

Host: Hello and welcome to today's webinar, *AI in Mental Health: Part 2 – AI in Clinical Decision-Making*. This topical webinar is paid for by Upheal, a secure AI-powered progress notes designed by clinicians to save mental health providers time. Some important points before we get started.

Leila Anderson, the presenter for today's webinar, has a working relationship with Upheal. As such, she has both a financial and intellectual interest in the webinar content. The purpose of this presentation is to provide a balanced view on the therapeutic and potentially harmful uses of artificial intelligence in clinical decision-making. APA does not endorse any products or services. The content was created by Upheal and does not reflect the views of APA or its editorial staff. Email advertising@apa.org to learn more.

Next, this program does not offer CE. However, we will email everyone watching live today a certificate of attendance. Certificates will only be issued to those who watch for a minimum of 45 minutes. A recording of this presentation will be emailed to everyone in two weeks' time. The recording will include the presentation slides.

During our time together, you will be on mute. You can communicate with us using the Q&A box located on your webinar screen. The Q&A is enabled for attendee interaction with submitted questions. Attendees can view all submitted questions and upvote those they find important. This helps the presenter identify questions that many attendees want answered. A link to the presentation slides will be posted in the chat box of your webinar screen. If you miss them, don't worry. We will email them to you with the recording.

Now let me introduce our presenter. Leila Anderson is a systemic therapist, supervisor, and clinical AI prompt engineer at Upheal, whose practice is at the intersection of clinical rigor, ethics, and emerging mental health technology. She brings over a decade of experience across residential, correctional, psychiatric, and community-based settings with specializations in complex trauma, shame, and high-acuity clinical presentations. Her therapeutic approach integrates IFS, EMDR, and DBT with particular emphasis on the self-of-the-therapist work and ethical decision-making. Leila serves on the TAMFT Board of Directors and speaks regularly at state and national conferences. Welcome, Leila Anderson.

Leila Anderson: Thank you so much. I'm really excited to be here. I don't know how many of you might have caught last month's webinar. We were talking about clinical decision-making for clients specifically. When you have clients who are already using AI, how can you, as a clinician, help those people interact with AI in ways that are effective and less likely to cause harm, more likely to bring benefit?

We had a really great conversation. So many of you submitted really great questions, followed up with very thoughtful emails. It's incredible to see how much care and attention and thought is being brought into these admittedly very complex questions. I feel like it's worth addressing the elephant in the room here in some ways, which is that we all hear so much about this particular topic. We're inundated with a lot of opinions, a lot of perspectives on this stuff. I think there's genuinely some complexity in what the most ethical way to engage with some of these tools is.

A lot of what we're going to be doing today is talking about clinical reality first, so how AI is already being heavily used by clinicians in using tools to shape clinical decisions. We're going to talk a little bit about a decision support framework, piecing apart where it is relatively safe with clinician sign-off for AI to support judgment and where it can't, where it's very important that something remain exclusively the realm of clinician or very primarily the realm of the clinician.

We'll talk about some decision-level risks, automation bias, fabrication, deskilling. We'll talk about consents, PHI, evaluating tools, what things you can look out for. I'm a huge fan of understanding exactly what questions should be asked of a particular tool that you're going to implement. If a company or a sales rep or whoever can't answer your questions about the way that your data is being used, the way that guardrails are put around the risks that we know exist with AI, that's probably not the safest tool for your clients and for yourself. We'll discuss that.

Then we'll talk about what it looks like to take next steps. Maybe you are already using AI in some capacity, but you would like to expand that or considering whether or not you should expand it. Maybe you don't currently use AI, but you're interested. Maybe you really feel uncomfortable with AI and don't ever want to use it. I think all those are valid perspectives here. I know there's a single right thing for us to come to.

I will say at the beginning of this, for me, personally, my background, as you may have seen in my bio, I've done a lot of different things. During my career, I have been a therapist since-- I was licensed in 2015, so it's been a little over 10 years, and spent about a couple years in corrections facilities, so working in the jail where there was a ton of emphasis specifically on very clear but very brief diagnosis and assessment and documentation that functionally mostly served to reduce the liability of the jail. It's a very, very narrow point to my documentation. I'm assessing an inmate at the jail and trying to make sure that everything I assess is around making sure the jail doesn't have liability if something happens to that person.

That was my first introduction to documentation outside of a school setting was the need for a very high level of precision and also a very low level of excessive detail, so making sure that I'm not putting something in that might later come back to bite somebody when they're going to court in a way that could potentially be harmful to that patient.

I then worked in acute psych, which again is a lot of documentation, but relatively quick, relatively brief. There's not going to be a ton of detailed clinical synthesis or conceptualization to a large degree. Then after that, I spent about five years working in residential treatment where a lot of the work I was doing was geared more toward "How can I write notes that help justify continued length of stay? I think this person with a history of chronic relapse needs more time in treatment, and the insurance company thinks they ought to be discharged after two weeks." How do we come to a conclusion here that supports the client's ongoing sobriety if that client wants to be in treatment and wants to be able to get well?

I started documenting in more detail and started understanding more of the compliance factors, the clinical nuance to how do I create documentation that really serves this client's interests in this particular setting, in this particular moment. That

was when I really got to appreciate documentation as a tool for us becoming better clinicians, but also as a tool for aiding in my own clinical decision-making. As I was documenting these sessions, I was realizing I do have good reasons to want this client to be here for an extra week or an extra two weeks, or however long I felt like they needed to stay in treatment because I was documenting it, because I was thinking through it as I was writing that note.

