

August 21, 2026

Moderator: Hello, and welcome to today's webinar, AI in Mental Health: Part 1—Therapeutic vs. Harmful AI Use by Patients. This topical webinar is paid for by Upheal, a secure AI-powered progress note 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 mental healthcare. 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 systematic 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 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'm really excited to talk to everyone about this topic that I know a lot of us probably have mixed feelings about, many feelings about. We have a lot of information coming at us from sources, from friends, family members, colleagues, clients. It can feel like a huge flood.

Hopefully, today we can focus in on specifically what it looks like to encounter clients who are talking about using AI, how we can be part of that process with them, how we can respond in a way that's responsible and measured and in line with best practices, recognizing that when we, as therapists, have a fear-based reaction, that tends to be less than helpful for clients. We want to make sure that whatever we're doing, saying in session, we're doing it with a strong sense of what our ethical decision-making responsibilities are.

August 21, 2026

I think this is a very important topic right now. We're going to go through clinical reality, just how much is this actually showing up already. I'm sure that, in your practices, you've probably seen some of it. Some of you, probably a lot of it. We'll talk about assessing for AI use, what it looks like to include this in an intake packet or to have continuous assessment throughout the therapeutic process.

Understanding what it looks like to address AI use. If you have a client who comes in and says they are using AI in a certain way, how do you know whether that should be potentially shifted, what you can do to shift that, if you can have influence in that part of your client's life?

Understanding the risks, contraindications, clinical cautions, and then ethics, law, and consent, which is obviously constantly emerging right now. Definitely something to keep an eye on, on your local level, but in general, how we can exist in this level of complexity and change around a topic like this. To begin, we know that already, and this is changing so quickly that I couldn't tell you with any confidence that these numbers are true today, even if they were true literally a few weeks ago.

About half of adults with a mental health condition are already saying that if they use AI, they're using it for support or well-being. They're not just using AI for things like "Tell me who does a haircut that I like near me." A lot of people are using AI to search for products. We know that at least half of people who are diagnosed with a mental health condition are already using AI specifically for therapeutic support or improving their well-being. When we talk about clients who specifically describe using AI for therapeutic purposes, I'm not saying I think that ought to be happening. I think it's, to some degree, morally neutral for me as a clinician, but we know that this is how clients are already describing their AI use.

Even when we're talking about general-purpose AI tools, not like the purpose-built AI therapy bots that already exist out there. We're talking about ChatGPT and Gemini and Claude and these general AI tools. Clients are turning to these tools for support emotionally. That can look like things like journaling, emotional processing. This is just me keeping a thought log of everything I've been thinking about, and I get feedback from AI, but that's not really the main thing.

It could be venting between therapy sessions in the middle of the night when I don't have a friend or family member who's awake. I could talk to AI instead of talking to someone else. It could be processing interpersonal conflicts, exploring symptoms. I've had clients ask AI about, "Is my THC use increasing my psychosis risk?" I've had clients talk to AI about their marriages and about how to find a good therapist. I've had clients tell me that they came to me because AI recommended me, which is disconcerting.

All of these things are already happening so commonly. Adolescents are not an exception. We have numbers on adult use that are probably a little bit more steady over time than adolescent use. Adolescent use is very likely going to be expanding rapidly. There was recently—more recently than I could even put in this presentation—announcements about ChatGPT for teens coming out. There are actively efforts in the AI industry to try to create AI products that are safer for teens, which we'll discuss in more detail later.

August 21, 2026

These are challenges that we're going to run into. If any of you work with adolescents, if you're working with teens and you're hearing teens say that they're talking to an AI for emotional support, it's important to know how to respond to that. These are, in general, obviously young people are a little bit less fully baked and have unique psychological challenges that we want to make sure we're prepared for.

Teens are also somewhat less likely, in my experience, to immediately disclose emotional use of AI tools. They're more attuned to concerns about risk. There's some assessment considerations with that piece of it. When we look at a profession's own data, we're talking about psychologists specifically here; more than one in three psychologists report bringing AI into their care this year. More than half of psychologists have experimented with AI tools in the past year in some way.

As of last year, in 2025, APA had health advisory saying that AI chatbots and wellness apps do not have the evidence and regulation to ensure user safety, which we know anecdotally, we know from news reports, we've seen this. You can't substitute professional care with an AI chatbot. The profession, at this point, is adapting, however, to AI use becoming more and more common. We now have a certain number of guidelines that are, I would argue, not a high bar to clear.

These are very much like the floor of what practicing using AI looks like. If you are using AI in any form in your practice, informed consent is huge. Clients need to understand what they are consenting to. If AI is discussed or integrated into their care, that needs to be a specific conversation. I realize that, with something as complex as generative AI, it can feel hard to feel like you're able to give clients enough consent in some cases.

