

Webinar Title: Clinical Documentation in the Age of AI

Date/Time: Friday, July 11, 2025 at 1pm ET

Moderator: Hello, and welcome to today's webinar, Clinical Documentation in the Age of AI. 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 navigating clinical documentation in the age of AI.

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, 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 health care. Welcome, Ted.

Ted Faneuff: Thank you very much for that awesome introduction. Hi, everyone. As they said, my name is Ted Faneuff. It's a real pleasure to be here with you today. I'm genuinely looking forward to this conversation. It's an important one, a really important conversation about how to navigate this digital shift in clinical documentation, especially concerning the ethical compliance and practical considerations with AI. Like they said, I'm a licensed clinical social worker. I've been doing this for a while. In my career, I've worn a lot of hats. I've been a frontline provider in community mental health and private practice.

I've managed a privacy program for a large health care system. For the past several years, I've led clinical operations and large teams of therapists and the mental health tech space across the world, and currently serving as head of clinical operations, as they said, at Upheal. My hope in speaking to that is, my experience is really across a lot of diverse roles, and it's really shaped my perspective on the challenges and the opportunities that we're going to be discussing today.

Next slide. Our roadmap for today, for our time together. We're going to focus on three key areas. First, I want to set the stage by exploring the challenge and the tool. We're going to look at the documentation burden that we all face, and really setting that stage for why this is such a pressing issue in our field, and get a clear, simple understanding of what AI is and isn't. Next, we'll dive into one of the most important parts of the conversation, which is the core ethical pillars, covering confidentiality, informed consent, and professional competence with respect to AI.

Then, finally, we'll move into some practical application considerations, focusing on some best practices for compliant and safe use of AI in your practice. Then we'll wrap up with some key takeaways and hopefully have some time for your questions. My hope is that you leave today not only informed but empowered to critically engage with AI in your work, whether you're already using it or you're just starting to explore those possibilities.

Next slide. Before we go dive into the details of AI and mental health care, I want to tell you a little story. I want to take you back to December of 2021, which was a moment that completely changed how I think about clinical documentation and technology's role in our field. It was a Saturday. I was sitting at my kitchen table, stacks of notes around me, trying to catch up on documentation like I always did. My kids were playing in the living room right in front of me, laughing, absorbed in their world, doing their thing. I really wasn't absorbed in mine. I wasn't even close to being present.

Instead, as an aspect of just the work that we do, I was drowning in that paperwork. I was stressed, exhausted, and honestly questioning a lot of things at that point. At that point, I'd been in private practice full-time since 2016, and from day one, documentation was always my Achilles heel. I really love doing clinical work with the clients, but when it came to writing my notes, I just felt paralyzed. I constantly worried about audits, about missing key details, whether I was capturing the right things. Even though I felt like I was a good clinician, I always felt like an imposter when it came to and putting it all down on paper.

Like a lot of us do when we feel overwhelmed, I procrastinated a lot, put off my notes for as long as possible, but really, that only made it worse because when I finally sat down to do them, I couldn't always remember everything clearly. That day at the table, it all hit me at once. I had spent hours catching up on documentation and was always still behind on notes and billing, which meant financial stress on top of everything else. I just remember putting my head in my hands and just sitting at that table, and I cried. I was thinking, I really can't do this anymore.

I seriously considered, at that point, shutting down my practice altogether. Again, not because I didn't love the work that I did, but because that admin side of it felt too much of a burden to carry alone. Interestingly, interesting timing, something unexpected happened. Around that time, I came across information about a beta test for a new AI product that could help generate clinical notes. Like a lot of people who hear this for the first time, I was deeply skeptical at first as well, so really skeptical. Not because, again, I didn't know a lot about it. I didn't know anything about it, but I was also desperate.

I reached out and I started testing it. Honestly, that's when everything changed for me. For the first time, I truly felt relief. That AI gave me a structured starting point for my notes, which really allowed me to be more present with my clients instead of mentally drafting them in my head. My anxiety about documentation really decreased. I finally felt confident I could focus on my clinical work without spending my weekends drowning in unfinished notes.

Next slide. I don't think my story is unique, though. As a clinician who's worked in both direct practice and now in the tech space supporting therapists, I've really seen firsthand how the demands of documentation can drain our time, our energy, and sometimes even the joy from the vital work that we do supporting our clients. Many therapists feel overwhelmed, not just by the notes themselves, but by that consistent balancing act between ethical pressures to be accurate and compliant and the need to be fully present with our clients. Now, on top of that, AI documentation adds another layer to that challenge.

