

The Good Bot: Artificial Intelligence, Health Care, and the Law — Colorado AI Act

Host: Brett Mason

Guests: Erin Whaley and Emma Trivax

Recorded: December 5, 2024

Aired: February 20, 2025.

Brett Mason:

Welcome to *The Good Bot*, a podcast focusing on the intersection of artificial intelligence, healthcare, and the law. I'm Brett Mason, your host. As a Trial Lawyer at Troutman Pepper Locke, my primary focus is on litigating and trying cases for life sciences and healthcare companies. However, as a self-proclaimed tech enthusiast, I'm also deeply fascinated by the role of technology in advancing in the healthcare industry.

Our mission with this podcast, is to equip to 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's being integrated into the healthcare, I recommend you start with our inaugural episode. There we lay the groundwork for understanding the technology that is the basis for all of our discussions.

I'm really excited to welcome to the podcast today two of my colleagues from Troutman Pepper Locke, Erin Whaley and Emma Trivax. We're going to have an exciting discussion today talking about some of the regulatory and legislative action that's occurring around artificial intelligence. So, Emma, Erin, glad to have you guys here. Thanks for joining me.

Emma Trivax:

Glad to be here. Thank you.

Erin Whaley:

Thanks for having us.

Brett Mason:

So, Erin, could you start us off by giving us an overview of how AI is currently being utilized in the healthcare industry?

Erin Whaley:

I'd be happy to. One of the things that folks don't realize is that AI has been used in healthcare for quite a while. When we think about AI in healthcare, we have seen it used in revenue cycle management. So, it has been used to run billing analytics, help improve billing platforms. We have also seen it for years on the clinical side in electronic medical records in the form of clinical

decision support tools, or in the form of drug-drug interaction alerts when a provider goes to prescribe a new drug or an allergy alert.

What is really getting a lot of attention and fanfare and appropriately so is the use of generative AI in healthcare, and that has numerous applications. One of the places we are seeing it a lot is in virtual scribes. So, AI that can listen to an interaction between a provider and a patient and create a patient note from that or help the provider to draft a communication to the patient as a follow-up.

There are a lot of other exciting applications of that, including predictive analytics, using various readmission risk scores, trying to predict where gaps in care need to be filled in order to improve health. We see it on the pharmaceutical side with respect to research and potential new drug discovery. We see it in imaging. And can AI read an image better than a clinician can? So, we are continuing to see that evolution, and we're really only at the beginning.

Brett Mason:

Well, Erin, thank you so much for that rundown. I love hearing you say that AI has been used, it's been in healthcare, this is nothing new. And that's something we're seeing across industries with artificial intelligence. It has been incorporated into industries that this Generative AI, the ability to create this content is the exciting new area that we're seeing. I actually heard this past week, an interesting statistic. The FDA has actually only approved around a thousand medical devices that have an AI-enabled component from 1995 until 2024. But the exciting thing is everything that's coming to us now, the things that we're looking forward to and that we're looking ahead to. So, thanks for that rundown.

Emma, what are some of the risks associated with integrating artificial intelligence into healthcare and care for patients?

Emma Trivax:

Yes. I don't want to fearmonger or make anyone feel uncomfortable, but there certainly are risks, just as there are plenty of benefits. But the risks, one of the main concerns that we are seeing is patient safety, right? So, AI systems, they are flawed, they can make errors, and when it comes to healthcare, sometimes these errors can have serious consequences. If, for instance, there is some sort of transcription software that transcribes incorrectly into a patient's chart and then a doctor later on relies on that incorrect transcription, that could mean improper healthcare as a result.

Another issue that we see is actually bias. So, AI systems, they are taught a limited set of things and depending on what they're taught, it can lead to creating certain biases, right? So, if there are the certain amount of patients, a certain subset of patients that the AI is relying on for its data and then a new patient comes in who's different than other patients, the data that it's relying on might not be correct, right? It is biased in any way that an AI system can be biased. So, it leads to these disparities in treatment and outcomes.

Last, there's this really big challenge we're seeing of accountability. Who is responsible when an AI system makes a mistake? Is it the healthcare provider? Is it the AI developer? We are seeing many medical malpractice insurance companies creating policies, either never permitting AI to be covered under their policies, certain use cases permitted. So, there is a massive question we're trying to untangle right now as to where does the responsibility lie when AI is utilized?

