

The Good Bot: Artificial Intelligence, Health Care, and the Law — State AGs' Continued Focus on Enforcement – With or Without AI Legislation

Host: Brett Mason

Guests: Gene Fishel and Chris Carlson

Recorded: August 15, 2025

Aired: September 16, 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 here 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 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 other industries, and the legal implications of integrating this technology into the healthcare sector.

I'm really excited today to be joined by two of my colleagues. First of all is Gene Fishel. Gene is a member of our firm's Regulatory Investigations, Strategy, + Enforcement practice, and he has the distinction of serving in the computer crimes section of the Office of the Attorney General of Virginia for over 20 years. Gene, thanks so much for being here and lending us with your expertise.

Gene Fishel:

Great. Thanks to be here, Brett.

Brett Mason:

I'm also happy to be joined again by my colleague, Chris Carlson. Chris is in our Regulatory Investigations, Strategy, + Enforcement group as well, with a focus on State Attorney General and also, has a background of working in the AG's office. Thanks for joining us again, Chris.

Chris Carlson:

Of course. I must not have done something horrible, because I'm back.

Brett Mason:

Well, as you could probably tell from the introduction of my two distinguished guests today, today we're going to continue a discussion about the current state of artificial intelligence specific laws, how those are being viewed by the various state's attorney generals and different enforcement and litigation actions we're seeing around those. Gene, for those who are not on the up and up on all of the changes and updates that are in this area, can you start us off by just

giving us an overview of the state laws that are being suggested, or enacted around artificial intelligence?

Gene Fishel:

Sure. I'll give this a shot. It's a ever evolving landscape, as I'm sure everyone knows, when it comes to AI, particularly generative AI. We start with the four states that have enacted specific AI governance laws. Those are Colorado, Utah, California, and Texas. Those four states have been the first to actually get legislation through, that specifically addresses generative AI. I'll give you a little bit of a high-level overview of each of these and generally, what they do.

We start with Colorado, which was the first to pass AI specific legislation back a couple years ago. This law goes into effect next year, February 1st, 2026. Colorado's law is relatively narrow in scope. It applies only to high-risk AI systems. Those are systems that are a substantial factor in making decisions, or that make consequential decisions, or decisions with some material or legal effect. The purpose of the Colorado AI law is to weed out algorithmic discrimination in AI systems. It's really only focused on bias and discrimination.

An AI system that creates differential treatment among protected classes, or produces outcomes that are different among protective classes, or it prohibits companies using AI systems where there is a foreseeable risk of some sort of discriminatory outcome, that is prohibited within the Colorado law. Of course, like these other laws, there are some entity-based exceptions within Colorado's financial institutions, for example, who are subject to similar restrictions and other laws. It doesn't apply to them.

Notably, the Colorado law applies to both developers of AI systems and deployers. Companies that are specifically deploying and covers them, even though they may not have developed the AI system. Also, importantly, the Colorado legislature in the coming days is actually going back and reviewing this AI law and are going to determine whether any changes need to be made before it goes into effect on February 1st, 2026. That's Colorado algorithmic discrimination, as what it applies to.

Utah, another state that passed an AI specific law. Essentially, Utah's law just brings generative AI systems within the purview of their Consumer Protection Act, essentially their UDAP statute. That's really all it does. It does mostly apply to, again, high-risk systems. Utah legislature passed some amendments earlier this year that narrowed the scope of what companies need to do when they deploy generative AI systems. Narrowed the scope of this act, meaning generally speaking, companies must notify under the Utah law when they're utilizing generative AI. Well, the legislature earlier this year narrowed that provision to include only instances where a company is specifically asked whether they are using generative AI systems. That's the only time now that this notification requirement is going to apply. There's also a safe harbor provision in there. If you conspicuously post that your company uses generative AI, then you're going to be safe under the Utah Act. Really, consumer protection-based law in Utah.

California, the third state that has enacted AI specific legislation. California has gone about it in a little different way. They enacted about 11 different pieces of legislation. Most of those pieces address very narrow specific situations and cases. For example, there are a few of these laws

only apply to political advertisements and election-related laws that really don't apply across the board. There are a couple very narrow healthcare related laws and certain circumstances.