The capabilities of these tools are evolving at a breathtaking pace, often moving a lot faster than our professional regulations and clinical or ethical guidelines can keep up with. This creates a gap where we, as clinicians, have to rely on the core ethical principles to navigate this new territory responsibly. That brings us really to a central theme for our discussion today, which is weighing the problem of documentation burden against the promise of AI assistance. While AI offers some very real potential to ease this burden, the foundational literature for this talk confirms the other side of the scale as well.

There's been overviews and ethical reviews ongoing done by researchers like those highlighted in the references at the end of my presentation, but they consistently highlight that AI tools simultaneously raise significant concerns around accuracy, data privacy, and ethical implementation. Our goal today is to explore that full picture, the problem and the promise, to find an ethical and responsible balance that's grounded in APA's guiding principles.

Next slide, please. Let's start by defining AI. In simple terms, artificial intelligence is not science fiction. It's not yet the *Terminator* or *WALL-E* or anything of that nature. What it really is is technology that has this remarkable ability to learn patterns from data and then to make predictions or decisions. Already in your life, you encounter this technology with tools like Netflix, which makes recommendations off of viewing habits you have, or Google Maps, which suggests the quickest route to your destination. In both of these cases, the technology works by identifying statistical patterns in vast amounts of data to make those predictions.

In our work as therapists, we're already familiar with some forms of AI. You can think of things like voice-to-text dictation software, or basic transcription services that turn audio into raw text, or even the spellcheck and grammar tools in our word processors. Those tools are really helpful for clerical tasks, but they primarily just capture words as they're spoken or written. The technology that we're focusing on today, and just in general in mental health, it's a little different. It's called generative AI, and it's significantly more advanced. Instead of just transcribing, generative AI, especially when designed for clinical applications, can help in a lot of very practical ways.

How does it actually work? To generate a progress note, the AI needs information about the session. This can be done in a couple of ways. Some platforms allow the AI to listen to a session recording, create a transcript, and then the generative AI completes that task of creating the progress note based on specialized instructions that are provided to that large language model. Other platforms, or combined with the platforms that allow it for the recording, also allow you to use AI by inputting session details or brief notes into the system. Then the AI can expand that into a full progress note without any audio recording.

In many of the platforms, all of the platforms usually have the option to delete the transcripts after use. In most cases, the audio of that session is never kept without strict permission from both the client and the therapist. For example, that process really allows the AI to analyze a conversation and create that structured draft note almost immediately after a session, which not only saves you considerable time but also significantly reduces documentation lag, helps you to minimize that need for late-night charting. Beyond just speed, some systems can really improve the quality of that draft by highlighting if critical information like a link back to treatment goals appears to be missing.

We've actually just launched some of that technology recently to an Upheal. On top of that, it also can assist in enhancing billing compliance by helping to align the note content with medical necessity standards to meet payer requirements better. However, and this is always a very critical point to keep in mind, AI truly does not understand therapy in a way a human clinician does. It operates by recognizing language patterns, not by grasping that clinical significance. While it can summarize key points in a conversation, it might easily miss important client nuances, overlook those subtle nonverbal cues because it can't see what's going on, or misinterpret the emotional tone of a particularly difficult conversation.

Therefore, it's really essential to view AI's role obviously not as the therapist, but as a sophisticated digital assistant, a powerful tool for efficiency that can help you draft and structure your notes, but it always needs your diligent oversight. You'll hear me say this quite a bit throughout this presentation, but every AI-generated draft requires a clinician's careful review, editing, and input to make sure that that final note is clinically accurate, relevant, and ethically compliant.

Next slide, please. Let's get to the heart of it. The ethical landscape of using AI. The good news is we're not exactly starting from scratch. It's still somewhat of a wild west in terms of regulations and ethics, but we have some guidelines to go off of. The APA ethics code actually already gives a strong foundation. The recent AI tool guides that APA has put out also gives practical roadmaps. I find it helpful in this conversation to break all this down into three core pillars that apply directly to your practice, the first pillar being competence and due diligence.

This one is all about doing your homework even before you ever sign up for a free trial of one of these products. It goes right back to that core ethical duty of professional responsibility and managing competence with new technologies. The APA tool guide gave a great checklist, which means basically to investigate the vendor and really dig into their terms of service. At the end of the day, you're responsible for the tools you bring into your practice. Also, let's talk about protecting

client welfare, which combines two of the most important duties for a psychologist or a mental health professional, which is confidentiality and informed consent.

When a third-party AI tool is involved, your responsibility requires a deeper level of due diligence, which we'll cover in detail on the next slide. That diligence is really critical because, as we've seen in some research studies, one in particular, 60% of Americans would be uncomfortable with their provider relying AI in their own healthcare. That was a study done two years ago. I think that although that public hesitation, it emphasizes how vital it is for mental health professionals to have a robust and transparent informed consent discussion with their clients. I think that we've also made some headway in the last two years about the public's comfort in accepting AI into their practice.

