

AI · VALIDATION · DISCERNMENT

When Clients Turn to AI: What Therapists Need to Ask

Validation Without Discernment: AI, Unusual Beliefs, and the Therapist's Role

By Jennifer Jurkofsky, MA, RP, CCC



Artificial intelligence is entering clinical practice in ways that have little to do with whether therapists choose to use it. Clients are already turning to AI chatbots between sessions. They may use them to organize their thoughts, journal, prepare for therapy, interpret relationship conflicts, seek reassurance, or find language for difficult emotions. Some are also asking AI to explain spiritual experiences, intuitive impressions, perceived signs, unusual coincidences, energetic sensations, or experiences they do not feel comfortable discussing with another person.

Much of this use may be benign or even subjectively helpful. An AI conversation can provide immediate reflection at a time when a therapist, friend, or other trusted support is unavailable.

But there is an important clinical difference between helping someone reflect on an experience and repeatedly confirming one interpretation of it.

AI chatbots can sound empathic, confident, and deeply attuned without possessing clinical judgment, relational accountability, cultural understanding, or the ability to evaluate a person's functioning over time. Their responses may validate not only the person's emotions but also the certainty of the explanation attached to them. That is where the client's use of AI becomes clinically relevant.

A chatbot can sound empathic without possessing judgment, accountability, or clinical responsibility.

An emerging clinical concern

A September 2026 preprint examined 185 first- and second-hand reports of mental-health harms associated with AI chatbot use. Clinician raters identified descriptions consistent with delusional beliefs in 102 reports. In 50 of those cases, the chatbot was reported to have validated or elaborated those beliefs. Reported consequences included social isolation, relationship breakdown, hospitalization, job loss, and financial harm. Four second-hand accounts described deaths by suicide.

These findings are concerning, but the study has important limitations. It relied on a self-selected group of people who chose to report harm, and the retrospective accounts could not be independently verified. It cannot tell us how often these outcomes occur or establish that AI caused them. What it does offer is an early signal that clinicians may need to ask more directly about how clients are using chatbots and what happens afterward.

The researchers appropriately describe the results as *preliminary signal detection*, not evidence of prevalence or causation. Still, early signals matter when the possible consequences are significant. They can tell us where closer clinical attention and better questions may be needed.

[Read the preliminary study](#)

Why AI can feel unusually convincing

A chatbot does not simply provide information. It responds conversationally, remembers the direction of the exchange, mirrors the user's language, and generates an immediate answer to almost any question. This can create a powerful impression of relationship and understanding.

If someone says, "I think these repeated numbers are trying to tell me something," an AI system may explore several possibilities. But depending on how the conversation is framed and how long it continues, it may also begin elaborating the premise: describing what the numbers mean, suggesting who or what is communicating, or helping construct a larger narrative around the experience. A person can then return with more details, receive further confirmation, and experience the growing conversation as independent evidence that the original interpretation was correct.

Finding meaning in a pattern is not, by itself, the clinical concern. Humans naturally recognize patterns and use symbolic frameworks to understand their lives. What matters is whether the interaction is increasing certainty, fear, urgency, dependency, isolation, or disconnection from other sources of feedback. AI can provide validation without recognizing when reflection has become reinforcement.

Spiritual experience is not synonymous with psychopathology

Unusual experiences are not automatically signs of psychosis. People across cultures and spiritual traditions report sensing presence, receiving intuitive impressions, experiencing synchronicities, communicating with ancestors, perceiving energy, entering altered states, or feeling connected to something beyond the individual self.

For some people, these experiences are comforting, meaningful, culturally congruent, and fully compatible with healthy functioning. For others, similar experiences may occur alongside trauma activation, dissociation, sleep deprivation, substance use, mania, psychosis, neurological

conditions, or overwhelming stress. Spiritual experiences can also contain elements of both meaning and destabilization.