Really, the two broadest provisions that were enacted in California is one, it states that any generative AI that produces, or generates personal identifying information, that is going to fall within the purview of the California Consumer Privacy Act, the CCPA. That's all that does. The second one, the second piece of broad legislation in California only applies to developers of AI, not deployers. It states that the developer must disclose on its website that it is utilizing generative AI system. It's just a transparency measure that only applies to developers of AI systems. Actually, there's a threshold there. It's a website. It's a company that that provision applies only to companies with a million, or more visitors. Again, pretty narrow measures in California.

Finally, the most recent and arguably the most comprehensive law is Texas. Texas passed earlier this year in AI governance law that does several different things for healthcare providers, specifically, healthcare services and people providing healthcare treatments who utilize generative AI, or automated decision-making systems in their practices must disclose the patients that they're using such systems. That's one point. It prohibits developing AI that causes harm to another person, or that might encourage one to inflict self-harm.

Also, the law prohibits deploying, or developing an AI system that infringes, restricts, or impairs any rights under the U.S. Constitution. That's a pretty broad provision of this Texas law. We'll see how that translates. Also prohibits discrimination against a protected class based on race, sex, age, disability. Finally, there's a provision in there that you cannot develop, or produce an AI system that's solely used to create deep fake videos, or produce child sexual abuse materials.

Texas probably covers the most different areas between these four. I would note that in all of these, the respective attorneys general in each of these states has enforcement authority under each of these laws. That's the general landscape of where we sit right now. It's not a lot of AI-specific legislation. I will note that there are dozens of states, 45 states right now, considering AI-specific legislation. This is going to change and it's going to change in the near future, these AI requirements.

Brett Mason:

Well, Gene, thanks for that rundown. You mentioned that the Colorado Act is going to become active in February of 2026. What about Texas, Utah, and California? When are those?

Gene Fishel:

Utah's actually came into effect last year in 2024. The new amendments take effect this year. I'm not sure of a specific date. Texas is January 1st, 2026, and the California laws are 2026.

Brett Mason:

It seems like 2026 is going to be a big year for us, watching how these new laws in various states are going to be enforced.

Gene Fishel:

Yup.

Brett Mason:

Thanks for that. One of the things you mentioned was states having authority under UDAP as well. Chris, I wonder if you could tell us a little bit about that, remind our listeners who may not be familiar with the acronym, what UDAP stands for and how our state attorney general is going to be using their authority under that to enforce around AI.

Chris Carlson:

Yeah, Brett. When Gene says it's not a lot for four laws, my head's spinning. I'm sure companies' heads are spinning thinking, how am I going to comply with this? On top of, I think we talked about last time and the state's view has not changed that their unfair, deceptive acts and practices law, each state's consumer protection law already applies. While we may have added enforcement mechanisms under state laws, including Utah, each state and we've heard this many times from state regulators saying, we don't have to wait for some state law to be passed on AI for us to regulate it. Those statements actually, and I think we've seen in addition, these state UDAP powers also give the authority of certain states to pass regulations to define what would be a violation of UDAP.

A perfect example is Rhode Island, essentially issued a notice and comment saying, "We're going to issue our own rulemaking on what we think is a UDAP under AI." Ask for notice and comment on what would be a UDAP as it related to automated decision-making and the potential for discrimination, essentially, two really major topics of concern. It's no surprise, the ACLU, for instance, at Rhode Island, submits comments and their first concern is in healthcare space. They're talking about potential discrimination in the healthcare space, potential harm in the healthcare space. When we think about enforcement actions and our last time we touched base, it's no surprise that the first AG action related to AI was a Texas case that related to healthcare outcomes and potential hallucinations of AI and patient decision-making.

This is the patchwork we're in, and Gene, it's helpful to know that background when you think about how we had this federal legislation that came by. There's no wonder that big tech and others may have made the decision, why don't we try to preempt all these 900 such laws that are pending. I want to hear from you, Brett, before we switch over, what are you thinking from your real healthcare experience, this patchwork of laws?