What does that conversation look like? It goes really beyond our traditional explanations. It means you have to clearly explain the role of AI as a tool for organizing notes and not making decisions. It involves openly addressing data privacy and security, even more so than we already do right now in our practice, including where's data stored, how it's used to train AI, if that's applicable. It critically means honoring your client choice and autonomy by giving them a clear option to opt out if they're uncomfortable with that. That kind of transparency is really an ethical obligation that we need to keep in mind.

The third and probably most important pillar, which is non-negotiable principle, which is really your professional oversight as a clinician. You can think of this as you're always in the driver's seat. That comes from our fundamental duty to do no harm and to take responsibility for the tools that we use. The ethics code makes it pretty clear that psychologists, again, are responsible for the appropriate application of any service they use, including automated ones. We'll talk about that a little bit further, too. Remember that the AI is your tool, not your replacement. You're always responsible for that final note.

Next slide. We've established our ethical duty to protect client confidentiality. What does that look like in today's digital landscape, where really our client data is no longer stored in filing cabinets? With the widespread adoption of electronic health records, cloud storage, and now the sophisticated AI-powered tools, our data is often processed and stored by third-party vendors. That shift, that fundamental shift, really introduces new layers of risk and raises some pretty critical concerns for us as diligent clinicians. To protect our clients, we have to be proactive in this. That starts with understanding the major legal frameworks.

Obviously, in the US, HIPAA establishes pretty strict national standards to protect sensitive patient health information. Under HIPAA, any third-party tool that we use, and this obviously definitely includes AI documentation tools, must also comply with those stringent security standards. In Europe and other parts of the world, the GDPR, or the General Data Protection Regulation, governs how their personal data is used and gives clients more extensive rights, including the right to request access to their records on demand, and in some cases, to have their client data deleted entirely.

I bring up GDPR only because it is a much more stringent privacy protocol than what we currently have in the US. It's important to note because these sort of things could

continue to enhance or change our HIPAA regulations eventually. This brings us to a really important distinction: the difference between a company that claims compliance and the reality of true robust security. Many companies out there will state that they're HIPAA compliant. That term can be a little bit misleading. In fact, anyone, private practitioners included, anyone can claim to be HIPAA compliant simply by doing their own risk assessment audit.

You can do it and say, "Yes, I check all the boxes. I'm HIPAA compliant." It's really simply not enough for a company to state that it follows HIPAA regulations when it comes to things like AI. As responsible mental health professionals, you have to take on that role of diligently verifying whether that AI tool we're considering truly protects client information in a meaningful way. How do we verify that? Some key things to look for when you roll up your sleeves and dig into the details, you first have to investigate their technical security. Encryption is non-negotiable.

Just like EHRs, the tool has to encrypt client data, both when it's stored, what we call at rest, and when it's being transmitted or in transit. Ideally, this should be done with a high encryption standard like AES-256, which, again, these frameworks are often guided by bodies like the National Institute of Standards and Technology or NIST. You also have to look for strict access controls. Is your account protected with multi-factor authentication as well? Beyond their claims, you have to look for independent third-party audits.

A certification like SOC 2 Compliance is a much more reliable benchmark of a company's security posture because that actually does require a formal, independent review of the policies and processes of a company. You want to make sure that they do those, like a platform, actually does those security audits regularly through a process like SOC. When you're looking at these particular platforms, be sure to ask for proof of third-party validation. At Upheal, for example, we use a platform called DREDA to continuously verify, monitor our security, privacy protocols, and we make all that information transparently available to users at any time.

You want to make sure that whatever tools you're looking at, that this is also available to you. Next, you need to become a detective and really scrutinize privacy policies, especially around data sharing. A lot of things can be buried in legalese. You need to find out if a company sells or shares client data, and even if it's de-identified. Some companies might use that data to improve their own AI models, while others might state in their terms that they sell aggregated data for purposes like marketing or research. That's really crucial information that I say can often be buried in the terms of service.

Sometimes they'll make claims such as, "We don't sell your data," but what they consider your data is different once it's de-identified and aggregated. They consider it their data. You have to make sure that your data is your data, and it maintains that throughout, and they can't sell your data for whatever purpose. Finally, you really have to investigate data ownership and retention. If you decide to stop using an AI tool, what happens to your client records? Can you permanently and completely delete them? A company's policy really has to explicitly state that you, the clinician, retain that full ownership of the data and can request that it's verifiably removed from systems. Any vagueness here should really be considered a major red flag.

