

Your Ambient Scribe is Listening- Is Your Policy Ready?

PLUS Staff: [00:00:00] Thank you for listening to this PLUS Podcast, *Your Ambient Scribe is Listening: Is Your Policy Ready?* Before we get started, we would like to remind everyone that the information and opinions expressed by our speakers today are their own and do not necessarily represent the views of their employers or of PLUS.

The contents of these materials may not be relied upon as legal advice. With the housekeeping announcements out of the way, I am pleased to turn it over to Dr. Jennifer Shannon

Dr. Jennifer Shannon: Hello, everybody. Thank you so much for joining us for this podcast today. I'm your host, Dr. Jennifer Shannon. I'm a practicing physician in Seattle, and Co-founder and Chief Medical Officer of GLACIS, where our work focuses on making AI systems observable and verifiable.

The topic today is close to my heart. As a practicing clinician, I'm using ambient scribes [00:01:00] in my practice every day. They're the most widely adopted AI tools in healthcare, but the best practices and liability frameworks are still trying to catch up with the pace of risk that generative AI brings.

I'm very excited today to bring together the unique perspectives from a real-world standpoint, as a clinician using these in practice, as well as the legal and insurance perspectives when it comes to ambient scribe liability.

The question for today is: When an ambient scribe gets something wrong and there's a bad outcome, who is responsible? Is it the clinician who used it? [Is it] the AI vendor [or] the healthcare organization who has deployed it? What might you need to prove in a post-defense claim?

To help us unpack that, I'm really excited to be joined by two experts who bring unique perspectives from both the legal and insurance perspective. I have Kyle Laudadio here from CFC, and Matt Keris with Marshall [00:02:00] Dennehey, who brings the legal perspective. I'd love to have Kyle please introduce yourself and tell us a little bit about your background. Then we'll turn it over to Matt as well.

Kyle Laudadio: Yeah, of course. [I] appreciate you having me here. I'm excited to have this discussion with you guys.

My name's Kyle Laudadio. I lead our US healthcare practice here at CFC. CFC is a specialty insurer, based out of London, and has been writing healthcare for about two decades. Specifically, we have what's called our e-health form, which is designed to provide bespoke coverage for telehealth organizations. We actually introduced that form in early 2017, so we've been doing this exact type of business for about 10 years.

The first few years of the product have definitely been focused more on telemedicine. As the market has evolved, we've evolved with it. It's really been interesting to see the last couple of years, especially this year [with] more and more AI coming into [00:03:00] the fold, right? Everything [has had] some sort of component or use of artificial intelligence in the course of care.

One thing we've done recently is affirmatively updated our language to include artificial intelligence as a coverage trigger within our policy. So, [there is] definitely lots that we could talk about and I'm excited to discuss it with you guys.

Dr. Jennifer Shannon: Thanks so much, Kyle. Matt?

Matthew Keris: Yeah. Thank you. Thanks Jennifer [and] thanks Kyle for having me.

My name is Matt Keris. I'm a medical malpractice defense attorney, so I'm one of the gentlemen with the white hat. I am not on the billboards. Hopefully, for those healthcare providers listening in here or watching, you'll never have to see either myself or a person like me in your career. But there's a good chance that you will at some point.

So how am I involved here? I've had the honor and pleasure of talking about electronic medical records for the last 20 years or probably over that. My discussion has evolved over the years from audit trail and now, obviously, AI is in [00:04:00] my opinion turbocharging the electronic medical record to a point now not necessarily with scribe. It is where the record is actually engaging with the clinician, [that] is very revolutionary.

I've specifically been involved with helping people navigate some of these new issues [and] developing or plotting out where these new landmines are going to be in our cases since approximately COVID time period when I was asked by

the American Society for Health Care Risk and Management (ASHRAM) Journal editorial board to lead an article on the potential risks of AI from a defense perspective. I wrote that, again, during COVID. It's been published. It's been peer-reviewed.

Despite that article being written even before ChatGPT, I've been told by many it still has legs. It's still good and some of the issues I was talking about over six years ago are still relevant.

