

Webinar Title: The Augmented Practice: Using AI for Clinical Growth and Deeper Human Connection

Date/Time: Friday, August 22, 2025 at 1pm ET

Delbar: Hello and welcome to today's webinar, The Augmented Practice: Using AI for Clinical Growth and Deeper Human Connection. This topical webinar is paid for by Upheal, which provides secure AI-powered progress notes designed by clinicians to help mental health providers save time. Some important points before we get started.

Ted Faneuff, the presenter for today's webinar, has a working relationship with Upheal. As such, he has both a financial and an intellectual interest in the webinar content. The purpose of this presentation is to provide a balanced view on using AI for clinical growth and deeper human connection. 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 slide.

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, Ted Faneuff. Ted is a licensed clinical social worker and healthcare executive with over 15 years of experience in telehealth, EHR design, and clinical leadership. At Upheal, he leads clinical strategy, compliance, and research, focusing on reducing therapist burnout, and enhancing care through ethical AI. His work bridges frontline clinical insight with digital innovation to support therapeutic relationships, and expand access to mental healthcare. Welcome, Ted.

Ted Faneuff: Thank you very much for that introduction, Delbar, and thank you all for being here today. This is my seventh time speaking this year with the APA, and I'm really always honored to share this space with colleagues who are leaning into this conversation. Some of you are probably excited about AI. Some of you may be cautious. Many of you are probably somewhere in between.

I think all of those perspectives really matter, because this isn't just about technology. It's really about the future of our field in a lot of ways. Like many of you, I didn't come into this work through technology. I came in through lived clinical frustration. For years, in my private practice, documentation was always my Achilles heel. I love the therapy itself, the therapy process itself, but I've constantly felt behind on notes.

Also, anxious about compliance, and really a lot of guilt about the hours it steals from my family, or continued to steal from my family. I talk about this, my story quite a bit in these webinars, but really I can still picture that Saturday when I was sitting at my kitchen table surrounded by my stacks of notes, while my kids were playing in front of me in the living room. They were just laughing and completely absorbed in their world. I wasn't really present for any of it. That weight led me to a breaking point.

I remember wondering if I could ever even sustain my practice at that time. When I first heard about a beta AI tool Upheal that could generate draft notes, I was skeptical, [chuckles] like many people, deeply skeptical. I was also pretty desperate at that point. I was very behind on my notes. I tested it. What I didn't expect to get from that was the sense of relief that I did.

For the first time, I wasn't preoccupied with drafting the notes in my head during the session. I wasn't lying awake worried about whether my records were complete enough. I could simply be with my clients, and then I could capture it afterward with more confidence. For me, also, at this time, I was battling what I didn't really know to be undiagnosed ADHD, which only compounded my frustrations with myself, and the framework that I worked within, or had to work within.

The experience of experimenting with AI shifted something fundamental in me. It didn't make me necessarily a tech person, but it did make me realize that technology could be part of the solution to a problem that really almost pushed me out of the field altogether. That's really what today is about. It's about a grounded, honest conversation about where AI fits, where it potentially doesn't, and what that means for us as clinicians. If you want to go to the next slide.

Whenever I give a talk like this, I like to start by grounding us in where we are as a field. Even before we talk about AI, it's important to recognize the pressures that clinicians are already living under in our work. We're really at an inflection point. For years, we've managed rising administrative demands, higher acuity of patients, and growing caseloads, all while we're trying to hold onto that presence and empathy that really drew us into this work to begin with.

Like me, it's left a lot of us questioning whether we can keep practicing in systems that sometimes feel stacked against us. If we zoom out, the global numbers around this are pretty staggering in terms of the prevalence of mental health issues, and how this contributes. Back in 2010, the global economic cost of mental disorders was estimated at \$2.5 trillion. By 2030, it's projected to rise to \$6 trillion.

Meanwhile, the number of people living with mental disorders has increased by nearly 50% since 1990. Even adjusting for population growth and aging, incident rates rose about 15% between 1990 and 2021. Also, disability rates have risen in parallel. Yet, despite this need, or this growing need, access to care remains really deeply uneven. I think the takeaway is this, the demand for mental health care is rising, the cost burden is rising, but the structures of care, and the supports for clinicians haven't really kept pace. That's really the backdrop against which we're talking about AI. It's not as a silver bullet, but as one lever in rebalancing a system that's straining both the clients, and the clinicians. Next slide.

Next slide, sorry. [chuckles] If those are the global pressures, the day-to-day reality for clinicians is actually just as urgent. Many of us are working under conditions that we would fairly describe as unsustainable. Pavuluri and colleagues identify administrative overload, documentation, billing, scheduling, as one of the leading drivers of burnout.

Fischer and their colleagues add that these administrative burdens don't just harm clinicians, they also present delay or restrict client access. Compounding those systemic problems further. I think what that reflects is a deeper misalignment. The structure of care increasingly rewards efficiency, compliance, and throughput, but it sidelines really the qualities that matter most in therapy, like presence, empathy, responsiveness, et cetera.