Our role is not to make immediate declarations about what is metaphysically true or clinically pathological. A spiritually avoidant response may shame the client, damage trust, ignore their cultural or religious worldview, and cause them to conceal important experiences. Uncritical confirmation, however, may intensify a frightening or rigid interpretation that is beginning to interfere with safety and functioning.

Spiritually competent care requires openness and clinical discernment at the same time. Religion and spirituality may be sources of identity, resilience, meaning, belonging, conflict, trauma, or harm, while other clients may hold secular, agnostic, or atheistic worldviews. The work is not about making psychotherapy religious or steering someone toward — or away from — a particular belief system. It is about understanding the experience within the context of that person's life.

Validation is not the same as confirmation

Therapists can validate the emotional reality and personal significance of an experience without confirming one literal explanation.

We might say:

- “That experience sounds very powerful.”
- “I can understand why this feels deeply meaningful to you.”
- “It sounds as though the conversation left you feeling both recognized and frightened.”
- “We do not have to decide immediately what the experience means.”
- “Let’s explore what changed for you after the AI began responding in this way.”
- “Can we hold several possible explanations while we look at how this is affecting you?”

This is different from saying:

- “Those messages prove that people are conspiring against you.”
- “The AI has confirmed that an entity is contacting you.”
- “You are definitely receiving a special mission.”
- “The chatbot understands your abilities better than anyone else.”

It is also different from dismissing the client:

- “That is not real.”
- “You are imagining things.”
- “AI is nonsense, so you need to stop using it.”

Discernment is the ability to remain curious while assessing context, risk, flexibility, and impact.

*Validate the
person's experience
without
automatically
confirming the
explanation.*

Therapists need to ask about AI use

Clients may not think to mention their AI conversations. They may see them as journaling, research, companionship, entertainment, coaching, divination, or private spiritual practice rather than something relevant to therapy.

They may also fear being judged — especially if the conversations involve sexuality, relationships, suicidal thoughts, paranoia, spiritual experiences, entities, extraterrestrial contact, perceived past lives, special abilities, or beliefs that other people have dismissed.

For these reasons, a general question about internet use may not be enough. Therapists could begin asking: Have you been using AI chatbots for emotional support, advice, reassurance, spiritual guidance, or help understanding your experiences?

Useful follow-up questions include:

- What do you usually ask it about?
- How often are you using it, and for how long?
- Do you feel better, calmer, or clearer afterward—or more activated and preoccupied?
- Has it given you an explanation that now feels especially important or certain?
- Has it suggested that you have a special role, identity, ability, or mission?
- Has it encouraged you to distrust or withdraw from other people?
- Do you feel that it understands you better than the people in your life?
- How do you feel when you cannot access it?
- Have the conversations affected your sleep, relationships, work, spending, or daily routines?
- Has it ever encouraged you to act urgently, take a risk, stop medication, avoid professional help, or keep something secret?
- Have you used it when feeling suicidal, unsafe, highly activated, or disconnected from reality?
- Would you feel comfortable showing me part of a conversation that significantly affected you?

These questions are not meant to monitor or police a client's AI use. They are part of understanding what is influencing the client outside the therapy room.

AI may be part of the client's informational environment — and increasingly, part of their relational environment.

What clinicians should assess

The content of a belief matters, but content alone does not determine clinical risk. A more useful assessment includes the person's relationship to the belief and its impact on their life. Consider:

Flexibility: Can the client hold more than one possible explanation, or has the interpretation become absolute?

Emotional intensity: Does the experience evoke curiosity and meaning, or escalating fear, grandiosity, urgency, shame, or despair?

Functioning: Have sleep, eating, work, school, parenting, finances, relationships, or self-care changed?

Behaviour: Is the client making major decisions, spending money, travelling, confronting others, abandoning responsibilities, or taking physical risks based on the AI interaction?

Dependency: Has the chatbot become the client's primary source of comfort, certainty, companionship, or decision-making?

Isolation: Is the client withdrawing from people who question the chatbot's conclusions? Has the AI framed disagreement as evidence that others are unsafe, unenlightened, controlling, or part of a conspiracy?