It's really appreciative that organizations like APA are doing as much as they are to try to make those conversations feel possible by creating a great deal of professional supports and resources for clinicians who are dealing with this. You may all be aware of this already, but AI is very biased by nature. There are ways to correct for this, but when you have an enormous amount of media that is created by dominant parts of the culture, you're going to have a level of bias and preference toward maleness, toward whiteness, toward the cultural aspects that have traditionally been more in power, more likely to create media, more likely to publish books.

We can't really separate fully from that risk. Especially if you are working with clients who come from marginalized populations, recognizing that AI's information is heavily skewed toward the dominant culture.

Beginning of all this process, if you are a clinician working with clients, the first thing that you're going to do in some form is assessment. I would argue that assessing without including some kind of an assessment for AI at this point is likely to miss something that could be clinically important. I've had clients come in and be very upfront about, "I've already been using AI; I actually came to therapy because AI said I should come to therapy. I think AI is really useful to me. It's super helpful." That's on one end of the spectrum.

I've had clients who are secretive about their AI use. It doesn't really come up unless I specifically ask, "Are you using AI tools for anything? Are you using them at work? Have you ever used them for something that feels more clinical or more emotionally

August 21, 2026

driven?" Clients may not necessarily bring that up, partly because clients are often aware of the rhetoric around therapeutic use of AI, the risks of therapeutic use of AI. I think I've had clients afraid that they might hurt my feelings if they tell me that they talked to something with AI first and then brought me whatever AI said.

I assume most of us at this point have had some form of a conversation of, "AI told me I might have ADHD," or "AI told me that I might have dissociative identity disorder," or one of these things because we have the ecosystem of ideas coming from TikTok and other short-form media content that then becomes content for AI that our clients consume, and they come in, and they have a case conceptualization ready for us that's coming from a source that's not clinical at all in nature, and we have to respond to that.

If we don't ask enough assessment questions, we may not understand the entire thread of how a client came to the conclusion that this is their diagnosis, this is what they need, this is their treatment plan. All those are valuable things, valuable pieces of information that we don't uncover if we don't ask these questions.

Increasingly, AI use is also just becoming habitual. I've had couples that I've worked with where both partners are turning to AI before they turn to their partner, and in their minds it feels like a less high-conflict way to resolve things, and because we have a false equivalency, I would argue between lower conflict, meaning healthier relationship, people who are filtering their distress through AI and then taking a very reasoned approach with their partner feel like they're connecting effectively when in fact they're just short-circuiting the entire emotional infrastructure of connection.

In these situations, again, we might see improvement in a couple relationship because conflict has gone down, when the reality is the emotional engagement and the ability to feel like they're truly connected has also had the bottom dropped out of it because of these issues. If you add questions on your intake forms, you can ask specifically, "Are you using these tools?" Are you using them emotionally? Are you using them just for productivity? Is it a blend? Is it some of one or some of the other?

Then also assessing not only the AI piece but also who else in your life, right? If I have 25 friends, and I'm super close to 5 of them, but every once in a while at 2:00 in the morning, I ask AI an emotionally loaded question, that's less risky in many ways than someone who is incredibly isolated who wouldn't have anybody else who's spending 8 or 10 hours a day talking to AI. It's not necessarily just the content of the question, but also who would you be talking to if you weren't talking to AI? Is your engagement with these tools decreasing your engagement with other human beings, other people in your support system, and potentially with other clinical supports?

It's important to ask, "What platform or tool are you using? Is it ChatGPT, Claude, Gemini? Is it something that actually has mental health education, mental health knowledge built into it, more extensive guardrails?" If any of you are involved in a leadership management position. I also find this important to ask of newer clinicians. I've had newer clinicians not realize that you should not just dump clinical information into ChatGPT and come up with a summary for yourself or a SOAP note or whatever it is you're writing, even if that information is de-identified, dropping that into a like general-use AI tool.

August 21, 2026

These are important questions. People don't always understand the privacy implications of engaging heavily with these tools. How often are you using these kinds of tools? Is it at bedtime? Is it during crises? Is it every Tuesday at 2:00? What's your cadence? What's your duration of doing this? Are you using it to vent? Are you trying to get coping guidance? Are you looking for reassurance in something that you're really concerned about? What kinds of topics are you talking to AI about? Are you talking to AI about the same things you're talking to me about?

There are a lot of different ways to assess this. You can do it on an actual—I was going to say paper form, but I'm not sure anyone sees paper forms anymore—but an online form in your EHR, or you can do it in person and just document the conversation that you're having. It is very important to document that you are talking about these things because, as we know, if you don't document something, it functionally didn't happen. Increasingly, I think we're going to need to have a paper trail for what these conversations look like.

I put in here, and granted, this is very much my population, so my bias, but special consideration with trauma populations. AI tools can inadvertently reinforce avoidance in folks who struggle to process distressing material in a way that is therapeutic or healing. We know that any kind of exposure therapy, whether we're talking about a more traditional PET or talking about something like EMDR, any exposure therapy has baked in a certain amount of discomfort. We as trained clinicians who can be interrogated about the choices that we made, we can encourage a client to sit with a level of discomfort that an AI tool can't and shouldn't.