Next slide, please. After all the discussion about technology and compliance, let's bring our focus back to the absolute core of what we do as mental health professionals. The work obviously is, and always should be, fundamentally rooted in human connection. This really aligns with that primary ethical duty to do good and to avoid harm, which we know is achieved through a strong therapeutic alliance. However, as we integrate AI, we potentially can face that subtle, but significant risk, which we'll call automation drift. When I was doing some research into this, I was looking at user stories and things about when people were just starting to implement AI into their practice.

You can consider the story of a psychiatrist, which was reported in a Business Insider article, who realized that because she knew an AI was summarizing her session, she was becoming less engaged in the session. She had to consciously fight that urge to drift and remind herself to stay fully present with her clients. That highlights the risk of letting technology mediate the therapeutic encounter. It shows that even when a tool is efficient, it cannot replicate the essential human act, though, of being present, which is the core of a therapist's role.

That brings us to the true goal of using AI ethically in a practice. It's not to automate the therapy, but to automate the burden. By handling the administrative load, AI really can free a psychologist's precious time and mental energy to reinvest back where it truly matters, and deep listening and strengthening the human connection with the client. Obviously, just some things to put into actual practice. It comes down to being incredibly intentional with your presence. When AI is assisting with note-taking, it becomes even more important to stay fully engaged.

It could mean consciously minimizing screen time during your sessions, making that consistent eye contact, actively tuning into both the verbal and the nonverbal cues, simple practical things like silencing notifications, or scheduling just a few dedicated minutes after the session to review the AI draft rather than during. Those things can make a significant difference in maintaining that connection and ensuring that the technology serves the alliance and not at the expense of the therapy.

A story that I always like to come back to highlight this. A clinician who's with us shared with us that they had a client who was making comments related to suicidal ideation in the session. In those moments, the focus quickly shifts to safety planning. Naturally, that's where we go. There's always that concern in the back of our mind of ensuring that you remember all of the details for the documentation after the fact. Using AI in this clinician's practice freed them up to stay present with that client, knowing that the documentation and the specifics would be there when they were done.

In this case, they elected to have the transcript stored from the session, so they knew it would be there. As such, that client did feel more tended to and seen, and thus the potential for that relationship to be strengthened was enabled by the use of AI. I would argue that although there's a potential for that automation drift that we talked about and the loss of presence, there's an equal opportunity that enhances the therapist's ability to stay present if we're mindful and intentional about it.

Next slide, please. Shaping the future of ethical AI, though, isn't only about how we use the tools. It's fundamentally about how we talk about them with our clients. I

often say to people that clients' comfort with adopting this technology rests on your ability to talk about it comfortably with them. That leads us to a conversation about rethinking and adapting our informed consent processes for this new era. As AI does become more embedded in our work when we make that decision, our absolute ethical obligation is to ensure that clients are fully aware of how that information is being handled.

That raises critical questions for us to consider. Do our clients have the inherent right to know when AI is assisting? Should they have the option to decline? How do we ensure that the conversation is meaningful and not just another form to sign? The ethical need really is underscored by that major shift in our practice of that increased client access to their own notes. That's really been a shift over the last couple of decades. Regulations like the 21st Century Cures Act have really changed the landscape of this, forcing us as mental health professionals to ask a critical question, which is "Who are the progress notes truly for?"

If we know that the client's going to read them, our entire approach has to evolve. That's where AI introduces a very real risk. Like we've discussed, while an AI platform can efficiently list topics, it can't capture the spirit of the session, although it continues to improve in this area daily. It doesn't really understand the subtle emotional nuance. You can imagine if your client is reading a note about a deeply personal session that comes across as sterile and cold through the use of AI, instead of strengthening that therapeutic alliance, that note could actually become a point of rupture.

Knowing that risk, it's important to get pretty practical. When a client is requesting their notes, transparency is your best friend. It's really a chance to reiterate the AI's role. It's something you should already be discussing at the beginning of treatment, but throughout if they question and they want to see their notes, you could say something like, "Recall that I use a specialized tool that helps me structure my notes efficiently, but I want to assure you that I personally review and edit everything to ensure it accurately reflects our conversation."

Or something that I've used, while an AI tool assists me in drafting the initial summary of the session, everything you read in your note has been personally reviewed and finalized by me. It's my note reflecting my understanding of our work. If a client reads their note and still disagrees with something, that's a crucial moment for collaborative care and conversation. Informed consent never ends. A helpful approach is a three-step process. First and most importantly, validate the feelings. Thank you for bringing this to me. It's really important that you feel understood and seen.

Second, gently clarify the purpose of the notes. Progress notes aim to summarize key themes, rather than being a verbatim transcript, but your perspective on what's really key is important. Third, you can offer to collaborate. Let's discuss how we can incorporate your perspective into this note, or let's discuss again, do we want AI to continue to be present in helping us work the notes? We don't have to do that. Just because they admit consent to one-time use of it, or they consent to a point that they're uncomfortable with it, they always have the right to withdraw that.

