

The Good Bot: Artificial Intelligence, Health Care, and the Law — Legal AI in Practice: Firm Governance, Build vs. Buy Decisions, and Vendor Due Diligence

Host: Brett Mason

Guest: Leigh Zeiser

Recorded: 11/5/25

Aired: 1/15/26

Brett Mason (00:04):

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 partner and 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 the legal implications of integrating this technology into the healthcare and legal sectors. I am excited to be joined today by Leigh Zeiser. Leigh is Troutman Pepper Locke's Director of AI and Automation. She's been working to bring AI enabled solutions into the large legal market since 2018. She has experience guiding the technology journeys of both law firms and in-house clients. So, Leigh, welcome. Thank you so much for joining us.

Leigh Zeiser (01:03):

Thanks so much, Brett. I'm happy to be here today and can't wait to get the conversation started.

Brett Mason (01:08):

Let's start off with talking about what we're doing here at Troutman Pepper Locke. Give us a quick snapshot. How are we currently using artificial intelligence across our practice groups and our business operations?

Leigh Zeiser (01:21):

Yeah, our firm actually launched Athena in 2023, and to my knowledge, we were one of the handful of early firms that offered a secure private chat agent, firmwide, that early. Since then, multiple different skills or capabilities have been added to Athena. Plus, the firm has invested in vendor-driven and practice-specific AI enabled solutions including research solutions and others. There's really a very rich AI enabled environment here at Troutman Pepper Locke, helping to support our journey into the future.

Brett Mason (02:00):

I was excited because we just launched the updated version of Athena, which is our internal GPT and now it's a GPT five, right?

Leigh Zeiser (02:08):

That's correct, we did, and what's super exciting about that, Brett, is that Athena 2.0, as we're calling it, has Smart Chat. Smart Chat brings together some of the skills and capabilities where folks would otherwise have needed to select which path they wanted to take. For the average user, that can be difficult to understand what are the nuances and differences between skill sets. For example, we had a capability that was "Document Insights" and also one that was "Chat with Documents". Users had to choose which path to go down, and many didn't understand that "Chat with Documents" was meant for multiple documents versus "Document Insights" was really meant for chatting with a single document. So, there was confusion there. Now with Smart Chat, users do have to select a skill. They can just enter their prompt and our agent will determine what is the best skill to apply based on their prompt.

Brett Mason (03:11):

I know everyone at the firm is excited to make it more user-friendly – lawyers, non-lawyers, everyone, right?

Leigh Zeiser (03:18):

Absolutely. We're going to continue to push the needle on making things user-friendly. That's actually a big part of our focus for 2026 as well. We want lawyers to focus on lawyering and not on understanding nuances between technologies and having to wear that technologist hat. We think that that will really enable our teams to produce more and better work product if they can accelerate through easier interfaces, skills that are combined like we're discussing with Athena 2.0, all of this will come together ultimately to create an optimal experience for our attorneys and our staff to be able to focus on what they need to accomplish with these tools, versus understanding the tools themselves.

Brett Mason (04:06):

I know we've talked about Athena, which is Troutman's GPT, and by the way, Athena is the name of my dog, so she gets very excited when I talk about Athena as well. But what other categories of AI tools does Troutman Pepper Locke use?

Leigh Zeiser (04:19):

We have, really, multiple categories of tools, from research, drafting, eDiscovery through our eMerge subsidiary. We have knowledge management and workflow related tools, and we're continually expanding our tool set to embrace a future-ready firm. I think a key element to note here is that we're not offering generative AI without fully vetting the solutions and understanding how they fit into the big picture. We look at things like data processing, security, governance

controls and more, and we require foundational training on generative AI concepts for each solution as well as ethics around the use of generative AI. We want to really ensure that our people not only have the tools available to them to expedite and create the best value in their work but are also confident in the fact that the tools are secure and know how to appropriately use them. It's really critically important from a change management perspective.

Brett Mason (05:27):

My understanding is that some of the tools we have here at Troutman were built internally and some were buy decisions or vendor-built tools. How do we balance the decision on “let's build it ourselves” versus “let's buy from an AI vendor” when we're looking at different AI tools for the firm?