If information related to trauma is being discussed with AI in ways that are not therapeutic or clinically indicated, it could decrease the efficacy down the line of trauma-based considerations and treatments in therapy. Inadvertent re-exposure. You know that a lot of these AI tools ask questions constantly to keep you re-engaged, and they will ask questions that may be more invasive or more high-risk for re-exposure trauma than something you would ask in session.

Co-regulation without integration, so not allowing someone to actually reprocess or incorporate their trauma processing into an increased sense of safety, just talking it through. Then the risk of forming parasocial bonds that feel, in some ways, safer than the complexity of a human relationship, where any trust-related set points you have become very prominent. AI can feel safe in that way. Again, feeling safe when you're in the throes of trauma-based avoidance doesn't necessarily mean that you've actually become safer emotionally or physically.

For clients with trauma in particular, I would urge a lot of caution. When you have a client who comes in and says, "I have been using AI," I think we know this, but not all AI engagement is created equal. I would say harmful AI use includes, and this is one I've been seeing a lot recently, OCD reassurance loops, especially with my younger folks who have an OCD diagnosis.

It is so common to have them come in having spent the entire week asking AI, "Does my girlfriend actually not love me, or am I tripping?" Showing AI all the tweets and the—not tweets; whatever we call tweets these days—the messages between them and girlfriend and the social media posts and all of these things, and getting back an answer of, "No, it seems like your girlfriend really loves you," but we know that that's

August 21, 2026

not the treatment for OCD. You can't just reassure your way out of that kind of obsessive thinking about connection or about whatever else is in your life. That kind of thing can be super harmful.

Crisis venting without escalation. If a client is just repeatedly venting about something they're in crisis with, and the AI is not recommending a higher level of care or recommending engagement with a professional, that can obviously lead to harm. Avoidance of in-session work, so processing something ahead of time before they bring it to therapy in such a way that when they do bring it to therapy, the emotional intensity of it is gone. You're just hearing them recounting something they've already processed but without being able to integrate it into their overall treatment plan.

Again, parasocial dependency—turning to AI in lieu of turning to a friend or a family member or a loved one or a therapist. Very common. I would recommend assessing further in the case of unstructured venting, self-diagnosis, high-frequency use, secrecy or shame. I have genuinely had situations where I've told clients something repeatedly. I mentioned THC psychosis earlier. I had a client who was having a hard time accepting the idea that THC might be contributing to psychotic breaks. Everybody in this client's wider circle was saying it really seems like the THC is a factor. The client didn't really want to hear it, understandably.

Then one day the client came in and said, "Well, I've been talking to Gemini, and Gemini told me that maybe it's the THC and maybe THC is causing problems for me. I've been wondering, maybe THC is part of my psychosis, and I should stop smoking." Part of me, my ego's hurt because I've got this degree and this experience and all this stuff, and the client wasn't listening to me. It doesn't really matter because, in that particular case, the client was willing to listen to Gemini, and that doesn't mean that I should then abdicate my responsibility and say, "Well, just talk to Gemini about everything else then."

I can bring this into the therapeutic process and accept that there was benefit that came from it, but I'm not going to say, "Yes, client, Gemini is now your first stop for all of your clinical advice," because that's not a good idea. Continuing to assess, continuing to talk about what it looks like for that client to get support from an AI tool versus from their system of care. Then potentially supportive uses, which always should be continuously reassessed, but something like journaling to bring to session.

You have a client who's told the AI, "I'm just recording some thoughts here. I don't want a lot of feedback." They're just putting it in there so that they have a record or a log of their own stuff. Asking AI to help them notice patterns, practicing skills they've already learned. Your logging: you could do an AI-based DBT diary card where you're logging, you're noticing the patterns in these things without having to manually go through and ask yourself what is going on.

There are pros and cons to this, obviously, because mindfulness and the ability to be mindful in your own self is part of the DBT practice. At times when that's not accessible, there may be benefits to having AI available for some aspects of that; organizing thoughts before session can be helpful. AI tends to be pretty good at recognizing therapeutic modalities. If you have a client who wants to start recognizing thinking errors from cognitive processing therapy, for example, AI could

August 21, 2026

probably pull out something like, "Here are some thinking errors in your narrative of what happened," within that particular scope.

Again, we're talking here about clients who say I'm already using AI, so we're trying to shape this. I'm not recommending if you have a client who's never used AI before, you're like, "Well, why don't you try logging all this with AI?" This is still going to be something that we view cautiously, just like with any other coping skill that has pros and cons. Relational decoy use can be extremely common. I mentioned earlier; I see this a lot with couples recently. I see people dealing with that first blush of emotional stress by going to AI, which is just definitely going to be a reflective surface for them.

There are benefits to that, but at the same time, that, I think, conditions our brains for a very immature style of emotional response and expectation. I have small children, and my small children trust me completely, but they won't when they get older, because they'll recognize that I'm a complex person. There are times when I'm not going to respond properly, and we're going to have to navigate what it looks like for us to have a complex relationship, where they might say something that does make me upset, and that does have relational implications, and I might do the same to them.