Teeing this up with my practice, I routinely give institutions advice. I speak and write on this [00:05:00] topic, often, from both where AI is coming in on the healthcare side as well as the legal side of things. I'm happy to be a resource here today for everybody.

Dr. Jennifer Shannon: Thank you. I'd like to start with a real case from my own practice, just to make it more concrete and real. Then we'll use this to frame the liability questions that come up when using ambient scribes in practice.

We know that these tools are highly accurate, but in a small share of cases where they aren't. They create a real liability through things such as hallucinations and even more importantly, and more common, omissions. I've had two such errors in my practice.

In the first one, the ambient scribe generated a note stating that I had prescribed Lamictal for PTSD. That is not an indication for Lamictal. It's not actually what I said because I was actually able to compare the ground truth transcript. This is not always available. If I hadn't caught that and the patient later had a bad outcome, the record would have shown that I had prescribed a mood [00:06:00] stabilizer, Lamictal, that has serious potential side effects, like Stevens-Johnson syndrome, for an off-label use.

In the second one, there was a speech recognition error. The patient had been talking more quietly and reported that they were on Ambien, which was generated in the note. Instead, the patient was actually on Albuterol.

Now, when I read the note, I thought, "I don't think the patient actually said Ambien." So, I was able to go back and actually listen to the audio and catch what they said. Again, one is a controlled substance, Ambien, that has very serious side effects, and Albuterol is an inhaler.

So, here's the question: As clinicians, we review enormous volumes of documentation that is generated by LLMs every day. As a human expert in the loop, we appear to be carrying the majority of the liability for what these scribes produce. Basically, if you sign it, you own it. Matt, I'd love to throw this to you. If we use that case study and let's say there was a claim with a bad outcome, what might discovery start looking like at the beginning? What would evidence be for a post-deployment defensible audit log?

Matthew Keris: Sure. As a defense lawyer, I have to prepare my clients as if I were in the mind of my opponent, and that is the patient or the plaintiff lawyer. With this scenario, I haven't had this yet. But these cases are starting to bubble and come to the surface.

The first thing I would want if I was a lawyer representing a patient in this situation is to try to get a copy of the audio transmission, whether it be in an audio file-- It's my understanding it could also be produced in a written form, similar to a deposition [00:08:00] transcript, where it'll say "Q for a question [and] A for an answer," something along those lines. That'd be one of the first things that I would want to try to get, that recording.

Depending on your jurisdiction and the laws in the jurisdiction, those recordings can either be kept for a short amount of time or for a prolonged period of time. There's no, as far as I know at this point, there's no uniformity in that respect.

Jennifer, you had mentioned audit trail. Audit trail right now is a major contention in our cases because the lawyers on the other side want to know what's going on with the record, particularly after healthcare providers know there was a negative outcome. If they were to ask for an audit trail of interactions with the ambient listening device and proofreading, they would want to see-- The real basic issue is [the following]: Did you have an opportunity to go in and [00:09:00] review your dictation or what the patient said?

Either the audit trail could show that you didn't, or it can show you how long you were in there. Depending on how complex or how long the interaction was, [it] can scrutinize the duration of your reviewing of that record. For example, if it's a 20-minute visit [and] if you're only reviewing the transcription from the ambient listening for five seconds, the argument would be, "Are you really looking at this? How hard are you looking at this?"

[This] goes back to almost like a 20th century liability concept when doctors used to pick up the phone and transcribe over the phone to a live person. Did

they proofread their materials? They wanted to test if you had that opportunity to do that. If you didn't, that could be some trouble for you.

It can also-- I don't want to go too much further on this too, but that would be a scrutiny to you as a clinician. It [00:10:00] could also lead to scrutiny to the institution for permitting physicians to not review transcribed patient visits, which could lead to institutional problems as well.

The audit trail and the transcript as well are the initial things, but there are a lot of other things that they can ask for too, which maybe we'll get into in a little bit. I just want to leave it there at this point.

