***The Good Bot: Artificial Intelligence, Health Care, and the Law* — Preventative Medicine: Health Care AI Privacy and Cybersecurity**
Host: Brett Mason
Guests: Brent Hoard and Morgan Hague
Recorded: 5/1/24

Brett Mason:

Welcome, everyone, to *The Good Bot*. This is our second episode of the podcast. *The Good Bot* is a podcast focusing on the intersection of artificial intelligence, healthcare and the law. I'm Brett Mason, your host. I am a trial attorney at Troutman Pepper where my primary focus is on litigating and trying cases for life sciences and healthcare companies.

However, I am a tech nerd and a tech enthusiast. I'm also deeply fascinated by the role of technology and advancing in the healthcare industry. Our mission with this podcast is to equip you with a comprehensive understanding of artificial intelligence technology, its current and potential future applications in healthcare, and the legal implications of integrating this technology into the healthcare sector.

If you need a basic understanding of what artificial intelligence technology is and how it is already being integrated into the healthcare industry, I highly recommend you start off by listening to our first podcast episode *In This Podcast*. There, I had a great conversation with Morgan Hague, who we have on again today for this episode, to just lay the groundwork of understanding what technology is being used and what we see is going to be used in healthcare moving forward.

I'm very excited about this episode today because we are focusing on two issues that mostly come up whenever you talk about artificial intelligence. We're going to be talking about the key security and privacy implications around the use or misuse of artificial intelligence.

As I mentioned before, we're joined again today by Morgan Hague, who is a manager in the IT Risk Department at Meditology Services. We're also joined by one of my colleagues, Brent Hoard, who is a partner at Troutman focused on privacy and data protection issues.

This is our first official podcast episode of our series where we're going to be talking about data privacy and cyber security issues. Because there are so many issues, we can't get it all into one episode for you. If this is an issue you are very excited about or interested in, please check out all the different episodes that we have rolling out, because there's going to be several on this topic.

I just want to start us off by setting the stage. Morgan, Brent, thanks for joining us. Morgan, could you provide us some background on how artificial intelligence is being used currently in healthcare beyond those that we just talked about in our first episode?

Morgan Hague:

Exactly. Yeah. Really, the sky is the limit I like to say with AI. I think a couple of the ones that we're really seeing fall into these two lanes, right? There's predictive analysis. Or, really, just

data manipulation, data forecasting. It's just taking these big troves of data that providers, payers and insurers have and applying some sort of logic to them, "Hey, look. Looking at this, can you tell me X, Y or Z?" That's kind of one of the major use cases we're seeing.

But the other piece really is automation. Companies want to become, as always, leaner, faster, more efficient. To they're using AI to drive a lot of outcomes on that end can we reduce cycles in the claims department, whatever it is. At a high-level, I think that's what we're seeing. That drives a lot of decision-making for providers and, hopefully, also helps us improve some of the patient outcomes from the healthcare side.

Brett Mason:

Thanks for that, Morgan. With that understanding, Brent, are you seeing these types of technologies being incorporated into the practice or into the work for some of your healthcare clients? And if so, what are some of the recommendations you might have regarding entry points for integrating these types of artificial intelligence softwares?

Brent Hoard:

Thanks, Brett. Yes. A lot of entry points. And even if you are not a large, sophisticated organization that has access to data lakes and wants to create its own AI models, there are plenty of different applications of AI that are being built into sometimes the products that you're using. You may not even know it. But places where you can look for entry points, there is natural language processing, which would be used for clinical documentations, records, note-taking, dictation, saving a lot of effort in those things. Think of that as your voice-assisted product that you have in your house, but this is doing it in a clinical setting.

You have generative AI, which creates responses. And that can be used for patient engagement, education, patient self-help, so you can offer a lot of different tools on your website, or some point for your patients to be more interactive with your practice, or to learn more about a condition.

You also have, as Morgan mentioned, analytics, and machine learning, and deep learning that are used for clinical decision-making. One cool application that I've used personally is for image analysis where I sent in a picture of a mole and it was analyzed using the AI technology and came back and told me, "Is there an issue or not?" And that was just really simply done on my phone. I did still follow up with a visit to the dermatologist. I live in Florida. But that was really just kind of a cool application. And so, other things, medical devices that can be tracking different things. Whether they're functioning properly. Whether there are issues to alert the medical providers that are there about the device.

And as I mentioned, this is being incorporated into a lot of existing products. Or you can look for it as potential add-on products as like a bolt-on to existing services that you already have. And looking at that, I think right now, the state of the state, there's probably less risk associated with some of the administrative functions that I mentioned, like dictation, or scheduling, or patient engagement, as opposed to clinical. But everything is developing so fast that it is a very interesting world. And it's just going to keep going.