That kind of approach really reinforces that transparency. It builds trust and ensures that AI remains that tool that enhances our work, while we remain the dedicated stewards of our client stories. It's also worth noting that AI platforms are getting more enhanced in addressing personalization here, too. A lot of them offering abilities like smart editing to enhance the tone of the note with the client in mind.

Next slide, please. Now that we're thinking about our notes as client-facing documents, we have to talk about really the craft of creating a compliant and defensible record. This starts with the core components. Every note has to include the essentials like client identifiers, session details, and the clinical information required to demonstrate that medical necessity. A great note is more than just a collection of facts. It tells a coherent story. That's where that concept of the golden thread comes in. It's a really critical principle for compliance, ensuring that everything connects logically from beginning to end.

The thread starts with your initial assessment and diagnosis. This is actually the first place where AI can provide some significant support. The intake and diagnostic assessment process is often where we're inundated with the most information all at once. An AI tool can create a detailed transcript of that intake interview, and then generate a structured biopsychosocial summary, pulling out those key details like the presenting problem, family history, and it can even identify and flag key themes or statements that might align with certain diagnostic criteria for your review.

Again, to be clear, the AI should not ever be making a diagnosis. That's the exclusive role of a psychologist. However, by handling that initial organization of the note it really allows you to focus more on the rapport building and the clinical observation during the assessment itself. From that strong start, the golden thread then has to flow directly into the treatment plan. Again, that should have measurable goals based on that assessment. Then again, another area where AI is helping here is to take things like a verified diagnosis and the intake, and creating a treatment plan that's specially tailored with a therapeutic modality in mind.

That actually makes it much more tailored to the client's presentation. Finally, every single progress note has to explicitly connect back to the treatment plan goals, documenting those specific interventions, the use, and what the client's response was. A broken thread is really one of the biggest red flags in an audit because it really does fail to demonstrate medical necessity for the services provided. Even here, in the progress note, AI can act as a helpful assistant, and a well-designed tool can automatically pull all of the client's established treatment plan and goals directly into the note you're writing.

Really, it can help you document interventions that align with their goals or even analyze documentation over time to highlight progress. However, even with these advanced capabilities, it's really critical, again, you'll hear me repeat this, but it's really critical to remember the role of AI is only an assistant. It can surface information and suggest connections, but you, again, as the mental health professional, have to always provide the core clinical content and use your judgment to ensure that that golden thread is not just present, but clear, logical, and accurately reflects that client's unique journey.

Next slide. Following that golden thread also means respecting documentation boundaries. This can feel a bit like a tightrope walk between two common compliance issues, which is over-documentation and under-documentation. Let's talk about the don'ts. It's critical that psychologists don't over-document. Progress notes, they're meant to be, I should say, clinically relevant summaries, not verbatim transcripts, like we said previously. Common mistakes include adding too much extraneous detail, or, and this is a big one, including your own subjective opinions that aren't really clinically grounded.

Your notes really should always strive to be objective and focused on the client's reported experience and the observable progress that you have. On the flip side of that, under-documentation is just as risky. You have to obviously justify the need for your services. One of the most common audit triggers is when a note fails to justify the service that was built. Auditors want to see clear evidence of active treatment, and they're increasingly focused on whether the session length is justified by that clinical richness of the note. That all comes down to the final and really most important point on the slide, which is accurate and thorough documentation is your best defense in the case of a complaint or a legal issue.

You have to always remember the rule in the eyes of an auditor or in a legal review, that old saying holds true: if you didn't document it, it effectively didn't happen. One of the things again here that I like to point out in AI's usefulness here is even after you draft your note and you put your clinical view on it and you vet it, you can use things like AI within a platform to ask it to see if it's actually too much information. Maybe you want to actually reduce the information, but keep the core clinical content that you have. You can utilize AI in that respect, which is another area that's very helpful to avoid things like over-documenting or under-documenting.

Next slide, please. As we prepare the notes that are both ethically sound and client accessible and all connected with a golden thread, we do have to talk about three specific risks that mental health professionals really have to actively manage when they're using AI. First, let's talk about AI bias. If an AI model is trained on non-diverse datasets, the output obviously has an opportunity to amplify societal biases, potentially misattributing symptoms. That also can lead to a loss of client nuance where the notes really feel generic and miss that subtle emotional tone or culture expression of distress that are really so central to the therapeutic work.

