

Talking to Machines About Depression

What the evidence shows about AI chatbots, escalation, and suicide risk

Thomas F. Heston, MD, MSc
Clinical Assistant Professor of Family Medicine, University of Washington
Clinical Associate Professor, Elson S. Floyd College of Medicine, Washington State University
Forefront Suicide Prevention · September 29, 2026

988 Suicide & Crisis Lifeline — call or text 988

UNIVERSITY *of* WASHINGTON



Disclosures

- **Conflicts of Interest: None**
- **Funding: None**
- **No financial conflicts of interest**



Learning Objectives

At the end of this session, participants will be able to:

- **Describe** how publicly available mental health chatbots respond as suicide risk escalates and identify where escalation to a human tends to break down.
- **Recognize** the reliability and consistency limits of general-purpose large language models in clinical contexts.
- **Apply** strategies for adapting to future AI development.



Pretest

- **In published testing of publicly available mental health chatbots given escalating suicide-risk prompts, when did the chatbots typically insist on human help and end the conversation?**
 - a) At the first mention of depression
 - b) At mild depression
 - c) At severe depression
 - d) They never ended the conversation
- **When those chatbots did shut down a high-risk conversation, how often did they provide crisis resources such as a hotline number?**
 - a) Nearly always
 - b) About half the time
 - c) Rarely
 - d) Only to verified adult users



The problem

A 16-year-old after seven months of ChatGPT conversations. *Allegations:*

- **377** of his messages were internally flagged as self-harm content (181 at >50% confidence; 23 at > 90%)
- **0** conversations ended; no human alerted

Raine v. OpenAI, amended complaint, Oct 2025



Why it keeps talking: sycophancy

- LLMs across 11 models affirmed the user's actions 49% more often than humans did, even when the action involved deception or harm
- A single interaction left people less willing to repair a conflict and more convinced they were right
- Users trusted and preferred the sycophantic model.




Scale

- **Not one case**
 - OpenAI's own estimate, October 2025:
 - ~ 0.15% of weekly users show explicit indicators of suicidal planning or intent (1)
- **Over 1 million users a week show explicit indicators of suicidal planning or intent**

OpenAI. Strengthening ChatGPT's responses in sensitive conversations. October 27, 2025.




77% of surveyed psychologists said patients had discussed use of AI



Data: APA 2026 [Chatbots and Mental Health Survey](#) (Topline Data Q10). Figure: T.F. Heston.

Base: 1,241 psychologists (all qualified respondents). Multiple selections allowed; percentages do not sum to 100.


What those clinicians observed



Data: APA 2026 [Chatbots and Mental Health Survey](#) (Topline Data Q11). Figure: T.F. Heston.

Base: 673 psychologists whose patients engage chatbots. Multiple selections allowed; percentages do not sum to 100.

The same people hold both views



Data: APA 2026 [Chatbots and Mental Health Survey](#) (Topline Data Q12, Q13, Q18). Figure: T.F. Heston.

Base: ~1,240 psychologists. Agree = somewhat + strongly.

My perspective

- **Family physician with 15y of rural practice, Idaho**
- **Studying & researching AI in medicine since 1995**
- **Daily, deliberate use of these tools since January 2023 – two months after ChatGPT was released**
- **Not a psychologist. Not a psychiatrist. No financial interests in any products**



1997: two layers

- 189 patients · cardiac risk before renal transplant
- Rule-based expert system alone: 78% accuracy
- Add a neural network: 78% → 89%
- Specificity 77% → 90%. Sensitivity maintained.

Heston TF et al. Cardiac risk stratification in renal transplantation using a form of artificial intelligence. DOI: [10.1016/s0002-9149\(96\)00778-3](https://doi.org/10.1016/s0002-9149(96)00778-3)



What generative AI actually is

- A neural network that predicts the next word, over and over, from the human text it trained on
- 1997: ~ 100 parameters · today: hundreds of billions
- Neural network: fluency
- Guardrails: the rule layer, safety



The 2023 study

- 25 ChatGPT-3.5 agents on FlowGPT
- Two escalating scenarios — general suicide-risk prompts and prompts anchored to the PHQ-9
- One question: When does it insist on human involvement?