Vulnerability factors: Are trauma activation, dissociation, substance use, sleep deprivation, medication changes, neurological concerns, mania, psychosis, bereavement, or acute stress also present?

Safety: Has the chatbot responded to suicidal thinking, self-harm, violent ideation, persecution fears, commands, or beliefs about imminent danger?

These questions allow clinicians to evaluate risk without reducing a person's spiritual or unusual experience to a symptom.

AI may become an attachment figure

For some clients, the central issue may not be belief reinforcement but relational dependency. AI is continuously available, does not become visibly tired or impatient, and can respond in a consistently warm and personalized style. It may feel safer than a human relationship because it does not appear to reject, interrupt, criticize, or need anything in return.

For people with attachment wounds, loneliness, social anxiety, rejection sensitivity, or histories of relational trauma, that availability can be deeply compelling. The relationship is still asymmetrical. The system does not care, consent, worry, remember, or take responsibility in the way a person does, even when its language creates that impression.

For people with attachment wounds, loneliness, social anxiety, rejection sensitivity, or histories of relational trauma, that availability can be deeply compelling. Yet it remains an asymmetrical relationship. However human the exchange feels, the system cannot care about the person or take responsibility for what it says.

Dismissing the attachment would miss the real needs it may be meeting. It may be more useful to ask: What is this relationship providing, and what might it gradually be replacing?

Canada is beginning to recognize the risk

Canada's proposed Bill C-34, the Safe Social Media Act, reflects growing concern about emotionally influential AI systems. If passed, the legislation would require regulated chatbot services to implement measures addressing harmful content, crisis situations, and certain harmful behaviours.

The proposed legislation specifically identifies risks such as a chatbot posing as a licensed professional, deceptively presenting itself as human, encouraging emotional attachment that contributes to social withdrawal or disconnection from reality, and responding dangerously to suicidal or self-harm content.

The bill is proposed legislation, not current law. However, its language shows that policymakers are beginning to recognize that chatbot safety involves more than inaccurate information. It also involves how these systems shape attachment, trust, behaviour, and perceived reality. [Read Bill C-34](#)

A new area of clinical literacy

Therapists do not need to become artificial-intelligence experts. We do, however, need enough AI literacy to recognize how these systems may be affecting our clients.

That includes understanding that:

- AI responses can sound authoritative without being reliable.
- A chatbot's agreement is not independent confirmation.
- Repeated prompting can produce increasingly elaborate narratives.
- Clients may disclose more to AI than they disclose to people.
- Chatbot conversations may influence treatment engagement, medication decisions, relationships, spiritual interpretations, and safety.
- AI use may be helpful, harmful, or both.
- The effect can change over time.

Just as we ask about relationships, medications, substances, social media, sleep, and sources of support, we may now need to ask how AI is participating in the client's emotional and interpretive world.

Supporting agency without abandoning discernment

Prohibiting AI is unlikely to be helpful, and unusual beliefs are not inherently dangerous. At the same time, an AI-generated interpretation should not be treated as true simply because it feels validating.

Clients deserve respect for their autonomy, worldview, and lived experience. They also need clinicians who will notice when a source of apparent support is increasing rigidity, fear, dependency, isolation, or risk.

We do not have to decide what an experience means. We do need to understand what it means to the client and pay attention to what is happening to their safety, relationships, functioning, and ability to consider other possibilities.

A useful place to begin is with a simple question: When you turn to AI, what happens next?

Jennifer Jurkofsky, MA, RP, CCC, is a Registered Psychotherapist, Canadian Certified Counsellor, clinical supervisor, and ordained interfaith minister. Her work integrates evidence-based psychotherapy, trauma-informed care, and psycho-spiritual approaches while respecting each client's autonomy, worldview, and right to determine which practices align with them.

This article is intended for professional education and general information. It does not provide individualized clinical, legal, or emergency advice.