Dr. Jennifer Shannon: Yeah, those are all great points. I think the challenge right now that we have, to your point, is there's such variability in the data privacy and storage and retention that some companies don't actually keep the transcripts and audio. So, we don't necessarily have a ground truth that as clinicians we can compare it against.

I think this also nicely leads us into a question for Kyle, from the insurance side. How would the AI liability in this case potentially be attributed, or what lines of insurance coverage would this trigger?

Would it be mainly under medical [00:11:00] malpractice, tech E&O, cyber privacy, product liability, bodily injury? I know that this touches multiple pillars within insurance and would love to have you talk a little bit more about how you would assess that.

Kyle Laudadio: Sure. I mean, the short answer to that question is yes. Right? They're going to find ways to find coverage triggered anywhere that they can. Matt would know that, right, looking at different areas.

It really depends on who is insured in the scenario. Right? It could be we're insuring the AI company, the tech company that developed the AI. That's one type of policy. Or [we can] actually have a policy for the clinician, the provider in that case, using the AI. Sometimes they can be both, right? That's where it's important to have one of these bespoke policies that's really custom fit for digital health.

There can be that blurring of, "Was it a good old-fashioned med mal claim? Was it something [00:12:00] that was a result of a cyber breach? Someone corrupted that medical record.

Maybe the scribe did its job, captured it correctly, and it wasn't actually a misheard word. It really did change it to Ambien, or a third-party bad actor changed it to Ambien and that was really as a result of a cyber breach.

Or, did the tech breakdown? Was the algorithm flawed and was hallucinating and caused the issue? [It] could even be hardware if it's actually a listening device that you're using in the course of the care.

The easiest place to have it is on a policy that has all those coverages together. You avoid the multi-carrier finger-pointing and the slowing of the claim being handled properly, right? That's really important.

If I was looking at it from your standpoint, as the clinician in this case, I would still think it's probably going to mostly come back on the medical professional liability. Ultimately, you're the one who should have been signing off that everything's correct.

What I worry about, [00:13:00] and Matt alluded to it a bit, is right now Ambien scribes/ AI is new. It's novel, right? People are going to be skeptical of it and take time to review the notes. On the whole, not everyone's doing this. [They don't] take the 15-20 minutes to really dig in and make sure it's accurate.

They're not really saving time, right? Maybe it's adding time to the visit? Months go by. Years go by. It's always right, and then all of a sudden it's, "You know what? I'm busy today," and it's the rubber stamp. You start to trust it more. That's when the mistakes happen. That's when the bad day happens where you prescribe someone a medication that they have an allergy to or something that doesn't coexist well with Ambien and they have a bad outcome, right? That is what I worry about the most.

It's not just, "Okay, the third-party tech had an issue." It [could] also be when you're onboarding them that you don't just trust what they're telling you, that you're verifying. You're almost doing some underwriting [00:14:00] as a clinician that we would do if we were underwriting.

What was their data set? How are they auditing? How are they measuring for drift and detecting that before it becomes an issue? If that all checks out for you as a clinician or a system, how are you training your clinicians to use it?

In that case, is there a certain expectation of how much time you're supposed to be putting in to review those notes or do spot audits across all of your visits weekly or monthly to make sure that it is working properly. Because like you

said, over time it's going to be more ubiquitous and we're going to all get very much used to it. That's when I think the real poor outcomes are going to result.

Dr. Jennifer Shannon: Yeah, you bring up a really excellent point about human in the loop which right now, as a clinician, is the default liability. But human in the loop is not human in control at this point in the sense that we don't have a lot of visibility into how the AI vendors are developing and deploying these tools [or] [00:15:00] how the healthcare organization is or isn't monitoring them post-deployment. We don't have an evaluation approach that we can use to really look at this.

Kyle Laudadio: Yeah ... that, that human in the loop is a good talking point we all use, but it means something completely different to everyone, right? So how do we define that? You could ask five different people [and] get five different answers. It probably means [something] different to the tech companies than it does to a clinician, right?