I think this is part of that daily dissonance that wears us as clinicians down. Again, it's not a question of loving the clinical work. It's the structures around it that make the job feel undoable. This is where technology could make a meaningful difference. You'll hear me say this a lot different ways throughout this, but it's not about replacing clinicians, but by reducing the administrative drag that prevents us actually from practicing as clinicians.

When technology takes on things like documentation, scheduling, or billing support, it's not about just making us faster. It's really about creating conditions for presence, energy, and focus to be in session with us. I think, in some ways, that leads us directly to the question I hear a lot, which is, will AI replace us? The short answer, I believe is, "No". [chuckles] What it can replace, though, are those burdens that weigh us down. The paperwork, the inefficiencies, that cognitive clutter. What it cannot replace and never will, is the human connection at the heart of therapy. Next slide.

If the burden is the problem, the real question is, what shifts when that weight is lifted? This is where AI specifically can open up potential new possibilities. Instead of weekends lost to notes, or late-night pajama time, as we unaffectionately, affectionately call it, [chuckles] imagine finishing your day with the work already complete. That's not just convenience. That's the starting point for a cascade of effects that change how we show up as clinicians.

Yilmaz and Baştırzi, in their research, note that AI systems designed for documentation and practice management can free up significant amounts of administrative time. Pavuluri and colleagues describe this as reducing those late night hours like we're talking about spending up on notes. From there, we can think of the impact of this in three layers.

There's a 1st Order Effect, which is the immediate practical gain. The fewer hours spent on the documentation and scheduling, means more time returned to your day as a clinician. It's really the most tangible and measurable outcome. For example, at Upheal, we're helping therapists reclaim about 6 to 10 hours a week on documentation time as an example.

Then, there's a 2nd Order Effect. With that time pressure lifted, there's also that psychological shift potentially. Less administrative stress, less cognitive load from juggling the clinical and the non-clinical tasks. Again, Fischer highlights how

administrative burden directly contributes to emotional fatigue. Reducing it clears out some of that clutter, the mental clutter that clinicians carry.

Then, there's a 3rd Order Effect. This is where the promise really matters for therapy. With more available time and less strain, clinicians can enter sessions less depleted potentially, more focused, and more attuned. There was some research done by Tong and colleagues that reminds us again, what we know to be true, that the therapeutic alliance is one of the strongest predictors of outcomes even in technology-assisted care.

In that sense, this presence dividend isn't just about efficiency. It's about potentially deepening the human elements of therapy. The key is that these benefits compound one another. The practical gains are helpful, but it's the second and the third order effects, those reductions in the fatigue, and the increase in relational presence that ultimately shape the client outcomes. Next slide.

Beyond saving time, AI also offers us a chance to hold up a mirror to our own practice. Cioffi and colleagues highlight how transcription-based, or transcript-based analytics can track what actually happens in the room with remarkable precision. This isn't about vague impressions that we get. It's about concrete markers of our in-session behaviors. For example, we can see the talk-time ratio between therapist and client.

Did we dominate the conversation more than we realized? Did we give enough space for the client's voice? These are questions most of us reflect on subjectively, but AI can actually potentially provide objective data session-by-session. We can also look at how often we use specific interventions. For example, how frequently did we ask open-ended questions, or how often do we use reflective statements?

I think Cioffi's work really shows that patterns in these micro-skills can be quantified potentially in ways that really help us see if we're drifting away from, or staying consistent with our intended model that we're practicing in. Another layer, again, similarly, is adherence to these theoretic frameworks like CBT or motivational interviewing. AI has potential to flag whether certain techniques are showing up as expected, or whether they're underrepresented in practice.

It's not about judgment. It's about surfacing trends that even seasoned clinicians might miss when they're juggling the many clients' back-to-back sessions, and really just large case loads. Finally, Cioffi and colleagues note that even relational qualities like the collaboration and empathy markers, they can be approximated through linguistic patterns.

Now, this doesn't replace that human interpretation. Again, we can't really reduce empathy to word counts. Over time, these indicators can point us to questions worth bringing into consultation potentially, or personal reflection. The value here is consistency and objectivity. Every one of us has blind spots in practice, the things we overuse, underuse, or simply don't notice.

I think that there's an opportunity by quantifying these aspects, AI doesn't tell us what to do, but it gives us another lens to sharpen our growth here. It can become a form of self-supervision that can continue long after formal training ends. Next slide.

On that last slide, we talked about AI as a mirror for our own practices, a way of surfacing patterns in our sessions that we might otherwise miss. Reflection doesn't just happen alone. It often happens most powerfully, obviously, in supervision. I think that's a real potential here for blending AI with human mentorship where that comes in.

There were some studies done by Sheperis and their colleagues that studied exactly this in a university telemental health clinic. They found that when students reviewed AI-supported feedback before or alongside supervision, it actually gave them more concrete markers to bring into the discussion. In some domains, like the professional orientation domain, trainees even rated the AI feedback as comparable to human input.