When AI-assisted tools started coming out, my initial thought was, honestly, skepticism. What happens if this tool takes away the ability of especially newer clinicians to think through those things in a deep and complex way? What would have happened to my professional journey if I hadn't had that opportunity? I think these are real questions and something that we can't ignore. We can't take these things lightly. These are genuinely important things to think about.

That's leading aside the question of whether or not AI should be helping in decision-making. I think through does this client need a higher level of care or not? Is this a CPS report or not, Child Protective Services? In most cases, if I really have a question about this, I want to be calling a human being and talking to a peer, talking to a mentor. I would expect my supervisees, the newer clinicians who I supervise clinically, to call me with those things. What kinds of questions? Where is the line where I might be comfortable with them instead consulting a purpose-trained AI tool versus calling me? Again, these are complex questions.

Part of why I ended up working with Upheal, I do feel strongly that Upheal does a really good job with these things. I wouldn't work for them if I didn't. Part of why I started looking into this was just recognizing as much as I have taken a lot of value from my documentation journey and have not felt like documentation was an overly strenuous burden.

I've also worked with a lot of clinicians, especially neurodivergent clinicians who have really, really struggled with getting notes on paper in a way that does accomplish those ends of helping them process through things, of justifying continued care, making sure that client's needs are actually communicated clearly to an insurance payer. It feels unfair that people who are incredibly good clinicians but struggle with that process that I connected with pretty easily of just getting words on paper to justify a client's needs. It feels unfair that they would have such a harder time than I did.

Creativity is an enormous part of the clinical process in my mind for me and my practice, thinking through even before AI, but then as the beginning of AI tools began to be developed, what if we could create something that did make the documentation piece easier while also not resulting in deskilling, not resulting in a lack of clinical depth, a lack of appropriate judgment being applied?

All that to say, I think there's a lot of value, not necessarily in every situation, not necessarily for everyone. I realize this is an enormously long introduction to this topic, but I don't want this to come across like I'm just shilling for a tech company here when I say I think these are very valuable tools and can actually create more depth and sophistication to the way that we document and interact with clinical decisions, in the same way that those of us who have been doing this for a while have always found ways to deepen our skills and find more sophistication.

Unfortunately, a huge number of mental health providers burn out before they get to the point where they've been in the field licensed for 10 years. That's not okay. That's an enormously wasteful thing. We need the minds and enthusiasm and passion of people who are coming into this. Newer clinicians, newer psychologists are such a resource because they do tend to be so creative and so willing to jump in and try things that other people haven't necessarily.

A lot of people who reached out to me after the last webinar were students or newer psychologists. It's so fun to see that level of enthusiasm. Again, all that to say, apologize for the super long introduction, but it felt like there was some groundwork that was beneficial to explore before we dive in.

Now, AI is probably already somewhere in your workflow or your colleagues workflows or the workflows of some of those psychologists you know and are connected with. More than half of psychologists have experimented with AI tools in the past year. That could be something like-- ChatGPT has a ChatGPT for Clinicians now that you can get if you have an NPI that is-- it's HIPAA compliant. You can play around with that. It might be an actual tool like Upheal, which is actually a fully AI-native EHR at this point. It's an electronic health record system that includes AI throughout it.

There's everything in between. There's a lot of different ways to engage with AI. There are skill-building platforms. If I wanted to get better at motivational interviewing, I could go and have AI-integrated practice in motivational interviewing. AI originally showed up through documentation. You all probably know this. The ambient note-takers, the things that can transcribe your session and then create a note, or you can summarize or dictate your session, and it will create a note that's well-formatted and appropriately structured. These things have been around for a while.

In general, any of these tools, they can't vouch for themselves. They can't say, "We're doing a good job. We're compliant. We are going to reflect exactly what happened in session." They still need to be clinically validated. Any AI tool should be clinically validated. Never put your name on something you haven't actually read through and haven't actually confirmed the validity of, because there are a million things that can happen, just fluke things.

Your client's mumbling and something gets misheard and mis-transcribed, and that means that it gets misrepresented in a note. Clinical judgment can never be replaced, but AI has been here for a long time. If we start with openness and clarity into what AI can and can't do, then we can find some safe and ethical uses for it.

Profession is catching up as far as providing guidelines for what ethical use of AI looks like. In June 2025, APA had an ethical guidance document that included things like you need to have informed consent and transparency with your client, so don't be using AI in any capacity if your patients aren't aware that you're doing that. That includes, in my opinion, what I tell the new clinicians I supervise, is even if you're talking about a case in a de-identified way with a tool like ChatGPT, it's important to have visibility into what you're doing when you talk with your clients about this, because this is a new tool and clients deserve to know what exactly they're consenting to.

Bias awareness. If I have a client who is from a non-dominant culture, AI is not necessarily going to understand the nuance of that particular system, that particular client in a very sophisticated way because AI is overwhelmingly trained on dominant culture, on cultures that are more white, on cultures that are more male. If you have any client who's not part of that, what AI tells you may not be as accurate.

I think about many, many clients I've worked with through the years who I've had to go through and consult with colleagues who shared that client's culture so that I could understand in a deeper way how to be competent with that client. It's much harder to actually trigger AI to do that kind of-- the genuine self-reflection versus the surface level, like, "Oh, no. You caught me. I wasn't being culturally competent," which AI will apologize for anything, but that doesn't mean that it's going to go back and change something.

Then accountability. You as a person with a license, a person who can be interrogated about the choices you make, you remain accountable for the decisions that you ultimately make. That's a huge difference because you can't just say, "Well, AI made me do it" and call it a day if your patient has a bad outcome.

