
The Good Bot: Introduction to The Good Bot: Decoding AI and Its Integration Into Health Care

Host: Brett Mason

Guest: Morgan Hague

Recorded: 5/1/24

Brett Mason:

Welcome to *The Good Bot*, a podcast focusing on the intersection of artificial intelligence, health care, and the law. I'm Brett Mason, your host. As a trial lawyer at Troutman Pepper, my primary focus is on litigating and trying cases for life sciences and health care companies. However, as a self-proclaimed tech enthusiast, I am also deeply fascinated by the role of technology in advancing the health care 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 the health care industry, and the legal implications that come with integrating this technology into that sector. If you need a basic understanding of what artificial intelligence technology is, and how it's being integrated into health care, you are listening to the right episode. This is our first episode of *The Good Bot* podcast, and I am very excited to have on the podcast, Morgan Hague from Meditology.

We are going to be going over the basics of artificial intelligence, some of the use cases that Morgan and his team are seeing in the health care industry currently, and where we think it's going to be going. We're going to lay a good groundwork for understanding the technology that's going to be the basis for all of our discussions on this podcast.

So, Morgan, thank you so much for being on the podcast. We're excited to have you here today. Can you introduce yourself to our audience and talk a little bit about Meditology and what your role is there?

Morgan Hague:

Absolutely. Thanks, Brett. Thanks for having me. So, Meditology is a top-ranked security and privacy firm, and we're focused on the health care space. That's really where we have a lot of our expertise to work with providers, payers, and everything really in between. We are OCR's expert witness. They come knocking. We've been there, done that. Myself, personally, kind of relevant to today's conversation, I do lead our emerging technology service line, as well as our strategic risk management service line, and I'm a contributor to a OWASP's AI exchange, and their security guide.

Like you, I'm a big tech enthusiast. That's been very interesting, seeing kind of the boom of AI since ChatGPT really burst onto the scene, at least for the general public. So, really interested in this kind of stuff and excited to talk to you today.

Brett Mason:

Fantastic. Let me ask you to back up and just tell us, you said you're a contributor to OWASP's AI exchange and security guide. Can you tell us what OWASP stands for and explain a little bit about that?

Morgan Hague:

Absolutely. So, OWASP is just an acronym, but as an organization, they are effectively one of the go-to authoritative sources. I call them an open-source organization for application security. So, if you think about any exploits, people can manipulate with code in an application or at a database level, whatever it is, they do a really great job kind of as an evangelical organization, trying to teach people the best practices when it comes from a programming perspective, or software development perspective. Even they've got kind of branches all across the system gambit.

It's a little bit more technical, but their AI exchange and the security guide, I'm happy to talk about those and would recommend anybody listening, that's kind of more on the technical side. Those are easily two of the best resources that are out there right now. If you're trying to understand how can I actually feasibly implement controls. So, the things we'll talk about a little bit today and in other episodes, to mitigate the risks associated with using AI and building AI systems. Those two alleyways out of OWASP are phenomenal. It's just a great organization. It's a little bit like an industry group that gives people insight into what can I do better? What are best practices around securing my systems and developing things in a responsible way?

Brett Mason:

Well, I'm glad we have you here. Because clearly, you are the expert on the tech side, much more of the tech enthusiast than I. I thought it's interesting that you said there's so much focus on artificial intelligence, since generative artificial intelligence really came to the forefront of public attention last year. I think you and others in the tech space know that artificial intelligence has been used in a lot of different technology over the years. This isn't a new thing. But for most of us, it's new. It's the first time we're paying attention to the use of artificial intelligence across industries.

So, can you just start us from the beginning? What is artificial intelligence? What are the definitions that are being used to discuss artificial intelligence today?