Mental health professionals, that manual review that we talk about is really essential to ensure that the note is not only equitable, but also reflects the unique, nuanced experience of the client. The second major risk is the phenomenon of AI hallucinations. It's important to remember that AI models, again, are not clinicians. They operate on complex statistical patterns. It's not true understanding or factual recall. Because of this, they can sometimes produce content that seems plausible, but is in fact entirely fabricated.

For example, if a client mentions feeling fatigued, an AI might, based on whatever patterns in its training data, incorrectly suggest or imply a depression diagnosis in the draft note, even if you, as the clinician, never made that assessment. Those kinds of errors really introduce significant clinical and legal risk, making your review for accuracy absolutely essential. Again, that third risk to watch out for is like we

talked about in the previous slide, that AI-prompted over-documentation. Sometimes, an AI tool can generate a lot more detail than is really clinically necessary.

Our ethical duty is really to create notes that are useful and relevant, which means that you have to act as the editor. Ultimately, the key to mitigating these risks is your vigilant oversight and choosing the right tool. To feel confident and comfortable, it's really important to use AI tools that have been thoughtfully curated and specifically designed or fine-tuned for clinical best practice, so ideally, with input from experienced clinicians on their leadership team or in their product team. That really helps to ensure that the AI is a genuine support, never a replacement for your irreplaceable clinical judgment and expertise.

Next slide, please. After covering all that ethical principles and the specific risks, let's pull it together again into a final practical three-step workflow. This is really the framework for how psychologists or mental health professionals can ethically engage with AI from start to finish. Again, step one is again about your due diligence before use. It's more than a security checklist. It's about assessing a vendor's alignment with your professional values. The APA guide suggests looking again for clinicians and leadership positions for a reason, as it's an indicator that the company may better understand that primary ethical duty to your clients.

After you've done that vetting and documented your view, you can then choose a responsible partner for your practice. At Upheal, I help lead clinical operations, but we also have several clinicians on the team that do clinical prompt engineering, or various roles throughout the company. Clinicians are well-represented in Upheal. Once you've actually chosen a partner for your practice, the next step would be responsible implementation. That's really best understood through the lens of the APA's ethics code standard on delegation of work. Although the focus of that standard is really more on the delegation of work to others, when you use AI to draft a note, you're still delegating a professional task.

That code really requires that you authorize only those responsibilities that the tool can be expected to perform competently, and that you have to verify that competence. That's what your manual review is. It's you, the psychologist, ensuring that the delegated task was done correctly to guard against any bias and error, and ultimately to uphold your authority as required by your ethical code. Then step three is really that ongoing professional competence. This is especially critical because, as we've established, AI technology is evolving much faster than formal regulations can keep up with.

That's where our codes of ethics come in to offer really direct guidance for situations like this. One standard specifically addresses emerging areas in which generally recognized standards for preparatory training don't yet exist. This very much pertains to new tech like AI. It states that in these areas, psychologists must nevertheless take reasonable steps to ensure the competence of their work and to protect clients from harm. Your ongoing learning here is that reasonable step. It's not just about seeking initial training.

The APA's AI tool guide makes it clear that competence includes periodically re-reviewing privacy policies in terms of service for the tools that you already use, as they can and do change, and then continuing to take those trainings, such as this

one, to continue updating your knowledge on the topic. That vigilance is our primary way to ethically and competently navigate that gap between technology's rapid evolution and the slower pace of formal regulation.

Next slide. As we explored today, AI is really rapidly becoming a part of the mental health landscape. It's already here. It brings both incredible promise and serious ethical considerations. If you remember nothing else from our time together, I hope that you leave with these key takeaways. First, as a psychologist or mental health professional, you have to remain fully firmly in control. AI can be a powerful tool to support your work, but it should never ever replace your clinical judgment or that essential human nuance that's so vital to effective therapy.

Second, ethical communication is not optional, it's foundational. It's especially true when it comes to having those transparent informed consent conversations with your clients about how their data is used. That open communication is really critical for maintaining the trust and protecting the therapeutic alliance, and to support this, some platforms offer built-in features to help manage the process. For example, at Upheal, we provide several options to help you collect and discuss informed consent with your clients.

You can use a prearranged statement to add to your existing paperwork, or you can generate a PDF document from the platform that clients sign, and then you can add to their chart, or you can have the platform collect consent directly through email. I also am going to be providing some handouts to you today that goes over consent language that you can use just in conversation with your clients about it as well. Third, as professionals, we really have a role to play in shaping the future of this technology. This really brings me to a final call to action.

Whether or not you plan to use AI immediately, I really urge you to get involved in the conversation. Ask questions, challenge assumptions, advocate for AI that's transparent, inclusive, and always client-centered. Now, some companies are examining how to use AI to replace aspects of therapy. We know this. However, I hope I've demonstrated that AI has profound limitations that prevent it from fully replacing the need for mental health professionals. Still, that doesn't mean that we shouldn't continue to advocate for that responsible development in this area.