Brett Mason:

Well, it's interesting, because there's so many different ways that artificial intelligence technology can be used in the broad healthcare industry. I think what we're seeing from these laws, as you express, Chris, that they're looking to protect the patients, right? Whether that's protect their privacy, protect their – anything that's protected under HIPAA, or protect them from potentially decisions made by AI that could impact patient care and treatment. Those seem to be top of mind when it comes from a statutory, or regulatory enforcement perspective.

At the end of the day, what we're seeing in the law is wanting to protect patients, which I think those in the healthcare industry obviously hold in the highest regard as well. There's a lot of different uses outside of just patient care. I think it will be very interesting to see how these statutes are enforced when you're talking about perhaps, using patient data for a hospital system to use an AI to help increase their efficiencies, or increase the way that they're billing. There's different ways that it can be used that maybe these statutes don't really take into account that companies, especially healthcare companies, are going to have to figure out how they play a role.

Let's talk about that. How do we advise companies to navigate this AI patchwork, both between the statutes that are being enacted by the states, but also the regulations that we're seeing states bring forward under their UDAP authority?

Gene Fishel:

That's a loaded question, right?

Brett Mason:

Give us the answers, Gene. We want them.

Gene Fishel:

Well, and of course, a lot of it is dependent on the kind of industry a company is operating within, because not only do you have these state laws, you have agency regulations that companies have to abide by. I think, and we'll talk a little bit more about some specific technical issues here in a minute. My biggest piece of advice that's been to the companies in this space is, and it's more of a procedural piece of advice and substantive, is companies really need to gather all the stakeholders together to discuss where their AI use sits currently and where it's moving forward.

What I mean by that is the C-suite executives, your IT professionals, and your counsel at a minimum, including experienced outside counsel, all need to come together and there needs to be a mechanism where these groups are communicating with each other about how AI is being deployed, how is this going to impact operations and the potential legal exposure that a company faces in its deployment, or development of an AI system. If you're operating on a national scale in all 50 states, or dozens of states, as with other types of these patchwork of

state laws, it may be, you need to adopt the operations and compliance measures that satisfy and comply with the most restrictive state laws that are out there.

Oftentimes, that's California, for example. If you comply in most of these areas in California, certainly as it relates to privacy, you're probably going to be in compliance in other states. This particular situation needs to be assessed on a case-by-case basis and stay tuned, because in a minute, we will give some specific advice on what you should deploy as it relates to a specific AI system to be in compliance here with these upcoming AI laws.

Brett Mason:

Now, Chris, you had mentioned that there were some federal efforts in the AI space. Did you want to touch on that a little bit, Chris, or Gene, what's going on on the federal side?

Chris Carlson:

I'll let Gene talk on the saga, which was the Big Beautiful Bill moratorium. One thing I will add to what Gene was saying is thinking about, what are your public disclosures about your AI. What you're communicating to the public typically is the source of UDAP laws, or even just business-to-business concerns regarding what the effects, what are the potential harms, what are your mitigation options, just being thoughtful about what you are communicating publicly about your use of AI as a company I think is really important.

Gene Fishel:

Yeah. As far as federal efforts, we certainly can't overlook what's going on in the federal sphere. As we've seen with the new administration, the Trump administration, they have been patently clear that they are for deregulation when it comes to AI. I mean, they've made that expressly clear. The most notable example of this was the budget reconciliation bill that was moving through Congress earlier this year. That bill originally had a provision in it that imposed a 10-year moratorium on states enforcing any sort of AI-specific law.

What happened there is just a really interesting process here, because, notably, those in favor and those in opposition to this moratorium did not divide along party lines. It was bipartisan on both ends. Of course, the Trump administration was in favor of this moratorium. You had Republicans who, of course, controlled both houses of Congress, you had Republicans split. Some Republicans were in favor of it. Some were not.

On top of that, a 40-state attorney's general submitted a letter to Congress opposing this moratorium. Of course, 40-state AGs, that's a bipartisan effort right there. Their rationale was, "We're opposed to this because we have consumer protection, duties to protect our consumers." This is going to hinder those efforts when it comes to oversight of AI systems. They also made, of course, federalism state sovereignty arguments as to why the federal government shouldn't impose on this in this area.

