The Good Bot: Artificial Intelligence, Health Care, and the Law —
Preventative Medicine: Health Care AI Privacy and Cybersecurity – Part 2
Host: Brett Mason
Guests: Andrea Bilbija
Recorded: 7/22/24

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 am 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 sector. If you need a basic understanding of what artificial intelligence technology is, and how it's being integrated into healthcare, I highly recommend you start with our inaugural episode of this podcast. There, we lay the groundwork for understanding the technology that is the basis of all of our discussions.

I'm excited to welcome to the podcast today, Andrea Bilbija, who is the Chief Compliance Officer at Sprinklr, and my colleague, Brent Hoard, who's one of our partners in our privacy and cybersecurity practice group. So, Andrea, why don't you just introduce yourself to our listeners and talk a little bit about Sprinkler and your role there?

Andrea Bilbija:

Yes. Hi. Thank you so much for having me today. It's a pleasure to join you both. I'm the Chief Compliance Officer at Sprinklr, which is a global cloud-based B2B company. We service brands with a unified customer experience platform. So, the goal of the product is to help our customers service their customers with social media outreach, contact center communications, advertising, and brand awareness tools. As you can imagine, a large part of that requires ingesting and processing vast amounts of data. One of the big value adds that the company provides is related to AI, and how we leverage AI to synthesize and organize that data.

It's a global company, so we're servicing brands globally, and have the joy of complying with various global regulatory frameworks in the space. And are also doing a lot of innovation in the AI space through third-party partnerships and building our own AI solutions.

Brett Mason:

Would it be safe to say that your business is in the business of data, managing data, collecting data, using data, and helping other businesses do that as well?

Andrea Bilbija:

Definitely. It's a very, very significant part of Sprinklr's value add to our customers.

Brett Mason:

Brent, you've been on the podcast before, but can you just remind our listeners about your role and some of your specialties at Troutman Pepper?

Brent Hoard:

Yes, great to be back. My name is Brent Hoard, I'm a partner in our privacy and cyber practice group. I started my career as a health lawyer, and over the years, that turned into a privacy and cyber practice with a lot of focus on HIPAA and emerging technologies. So, it's great to be back and talking about AI again.

Brett Mason:

I'm excited to have both of you here. When we were planning for this episode, I know the three of us really noted that there's so many different topics we could talk about in this area that the two of you are focusing on. There's some key themes that we wanted to hit on today. I think our listeners are going to hear those themes as we go throughout the different areas and talk about the different details. But Andrea, can you just talk a little bit about transparency, bias, and discrimination, and risk and safety? In a big picture level, what's the importance of those three areas for what you're doing with data at Sprinklr and how you're interacting with your customers?

Andrea Bilbija:

Sure. So, I think over the years, security, data privacy, and now AI governance were nice to have. They were things that you would put up in an RFI process, but weren't maybe necessarily the focus. I think that there's been a huge shift for that being a critical component to servicing businesses, sort of a make or break it, table stakes discussion during RFI processes. A lot of what you have to be able to show as a reputable B2B company is transparency with data usage, what you plan to do with it, how you plan to safeguard it, how you help your customers comply with their obligations as a downstream vendor.

A lot of that is now spilling into the AI space. It's not just about data processing, and data storage in the most basic capacity, but also the transformation of that data, how we plan to use it for AI models. Are those AI models shared with other customers? And if so, how does that expose customer data? These topics feed into the overall concept of risk and safety. I think it really comes down to what can you show your business customers to demonstrate that you're a trusted partner to whom they can give or send vast amounts of data. Where they feel like they know what you're going to be doing with it, and conversely, you internally can comply with their various policies, procedures, and safeguards for data use.

Brett Mason:

For our less technologically savvy listeners. Could you explain sort of the difference between what was occurring before AI was being more used heavily in data collection and use? Is there a really big difference? I guess, what I'm getting at, now that artificial intelligence technology is being used much more widely.

Andrea Bilbija:

It's a great question. I think that if we were to take a very, very zoomed out, 100,000-foot view, you could see a lot of through lines. In the weeds, though, I do think there are some differences. So, pre-AI focus, I think it's worth saying a lot of companies have been doing AI for a very, very long time, it just hasn't been as prominent of a discussion point. So, a little bit of it is sort of going backwards over well-treaded ground. But previously, I think the primary focus from a privacy and security risks space was compliance with global privacy laws.