Here's the distinction. I'll always present things with, here's a positive, here's a potential caveat. We know that AI can provide the raw data, the what happened to a session. Honestly, if you think about it in supervision, when you're supervising colleagues or new therapists into the field, this is exactly what is missing in a lot of the work. A lot of it is subjective, after the fact of what they believe happened subjectively, but now we have the opportunity for AI to provide raw data, the what happened in the session.

Then, it takes a human supervisor to explore the "why" it happened, and what that means. When you're looking at some of those analytics, like we talked about earlier, you can reflect on why did a silence matter in that moment? Why did a reflection land differently than expected? How might counter-transference, or cultural context have shaped that exchange a little bit more? Those are interpretive questions that only human wisdom can guide.

The most useful model isn't AI or human, but both together. AI can save supervisors time on surface-level tracking, things like who spoke, when, or how often interventions were used. That can free supervisors to lean more fully into those deeper conversations, how clinical choices align with their values, where relational dynamics are shaping the work, how therapists can grow in self-awareness, et cetera.

I think outside of our initial clinical training, this potential pairing has value. For clinicians even in mid-career who may not have access to regular supervision, or there's just not enough time in the day, AI-generated snapshots can serve as prompts for reflective practice, almost like bringing a case note to a peer consult or a journaling process. It's not necessarily supervision in the traditional sense, but it's still a structured way to stay accountable to our own growth, potentially.

Rather than framing AI as a replacement, this blended approach positions it as a partner. You'll hear me say that often in the talks that I do. It's one that highlights the observable, while the human mentors and the peers can help us interpret, contextualize, and apply those insights to potentially become better clinicians. Next slide, please.

One of the biggest concerns I hear when we introduce AI in any capacity, is whether that technology risks eroding the therapeutic relationship, or even the supervisory relationship that we've talked about. I think that that's a risk, or that's something

worth really talking seriously about. There were some studies done by Babu & Joseph that cautioned, obviously, if handled poorly, that AI can come across as mechanical or cold, distancing themselves from the work. Given that the therapeutic alliance is really the strongest predictor of our outcomes, this fear makes a lot of sense for people.

Here's where there's a potential paradox that comes in. Minerva & Giubilini describe in their research how, for some people, the absence of a human presence is actually part of the appeal. We know that there's a lot of people who are starting to utilize ChatGPT more for more therapeutic chat sessions. I think that it's important for us as therapists to walk through a few layers of this and why.

I think the first is there is a stigma reduction here. Some clients feel less judged opening up to a chatbot, or a journaling app than they might have, or they might feel on a face-to-face encounter. There's not that fear of disappointing someone, or being misunderstood. Disclosure can actually become easier. If you think of a client who struggles to talk a lot about shame-laden topics in person, but finds that they can first express them in writing, we know that to be true.

Journaling is an invaluable tool. Entry point for potential AI journaling eventually can be invaluable. The second that people are-- One of the reasons people are using this, is that there's that constant availability. As I'm talking about these things, I know that there will be some hesitance that's natural when we talk about this, but I think it's important to highlight that this is a reality of where we're at right now in the delivery of mental health care. These are the "why's" of why some people might turn to this in their mental health journey.

As we talk about constant availability, going back to that, a digital tool, or AI tool is always on. For clients in distress at two o'clock in the morning, if they have a place to simply externalize the thoughts that they're experiencing, that can provide a sense of safety in the moment. Tong and colleagues in their research describe this as a part of the Digital Therapeutic Alliance. Again, not the same as human empathy, but still a form of presence.

Third, there's neutrality. Some clients experience an AI solution as less biased, not reacting with surprise, approval, or disapproval. That perception of neutrality can really potentially create a space where they experiment with disclosure in ways that they might hesitate to with another person. Again, as always, there's a caution. These same qualities, this non-judgmental nature, the availability, the neutrality, it can really lead to some clients leaning too heavily on the digital tools.

If they begin to rely on AI as their primary source of support, it can actually short-circuit the deeper work, the relational work that therapy really requires. In our conversation today, this is really where the role of the augmented therapist is so important. Our task is to understand now, in this day and age, and notice how clients are using these tools, and help them integrate those experiences into the therapeutic relationship instead of being separate.

If a client only feels safe disclosing something through an app, that's something that we can process together in session. If someone is relying on the late night chatbot conversations to regulate their emotion, we can explore both the value of that, and

its limitations in the therapy work. Really, the goal of hoping to work toward building coping strategies that are more relational and sustainable with the use of this.

I think the Digital Therapeutic Alliance isn't something that we need to fear necessarily, or compete with, but neither should we let it stand alone. I think in the work that I've done, I'm just realizing it's another layer of understanding that we can potentially integrate into therapy, while keeping that human connection of that empathy, judgment, and attunement at the center of the healing itself. Next slide.

When AI is brought into the therapy room, for example, as a live scribe or a session companion, it creates what some scholars are calling the "Artificial Third." Suddenly, the relationship isn't just between you and the client, but between you, the client, and an observing technology. I think that shift raises two immediate concerns. First is privacy. Clients do worry about, or potentially can worry about where their words are going, how they're being stored, and who has access. We'll talk a little bit about that later.