Human in the loop could be you just signed off on it and approved it. It could be that you're the one that turned it on. It could be that you knew that it was happening in the background and is autonomous. I think using that just as a blanket term-- I think that's also something that the users of the technology need to understand to make sure that they realize [and make sure] what's in their services agreement lines up with what their expectations are.

Dr. Jennifer Shannon: Exactly. And from an insurance perspective, you [00:16:00] mentioned it could have had a failure at multiple points. Are there things that, from an insurance underwriting standpoint, you are looking for? How is this changing with gen AI risk when it comes to, say, insuring an Ambient Scribe company versus a healthcare org? Are there things that these companies and healthcare organizations can be proactively doing in terms of testing, evaluating, and showing that controls are working?

Kyle Laudadio: Yeah. I think to take it from the AI standpoint, from the developer side, we are really underwriting to where they're getting the data, number one. Is it their proprietary IP that they're feeding into that? Did they get those proper permissions and privacy sign-off from the patients that the model was trained on?

How wide is that data set? You know, a couple hundred patients' data being put in there is probably not going to be enough. [If you have] hundreds of thousand[s], you're getting a little bit [00:17:00] better. It's going to be probably more accurate. Looking at the diversity of that patient population or information

that went into that, what frontier model [was] it built on [is important], right? [You need to know] that they're not just taking ChatGPT and retrofitting it for their purposes. [It is important] that it really is something that's proper for use in healthcare.

Once it's implemented, how do they make sure that it continues to be accurate? This cannot be set it and forget it. We all know about the hallucinations and the drift that go on. That's fine when you're using it to come up with a great recipe for when you're cooking dinner, but when it's people's lives on the line in healthcare, the stakes are much higher. We can't have even 1% go sideways here. We want to make sure that they're doing those continuous audits. They're updating it. They have their finger on the pulse, no pun intended, [and] that this is continuing to work properly or improving. We want to see the accuracy rates upfront and ongoing accuracy rates for that.

[00:18:00] From the clinician side, one thing we look at is [that] you're using this. Hopefully it does make your practice and your care more efficient. [Hopefully it] takes that administrative burden off, which is great.

What I don't want to see is all of a sudden your patient encounters skyrocket. It's good if you can increase it--I don't have any set number. A decent percentage is fine. But if all of a sudden, you're doubling the number of encounters you're doing a day, that's probably not great.

What I'd rather see as an underwriter is that you're using this to have more meaningful, thoughtful visits with your patients. You can put the pen down and not [have] to be taking furious notes. [You can] really listen to them [and] have a conversation. Maybe that means you have an extra 5 or 10 minutes with the patient so the ongoing outcome is better rather than saying, "Oh, I can fit five or ten more visits in today and bill for those."

I think Matt would know. You see that in discovery that the visits went up. Really, it's just the kind of the old profits over people angle, right?

Matthew Keris: Right. They're going to go look into what the quality [00:19:00] is and if there are any type of responsibilities they can ask for contracts or documents, as to the individual provider, in terms of limitations or goals in terms of how many patients they should be seeing.

I do know from my own personal experience that some clinicians were told to limit their complete visits, from intake through vitals and weights and measures

and history and physical and diagnosis, into 15 minutes. And if they didn't do it in 15 minutes, they were going to get docked.

So yeah, that's always something ripe. That would go, again, to the institution for having such a thing. And then potentially as an argument that you're trying to discourage good care. And again, it goes to the clinician.

Patient care is paramount, correct? You're not even following that rule when it comes to patient visits because you're limiting yourself to an arbitrary rate of 15 minutes. So, it cuts both ways to the institution and the clinician.

Dr. Jennifer Shannon: Those are great points. [That's] the difference between what is efficiency versus [00:20:00] productivity and what you do with that additional time. Often, I find that I'm spending that time reviewing the notes that the ambient scribe generated to make sure that they're accurate, and so it may be a wash in many situations.

Matthew Keris: If you don't mind, Kyle had mentioned the oversight before.