So, think GDPR in Europe, CCPA, and CPRA in the States. Many of those requirements were fairly high level. I think companies had room to sort of come up with security frameworks that were commensurate to the data processing that they were doing, or the risks that they were incurring. In the privacy space, it was really focused on that controller processor distinction. Where am I your business partner? And where am I receiving this data as an independent controller?

I think that those themes are still very present in the AI conversations. I think they very much infuse those AI conversations. What's different, though, I think, is some of what we're seeing with the EU AI Act, which is this concept of high risk, medium risk, or low risk models. Which is much more focused on the end output of the AI model, what it's doing with data, and how that model is being leveraged by a business for critical decision making. Again, some of that existed in the GDPR with the parameters around automated decision making, but it's becoming much more focused. At least on the business side, we're definitely seeing more AI forward or tech savvy customers, include AI addenda in their commercial terms that are fairly prescriptive about how they want you to be managing data as an AI partner. So, a lot of through lines, but definitely starting to get a little bit more particularized, and definitely much more technical in the obligations that we're seeing flown through commercially.

Brett Mason:

Thank you for that. I appreciate that explanation. So, let's go ahead and start at, kind of what I see is the first step, which is the data collection. I know, Brent, you have done some advising for customers and clients that you have about ethical data collection and guidelines for doing that. So, could you just describe, for the listeners, what are some of the big considerations when you're doing ethical data collection? And what guidelines would you recommend for entities that are going to be collecting data using AI or encompassing AI?

Brent Hoard:

Definitely. It all ties back, I think, to a lot of those key themes that we were talking about. If you look at the laws, right now, and the legal landscape in the US, there really isn't any type of overarching AI law. Yet, there are bits and pieces around automated decision making in the

CCPA. Colorado passed its AI Act, but that doesn't go into effect until 2026. We've had some different frameworks through government entities and other industry groups related to AI. But those all really tie – without getting into the details of the laws, regulations, different guidelines, they all tie back to those key themes.

I think that's the key driver for understanding how you want to collect, use, and share AI. If you can remember transparency, bias discrimination, engaging in automated decision making, and then risk and safety, those are big themes to think about. As you're looking to direct collection, use, sharing, or focusing even on collection here at this point, I think it's very helpful to understand and articulate a final purpose. What do you want to do with AI? I mean, AI could be anything, but what is the desired output? What services do you think about offering or enhancing through the use of AI?

So, if you can articulate those and figure out what the end game is, that can be a guide for what kind of data do you want to collect and how you go about it. From a privacy perspective, as a privacy lawyer, we're always talking about data minimization. and trying to limit the amount of data that you're collecting to serve a purpose. It reduces risk, it's incorporated into a lot of global privacy laws, and it's just the best practice. So, if you have that goal, you can minimize data collection or target it, rather than just taking in the kitchen sink. And that's in a sense ethical data collection: what are the sources, where are you getting that from? Are you scraping data off of public sources? That presents potential IP risk, depending on what the terms ar e, where those websites are defining what you can and can't do with the data that's there.

Is it third-party data sources? Is it our own data?

So, that kind of shapes the overall scope of what you want to do. That really lets you drill down, and say, we've got our outcome that we want. We know what we're looking to do. How do we go about getting that data?

Brett Mason:

Andrea, I want to get your thoughts on that as well. But before we go there, I would love Brent, can you explain for our listeners, what do you mean by scraping or the IP scraping for data? Because that is something that I don't think a lot of people are familiar with what that is and how that actually is accomplished.

Brent Hoard:

Oh, sure. There are public-facing websites, so there is code well beyond my technical capabilities. But you can write code to go out and basically gather data from public-facing websites and ingest it into your own systems. One of the few enforcements involved a collection of images and other data off of public websites, that was used for law enforcement purposes. That was the case where they just – the data is public, you go in and take it. But websites should all have a terms of service and a privacy policy. For the terms of service, it describes what you can and can't do on that website. So, before you go take something from somebody else's website, you need to make sure that you're actually allowed to do that.

Brett Mason:

A lot of what I'm hearing from you, Brent, and this reminds me of the conversations we had before, before the podcast, was that, before you get started with this data collection, you actually need to understand why are you doing it, what are you doing it for, and what's going to be the end purpose in order for you to construct an ethical and responsible manner in which to do the data collection? Am I hearing that correctly?