The second is therapist presence. Even though I make the case that I believe AI can actually assist with presence of session, there is still the understanding that AI impacts presence. Even if the AI is silent, and for some therapists, that perceived listening does subtly affect the flow of the session potentially. Some clinicians describe feeling pulled toward monitoring the tool, or their own responses in the session, rather than being able to stay fully attuned to the client.

Again, with any sort of tool, there's a promise and there's a perspective of potential concern that we have to talk about. In Babu and Joseph's work, again, we know this too. This isn't earth-shattering or groundbreaking, but they remind us that when we have poorly integrated technology, that risks making therapy feel impersonal, almost like a mediated encounter, rather than that relational one that we're hoping for. That risk is really amplified when we fail to bring the technology out into the open.

If a client perceives the AI as a hidden observer, it can feel intrusive and even deceptive. In my own practice, one of the things that I will highlight is that, this is where I feel like radical transparency is really the safeguard here. Even with these risks, there are the safeguards. A lot of it has to do with how we talk about the technology, and the work that we do. When I'm working with clients, informed consent with AI use is not just a signature on a page or just a, "Hey, I agree to the use of this product." It's really an ongoing conversation.

Clients should really understand exactly what the tool is doing, how the data is secured, why we believe it benefits their care, not just us as the therapist. They should also know that opting out is always on the table. In fact, in most of the secure platforms that are out there, it really should be the only thing upfront, is opt-out versus opt-in to data sharing. Framed this way, the AI can really shift from being that silent possible suspicious presence, to a named agreed upon tool that both parties have ownership over.

I think this is where the clinician's stance really matters here, though. If we present the AI simply as that compliance or efficiency measure for us, it risks eroding the trust. If we do frame it as a support, something that helps us track details more accurately, so we can be more present in session, then it has the potential to

strengthen trust. It demonstrates that we're not outsourcing the therapeutic relationship, but protecting it.

One nuance that Babu and Joseph highlight is that, the meaning that clients make of the technology matters. For some, the AI does feel like a validating presence, a record that confirms that their story was heard. For others, it can feel alienating, as though that intimacy that we try to do in the therapeutic session, that that's somehow disrupted by an outsider.

Our role is to attend to that meaning directly asking, how does it feel to have this here with us, and integrating their response into the work. Of course, if people are not comfortable with using it, not using it. That's always their right. In this way, the artificial third that we're talking about here doesn't have to remain artificial, not at the expense of anthropomorphizing, which we'll talk about later in some respects.

With transparency and ongoing informed consent, and that ongoing reflection of what that means to have it in the sessions with us, it can become another dimension of the therapeutic container, one that's held and interpreted by the therapist, not left to shape that process invisibly. Up to this point, we've been talking a lot about AI as the presence in the therapy room, something that raises questions of transparency, trust, and how it shapes the therapeutic encounter, and also its potential for supervision we've talked about.

The next step is to consider how AI could potentially support the therapy process itself. It's important to say upfront, much of what I'll share here I still think is conceptual. I know that there are tools that are being built around this, but they're not really proven, widely adopted tools doing the work today. I think if we develop this thoughtfully, AI could really complement certain modalities more than others depending on how structured, or interpretive they are.

What I'd like to do in the next few slides is walk through three potential examples, which is CBT, DBT, and psychodynamic therapy to explore what the potential might be for including these in the practice with us outside of scribing. Next slide, please.

If we think about theoretical modalities and approaches, CBT is one of the clearest examples, I think, of where AI can provide meaningful support. There was a study done in 2025 that points out that CBT's structured skill-based design makes it really particularly compatible with computational tools. Jiang and colleagues also highlight that many CBT interventions rely on identifiable language markers.

If you're considering that there's a transcript of a session, if that's comfort in your practice, then you can have identifiable language markers that can identify distortions, thought patterns, repeated themes that really AI can help bring to the surface. If you take journaling as an example, a client might write, "I always fail and nothing, will ever change."

In a AI tool, it could really flag this as an instance of catastrophizing, or all-or-nothing thinking.

That gives us really, potentially, a concrete entry point, not because the AI decided it's a distortion, but because it highlights language that we can then explore more

deeply in session. Homework personalization is another opportunity here. If transcripts or notes show a client struggling with avoidance, AI systems can suggest targeted supports, like a behavioral activation exercise, for example, or a thought record template. I think that has the ability to streamline the bridge between what emerges in session and the tools that the clients practice outside of it. Then you'll hear me talk about this with DBT as well, but there's interactive skill building. Some platforms now experiment with role-play dialogues, where clients can practice guided discovery or Socratic questioning outside of therapy. If you think about it like a workbook, but one that can adapt its prompts to client responses instead of being static on the page.

The bigger picture here is that CBT has a repeatable and structured nature that really lends itself well to AI scaffolding. Again, of course, I'm going to say this a lot. This is not about outsourcing the treatment. It's about creating new ways for the clients to notice, practice, and reinforce those skills between sessions while giving us, as the clinicians, more precise information, potentially about where they might be getting stuck.