Brett Mason:

That's a difficult question that's going to continue to be an area of controversy, I think moving forward, Emma. I always think it's interesting when we're talking about how accurate the AI is or could be and what if it makes mistakes when we know that human beings make mistakes and perhaps the AI, while it may not be perfect, it could be more accurate than human beings. So, I think it'll be interesting as more technology is adopted and we can actually run some studies and look at accuracy of AI versus accuracy of humans manually inputting, for example, the different medical information and a medical record.

So, Emma, you and I have talked a lot in the past couple months about different things, the complexity of AI integration in the healthcare industry. We've talked about the newly enacted legislation in Colorado, which is kind of being looked to as a model or a potential playbook for other states in thinking about how they want to deal with AI from a legislative perspective. So, can you just talk to us about what are some of the main takeaways initially from the Colorado Act?

Emma Trivax:

Yes, that's a great question. So, the Colorado Act, it's called the Colorado Artificial Intelligence Act, or I keep calling it CAIA, it's an easy way to remember it. It's a comprehensive law that focuses on the development and deployment of high-risk AI systems. What it really does is it's trying to regulate AI to prevent what I just mentioned that algorithmic discrimination or bias, and it promotes transparency and accountability for patients because one of the main focuses here within that act, within CAIA, was informed consent and transparency.

I think a lot of states are looking towards this because the question of informed consent has come up a lot, as well as other questions like ethical consideration. So, this CAIA, this Colorado Act is a great framework for other states to start building off of like you just mentioned, Brett.

Brett Mason:

I know one of the things that they talk about in that Act is they talk about defining different types of artificial intelligence systems. My understanding is that there is a definition for a high-risk artificial intelligence system under the law. Erin, I wonder if you could just talk us through, how are they defining and what qualifies as a high-risk artificial intelligence system under the Colorado law?

Erin Whaley:

The Colorado law is defining that high-risk AI system as a system that makes or is a substantial factor in making consequential decisions. It's a really interesting definition because there is a lot of room for interpretation there.

Brett Mason:

Absolutely.

Erin Whaley:

If we think about the example of an AI scribe in healthcare and Emma's example of what if that AI scribe mistranscribes something that then leads to errors and reliance and incorrect information down the line, a lot of what we're seeing is those AI scribes, those assistants in healthcare simply provided the starting quote for the clinician. But it remains the clinician's responsibility to confirm that the information that they are capturing in the EMR is correct.

So, it is not acting independently to create the final documentation. It is creating a draft, whether it is a draft of the office note or a procedure note for a communication to a patient. It is simply that. It is a starting point; it is a draft. So, I think some would argue that that would not qualify as a high-risk AI system because you do have that human intervention, and at least currently, we are seeing a lot of that human intervention. So, I personally wonder how many of the current systems in healthcare would meet this high-risk definition.

Brett Mason:

Okay. Yes, that makes sense. I mean, again, a lot of what we're talking about right I was looking forward, we're looking at technology that's being developed right now and thinking about how it's going to be used. This definition of high risk could be somewhat anticipating different products that are going to be coming out, I would think. What do you think, Erin?

Erin Whaley:

I think it could be. I also think that the stated goal, one of the stated goals of the law to prevent algorithmic discrimination is a really tough one for healthcare. I think that is what we are all aiming for. We all agree that is the goal, but it is particularly hard to capture in healthcare because so many of the models are trained on existing healthcare data, which many would argue has implicit biases and represents historic discrimination and issues. So, I think that goal is a good one, we all support it, but it will be difficult to measure and define whether they are actually meeting that goal with the law.

Brett Mason:

So, Erin, I appreciate you pointing that out, that the issue with bias is even more heightened in the healthcare setting, giving the historical patterns of data that's available. And so perhaps

that's why they wanted to make sure include high risk as an area that needed to be looked at and regulated a little bit more closely.

Emma, let me ask you about this. One of the things that we talked about is that there's developing and then there's deploying these high-risk systems. That is something that they distinguish between in the Colorado Act. Can you talk about that and what exactly they mean by developing and deploying and how they made that distinguishing a factor?

Emma Trivax:

Yes. Before I jump to that, I wanted to go back to what Erin was just talking about for a second as I use this a lot just in conversation, but when you're thinking high risk versus not high risk, something I like to think of as AI. Okay, it's artificial intelligence, but really if you apply it like assistive intelligence or augmented intelligence, it's not the final intelligence, it's just something that's helping you. I find that to be a very helpful perspective when trying to figure out if this is high risk or not. Is it just assisting you or is it actually artificial substituting? Sorry for that tangent, but –

Brett Mason:

No, I appreciate that. That makes sense. I think that's helpful because sometimes when we call something artificial intelligence, we're not being specific enough as to what the product is actually doing. And so, what you're saying there is a way to think about it a little differently.