Brent Hoard:

Absolutely, yes, that's a huge help. If you're collecting data, I mean, at the end of the day, the output from any kind of AI model, and machine learning is only going to be as good as the data going into it. So, if you have a clear purpose, you're going to get the right data to accomplish that without ingesting a whole bunch of other data that you don't need that takes up space and costs money to store. Then, if it's personal information, can significantly increase potential breach risk, and other things that you just don't want to have happen. So, I think that is great advice.

Brett Mason:

Andrea, can you talk about how you work with your teams at Sprinklr to be intentional and thoughtful about what you're planning to do with collecting and using data, especially when there's going to be AI used as well?

Andrea Bilbija:

Yes, definitely. I think Brent's point is exactly right. I think this is an area where, as in-house lawyers, you have a real obligation to your internal stakeholders to push on a business plan, or a roadmap. Because so many of the decisions you make around what data is going to be used, or how that data is going to be used will have long-term consequences if you're sort of thinking about it within the framework of global privacy laws.

So, having the forethought to not just say, "Well, let's get all the data we can for any purpose we can," but also being broad enough that you enable your business team to have multiple potential avenues, and use cases available to them, I think is really important. Some of the things that we talk about frequently internally are related to third parties, like Brent mentioned.

So, for partnering with a data source of any kind, what are the terms of those contracts say? Are there purpose use limitations? Does the use of that data in your own AI model constitute product development, which is oftentimes a prohibition on data use when you're getting it from a third-party source. So, being really crisp on the existing sources of data that you already have and how you can repurpose the use case if it varies from the initial reason that you got it. The second angle of that is your direct customer contracts.

So, I think regardless of whether you're a B2B company or a B2C company, you have to think about your first party data that you're obtaining from either your businesses or your consumer end users. And whether you have previously given sufficient notice for the use cases that you have in mind, or whether you haven't. If that means that you have to do some sort of notice campaign, maybe it's just an alert, maybe it's opt in, maybe it's opt out, I think that all three are possible and worth talking through. But from a very, very practical standpoint, that can cause business disruption with your customers, with your sales teams, and you have to be really clear about why it is important enough for you to have that data to merit that campaign.

I think the other sort of practical angle of that is, if you allow customers, again, regardless of whether it's business or consumers to opt in or opt out. How do you track that internally on the relevant data set, so that you're only using the data you're permitted to use? That's a real data governance conversation that you have to have about how data is classified, where it's stored, and how people are trained on how to use that data.

Then. the last thing I'll mention. I think, a lot of companies, AI becomes a part of the overall privacy and security conversation. But one thing that we're really mindful of is whether the commitments that we've made as a company over privacy and security extend equally to AI. Or are there places where perhaps we need to take a slightly different approach, or we just need to qualify certain controls or certain statements, because it's different for that kind of data. So, just making sure that if we have commitments out to our customers, or on our website, or in any of our materials, where they extend to net new uses of data, or could potentially extend to net new uses of data. Are we really comfortable with how those commitments are framed, and do we want to make adjustments so that we're giving our customers and end users the transparency they need to know what we're doing with that data.

So, it's definitely a planning exercise. Do work to save work down the road exercise, which can be hard when AI moves as fast as it does. In the long term, gives your teams a lot more flexibility when there's more of a sandbox within which to operate, rather than enabling one singular use case, which more often than not ends up being way too much of a construction down the road.

Brett Mason:

I know you just talked about a lot of different things there, but one of the things that you mentioned, again, was in the contracting. I think you mentioned earlier that you're seeing a lot of your customers wanting to have that AI addenda. What do you recommend around that? Do you recommend that your company be proactive on drafting those addenda and presenting them to your customers? Or what are your thoughts on that?

Andrea Bilbija:

I think it depends on what kind of AI the company is doing. I think that where it's mostly an integration with third party AI, you probably don't have that much control over some of those things. It's worth being receptive to your customer's paper. Where AI is a critical part of your platform, though, and you're building your own AI, I do think it's worth being proactive and coming up with your own paper. I think much like in the security space where it's very challenging to maintain thousands upon thousands of bespoke security addenda that your customers could be flowing down to you. The same is true for AI, especially if you have a platform-wide AI model that all of your customers use.