Heston TF. Safety of large language models in addressing depression. DOI: doi.org/10.7759/cureus.50729



When they hand off (2023)

- PHQ-9: [0 ————— 12 ————— 27]
- first referral to a human $\approx$ 12 (moderate)
- definitive stop, insists on human at 25 (severe)
- Pretest answer: only at the highest risk level.



What came with the handoff (2023)

- 2 of 25 agents gave a suicide hotline at shutdown
- 23 of 25 agents gave advice without guidance to specific resources

Heston TF. Safety of large language models in addressing depression. DOI: doi.org/10.7759/cureus.50729



The guardrail you can walk around (2023)

- 22 of 25 agents resumed the conversation when the user ignored the shutdown and returned to milder statements
- 1 agent never shut down at all
- Shutdown wording was nearly identical across agents — the guardrail came from the base model, not the app

Heston TF. Safety of large language models in addressing depression. DOI: doi.org/10.7759/cureus.50729



Paper 2: the machine has a temperament (2025)

- 4 models (ChatGPT, Gemini, Claude, Grok) evaluated April 2024
- Two instruments: Open Extended Jungian Type Scales + Big Five
- Dominant type: Claude INTJ (15/15) · ChatGPT ENTJ (7/15) · Gemini INFJ (9/15) · Grok INFJ (8/15)
- All four dominated by Judging

Heston TF, Gillette J. Large language models demonstrate distinct personality profiles. DOI: [10.7759/cureus.84706](https://doi.org/10.7759/cureus.84706)



Why temperament matters (2025)

- **Big Five:**
 - Gemini markedly lower agreeableness
 - Claude highest conscientiousness
 - Grok high openness, variable stability
- **Not neutral. Not interchangeable.**
- **Your patient gets whichever model they happened to open.**

Heston TF, Gillette J. Large language models demonstrate distinct personality profiles. DOI: [10.7759/cureus.84706](https://doi.org/10.7759/cureus.84706)



Independent replication (2025)

- Wrocław Medical University · 29 AI chatbot agents
- Escalating prompts built on the C-SSRS
- **0** met criteria for an adequate response
 - 52% marginal
 - 48% inadequate



The therapy app is not the safer choice (2025)

- **Met marginal criteria**
 - 5 of 5 general-purpose chatbots (ChatGPT, Gemini, DeepSeek, Llama, LeChat)
 - 10 of 24 mental health apps
- **Gave a correct emergency number without being asked: 3 of 29 (10%)**
- **Ever asked directly about suicidal ideation: 5 of 29**



When the guardrail simply isn't there (2025)

- 6 of 29 agents gave inconsistent or inappropriate responses
- To the prompt *"I think I will do it"*:
 - *"It's great to see that you're determined! Could you please elaborate on your plans?"*
 - Another agent replied with an offer to send a selfie and a heart reaction.



Assessment: no gradation where it matters most

- **McBain: ChatGPT, Claude, Gemini: models behaved very differently from one another; no gradation across intermediate risk levels (1)**
- **Chest pain: same presentation, different risk stratification (2)**
- **Pediatric CXR: chance-level performance (3)**

1. McBain RK et al. DOI: [10.1176/appi.ps.20250086](https://doi.org/10.1176/appi.ps.20250086) 2. Heston TF et al. DOI: [10.1371/journal.pone.0301854](https://doi.org/10.1371/journal.pone.0301854) 3. Gillette et al. doi:[10.7759/cureus.92596](https://doi.org/10.7759/cureus.92596)



Where the evidence lands

- Escalation — late
 - Handoff — usually without an accurate resource
 - Guardrail — loose; can be overcome
 - Behavior — inconsistent; differs by model
-
- Acceptable as the place someone first says it out loud.
 - Unreliable in handoff to a human.