Morgan Hague:

Absolutely. Yes. So, to your point, there's a wide gambit. I like to use two definitions, basically. So, there's one that's very, very, it's kind of logical in nature. So, really, AI system artificial intelligence, and you'll see a lot of parallels with things like machine learning. That's kind of all under the umbrella. But artificial intelligence, on one hand, is really just a system that's engineered in a given way to provide a set of objectives, and it's going to give you an output and it can improve that output time over time. That's kind of where you'll hear people talking about learning, training data for AI effectively. It's in the most basic sense. I've got a question I need an answer. AI as a principle should be able to help you reach that answer without you having to do everything kind of in between.

It's not as A to B, as a lot of typical processes are. AI can kind of take you from A to Z effectively is what we like to say. Just kind of a more academic definition is, and this is kind of what I call it, the James Cameron version of AI in the future. Maybe we'll get there. It's kind of the branch of computer science that really is meant to intelligently or I guess, drive human intelligence or potentially beyond human intelligence from a system. So, everything that we do, we wake up, and we have our specific thoughts, we have dreams, everything. There are people that are really on the computer science side of it, trying to kind of mimic that from a system perspective. There's a wide umbrella there. But I think those are the two common threads that people will pull on.

Brett Mason:

Talking about the technology that's already being used today, and what we're seeing being developed right now, does that fall more in that first camp of the AI systems where you have a set of data, the AI is taught to generate certain outputs, and then it can use that data to come up with solutions and answers. Is that more what we're seeing right now?

Morgan Hague:

Absolutely, yes. So, the big boom that we'll see is what people call large language models, or LLMs. So, ChatGPT kind of falls into that category. I think the easiest way to define it is, the more what we call training data, that a system is going to get, the closer and closer as long as you kind of give it feedback, and you give it prompts to the desired result that it can get. If we're using ChatGPT, as an example, you see, they're different versions. You have GPT-4, GPT-3, whatever it is. Each iterative version of that model, and kind of the engine feeds that along with the data, gets it closer and closer to kind of replicating deliverables that somebody would actually create. If you wanted to give it a prompt and say, "Hey, look, what are the top 10 cars being sold in 2024?" The first version of it might have just made several mistakes, and it doesn't write in a way that's consistent with what a human right. But as it's able to kind of take in more information, it gets closer and closer.

So, that's how most AI models work now. Same thing with pictures, you see things like Sora, that can make videos and pictures, it does the same kind of piece it. It makes assumptions kind of about what should be there based on what it's seeing.

Brett Mason:

I hear that ChatGPT-4 just passed the bar exam. I know, myself as an attorney, and others are getting a little bit nervous. But at the end of the day, there's a lot that I think artificial intelligence is going to do to help us be more efficient and take away routine tasks. Are we already seeing different health care institutions use artificial intelligence software?

Morgan Hague:

Absolutely, yes. So, there are a couple of use cases and I know we'll probably break these down a little bit. But one of them I'll like to pick on is medical imaging analysis. So, that's kind of a fancy way of defining, if you go in and get a sonogram, or a CAT scan, those kinds of gray, white, black images that you see coming up on a screen, even since the eighties, those have been using AI or versions of that.

In the health care space, at least, there's definitely been a use case for that. And that kind of fits the bill, like I mentioned, the system is trying to infer what it's going to see. You can't literally just kind of see through the brain in the way that kind of an x-ray can. It's a little bit different, just like in the womb, right? We don't want to expose kind of as sensitive areas, radiation, and things like that.

Those systems basically infer what they're going to see, and then they, years and years get better and better and better. That's kind of a popular anecdote. If you went and got an ultrasound just 20 years ago, what you're seeing today, if you do like one of those 3D scans is lightyears ahead of what they used to have, and that's really just powered by AI. So, that's kind of one use case. Long story short, several different ones, or several different examples, organizations are using, and there's a significant investment there, too.

I think in 2023, Stanford did a study and it was like roughly six billion or so dollars in the health care space for AI and that was actually the leading industry for investment. I'll tell you anecdotally, too, I'm getting a lot of interest from our clients in AI, how can they secure their AI systems, and we can talk a little bit about that, but it definitely is very prominent.