I think that's to the extent-- it's not going to be out there immediately. That oversight is critical because we all know how there's the drift [and] how accuracy degrades. There's even publications out there indicating that institutions know they need to be monitoring the drift, et cetera, and they're simply not. They don't have the money, resources or whatnot.

I definitely, going back to Jennifer's--- one of the initial questions she had for me was, "What kind of discovery's out there?"

There's going to be requests about monitoring the accuracy of these platforms, whether it be ambient or otherwise. It's going to be there. The answer right now is it's not being followed. It's not good, which [00:21:00] won't necessarily translate well in the defense of these cases.

Dr. Jennifer Shannon: Yes, we just got back from a legal summit, at Coalition for Healthcare AI, where the topic really was shifting from not just point in time, approve, deploy and then not monitoring post-deployment as much, but really focusing on the post-deployment. Not just even model accuracy, but actual operational accuracy. Is it behaving as designed? What controls are in place? How do we show and prove those controls are actually working?

I think this leads us into a topic everyone's interested in when it comes to Ambient Scribe. Some of the first lawsuits that are coming out are related to consent, or shall we say lack potentially of consent, versus informed notice.

Matt, I think when we had spoken [about this] previously. You had made that great distinction. Can you talk a little bit about the importance and distinction between [00:22:00] consent and informed notice? What should we be aware of when we are using these tools?

Matthew Keris: When I was training as a young lawyer, informed consent was a concept about surgeries when physicians went through the risks and benefits of a surgical procedure as well as the risks and benefits of the alternatives. The way I understood it and the general rule of informed consent across this country is you need to have it before you do a surgery. Spell it out. Write it out. Have a dialogue so that the patient has this, again, informed decision about a surgical procedure.

What troubles me, and I've done it myself and I started that way until I heard others mention that maybe we shouldn't be using the term informed consent with AI use. Where in the surgical context, we're asking for permission to do the surgery. With the term informed consent, you're almost asking consent to use AI.

As each day [00:23:00] goes by, each month and week, we all know AI is being used more and more on a granular level that perhaps even the clinicians don't even understand. The informed notice is a better way of that I see in describing or having a discussion with the patient because it's notice not consent.

Hypothetically, if a patient says, "No, I don't want to give you my informed consent for any aspect of AI in healthcare," that's almost impossible now between scheduling the administrative scribe. These systems are going to go down. So, consent, it's out there. It's being used a lot.

Furthermore, when doctors are giving or trying to obtain informed consent of a patient, they're explaining the risks and benefits of the surgery. Do you want your really super heart surgeon talking about [00:24:00] the risks and benefits of AI? Do they know? Would you trust them? Are they articulate and knowledgeable enough to talk about those types of things? [Do they] understand, truly, how it works? My bet would be not really.

Do you want these clinicians to be the spokesperson for how AI is being used in your system when these systems are updated every day? There's the latest and

greatest in everything. Bells and whistles change all the time. Consent makes you think that you need to actually have that. Again, because AI is used so granularly, you want to let them know that it's being used. But the consent part of it, be careful about.

And one other thing, I know I've been taking a lot of time. I'm telling you not to get informed consent, but for the clinicians out there, use your best judgment. If a patient's at a major crossroad, surgery or not---this type of chemo versus this type of cancer treatment--- If the physician is truly [00:25:00] using artificial intelligence, whether it be from the system or on their own research and they're using informed consent in order to give what they best recommend, I would have the conversation with the patient that you did that.

The last thing you want is for the patient to learn for the first time you used some type of platform after a bad outcome. I almost can guarantee that the response will be, "Had I known that you trusted the computer in giving me a number one thing to do with my treatment and care, I would have never agreed to the computer telling you what to do."

I can't give a bright line test or example. Use your judgment on that because this is still new. This is still a concept. Data centers in AI is not a popular thing right now. So, I don't anticipate that changing anytime soon as well. Yeah.

Dr. Jennifer Shannon: Yeah, those are great points.