So, I think that an approach that I would advocate for and that we've seen success with is being proactive about the controls that we have in place, having a negotiations playbook on standby where it's very clear about where we can budge, and where some things are just inherent to the platform, and really aren't going to be changeable. Then, tailoring that so that we can address customer concerns if we're not able to accept their edits on those provisions.

So, we're definitely advocating for more of that proactive approach. But understanding, to Brent's point, this is a huge gray area. I think we're all building the plane while we're flying it. So, where we've really found the most success is just having open candid conversations with our customers, hearing their concerns, explaining our product and our processes. That often tends to resolve most questions that get lost in translation, I think when you have highly technical addenda, like we're seeing with the AI terms.

Brett Mason:

Andrea, sounds like Sprinklr, it has its own AI tooling that it's including in the products. Is that a fair description of some of the products that you guys are rolling out?

Andrea Bilbija:

Yes. Though Sprinklr has its own AI tooling, we also have partnerships with certain existing LLMs or large language models as they're called. So, think OpenAI or Google's Vertex. We support both third-party integrations and our own internal AI.

Brett Mason:

Brent, thinking about some of our healthcare clients and the health care industry. Are you advising folks who are not building their own platforms but are going to be using tools from vendors? And if so, what recommendations are you making to them about using those kinds of third-party tools to collect data and use data?

Brent Hoard:

It's interesting because healthcare, protected health information under HIPAA or just health information generally, is sensitive information under state laws. It's a whole other wrinkle of everything we've been talking about, just because it's sensitive, it's regulated. So, that is something, anytime there is third-party sharing, that's something you should perk up - you hear that. So, what I would be thinking about in that situation is, what does that third-party data sharing entail? If it is PHI, you definitely need to be thinking about a business associate agreement and putting that in place for third-party data sharing.

Then, other components of HIPAA are going to come into play there too. We were talking about data minimization, having a plan, knowing what you're doing. HIPAA has the minimum necessary requirement that you really are only supposed to use and share the minimum necessary protected health information to get something done. It adds a whole different layer. Not saying you can't do it, but you just need to be cognizant of it in the healthcare world. I would also think about it too, in terms of data collection, and where that data is coming from. Are you actually able to share that with third party?

So, you want to look at both ends of it, where the data is coming in from a collection standpoint? Then, also, who are you sharing with? What are they doing with it? Then, from there, is there any other sharing going on behind the scenes with that third party? So, it does add a little bit of complexity to it, and I think it demands more rigor and understanding what you're trying to do and having a very clear and distinct plan of how you're going to do it. But yes, I think, there are ways to get there for sure, and plenty of interesting applications in the healthcare world. One thing I would add in there, too - deidentification from that minimization standpoint. Can you do it with deidentified data? So, also, consider that.

To deidentify the data, two methods. There's a safe harbor method where you take out certain identifiers from the information, expert determination, which is where you have somebody statistically, again, math beyond my capabilities, but they can tell you there is a very minimal probability that this data can be reidentified. You have much more flexibility, and ability to use and share that data. It's outside of HIPAA.

But, I would also think about for sharing with third parties too, don't lose sight. I know I mentioned it, but don't lose sight of, it isn't PHI, subject to HIPAA, don't forget about state privacy laws too. There are restrictions on how that data can be disclosed, what can be done with it. Then, also, potentially, enhanced reporting requirements for breach and other things. So, you want to look at that, just don't lose sight of, its health information. No, it's not HIPAA, we're not necessarily free and clear. So, look at that too.

Brett Mason:

One thing I wanted to follow up on, kind of in the middle of that explanation, which was great. You mentioned that there are so many advancements that can be made and exciting, maybe even treatments for patients, or ways that we can more effectively care for people in healthcare once we incorporate these AI tools. So, just for our listeners, Brent, even though there are security concerns, privacy concerns, there's a lot of different laws to deal with, whether they're federal, international, or state. At the end of the day, do you think that the incorporation of these AI tools using healthcare information is a positive development if we can do it safely and carefully?

Brent Hoard:

Oh, absolutely. Yes. We have, just working with clients, particularly, just in my personal experience, a lot of imaging related tasks that have been accomplished through AI. They give wonderful reads and feedback. That's actually just coincidentally been the last couple of AI related questions and clients that I've worked with, have been in the imaging space. I personally used some of that for, I think I mentioned this on the first podcast. But take a picture of something weird on your skin – I live in Florida, so I took a picture, send it in. It doesn't replace going to the dermatologist and getting a read, but it can at least give an idea of, "Ah, it's just a benign mole, get it at the next check." Or, well, it could be something odd. So, beyond that, I use that as an example. But I think all of those particularly diagnostic type functions are where I've seen a lot of really cool applications for technology with our clients.