Brett Mason:

Going back to what you just talked about the medical imaging analysis. In looking at what the FDA is doing around approving medical devices that have artificial intelligence software incorporated, the vast majority of those are what they would call diagnostic tools that use medical imaging. From my understanding, and you can tell me if this is correct or not, that currently is one of the best examples of how artificial intelligence can be used. You have a clear dataset, for example, ultrasound images of pregnant women, and you have a clear objective when we compare this ultrasound image for this patient to this large dataset that we know is reliable. Can we identify that there are any issues? Can we identify if we're seeing anything that needs to be addressed by the doctor?

Clearly, that's the technology that is being used the most in medical devices currently, as the vast majority of products the FDA has approved, really are medical imaging analysis or diagnostic tools. In addition to that, what other ways are we seeing health care industry individuals or companies use artificial intelligence?

Morgan Hague:

One of the more interesting pieces that we're seeing, and personally, I think one of the more sci-fi-esque use cases that we see is what they call drug development and discovery, right? So, that as a discipline exists. It's basically if I'm a company that's developing a new drug. During COVID-19, there's a big pandemic that kind of came out of nowhere and I need to find an answer quickly. AI now is being used to streamline that in a major way. So, it basically takes all of the data that we have from the pharmaceutical company. AstraZeneca has made huge investments. Really, anybody in that space is making material investments in AI modeling to take all of the information they have, take that data, and literally at a molecular level, try and forecast what needs to go into the drugs that they're developing. What needs to go into these compounds to react and respond to viruses, bacteria, whatever the illness might be, wherever the source might be.

So, it's reducing cost, or it's anticipated, at least, it's still fairly early. To be honest, it's anticipated to reduce costs. One of the major benefits outside of that it's just time to deliver. You're trying to test. So, there's obviously going to have to be trials that are conducted with any kind of drug. But the big timesaving that you'll see with a lot of this AI development is, how long does it take us to get to that point where we feel comfortable executing some of those trials. It's really, really interesting. That's a little bit more on the pharmaceutical side, but it obviously has implications really everywhere. Insurance companies are implicated; providers, obviously. That's one use case, I think, is very, very compelling and unique.

Brett Mason:

Yes. We're trying to keep these podcast episodes a little bit short, so that we don't bore everyone. I know we have so much more we could talk about. But in addition to those use cases, are we also seeing companies looking to incorporate artificial intelligence software into personalized medicine, remote patient monitoring, robotic-assisted surgeries, clinical decision support, and other types of diagnostic or health care provider support technology?

Morgan Hague:

Absolutely, yes. All of this. So, I think to kind of put them under one category. I think in every company, especially providers, they are looking to increase contact with patients to try and make the experience a little bit better, and make sure that they're getting as much care as they can just because resourcing can be a little bit difficult. So, you will see things you mentioned, like the remote patient monitoring, that goes a long way in making sure that they can keep visibility in terms of how the patient is doing and make sure that treatment plans are going accordingly.

Then, another kind of interesting use case on the same vein is clinical decision support, right? Treatment plans right now are going to be typically derived by your physician or by the NP, whoever's kind of helping you with your treatment case. Clinical decision support provides a really interesting use case where they can pull from these massive pools of health data that we have and say, "Hey, look, if somebody fits a very similar profile, the demographics line up, and they want the similar outcome, what is something that I can use to inform my plan that I'm developing for my patient?" Hopefully, that pulls on from different regions.