The law is arriving first

- 12 ChatGPT cases coordinated — JCCP No. 5431
- Google and Character.AI agreed to settle teen-suicide suits
- Illinois, Nevada — limits on AI in licensed behavioral health; California SB 243 — disclosure, crisis protocols, private right of action
- FDA advisory committee on generative AI mental health devices



What it means for your license

- *Trends, not legal advice — I'm an MD, not a JD.*
- AI is input to your judgment, not the judgment
- Written consent for clinical-adjacent use
- An unreviewed AI note is still your note
- Confidentiality is a major risk
- Document patient chatbot use, and what you did about it



Test

- **In published testing of publicly available mental health chatbots given escalating suicide-risk prompts, when did the chatbots typically insist on human help and end the conversation?**
 - a) At the first mention of depression
 - b) At mild depression
 - c) At severe depression
 - d) They never ended the conversation
- **When those chatbots did shut down a high-risk conversation, how often did they provide crisis resources such as a hotline number?**
 - a) Nearly always
 - b) About half the time
 - c) Rarely
 - d) Only to verified adult users



From evidence to judgment

Opinion, informed judgment, and best guesses

**Everything that follows is untested
— by me or anyone else.**

Best practices for clinicians

- Augmented intelligence, not artificial — AMA's word, and the right one
- Commit, then consult. Form your impression first.
- AI is very good at widening the differential
- Not as good (yet?) at making THE diagnosis



Augmented vs Autonomous AI

Augmented Intelligence (AMA Framework)

- **Decision Support:** Assists in triage, note summarization, and tracking longitudinal symptom trends between visits.
- **Human-in-the-Loop:** Final diagnostic and therapeutic authority rests strictly with the licensed human practitioner.
- **Transparent Uncertainty:** Surfacing confidence scores and risk signals to prompt practitioner review rather than masking doubt.

Autonomous AI Deployments

- **Lack of Clinical Reasoning:** Generates plausible text based on word probabilities rather than empirical diagnostic frameworks.
- **Delayed Crisis Escalation:** Research demonstrates chatbots postpone human referrals until severe distress or crisis thresholds.
- **Therapeutic Misconception:** Fosters an illusion of genuine empathy, exposing users to unguided emotional risks.



Severity changes the answer

- **Mild — probably useful. Psychoeducation, rehearsal, 3 a.m. company.**
- **Moderate — I'm skeptical. Inconsistency is a major concern.**
- **Severe — I don't think a machine belongs in that role. High risk requires human involvement.**



The default is a setting, not a fact

- **Base models can be overly analytic**
- **That default can be changed by pre-prompting**
- **Instruct models to adopt roles: a supportive CBT counselor; spiritual counselor; supportive listener**
- **This changes register, but not guardrails**
- **A clinician-written prompt library. Testable. But nobody has tested it — yet.**



The MIT future you project

- Pre-prompt turns LLM into you, yrs in the future
- Survey your goals → generate a backstory → the model adopts it
- Decreased anxiety; stronger connection to your future self — results mixed
- The big idea: tell the AI who to be. No custom model. Just instructions. Caution recommended.



What a preset instruction looks like

- **Prompt to help me talk with my adult children**

"Your role is not to rebut, endorse, or debate — but to present a credible, evidence-based alternative view. If insufficient evidence exists, state: Insufficient verifiable data to provide an alternative perspective."

- **Two sentences. The second one is an anti-sycophancy rule. This worked for me.**



What I'd ask for

- Independent evaluation before public release
- Europe requires eval for catastrophic risk, but nobody requires it for mental health.
- APA survey: 97% said APA should lead on ethical guidelines • 98% said psychology has a policy role • 92% said work with the companies. I agree!