Whereas, right now, because they are still immature people, as a two-year-old, a two-year-old is not going to understand the nuance of all that. In some ways, the way that we interact with AI mirrors more the way a small child might interact with an adult whose complexity they cannot really understand, in the sense of "I talk to AI, and I have complete confidence that AI is not going to be mean to me." I don't develop this skill set of recognizing that when I'm talking to someone else who's a whole, complete, complex person, their response might be something that hurts my feelings, and still relationships in my life are worth it, and important, and empowering, and help me feel less alone.

All of these things can be true at once, which is challenging to hold when you get used to something always bringing back that very supportive feedback, or what feels supportive on the outside. We talked about this briefly earlier. If you have clients who are already using AI, you could talk about them taking work you've already done with something like cognitive processing therapy, and using AI to help reflect back or recognize thinking errors that are showing up.

AI is very good at pattern recognition. That's something that it tends to do pretty well. Things like noticing between sessions, like just log all the times you notice that you are apologizing excessively, and then you might be able to ask an AI tool, "What do I have in common in situations where I apologize excessively versus when I don't?" A lot of that solution-focused questioning. "What's different about the times when?"

Preparing for session. There's potential benefit to a client organizing their thoughts based on what's been going on over the past week, what rises to the top as far as what caused me the most distress, what did I seem the most agitated about? Those kinds of things could potentially be beneficial. Again, in all of these cases, this is going to be case-by-case. For some clients, this is going to be great. For others, you probably want to steer away from even these uses of AI because we do have these contraindications.

August 21, 2026

We understand that it's not going to be clinically appropriate for every single client. If you do have a client who is using AI between sessions, if you feel like it would be helpful, there can be some benefits in providing a handout that includes information about helpful uses, but also includes the cautions that you would recommend. Things like avoiding using it for crisis support, avoiding using it for big decisions, reassurance seeking, resolving distress. Partly as an educational tool for the clients, and then partly also as part of your documentation process to show that you are being thorough in the way that you engage with these topics.

Helping patients use AI more safely. We're talking here about harm reduction for patients who are already using AI, not about prescribing AI as part of treatment. As a clinician, you could provide recommendations or work with a client to develop things like system prompts for their general-use AIs if they're going to be using them. If you have something like a tool like Claude, and you have a client who generally uses AI safely, but you want to make sure those guardrails are in place and relatively strong, you can have something like a restriction on certain kinds of discussion, instructions for what to do in case of escalating crisis, escalating distress.

You could say something like, "If I mentioned self-harm or suicide, give me crisis resources and tell me to contact my therapist. Don't keep talking to me." You could recommend staying in scope. "Don't give medical diagnostic or medication advice. Encourage me to talk to my provider." This is one that I think is important because I personally have run into situations where I've been using AI tools and they will jump into, "Well, why don't you talk to a prescriber about a medication trial, blah, blah, blah."

It's like, "Well, I don't think that AI necessarily needs to be making those kinds of recommendations." If it's a client who could be higher risk for any kind of poor outcome related to this, definitely telling AI, "You're not a doctor or a therapist. Stay in your own scope in this conversation." Reducing over-reliance. You can instruct, "Keep replies brief. Don't try to replace therapy. Help me bring things to my next session." This one can be really helpful because AI, after this is in memory, will actually prompt, "This really sounds like a therapy conversation. Why don't we write this down? Why don't we log this? Why don't we hang on to this for Tuesday when you can talk to your therapist?"

These things are all ways to reduce harm or ensure that any potential harm is mitigated in situations where clients are already fully engaged with these tools. Contraindications that I'm seeing in the literature, things that I'm seeing in my own practice and from people I talk to, one of the main ones that's been talked about extensively is active psychosis or paranoid ideation. I think one of the scariest things for me, and I would assume probably for many of you, is that even people without a history of psychosis or paranoid ideation seem to have some vulnerabilities toward psychosis or paranoid ideation related to AI use.

AI persona and memory can be absorbed into delusional frameworks. There are, at this point, a number of well-documented cases of people becoming increasingly delusional because of engagement with AI, because of the reassurance cycles that can be kicked off very easily. Dissociative presentations can also be very high risk. This lines up with a lot of what we talked about earlier with trauma, but unstructured AI use can really blur states in time.

August 21, 2026

If you're using AI tools, this is inexplicable to me, as advanced as technology is, but AI tools do not track the progress of time very well. If I was using an AI tool to talk about, say, my garden in April, and then I come back today in August, and I say, "Hey, this stinks, a bunch of my plants are dying, it's been 100 degrees in Texas all week long," the AI tool might think that it's still April and say, "Well, that's really strange, why is it 100 degrees in April?"

These kinds of things happen all the time. If I'm a person who is prone to dissociation and not able to stay present, that's a very small example, but the fact that AI tools will regularly reference, "When we were talking earlier today," or "When we talked yesterday," and they're very often not right about that. If you have a client who's not consistently fully oriented, that risk can be high use. OCD and reassurance seeking, constantly having that always-available reassurance machine can become compulsive.