Here we're going to be talking about a decision-making framework for working with AI. AI is never going to be more than decision support. It will never be the decision maker. AI is really good at organizing information, surfacing patterns. It doesn't carry accountability. We were just talking about this. It's a pattern-level thing. If it's like a-- I've seen it be super useful for things like I might say, "I just noticed in this most recent session that the client's current partner has a lot of characteristics of the client's mom. Has that come up before?"

Maybe I've been seeing this client for a year, but with modern search tools, with what we call like "RAG," Retrieval Augmented Generation, it can actually search back through my previous sessions and be like, "Oh, well, yes. Well, actually, this has come up three other times, and here's what was going on at these times. Here's some context for the similarity that you're just now noticing."

I think I'm good at noticing patterns. However, I see a decent number of clients. I do other things in my life as well, and there are times when I might miss something and having that ability to go back and say, "Well, hey, I already noticed this. Has this come up before?" That can be super beneficial. It can be a second opinion. There are times when I might question, "Is this truly anxiety or is this OCD? Do I need to be working with this client on just general de-catastrophizing and thinking about the prevalence of negative things happening, or is this OCD where I need to be preventing the response by not dipping into reassurance-giving?"

Those kinds of things where I'm thinking through, I think, based on X, Y, and Z, this feels more like OCD to me. I could ask an AI tool potentially, "What are the features of OCD versus anxiety? How do each align with this client's presentation?" I can consider it. I can go through pros and cons personally. I think if it's a tricky case, it's still always worth consulting with a human. I'm, as a clinical supervisor, a huge fan of talk to your people, talk to the actual humans who share your license and understand what your responsibility is in these situations but it can be helpful.

AI can augment effort and documentation on certain tasks, any decisions that are defining care. If I'm making a choice about do I proceed with treatment as though this is OCD or as though this is anxiety, that's my decision, and it can't ever not be my decision.

I've got these divided into green, yellow, red level here. Obviously, green is safer. Red is less safe to just proceed without some more support. This is like AI drafting or supporting. AI is summarizing my notes. It's organizing my data. It's surfacing themes. It's drafting measurable goal language to review. I don't know if any of the rest of you just pull your hair out of a treatment plan sometimes, but just having the like, "Here are the goals that I have for the client has that we've worked on together. Can you make sure that we're putting these in a measurable way?"

Then I review them with my patient, and we make sure that we're in alignment with them. We're never signing off on anything without reviewing them. Just knowing I don't have to go back through and flip through my Wiley handbook to find measurable language. AI is actually pretty good at that stuff.

Psychoeducation scaffolding. I hate worksheets. Just like as a human, I hate worksheets. I have some since I was in school. That means that I'm not great about remembering to give them, but some of my clients really benefit from worksheets. There are times when I'll ask an AI tool. There's an AI assistant within Upheal. There are other tools that are similar.

I can just ask, "Hey, can you create me a homework worksheet for Bob based on our last couple of sessions?" It will pull up a CPT-informed homework worksheet that is genuinely really good. Then I can send it to my client. I don't have to get my hands dirty messing with worksheets because, again, I don't like worksheets, but I respect the fact that some clients do.

Yellow tasks, assist and verify, might be things like trajectory or missed language analysis. If I'm asking, like, "Does it seem like this client is improving over time or not improving over time?" My conceptualization of the case is I think this might be OCD and not anxiety. That's something that I can pitch. I can go back and forth with AI on it. I'm not going to go along with an AI conceptualization that disagrees with my clinical conceptualization without a lot more research, a lot more consulting with other people.

Any output needs to be checked against a source. Any AI tool with itself will be able to say, "Here's where I got this idea." As an AI tool, I'm thinking this is anxiety and not OCD based on this particular article. I go look at the article, and then I can agree or disagree based on that primary source. Red, clinician only. These are things that AI does not need to be having a hand in.

Actual diagnosis, I think AI can probably surface differential diagnoses, but AI can never diagnose. You have to have a license to diagnose. Risk determination, I'm not going to trust AI saying, yes, this client is or isn't safe. Crisis response in particular, I think AI is pretty bad at, as we know from stories in the media, from people's risk increasing after conversations with AI.

In my personal experience, there have been times when I've investigated what does an AI tool say in response to these particular kinds of client situations, and I've been pretty concerned, honestly, because there are things like AI is always going to err on the side of what looks the most perfectly accurate on a transcript. That means that an AI tool is going to immediately launch into administering the Columbia Suicide Risk Assessment if a client mentions something that indicates hopelessness or passive suicidal ideation or endorse suicidal ideation.

That is on paper, the correct response. However, we all know that there are situations where a client might be extremely anti-authority and might immediately start going sideways or try to leave a session if you start assessing, or a client might have had really harmful experiences in a psychiatric institution before and be really opposed to the idea of possibly going back.

That client, you might need to go through the back door. You might need to be very gentle with the way that you assess. I'm very familiar with the literature that says, "You do not increase someone's suicide risk by assessing for suicide." However, if your client has other things going on that are extenuating, that are creating more complexity, you, as the person with the license and the relationship with the client, can bend those general guidelines because you know your client, because you know that the effects of your words are going to be different on that particular person than they would be on a different person.

AI can't do that, and it won't do it because they want the most perfect-looking transcript they can have for their own liability reduction reasons. That's fine and justifiable, but it's also a very strong reason never to rely on AI for any kind of crisis decision-making.

Treatment decisions, therapeutic relationship, all these things, these are human things. We know the research into interpersonal neurobiology. We know that our bodies physically affect each other. We can help clients calm down. We can help clients make progress, in part, because of the way that our voices sound and the way that our bodies present and the way that we connect with people. That's not something that's going to be replaceable.