Brett Mason:

So, it may be difficult to thread the needle of all the various compliance and regulatory requirements, but sounds like it's still worth it for us to try.

Brent Hoard:

It is worth it. Yes.

Brett Mason:

Once you do have the data, it's been pooled. Andrea, can you just talk with us about changes you've seen since AI has really entered the conversation and how that's guiding, how you work with your teams on data use once you have it collected?

Andrea Bilbija:

Yes, definitely. I think the biggest thing to talk about in this space is again, data governance. It is highly unlikely, I think, in today's world for you to have consistent data use rights over every single data use set that you get. You can try, you can try to negotiate for it. Sometimes you'll succeed, sometimes you won't, and you'll be faced with that business question of, okay, there's a limitation here. Is it worth it for me to proceed? For all the reasons Brent mentioned, chances are, you're going to say yes, because they're still worthwhile applications of that data, even if not every single use case you had dreamed of is enabled.

So, internally, you have to figure out how you differentiate between those datasets. What can you use for general product development. I'm speaking again, from a SaaS company here, but what's available for general product development? What can you share amongst all customers? What do you need to limit to one customer if it's their particular data set? How do you create an agile environment for your engineers and developers who are working fast to release new features and new products, where they don't have to go through a crazy bottlenecked approval process just to get access to certain data feeds or datasets?

I think that's where really putting the time in to figure out, are you going to leverage metadata tagging? Are you going to have certain in product notices that are available to your tech teams that tell them, "Hey, this data is not available for X use, or it is available for Y use." Those are things that are really hard to do, because most companies already have a lot of data in their systems. I think, for the vast majority of companies, it's often not net new data coming in, but figuring out how to have new uses of existing data. So, it's a really difficult exercise, but one that we think is worthwhile. We think it's important to have that sense of data provenance, and data use and limitations, on that use.

Then, the other challenge, and it's true for so many compliance pieces is that, that stuff becomes stale almost the second you finish making it. So, you have to have really good operational hygiene and processes to keep it up to date for your teams. We were working through that internally, trying to come up with the most efficient solutions that we can. Our goal is always, we call it the sandbox. We want to give our product team and our engineering team, the biggest sandbox within which they can operate freely, within those two parameters. And anything that's high risk, or novel, and ways that present high legal risks get escalated.

So, we're constantly trying to find that right balance of, if you stick to these five or six requirements, you're free to innovate, you're free to develop, you're free to be as creative as you want to be. When it goes above that threshold, that's where you have to pull in subject matter experts on the legal team, or security team, or otherwise. So, we're definitely focusing a lot of energy on the data governance pieces.

Brett Mason:

Brent, I know one of the things that you've done is advised some clients on this oversight and governance piece. Can you talk about recommendations that you make to companies when they're looking to make sure that they are getting it right on the compliance end?

Brent Hoard:

Yes. I think one of the key pieces, just kind of adding on to discussion is with potentially data protection impact assessments. So that, when we talk about high-risk data use, it's not necessarily governance from an organizational standpoint, but enabling higher risk potential data uses. I deal in healthcare, so you're wanting to be looking at those types of use cases. Again, you don't want to stop potential use cases, you just need to be cognizant of what are the risks, examine them, and then figure out how to mitigate against that, so you can do it.

I think that if you build a culture around understanding the risk, it can go a long way. Get the right people involved, perform your data impact assessment, or whatever name they're going by these days. There are a bunch of different variations, but just document it, make sure the right people are at the table. That, to me, has been one of the things that I've seen, is just making sure you have the right people involved. That can be through various governance charters, committees, but you want to have the right stakeholders at the table for the overall program itself.

But also, to the extent there are things, to Andrea's point, do you need to escalate something? Who are you going to have looking at these higher risk data uses? So, it's really understanding within the organization, who do you want to have at the table, who do you need at the table, and then also documenting that. Rather than ad hoc. This isn't a place that I would necessarily want to ad hoc things. So, document it, form your committee, come up with some very clear roles and responsibilities for people, what they're looking at. Whether it's privacy and legal, whether it's the IT team and information security. If you can get all those people together and working, you're going to create a safe environment and be able to accomplish what your goals are for your AI program.