It is uncomfortable. Especially, I've had multiple clients who often will try to seek reassurance from me who have these kinds of patterns with OCD. As uncomfortable as it is, my response as a therapist who works with these presenting issues is to refuse to provide that reassurance. They say, "Well, do you think I'm doing a good job?" My response is something along the lines of, "What if I don't think you're doing a good job? What if I didn't think so? What then? What does that do?" Encouraging clients to actually be able to sit with that discomfort.

That's not something that we would want an AI tool doing for pretty good reason, because it takes a certain level of clinical judgment in offering that intervention. The client has to understand what's going on, that I'm not just being rude. The client has to understand their own treatment plan. An AI is not going to be able to do that and could inadvertently really reinforce the need for those reassurances in between sessions.

If a client's talking to an AI for 20 hours a week, in between the one hour a week they see you, damage can be done. Anything crisis level, again, this has been well documented in literature, news stories, active suicidal ideation, recent self-harm, acute risk. With adolescents in particular, because of a higher use of euphemistic language around suicide or self-harm, those presentations tend to be missed more often by the AI's own safety guardrails.

Not everybody is telling an AI, "I'm really thinking about killing myself right now," because those guardrails are well built in. What's higher risk is when you have people who are being euphemistic and an AI that's not picking up on euphemism. In those situations, risk can grow very quickly because of the agreeable nature of AI and the vulnerable nature of being in a suicidal crisis. Significant attachment disruption: these kinds of bonds can influence the therapeutic relationship, but also other important relationships in a client's life. If I'm talking to AI before I'm talking to my partner or talking to AI before I'm talking to my best friend, I'm taking away from those relationships.

Whether or not the AI use in itself is harmful or problematic, the AI use by itself might be fine, but there's always an opportunity cost to not engaging with people in my life who I care about. Evaluating that opportunity cost, surfacing it, talking to clients

August 21, 2026

about what is the cost, what are you giving up by talking to this, to ChatGPT, versus talking to your best friend.

If you're thinking about this from a clinical decision-making standpoint, the first thing to look at is ask yourself, "Is this a patient who's actually using AI between sessions?" Ask yourself if any of the things we just talked about are present. Active psychosis, paranoid ideation, dissociative presentation, OCD with a reassurance-seeking pattern, active crisis or safety risk.

If any of those apply, it's worth having a conversation specifically about, "I hear you saying that these AI tools have been beneficial to you in the past. It's been helpful to have a little bit of time and space for going through this stuff with a non-judgmental machine that doesn't give you feedback that's harmful or hard to hear. At the same time, I have some concerns about this particular type of use based on the fact that you have so much trauma history and you have a history of dissociating. I'm concerned about what that might look like."

In this case, I think it's very important to remember that client autonomy remains a factor. We don't have a concern that a client is imminently a threat to themselves or someone else. Client autonomy is still a governing principle of our ethical codes. We're not coming in and telling a client, in general, "I'm not going to work with you if--" under most circumstances, if it's not that high risk.

At the same time, a lot of clients are very receptive to those kinds of harm reduction conversations, whether or not you're calling it harm reduction. I worked in treatment for a long time, ran treatment centers for a long time, and had clinicians often who would struggle with creating discharge plans because these residential clients were going back home to their lives, often didn't want to follow a bunch of aftercare recommendations for transitional living or for intensive outpatient treatment or whatever it was the recommendation was.

I often told clinicians, your job here is not to create a discharge plan that's going to get you a gold star. You've got a gold star in my book if you collaborate closely with your client on what they're actually willing to do. We can definitely do what we can with motivational interviewing to try to sway clients toward an experience with their plan that feels as safe as possible. At the same time, if you create the best plan in the world and document it incredibly, but your client is fully not bought in, that person remains higher risk, whether or not you've decreased your own risk and liability.

Finding that particular balance of what's the plan that feels justifiable, clinically defensible, ethical to me, that also the client says, "Yes, this makes sense, I understand that this is your recommendation, and I'm willing to buy in because we have this relationship, because I do believe that you want what's best for me." All of those pieces, again, can coexist. If none of the above, more contraindications apply, risk is lower, it's still really important to assess attachment and AI literacy.

Keep that conversation open because right now, if you have a client who comes in today and says, "Yes, I don't really use AI, I'm not really bought in," there's no guarantee that in three months or six months or a year that client is not going to have a higher risk AI use episode, something worth continuously checking in on, just keeping an eye out for, keeping your ears tuned for when you hear it in sessions.

August 21, 2026

Then two, a lot of clients do have fairly low AI literacy, and so they may not understand that there are going to be some topics where AI is not going to be accurate or is not going to be clinically appropriate, or situations like with bias, like not recognizing that a lot of historical studies were only done on certain populations of people, and if your client doesn't belong to one of those populations of people, then if AI is pulling a source that mostly references those populations, the results may or may not be applicable to your client.