Kyle, from an insurance perspective, what would you [00:26:00] want to see around notice or informed consent [and] documentation of this, especially with some of the recent legal cases that are still in progress?

Kyle Laudadio: Yeah. I think for our purposes, Matt just gave a really thoughtful answer to it. I agree with him, at least for the very basics for the conversation we're having on Ambient Scribe. At least they were given notice that this is being recorded or not and they could opt out if they wanted to, right? Given the opt in or opt out, it's very clear.

We've heard varying degrees of what that means, especially in digital health or telehealth. When you log on it gives you a banner that you have to click and say you understand that it's going to be recorded or you can opt in [or] opt out there.

In person, it's whether there's signs. They just verbally attest that you have their approval to use the scribe or have to sign off on it. We just want to see that it is

[00:27:00] documented somewhere that they were aware that it was happening because it is a huge privacy concern.

One thing that we really haven't talked a lot about, and Matt alluded to it earlier, [is] how long they hold on to the transcripts or the recordings. I get it, for discovery purposes to defend that maybe the clinician did make all the right decisions, [it] would help you in your defense.

From a privacy standpoint, that concerns me. What you don't want is a third party being able to have access or hold on to PHI, especially recordings of healthcare visits. If that third-party tech company that's developed the AI somehow has a breach, that could be a systemic event across multiple health systems, hospitals [and] clinicians. That's a really bad day. Th[ere] could potentially [be] millions of impacted individuals there.

At the same point, even for you as [00:28:00] the direct user, if you have a breach, now you're not only holding on to the written medical record. There is this rich data that's in a fully recorded visit. So, I do worry about that.

That's something that when you're entering into your contracts with a tech company to have these scribes or other AI uses [is important to have an understanding [of] what kind of access do they have to this, if any. How long will it be held in the cloud before it's deleted? If it's deleted, is it fully deleted? What's retained? Is it just high-level notes and takeaways from the visit, or is it a full transcript, right? It really depends on the scribe you use and the contract that you develop.

So, you have to understand very clearly what you're entering into and how that data is going to be used or held in the long term.

Matthew Keris: And you know what, running with that, Kyle, too--

Kyle Laudadio: Mm-hmm.

Matthew Keris: The people who are requesting this information, the plaintiffs' lawyers, communicate well. They coordinate their requests, and [00:29:00] they sometimes give the perception to courts that one size fits all in terms of what a system can keep, how it maintains things, what can be produced. And if health system A can do it, why can't health system B do what health system A does? To the extent health system B can't do what A does [is] they allege conspiracy, or cover-up, dragging their feet and those types of things.

That's a critical piece what you shared with in terms of what each system does, and knowing what your system can do. Your system may be able to do X, Y, and Z this year, but next year it changes. The year after that it changes again.

A constructive criticism that I have amongst my peers and my clients is that we don't do a good enough job being historians on our systems, whether it be audit trail, scribe, EMR. Things change constantly. We're being held, in our cases, we're being held accountable to one year. What was proper in that year? When discovery [00:30:00] and these things change so rapidly, like they do, it's very difficult for us to reconstruct four or five years down the road from that from what we do now.

If we can all be better historians, and I can't give you another bright line test on how to do that, that would make things easier for us.

Kyle Laudadio: Yeah. And Jennifer- Yeah ... do you mind if I kind of turn the tables on you with a question?

Dr. Jennifer Shannon: Absolutely

Kyle Laudadio: I think in your examples where the scribe had a couple inaccuracies that you were able to catch, what kind of opportunity do you have to provide feedback on that to report this to the tech company? That's something that I [would] like to know because you're going to make them better, but also document that you caught it and called this out. Is that something you were able to do? Was it easy?

Dr. Jennifer Shannon: Excellent, excellent point. You honed in on, I think, the exact gap. In these two cases I did actually have the audio transcript that I was able to look back at just because it was that [00:31:00] specific ambient scribe.