It still may always need clinician oversight and training, but that doesn't mean that we do not need to continue to resist that technology's advancement in these areas. The future of mental health really is being built right now, and we need to ensure that it reflects our deepest professional values and supports clients. Next slide. That actually is the references slide for this. If you have any questions about some of the material that I had in the presentation today, you can go directly to these links for these articles from the references. That's all I had today for you in my presentation. I guess we'll turn it over to audience Q&A.

Moderator: Wonderful. Thank you, Ted. Now we're moving into the live Q&A. Thank you to everyone who submitted questions during registration and to those adding them in the Q&A box today. Ted will do his best to answer as many questions as possible, though he may not get to all of them. Your questions are appreciated. Let's dive in. We have one question asking, "Are there any climate impacts of using AI for our work as psychologists?"

Ted: That's really an important question, and it's gaining a lot of steam in the industry. The short answer is yes. AI tools, just like any computer tools, they do have environmental impact. It really comes down to the massive amount of energy that's needed to power the data centers that train and run the complex models. Now, there is a difference between training in terms of consumption, between training the AI, which takes a huge amount of power up front, and the energy needed for that daily use, but it all does add up. On our end at Upheal, and I think, there's other companies that can maybe attest to this, we really are taking this pretty seriously.

One of the ways that we tackle this is by partnering with cloud providers who are committed to using renewable energy and by supporting carbon offset initiatives. What's really interesting that I like to point out to mental health professionals is the crucial role that clinicians like us play in making AI more environmentally sustainable. You can think of it this way. When an AI model has to guess, it wastes energy. When you, the expert, provide real-time feedback right in the product, or when you have clinician-led teams, you're teaching the AI to be more precise from those interactions.

Our internal clinical team also helps refine the AI prompts so that it doesn't have to infer as much. The result is a more efficient AI that avoids the unnecessary energy-intensive retraining. Your clinical expertise not only improves the relevance of the nodes but it also helps reduce the environmental footprint. In short, your involvement when you're using these tools, especially ones that have that human-in-the-loop and more direct clinical involvement in training without using data, it helps to make the AI smarter and more sustainable.

Moderator: Wonderful. Thank you. We have another question that asks, "What is the best way to identify within clinical notes if AI was used to ensure ethical and legal compliance?"

Ted: That's a really great question. Again, I think that goes back to the heart of transparency. There's not a universal mandate yet as to whether or not it's required. One emerging, though, best practice is for the clinician to be proactive. Instead of having someone reviewing the chart, trying to guess if AI was used, the most ethically sound approach is to add a clear statement to your own notes.

As we talked about, we have to be stewards of the client story. You can add a simple footer at the end of your note, which is something like, "This note was drafted with AI assistance. It's been reviewed, edited, and verified for clinical accuracy by the undersigned psychologist." That creates that clear record. That transparency is really your best practice for both ethical and legal compliance. Within Upheal and other competing products, we even allow for those static statements to be created and attached to each note, should you choose to use it. Always, anytime there's a legal question that comes up, I will just also defer to making sure that you're getting guidance from your legal team.

Moderator: We have another question that asks, "Is the AI transcript a part of the clinical record that clients have access to?"

Ted: Yes. That's a great, again, critical, legal, practical question. Generally, the official clinical record consists of the final signed progress note that you verified and

placed in the client's chart in your HR. The raw transcript or the initial unedited AI draft would typically be considered underlying work material, similar to your own personal process notes. However, laws like we talked about that 21st Century Cures Act are pushing for greater transparency, and the legal interpretations are going to vary here. The most risk-averse and transparent approach is to operate as if any data related to a session could potentially be requested.

That's why it's important to use HIPAA-compliant platforms where the underlying data is securely stored and protected. I would confirm again, though, with your legal representation to make sure. The one thing I will say, too, is on our platform, we allow the ability for the transcripts to be deleted, either automatically after the session is transcribed, or you can set a predetermined amount of time, like after seven days or something of that nature.

Moderator: Brilliant. We have another question asking, "Why would a client consent to our use of AI? It seems there's no benefit to the client. Why would they consent?"

Ted: I think that's a really fair question, and it's one that we should all be prepared to answer. I think the key is to frame the benefit in terms of what does matter most to the client, which is the quality of their care. Obviously, the primary benefit to the client is that by using an AI tool to handle the administrative burden, you, as a psychologist, can be more fully present with them in the session. In the conversation, I wanted to-- When we were talking about presentness and automation drift, at the beginning, my goal was to juxtapose that there is an opportunity for drift, but more often we're seeing that it actually helps someone maintain presentness.

