The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices
Host: Brett Mason
Guests: Alison Grounds
Recorded: 7/23/24
Aired: 11/19/24

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, 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 first episode of this podcast. In that episode, we lay the groundwork of understanding the technology that is the basis for all of our discussions.

I'm really excited today to have on the podcast, Alison Grounds, who is the managing partner for Troutman Pepper eMerge. She is also the chair of the Innovation Committee at Troutman Pepper and one of the leaders of the Generative AI Task Force at Troutman. So, Alison, welcome to the podcast. Thanks so much for joining us today.

Alison Grounds:

Well, thanks for having me, Brett. I'm excited.

Brett Mason:

Now, Allison, since you have been one of the people in the forefront here at Troutman Pepper on us incorporating the use of Generative AI, I just want to talk with you first about a little bit of the background of what that work was and how we're using Generative AI here at Troutman Pepper to begin.

Alison Grounds:

Yes, it's a great question. But when Generative AI first became a thing, if you will, I would say, that really sparked the attention of everyone was when ChatGPT launched. We got so many quick users. We quickly formed a task force because we said, "You know what, AI is not new. We've been using it in eDiscovery and other aspects of our practice for years, and we've been advising clients about that."

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

So, at the firm, we took that as an opportunity to sort of consolidate resources across practice groups, across administrative support teams, business professionals. We formed a task force focused on how we could make sure we were deploying Generative AI and AI generally, most effectively within the law firm.

Our key use cases, I would say, fall into, we built out internally, our innovation team, Athena which is a GPT-powered AI assistant that is in our secure firm environment, so we can use that for all sorts of things. Improving our email drafting. I've used it to help me do outlines for things like podcasts and in CLEs. So, lots of good use cases. Then, the task force has also been testing other practice-specific tools. So, we've deployed eMerge, which is our eDiscovery team and data management team in the firm, Generative AI solutions for reviewing and analyzing documents for litigation or investigations. And we've been pretty impressed with how it works compared to prior iterations of AI. And then we're also using it to extract information from documents for clients.

So, creating chronologies, deposition outlines, analyzing transcripts after depositions are taken, and we're really seeing some great results from that and are excited about what may happen next. But those are just a few things. We're also testing software for tax, legal research, really anything, and as you know, every software that exists in the world today is coming out with some Generative AI version. So, we're just trying to stay on top of that and see what will make the most sense for our firm and our clients.

Brett Mason:

I know throughout that process, there's been different approaches or what I say, evolving approaches as the technology is evolving. From your experience working on the innovation team and the Generative AI task force with Troutman Pepper, do you have some best practices or policies that you would recommend to companies that are looking to use Generative AI for their own business functions?

Alison Grounds:

Yes. It's been exciting. I mean, one of the reasons we kind of got together and made sure we were thinking about all the different aspects of Generative AI in particular, we thought, this is a great time to go through this journey with our clients. This is still a new iteration of AI. So, we've taken a lot of our own lessons and been able to go through this with clients. Figuring out who the right stakeholders are within an organization and most of our points of contact are in the legal department. Historically, sometimes AI and technology is happening in the business side and not in the legal department. So, some of the lessons that we've learned is we very quickly connected those dots, and we've done that by pulling in our business professionals and our general counsel's office, our ethics council, our IP professionals, to understand both the legal risks and the technical risks. So, having data scientists involved as well.

I'd say the most helpful thing we're seeing with our clients is helping, again, our point of contact is often the legal department. What role does a legal department play? How can we help coordinate and not stop innovation, but empower it and do it in a way that's going to be beneficial to the business with reduced risk?

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

We're counseling our clients to sort of look at their policies and see. I always say, no policy is better than a bad policy. Helping clients to think about what AI strategy looks like within their company, what risk tolerance they have based on their industry and profile, and then pulling the right stakeholders together to implement those policies based on their culture and their business to make sure they're taking the right steps. So, we took some of our own lessons and then, of course, consolidated the experiences we were seeing among our clients to help advise in that space.

Brett Mason:

Have you seen clients' perspectives on using Generative AI change over the past year?

Alison Grounds:

Oh, my gosh. Absolutely. I think, at the very beginning, the knee-jerk reaction was absolutely not. You had some examples of bad use cases, right? The lawyers citing it incorrectly for cases it didn't exist in briefs, engineers putting in proprietary source code into public models. So, the very first reaction I think we saw from a lot of people, as you see with any innovation, is fear. Let's just block the sites. To be clear, we also blocked ChatGPT and built Athena so that we would have a safe place to use this technology.