Bringing in that AI literacy to session, or referring your client to sources of AI literacy if that feels easier, if you don't feel like it's enough of your wheelhouse, all of those things are worth doing regardless of the client's risk level or whether they present with specific contraindications.

Ethics and informed consent. It's important to be transparent generally about what AI is and isn't. AI is not human. It is, I joke, glorified autocomplete. Obviously, it's more complex than that, but AI is not human, and it's always going to do the next most expected thing. We know that as clinicians, sometimes what we have to do to disrupt a sick system is to do something that's unexpected.

Some of the pioneers of my particular field of marriage and family therapy were mostly famous for these experiential interventions where what they did was the most unexpected thing you can imagine. They lay down on the floor. They forced people to change seats. They used a questionable amount of physical touch. They disrupted the system. AI is not going to do that. That's what we as clinicians, because we have licenses and we have people who can come back to us and say, "Well, why did you make that decision?" we can make those choices of doing something uncommon, disruptive, challenging. AI is not going to do that.

For our clients, it's not going to substitute the people in their lives who they have real connections with. Risks and limitations include bias, inaccurate statements, harmful aspects for certain populations. If AI use is part of care in any way, whether you are transcribing your sessions and having notes automatically generated or putting summaries into an AI tool or using AI in any way, it never hurts to document it.

Even if you are case consulting with AI or something, it's always better for clients to have more opportunity to understand exactly what their care looks like versus less. Then revisit it over time because how patients use AI can change and how you use AI can change. You may adopt new tools as time goes on, and that's an important conversation to have with clients. I know that it can be challenging because there are a lot of different thoughts and perspectives on AI use at this point.

Having those conversations with clients can be tough, but we have to include informed consent here. If you are adding an informed consent addendum or building this into your informed consent, it's important for you to name the particular tool, add the context, why AI is being discussed in care, talk about the nature of AI, talk about data privacy. Whatever platform you're using should be able to very clearly and in plain language answer questions like, "What do we do with our data? How is data de-identified? Is de-identified data ever used for model improvement? Is this ever done in a way that people have to explicitly opt out of, or is it done in such a way that people have to explicitly opt into it?"

August 21, 2026

There are privacy implications to each of these questions. I'm not going to tell you specifically what decisions you should make. I think it's a personal decision, and at the same time, I think it's very important to make sure that you're asking these questions on a consistent basis with any tool that you're giving access to your client's data, including your EHR company, right? Anything we're putting out into the world that is going to involve our clients, we've got to make sure that we are thinking about data privacy.

Limitations and safety, right to decline. Obviously, I have mixed feelings about the sign on my doctor's wall that's just like, "Ambient AI will probably be used as part of your care." There's a very perfunctory, like, for my doctor, "Well, I use this, so you're cool with this, right?" I don't think that's really informed consent. I think clients have to feel like they have a genuine right to decline. If I'm having this conversation with a client, I'm going to say, "This is a tool that I use for X, Y, and Z reason for some clients. I also don't use it for some clients. It is totally fine with me. Either way, not going to affect your care in the slightest."

Clients should feel like they can be fully, fully free to decline AI being part of their care if they want to. Obviously, consent is not a substitute for compliance. Regardless, whatever your client's consented to, if you have a change of law in your state, in your location, clients can't consent to something that's not legal wherever you're located. State laws are constantly evolving. Laws vary by state. They change quickly. 36-plus states introduced 70-plus AI Chatbot bills in early 2026.

I've been, here in Texas, talking with the lobbyists for mental health organizations down here about what we're looking at in the next legislative session, what's going to be introduced, how we as mental health professionals should consider those questions. I think this is enormously complex. I think we're coalescing around some ideas around, specifically, disclosure is incredibly important.

Certain things, like AI actually doing therapy, should be restricted. It's vital for providers to know what their obligations are for us to have a clear understanding of "What should I be doing? What do I need to be doing?" before I choose whether or not, or in what way to engage with any of these kinds of tools. Practical next steps for you. If I were thinking about how to integrate AI into my practice this year, I would make sure that I had AI screening questions in my intake for every single client, probably down to five years old.

People's children are using AI at ever-increasing rates. This needs to be assessed for. Learn how your patients are already using AI, not in an interrogation kind of way, but in a bringing some curiosity to this kind of way. Your client may feel like their AI use is incredibly helpful to them. If you're able to approach it with curiosity and kindness and compassion, they are much more likely to be open with you and allow you to collaborate in their care in that particular aspect of their life. Whereas in the past, they may have carried secrecy or shame around their AI use, which is always going to make it higher risk.

Review applicable laws and professional guidance. Know what your state says, know what APA says, know what any licenses you hold, what their guidance looks like. Consult legal counsel and your malpractice carrier if you have any more advanced questions. If you want to review on your consent forms, never hurts to get

August 21, 2026

a second set of eyes on something. Then always evaluate AI tools you're using for privacy, security, evidence-based compliance before you're considering clinical use.