Whereas, as we were saying, there's such wide variability in data privacy and retention. I know some ambient scribe companies just automatically delete it or they might just store it for two weeks or 30 days. In this case, I had a ground truth to look back and say, "Well, did the patient really say that? How did the LLM actually come up with that connection [that] I was prescribing Lamictal for PTSD? Is there anything in the notes that would even allude to that?"

For me as a clinician, having that transparency to compare was huge. But in a lot of cases, and again, imagine a clinician reviewing 500 pages of this type of content, you're not necessarily going to go and check against the real transcript for each of these.

The other important point that you bring up, Kyle, is we want to be able to improve these systems with the [00:32:00] feedback. If there is a hallucination or an omission, we want the AI vendor to be able to go back and look at their chain of custody, audit trail, or [have] the ability to trace back how that AI system actually came to that conclusion. These systems are not deterministic.

[We want to] be able to improve and hopefully fix that with the data that we provide. There is a little box where we can put in our feedback. And a lot of times it's more basic, like just a thumbs up or thumbs down. Then it asks you, specifically, why and you have to do it. But clinicians are really busy, to be honest. I don't always fill those out. So, I think there's a real gap between how AI vendors, clinicians, and healthcare systems can really work together to ensure that post-deployment monitoring and improvement are occurring.

Matthew Keris: You know that's an interesting topic in terms of the [00:33:00] duration on how long these recordings are lasting. I mean, a lot of health systems that I'm involved with are deleting these audio transcripts as soon as, they can under the laws in that particular state.

But Jennifer, you had a great point. You went back and you actually listened to the transcription and that was beneficial to you. It could be beneficial down the road. So, I can argue it both ways. I'm a lawyer, that's what I'm trained to do.

Dr. Jennifer Shannon: It's a little “[You’re] damned if you do [and] damned if you don’t,” right?

Matthew Keris:

Yeah, “damned if you do [and] damned if you don’t.” But, just as much as it can hurt you in a particular case, just like audit trail, audit trail can help and hurt. Oftentimes our opponents like to use it against us. If we don't have it anymore, they'll give the illusion that we have something to hide and that's why they were destroying it. It could be, again, compliant with the law and the privacy rules in that jurisdiction too. So, it's an interesting debate right now.

Kyle Laudadio: Yeah, I suppose it could be-

Dr. Jennifer Shannon: Yeah

Kyle Laudadio: Jennifer was finding a mistake once a day or once a week. Great, [00:34:00] she's reporting it. And then [one day she] doesn't catch one.

They're like, "Well, you knew that's always making these mistakes." It's almost less defensible to not catch it, right? So-

Dr. Jennifer Shannon: Right

Kyle Laudadio: ... so the good thing is that she is reading back, but now it's like it almost is a higher expectation to be looking out for those mistakes if she has been catching them before.

So yeah, you can't win either way, I suppose.

Dr. Jennifer Shannon: No, it's a great point.

Kyle Laudadio: With these lawyers getting involved.

Dr. Jennifer Shannon: Yeah, the lawyers. This came up at the legal summit, which was, "What does that standard of care right now look like for a clinician and a AI interaction? And if 80% or more doctors are using AI now is it the standard that you should be using AI?" If you don't, then you're actually not following the standard. And same thing for ambient scribes in terms of how they ensure accuracy.

I feel like we could talk forever, but I'm aware that we're already coming to some concluding thoughts. I would love to kind of have each speaker [00:35:00] think about one big takeaway or recommendation that they might make to, you know, Matt, from your standpoint, a clinician using these tools, ambient scribe specifically in practice. What are some proactive things that we as clinicians can be doing to protect ourselves, if you will, if there were to be a bad outcome?

And then Kyle, from your standpoint of insuring clinicians, AI vendor s[and] healthcare systems around this, what types of things and recommendations might you make to them around ensuring that they could have a post deployment defensibility stance if necessary or to be able to protect themselves?

Matthew Keris: I'll start. You know, I've been, I've been practicing for about 30 years. When I first started, we had paper records and EMR till now. And I'm so glad for ambient listening for clinicians now because physician [00:36:00] burnout was so bad for so long. Everything I'm hearing about ambient listening is that people are starting to get their lives back again because they're not spending two or three hours at night, at home on their time when they could be with their spouses and their children, going through charts.