But you're certainly seeing a pivot and you're seeing clients embracing it. We also are tracking all of the outside council guidelines that come in from our clients. We have yet to see a client forbid us from using the technology. Some are asking us to let them know when we do, and we're also seeing a shift of, I've got several clients that are excited to be in pilot programs with us in using this technology so that they can be thought leaders and have hands-on experience within their own organization.

So, definitely seeing a shift in excitement, and also moving from policies that say no to policies that are more flexible and factor in the risks, but also empower the use cases. So, certainly seeing that shift.

Brett Mason:

Now, you talked a little bit earlier about the fact that artificial intelligence has been a part of eDiscovery for forever. So, what's new? What's changing? Are there exciting things that are being done that are going to make eDiscovery more efficient and effective for companies if they are in litigation?

Alison Grounds:

Absolutely. I think, as I always tell people about eDiscovery in general. I mean, it was one of the first legal problems that required technology. When I first started practicing as an IP litigator and handling lots of large pharmaceutical litigation matters, I would review documents in paper, or in a database. Certainly, we're excited about things like search terms and conceptual mapping and using technology and AI to find similar documents as it has evolved. But now, I really think we're finally seeing a pivot with Generative AI in using AI in a way that's much more transparent and reliable.

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

I think, similar to ChatGPT kind of going directly to the end user, right? You can communicate with it. It's conversational. It gives you feedback. The Generative AI tools that we're using in Relativity through their aiR products do exactly that. You give it prompts. You explain what you're looking for. Because it's got that large language model on the back end, it already has a baseline of understanding about how things occur. Then when it codes the documents, it tells you why. It gives you its rationale. It points to the citation within the document that was the reason for its decision, which is much more transparent and interactive than what we've seen with prior AI and eDiscovery.

Brett Mason:

Alison, when we were preparing for this call, one of the things I remember we talked about is using Generative AI tools to help with the creation of things like privilege logs, which as a litigator, I know can be extremely extensive and tedious and time-consuming to make. So, can you talk about that efficiency and what you're seeing and using in that way?

Alison Grounds:

Absolutely. Privilege review is probably just manual document review is the most expensive part of the entire discovery process, as you know. Privilege in particular is one of the most sensitive areas that I think the AI in the past has struggled with. So, we certainly use tools to automate the crafting of privilege logs. We can do some cool name normalization. We can extract some basic data to make some of that easier. What we're seeing with the Generative AI tools is an improvement in their ability to identify privileged documents. In the past, we relied on search terms. We were really still heavily dependent on human review to really get those final calls.

We have a committee dedicated to privilege protection at eMerge and the head of that committee, Eric Chapman, has tested all the tools. And he's just a curmudgeon and he's like, "Rah, it's no good." We pilot with aiR's product for privilege review, and he really had to kind of take a step back. This is pretty impressive. It still misses some of the context, right? It doesn't quite understand the email and its attachment and their relations, but that is vastly improving quite quickly. I think because it is based on a large language model and is a more advanced form of AI, it does get the context of the document better and it understands more nuance in the privileged calls.

So, I think we're going to see an improvement in the use of technology to make privilege calls. Because it's generating those explanations, it's possibly going to help us also generate the logs a little bit more accurately.

Brett Mason:

What do you have to say to attorneys who are genuinely afraid that the use of these artificial intelligence tools are going to take their jobs or make attorneys not necessary throughout this process?

Alison Grounds:

This same discussion happened when we first started using simple AI, if you will, for document review and for eDiscovery, and for other legal solutions. My thought on whether or not AI is

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

going to take over our jobs, and of course, Generative AI is different. It is generating documents. It's generating text as everything has done before.

But I've got two thoughts on it. One, in the eDiscovery space, a lot of what we've done has been limited by the sheer volume of information. So, maybe we only collected five custodians. Or maybe we only decided to pursue our contractual rights on that one big contract because it was cost-prohibitive to go through discovery. It seems like we could do more. Maybe we can review more documents faster and more accurately. We can have more litigation if we need to, to resolve our disputes. Maybe we can prevent more litigation because we're analyzing the documents ahead of time and we're catching business issues in advance.