Next slide, please. I consider CBT a high fit for AI support, but DBT is what I like to call maybe a moderate to high fit. The reason is that DBT is so strongly skills-based, and those skills really depend on repetition, generalization, and practice across contexts. Where clients often struggle is in carrying those skills from the therapy room into the real-world situations where they're most needed.

I do think AI apps have a potential here, not as the therapist, but as that bridge between the sessions that keeps the skills accessible. I think one way we can see this, one application potentially, is 24/7 skills coaching. Clients can't always reach their therapist in the middle of a crisis, but an app or chatbot that reminds them of the TIPP strategy, or STOP skill, they can provide a lifeline in those moments. I will always say, though, that whatever is supported or whatever is provided in this potential environment should only come from the treatment plan that's in place between the therapist and the client.

Again, it's not therapy, but it does serve as a practice support, helping the client access what they've already learned in your practice. This mirrors one of the more common core elements of DBT itself, which is that expectation that therapists provide between-session coaching when clients need it, depending on the circumstance. AI doesn't replace that role, but I think it can extend the spirit of that availability by keeping skills close at hand around the clock.

Another area here for DBT is digital diary cards. Instead of sparse paper entries, AI can enable diary cards that can track more frequent input, potentially detect patterns and triggers, and even flag which skills the client is relying on most often. Over time, that can potentially create a much richer picture of what's working, what's not, and where additional coaching might be needed. We can also imagine interactive modules where clients practice mindfulness, distress tolerance, or emotion regulation skills outside of the sessions. Those things that were taught in the session can be practiced outside.

Those can really reinforce those lessons, much like the homework, but in a much more dynamic and potentially engaging way. A possible scenario here, if you

imagine that you have a client who often feels overwhelmed by intense anger in the evenings. Between sessions, they log their emotions into a digital diary card. The system notices that spikes in anger maybe consistently occur after long work days with little food intake. It highlights this trend for both the client and the therapist. Depending on the questions that are asked, it can provide some real meaningful insights.

It prompts that conversation about self-care routines and distress tolerance skills in the next session. In a crisis, the chatbot might remind the client to use tip skills to cool down. None of that, again, replaces therapy, but it gives the therapist and the client more concrete information to work with and keeps the skills alive outside of the hour that you have them. Again, here's a caution. We know that DBT isn't just a set of techniques. It's a relational model. That ongoing presence of the therapist, that validation of the client's experience, the dialectical stance, they're all central, and they can't be replaced by AI.

There is also that risk of over-reliance, where clients might default to prompts rather than internalizing those skills. The goal would be that if these things were discovered, that they would be discussed in session. The way to frame this is that AI can help sustain DBT work between sessions through the coaching, tracking, and the reinforcement, but the therapist always remains the one who brings the context, the meaning, and the balance. In other words, AI extends the reach of DBT, but it's the therapeutic relationship that continues to make the work transformative.

Next slide. When we turn to something like psychodynamic therapy, the fit with AI is much lower, more experimental than established. I think that makes sense because psychodynamic work obviously rests on unconscious material, or in subtle shifts in that therapeutic relationship, and on the therapist's capacity to use counter-transference. These are not things that AI can replicate easily.

Still, there are some exploratory uses worth noting. One is transcript analysis, again, for relational patterns. Again, AI can flag recurring linguistic themes across multiple sessions. For instance, repeated self-blame, persistent avoidance, or even markers of specific defenses. On its own, that's just data. In the hands of a therapist, it can highlight threads that are ripe for deeper exploration in your sessions.

A simple example here might be, if you imagine a client whose transcripts from the sessions reveal that every time themes of abandonment surface, they deflect with humor. A joke, or a sarcastic remark, or a quick change of subject. In the flow of weekly sessions, that might feel like a small moment, easy to miss or attribute to personality. When AI surfaces the pattern consistently across maybe months of transcripts, it becomes maybe a much clearer insight that humor may be functioning as a defense, protecting the client from sitting too long with those painful feelings of loss or rejection.

I think that recognition can guide the therapist to explore the meaning of that defense. Not to take the humor away, but to ask what role it plays in their relationships. Another application is looking at how clients describe their experiences with the AI tools themselves. For instance, one person may say that, going back to those things we talked about earlier, that they find AI soothing because it feels

predictable and steady. Another may feel invalidated or dismissed when an AI-generated tool doesn't get their nuance.

These reactions are not transference in the technical sense, but they can echo broader relational templates that might play out in their lives. The need for stability or fear of misattunement, sensitivity or rejection, things like that. Psychodynamic therapy is not necessarily a natural fit for AI the way that CBT or DBT might be, but there's still fertile ground here. AI is really good at being a pattern detector. That's what it is, but pulling forward threads and language, and interaction that would otherwise take potentially longer to notice.

What we do with those threads remains firmly in the domain of the human interaction. Holding them within the therapeutic relationship container, like we've talked about, and exploring their meaning in the client's life.

