already has in it based on its training, as well as what kind of data do your employees and users in the company want to be providing to the model**.** And then I think the other aspect to keep in mind is, you know, having kind of a big picture strategy and principles around how you want your teams to be using these tools. So that might mean making sure that you're focused on the accuracy of outputs to the extent that you might be using customer data or, you know, in the health care space patient data to make sure that you're maintaining the trust that those people have given to you for using this information. As well as to make sure that you're kind of managing any risks associated with bias when it comes to models that might be trained on one subset of data but might not necessarily map well to the person for which you're implementing these tools.

So starting from that point, I think one thing to try to do is help your teams get onboarded onto commercially provided systems because at least right now those tend to have pretty good protections for data when the models are used. Typically they'll say "we don't own any inputs that you're providing to us, we don't own any of the outputs." And so that kind of creates a phase one level of security where, even if our users are putting information in that might have business or other kind of confidential value, we know that the model is not going to retain that. Similarly, if they use the model and come up with some sort of innovative solutions, some sort of new technology, you can kind of rest more assured that whatever outputs that are derived from the model will be retained by your company and your users, rather than being used by the model to update itself or in a way that might be exposed to other users from other enterprises.

And then again the other thing to think about is understanding what was really used to train these models, right? There's a lot of different cases going on right now trying to answer these questions, like the New York Times case for example when they're going after OpenAI and saying "hey, you took a bunch of our articles and now you're exposing outputs that look a lot like (if not identical) to what our writers created" and so just having some kind of awareness that other people might have ownership over those outputs that you're receiving from the model.

Another example, let's say your team wants to use some sort of automatic code generator. It is helpful to remind them that even if the model itself seems to be creating new useful code from scratch it's been trained on all kinds of different data that might have been provided through the internet to various databases and that code might have open source obligations on it, might have bugs, or other errors or security concerns and so even if it seems like a really valuable tool to quickly stand up specific functions, help speed up the development process the team should be aware that what they're getting might have some risk factors embedded right into it. They need to be careful when they actually implement it that they're checking if it looks like something that's already out there or running the standard kind of bug testing and everything that they would do otherwise.

Xiaochen Zhang

There are tons of different types of apps out there and when you talk to your clients, so far that where are the gaps and it's more as, you know, if again that's you don't want to just, you know, exposed to the risk and then, you know, to the extent that, you know, it's not manageable so for them that the

gap is more as, you know, they don't have the bigger picture or they don't, you know, have the checklist to ask the right question or is, you know, they don't have the capital or, you know, or they don't have the toolings. So where are the main gaps that you have seen?

Nikhil Pradhan

That's a good question. I think a part of the gap is that within any given enterprise there are a lot of people who have different small views of why this technology is valuable sosome teams want to use it because they know it'll help them create the output the work product that their roles need much more efficiently, you know, save them time help them kind of take gone more and develop more but they might not necessarily be thinking through in terms of the big picture like, you know, for my company as a whole how am I supposed to be, you know, implementing this tool effectively, you know, keeping those principles in mind from transparency to trust and accuracy and managing bias.

On the other hand there might be, you know, higher level leaders who know at a strategic level okay it's important for us, you know, for marketing purposes for external credentialing as well as business purposes internally to be implementing AI tools somehow but we're not sure exactly what that looks like or, you know, what's the good way to do it and what's not.

I think the simple answer to your question is there's, you know, a bit of a knowledge gap in different areas where everyone wants to use these things but there is any single person necessarily who has kind of a complete view of what that looks like. and so, you know, I think part of the solution there is just having open dialogues between different teams about, you know, what they're trying to achieve and and what are the implications of using these tools in different ways and then kind of keeping in mind the company's overall principles for why this technology could be valuable, you know, when they have these conversations so everything kind of stays consistent with the overall goals.

Xiaochen Zhang

Awesome. If I understand it correctly create a culture and also empower your teams. I think, you know, those are the key message where that's where you can just try, you know, as a company or organization that this can be addressed, you know, in a more systematic way rather than just no matter you are junior or senior you a company and then there's there gaps yeah.

Nikhil Pradhan

Exactly. I think that's a perfect encapsulation.

Xiaochen Zhang

Perfect and when again this you mentioned about just, you know, when you use it and then try to use those which already are like commercial apps where, you know, all those potential risks are being managed, you know, as that during the procurement process and there that's maybe you can further explain, you know, to us that when a company enter into a agreement, you know, with a service provider on Generative AI what, you know, what make this kind of procurement different from other type of corporate procurement and what are the key points that, you know, those procurement specialist they need to be aware of?