Leigh Zeiser (05:48):

That's a great question Brett. Historically, firms that have built technology tend to run into challenges supporting the technical debt long term. Plus, when we're talking about generative AI, people throw around the term “build”, and that's not really what most firms are doing. It's incredibly expensive to build a language model. Firms are not really “building”; they're leveraging a language model to create or bring to life a use case. When we're choosing to leverage these capabilities, we're often looking for use cases that are truly bespoke and can't be provided from a vendor, because let's face it, most vendors have deep development pockets, and they want to partner with cutting edge firms like ours to advance in the marketplace and expand their portfolio. If a vendor has the ability to collaborate with us to co-develop a solution, or they have a solution that already exists and, of course, meets our requirements for security, for accuracy, and all of the other things that go into analyzing is a tool of fit for a particular use case, then we'd be more likely to buy versus build.

(07:06):

But increasingly, especially with no code and low-code solutions expanding into market and tools like Power Automate becoming more frequently utilized for automation capabilities within the Microsoft tech stack that many law firms employ within their infrastructure, it's becoming more and more frequent that we can bring forward use cases, leveraging language models to deploy our own solutions without major investment and longstanding technical debt. It's very feasible that we're going to see an increase in law firms leveraging that technology to create bespoke solutions, and I think that that will bring competitive advantage. But it's important to also keep in mind that generative AI is not a magic wand, so when we're leveraging a language model to deploy a bespoke solution here at the firm, that may not be enough. There may need to be quite a bit of fine tuning. There may need to be code underlying that process, automation or other capabilities. You want to level set what's possible currently along with where we're going and what that future possibility looks like.

Brett Mason (08:25):

You talked about vetting and looking at different tools. Can you walk us through the firm's evaluation framework for AI tools? What are the criteria that are non-negotiable?

Leigh Zeiser (08:35):

Yeah, so first and foremost, the security analysis. It doesn't matter how interested we are in a solution; how flashy it looks in a demo. Let's face it – all the vendors make their demos look good. The security and technical analysis to ensure that we understand how the solution operates, where our data will be processed to and from, and how it fits into our technical infrastructure is, first and foremost, one of the most important things. We then need to also look at the fit for purpose of the tool. How well does it do what we need? Every solution on the market will, at this point will say, oh, we have a generative AI component, or we're generative AI enabled, but they all operate really a little bit differently and therefore, the results are different. Depending on what you're really trying to use that tool for, you can have wildly different outputs.

(09:37):

For example, a solution that purports to allow you to analyze multiple documents and create a matrix output, may be really, really great at some of the traditional clauses like extracting notice and assignment using generative AI, and the user could just input a prompt and then the tool will provide that information back out to them in a grid format. But if you're looking for points that are not usually at the clause level, but embedded in the text of documents, it may not do as well. How well it does varies tool to tool. You really have to examine how accurate a solution is and is it really fit for the purpose that you are going to be using it for. The other piece of this is contractual obligations. What is the liability coverage? How are the vendors providing notice? Not to the users, it doesn't help the business or the firm.

(10:42):

If a user gets a notice that a new generative AI component has arrived in that solution because likely the user is going to say, okay, great, and then move on with their day. It's all about how are the administrators, the business owners of the solutions, notified of that change in the updates because a vendor may introduce a new model and that could trigger the need to go back through the security and the technical analysis. It's truly important to look at evaluating AI tools from many different angles, and I think that's where right now with so many tools proliferating the market, we're trying to slow down to speed up and make sure we get the right combination of solutions and the best vendor relationships moving forward.

Brett Mason (11:37):

When your team is doing that vetting process, do you look at and assess vendor claims around accuracy, recall, hallucination rates, and how does your team do that if so?

Leigh Zeiser (11:51):

That's also a great question. I think that not enough people in the market are doing that. I have always said it's important to not take things at face value. So, what do I mean by that? There have been solutions that people in the industry get a big buzz about and they say, oh, everybody's using this and it's so great. Then if you sit down and actually do a controlled test where you're not just throwing something in and saying, yeah, this output looks great, but you're

line by line analyzing, what did you ask the generative AI solution for, and what did it provide to you as an output and how accurate is that versus what you really needed in order to satisfy that use case? That will really change a user analysis. For example, I had a scenario where attorneys at my prior firm got approval from their client for us to test running a DPPA agreement through a generative AI solution to see if that solution could draft a new DPPA from the template.