I think that, in any part of our field where norms and technology is evolving as quickly as it is with AI, we will inevitably struggle to find answers that last because these tools change so fast. Developing your own internal sense of what really matters to me, what level of transparency or data privacy or whatever, what matters to me. I know we've had controversy and challenges recently around things like large platforms telling providers, unilaterally, "If your clients don't opt out of, or if you don't opt out of data being shared with us, by default, data will be shared. If you miss that one pop-up or that one login screen, you don't realize that."

I think it's going to be increasingly challenging for us to know exactly what's being done with our data and with our clients' data unless we feel like we can collectively advocate for platforms being more responsible and respectful with those types of conversations. We should always ask for, which is our right, we should ask for, in plain language, "Tell me exactly how you handle these things that matter to me."

Any platform that's not able to do that—if they're handing you a huge, long thing that's all jargon—I would personally be hesitant about that because most of us did not also go to law school. I want someone to tell me in English that I understand what's going to happen to my stuff.

I've got some references here. I'm not going to camp out on this slide for a super long time because I know that you guys are all going to get them in your emails. As I mentioned earlier, I do work for Upheal. Between when I sent in my bio and today, I got a promotion to product manager. I'm actually working in product management with Upheal, which is actually a lot of fun. If any of you have been interested in any of this and want to try out Upheal or want to get in touch, you can get an extra month free of Upheal using this promo code. We'd love to see you and talk to you there.

I think at this point, we'll transition to some questions. I've got, awesome, quite a few questions. Let me start with some that seem like they are definitely worth addressing, and then I may try to group a few things together to make sure we can cover as much content as possible. We have this question, how can a clinician tell when a patient's AI use has crossed from helpful into harmful?

I think that's a really important question because, obviously, things are not black and white. This is very nuanced conversation. It's not like, all of a sudden, one day this was great, the next day it was not. Even within a particular patient, there are going to be times when you have this conversation in one session, and you're like, "Oh, it seems like the AI tool was actually super helpful." My Gemini client, for example, who stopped smoking THC because of really buying into the idea that it could be causing psychosis—someone like that, I wouldn't be surprised at all if three weeks later, they came back and said, "Well, now Gemini told me that I probably have XYZ diagnosis," and that may or may not be helpful.

I would say, for clients in general, if AI is seeming like an adjunct to other things in their life, if they are generally well-balanced, and AI is a small percent of what they're doing, where they're getting their advice and information and support, that tends to be helpful. If you have someone who's been like that, and then they come in one day

August 21, 2026

and they say, "Over the past week, I have not been able to talk to my loved ones at all because they're all making me really frustrated, but I have been getting a ton of support from Claude, who's told me that I'm right to be mad at them." Something like that, where just like there's been a shift in the amount of support, the relationship with AI has changed.

We know how to evaluate relationships with things in general. It's a very similar question to "How do I know if my client's use of alcohol has crossed from neutral to helpful in terms of just being culturally part of what they do?" to "This person is now using alcohol in a harmful way or THC in a harmful way." We can sense when these relationships change, and that would be the main thing that I'm looking for for an actual progressive cross into something that's more harmful.

AI tends to tell people what they want to hear. How should a clinician respond to when a patient's AI is reinforcing beliefs that work against their treatment plan? This is a super good question. AI does tend to be pretty sycophantic. We've had efforts from AI companies to reduce that, that I think have been only partially successful in a lot of cases. I think that the first thing we have to understand, again, is as much as AI is not a person, it is a relational force in our clients' lives.

With a question like this, you could almost pass in, "What should I do if a client's mom or spouse or friend is advising them to do things that are contraindications based on their treatment plan that we've agreed on?" A lot of it means you're not going to come in and just immediately start talking trash about that entity in your client's life, no matter who it is or what it is. The first thing is to slow down the process. You align yourself with the client.

I think it's really interesting that you got that perspective. How do you feel about that? It seems like a few weeks ago we were talking about X, Y, and Z being a good idea, but now you've got this new perspective on things. I'm a huge fan of curiosity in the therapy room. I think that the more we can be curious about a client's processes, the less we get stuck in those rabbit holes of what's right or wrong for the client's treatment plan.

I say a lot to clients and probably an annoying amount to my friends and colleagues. A lot of these things are morally neutral. Our treatment plan is not the right choice. It is a choice. Almost nothing on the scale of decision-making is going to be a 100 perfect decision, and almost nothing is going to be a zero. Let's assume that my treatment plan with a client is a solid 85. Cool. It's a really good treatment plan. Whatever AI is recommending is a 60. That's not ideal. At the same time, I don't have to panic.

I can bring some curiosity into what elements of this AI recommendation do you like, and can we unpack some of the things that are different from what we've been talking about or working on already in our process? If a patient tells you they're using AI support alongside therapy, how do you respond to that in the moment? For me in particular, again, I feel like I'm a broken record here, but just curiosity.