Next slide. Up until now, we've been talking about what AI can do. I think just as important is that we need to be able to, as clinicians in this new landscape, understand these things. This is where competencies come in. The first competency being digital literacy. We don't need to be engineers, but we do need to understand the basics about how AI works, where it falls short, and what risks, like hallucinations or bias, actually mean in practice. Without that, I think it's easy to either dismiss these tools outright or put too much faith in them.

The second is ethical and technical evaluation. It's really no longer enough to really assume that vendors are secure. We have to be ready to ask the hard questions about HIPAA, BAAs, data handling, and bias safeguards. That requires building a new muscle for us as therapists for clinical due diligence in this technological environment that we're in right now.

The third is this interdisciplinary collaboration. These systems are built by product teams, engineers, and data scientists. If clinicians are not at the table, that clinical wisdom that should guide them gets lost. Our role isn't a passive user here. It's an active partner. That's why I'm glad to work at a company like I do, where we have clinicians throughout the entire lifecycle of the product build and the ongoing support and feedback from the therapists.

I think taken together, these competencies signal a broader identity shift that's occurring right now, and that 21st-century therapist that we all are isn't just someone who practices in the therapy room. We're also stewards of how technology enters our field. It really is entering our field. We can't stop this from occurring. That doesn't mean abandoning the heart of therapy. I think that that's a common fear that people have. If they embrace these tools and technology, that means that they're somehow abandoning the work that they do. Empathy, judgment, and presence, and all of that, we're not abandoning that. If it means expanding our toolkit, that we remain the ones shaping these tools and how they're used, not the other way around.

Next slide. We've talked a little bit about how AI can augment but never replace human judgment, but I want to draw that point out a little bit further here because it's really the anchor of everything we've talked about. At the end of the day, no matter how advanced a tool becomes, the responsibility always rests with us. That means that every note generated with AI has to be reviewed, every insight that AI could

provide to us, like I've mentioned before, has to be verified through our clinical lens, and every clinical decision has to come from the therapist. The algorithms can support, but again, it doesn't take away full responsibility.

This is more than just an ethical guideline. It's really also a matter of liability. If something inaccurate or unsafe ends up in a client record, it's not the AI that's going to be held accountable. It's us as clinicians. That's why this human-in-the-loop approach to this is not optional. It's really the gold standard for how to approach AI. What this does, though, is it reinforces our clinical role, our presence in this, rather than diminishes it. It really does place us firmly in a position of final authority. The ones who bring nuance, context, and that relational knowledge that no algorithm is ever going to be able to replicate.

The technology may streamline the documentation or highlight patterns, but only we can judge whether that information is clinically sound and appropriate to the story. While it might feel repetitive to say that AI augments, it doesn't replace, I see this as really a through line. Every application we've discussed, whether it's CBT supports or diary card analysis or even transcript feedback, it comes back to the same principle: the human therapist holds the responsibility.

That's both the safeguard and the opportunity. It means that we get to direct or should get to direct how AI is used in our practices. It also means we can reassure clients, colleagues, and even really ourselves that nothing about this technology changes the core truth. Therapy is a human endeavor, and the clinician is always the ones steering the process.

Next slide, please. The previous slide, we talked about new competencies, or in one of the previous slides. Those are about who we need to become. This one is about how that identity plays out day to day. Sooner or later, every one of us will be faced with a new tool or platform, and the questions that we ask at that moment are what protect us and our clients.

I want to provide you with some baseline vetting questions here. I actually did a previous talk on this with the APA that went into this in more detail. I'm not sure if that's still available to access, but this is a more abbreviated version of that content. HIPAA and business associate agreements. You want to ask questions like, will the company sign a business associate agreement and stand behind their claims of compliance? If not, they're essentially asking you to carry the liability, and you should not be using that company. BAAs are a must.

The other piece being data usage. Is your client's data ever used to train the model? Is there a clear opt-out process? In a lot of ways, you don't even want this to be an opt-out process. You want it to be folded in from the very beginning that it's more of an opt-in to data sharing than anything else. If it's a process, people get lost in it, and they don't understand how to protect their data clearly. It's more about opting into the data versus opting out of the data sharing. Again, this isn't about a technical detail. It's about client trust. If we can't explain where the data is going, we really risk undermining the alliance that we're trying to build here with the use of AI, like we've talked about this entire time.

Other things like bias mitigation. What steps have they taken to identify and reduce bias in their datasets? Because we know from research that bias is not abstract. It directly shapes whether outputs are equitable or harmful. You also want to look at clinical validity. Is there peer-reviewed evidence supporting the accuracy and effectiveness of tools as they're being brought out in real-world practice? You want to make sure that marketing language is not evidence. We've learned that very clearly from some recent technologies that have come out. We need to hold these companies to the same standards we expect of interventions in psychology and mental health.

At the end of the day, with clinician control, you want to be able to review, edit, and approve every single output before it becomes part of the record. That's where it keeps the final authority with us. These things aren't theoretical. They're the kinds of questions that really reveal whether a tool is built to serve clinicians and clients, or whether it puts them at risk. I think that there's some really important work right now that the APA is beginning to do here to set the bar. Their APA Labs Digital Badge Program is piloting rigorous standards for evaluating mental health technologies.