Nikhil Pradhan

Yeah I think, you know, on the first part of the question what makes things different I think part of it is that a lot of this is so new that, you know, new ways of using it are arising all the time and so things that, you know, with more with more mature technologies people have already worked through some of the implications and have seen, you know, how things play out there's just, you know, there's a better understanding in the market of okay if I'm, you know, if I'm using a cloud service I know where my data resides, I know, you know, if there if I have like export control issues with, you know, sensitive information I can try to negotiate for making sure that all my data stays in in the country that it needs to right?

Whereas I think with even though AI has been around for a very long time with Generative AI, you know, there's both a broader group of people who are who are using it already even just in last year and a half or so and I think also more implications in terms of, you know, what data has been used to train these models and that might be a bit opaque to the end user or to the company procuring the technology. And so I think that again kind of starting with the data point really that really leads to figuring out okay when I'm looking at negotiating for purchasing or licensing an AI tool, you know, I need to track the data both for how the model was trained as well as at run time when my users are using it so if we take for example like in the healthcare context. Let's say a healthcare provider is, you know, licensing a tool that helps them come up with recommendations for, you know, here's a treatment that we should be giving to a given patient based on information we've, you know, collect it about the patient so that their physicians can kind of make better recommendations about what next steps are. Starting with the data in and then looking at the agreement thinking about okay, you know, our physicians are going to be providing sensitive personal information about the patients into the model, there should be clear statements in the agreements that know nothing that is being provided at least about patients if not more generally can be retained by the model or by kind of any other functions that go along with it. And then again more on the data side thinking about bias considerations so, you know, ideally there's a way to kind of understand when the model generates its outputs it should have sufficient transparency to tell you okay, you know, here's information that was used to support the conclusion about what kind of recommendation for therapy to give to some patient.

Here's information about distribution of existing information that was used to then train up the model and help it get to that recommendation so then the physician can look at that and, you know, make a comparison and validate okay this makes sense we haven't inadvertently provided some recommendation that's, you know, based on one population segment that doesn't actually reflect the person that I'm trying to treat here. And then, you know, continuing from there thinking about okay if things go wrong whose responsibility is it to, you know, handle any liability that comes out of that. So, you know, typically what you'll already see is that if, you know, if the users are kind of using models in a way that complies with their basic terms. So not trying to jailbreak things not trying to push the edges of what the model can do or should be doing, you know, typically the model provider will say, you know, we'll take on we'll have indemnification so we'll take on the risk if something goes wrong. And they might even say things like, you know, especially if like if the user puts like risk mitigation into their prompt so, you know, if they say things like please, you know, provide me like the most accurate answer or provide me the answer that you determined with the highest confidence you might see those in the agreements or you might ask for them in the agreements if you're the technology provider because, you know, at that point like you want to you want to point out that, you know, if the person buying the technology from you assuming they use it properly it makes sense that like, you know, you would be taking on liability if things go wrong.

The flip side of that being, you know, they might say if, you know, if you enter prompts that go against the terms that we set out then, you know, you're responsible for what happens with those outputs

rather than, you know, us as the technology provider. I think another important thing that might still be, you know, for the future rather than currently existing but if you think about like again kind of in the healthcare context. Any given physician using this type of tool might not have that global view of okay what are what are the outputs looking like for like all the people in my practice or across like our entire health care system and, you know, do have a good sense the recommendations it's giving, you know, make medical sense, are appropriate to the populations that we're using them for. They might only have that view of the patients that they're working with and so having, you know, kind of like a higher level user in the loop who can evaluate all of the apps that are coming out and kind of consider bias and accuracy at a more global level I think that'll be something that would be really helpful for companies to consider so that, you know, it's not necessarily the responsibility of each end user to be doing all these, you know, bias and accuracy checks every single time that they run the model but there's someone with a specific role who can help them, you know, manage that more comprehensively.

Xiaochen Zhang

Thank you and obviously as what you mention as, you know, there are tons of potential new risk where that a typical, you know, IT or application procurement may not involve and it may just, you know, try really change your relationship, you know, with your customers, change your relationship with your regulators and definitely I think it's a lot more complicated than the other type of, you know, application procurement. So maybe in the future we can just working on, you know, responsible procurement related, you know, Workshop or something to help, you know, those folks in the future to understand what are the new elements which are, you know, introducing into that. And of course that, you know, that also changed some, you know, the nature of the procurement specialist job, you know, where that's you need to really just address tons of new things which are, you know, you probably have not been thinking of in the past, you know, no matter how many years you're doing it and then you thought that, you know, you have you have all the skill set needed to do it in the right way but now that with the, you know, AI procurement is really a new game.