Brett Mason:

Right. And the beauty of this podcast is we're starting towards the beginning where there's a lot of looking forward and thinking about how it's going to be integrated. And, hopefully, we can provide some value to be thinking about the legal issues around that. I do have to ask, Brent, was everything okay with the mole that you sent in?

Brent Hoard:

It was. Yes. Thank you. Yeah. All good.

Brett Mason:

Okay. Good to know. And if it later turns out not to be okay, then we may be having a different discussion about the accuracy of artificial intelligence tools for giving you some research on your potential mole.

We're talking about a lot of different ways that artificial intelligence can be used both from using the external tools like you talked about, from different vendors, add-on services, or potentially internal tools. What are some of the more significant security and privacy challenges we're seeing with this widespread adoption of artificial intelligence? Morgan?

Morgan Hague:

It's a little bit interesting, because I think we're starting to see the kind of boundaries form around the different kinds of attacks. If you're in security, a lot of people like to form what they call attack chains. Basically, if somebody wanted to steal your TV, they could come in through the window, they could come in through the front door. Whatever it is. There's different kind of vectors and different things they would have to do for that. Although, obviously, they're probably not stealing a TV. They're stealing records and databases. Right?

With AI, what we're basically seeing is three major categories. One of them is what I call attacks targeting AI behavior. You're a company, you're investing millions, potentially hundreds of millions of dollars in your model. If somebody can infiltrate the environment and change parameters, poison your data, that's what they call it basically, it really nullifies the ability for your model to produce tangible outcomes. And so, it has a material impact on the viability of that business function.

And so, there's a legitimate risk there that's a little bit different than traditionally people are going to get in and lock your data via ransomware. Or they'll exfiltrate it and sell it on the dark web, whatever it is. With AI, there's a little bit of a different outcome that's kind of unique.

I think the other piece that that people have started to think about is what we call AI augmented attack. And so, that's going to be – there was a toolkit named Black Mamba that was built in a lab, but it was basically a malware that could evolve using AI to become more and more potent, basically.

It was effective in the lab. Knock on wood. We haven't seen anything out in the wild that's AI-driven. But it's on the horizon potentially, which isn't a great thing. But in the same vein, we'll

have AI-empowered defense tools as well. That's the other piece that, like I said, really isn't as common, but it is something we'll need to be kind of wary of.

And then we've also got just some of the more AI kind of liability and privacy pieces. With AI, the quantity of data that we're using is so vast. Because without it, the models really aren't super effective. Especially in healthcare, a lot of that data that's being used is PHI. It's relevant to personal health records. And with kind of this boom in the industry, we are starting to see a lot of vendor-driven capabilities. And sometimes they're not really as mature in the privacy and security departments as it probably should be. And so, there's always going to be risks there along with new legislation. There's new acts being proposed every day at the state level, at the federal level, and internationally as well. That's something organizations are going to have to contend with. And so, those are a few high-level examples. But, Brent, I don't know what your thoughts are on some of the specific risks or other examples you might be seeing.

Brent Hoard:

Yeah. I think kind of building from that, whether it's a malicious act or just a bad model, you do have some kind of safety risks potentially that you need to be on guard. That could be something that creates a bad clinical recommendation. Or if you're using it to design in pharma and life sciences, developing molecules, you could end up with bad or unsafe drug designs and other things. You have that safety risk that's present.

As Morgan mentioned, security risks on top of that. At the heart of AI, it's software. And it can be exploited. It can be attacked. So, you need to secure it. I think if you're using AI in an administrative type setting, you need to be on guard for a lack of resilience or a dependence on that technology. If you're doing all of your schedule using an AI tool and the tool goes out, you might be in trouble. Think about ways that you can build in some safeguards or backups. Just like you backup data, have a backup process in case things go bad.

And then, finally, for the data output, one of the focuses for regulators is around bias and discrimination in the outcome. Looking at data that is training the AI model, you can end up with certain biases or discriminations depending on what's in that model. That's why it's really important to make sure that you're using a reputable vendor. If you're using data, look at the data and be on guard. Have some checks and other balances for identifying and addressing potential bias or discrimination in the output data and your underlying data set.

Brett Mason:

And this reminds me, Morgan, what we talked about previously, that your artificial intelligence really is only as reliable as the data it is based on. And you can have those cybersecurity concerns that the data itself may be messed with and now it makes the total artificial intelligence tool worthless.