(13:09):

At first blush it looked really good, but when you started going clause by clause, then you'd notice that even though the instruction for the language model was to follow the template and to use the language verbatim only changing specific items that had been called out, ultimately the tool decided to change the names of clauses. It changed the security clause to "other" as the clause name for no clear reason. It decided to truncate section, so there was a seven-subpart clause that it turned into a one sentence clause. The accuracy is really all about not just taking a quick pass and saying, yeah, this looks good. It's about analyzing things line by line, creating an audit trail so that you can understand how did this tool perform versus the next closest competitor and what are the watch outs. That also enables us to advise our constituents on what they need to look out for as they perform their human in the loop audit and review of any output that they're working with.

Brett Mason (14:26):

Well, I know I appreciate the vetting that your team is doing because at the end of the day, I think it would be very frustrating for anyone at the firm to be using a tool that then is creating work products that requires even more intense vetting to make sure that it's accurate. That's not very useful or efficient if we're having to triple check because the tool isn't working the way that the vendor has claimed it works.

Leigh Zeiser (14:52):

Yeah, I do also think that there's a lot of concern about hallucination and the time that needs to be invested in order to validate an output. If you think about it, it's not incredibly different from what is already being done today. If you draft something, you typically don't just turn it in without looking back over it to make sure that you got the flow right, that there weren't errors in punctuation, that concepts were placed in the right location to make an argument really persuasive. There already is vetting of attorney work product at the major law firm level in order to ensure the best quality product. With these solutions, if you understand how they work and what their risks are, then you know what to look out for when you're doing the analysis and vetting that you probably already would've done, right?

(15:56):

We talk about human in the loop as if it's this really new thing, but ultimately we've always been the human in the loop and great news: we continue to play a really important role in creating the work product. If you look at a lot of the news headlines that have caused fear and concern in the market, it's not about how the tool operated. The generative AI tools did what they were built to do in the way that they were built to do them, but the human that was using the solution just

didn't understand what those risks were and how to manage to them. Anything we can do to create that validated, secure experience and provide the necessary foundational education to arm our people with the ability to spot and triage any potential concerns with a generative AI output as they're doing the review that they already would've done, I think is a critical thing that we can deliver from our AI solutions team.

Brett Mason (17:03):

Let's pivot a little bit now and talk about some advice you have for in-house legal teams that are looking at AI tools. When in-house teams evaluate AI tools, what are the top five questions they should be asking AI tech vendors?

Leigh Zeiser (17:20):

I really like this question because more and more we're seeing in-house teams reach out to us, to their attorneys, and to vendors and consultants in the marketplace to ask exactly this question: what should I be doing? How should I be looking at this? Help me get up to speed and understand how I can better operate within my organization to leverage AI tools and evaluate what my outside counsel may be doing with AI tools as well. The first question is, if an in-house counsel team member is talking to a vendor, what model or models are they using? It's important to know that, because if they're just telling you "it's generative AI" or "we built our own solution", you really need to dig into that. A lot of times when the vendor or the salesperson says, "we built our own", they just mean that they're leveraging a third-party language model to create a fine-tuned solution like we talked about a little bit earlier.

(18:26):

Ultimately the model underlying the solution is really, really important because it is a third-party provider that is going to be processing your data, so you need to understand where, and this is maybe question number two, where do those models sit? Not just who is providing them, but where do the models sit or essentially where's your data going to be processing to and from? There's a research vendor that I evaluated early on in early 2024, and when we asked these questions, we determined that in fact, even though this was a very well-known provider, they were processing data to France. The models that they were leveraging in their solution were physically in France and they did not at that time have a US alternative. Frankly, that is a huge risk because then there's implications around outside counsel guidelines that may prohibit the processing of data overseas.

(19:32):

There's GDPR implications, so we really need to understand who's taking ownership of the model, or providing the model, where does the model sit? Then third question, how is the vendor that is deploying this model, auditing the model and the model provider, to ensure continued accuracy, to eliminate bias and to prevent drift? Over time, as models are exposed to more and more data, they can drift and become less accurate than they previously were. Bias can be introduced. You want to make sure that the vendor that's responsible for bringing this model into the tool that they're selling to you is not just relying on that third party model provider, delegating their responsibility to Anthropic or Open AI to worry about things like model drift and

accuracy. They need to have responsibility in auditing that as well. Maybe fourth question is, what's the liability coverage in the underlying vendor agreement?