And related to that, you know, want to just ask you another question and which is around, you know, the IP protection and of course that's one excitement, you know, for the enterprise no matter big or small to use AI is, you know, to trying to just enhance their power in innovation and there that's naturally that's again because of Generative AI is different from any other type of, you know, foundational infrastructure to enable innovation where that's a you there are you really just bring in a new element which can just, you know, have a lot of consequence implication. Which is, you know, a lot more simplicated or complex compared with the, you know, the other, you know, application or foundational infrastructure you brought into the system in the past. So IP is, you know, one of the key thing within the, you know, Generative AI related know where that naturally when you want to leverage Generative AI you need to put your data, you need to put your idea, you need to put your IP into an a platform where that's highly possible others are able to access to it. So, you know, how when you are, you know, innovating with the Generative AI how do you protect your IP?

Nikhil Pradhan

Yeah great question so on that specific point you raise about, you know, putting your confidential information, your innovations into models, you know, I think that's something that kind of going back to what we've already talked about it just really requires coaching on I guess avoiding that as much as possible or, you know, making sure that you're doing it with models where, you know, that you have agreements that your data is not going to be used by the model. In terms of protecting innovations that are being developed in this space typically the two most useful tools that will come

up are patent protection and trade secret protection. At high level what patent protection gives you is the ability to prevent others from doing something that you've claimed in which you've shown is new relative to what else is out there. And it's also limited to a 20-year term basically and so it's kind of, you know, you get you get 20 years of preventing other people from being able to do something in exchange for by making your patent public, you know, teaching the world about the innovation that you've come up with.

So typically we'll find that patent protection is valuable for kind of high level functions and structures that explain, you know, a solution to some sort of technology problem and that show what the performance benefit that is. I'll circle back to that but on the flip side trade secret protection is a way for you to ensure that if your confidential information somehow is, you know, stolen, misappropriated that you can kind of go after the person that kind of took it away from you. So typically what you'll see is, you know, especially in like the AI software contexts people might establish trade secret protection around more kind of low level specific ways of implementing things that, you know, just looking at their product someone might never actually be able to know if that's how it works but it still has business value to you to, you know, deploy something in a specific way, use a particular, you know, software architecture and so on. So, you know, in terms of patent protection I'm sure you see all the time people say, you know, data is our mode and I think the patent protection gives you kind of another pillar in in building up that mode both for companies that are, you know, developing fundamental AI technologies as well as for those that are, you know, taking maybe to some extent taking what's out there but then, you know, tweaking, fine-tuning, using it to solve a specific business problem that their company has or that their customers have. , you know, in in both those cases it's really valuable to think about what exactly is the problem that's being solved and especially if you can cast that in terms of a technology problem so, you know, are you using or your customers asking for you to use AI because they want things to be faster to have quicker turnaround times use, you know, less network information anything like that.

These kinds of like specific computer technology problems and then from there kind of building to explain okay this is why we need to use this type of model or this type of machine learning architecture to solve those problems. Kind of similarly for companies that are, you know, deploying off the shelf AI solutions understanding okay why is it important to use a language model in this in this case right? Again are you are you relying on it to get kind of more accurate information are you doing in a way that kind of, you know, speeds things up or reduces inaccuracies or lets, you know, access data that you or make use of data that you couldn't afford. So, you know, a lot of companies now are turning to language models because they have, you know, they have this wealth of unstructured data, you know, whether it's reports or, you know, clinical notes anything like that and they can use the language models to figure out insights from those and then kind of level up the advantages they're able to provide to their customers and users.

Xiaochen Zhang

Thank you and I think, you know, we have a very clear picture that, you know, all this very important questions which are raising Generative AI and definitely I think, you know, there are a lot of space where that we can continue the conversation and then to broaden it and also provide a lot of very concrete guidance how that can be addressed so this is just, you know, such an exciting topic and I'm looking forward to just have a further conversation and then to just, you know, dive deeper in into, you know, all the aspects we have talked about.

Nikhil Pradhan

Absolutely. Yeah it was great talking this through with you. I think, I agree that, you know, things are just getting started and all the innovation out there is going to be very exciting so I'm looking forward to continuing the conversation with you as well.

Xiaochen Zhang

Thank you and thank you everyone who joined us or listen to the recording in the future and yeah we are looking forward to continued conversation with Nikhil and dive into procurement and also, you know, Generative AI related innovation related challenges in the future with you. Thank you for your time.

National Law Review, Volume XIV, Number 110