The moratorium provision, it went through some changes. I think it went down to a five-year moratorium at one point. Ultimately, there was not agreement on this, and it was completely

stricken from the bill. As it stands right now, states are free to enact and enforce any sort of AI-specific legislation. Now, a couple of weeks ago, two or three weeks ago, the White House released its "AI Action Plan" to basically, what this plan is doing is it's trying to promote innovation and AI development with the least amount of interference from the government as possible. If you read through the plan, there are essentially three pillars. Again, the goal is to foster American development of AI systems and promote America internationally as a leader in AI.

It does several different things. It sets up regulatory sandboxes, or what the plan calls AI centers of excellence to be established around the country, where companies will be able to rapidly deploy and test AI tools, while sharing data among these test results. As many folks know, one of the more controversial aspects of AI proliferation is the building of data centers. AI systems require a lot of energy, depending on the size of the system. They require these massive data centers to be built that take up a lot of energy, that take up a lot of land. This White House plan seeks to basically, make the permit process for building these data centers easier.

It also, on the international side of things, the government is going to start assisting companies in delivering full stack AI systems to companies overseas and allies, American allies. That includes, when I say full stack, that includes hardware, AI models, large language models, software applications, and even policy and technical standards to be delivered to our allies, and maybe including developing countries. One of the impetus for that is actually related to China, because China has been really out front in developing AI. China's goal is to develop AI systems without Western-made microchips. Also, they're trying to thwart China's influence over international governing bodies when it comes to AI.

I say all this, because it does relate back to state enforcement in that the White House does not want states passing legislation that interferes with any of this, and potentially it could. As a corollary to this White House action plan, or in conjunction with it, I should say, the reports are the White House is going to continue to lobby Congress to pass some legislation that is similar to this moratorium that was stricken from the budget bill in the near future. That's where the federal efforts are headed. But right now, there is no superseding federal legislation that will supersede state AI laws, and the moratorium would not have had an effect over the UDAP laws anyway.

Brett Mason:

That was going to be my question. Would the moratorium stop the states from exercising their authority under UDAP?

Gene Fishel:

No. In fact, states' AGs have been coming out saying, as Chris Mitchell, we're going to enforce under UDAP privacy laws, those sorts of things. It remains to be seen whether there can be any agreement on federal legislation. Over the past 15 years, they've tried it with data breach notification laws. That hasn't happened. They're 51-plus data breach notification laws. They've tried it with privacy laws. They're now 20 states with comprehensive consumer privacy laws, so there's a patchwork of those laws. I don't know. It hasn't been success when the past, because of the parties can agree, or all the interested parties can agree, so we'll see what happens.

Chris Carlson:

Well, let me push back a little bit, though. I think that the potential moratorium really was trying to be sweeping in terms of preempting state regulation of AI. I would have expected companies to the extent, states sought to enforce their UDAP laws against AI-specific things. Very similar to the Airline Deregulation Act, which was passed, the airlines have been able to say, UDAP laws don't apply to the aviation industry based on that law. I think that there very much was perceived from the states a threat that the moratorium would have preempted any authority to take enforcement actions.

Gene Fishel:

That's true. It's arguable. I think the arguments on the other side were the way it was worded, it was AI-specific law that was prohibited. Of course, states capacity laws, they just couldn't enforce them. But I have no doubt that had that moratorium passed, it would have been litigated in some form in court.

Brett Mason:

Now that we have that background, Chris, can you talk to us about the recent attorney general's AI enforcement actions?