Emma Trivax:

Yes. To answer your other question about developers and deployers, I think the way Colorado is defining it kind of goes back to that point. We talked about earlier regarding who's responsible for things, who's responsible for this AI that is being implemented. So, what Colorado does is they have these two terms. The first would be a deployer. This is a person, usually it would be the individual working in the health system who is doing business that deploys or uses that high-risk AI system that Erin just defined, and then the developer would be the person doing business in Colorado that actually develops or takes an existing AI system and substantially modifies that system.

So, really, it's the users versus the vendors or the developers, or the deployers versus the developers.

Brett Mason:

Interesting, and I think one of the things that the Colorado law does not make clear, and perhaps, you can correct me if I'm wrong about this, but there are still a lot of questions about who is responsible and when, even though you might have that definition of deployer versus developer. The question of who and when is still one that's up in the air. Would you agree, Emma?

Emma Trivax:

Completely, would agree. There are different rules for the developers and deployers and what they can and can't do, but at the end of the day, if there's an injury resulting from the AI that is used, we still don't fully know exactly what you said, who is responsible, and when.

Brett Mason:

Does the law break down, and perhaps Erin, you can jump in here, does the law set different standards for the developers and deployers as to what they're expected to conform to once the legislation does take effect, which I don't think we've said, but it should be in May of 2026, correct?

Erin Whaley:

Correct. Yes, Brett, they do set different obligations for developers and employers. So, I'll start with the developer side and then let Emma chime in on the deployer side.

On the developer side, they are looking for the developers to avoid the algorithmic discrimination that we talked about, to make sure that developers provide detailed documentation on the AI systems, performance, data governance, and measures. Those are all the measures to mitigate discrimination risks. Developers have to disclose within 90 days any known or foreseeable risks of algorithmic discrimination to both their deployers and the Colorado Attorney General.

Then the developers are responsible for making public statements that summarize the types of high-risk AI systems that they've developed, how to manage the risks associated with those systems, and then details about the information collected and used and they have to keep that statement current.

What I think all of these obligations have in common is transparency and they're really aimed to make sure that deployers and the public, for that matter, understand what the AI system is, understand how that AI system was trained, where the risks are, so that those can be mitigated. It's interesting. I was recently at a conference, and they were talking a lot about the importance of transparency around AI models. I was listening to the head of the Coalition for Healthcare AI talk about the fact that they are releasing what they're called, what they're likening to a nutrition label for AI.

Brett Mason:

Wow. That's, wow.

Erin Whaley:

Yes. It's really interesting and exciting to help create a standard format for AI developers to communicate the really important details of their AI models so that everybody knows what they're dealing with.

Brett Mason:

I can already see in my area, pharmaceutical and medical device products, some people likening these to warnings, these nutrition labels.

Erin Whaley:

I think it is similar to the warnings and making sure that there is that full transparency as to the benefits, the risks, the development, the data sets that the models were trained on, so that folks can determine whether it's a good fit for their populations, whether they're going to need to tweak it for their populations, and make sure that everybody is going into this, eyes wide open.

Brett Mason:

Wow, that's really fascinating, especially considering that the developers have to provide that information both to the deployers, so those who will be using their systems, but also to the public who may be impacted by those systems. So, it'll be interesting to see how companies address that and what type of statements they put together.

Emma, why don't you talk to us now about the deployers? What does the Colorado Act require for deployers at these AI systems?

Emma Trivax:

Yes. So, it's almost twofold here. So first, they're requiring certain notifications be made. Then on the flip side, they also have different ways to implement policies about AI.

First, I'll talk about those notifications. To the idea of transparency and consent, they are going to be required to make pre-decision notifications to the consumers. If one of those high-risk AI systems are used, they must be informing the patients or the consumers before any kind of consequential decision is made using that system. Then also, within these disclosures and notifications, there's just a general information disclosure that deployers have to give to these consumers just essentially saying why we're using this AI system, the type of decision that will be made, various contact information for the vendors and for themselves, and then a plain language description of that AI system. Maybe that'll be that nutrition label, 10 % of your daily use of AI.