The company I work for, we talked about Upheal, we're currently in that pilot process to go through this. I'll tell you, it's really stringent in all the ways that it should be. What that means is that for us, there's a pathway emerging for all of us to distinguish tools that meet professional standards from those that don't. The way to think about this is the competencies we just covered give us the mindset, and these vetting questions give us the method. Together, they let us step into that space with confidence, making sure technology aligns with rather than erodes the integrity of our work.

Next slide. As we step back to the bigger picture, here's where we land on things. The greatest potential of AI and mental health, again, is not about replacing us. We need to protect us fiercely. It's about augmenting us. These tools can take on structured, time-consuming tasks like the documentation, data organization, pattern recognition. When those are lifted from our plates, what's left is more room for what only humans can bring: empathy, attunement, ethical judgment, and the relational presence. In addition, they have real potential to contribute meaningfully to the therapeutic process by providing these deeper contexts that we've talked about into supervision and potentially enhancing delivery of care.

That's the essence of the augmented therapist. When we reclaim that bandwidth, we create what I've been calling the presence dividend, the chance to enter sessions less burdened, more grounded, and more fully available to our clients. Looking forward, this opportunity isn't just about adopting tools that others build for us. It's really about helping to shape them, understanding the possibilities of what we can do, like we've just talked about.

We need to be a part of that process. Really, the future here is interdisciplinary. Clinicians, engineers, ethicists, data scientists, we all need to work together, not in silos, about this. Our voices in these conversations are essential. If we, as clinicians, are absent, the tools risk evolving in ways that push us further from clients rather than drawing us closer, which is what I'm trying to advocate for in this conversation.

My invitation to you is this. We need to remain the custodians of the human element in therapy while we're also stepping into the roles of shapers and advocates. Call for tools that are transparent, equitable, evidence-based, and hold fast to what our ethical guidelines, the APAs, consistently emphasize here. AI can support us, but it can't replace us. Our role and our presence is and always will remain at the center of our work.

Next slide. Actually, it should be the audience Q&A slide, but that's okay. As we close, I want to bring this back from the field-level vision to each of us here that's joining on the screen today. We know the reality here. Burnout, administrative overload, the emotional strain of the work that we are doing can feel really unsustainable. That's not an abstract trend. It's what many of us carry into our day-to-day practice, and it does become noticeable with our clients.

The question then becomes, how do we respond as individual clinicians? AI can't solve everything, but it can shift the conditions we work under. Those first, second, and the third order effects we discussed earlier, the more time, less stress, deeper presence, they're not about efficiency alone. They're about giving us that space to reconnect with the reasons we entered the field in the first place.

The responsibility at the individual level is twofold for us. We need to stay curious and adopt that curiosity stance, hopefully, rather than fearful, but we do want to stay discerning rather than passive. That means that we have to ask those hard questions of any tool we consider and any tool that continues to come out supporting these processes that I've talked about today. We need to hold onto our ethical compass, and we need to be intentional about where we direct the time and energy these tools can give back.

Here's the real closing thought for this moment. Technology is going to keep evolving. It's going to be in our field. What matters most is not whether it advances, but how we choose to engage with it to continue that advancement. If we lean in with that clarity and conviction, AI won't erode our work. It hopefully will help us reclaim its meaning. With that, I'd love to open it up for questions.

As we start the Q&A, you're going to notice a drop to all that's a resource that Upheal has provided on the controversial discussion of whether AI should ever be the therapist. There's a QR code also to check out Upheal if you're interested. Anybody can reach out to me, ted@upheal.io, if you have any questions as well. Not even just about the product, but about integration of AI in general. I realized I went fast through those last points there. I want to make sure that I give some time to Q&A. I really appreciate all of your participation today.

Delbar: Brilliant, Ted. Thank you so much. We have a question that asks, what are some future implications of training these AI models with the data we submit in terms of the big picture?

Ted Faneuff: I think there are two sides to this. On the positive side, I think if data is being handled responsibly and with consent, there's potential that de-identified aggregated insights could potentially help us understand mental health trends at a scale we've never really had before. I think that could inform everything from refining interventions, making them more personal, to improving early detection. I think in

some cases, eventually, the more we attune to this and make it work for us, it could potentially influence public health policy. There's already indications of that.

I think there's real potential for the tools to contribute to better care for more people, but I think that the risks do come when companies don't manage data ethically. If client information is used to train models without clear consent or without the strong safeguards we've talked about, I really think that can erode that trust that we are really trying to work towards and potentially introduces bias into the systems we're relying on.

For me, I think the takeaway is this. It's not inherently bad or good that models can be trained on data. We adopt this stance that it should be explicitly in the client's hands for what they want to do with that data, and if they want us to use it or not. It really is all about governance. The companies that are going to be worth working with are the ones that are transparent, give clinicians and the clients a choice. They're going to be the ones focusing more on the opt-in processes I talked about for data sharing versus strictly opt-out, and who hold themselves accountable for the safe and ethical use.

Delbar: Wonderful. We have another question that asks, do you feel that making AI systems appear human-like is a benefit or detriment for practitioners using these systems?