AI, conceptually, is about providing the next expected response based on an enormous amount of analyzed data. Sometimes the process of therapy or assessment is an unexpected process. Sometimes you're taking a very complex system, and you're making meaning out of it with the person you're working with. That's not something that can be replicated by glorified predictive text, as valuable as AI tools can be.

Specifically speaking on conceptualization and planning, we've talked about a lot of this, but it is always worth, before you lean on an AI tool, start by originating your own ideas. Don't start by asking the AI tool, "What should I think about this?" Start by thinking, "Here's how I conceptualize the case," and then you can pressure-test it with a tool. Draft your own ideas. Don't say, "Create this treatment plan for this client."

If you haven't given it the information that is between you and the client that you've already agreed on, because treatment planning is a collaborative process between

you and the person you're working with, not between you, the person you're working with, and AI. These things obviously need to remain mutually agreed upon within the provider-client relationship.

Keep your own read. AI tools can augment or advance your thinking. They can also challenge your thinking, and there might be times when you temper what you're thinking based on the AI conclusion. However, if you go in with no idea at all and you're a blank slate, it's a lot easier to have your thinking be swayed by an AI tool.

Used well, if an AI tool does challenge you, it could prompt you to return to the original source notes, consider whether you need to reevaluate something. Used poorly, you might just accept an AI recommendation without independent review. The example I have here is like an AI summary might flag for me, like, "Hey, just a heads-up, your client, Bob, has been seeming a little more hopeless over the past five sessions. You want to check in with him on that."

I might think through it and be, like, "Well, yes, I hear that, but he's also got a ton of protective factors, and a lot of it is very closely linked to a recent loss, or it's very closely linked to a job going sideways or whatever." It's not that the AI flagging that is good or bad on its own, right? It's morally neutral. I can think through it and decide, "Is the AI tool surfacing something that I genuinely haven't considered, or am I aware that Bob is more hopeless?"

We're already doing what we need to do, and also I've made a risk determination that I'm going to stick with. Either of those things is fine. AI is just surfacing something. I'm not going to make a blind decision based on any of those facts. Same tool, opposite outcomes in these cases. The difference is whether the clinician stays in the loop and continues to be part of that treatment process.

Some of the risks of AI, automation bias. AI is very confident. It's very fluent. Everything it says is going to sound confident. Usually, it's not going to be altogether wrong. Usually, it's going to be partially correct, which is the worst. We don't want to make decisions based on AI just because it sounds confident. Your messy thinking, your messy session, your messy conceptualization, that is the beginning of a process of creating a sophisticated case conceptualization.

If you're newer and you're a career clinician, you're probably not as good at this as you will be someday. However, you're probably also less prone to those biases from your own experience because you're thinking about all the aspects. You're not running your thinking through a familiar funnel based on your own experience.

A lot of research shows that newer clinicians are actually slightly better at some of these aspects of clinical care than more experienced clinicians, because automation bias shows up even for those of us who haven't touched AI, just because we've been doing this for so long. Being willing to sit in the mess for a little bit and piece together, like, "What do I actually think about this case?" Is super beneficial.

Again, always form your own impression first. Treat the AI's opinion as a second opinion. I think we're all aware of this at this point. AI can invent plausible clinical detail that was not in the session. AIs can fabricate things. If you're using any kind of

a transcription-based AI tool, transcripts also aren't perfect. Things like speaker identification aren't perfect.

There was a time when I've seen speaker ID get totally swapped. The client utterance says were attributed to the clinician and vice versa. That's an easy fix in most platforms, but if you're not paying attention, you're just like, "What do you mean this note says that the client spent 50 minutes comforting me over my breakup?" Obviously, that didn't happen. There's a lot of things like this that you got to keep an eye out for. They don't happen often necessarily on these platforms, but again, always read things, always sign off on them.

Algorithmic bias, we've talked about this a little bit, but just in general, because of the way the AI models are trained, if you've got a client who's not part of the dominant culture, validity is not necessarily going to be equally weighted for every single person who sits across from you in your office. Then, deskilling, we know that at a certain point, if our brains are relying on something that's not our own brains to perform the skills that we've been trained in performing, there's a risk that we become less good at them. That is not at all the end that I want to see with the use of these tools.

However, if you're careful, you can use AI to reduce some blind spots. Used deliberately, AI can recognize patterns very sophisticatedly. It can surface themes, shifts in language, goals that stalled. As a person with a pretty strong ADHD streak, there are times when it doesn't really occur to me. Again, I have a hard time with treatment plans, I do my best, but it doesn't occur to me. I'm like, "Oh, yes. My client did say six months ago that one of their goals was to work on getting over their ex, and we've barely touched that."

An AI tool could say, "Hey, just a heads-up. You've got this thing on your treatment plan that you haven't really been addressing. You've been doing a great job addressing the trauma, but you're not really addressing the client's desire to get over the breakup." That's a good reminder for me. I can go back to that client and say, "Hey, [chuckles] I'm realizing that you had this goal six months ago. Is that something you still want to work on?" That helps me maintain my therapeutic relationship with that client, helps me respect their time. That can be beneficial, things like--

A lot of my research in my graduate work was on the use of therapeutic metaphor as a guiding force for clinical change. If you have a metaphor that you hold on to in clinical change, does that help the client track their own progress, track their affinity for the metaphor, the growth within that metaphor? I know people who use things like tarot cards, things like images of seasons changing. There's a lot of different ways this goes.