And, of course, now all these things that you guys have just run through are all of the scary things that everyone's considering that makes them want to stay far, far away from using any artificial intelligence. I'm sure that's how most lawyers feel at different organizations. But are there safeguards and controls that can help address these challenges that we've just laid out?

Morgan Hague:

Absolutely. Yes. There's a wide variety them, thankfully, right? Like we mentioned, there's a lot of scary stuff out there around AI and just system security in general. But I think one of the more interesting ones that is specific to AI to your point is what we call discrimination controls. Brent mentioned bias and models. That's becoming something that you'll see like with ChatGPT as an example. It's trying to prevent bias in some circumstances. And so, what it does is it's maybe a little bit too strong one way.

There are some interesting examples where, looking back in history, some things might not be represented in a way that's contextual. But it's because they have bias controls and they're trying to eliminate something where – especially in healthcare, certain demographics might be underrepresented from a data point perspective. And so, that results in maybe errant personal care plans. Those discrimination controls are very, very important to making sure that models are working effectively. And, again, you can kind of trust the outcomes that come from that. That's a little bit of a new discipline. It's kind of like data integrity, but a step up.

And so, I think one thing that's important for people to understand is this new class of what they call responsible AI. Similar to how privacy historically has had a lot of principles around fair use and making sure you've got proper relationships with your data subjects, and disclosures, and all of that. Responsible AI is hoping to drive to similar discipline around how people are building their AI functions.

And so, that's become a little bit of an agnostic term. But I definitely recommend taking a look into that if you are interested in AI and how you can kind of make sure that the discrimination within a model is contained and even the fair use of that model is appropriate. And so, that's one of the more specific pieces.

I think outside of that, we've been mentioning, at its core, AI is really just any – I won't say it's any system. But it's a system. It's software. It's hardware. It's a combination of those two things. Just because of the level of resource that's required, a lot of AI is being driven by what people call the cloud. People are typically going to be running these systems on their own in their own data center. There's always going to be some complication associated with that.

And so, all of the controls you would typically need to have built around security. Access controls, encryption, data backups, those all have to be in place. And then there are a couple of kind of above and beyond security controls that are really just more tailor fit to AI.

We mentioned several times the data that you're receiving is really the most important thing with the model, arguably. The pipeline to receive that data is really, really critical as a result of that. And so, typically, organizations have to – we call securing interfaces. If you're communicating with somebody, your interface is basically the telephone wire that goes between companies. If somebody can intercept that and change the data that's coming from one side to the other. If you're talking to your neighbor, they say something rude. Maybe it wasn't them saying something rude. It's kind of a labored metaphor. But it's the same thing with the data pipeline for AI.

There's a new discipline around data pipeline security that's meant to evolve that just because the level of training data that comes in and out of an organization that's using an AI model is

much more severe than typically we would see. That's something that I think a lot of organizations have to contend with.

And then I think the other kind of interesting novel capability that we're seeing is much more robust certificates or digital artifacts to maintain the authenticity of digital communication. We've seen what they call vishing. It's fishing with a V basically. Events that have started to occur. If you're familiar with deep fakes, they're very realistic, basically, videos that appear like somebody is saying something or doing something, but they're not legitimate. They've been altered.

And so, there's actually a couple of highly documented incidents where somebody was victim to these because they thought, "Hey, look, my boss told me to transfer this money to an account." Wasn't legitimate. There's a lot of interest. And then some interesting organizations coming out with different kinds of certificates and artifact, labels that will help ensure that things are legitimate.

Inter-organization communication, there would be a checkmark or something like that to make sure, "Hey, this hasn't been messed with, basically. And so, those are a little bit more of the novel use cases we're seeing from a security side. I think the only other piece that organizations always need to keep an eye on is governance at a high level from a security and privacy side.

There are some new frameworks coming out around AI. NIST has their AI risk management framework, ENISA, which is out of the European Union actually has a very, very solid framework as well. And there are new ones coming at really every month or so. But it's an organization's responsibility, obviously, to make sure that they are up to date from a governance standpoint. Their policies, and their procedures, and everything under that umbrella are all updated.

And so, that's one thing I think that's important and goes a long way for several reasons. And so, there are several other ones. And, Brent, I'd love to hear your thoughts as well.

Brent Hoard:

Yeah. I think, from a legal perspective, as a transactional lawyer, I like to write contracts. I do it all day. But in this case, contracts can be incredibly important. And so, I thought I would start there. Whether it is with a vendor who is deeply involved in creating AI models for you or it's just AI that's part of a service, having that contract in place with appropriate reps and warranties that are made regarding the data, the data sources, the rights to use that data and the way that you want to use it. Potentially, data protection agreement or data processing addendum, whatever you call that. Business associate agreement, because we're in healthcare and it might have PHI involved.