So, there will be underfunded providers, certainly in different regions of the world that they haven't seen something basically, as frequently as maybe providers have in large cities, that goes a long way in that country that they can see positive outcomes. That's a big theme, though, I think that we're seeing is people want to be able to, or providers want to be able to maximize kind of the care and the value at which they're delivering that care. In the same vein, you see insurance companies, payers, using AI for things like revenue cycle management. That goes a long way in making sure that they can process these things as efficiently as possible, they can even forecast some decisions in terms of whether certain claims are going to be approved. Hopefully, the idea is that it speeds up delivery of care. We know there has been a little bit of stress on our health care system since COVID, really, and I think everybody is looking for a little bit of a relief valve. So, those organizations are looking to leverage AI, for that, and several other outcomes, but it's a lot of compelling use cases for sure.

Brett Mason:

Now, would you agree with me, there are just tons and tons of different ways that artificial intelligence is going to be incorporated into health care as we advance?

Morgan Hague:

Absolutely, yes. Really, the sky's the limit, and it's funny because it's been so exponential since the boom of ChatGPT a year or two ago and it's kind of interesting. That's really the switch that flipped everything. But really, just at the beginning, I think there's so many things that we'll see impact on so many levels in the health care space.

Brett Mason:

I'm hoping that a lot of our listeners, this podcast will be kind of more on the legal side. So, that's why we love having folks like you who have more of the technology and can let us in on how this is all being done. You've just talked about a bunch of different ways that AI is going to be used in health care. Are there different vendors and different ways that the AI itself is going to be incorporated?

Morgan Hague:

Absolutely, yes. That's honestly what I would say is significantly more common in terms of how people are actually going to use AI. So, I'll pick on like AWS. Amazon has what they call HealthScribe, right? If I'm going to an ophthalmologist, and I'm getting a new prescription for my glasses, whatever it is, typically, there's going to be somebody in the office there that's writing down whatever the doctor says. HealthScribe, in effect, replaces the need to have a scribe in that environment. That's one kind of example of a vendor, providing a service that a lot of providers are kind of scrambling to offer internally. For them, it's a little bit of a cost savings.

But that's the delivery model we're going to see most common. It's going to be a vendor providing some sort of a service. A lot of that is driven by the cost-prohibitive nature of AI. If you're a kind of a small shop, and you want to build this data model, data is kind of gold from an AI perspective. So, it's really difficult to build a really effective model unless you kind of have that funding at baseline data that's required. That's a lot of what we'll see several different examples, anything you can think of, really in the health care space and beyond. I think vendors are going to drive a lot of that adoption. It's not going to be as common that we would have an organization building their own. We call it data science function internally to build that AI.

Brett Mason:

So, just so we're clear what we're talking about, what we're saying here, when we're talking about something like Amazon's HealthScribe, that is a technology that would be an external tool that is going to be used by a health care entity, like a provider, for an internal goal. For example, the internal goal would be to streamline taking notes when meeting with patients, uploading those notes into medical records, getting responses to patients' questions, especially if they're online questions. So, you've got that external tool being used for an internal goal. That's a little bit different than what you were just talking about, which is health care entities that are building

their own internal tools that leverage artificial intelligence software. Am I understanding that right?

Morgan Hague:

Exactly, yes. So, as you start talking about kind of the risks, and even the legal side of this, too, it's important to kind of note those distinctions. If Amazon is building this huge data warehouse, and they're the ones that are kind of applying these models, they're set to be a lot of boundaries in terms of how they can manage some of that data that they're effectively beholden to. As the service user, you might not be, and so most organizations probably won't fall into some of those categories from a legal perspective, from a security control perspective.

I think one other kind of use case that most organizations are probably fitting in now too, is kind of a non-sanctioned use of a tool, right? There could be that you're using a tool that's key to your organization. It's aligned with business objectives. But it's an external tool. It doesn't make sense for you to try and do that internally. A lot of times, what companies are dealing with now and trying to get around with, or trying to get around is Joe Shmoe is using ChatGPT to write all of his executive summaries. It could happen.

Brett Mason:

I'm sure it is happening.