If I tell an AI tool, "Hey, this is the metaphor we're using to track change for this particular client," it can help me keep an eye on that metaphor, again, because I'm not a person who's necessarily going to think about these things at all times. It can help me draw those connections and those conclusions and bring that metaphor that the clients identified as being important and impactful for them into the work more strongly.

Again, always as an adjunct consultation, not a substitute, not a source of clinical truth, but super beneficial for things like self-of-the-therapist reflection and thoughts on how the client's doing, how I'm doing, how our therapeutic relationship is, how do I know that we are maintaining focus on the work that is most beneficial to the client as evidenced by their decrease in symptoms, as evidenced by their increase in ability to experience Compassion, whatever it is that we're looking for. AI is good at surfacing that stuff.

Ethics and privacy law, I'm not going to camp out a ton on this mostly because this is so in flux. These laws are coming out constantly. It is very possible that in the last few weeks since I made these slides that things have changed. When AI touches care, always be transparent, always account for bias, always stay accountable, and don't treat consent as a shield.

Just because your client consented one time to you using AI in whatever form you're using it, whether as an ambient notetaker, whether as a clinical decision-making support, that consent needs to be an ongoing process. Maintaining the client's awareness, like, "This is still part of your care. Are you still okay with this? Are you still feeling comfortable with this?"

Then maintaining your own strong sense of your own accountability. If you're ever considering changing a conceptualization or a decision based on AI, cool. It's great to consider that. It's great to have an open mind. However, remember that you are still accountable for that decision.

Any tool that is touching PHI in any form, and I tend to be pretty conservative with a lot of this. My preference is that even pretty significantly de-identified information is not hitting a consumer-level tool like Claude or ChatGPT that doesn't have any BA that you can sign. If you're using a tool that's within an EHR, it's a designated notetaker, it should be HIPAA compliant. You can check and see that in general, those tools are HIPAA compliant.

Obviously, a tool being HIPAA compliant only means that the tool has the appropriate technological safeguards. There's always still a responsibility for the clinician to make sure that they're using that tool in a HIPAA-compliant way. Always make sure that you sign a business associate agreement. Make sure that any tools you're using are very transparent about what those safeguards look like on their end.

Any tool you're using should be able to tell you exactly where is your data being stored, how is it being handled, is it being used to train models, is that being done with the client's explicit consent, or do people have to explicitly opt out of it to keep it from being done? All those questions should be easy to answer by any tool you're considering using.

Vet before you adopt. The APA has a great companion checklist that I've got referenced at the end, but it goes through some guidelines for just how to make sure that what you're using is going to be up to snuff, that it's going to meet those guidelines. It's a great way to structure your conversations with whatever onboarding person you're working with from the tool you're evaluating.

Evaluate for privacy, security, evidence-based, compliance before any clinical use. Before any client data touches that, you want to make sure that it's actually going to be compliant.

From a legal standpoint, state laws are constantly evolving. I'm on the Texas Marriage and Family Therapy Board. We're working with our lobbyist on what it looks like to try to shape clinical guardrails around clinical AI use here in Texas. I'm sure that many of you are part of similar conversations in your own areas, in your own professions.

This is genuinely tricky stuff. There are things that are very hard to regulate. There are things that you could regulate, but that doesn't necessarily remove the potential for harm. 36-plus states introduced 70-plus AI bills in early 2026. This is everywhere right now. Before you're using any kind of tool or implementing AI into part of your clinical workflow, make sure you look up your own personal state's rules so that you're not running afoul of anything that you wouldn't want to run afoul of.

You can also contact any professional boards are great. I'm sure the APA has more resources than I'm even aware of. If you have a lawyer retained or available through your professional organization, your own malpractice carrier will more than likely have good guidance for you on specifically what you need to be looking for. All of those are great resources for touching base before you adopt a new tool.

Practical next steps. I'm running through this a little fast because I want to make sure-- we've got some really good questions already that came in before we started. I want to make sure we have a little more time than I would normally leave for questions because this is a complex topic. I recognize that it's controversial. I want to make sure that we have plenty of time.

The practical next step. Audit where AI already touches your decisions. If you are already using AI as a tool on any level, again, whether this is just generally like you're putting the vaguest possible things into ChatGPT, like, "What are the best practices for working with reassurance-seeking OCD?" That's great. That's not a very intensive use of AI. It's still worth auditing. It's still worth asking yourself, "Am I double-checking this against other sources? If I ask ChatGPT what its sources are, does it pull up reputable articles? When I click on those articles, do they go to the actual article, or is it a dead link that goes to some completely other thing?"

I think this has gotten a little bit better, and ChatGPT for Clinicians has gotten quite a bit better with giving you real articles. For a while, AI tools were notorious for giving you very real-sounding APA citations for articles. Then you enter that DOI into any search, and it's a completely different article. It was just fully fabricated, which is wild. Uses like that, these tools are improving. They're not there yet. It's always worth auditing, double-checking.

If you're doing something like using a tool like Upheal or any of the other scribes or AI-enabled EHRs that are out there, think through, like, "Where in my workflow does AI already show up? Am I confident that within my workflow, everywhere AI touches my workflow, I feel good about this? I feel like this is ethical. I feel like this is augmenting versus replacing. I feel like this is enhancing my clinical skill versus de-skilling me."

All of those questions are super important because the landscape is changing so quickly. I think it's more important to have the right questions versus having the right answers. It's very hard to have answers that are going to feel evergreen in a system as complex as this one that you're asking-- the kinds of questions that are being asked by things like the APA Companion Checklist, it's going to be much easier to feel confident that you can find the right answers over time, even as these systems change and the capabilities of these tools continues to grow or change or whatever happens with these tools that are still very much developing.