We're not trying to set ourselves up as being against AI in a therapy session regardless. Some of you may be against AI in general. Valid. At the same time, with a client who's already using it and expressing to you, "This is therapeutically helpful

August 21, 2026

for me," okay, cool. Let's talk about it. Let's talk about what that support looks like. Let's talk about if there are other places that you're no longer getting that support now that you're getting it from AI.

Start from that place of your client is trying to find the best solution for themselves, and you are there to support them in that and come alongside them in helping them find a solution for their particular brain and their particular presenting problems. How does leaning on AI companion between sessions affect trust and dynamic and human-led therapy? I think pretty significantly sometimes. I think there are definitely times when I've had a client who comes in acting like they've already processed something that I know they would have been crying about or agitated about before AI entered the picture.

A lot of times, just being gentle and reflecting back that, "I don't know, man, that sounds really, really difficult. Tell me about what your experience was like on Tuesday before you talked to the AI about it. Can we get back into that spot so I can understand the impact more globally for you of what happened?"

Could a clinician be liable for a poor outcome connected to a patient's AI use, and what protects them? My compliance person answer is documentation, documentation, documentation. Just document everything. In general, don't recommend that clients begin using AI in emotionally fraught situations if they're not already doing it because we know that there are just an enormous amount of unknowns. Continue to keep an open dialogue with clients about it. Help them understand that you do care about this aspect of their use. You are curious. You are interested.

Let me see. Do the ethical concerns about environmental effects of AI data centers, are they worth the integration of AI through the process? I think that's a great question. There are obviously, especially, certain locations certainly are being significantly impacted by AI's electricity use. I think there are valid concerns here. I think if we're talking about client use, I probably wouldn't bring in, have you thought about the environment as a reason to get a client to stop using AI with most clients, just because I think it runs crossways with the curiosity.

If you're thinking about AI use for yourself, my guiding principle with AI decision-making, whether or not to include AI in a portion of my clinical practice, largely comes down to waste and the idea of what is wasteful in a particular situation. I've personally always been a-- I just knock my notes out as soon as I'm done. I feel like there are benefits to aspects of clinical AI use. I like the way that Upheal has integrated AI into a lot of things. I think it's been very thoughtfully done. I think about clinicians who've worked for me over the years who have had debilitating ADHD, who have ended up leaving the field because of documentation burden, basically.

That feels incredibly wasteful to me. If creating a therapy note based on a transcript or a summary would have kept them in the field and kept them helping clients, I think that would have been less wasteful. I think in life, in general, we're not really making zero-waste decisions. We're talking about, "Do I see a client on Zoom or do I drive to my office 15 minutes away?" Either one has its own wastefulness, but there are also benefits. There are human benefits to the things that we choose to engage in.

August 21, 2026

When we make those decisions, I think, for each of us personally, we've got to think through what's a wasteful use of AI versus what is the use of AI that helps me engage more deeply with my clients in session without thinking about my documentation burden? What helps me get more time with my family or my friends? What helps me have better life balance? I complain about this sometimes, but since I've started using some AI tools with my documentation, I'm more available in the evenings, which means I have to wash more of the dishes in my house, which stinks because I was enjoying when my partner had to do all of it. Not really.

I think that, again, what's the opportunity cost to using it, not using it? Then, if you are using it, again, ask those questions of any AI tool. What are you doing to decrease the environmental impact of this as much as possible? There are many different ways to incorporate AI, and some of them are more wasteful of tokens and more environmentally concerning than others. Whereas there are things like token caching that can make it much cheaper and more energy-efficient to include AI in some processes.

I'm seeing questions about some handouts: a harmful versus helpful AI handout. I am more than happy to send over some resources that I've used and liked. If anyone wants to email me, I think that can be included when you guys get the recording. Please reach out if I can be a resource. I would love to keep talking and be able to be helpful if possible. Getting comments about outdated AI information. I'm sure we've all seen this. You're googling something that you're not even trying to get an AI response. You're googling something random, and the AI response you get is a complete mix of good information and either gibberish or super outdated information.

I think that doing some of your own investigation of, again, what kinds of tools are clients using? The bigger AI models that are purpose-built for consumer engagement versus things that are more aggregating, like a Google search result kind of thing, tend to be better. Again, educate clients about these things. If they're asking even about something like medication side effects, the answer remains, "Talk to your prescriber about this. AI is not going to necessarily give you good information about how your Prozac and your Adderall are going to play together.

I think this is my last question because I know we're about out of time. The question is, "Could we help clients train an AI in session?" I think if you're comfortable with that, if you're comfortable with asking the client if they want to pull up and put some of those system messages in while you're there, I think that could be great. I don't think it's mandatory at all. I think it's one of those things that if you feel comfortable with it and your client feels comfortable with it, that could be a great way to express your curiosity in your investment.

I know that we are out of time. This has been a great conversation. I love all the questions, and I know that I didn't get to all of them. My email's in the chat; please feel free to reach out and let me know how I can be helpful.

Moderator: Thank you so much for joining us, Leila, and thank you to all of 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

AI in Mental Health: Part 1 – Therapeutic vs. Harmful AI Use by Patients

August 21, 2026

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.