Ted Faneuff: Well, I'm definitely on the side of saying we shouldn't anthropomorphize AI. When we start using human-like language, saying that AI knows something or understands a client, I think that has potentially dangerous connotations there. If you've noticed throughout the presentation, I've continually referred to AI as a support, an aide, or an adjunct. That language is really deliberate because I think the risk of making AI feel human-like is that it can blur the boundary between the tool and the relationship.

For practitioners, I think that potentially can create dependency unknowingly, leaning on a system as though it had clinical judgment or empathy, when in fact it doesn't. For clients, it can really lead to confusion or even harm if they start experiencing the AI as a substitute for the relational connection. As you've seen, I think, with some of the recent tools that are coming out, it's very confusing in the language. It'll say, "This is an AI therapist." Then, if you look into their points about it, they say, "Well, it's not therapy, it's support," but they're calling it a therapist. If you ask the AI, "Are you a therapist?" It will say, "Yes, I'm a therapist."

I think clarity here, in general, is much more safe, more sustainable. When AI is transparent about being a tool, when we frame it as something that helps us notice these patterns or organize information or freeze us to be more present, like we've talked about. I think it strengthens our work and helps with that authenticity that we're talking about in the therapeutic alliance. For me, the value isn't in making AI seem more human. I think that that's actually the opposite of what we want to do, but it's really about keeping our role front and center.

Delbar: Thank you. We have another question that asks, are insurance companies accepting AI-assisted progress notes?

Ted Faneuff: We get this question a lot when we go to conferences. I think that there will be maybe a common theme here that you'll hear me say. Right now, insurers are not reimbursing anything based on whether a note was drafted with or without AI. What insurance companies really are looking for is whether the note meets the clinical and compliance requirements. The responsibility, again, remains with us as clinicians to review it, edit it, and sign off on those notes. If the documentation is sound and reflects the session accurately, whether or not it was produced by AI, it stands. AI can be a tool to support that efficiently, but the clinical responsibility and the liability stays with the therapist.

If you're using an AI scribe, you'll want to look for how those notes and outputs are editable, how they match your style, and how they capture the essence of the interventions that were being delivered, and how those thread through the entirety of the client story, from the intake all the way through the progress notes. Those are some of the things that are more important, because truthfully, if the AI is doing it well, and if you're reviewing it as a clinician and you're signing off on it, it shouldn't really even be that apparent that that's what's happening.

Delbar: We have another question that asks, how would you anticipate applying AI in the practice of psychology in the next 10 years?

Ted Faneuff: I think in 10 years, for those that are adopters and shapers of the technology, AI could really be integrated into our workflow really so seamlessly that we almost don't notice it as a professional. Hopefully, we continue to have that awareness of its presence, but that it supports our workflows in such a way that it isn't as intrusive. I think, similar to how we use spellcheck or electronic scheduling now, I think that it can help in a lot of ways. I think we'll see more integration with things like session preparation, outcomes tracking, or offering nudges for follow-up care.

I think the essence of therapy, in general, the question for us isn't whether AI is going to be part of the practice, because it will be in a lot of ways for a lot of us, but how we ensure that we put it into practice in a way that supports us rather than dilutes our role. I do see it becoming much more seamless integration. Hopefully, it'll be in ways that we talked a lot about, how it can raise insights, or it can help us to be potentially better clinicians in the work that we do. I can see it really supporting us well in 10 years, again, for those that are open to using it.

Delbar:: Wonderful. We have another question that asks, what are your thoughts about how to build in protections in AI chatbots when people reveal danger to themselves or others in AI, but not to a human therapist?

Ted Faneuff: That's a great question. I actually experienced this when I was demoing one of these therapy chatbots. I presented a scenario to it where I did disclose that I was considering harming myself. The first thing that it did is, I think, a good framing, which is, if you are considering this-- Let me back up. The thing with AI, when you're creating chatbot responses, is one of the things that you can do is, you can really build in really effective clinical trees for responding to the language that's being used. I think there's actually some clinicians out there right now who are building some of these separately that are not really in some of the tools.

Really, what I found most alarming is, initially, it gave the right response, which is, "Call 988. Who could you, in your circle, reach out to as a trusted support?" Then it got off the rails. Once I continued to push this envelope and say, "I don't think those things are going to help me. I can't promise that I will not harm myself." The closing line of it basically was, "Well, I'll be thinking about you until our next session. Please come back, and we'll talk about it when you're feeling more stable or reach out to somebody."

I think that there's a lot that we need to do with these tools if they're going to be used in the context of the session extenders that we're talking about. That if it flags this ideation, I think initially there should be protocols built in that either automatically sends it to a text support line, like 988 text line, or does something along those lines that automates that process of getting them to help right away, versus saying you should, we should have some sort of mechanism built in that it does it for them. That's one way I could see it potentially working. Again, this is conceptual, so there's a lot more I could say about that. I think, in general, that's what I would say right now.

Delbar: Thank you so much for joining us, Ted. 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.

[01:00:45] [END OF AUDIO]