In that respect, I'm so glad because this is finally, at least in my mind from a person who's been observing healthcare for close to 30 years, we finally get some benefit of EMR that was promised us back in 2005 of all things. I'm thrilled about that. So again, showing my age, it was a problem in the 20th century about proofreading. It remains the issue now.

For clinicians, you have this extra time. Take the time to just make sure it's right. Kyle talked about this, there's going to be a trust in it but still [for] the very best ambient listening, it's never going to be 100% correct, never.

It'll get a little bit better every year, [00:37:00] but some of the statistics that I still see is that it still can get it wrong 25 to 30 percent of the time. We shouldn't expect that to change that much. They've been trying to do, even in my practice, they've been trying to do that. There's a big difference between afebrile and febrile, folks, right? We're no pain and in pain. It's those little things that are incredibly important.

I can also tell you from a clinician that the patients love the detail of these notes because they look much more robust. They have access to these records too. So, this is going to be a good thing for everybody.

For you and your patients, we just need to make sure we get through the bumps and the potholes and we learn to do the right thing the first time. Again, proofread, proofread, proofread. Make sure it's good and you may not even need to see me, as a defense lawyer.

Dr. Jennifer Shannon: That's always the hope.

Kyle Laudadio: From my side, at a very basic level for the insurability [00:38:00] and defense is making sure that your underwriter and insurer understand what you're doing and what you're using, whether you're a tech company providing developing these tools or insuring the clinicians that are using those tools.

Most policies in the market right now are silent on AI. I'm hearing that at the beginning of the next year, there actually could be exclusions for AI. There are going to be varying degrees of whether it's an absolute exclusion for the use of it or if there's going to be carve backs for certain types of use.

The important thing is, and Matt sort of mentioned this earlier, AI means a lot of things. So, if they're excluding AI in healthcare, does that mean you can't use it in your billing or in the intake or scribes for taking notes for the visit? Or does

that just mean aiding in clinical care? If you have those exclusions, understand it and make sure that it's not taking away a huge part of your liability,
[00:39:00]

Like I said, CFC has introduced affirmative language for the use of artificial intelligence and development of it. There [are] some other markets that are following us, not just to sit here and sell us, but we've leaned into it. I suppose that there's going to be more competition in that space in the next couple of months.

The most basic level is just to make sure that you're transparent and detailed about what you're using and developing when you're up for renewal or getting your first application in for insurance. [You must be] having these conversations with the underwriters to make sure that they understand it and that their intent is pretty clear. I would even be a little bit nervous with the silent policies on it right now, unless you have some sort of documentation that they understand that you're using these scribes in the course of your practice.

Dr. Jennifer Shannon: All excellent points. With that, I will conclude by saying the big takeaway for me is that ambient scribes and their outputs can and are really becoming part of [00:40:00] potentially an evidentiary record.

Should these tools need to be deployed with the same kind of risk management evaluation and post-deployment monitoring as any potentially high-risk clinical AI tool? And, more specifically, as a clinician signing the note, we often say, "You sign it, you own it." The defensibility of that really depends on me, as a clinician, doing meaningful and careful review of the outputs, knowing that, you know, these tools will never be perfect.

Ensuring that healthcare organizations have the appropriate AI governance policies and actual post-deployment monitoring in place, and then the ability to really reconstruct and replay what happened if there is a claim, ultimately, we can all work together to define and work on what this looks like in practice.

I want to thank both of you for joining us today. This was an excellent conversation. Really, really loved it.

Matthew Keris: We could have gone on. Thank you. I- honestly, I could go on and [00:41:00] scratch the surface- Okay. We need we need a part two

Dr. Jennifer Shannon: So good.

Matthew Keris: Yeah. So good.

PLUS Staff: Thank you for listening to this PLUS Podcast. If you have ideas for a future PLUS Podcast, please complete the Content Idea Form on the PLUS website.