But I think, ultimately, when you start, the contract is a key component to all of this and making sure that you have the right contractual safeguards in place can help immensely. That in a way flows to my second area, which is third-party or vendor-risk management. And when you are either taking on a third-party service provider or looking for somebody to, again, store your data, you want to do some due diligence on them.

And given AI and potential volumes of data involved, it becomes even more important to know who your partner is going to be. And make sure they don't have issues. Again, get them under an appropriate contract. Understand what's going on. Possibly have an audit right or other pieces that can help you feel more comfortable that they're handling your data properly.

If you are collecting data, it's important to make sure that you have the rights to use it the way that you want to use it. There have been cases where AI models have been basically destroyed because they were based on bad data. Data that didn't have the right to use it in the way that it was to make the model. The model, the FTC, came in and made the model go away. You don't want to invest in something like that just to have it lost. Make sure you have the rights.

Some of the ways that we do, we create ethical collection and use guidelines just based around – there's FTC guidance. There's the AI regulation in the EU. As Morgan mentioned, there are a bunch of different frameworks that are out there. You can create guidelines for your organization to collect and use the data.

Potentially, if you're looking to gather a lot of data, IP and scraping type issues, if you're collecting sources on the web, making sure that's being done correctly. And then I think, finally, if we're looking at it just purely from a privacy perspective, there is a concept of a data protection impact assessment that you might be familiar with from the GDPR, CCPA sometimes. But you can really sit down and look at what is our goal for our AI tool? How are we going to use the data? How are we protecting it? And address any potential issues that come up as you analyze that.

And then you have a record that you have looked at your AI tool, AI model, whatever it is and have documented that. And you can continue updating it as things change, because things are going to change here. It's so fast. But I think those are some great starting points from a legal perspective that you can take for your organization.

Brett Mason:

Brent, I really appreciate you breaking it down from the legal perspective after we had the scare from Morgan explaining to us all of the scary things that can happen with this. And for our listeners, if you're really interested in hearing more about data collection and use guidelines, or about the regulatory activities that are occurring it seems like every week with a new regulation coming out on data privacy and cybersecurity in this space, we are going to be having some more episodes on that. We're going to be having Brent back to talk with us about best practices and what you can do to help protect your organization if you are going to be collecting data with AI software or using that data in some way in an AI, in an artificial intelligence way. Brent, I look forward to diving deeper with you on those subjects in some future episodes.

You just went over quite a variety of best practices. With the defense of artificial intelligence in mind, who would you recommend be involved from an organization when the organization is looking at AI development and implementation, Brent?

Brent Hoard:

Yeah, I think it's going to be a group effort. And I think depending on the size of the organization, you do want C-suite visibility at a minimum just understanding what you're doing.

It's new technology. You want C-suite to understand the risks. And then as part of developing, deploying the technology, the groups that we've talked about, you have legal, contracting, guidelines, policies, procedures, that type of thing, IT and information security, everything Morgan talked about. You definitely want IT and information security involved. If you have a procurement department, somebody handling the contracting, either with vendors, so that addresses contracting vendor risk management.

And then depending on your use case, I would suggest getting involvement input buy-in from your key clinical and business stakeholders, the people who are going to be using these tools. Make sure you've got it dialed in and you're doing the right thing.

Brett Mason:

Morgan, any last thoughts on this topic before we close out for today?

Morgan Hague:

I think just to kind of echo the sentiment, it's evolving so rapidly. If you do kind of take the plunge as an organization and you're looking to build a model internally or leverage some high-impact services through a vendor where you're sending data, it does take a little bit of time and investment to kind of get caught up and make sure you're informed. I think just be prepared for that. Make sure you've got a couple of champions at the organization that are leading a lot of that heavy lifting. Yeah. Again, thanks so much for having me on and appreciate it.

Brett Mason:

Well, thank you, Morgan. Thank you, Brent. And thanks to our listeners. Listeners, please don't hesitate to reach out to me if you have questions comments or have any topics that you're dying to hear us talk about from artificial intelligence and the legal impact in the healthcare space. You can reach me at brett.mason@troutman.com. You can also subscribe and like this podcast or other Troutman Pepper podcasts wherever you listen to your podcasts, including on Apple, Google and Spotify.

Thanks so much for joining us. And Brent, Morgan, I look forward to diving deeper with you guys on data privacy and cybersecurity issues in future episodes.