Make sure that you form your own impression first. It's tempting. As someone who uses Upheal myself, it's tempting sometimes to get out of session and immediately go and read through the note that Upheal generated for me. I don't do that because I think that it's a form of de-skilling not to think through my own session first. I conceptualize my case first. What do I think worked well? What changed about my clinical formulation based on this session? What interventions did I do?

Then, when I look at the AI-generated note, I've already got a pretty good understanding of does this feel right to me based on the session that I just experienced because it could be the best AI tool in the entire world. If it's only going based on a transcript, it's going to miss things that I didn't miss. Honestly, that's great. That's job security. Our brains are very, very good at understanding what's going on in the room, in therapy, in a way that AI is not fully going to be able to replicate.

Vet all your tools. Run them through the Companion Checklist. Confirm PHI safeguards. Confirm that you have a BAA. Even if a tool is HIPAA-compliant, if you don't have a business associate agreement with that tool, you are not covered. You have to have that agreement that outlines exactly how you're using the tool, what your responsibilities are, what theirs are, what happens in the event of a data breach. All of these things get laid out in that agreement. It protects you from anything that might happen as much as possible.

Make sure that your consent language in your actual informed consent documents discloses AI use, ideally as a separate document. You don't want that just folded in as a footnote to your overall informed consent at this point. It's possible that as this gets more normalized, that might end up being more clinically appropriate. Currently, I very strongly advise people, have it be a separate document and have a clear conversation with clients about what it means for them to consent or not consent to that.

Even as someone who at this point is invested in this process of trying to make clinical AI be something that's genuinely useful to most people, I have a pretty tentative conversation with new clients I have where I tell them, "This makes no difference in your care and my willingness to work with you and the quality of care I feel like I can provide."

I actually do think that using AI tools can make it easier for me to focus on what the client's actually saying in session versus thinking, "Oh, I should write that down because I might forget that they told me this exact detail about their mom's chicken's name or whatever."

I like that aspect of it, but what I'm going to tell a client is it makes no difference at all. I can take this or I can leave it. I can write my own note. "If you are not comfortable with this, you have an absolute right to decline it." It makes no difference. Versus what I've seen a couple of times is more of a, "Well, this makes my life so much easier that I don't really take on new clients who aren't comfortable with it."

I'm not going to tell you how to live your life, but we need to be careful about coercion and about those power dynamics, which are challenging, and I think our clients I mean I know. Our clients have an absolute right to decide whether or not a tool like this gets used in their care. I do strongly encourage making sure that those consent forms are really well-surfaced.

Check your state rules. Your state rules might actually have specific guidance for exactly this scenario, saying, "Here's how you make it transparent enough based on our particular state." Then consult when unsure. Your board, your council, your law practice insurance carrier, all of them are going to have, most likely, opinions about what this looks like.

We've got references here. I know you're all going to be getting information, getting the actual webinar mailed out in two weeks. You're also welcome to email me. I think we'll have my email address put in there in a sec. We will get to Q&A. I did want to say first, if anyone's interested in trying Upheal, we do have an exclusive promo with code APA2026. I'm always happy to have you on board. If you decide to join, shoot me a message, and I'll make sure you get the right carpet treatment as much as I have any control over that.

That is it for the main part of the presentation. We've got about 15 minutes for Q&A. I've got already some questions coming in the Q&A here. I've got some that were submitted beforehand that I'm going to run through first. We will just get through as much as possible. I am always happy to hear from people. I'm always happy to talk through other questions if you want to reach out afterwards.

First of all, can we consult AI when forming a diagnosis or building a treatment plan? Is that ethical? I think this is one of those order of operations matters. Can you consult AI when forming a diagnosis? I would say if you already have a conceptualization in your head and you're trying to think through differentials, sure, AI is going to have an opinion. It may or may not be the correct opinion. A lot of times what I'll do is try to think through for myself. I think this is what's going on. I will consult my handy DSM-5 and make sure that I'm grounding it in some reality.

Then I might ask an AI tool, "Here's my thought that I recognize this and this might be differentials." I've noticed that a lot of AI tools tend to over-privilege, especially certain diagnoses, and which diagnoses they privilege depends some on the tool. I'm thinking about an example where I was using a tool that was really over-diagnosing oppositional defiant disorder and didn't seem to understand even age criteria for oppositional defiant disorder. The diagnoses weren't even valid.

That's the kind of thing that you, as a person who knows these diagnoses, can look at and say, "Okay, well, you don't really understand what you're talking about here. My judgment is going to be more valuable." I know I keep going back to this anxiety versus OCD example. I see quite a few clients who fall on that challenging middle

ground where they have a lot of anxiety, but they also have some things that generally seem like they're more of an obsessive-compulsive framework.

That's the kind of thing where I do find it valuable to consult, like, "Can we pull back? How many times? Show me the exact examples. Give me the actual transcript sections where something that seems like anxiety in the moment, the client was self-reporting as anxiety, really mapped better onto an obsessive-compulsive structure." That can be beneficial that I'm actually reviewing that transcript. It's pulling it back to me engaging with the primary source versus me just getting an AI opinion, something like building a treatment plan.

Again, I think if you already have discussed with the client, I think you can consult. The client's telling me that their goal is-- well, with the OCD example, the client's telling me their goal is to feel less anxious. I'm understanding their anxiety as being more obsessive-compulsive. They're not interested in exposure response prevention interventions because they're framing it more as anxiety. Those are challenging clinical situations where your understanding of best practice for treatment might not line up with what the client is most willing to do in their treatment plan, and because you have to be collaborative there.