Andrea Bilbija:

If I may add to that, Brett, I think it's such an important point. I do really think we often discount the importance of training sometimes in this space. I mean, I think the classic compliance line is, we don't need you to be an expert on the law, we just need you to know who to go to when you have a question. I think the same is true for AI. I think that, historically, it's been a place where it's been much more engineering product, computer science focused. But at the end of the day, inevitably, a complex regulatory framework is going to be surrounding this.

You should be doing training with your engineers, especially your early joiners who have less experience on, here are some things that you should issue spot for, and here's where you can go if you have questions. And to Brent's point, here's the escalation path for you to follow. I think it's worthwhile, and I think that the more eyes you have sort of overseeing a lot of different complex processes and a lot of different parts of the organization, the better you'll be able to surface the truly high-risk things that you want to be looking at.

So, training during onboarding, putting this as part of refresher training annually, or even just sending very targeted team training out. Whether it's a newsletter, or just a five-minute quick reminder of some critical concepts I do think can go a long way to making people feel like they have a place to go if they're not sure about how to handle a particular AI issue or data use issue.

Brett Mason:

Now, combining with what you both just said, because Brent, you talked about making sure the correct people at the table. And Andrea, you talked about building this sandbox. Who are the people who should be at the table as part of this government committee that's dealing with the AI issues that come up? Andrea, why don't you talk to us about your approach?

Andrea Bilbija:

Yes. I'd love to hear what Brent thinks should be, or who Brent thinks should be at the table. I can talk about who's at our table. So, we have our legal team, which includes our privacy attorneys. So, they're probably the key drivers, I would say of our committee. We have our engineering team who's touching the data day in and day out. As well as our product team, who's sort of thinking long term, how are these AI processes going to be integrated and embedded into our product offerings, what is the story we're telling our customers, what's the value they're serving. To Brent's point earlier, what is the point of the AI model, and why do we need it.

Then, we've also included our solutions consultants. These are folks at our company who meet with prospects and customers, to basically troubleshoot and problem solve for them. So, customer share is an issue, and this is the sort of solutions engineer who's going to come up with the right workflows for them. So, they often have a real eye to what customers are utilizing our technology for, and the exact use cases that our AI is being applied to. We've also included our security team, our product security team. Then, finally, product marketing. We want to make sure that all the work we're doing internally translates into how we're talking about everything externally. So, that's been a really great way to have a consistent through line between internal efforts and external collateral.

Brent Hoard:

I completely agree. Those are the right people at the table. Some things you wouldn't necessarily think about, like marketing. There have been companies that have gotten in trouble with the SEC, because of saying that they have products that are doing certain things, and they aren't, and it involves the use of AI. So, you want to make sure everybody's on the same page.

I think one other potential group, depending on the size of the organization, how you're involved could be procurement. If you have somebody handling contracts and contracting outside legal, just to make sure that depending on the nature of the contract, that you've got the right language out there, either for the customer or if you're looking to collect and ingest data, that you have the right provisions and the right contract that's going out. But I think that is a perfect setup, and that's who you want at the table. The other piece too, I think it helps getting that business input too, because that is, if you have the vision and the goal of what to do, that's going to shape, for legal privacy, how you're building contracts, how you're setting up privacy policies, and other notices, and other things. So, getting all that dialogue in one place is ideal.

Brett Mason:

I so appreciate y'all bringing up the marketing folks and the folks who are crafting the external messages. As a medical device products lawyer who just finished a trial, where fraudulent misrepresentation, and fraudulent concealment were part of the claims about the device, the marketing is important. As we know in healthcare, that can be governed also by the FDA, or the FTC, or different things that are being said.

So, making sure that the communications about a product are consistent and compliant across the board as the person who comes on the back end when lawsuits have been filed. I appreciate that so much. Well, we've talked about a lot of things, and we've kind of mentioned the different regulatory frameworks that are out there. We want to talk some about specifically what the EU is doing, or where companies can turn to guidance, maybe talk about what NIST is doing around these issues? Andrea or Brent, I'd love for either of you to speak on those.

Andrea Bilbija:

Sure. So, I'll start, and Brent, as the resident expert, please do jump in. There's definitely a lot of movement within the AI regulatory space. But to Brent's note earlier, not much of it is completely finalized yet. In the states, we've seen President Biden's executive order on AI development. There's certainly guidance coming out about federal agency use of AI. We've seen the Colorado Bill; Illinois has a biometric processing law that I think very closely dovetails with a lot of these concepts.