Then the last kind of notification that has to be made are adverse decisions. So, if a consequential decision or a healthcare outcome comes out that is adverse or negative for the consumer or patient, the deployers have to disclose the reasons for that decision, what data was used, what sources were used for that data, and then consumers can correct information like their personal data that maybe would have impacted the consequential decision. So, these are all just different ways that the deployers are talking to and working with the consumers to make sure that AI is being used responsibly and correctly and that they know it is happening, because it would be bad if some high-risk AI system was used and the patient or the consumer did not know about it.

That's on the front side with the notifications. Then on the flip side, we have the policies. And how do we actually implement these systems? Because deployers, particularly in the healthcare field, healthcare providers, we have to make sure that policies are always very clear, written out in advance. The deployers have to implement risk management policies and programs to govern these deployments of these high-risk systems, and again, like Erin was saying with the vendor or with the developers, they will have to disclose any kind of risks or known iterations of algorithmic discrimination, and the policy has to be reviewed, and it has to be considered against various guidance and risk management frameworks. But ultimately, as long as that deployer is using a policy that meets these requirements that meets this Colorado statute.

Brett Mason:

Thank you for bringing that down. I find it fascinating that the law is requiring deployers to provide plain language descriptions of the AI systems. I'm curious if you think that deployers and developers will end up working together to make sure that whatever language the deployer is including in their information disclosure is first of all, accurate, because I know a lot of doctors probably would have difficulty describing different technologies that they use.

Second of all, that it provides the amount of language that's needed. What do you think, Emma? Do you see there being some coordination between employers and developers over these information disclosure and transparency requirements?

Emma Trivax:

I think in an ideal world, yes. They will work together and come up with these great plain language descriptions that would match what the deployer and the developer think this AI is. I think in reality, it will be a bit more difficult unless the developers are explicitly required to provide the deployers with this language because, like you just mentioned, sometimes it can be hard for physicians to even take the time where they need to sit down and think, "How do we explain this in plain language?" Attorneys sometimes have to get involved if the developers don't give it to them. So, in an ideal world, yes. In reality, they probably won't be working together as much as we'd like.

Brett Mason:

And I would say as a products litigator, the potential failure to warn implications for all of these are already scaring me. So, let's talk a little bit more about moving forward. As AI is consistently evolving, Emma, what are some of the policies that have to be continuously monitored and reviewed to keep up with technological advancements?

Emma Trivax:

Yes, that's a great question. There has to be an assessment done of these policies and the systems at least annually or before you deploy it and then at least annually after. Then an assessment is also required additionally within 90 days, or any substantial modification is done to the AI system. So basically, you've got kind of three requirements, before deployment, annually, and then 90 days after any big change. These assessments, it's full thumb and it

covers everything, but some of the things would be the purpose, use cases, analyzing risks of algorithmic discrimination, and how to mitigate that. The type of data categories processed, and the outputs produced.

So, there's a litany of things that are reviewed, but ultimately, yes, at least annually. And then, deployers have to maintain this impact assessment and all those records for three years after the final deployment of the AI system.

Brett Mason:

Emma, I find this interesting because a lot of AI systems, the point of the AI is that it's taught to be continuously learning. So, if you have a system that's learning based on information that it is obtaining just through use of the system, how are we going to know if there's been a substantial modification to the AI system, if it's learning and growing in and of itself such that that 90-day notification time period begins? I think that's one of the things that's really interesting about these systems is that we maybe don't always have a red flag, "Hey, a substantial change has been made." Do you think that there's been consideration of that when they were making these requirements?

Emma Trivax:

So, I do think that this 90-day requirement could be problematic, but it goes back to, it's very subjective. What does a substantial modification mean? The deployer could have an idea of what it is, is just learning as AI goes, is that substantial after a certain amount of time, or does it have to be an actual, you go in, make a change to the system, and then that's substantial. We don't have an answer, and the Colorado law doesn't define that for us.

Brett Mason:

So basically, get with your lawyers?

Emma Trivax:

Absolutely, Brett. And Brett, I would say that large academic medical centers that have implemented various AI algorithms have dedicated and are dedicating a lot of resources to queuing those algorithms in that AI. So, they are constantly looking at things like drift. They are looking at outcomes and trying to measure that. Really, that's all in the name of patient safety and making sure that there is good output from that AI.

At this conference I was at, I heard the term algorithmovigilance. So, similar to pharmacovigilance that we see on the pharma side, looking at that on the AI side, which I thought was just a fascinating concept, and I'm sure we'll see a lot more of.

Brett Mason:

Absolutely. Well, with all of these requirements in the law, Erin, can you talk about, are there exemptions that the law contains for deployers or more developers?