There could be benefit in an AI tool that can actually search through the research on these treatments and say, "This would be an intermediate. We're not going full-on ERP, but we're also not treating like this is straight anxiety. Here are some options you might consider." Then you can take it back to the client and discuss them, make sure that you're being collaborative with that kind of thing. Long answer to that question, but I would say in general, it depends. Start with your own brain, end with your own brain. Consult with AI in the middle if you need to or if it's beneficial.

How do we know an AI tool is HIPAA compliant? Where is our session data actually stored? Any out-of-the-box AI tool, your Claudes, ChatGPTs, Geminis are not. ChatGPT for Clinicians that you actually apply for through an MPI, you have to have an approval process, can be HIPAA compliant, but it does require a BAA first, and they hide it in the settings. If you're ever looking into that, I'm all for it. It's got a lot of access to a lot of primary sources, a lot of articles. There are great things about it. However, make sure that you actually sign that BAA if you're going to put any patient information at all. Again, it's buried in the settings. Make sure you find that.

Any AI tool that's embedded within an EHR is going to be HIPAA-compliant. In general, make sure they can answer the questions, session data, where it's stored. That's a great question. In general, places are going to be required to store it in the United States. If it's a tool that's international, it might have some data stored in the EU. You can always ask those kinds of questions. I would be very skeptical of a tool that can't readily answer them. These are things that people need to be asking when they're signing up for stuff.

Are clinicians at risk of eroding their own competency by offloading cognitive work to AI? Frankly, I'd say yes. I think that there's a real tension between taking advantage of tools like this and protecting our own ability to think through things in a complex way. I am personally part of multiple consultation groups. I get a huge amount of benefit from these consultation groups and just hearing from actual human clinicians.

I also consult with AI at times. They don't really feel like the same thing. They don't feel like they necessarily meet the same need. Again, I think starting the process with your own conceptualization, making sure you understand what you're doing before you have that conversation with AI, those kinds of things all significantly decrease the risk. I think where it can actually increase your competency is by using it more as a pressure test, like "iron sharpens iron" kind of thing.

A lot of AI tools now, you can actually give them system messages like, "I don't want you to disagree with me on everything. When I propose a potential conceptualization, push back. Make me defend it. Make me explain to you, 'This is how I know that this is what's going on based on this dynamic, based on this diagnosis, based on this medication interaction,'" whatever it is that your particular context asked for. You can use those tools to sharpen you. They often won't just right out of the box. If you're using a purpose-built tool, you have a much better chance of it being effective in that way.

What are the most beneficial uses and the biggest risks of integrating AI into clinical practice? I feel like we've hit on a lot of this already. Most beneficial uses, I would say, especially for those of us who do struggle with getting notes done in a timely manner, getting notes done in a way that's comprehensive enough to support something like a medical necessity determination by insurance.

None of those are things that we probably went to grad school for. I almost quit when I figured out how much paperwork being a therapist is going to be, honestly, because I'm not a paperwork person. Be that as it may, I think there are a lot of benefits in using AI tools for these purposes. I think about the benefits of AI as a function of waste level.

I think that there are definitely uses of AI that can be wasteful. I think that it's also incredibly wasteful for us to waste human talent and human energy in educating someone, putting them in school for years and years, putting them in front of a client, realizing they're an excellent clinician, and then having them burn out in three years because they're not paid enough and they're overworked and they keep going home and writing notes instead of going to their kids' soccer game or sitting at dinner with their spouse. These are all hugely wasteful things.

If we can use AI to offset some of that waste, I'm all for it. I think that's very beneficial. I think the biggest risks of integrating AI into clinical practice probably are, in my mind, I think the biggest risk that's difficult to mitigate is probably the bias aspect, the fact that, in general, things are going to be biased more toward dominant culture. This is true even of the tools that are pulling from literature because the literature itself is often biased toward dominant culture, just as a volume play.

There's so much more research into the dominant culture. If you're using AI tools, directing it directly toward these journals, these sources are more likely to have the kind of information that I'm looking for. That nuance, if you're working with a tool that allows it, is beneficial, but I do consider that a pretty significant risk.

Should we document when AI helped generate a note or report? How do we distinguish AI-generated from clinician-generated content? I think documenting that kind of thing is beneficial. I think if we're talking with – Upheal, my platform, has the

ability for me to send a client summary to my clients after a session. I'm always very clear with my clients, like, "I don't have the time to write up a separate summary for you. I'm not personally writing this."

If it's a client who's consented to the AI tools, then we'll have a conversation about, "If this is something that you want, I'm more than happy to send it to you. I will always review it, but it's AI-generated. This is something that an AI is pulling from our session. If you ever have questions about it, concerns about it, let's get into it." I think it's beneficial, but I do think it's important that they always understand these are not my original thoughts. They're drawn from my thoughts in the session, but it's AI-generated.

If you're writing something like a psychological report, those tend to be somewhat more formulaic. I know the last psychological report I was reviewing definitely had some copy-paste inaccuracies, which gave me confidence that the psychologist was not using AI but did not give me confidence that the psychologist was looking through their copy-paste blocks very carefully. I think we've always been vulnerable to the-- we allure of a shortcut in doing work that can be very tedious.

I think as AI use gets more normalized, I can see it being increasingly used in things like psychological report generation. It's not going to hurt anything to have something at the bottom that just says, "X AI tool helped me generate this," or leave a watermark on it, or having your informed consent, "I use this particular tool to generate these reports. Here are the benefits. You can always opt out," or you can opt in. Having that kind of language, even as this becomes more normalized and more part of practice, is not going to hurt anything.