(20:42):

If something does go wrong, how secure is that agreement in order to protect you any costs associated with error? Because frankly, if the vendor that's providing you this tool is not auditing the model provider and there is some sort of drift, there could be inaccuracies that proliferate into the tool and you want to make sure, just in case of a rainy day, that your coverage, contractually, matches the potential liability. I think finally, fifth question. We mentioned this a little bit before, how are leaders or business owners of the tools, not the users, notified of changes to the platform or model? Because that might trigger an additional security review and ideally, it'd be great to get that into your contract, that the vendor is required to notify a designated business leader or business team versus just notifying users of changes that they're making that could impact your security analysis.

Brett Mason (21:53):

Well, thank you so much for those questions. I think that'll be really helpful to our listeners when they're thinking about different tools that they're bringing in. Let's talk a little bit about collaboration between outside counsel and in-house counsel. When it comes to discussing and using AI tools, what division of responsibilities works best when outside counsel uses AI? Who validates the outputs? Who owns the risks, and how are the processes documented?

Leigh Zeiser (22:19):

The person using generative AI must be, needs to be responsible for their work product. But the ABA Formal Opinion 512 also references the fact that any supervisor is responsible for the work as well. If we take a typical scenario at a firm where the associate perhaps is performing the work with a generative AI enabled tool, they need to thoroughly vet the output and ensure that there is accuracy in that output. They need to disclose to their supervisor or the partner overseeing that work that they've used, generative AI so that person is on notice and can understand what kind of risks that they're going to look for when they perform their standard oversight of that work product. But then it also, as the work passes to the client, there needs to be transparency to the client that there was some element of generative AI used in order to create that work product so that the client can also do their due diligence.

(23:32):

Ultimately, if there is a group of people that are working on a matter or a use case and they're employing generative AI, it's important for everybody to know that generative AI was used to create that work product and to make sure that they're carefully reviewing that output. Now, we're not going to say to a client, "Hey, client, make sure you review this". But I think that as attorneys, a lot of our outside counsel understand that if they're receiving that work product and they're going to use it in their organization, they also have a responsibility to make sure that they are comfortable with that output as well. I think that there is no one person that can be the individual responsible for validating and owning that risk.

Brett Mason (24:32):

I appreciate everything we've talked about today. It's given a lot of insight on what we're doing here at Troutman and what other firms and in-house counsel teams can be thinking about when they're looking at AI tools. We're winding down here, and so with your expertise in this area, especially working with law firms since 2018, tell us one misconception about artificial intelligence in the practice of law that you'd like to retire, that you'd like to put to bed.

Leigh Zeiser (24:58):

Oh gosh, just one?

Brett Mason (25:00):

You can do more than one, but what are you seeing, the misconceptions and the conversations you're having? I think this podcast is pro-technology. We're pro-figuring out how artificial intelligence use can advance both the legal industry and the healthcare industry for cost saving purposes, for better patient outcomes, for better legal outcomes for our clients. We're pro-use of technology here, but what are some of the misconceptions that you'd really like folks to hear from you that they need to let go of?

Leigh Zeiser (25:29):

I'm maybe going to hit two things, Brett. I think the first thing is that AI is in some way humanoid and it's going to take our jobs. Currently, it certainly is not.

Brett Mason (25:43):

So that's a misconception, you're saying. AI is not going to take our jobs. Alright.

Leigh Zeiser (25:47):

That is a misconception, and if you look at history over and over again, that's a conception that arises when there's a really cool technological advancement and yet jobs may change. The focus of how we do them may shift, but those jobs have continued to persist and exist. I was in a presentation the other day and an attorney said, yeah, I get it. I can remember when clients told us not to use email and people were concerned that email was going to eat away at our need for legal professional assistance, and yet we still have those. They're doing different work. Their work is not eliminated, they're just tech enabled. This is, at its core, it is a computer. It is an algorithm. If you understand that language models learn from vast data sets and what they're learning is the pattern, then you can understand that, essentially, they're applying their understanding of the English language, or grammar, and how words fit together, and mathematics, probability to be specific, in order to determine the best output for a given prompt.

(27:07):

Let me boil that down. If a data set that was used to train a language model indicates that most dogs are brown and you ask the model to describe a dog, it will say the dog is brown. But if you specify in your prompts that the dog lives at a firehouse, the model will understand from its data set that most dogs living at firehouses are dalmatians and it will say instead that the dog is black and white because it's looking for data patterns in order to provide a response. This is about the language model, interpreting patterns of information and responding to your prompt with the most likely output that it is aware of. It is not by itself really creating new information. It's just fitting information together based on its understanding of probability, its understanding of the dataset it was trained on, and its knowledge of language.