Morgan Hague:

Yes, exactly. It could help him get to a certain outcome. The reality is, there's a lot of gray area in terms of where that data is actually going. So, you're putting a lot of trust into the tools. GPT is an easy example to pick on. They are backed by some large organizations. So, the risk there is maybe a little bit lesser, but everybody now is coming out with AI tools. You'd have to provide your data for them to work really in the way that you want them to. So, there's a big third-party risk, is what we call it, associated with that and that's interesting. That's something we're seeing a lot of companies come to us and say, "Hey, look, do you guys have any tips on policies and procedures to deal with that?" Because it's kind of the Wild West a little bit right now with so many AI tools out there.

Brett Mason:

I am excited to let everyone know that we are going to be having Morgan back on another episode very soon to talk specifically about data privacy and cybersecurity issues. Both if you're using an external tool, like we've been talking about, or if your company is building an internal tool that's leveraging artificial intelligence. So, we're not going to get too deep in that right now. But if that is something you're really excited about, we are going to be talking about that soon. So, keep an eye out for that episode.

Now, Morgan, a couple other different things I wanted to touch on since this is our episode to try to understand and explain the technology, at least where we're at right now. When you are an organization, and you're thinking about whether you want to use an external tool, whether you want to build your own, what are some important things to be thinking about in that consideration?

Morgan Hague:

I think really, for most organizations, it just comes down to cost and I think, purpose. So, if you don't have the data available to make the model, honestly, that should be rule number one. I think that's something that we're going to have a lot of organizations struggling with. If you don't have proper data, your model is effectively worthless. You could spend millions and millions of dollars if you're just making up data points to kind of test that the model will put out something, you basically start with all of that money, which is something that a lot of organizations are going to kind of reconcile with, which is tough because the capabilities are so compelling, and they're so interesting, but that's just kind of the reality we have to contend with.

I think outside of that, when you're dealing with a third party that's going to provide you a tool with a key service for your organization, just make sure you're doing the proper checks. I think, make sure that it's not some really infantile organization that doesn't have some of the basic needs you'd expect to make sure the company is going to exist, because there's always some risk with them kind of going out of business as soon as you have a relationship, and you deliver your data to them. I think really, those are the two major things, honestly. And it is interesting to think about how this boom is going to continue as organizations to try and build their own functions internally.

Brett Mason:

Yes. I think so many people have talked about ChatGPT giving results that are completely bogus, or are extremely skewed, and that goes back to what you're saying. The strength of the artificial intelligence tool you're using, is really only as strong as the data that it is pulling from. We know that ChatGPT is pulling from the entire Internet, and we know how much ridiculous things are posted on the entire Internet. So of course, that's the data it's pulling from. It's not going to be reliable.

So, I do agree with you, it's very important for any of these organizations, if they're looking to use an external tool. They need to have a good understanding of what is the data that that vendor has incorporated into its tool to ensure that the artificial intelligence is actually pulling from reliable, accurate data. One of the things I like to think about is if you are a physician, and you want to review the peer-reviewed medical literature, to understand about a new issue that's come to light, you would look at highly reliable and authoritative medical journals. You wouldn't go look at what high school seniors have been writing about potential diseases that are coming out. That's a different data set.

We do see a lot in the news today, a lot of press releases about different companies that are launching AI tools. But what are they relying on? What data are they using? How are they really making sure that this is an algorithm that is going to be authoritative? And that's a really important consideration.

So, Morgan, you've been doing a lot of speaking and educating on the use of artificial intelligence in health care, which is one of the reasons we wanted to have you on the podcast. In that, have you been able to have conversations with folks who are attending those and see how different organizations are actually using AI right now?

Morgan Hague:

Absolutely, yes. I think, again, it's so varied. But I will say a lot of organizations are starting to lean into outside of medical imaging, which to your point, medical devices have started to really explode on that end. A lot of the functions are starting off and more kind of the analytics space.