So, I think it's got a lot of utility, we're creating more and more data every day, including Generative AI being used for so many things. That's creating a whole new set of data and information. I think our attorneys will still be needed in figuring out how to best use the data we create, how to analyze it, and we should still be, hopefully, using this information in legal disputes.

The other area that I see is I'm finding an emphasis or at least the increased importance of the human aspect of things, right? I mean, you're a trial lawyer. You know how important it is to have that connection with the jury. As long as decision-makers are juries and judges and other humans, you still need humans to provide that layer. So, the AI hopefully takes out some of the stuff we didn't like doing anyway, the redundant task, threw some of the junk, and then it lifts to the top of the things we really want to see. Then, you still get to use your big lawyer brain to come up with the strategy of what's going to be most impactful to the other human on the other side of the transaction, or to the other side of litigation, or internally within the compliance department, right?

I think there's still a role for us to play, and I'm hopeful that that role is much more interesting. I mean, right now when we're in a phase where there are so many AI tools, I mean, in our firm alone, I've just rattled off several of them, we're still in the phase where it's valuable as sophisticated attorneys that love technology, how to architect the right combination of technology tools to get the result that we need.

At some point, I'm sure there will be some dominant tools that just take over and there's not as much of a nuance in that piece, but there's still the nuance in figuring out what to do with that information and how to use it for the human-to-human interactions that we are still having across the legal landscape and the business landscape.

Brett Mason:

Talking about that right combination of tools, what's looking most promising right now and what are you saying that needs some improvement?

Alison Grounds:

I'm definitely impressed with certainly eDiscovery and legal tools that are helping us to review documents, like I mentioned, aiR for Review, being very good at what it's doing. I think they're still struggling with how to price that and we're still going to struggle with how to, if we get in fights and litigation about whether or not we need to validate. Do I need to disclose that I used

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

it? Do I need to disclose what my aiR rate was or my turnover rate? I would argue we should not have to. We didn't do that for manual review.

So, you have some potential drawbacks in the adoption, as we always have with people who may be questioning or wanting to have more to say about the use of the technology. I have a lot of hope and I'm very interested in the Generative AI tools that are helping us to do more document analysis, not just review, but synthesizing and finding themes, right? So, if you get a big old dump of a thousand or a million documents, I see this all the time in some of these cases where the other side maybe doesn't want to spend the time to really care about what's responsive. They just want to give you a bunch of junk. So, being able to use Generative AI and telling it, "Here's the complaint, here are my claims, here are my defenses. Tell me what's in this big pile. Tell me the things that are most likely to support my claims. Where's my risk?" You could do that with your own documents. You could do it with productions received. I think the early – it's not there yet, but the potential is there, and we're seeing some pretty impressive results.

I think that's what I hope it will help us to find the information that we need faster. Like I said, potentially, even before litigation is necessary, we might be able to resolve some disputes because we could do some analysis on the front end. So, I'm most happy about and excited about the Generative AI platforms and tools that help us synthesize information and prioritize what we need to actually take a look at.

Brett Mason:

Now, we've talked about integrating Generative AI at Troutman Pepper and the advising we've been doing on clients who want to do the same. Let's talk now about on the back end, once our clients or companies do start using AI in various forms to their business, as a litigator, I'm thinking about what is that going to look like once we're in litigation? Have you given thought to what it's going to look like to do eDiscovery of AI info or AI business-run software itself?

Alison Grounds:

Absolutely. We are actively consulting with clients on this issue who are already deploying and using AI classic kind, as well as Generative AI. So, we're certainly already seeing a trend, and I would say, this happens a lot in class actions. It comes up with financial institution clients and insurance clients. Any industry that is using algorithms and AI to help make business decisions, you're seeing requests for that. Did you use AI for your underwriting? Or did you use AI to decide what the interest rate was going to be? We want to see the algorithm. We want to kind of pick. That's even classic AI.

So, I think as people begin to use more Generative AI, the back-end questions about how did you come up with this result? Who made this decision? Was there any bias in your dataset or your algorithms? I think you're going to see increasing requests for that information, and our clients are having to decide from a litigation readiness perspective and a business perspective, what do we need to keep, right? You're going to have Microsoft Copilot. That's not a big deal. Everyone probably is going to have an opportunity to have that, right? You don't have to – so what's even Microsoft's default settings for how long it keeps the information, the prompt, and the answer, and the refinement? And then do you want to change that default based on your own business needs?