(28:12):

In its current state, it is mimicking human thinking, but it is not really thinking. If, for example, you interact with a chat agent and instead of providing you an output, it asks you additional questions, it's doing so because a human programmer specifically coded the system to gather more information on that topic or to gather more information if it didn't have a clear response that it could identify within its dataset. It's not that it looked at the prompt and independently decided, oh, I have to ask more questions about this. It's being prompted on the backend by the developer or triggered by the developer to give you that next question so that it can better provide an answer. Again, using patterns, probability, and understanding of the English language. I think that's point 1 on misconceptions that I would like to retire. I'm going to pause and let you react to that before I go into point 2.

Brett Mason (29:21):

Yes. I was just going to say, so what I'm hearing you say is, although this technology is going to help us be more efficient and help us to gather thoughts, we still need our lawyer analysis judgment to bring asking the right prompts to the table in order to actually use artificial intelligence technology.

Leigh Zeiser (29:43):

Exactly, Brett, and I think the other piece of that is we need to be looking at that output, as we've talked about so much today, to make sure that the AI got it right because it doesn't truly understand how concepts fit together. Because it's really focused on pulling together patterns of information, without that depth of understanding of how things fit, it's really easy for it to misconstrue, to maybe connect to topics that don't really fit together, or to try to analyze something in a way that isn't exactly right according to legal precedent. We need to keep our thinking caps on. Ultimately, the system is processing information and providing a great starting point to help us be efficient and effective, but it is not doing the deep thinking and the deep analysis anywhere close to what a human would do.

Brett Mason (30:44):

Okay, so our jobs are safe. That's good. Let's not be afraid of the technology because of that. What is the second misconception you'd like to put to rest?

Leigh Zeiser (30:52):

Yeah, so we're talking about efficiency, and that's been something that's discussed quite frequently in the market, but I think we really need to start shifting the conversation from efficiency to effectiveness, because while many of these tools do help us be more efficient, it's a misconception that that's going to result in a time savings. What we're looking at, what we're seeing is that practitioners that really use generative AI tools well, that are not out there trying to get generative AI to do the job for them, but are trying to interact with generative AI in order to improve their work product, in order to make sure they've covered all the bases. They may not save time, but they'll produce an optimal output that is much higher in quality than they could have provided within the time available without generative AI. I think what we really want is for attorneys to be lawyering and for clients to be paying for lawyering.

(32:03):

If we could use generative AI tools to give a lift on things that, otherwise clients would likely write off anyway, or really rote processes, like if you are parsing and extracting information and diligence, there's no reason why you shouldn't be using an AI enabled solution to do that. Paying an attorney to copy and paste information in order to pull together a matrix is just no longer necessary. However, once the information is pulled together in that matrix, there's a lot of thought and analysis and validation work that the attorney needs to do in order to make sure that they've got all the complete and thorough pieces of the puzzle and pull it together into a great analysis or memo for the client.

(33:15):

I'd love to see the conversation shift from this idea that efficiency is all about saving time and money, and these tools are just about efficiency. They're not, they're really about being more effective as attorneys. In fact, we are seeing many attorneys draft their own content and then ask a generative AI tool if they were arguing this draft with opposing counsel, what would the counter arguments be so that they can tighten the rhetoric and they can ensure that they've maximized the model argument in their document. I think that exercises like that potentially take more time than just writing something up and pushing it out the door but ultimately protect the client's best interest much better.

Brett Mason (34:06):

Well, I love that. Thank you so much, Leigh. I think this has been a great conversation, and I know that our listeners will appreciate your insights on how they need to be thinking about incorporating AI technology into their legal practices, whether they're outside counsel or inside a business. I also want to say thanks to our listeners. 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 our other Troutman Pepper Locke podcasts wherever you listen to your podcasts, including on Apple, Google, and Spotify. So thanks again, Leigh, for joining me.

Leigh Zeiser (34:42):

Thank you, Brett. Thanks for having me.

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.

DISCLAIMER: This transcript was generated using artificial intelligence technology and may contain inaccuracies or errors. The transcript is provided "as is," with no warranty as to the accuracy or reliability. Please listen to the podcast for complete and accurate content. You may [contact us](#) to ask questions or to provide feedback if you believe that something is inaccurately transcribed.