So, I mentioned data science a little bit earlier. A lot of companies are starting to build their own to large providers, especially that have the resources. What they're using that for is kind of twofold. If some of it is going to be in the patient outcomes front. So, for them, it's a lot of, "Hey, look. Let me figure out how we can deliver treatment in a more effective way, both in terms of cost and in terms of patient outcomes and experience." Then on the other end, business use cases are really driving a lot of value there, too.

Let's say I wanted to check in on a claim that I have that needed to be processed. Right now, you probably need a person to be able to do that for you. A lot of companies are exploring having AI called chatbots, whatever it is, kind of interfaces to provide consistent contact with patients as they're working through that. I think those are the probably two most common areas that I'm seeing with a lot of the people that I work with. Data analytics, which is really just the next step. You hear people talk about big data, data science. AI is kind of the next step in that evolution. Then, AI, which in a similar vein, you hear a lot of people talking about automation over the last decade. That's really the next stage there.

Brett Mason:

It's fascinating that you're already seeing so many organizations use this technology. But that's why we wanted to have this podcast. So, we're winding down on this foundational episode. I really hope everyone has gotten a good understanding of what the AI software is and how it's currently being used in health care. Additionally, some of the future technology that we're seeing may come into play in the next years as the technology advances. So, I want to wind down and just ask you, Morgan, one last thing. What's one thing you would want to make sure our listeners know about the use of artificial intelligence tools in health care?

Morgan Hague:

I think, honestly, and this is kind of simple, but have a plan. This is something that is easier said than done sometimes, because you have a lot of competing interests, right? Like I mentioned, the data team, they're going to love AI in most circumstances. They are going to be like, "Hey, let's get whatever we can. Let's push the agenda." The security grouch, as I like to say, is going to be like, "Please don't do that." So, there needs to be consensus of the leadership team in a lot of ways, and the organization needs to take a stance. The first step of that could just be an acceptable AI use policy that says, "Hey, look, please don't use ChatGPT for everything. Use it maybe for these specific use cases or don't use it at all." Whatever that is, have a plan, make sure it's documented somewhere, and that helps to hopefully curb some behaviors and start to drive the culture of the organization around how people are using AI.

Brett Mason:

I think most people would point to someone in your role as the AI enthusiast, as the IT manager, and then point to the lawyers of the organization as the security grouch. So, I hope on this

podcast with the lawyers, we're going to have on to talk about artificial intelligence and the legal impacts, we can dispel some of the grouchiness of our fellow lawyers. But hopefully, they will be more excited to talk with you about AI's possibilities once they get some information from us on how they can do so safely and effectively.

Morgan, thanks so much for being on. I've really enjoyed chatting with you.

Morgan Hague:

Thank you so much. It's been a pleasure.

Brett Mason:

I know we're going to talk more on some other episodes, but I also want to thank our listeners for tuning in to this first episode of *The Good Bot*. Please don't hesitate to reach out to me, Brett Mason. You can email me at brett.mason@troutman.com with any questions, comments, or topic suggestions. I also will take your complaints. This is new to me, so I'm excited to hear how we can make sure that this podcast is extremely valuable to you.

You can also subscribe and listen to this podcast and other Troutman Pepper podcasts wherever you listen to podcasts, including on Apple, Google, and Spotify. We look forward to bringing you more information about artificial intelligence technology and the legal implications of integrating that technology into the health care space. Thanks for joining us.

Copyright, Troutman Pepper Hamilton Sanders LLP. These recorded materials are designed for educational purposes only. This podcast is not legal advice and does not create an attorney-client relationship. The views and opinions expressed in this podcast are solely those of the individual participants. Troutman Pepper does not make any representations or warranties, express or implied, regarding the contents of this podcast. Information on previous case results does not guarantee a similar future result. Users of this podcast may save and use the podcast only for personal or other non-commercial, educational purposes. No other use, including, without limitation, reproduction, retransmission or editing of this podcast may be made without the prior written permission of Troutman Pepper. If you have any questions, please contact us at troutman.com.