Erin Whaley:

Yes. Like every good rule, there are exceptions. Whether the exceptions, whether you qualify for an exemption requires you to really examine those and examine your business. But generally, at a high level, there are exemptions for deployers with fewer than 50 employees, those who don't use their own data to train the models. That's when the AI system is used for its intended purpose.

There are also exemptions for developer and deployer systems that are approved by federal agencies and entities that are subject to existing laws like HIPAA, which may cover a lot of our standard health care, traditional healthcare providers, as long as they are not using an AI system that would be considered high risk.

Brett Mason:

That makes sense. So, if they're using a system just for scheduling or for just simple kind of maintenance or administrative things, that wouldn't fall under all of these requirements, right?

Erin Whaley:

Arguably, yes.

Brett Mason:

Arguably, yes, of course.

Emma Trivax:

That is how we are interpreting it currently, but the devil will be in the details, and we'll have to see how the interpretations play out once the law goes into effect.

Brett Mason:

So, you just mentioned, and as you said, the devil's in the details, it always is with this type of legislation, but you mentioned that if the HIPAA-covered entity is using a system that's not considered high risk, they would be exempt. So, Emma, can you talk to us about what types of AI systems used by HIPAA-covered entities might be considered high risk?

Emma Trivax:

So, the key here is might be.

Brett Mason:

Of course.

Emma Trivax:

So, this is our interpretation. But things that we think could be considered high risk here would be things like actual disease diagnosis, for example, or clinical outcome predictions, coverage determinations because that can be very detrimental if you get it wrong, imaging. I think Erin mentioned that earlier, how a lot of AI systems now can review images and give diagnoses based on that. So, things in that nature, right? Stuff that would directly impact the health or treatment of a patient.

Brett Mason:

On the flip side, what would be something that would be considered a non-high-risk AI system use in health care that might meet these exemptions?

Emma Trivax:

Yes. So, clinical documentation, note-taking, billing, or the one we see most often is appointment scheduling. So, all these kind of things that just kind of keep the background of the business running would typically be administrative in nature and thus not high risk.

Brett Mason:

Obviously, we've talked a lot about the Colorado law today. What are key takeaways we can learn from it and apply on a broader scale?

Emma Trivax:

I think a lot of states are going to look at Colorado as maybe a baseline for how they draft their own legislation around AI and transparency because, well, really, that's the problem. Colorado is one of the first states to have a comprehensive AI law in place. There's some other patchwork laws here and there, but Colorado is probably one of the most lengthy and comprehensive.

So, I do think a lot of states will look to Colorado. Do I think they will directly take some of this? I don't know because like we had talked about, there is a lot of room for interpretation, and I think there can be changes made in other states that make more sense for those states or maybe what Colorado has makes sense for other states. Ultimately, there's a lot of work that states have to do to ensure a proper implementation of the use of AI, but Colorado's a good starting point.

Erin Whaley:

Emma, I think you're right that in the absence of overarching federal legislation, the states are going to look to legislate. From a, I think for many developers and deployers, thinking about dealing with a patchwork of state laws is some cause for concern. Well, many of them may want to be unregulated entirely. I do think that they may prefer one overarching federal law that they have to comply with.

Similar, we see the patchwork of state privacy laws. We know how difficult it is for folks who work across states to comply with all of those. It's going to be just as hard to comply with the patchwork of states with respect to AI development. So, it'll be interesting to see who gets there first, and if federal law does curtail some of the state efforts, but in the end, developers and deployers, I think transparency is going to be key.

Brett Mason:

Yes, I appreciate that. And transparency and working hard to eliminate bias in the system seems to be the two things that are really important to legislators across the country. In European Union, we've seen there's a lot of discussion about that in their AI acts as well.

So, Emma, Erin, thanks so much for being on. I really appreciate it. Thank you to our listeners. I hope everybody learned a little bit about the Colorado Act and the way that Colorado's looking to regulate AI use, whether by the developers or the deployers in Colorado. Please don't hesitate to reach out to me at brett.mason@troutman.com. If you have any questions, comments, or there's a topic you'd love to hear us talk about on the podcast. You can also subscribe and listen to other Troutman Pepper Locke podcast, wherever you listen to your podcast regularly, including Apple, Google, and Spotify. So, Emma, Erin, thanks again for joining.

Copyright, Troutman Pepper Locke 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 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 Locke. If you have any questions, please contact us at troutman.com.