Chris Carlson:

Right now, you're seeing a lot of nastygrams being sent to companies regarding AI. Some of those were aware of as more confidential nastygrams, but we haven't seen as many settlements. I think quite honestly, that's because the wave is coming in, not because the actions aren't being taken. Obviously, we spoke on the pieces technology settlement that was an assurance of discontinuance with Texas, and it was the first of its kind on generative AI, and the healthcare space that was last year. It was important to note, there are zero penalties, a recognition that this is the first of its kind action. We're seeing states be more and more in terms of saying, "This is our UDAP laws, and we're not just going to have a slap on the wrist here. We're treating this as a true violation, and you're misrepresenting things," especially in the healthcare space. We're seeing that in terms of big data and making decisions, and you're not being thoughtful, nor are you disclosing how those decisions are being made. Today, I don't have anything. When you put me on podcast number three in six months, just wait.

Gene Fishel:

Just to tack on to that, I think you can look at it because in the last two years, the AGs, several AGs, have issued advisories and almost warnings –

Chris Carlson:

It's a great point.

Gene Fishel:

- as it comes to AI. Oregon did, New Jersey did, Massachusetts did. I think you can look at it. There are three buckets of laws that AGs are gunning, or looking at when they're assessing AI use right now, and those are, as we've already said, UDAP, of course, consumer protection laws, but also, privacy laws, including the comprehensive consumer privacy acts that have been passed in 20 states. Finally, anti-discrimination laws. AGs have signaled that systems that are producing discriminatory, or biased results, they will enforce anti-discrimination laws against those systems. Those three buckets where AGs have enforcement authority are what AGs are looking to utilize, even in states that don't, of course, have AI-specific laws passed in them.

Chris Carlson:

That discrimination aspect isn't it, when you hear discrimination, you typically think that's a blue state issue. Missouri sent a letter to four companies related to allegation that their chatbots were trained to distort historical facts, or produce biased results. This is something that both sides of the aisle are thinking about.

Gene Fishel:

That's a good point. Even traditional "business-friendly states" are examining this. This goes across party lines, this AI assessment, because I think there's a lot of trepidation. Folks are very unsure about the power of AI systems and what could happen in the deployment of these systems.

Brett Mason:

Just to wrap up for our listeners, what are some of the best practices you would recommend for companies using AI, whether they're developer, or deployer companies, to make sure that they're complying with existing law?

Chris Carlson:

Gene will go deeper into this, but I think the most important thing is for a company to understand what type of AI they're using and being able to assess from there. I think AI is such a jaded word, and understanding how they're using it internally is the most important step, so that when you have the sweeps that we expect to come, especially if you are in highly regulated spaces, especially healthcare, you will get inquiries in this regard.

Gene Fishel:

That's absolutely right. That was actually going to be my first point. Companies need to conduct an impact assessment of their AI system. After that, they need to develop and deploy a risk management system. These are questions regulators, particularly AGs, ask right off the bat. You have a risk management system. Have you conducted an impact assessment? Those two are very important, as they relate to compliance with AI-specific laws, but also, as we mentioned, the privacy-related laws. What kind, and this is related to the impact assessment,

but are you using personal identifying information, is your AI system touching PII at all? That is very, very important to AGs, and how is it touching it? Are you using PII to train your data models? Regulators look down on that. They do not want PII to be utilized as you train your AI systems.

How are you effectuating data requests related to AI systems? If you have a customer, or consumer that comes to you and says, exercising their rights under a privacy law and requesting that you delete their information, are you able to actually effectuate that request, if that information is touching an AI system, how do you train your AI system to forget that? AI systems, that is one of the problems with AI systems. They have a hard time forgetting. Once they have that data in them, how are you effectuating those requests? Those are probably the most important, maybe more technical aspects of AI utilization right now. If you can address those on the front end, you really minimize your risk, your regulatory exposure, and even litigation exposure down the road.

Brett Mason:

Thanks, Chris and Gene, so much for joining me. I think this is so interesting as we continue to see the evolution of the laws around AI and how different enforcement entities are thinking about AI. I really appreciate both of you being on.

Gene Fishel:

My pleasure.

Brett Mason:

Thanks again to our listeners for *The Good Bot*. Please, don't hesitate to reach out to me at brett.mason@troutman.com with any questions, comments, or topic suggestions. You can also subscribe and listen to other Troutman Pepper Locke podcasts wherever you listen to podcasts, including Apple, Google, and Spotify. We will see you next time.

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, expressed 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.