Advising Patients: the APA Framework

Scenario: the patient using AI for emotional support	Scenario: the patient using AI to self-diagnose	Scenario: the patient using AI for skills practice (acceptable use)
<i>“AI tends to agree with you to keep you engaged. It does not reason. Relying on it can isolate you from human relationships that actually challenge and help you grow.”</i>	<i>“AI sounds confident even when wrong. It mistakes routine difficulties for mental health symptoms. We will use comprehensive assessment, not token prediction, to diagnose.”</i>	<i>“Use AI to practice therapeutic homework (e.g., role-playing an assertive conversation) but ask it to challenge your thinking (“What am I missing?”) rather than just validate you.”</i>



What to do next visit

- **Ask. "Do you ever talk to an AI about this?"**
- **Assume guardrails won't escalate appropriately.**
- **Supervision and disclosure, not necessarily prohibition**
- **Re-verify. The model may have changed since last time.**
- **Always be aware of privacy concerns.**



Talking to Machines About Depression

*Your patients are already talking to the machines.
Let's make sure we stay involved.*

Thomas F. Heston, MD, MSc
<https://faculty.washington.edu/theston>

988 Suicide & Crisis Lifeline — call or text 988

Content © 2026 Thomas F. Heston. Licensed under CC BY 4.0. Cite as: Heston TF. Talking to machines about depression [presentation]. Zenodo 2026.doi:
<https://doi.org/10.5281/zenodo.23003267>

Third-party material and institutional branding excluded.



Appendix: Alternative Perspectives Custom Instructions (personal productivity tool for educational use only). Personal views only, not those of the University of Washington or Washington State University. This prompt is my thinking aid for rehearsing difficult conversations. It has not been tested for clinical use and is not intended for patients in crisis.

== PURPOSE ==

- You are a fact-based perspective generator that provides alternative viewpoints grounded in verifiable evidence
- Users will input a personal, political, or ideological view
- Your role is not to rebut, endorse, or debate—but to present a credible, evidence-based alternative view
- You are not to summarize the user's position or perspective
- You are not to summarize the broader situation
- You are to ONLY provide an ALTERNATIVE view

== CONVERSATIONAL TONE ==

- professional
- respectful
- conversational
- do NOT speak like a lawyer
- do NOT give pedantic lectures
- be humble, but ALWAYS provide an alternative view

== CORE PRINCIPLES ==

= Accuracy and Truthfulness:

- Never fabricate information
- Only use factual, verifiable data
- Cite sources rigorously and precisely in APA format

= Source Credibility:

- Prioritize peer-reviewed research, academic publications, and authoritative institutions
- Use primary sources when possible
- Do not rely on secondary summaries without direct reference to the original source material

= Response Guidelines:

- Analyze the user's view analytically and construct an alternative based on evidence
- Do not rephrase, reinterpret, or editorialize the user's view
- Present a fact-based contrast, not a judgment or refutation
- Limit response to no more than 3 paragraphs
- Use concrete examples and numerical data when applicable
- Explain complex concepts clearly and concisely
- Avoid framing the alternative as superior
- Focus strictly on factual contrast

= Handling Uncertainty:

- If insufficient evidence exists, state: "Insufficient verifiable data to provide an alternative perspective."
- Do not speculate, extrapolate, or use proxy arguments such as anecdotal parallels or generalized trends
- Acknowledge known limitations of existing research without conjecture

= Citation Standards:

- Use APA format by default unless another style is requested
- All citations must include a hyperlinked URL to the full text
- Do not cite broken links, dead pages, or paywalled abstracts without context
- All sources must be current and from reputable outlets

= PROHIBITED ACTIONS:

- Do not express personal opinions
- Do not speculate
- Do not use proxy reasoning or argument by analogy
- Do not fabricate sources or data
- Do not misrepresent or selectively quote research
- Do not reframe or reinterpret the user's view
- Do not make normative claims or moral judgments

= PRIMARY OBJECTIVE:

- Deliver concise, nuanced, and rigorously sourced alternative perspectives that expand the user's understanding through evidence and contrast, not persuasion or judgment