You can focus on listening and engaging rather than mentally drafting the notes while you're talking. When we explain it, we can say, "Again, I use this tool because it helps me focus 100% of my attention on our conversation, ensuring that you get that best possible attentiveness and care. Certainly, the primary benefit is to the clinician, but I wouldn't say there's no benefit to the client. I know that there's some proponents of AI in session saying that there's no tangible benefit to the client, and as such, we shouldn't really want to use the technology to optimize our well-being.

If we're honest, medical doctors have been using similar technology for years, and many are also now embracing AI in their own practices. There's always a primary and a secondary gain to such things, really, and I see this really as an access issue. If we can help therapists stay in the field and feel less burned out, more clinically present from the documentation burden, I see this collectively as a major benefit to the clients we serve now, or we may serve in the future.

Moderator: Wonderful. We have another question that asks, "Is it okay to use ChatGPT to help create progress notes as long as there is no identifying info?"

Ted: We get this question a lot. I think it brings up some pretty important points from our discussion today on data privacy. While it might seem safer if you de-identify the information, using a general-purpose tool like ChatGPT carries pretty significant risks. First, true de-identification is very difficult to achieve. Subtle details that you put in there can really potentially still identify the client.

Second, and I think very importantly too, the terms of service for many general AI tools allow that company to use your input to train their models. You lose control over that data then. The APA guidance and more ethical codes guidance points you really towards using tools that are specifically designed for healthcare, are HIPAA compliant, and that you again, have that clear ownership over the data that you've thoroughly vetted. Using a public tool for PHI, even if you try to de-identify it, isn't really the best practice. I would say, again, the outputs that you're going to get are going to vary greatly if they're not curated by a mental health professional on the backend.

Moderator: We have another question that asks, "As a clinician with dyslexia, I see AI reducing documentation barriers. Is its role in equity part of ethical discussions?"

Ted: I really thank this person for bringing up a very important and often overlooked point. Absolutely. While we focused really on some of the risks today and also the promises, the role of AI in promoting equity and accessibility within the profession is really a huge part of the ethical conversation. For clinicians with dyslexia, other learning differences, or physical disabilities, AI documentation tools can really be powerful accessibility resources, ultimately reducing barriers to completing a core professional task of documentation. It really does align perfectly with the goal of using technology to support clinicians so they can focus on their strengths.

It's really a vital part of the ongoing discussion about how to build a more inclusive and sustainable profession. I know for me, the use of AI has been immensely beneficial in my practice, having ADHD. When I was going through all the stuff I talked about in my starting story, that was actually the result of a lot of undiagnosed ADHD. I knew I had it for years; it was just never diagnosed.

Utilizing AI in my practice has served a very real accessibility issue for me. I very much wish it was around in 2018 when I broke my dominant arm. I couldn't write or type for weeks while I healed. I really wish it was available. I think that something to be careful of, and this is a call and urge to anybody who's maybe resistant and still after this presentation, or just in general, just because it might not be for you, AI does serve a very real, accessible issue for a lot of clinicians. I would encourage you to consider that in your responses when you're talking about AI.

Moderator: Wonderful. We have another question that asks, "How good is AI at processing accents or talking together, soft voices?"

Ted: Surprisingly, really, really, really well. I have found that this is an area that continues to evolve quite a lot in the ability to transcribe based off of very soft or accented voices. Even outside of accents, a lot of the programs out there, they will allow for any language to be able to be spoken in the session and then have it transcribed into English or et cetera. I've found that this is one area where it continues to improve.

The one thing I would say is that it's all about placement. If you're in person, you can utilize apps or different things that are still HIPAA compliant, still have the same security protections, and you can use them, just making sure that they're close enough, so that they can capture both voices. It does a really good job of collecting information, even if people are soft-spoken.

Moderator: Brilliant. Looks like we have time for one last question. Someone asks, "Which AI tools are most reliable and useful for clinicians?"

Ted: Again, I'll go back to the fact that obviously, I work for an organization called Upheal. I'm going to put out there that I truly do believe that Upheal is probably one of the most secure private platforms. We hire a team of clinicians, not only in the output of the notes themselves, but also to support the clinicians. We have people on the account executive team, we have people in our clinical prompt engineering, myself leading, and we have a clinician leading our clinical product.

I know that Upheal does a really great job of ensuring that clinician voice is heard and respected. I do know that there's other platforms out there that are very similar. I would say that if you're doing any review, you really want to make sure that you're looking for companies that don't just talk about a clinical presence on their team but truly demonstrate it, because those are going to be the ones that are the most effective, compliant notes and really tailored to you as a mental health professional.

Moderator: 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.

[00:57:44] [END OF AUDIO]