As with many things in this space, the EU is certainly taking a much faster approach with the publication and finalization of the EU AI Act. I think that in many ways, it gives companies a really interesting framework within which to operate because it creates tiers, essentially, of AI models based on the type of risk that it presents. And depending on how you classify your AI model, certain compliance obligations apply. So, where your AI model is zero risk or no risk, you have very sort of basic transparency documentation requirements. Then, it goes all the way up to models that they've outright prohibited because of the risk that they present.

I'm sure that we'll see many variations globally about what AI regs look like. But I think this is a helpful framework to start with, as companies are figuring out their AI governance structure. Because I think it's important to go through that risk stratification exercise, not only to figure out where you're going to have to invest the most time, but also to make sure that the use cases you envisioned, as well as the potential use cases down the down the road that your customers could be taking with them are part of how you're governing your AI.

So, I think a lot of companies are using the EU AI Act and the NIST framework to assess their current AI models, trying to be as proactive as possible about what we think other countries or states might include down the road. But they've served as a good template, I think, for everyone to start with, and to at least start talking about AI and risk. Which I think is the critical discussion around AI. AI can have a lot of amazing use cases and a lot of bad ones, if it's not built well, or if you're not clear about how it should be used. So, I think starting our assessment with that terminology in mind is going to lead to much better outcomes, and hopefully better innovation within the AI space.

Brent Hoard:

I think that was a great summary. It is so fast paced and changing. It feels like if you have alerts turned on, it feels like you're getting two or three different alerts about different things going on or proposed. So, it definitely pays to keep up with what's going on because it is very fast paced. I think you can look at it from the perspective of what you're doing. There are those key themes that we've mentioned. They're pretty consistent across all of the different frameworks and regulations that are out there. So, are you engaged in automated decision making? Do you want to use AI to make potentially impactful decisions about individuals, financially, jobs, different things without involving a human in the process? That's going to trigger different state laws, it might be higher risk.

So, I think understanding what you're going to do, but really, at this point, if you're focused on those key themes, and understanding what those are in working to build your model around that it's a very good starting point. As things develop, there might be different nuances or other requirements, depending on the jurisdiction or what industry you're in. I think those key themes are going to stay pretty consistent as we move forward. So it's really a pretty solid base to build your program and to be thinking about it from a governance perspective, what you need to do, what your obligations are. And if you do that - it’s a quickly emerging, changing area, but those are going to be things that -- you can change and adapt, but you've got the foundation.

Brett Mason:

For our listeners, can we remind them when we say NIST and looking at the NIST framework, what that is. Just so they know what to Google so they can find out more if they want to look at the framework that's been put out by NIST.

Brent Hoard:

It is a government affiliated group that publishes a number of different frameworks. There's one that is used for healthcare, and healthcare data for security requirements. It's meant for government entities to use, but widely applied. It's just good advice. A lot of expert thought and care goes into the development of these. I think if you've probably searched NIST, N-I-S-T AI framework, it will get you to the right place.

Andrea Bilbija:

I think another interesting standard worth considering in this space is the new ISO standard on AI. So, the International Standards Organization publishes various standards that companies can certify to, and it's a hallmark of your compliance with certain initiatives. The ISO security standards are very much considered, I think, table stakes for SaaS companies to demonstrate their security efforts. ISO published its standard 42001. It's relatively new, but that's another one that we're sort of reviewing and looking at the controls to understand. For ISO, what does a well-managed AI program look like? What are they looking for when it comes to demonstrating compliance with certain controls? It can be a helpful place to start, especially if your organization is one that's going to pursue that certification anyways. So, it's worth building against now if you have a blank slate to work with.

Brett Mason:

Andrea, Brent, thank you so much for joining me today. I've definitely learned a lot about the ways businesses need to be thinking about collecting, using, and working with data, especially if there's artificial intelligence being involved. Thanks to our listeners for listening in. 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 podcasts wherever you listen to podcasts, including Apple, Google, and Spotify. Be on the lookout, I will probably lean on Brent and Andrea to jump back in and get us updated as the artificial intelligence regulatory framework continues to change. Thanks so much, y'all.

Andrea Bilbija:

Thanks for having us.