Awesome. We've got so many good questions. I use AI a lot, now I write or sound like an AI, even though I don't use AI a lot in my work, just the exposure to AI, what would you do? [chuckles] I'm answering this one partly because I ask myself the same thing. I second-guess myself sometimes. Do I sound more like AI now because I am exposed to this so much?

I have actually started doing a lot of writing that doesn't involve AI at all to counterbalance. I'm not recommending this. I'm not saying all of you need to go out and start a Substack, although I mean, sure, do it if you want to. I think there's a lot of benefit to just staying connected to your own voice, both as a clinician and as a human. I have clients who don't consent to the use of AI. I do write some of my notes manually, but I also do a lot of my own writing to just stay connected to, like, "This is the person I am. This is my voice. These are my opinions," separate from AI.

That would be my two cents. I think it's not like-- Most of us are not sounding fully like all the little AI-like things that we can all probably recognize at 500 feet now. It is tricky because I know that I've done the thing where I put one of my college papers. I went to college in 2010. I put one of my college papers into an AI detector, and it was like "30% AI generated." I was like, "No, it's definitely not." What are you going to do? You can't do all of it.

Some research supports objective appraisal methods as often more accurate than clinical judgment. Will AI be seen as more akin to objective appraisal? 9% expanded risk for clinicians when our clinical judgment conflicts with that data. If AI

recommends an action to ensure safety and we override, do we have more risk for malpractice? This is a super good question.

I'm personally thinking about this in terms of emerging ethical frameworks. There's an ethical framework that I've been thinking about myself based in speech act theory, that I don't think is at all. The only way to think about it, it posits that we have to be able to explain our clinical judgment in terms of three different levels of thinking. Both what I say, how it affects someone, and how I mean it. All three of those matter.

If I can't articulate my clinical judgment or my clinical intuition in terms of "Here's what I want to say. Here's how I intend it. Here's how I think the client will receive it." If I can't proactively articulate something, it's probably not an ethical risk-related override. If I'm just going based on vibes and I can't articulate it, I would have more concerns.

That's the conversation I have with my supervisees and newer clinicians I work with is don't assume that because something feels like the right clinical decision, it is the right clinical decision. If you can't pause and answer those questions before you make it, it's probably more suspect. Afterwards, when we're talking in supervision, if you can't explain to me, "Here's why I made the decision that bent the rules," we've got concerns.

Yes, I think there is risk to that. I think that, again, humans are very, very complex systems. I think that we need to have ethical decision-making practices and processes for crisis situations that are not so rigid that we can't understand the need to bend them in some cases. At the same time, the research exists for a reason, and objective appraisal is always going to have a place. I think when in doubt, when you're not sure if you're going to be able to articulate an explanation, it is not going to hurt anything. I shouldn't speak in absolutes. When in doubt going with the research, the objective appraisal answer is generally a pretty safe thing.

Let me see. I think we've got about three more minutes before we need to call it. We've got a question about whether I've noticed better accuracy in any of the platforms that clients might be engaging with, which is a great question, like ChatGPT, Claude, Gemini. Unfortunately, my answer is it really depends on what release they're currently on.

The most recent release of Claude was notoriously bad. I think they just came out with a new one, 5.5, Opus 5.5, that's supposedly better. I haven't really tested it yet, but it depends so much. Unfortunately, I wish I could say, "Yes, tell them to use Gemini. It's going to be great," but I actually can't because these things are constantly being tweaked. A later model does not mean a better model.

If you have clients who are using these tools, go through with them, what are your red flags that this tool might be beginning to hallucinate, fabricate, be sycophantic, go through those risk factors, and help the client recognize if these things are happening. Maybe it's time to switch tools or change system messages, that kind of thing. If you haven't seen my webinar from last month, that also talked more about client-facing tools, feel free to check that out, but it's a great question.

We've got a question about whether sessions are being videoed or audioed for use of AI. AI notes, to the best of my knowledge, every AI tool out there, what they'll do is they'll record a session, transcribe it, and then upload the transcription. In some cases, you can also save the audio, I believe. Then based on the transcript, a note is generated, or you can get something or do a summary.

I don't think anyone is doing the video stuff right now except for some slightly more out there things. I think there's someone who's trying to work on more AI therapy stuff, which is a whole different can of worms that I don't really want to get into. [chuckles] Not having video does decrease the accuracy. A lot of AI tools are going to rely on default language.

For example, a mental status exam is going to assume that appearance is appropriate unless you specifically say that appearance is not appropriate because it doesn't have access to video. The assumption is if I, as a clinician, my client comes in and they're obviously super disheveled and look like they haven't slept in five days, I'm probably going to say something about that verbally. Then my verbal statement on that will go into the appearance section of the MSE, not because the MSE can actually see my client in any way.

That's always going to be a tension. I'm not sure that there's a resource-efficient way to handle video in that way. We may well get there, but we're not really there yet. Okey-doke. I know that I didn't have a chance to get to everything, but please feel free to shoot me a message. I'm happy to stay in touch. Thank you all so much.

Host: Thank you so much for joining us, Leila. Thank you to all our listeners for your participation. We'd also like to thank Upheal for making this webinar possible. A recording of this presentation will be emailed to everyone in two weeks' time. The recording will include the presentation slides. As soon as the webinar has ended, a short survey will appear on your screen. Please take the survey and give us your feedback. We thank you for your attention and hope you have a great day.