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

I think you have certainly a line of thinking that people should just keep everything. We're not required to keep everything. In fact, at this point, you've got a risk. There's privacy concerns and data security concerns with over-preserving and keeping more information than you need. So, we're definitely seeing clients think proactively about what tools are we using? What data are they generating? How are we preserving it? How would we preserve it if required to because there's litigation that it might be relevant to? What types of litigation would it be relevant to?

I mean, if you're chatting on the side in Microsoft Teams and just having a conversation with another human, people like that, they want that. If you're using Microsoft Copilot to help you draft an agreement, do they want to see that iteration? Does that matter? Is that the same as something you might have drafted and discarded before. So, it's a current active, proactive thought that everyone has on their minds. What is the kind of information that's being created? What do we want to keep around for our business purposes? What are we obligated to store and preserve if there's litigation?

I love it. I love when there's a new problem to solve. It's so interesting to learn about the different large language models, the different options, clients that are on the full spectrum from dipping their toe into these solutions to really being high-level technology companies that have a lot of different platforms that may not all communicate with each other and may not be as easy to preserve and collect as you would think.

As with most technology that our clients use, rightfully so, it's not built for eDiscovery. It's built for a business purpose. The reason they're using it is to have better patient outcomes or to improve their delivery of services to their own customers. Then the back end, it usually is a, well now, as an afterthought, let's think about the eDiscovery implication and that's how it should be. It shouldn't be the reason that's driving your decisions, but it is exciting to proactively think about what if? How is this implicated in eDiscovery, and the process, and how we're running our business from a data governance and information management perspective?

Brett Mason:

One of the things that I feel has come up throughout all these episodes I've been doing, and I'm curious in your thoughts, is just the theme of transparency. We know that many AI software perhaps have a black box where it can't actually show you how it did its work, how it got to the answer. Are you advising clients on thinking about that when they're looking at different vendors and tools to use, when it comes to a, you might need to be able to show your work later on down the road perspective?

Alison Grounds:

I think one of the things we're trying to help consult clients around is obviously the AI ecosystem includes off-the-shelf third-party applications that you may not have access to the black box. It also includes some certainly proprietary tools that our clients are building themselves internally. Then, in most cases, it's some sort of hybrid combination. So, figuring out what's your overall AI policy, what are the factors that you're considering? And then how are you ensuring that the AI is performing the task it was intended to do? So, that validation process.

Another way of thinking about that, and that goes to your earlier question is knowing which model was used at which point in time and kind of tracking those versions and what dataset was

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices

being used to train the model, understanding what the final use case was, and is it delivering as intended? What are your compliance and or overarching protocols for ensuring that whatever AI you're deploying is doing its job in a way that is at least as good as the alternative and that you're able to confirm and validate the results for.

So, I don't know that you necessarily have to be able to unwind the entire black box. If you hear interviews from the largest leaders in the tech space, they can't even tell you how some of these models are doing what they're doing. But I think the key role that we can certainly play as legal advisors, whether you're in-house or outside advising clients, is to get a better sense of that big-picture goal. What is the tool doing? How can we document which version we're using and what we understood it to be doing at the time? And how can we kind of confirm that we continue to monitor and use reasonable efforts to oversee the AI, depending on the legal or regulatory framework we're operating in.

Brett Mason:

Alison, thank you so much. I feel like you and I could talk on this topic for a very long time, but we really appreciate you being on the podcast here today and sharing some of what we're doing here at Troutman Pepper and how eMerge is helping our clients think about and start using Generative AI. So, thanks for joining me.

Alison Grounds:

Well, Brett, it was super fun, and happy to do it, and I will continue to listen and learn more from your other guests and you as well. So, thank you so much.

Brett Mason:

Thanks, Alison. Thanks so much 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 podcasts, wherever you listen to podcasts, including Apple, Google, and Spotify. I believe, Alison, correct me if I'm wrong, but our troutman.com/ai is where you can go online if you want to look at all the different exciting things that we're doing and the information we're providing on Generative AI.

Alison Grounds:

Yes. There are lots of podcasts, thought leadership pieces, case studies, lots of good stuff there. So, thank you for that plug. We'd be happy to direct people to see more of the great things we're doing.

Brett Mason:

Right. Thanks so much, Alison. Talk to you soon.

Alison Grounds:

All right. Thanks, Brett. Talk to you later.

The Good Bot: Artificial Intelligence, Health Care, and the Law —
Harnessing Generative AI: Innovations and Best Practices